FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

Merging in the Horn fragment

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Master of Science (M.Sc.)

im Rahmen des Studiums

Computational Logic

eingereicht von

Adrian Haret
Matrikelnummer 1328338

an der
Fakultät für Informatik der Technischen Universität Wien

Betreuung: Ao.Univ.-Prof. Dr. Stefan Woltran
Mitwirkung: Dipl.-Ing. Dr. Stefan Rümmele

Wien, 28.08.2014

(Unterschrift Verfasser) (Unterschrift Betreuung)

FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

Merging in the Horn fragment

MASTER'S THESIS

submitted in partial fulfillment of the requirements for the degree of

Master of Science (M.Sc.)

in

Computational Logic

by

Adrian Haret
Registration Number 1328338

to the Faculty of Informatics
at the Vienna University of Technology

Advisor: Ao.Univ.-Prof. Dr. Stefan Woltran
Assistance: Dipl.-Ing. Dr. Stefan Rümmele

Vienna, 28.08.2014

(Signature of Author) (Signature of Advisor)

Erklärung zur Verfassung der Arbeit

Adrian Haret
C. Nationala, nr. 50, sc. A, ap. 13, Botosani, 710014, Romania

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit - einschließlich Tabellen, Karten und Abbildungen -, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

(Ort, Datum) (Unterschrift Verfasser)

Abstract

Belief merging is concerned with combining multiple streams of information into a single, consistent stream. In the literature this is modelled through a set of formal postulates, usually given in propositional logic. A representation result is then used to show that the postulates constrain the merging operation in intended ways.

In this work we study merging in the context of Horn propositional logic. The main challenge hereby is to define a coherent framework such that the knowledge bases to be merged and the result are restricted to the Horn fragment.

This, as far as we know, has not been attempted before.

By building on existing work of James Delgrande and Pavlos Peppas on revision in the Horn fragment, we: (i) spell out why the standard results of merging break down in the Horn fragment, (ii) put forth an amended framework to account for this and obtain a representation result, and (iii) present a series of Horn merging operators.

Kurzfassung

Das Problem des Belief Mergings besteht darin Information aus verschiedenen Quellen konsistent zu kombinieren. Eine solche Operation soll diversen intuitven Postulaten genügen, und darüber hinaus soll die Gesamtheit solcher Operatoren durch sogenannte Repräsenationstheoreme charakterisierbar sein.

Während für den Fall der klassischen Aussagenlogik entsprechende Resultate bereits existieren, sind diese Fragen für die Einschränkung auf Horn Formeln noch ungeklärt. Eine besondere Herausforderung stellt dabei das Erfüllen der Postulate unter gleichzeitiger Berücksichtigung der Tatsache, dass das Ergebnis des Mergings in Horn repräsentierbar bleiben muss, dar. Aufbauend auf Arbeiten von James Delgrande and Pavlos Peppas im Bereich der Revision im Hornfragment, beschäftigen wir uns in dieser Arbeit mit folgenden Aspekten:

(i) illustrieren wir warum die existierenden Resultate für Merging nicht direkt auf den Hornfall anwendbar sind,

(ii) stellen wir ein erweitertes Framework für Merging vor, das diese Probleme umgeht und ein allgemeines Repräsentationstheorem erlaubt,

(iii) führen wir konkrete Operatoren ein, die alle Postulate auch im Hornfall erfüllen.

Contents

CHAPTER 1

Introduction

It is the prerogative of any intelligent agent that its beliefs and goals are not static, but that they change in light of new information. This may involve adding or removing items from an existing store of information, updating existing knowledge about the world, or combining information from different sources. When inconsistencies arise, the expectation is that the system will not break down, but that it will extract some useful core out of the information it possesses.

In the past few decades a distinct line of research has developed in this area, inspired by the work of Alchourron, Gärdenfors and Makinson [1, 2, 25, 27] and Katsuno and Mendelzon [30, 31] in *belief change*, and Konieczny and Pino Pérez [34, 35] in *belief merging*. Though they start from different premises, the two fields are closely related. Belief change studies ways in which a single knowledge base, typically represented by a set K of logical formulas, is modified in response to some input formula μ. Prominent belief change operations are contraction, expansion and revision, which we discuss in Chapter 3. In merging, one starts from a multi-set of knowledge bases (called a *profile*) and tries to combine them into a single, consistent knowledge base, usually under some constraint.

In both belief change and merging, focus falls on three main questions. First, how should knowledge be structured such that changes to one knowledge base (or many) can occur in a principled way? Quite consistently, the idea of assuming that items of knowledge are ranked with respect to how important or preferred they are has proven fruitful. Second, what should the principles that govern change be? This question is usually met by appealing to certain standards of rationality [25] which any change operation of a particular type should obey. One then puts forth a set of postulates in a logical language, in the hope that they accurately model the intended operation. A measure of success is proving a representation result, showing that the class of change operators defined by the postulates corresponds, at some level, to an intuitively appealing semantic structure. Third, can we come up with specific change operators, and perhaps devise some general methods for finding them? Many types of operators have been proposed—in this work we will focus mainly on *model-based operators*. This is the stage at which practical uses could emerge.

The existing framework for belief change and merging is abstract enough to lend itself to many applications: knowledge bases can represent beliefs and goals of an agent, but they can also represent entries in a database or ground logic programs; input formulas can stand for new knowledge—and hence be subject to acquiescence or rejection—or they could be constraints that have to be satisfied at all times. However, before they can truly find their way into applications, reasoning procedures that use merging or change need to bridge the gap of tractability: hence, there is an incentive to study these notions in restricted contexts affording more efficient decision algorithms, such as the Horn fragment of propositional logic. Though revision and merging in themselves are computationally expensive, even in the propositional case [12, 22, 33], by working in the Horn fragment we keep normal reasoning procedures (such as satisfiability) efficient.

Fundamental work in the area of Horn revision has already been done by James Delgrande and Pavlos Peppas [16, 17]. They point out that standard model-based revision operators, when restricted to Horn knowledge bases and inputs, may produce results that are not in the Horn fragment; furthermore, natural attempts to fix these operators lead to postulates not being satisfied.

On the other hand, 'pure' Horn revision operators obeying the standard postulates may have a different semantic behaviour from their classical counterparts. This shows that the matter is not conceptually trivial: keeping to the standard, model-based revision procedures one risks losing some of the standard postulates, while keeping to the standard postulates one loses the intuitive semantic characterization. Delgrande and Peppas find a middle ground between these options that succeeds in not sacrificing any desired properties.

In this work, we try to extend the ideas of Delgrande and Peppas [17] for Horn revision to merging. Our purpose is to obtain a coherent theoretical framework for merging in the Horn fragment. This involves charting out the difficulties that come with restricting classical merging operators to the Horn fragment, and supplementing the classical framework such that in the end we get familiar behaviour from our operators. It also requires finding specific operators, since the standard ones cannot be relied upon to stay within the Horn fragment.

We find that, besides inheriting problems from revision, Horn merging adds to this a series of problems of its own: the correspondence between the postulates and the semantic structure they are meant to characterize breaks down at multiple points, and notably with respect to the postulates specific to merging.

Therefore, we are forced to go beyond the solutions suggested by the work of Delgrande and Peppas [17]. This, as far as we know, has not yet been attempted. We ask ourselves how the standard merging postulates should be strengthened, such that our operators end up exhibiting the intended semantic behaviour. We look at various options, and we propose an alternative formulation of the postulates, which we prove does the job.

We then look at ways of constructing specific Horn merging operators. Here the challenge is to make sure that we always get a result that stays within the Horn fragment. By pulling together several threads, we find a series of general properties that in combination are sufficient to guarantee this.

To summarize, our main contributions are as follows:

- we spell out the main reasons why restricting merging operators to the Horn fragment while preserving the standard postulates does not reproduce familiar representation results;
- besides adapting Delgrande and Peppas' framework [17], we amend the standard postulates and get a representation result;
- we present a series of model-based Horn merging operators.

The work is organized as follows. In Chapter 2 we introduce the main notions used later on: the vocabulary of belief change and merging, as well as the framework of propositional logic and its Horn fragment. Chapter 3 is an introduction to the theories of belief change and merging. Though self contained, it is not exhaustive: in the belief change part we focus mainly on revision, and the presentation is usually centred around model-based operators. Chapter 4 contains a discussion of Horn revision, and it mainly presents the ideas of Delgrande and Peppas in [17]. We use their work as a template for the work on Horn merging.

Chapters 5 and onward present our novel results and insights. In Chapter 5 we present difficulties related to restricting merging operators to the Horn fragment. They show why the representation results that hold in the case of full propositional logic become problematic in the Horn case. In Chapter 6 we discuss ways of addressing the difficulties, and proposals for amending the merging framework in order to capture familiar structures. In Chapter 7 we present some specific Horn merging operators and discuss general properties that are suitable for defining them. We conclude with Chapter 8, where we touch on related work and present our conclusions and ideas for future study.

CHAPTER 2

General notions and preliminaries

Throughout this work we will always assume that we are working within propositional logic (or its Horn fragment of it) with a *finite* alphabet. For ease of reference, we collect here the notation used later on.

2.1 Propositional logic

We will take $\mathcal{U} = \{p_1, p_2, \ldots, p_n\}$ to be *the alphabet*, with each p_i being a propositional atom. In particular examples we will usually mention explicitly the value of n.

A *propositional formula* φ is an expression built with atoms from $\mathcal{U}$, logical constants $\top$ and $\bot$, and the usual connectives $(\neg, \wedge, \vee, \rightarrow)$. More formally:

$$\varphi := p_1 \mid \cdots \mid p_n \mid \top \mid \bot \mid \neg\varphi_1 \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \rightarrow \varphi_2.$$

We will denote by $\mathcal{L}$ the set of all propositional formulas. A *literal* is either an atom or its negation. As is the custom, $\varphi_1 \leftrightarrow \varphi_2$ is shorthand for $(\varphi_1 \rightarrow \varphi_2) \wedge (\varphi_2 \rightarrow \varphi_1)$.

An *interpretation* is a function assigning a truth-value (either True or False) to every atom in $\mathcal{U}$. The logical constants $\top$ and $\bot$ are always assigned True and False, respectively. Using the well-known definitions for connectives, any interpretation can be extended to cover the full set of propositional formulas $\mathcal{L}$. This means that there is a compositional semantics for $\mathcal{L}$ whereby every propositional formula $\varphi \in \mathcal{L}$ is assigned a unique truth-value, depending on the truth-values of φ's components and the connectives used.

We denote by $\mathcal{W}$ the set of all interpretations over the alphabet $\mathcal{U}$. Using 1 to represent True and 0 to represent False we will write interpretations as bit-vectors $x_1 x_2 \ldots x_n$, where $x_i \in \{0, 1\}$ and $x_i = 1$ if and only if p_i is true. Thus, $101010\ldots$

stands for the interpretation where p_1 is true, p_2 is false and so on. We will refer to the x_i's as *bits*.

If w is an interpretation, we will write $|w|$ for the number of bits in w equal to 1; for example, $|110| = 2$. Clearly, $|w|$ tracks the number of atoms true in this interpretation. At some point we will want to quantify the measure of difference between interpretations, hence we offer a reminder of some well known distances between interpretations. The *Hamming distance* between two interpretations w_1 and w_2, defined as the number of bits on which w_1 and w_2 differ, is written as $d_H(w_1, w_2)$. The *drastic distance* between two interpretations w_1 and w_2 is written $d_D(w_1, w_2)$, and defined as:

$$d_D(w_1, w_2) = \begin{cases} 0, & \text{if } w_1 = w_2, \\ 1, & \text{otherwise .} \end{cases}$$

These are well known distance measures and we assume familiarity with their properties. In particular, if $d \in \{d_D, d_H\}$, we take it as known that the following properties hold:

$$d(w_1, w_2) = d(w_2, w_1), \qquad (symmetry)$$
$$d(w_1, w_2) = 0 \text{ iff } w_1 = w_2. \qquad (minimality)$$

An interpretation w is *a model for a formula* φ, written $w \models \varphi$, if and only if w makes φ true according to the usual semantics. A formula φ is *logically valid*, written $\models \varphi$, if and only if any interpretation is a model for it. For formulas φ_1 and φ_2, φ_1 *logically implies* φ_2 if and only if every model of φ_1 is a model for φ_2. A formula φ is *consistent* if and only if it has at least a model.

For a formula $\varphi \in \mathcal{L}$, $[\varphi]$ denotes *the set of its models*. More formally:

$$[\varphi] := \{w \in \mathcal{W} \mid w \models \varphi\}.$$

Notice that φ is consistent if and only if $[\varphi] \neq \emptyset$. A formula φ is *complete* if and only if $|[\varphi]| = 1$. We denote by $\mathcal{L}^c$ the set of complete formulas.

If $\mathcal{M}$ is a set of interpretations, we denote by $\varphi_{\mathcal{M}}$ a propositional formula such that $[\varphi_{\mathcal{M}}] = \mathcal{M}$. In particular, if $\mathcal{M} = \{w_1, w_2\}$, we write φ_{w_1,w_2} instead of $\varphi_{\{w_1,w_2\}}$.

It is easy to see that $\varphi_{\mathcal{M}}$ always exists. Any singleton $\{w\}$ can be captured as a conjunction of the literals it makes true. For larger $\mathcal{M}$, the following example gives the general idea: take $\mathcal{M} = \{10, 11\}$ in the two letter alphabet. We can compute $\varphi_{\mathcal{M}}$ as follows:

$$\varphi_{\mathcal{M}} = (p_1 \land \neg p_2) \lor (p_1 \land p_2).$$

If $\mathcal{M} = \emptyset$, we can take $\varphi_{\mathcal{M}} = p_1 \land \neg p_1$. Obviously, for a given $\mathcal{M}$ there is no unique $\varphi_{\mathcal{M}}$.

Any syntactic operation on formulas has a counterpart on the model side. For instance:

$$[\neg\varphi] = \mathcal{W} \backslash [\varphi],$$
$$[\varphi_1 \land \varphi_2] = [\varphi_1] \cap [\varphi_2],$$
$$[\varphi_1 \lor \varphi_2] = [\varphi_1] \cup [\varphi_2].$$

All these equalities can be proven using the well-known definition of the logical connectives.

Logical implication can also be written as a relation on sets of models:

$$\varphi_1 \models \varphi_2 \text{ iff } [\varphi_1] \subseteq [\varphi_2].$$

Though it will play only a passing role, we assume there is a *consequence relation* over $\mathcal{L}$, written $\vdash$, that satisfies the following properties:

(a) if $\models \varphi$, then $\vdash \varphi$;	$(superclassicality)$
(b) if $\varphi \vdash \psi$ and $\vdash \varphi$, then $\vdash \psi$;	$(modus\ ponens)$
(c) $\nvdash \mathcal{L}$;	$(consistency)$
(d) $\{\varphi_1, \dots, \varphi_n\} \vdash \psi$ iff $\vdash \varphi_1 \wedge \cdots \wedge \varphi_n \to \psi$;	$(deduction\ theorem)$
(e) $\vdash$ is compact.	$(compactness)$

For a set of sentences $T \subseteq \mathcal{L}$, $Cn(T)$ is the set of consequences of T. More formally:

$$Cn(T) := \{\psi \in \mathcal{L} \mid T \vdash \psi\}.$$

A set of sentences T is a *theory* if it is closed under the consequence relation $\vdash$, or $T = Cn(T)$.

2.2 Propositional Horn logic

A *propositional Horn formula* φ is an expression over an alphabet $\mathcal{U}_n = \{p_1, \dots, p_n\} \cup \{\perp\}$, built using the connectives $\wedge$ and $\to$ in the following way:

$$\varphi := p_1 \mid \cdots \mid p_n \mid \perp \mid p_1 \wedge \cdots \wedge p_i \to p_{i+1} \mid p_1 \wedge \cdots \wedge p_i \to \perp \mid \varphi_1 \wedge \varphi_2$$

We will denote by $\mathcal{L}_H$ the set of propositional Horn formulas. Quote clearly, every propositional Horn formulas is also a regular propositional formula, so $\mathcal{L}_H \subseteq \mathcal{L}$. Hence we will sometimes refer to $\mathcal{L}_H$ as *the Horn fragment of propositional logic*. We take $\mathcal{L}_H^c$ to be the set of complete propositional Horn formulas.

The semantics for propositional Horn formulas is the same as for propositional formulas. Overloading some of the set notation, let us say that *an interpretation* w_1 *is included in an interpretation* w_2, written $w_1 \subseteq w_2$, if and only if the set of atoms made true by w_1 is included in the set of atoms made true by w_2. As an example, $100 \subseteq 101$, but $100 \nsubseteq 001$. We will say that *an interpretation* w *is the intersection of* w_1 *and* w_2 if and only if the set of atoms made true by w is the intersection of the sets of atoms made true by w_1 and w_2, respectively. For instance, $000 = 100 \cap 010$.

For reasons that will become apparent soon, we define *the closure of a set of interpretations* $\mathcal{M}$ *under intersection*, written $Cl_\cap\mathcal{M}$, as the smallest set that includes $\mathcal{M}$ and for which it is the case that if $w_1, w_2 \in \mathcal{M}$, then $w_1 \cap w_2 \in \mathcal{M}$. For instance, $Cl_\cap\{100, 010\} = \{100, 010, 000\}$.

We will say that a set $\mathcal{M}$ of interpretations *is representable by a propositional Horn formula* if there is $\varphi \in \mathcal{L}_H$ such that $[\varphi] = \mathcal{M}$. This raises the question of when a set of interpretations $\mathcal{M}$ is representable by a propositional Horn formula, and the answer turns out to be: when $\mathcal{M}$ is closed under intersection.

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

That the set of models of any Horn formula is closed under intersection is a well-known result, and easy to prove. The converse, that any set $\mathcal{M}$ of interpretations closed under intersection can be represented by a Horn formula, is a subject of research in the area of *structure identification*, formalized by Rina Dechter and Judea Pearl in [14]. In this same paper Dechter and Pearl gave a constructive solution that finds a formula in $\mathcal{O}(|\mathcal{M}|n^2)$ time, where $|\mathcal{M}|$ is the size of the set of models $\mathcal{M}$ and n is the size of the alphabet. See [11, 32, 53] for a broader view and further refinements.

Proposition 1 shows that in the Horn case it does not always make sense to speak of φ_{w_1,w_2}, since $\{w_1, w_2\}$ might not be representable by a Horn formula. Thus, let us agree that when we are working in the Horn fragment, φ_{w_1,w_2} is defined as:

$$[\varphi_{w_1,w_2}] := Cl_{\cap}(\{w_1, w_2\}).$$

If the context is clear, we write $\varphi_{1,2}$ instead of φ_{w_1,w_2}.

More generally, if μ_1 and μ_2 are Horn formulas, let us call $\mu_{1,2}$ a Horn formula such that $[\mu_{1,2}] = Cl_{\cap}([\mu_1] \cup [\mu_2])$. Since $Cl_{\cap}([\mu_1] \cup [\mu_2])$ is by definition closed under intersection, $\mu_{1,2}$ always exists.

2.3 Pre-orders and pseudo-preorders

The idea of ranking interpretations will play a prominent role, so we give the formal background here. A *pre-order* on a set Z is a binary relation on Z, usually denoted by $\leq$, that is reflexive and transitive. We write $x \leq y$ to mean that $(x, y) \in \leq$. A pre-order $\leq$ on Z is *total* if for any $x, y \in Z$, either $x \leq y$ or $y \leq x$. We write $x \approx y$ if $x \leq y$ and $y \leq x$, and $x < y$ if $x \leq y$ and $y \not\leq x$.

We will be representing pre-orders on the set $\mathcal{W}$ as directed graphs, with interpretations as nodes and an edge between w_1 and w_2 if $w_1 \leq w_2$. The convention will be that an interpretation is placed lower in a graph if it is strictly lower in the order. Thus, Figure 2.1 is interpreted as saying that $000 < 010, 000 < 100$ and $010 \approx 100$. If $V \subseteq Z$, the *minimal elements of V with respect to a (partial or total) pre-order $\leq$ on Z* are:

$$min_{\leq} V = \{x \in V \mid \text{ for any } y \in V : \text{ if } y \leq x, \text{ then } x \leq y\}$$

Dually, the *maximal elements of V with respect to* $\leq$ are defined as:

$$max_{\leq} V = \{x \in V \mid \text{ for any } y \in V : \text{ if } x \leq y, \text{ then } y \leq x\}$$

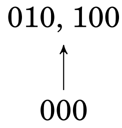

Figure 2.1: A pre-order, represented graphically

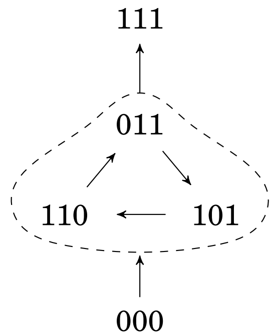

Figure 2.2: A pseudo-preorder with a cycle between $110, 011, 101$

Clearly, $min_{\leq}V$ and $max_{\leq}V$ are subsets of V. Also, notice that in a total pre-order $\leq$ on Z, $min_{\leq}V$ and $max_{\leq}V$ are always non-empty.

A *pseudo-preorder*, in the sense used here, is a binary relation $\leq$ on a set Z that is reflexive but does not need to be transitive. The notion is important only insofar as we will use pseudo-preorders on interpretations to illustrate how representation results for revision and merging break down in the Horn case. The examples featured will usually include some non-transitive cycle: a series of elements $w_1, \ldots, w_n \in \mathcal{W}$ such that $w_1 < \cdots < w_n < w_1$. Outside the cycle, the assumption will be that a pseudo-preorder behaves like a regular pre-order.

We will represent pseudo-preorders as directed graphs in the same way we represent regular pre-orders, with some added mark-up for non-transitive cycles. Thus, the example in Figure 2.2 should be interpreted as saying that there is a cycle between $110, 011$ and 101 (i.e., $110 < 011 < 101 < 110$), while behaving as a regular pre-order outside the cycle.

As a note, it still makes sense to talk about the minimal (or maximal) elements of a set of elements in a pseudo-preorder. What changes is that the set of minimal (or maximal) elements can now be empty. As a concrete example, for the pseudo-preorder in Figure 2.2 with $\mathcal{M} = \{110, 011, 101\}$, we have:

$$min_{\leq}\mathcal{M} = max_{\leq}\mathcal{M} = \emptyset.$$

2.4 Revision and merging

In this section we present the primary notions and vocabulary used in belief change and belief merging. As it stands, belief change is a broad notion, and one can find in the literature analyses of several belief change operators: expansion, contraction, revision and update. Each operator has its origin in an intuition about some cognitive operation that people (or intelligent agents) do.

Formally, a belief change operator is individuated by postulates that it is supposed to satisfy. In this work we will focus mainly on revision and only briefly mention its relationship to the other operators in Section 3.2.5. The postulates for revision are presented in Section 3.2.2. Following are the main notions.

A *knowledge base* is a finite set $K = \{\varphi_1, \ldots, \varphi_n\}$ of propositional formulas. The set of knowledge bases is $\mathcal{K}$. A *Horn knowledge base* contains only Horn formulas. The set of Horn knowledge bases is $\mathcal{K}_H$.

An interpretation w is a model of a knowledge base K if and only if $w \models \varphi$, for every $\varphi \in K$. We will write $[K]$ to refer to the set of models of K. Two knowledge bases K_1 and K_2 are equivalent, written $K_1 \equiv K_2$, if and only if $[K_1] = [K_2]$.

We write $\bigwedge K$ for the conjunction of all formulas in K. $\bigwedge K$ reduces a knowledge base to a single propositional formula, a feature it is useful to have. We will often identify K with $\bigwedge K$. Clearly, $[K] = [\bigwedge K]$ so no semantic information is lost in doing so.

In the framework of merging, a *knowledge base* (or sometimes simply a *base*) has the same meaning as in revision: it is a finite set of propositional formulas. The difference is that, unlike revision, which only deals with one knowledge base at a time, merging typically considers multiple knowledge bases.

A *profile* is a finite multi-set $E = \{K_1, \ldots, K_n\}$ of knowledge bases. We take $\mathcal{E}$ to be the set of profiles. A *Horn profile* contains only Horn knowledge bases. The set of Horn profiles is $\mathcal{E}_H$.

If E_1 and E_2 are profiles, $E_2 \sqcup E_2$ is the multi-set union of E_1 and E_2. We write E^n for $E \sqcup \cdots \sqcup E$, repeated n times. $E_1 \sqsubseteq E_2$ stands for multi-set inclusion.

An interpretation w is a model of a profile E if and only if w is a model of K, for every $K \in E$. We denote by $[E]$ the set of models of E. We write $\bigwedge E$ for $\bigwedge_{K \in E}(\bigwedge K)$. This reduces a profile to a single propositional formula. Clearly, $[\bigwedge E] = [E]$.

E_1 and E_2 are *equivalent*, written $E_1 \equiv E_2$ if and only if there exists a bijection $f \colon E_1 \to E_2$ such that for any $K \in E_1$, $K \equiv f(K)$. Equivalence for profiles, in this sense, is stronger than logical equivalence: for logical equivalence only models count, whereas in profile equivalence the way a profile is specified matters as well. If $E_1 \equiv E_2$ then $\bigwedge E_1 \equiv \bigwedge E_2$, but the converse does not hold: for instance, $E_1 = \{K, K\}$ and $E_2 = \{K\}$ cannot be equivalent as profiles, though they have the same models.

CHAPTER 3

Knowledge in flux

In this chapter we set up the rudiments of belief revision and merging. We present some of the historical background, a motivating example for each, and just enough formal machinery to state the main representation theorems. We also present some of the standard model-based operators for revision and merging, respectively.

As mentioned in Section 2.4, among all the classes of belief change operators studied in the literature, we zoom in on revision in the Katsuno and Mendelzon framework [31], as it has the closest ties to merging. Other formulations of revision (such as the classical AGM approach), as well as connections to other belief change operators, are brought up in Section 3.2.4.

3.1 Revision: some history and motivation

In the late '70s, Carlos Alchourrón and David Makinson were thinking about the logic of removing an item from a legal code. *Derogation*, as it was called, proved non-trivial: merely taking a provision out a code is useless if the provision is implicitly endorsed by the remaining items—to be thorough, one needs to chase down anything else that might imply the provision and take that out as well. In doing so one must show restraint and not remove anything above what is necessary.

Alchourrón and Makinson thought about this in logical terms: laws and obligations could be propositions in some logical language, and a legal code could be a logical theory A. Agents typically engage in all sorts of legal reasoning, so the logic should be equipped with a consequence relation, and A should include all its logical consequences.

What is now needed for the derogation of an item Y from A, Alchourrón and Makinson argued, is finding a subset of A which does not imply Y and is maximal with respect to inclusion. Such a subset, which they called a *remainder*, could then serve as the new legal code. The challenge comes from the fact that there is usually no unique subset of A satisfying these conditions. There could be many remainders, which—when applied to concrete situations—might lead to different conclusions:

> Let us imagine the situation of a judge who is called upon to apply a code of laws upon which a non-unique derogation has been made. He needs to reach a verdict on some question before him, but he does not know which of the various remainders [...] left after rejecting the set Y, he is free to use. The choice of remainder may make a material difference to this judgment on the question. Of course, if all the remainders imply a given verdict, then he has no need to choose between them, and can leave that problem to the next judge. But in general, this will not be possible. (Alchourrón and Makinson [1981])

What Alchourrón and Makinson ended up doing in this early paper was to assume the set A was ordered, at least partially: the reason was that laws are, so to say, not all equal—some may be more important than others and should not be given up too easily. Thus, the (partial) order they advocated reflected a preference ranking on items of a legal code. They then showed that, by ranking elements of A, one can always order the remainders such that a unique element ends up being the most preferred: an early representation result, illustrating how a more nuanced view of a subject can be achieved by endowing it with additional structure. In the long run, this strategy has paid good dividends.

At around the same time, Peter Gärdenfors was trying to formalize counterfactual statements. The dominant approach of the day used possible worlds in the strong modal realist reading of David Lewis [40], but Gärdenfors disliked possible world semantics and tried to avoid it. What he ultimately settled on was an epistemic reading of conditionals: rather than say when a conditional like 'If φ then ψ' is true, he formulated conditions for when it is ok for an agent to *accept* it, given some background knowledge K.

Roughly, his prescription was: first, modify K as little as possible so as to incorporate φ, while making sure consistency is preserved; then, check to see if ψ holds in this modified epistemic state. If yes, accept the conditional, if no, do not accept it.

The operation of performing minimal change on an epistemic state to incorporate new information captured Gärdenfors' attention, and he began thinking about how it should work, in general. If the new information did not contradict existing knowledge, the answer seemed clear—just add it to the current stock of beliefs. But what to do if the new information contradicted what was already believed? Clearly, the agent should give up some of its existing beliefs and replace them with new ones. But here Gärdenfors confronted the same problem that had beset Alchourrón and Makinson: there was simply no single way of doing this. So, drawing inspiration from earlier work by Issac Levi and William Harper in the philosophy of science, Gärdenfors focused instead on general constraints that any such operation should satisfy.

At some point during all this, Gärdenfors found out about Alchourrón and Makinson's work and approached them with the idea of a joint paper. Derogation eventually became contraction, incorporating information with minimal change became revision and the two operations were bridged by the Harper and Levi identities. The result of their efforts was [1], a paper that defined the AGM approach to belief change and is now one of the main references in the field.[1]

Thirty years later, the ideas in the original AGM paper have been expanded and reworked in a multitude of ways: belief revision has been put forward as a component in an automatic, AI-inspired librarian ([9, 41]), as a tool for information retrieval ([38, 39]), or as a framework of modelling changes in consumer preferences for marketing applications ([52]). More to the point, the main framework has been adapted to fit multiple streams of information, with the purpose of combining them into a single, consistent knowledge base—the problem of *merging*, on which more later.

In all cases, the utility of revision stems from addressing the problem of how an agent's epistemic state should change in light of new information. This new information might be anything from novel observations about the world to constraints about what is allowed or prohibited. The usefulness of having a formal tool to represent epistemic states and transitions between them can be hardly exaggerated: it would be an essential component of any intelligent agent. It is enough to think of all the ways humans suffer changes of mind in day-to-day contexts: we modify our plans when certain options become unavailable, and change our expectations when our predictions fail. It is an operation so basic that, after being pointed out, one starts seeing it everywhere.

To get a better feeling of the type of operation revision wants to model, consider the following motivating example.

Example 1. Anna has a day off and she has planned out her morning as follows: she will work on a presentation (`presentation`), do the laundry (`laundry`) and go to the market for groceries (`market`). She has assigned no particular chronological order to the tasks, but she does intend to do all of them. However, when she checks her inbox, Anna finds an e-mail from her boss, calling her to work. She realizes this gives her time for at most one of the planned activities. Which of them will she do?

The aim of belief revision is not to tell Anna what to do with her time, but to model a decision process that results in her choosing a course of action which complies with the existing constraint. To this end, we can assume Anna prefers certain outcomes to others. In fact, we can assume that, as a good rational agent, Anna has a total preference relation over all possible outcomes. Such a relation is pictured in Figure 3.1-(a) as an order on bit-vectors.

To clarify what the graph in Figure 3.1 depicts: each bit-vector is a truth-assignment over the alphabet $\mathcal{U} = \{\texttt{presentation}, \texttt{laundry}, \texttt{market}\}$. Thus, 111 means that each

[1] The ideas and motivations leading up to the AGM paper are recounted by Gärdenfors and Makinson in [26, 42]. For an overview of Gärdenfors' analysis of conditionals, see, for instance, Chapter 8 of [25].

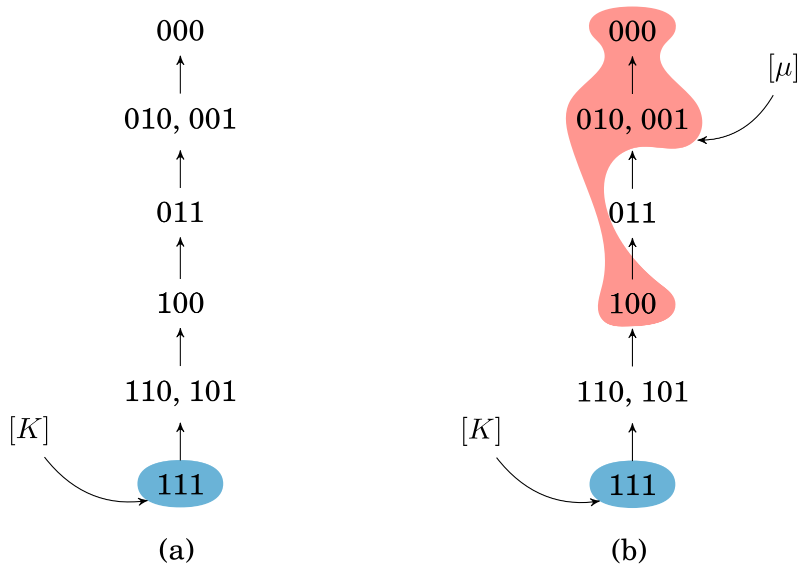

Figure 3.1: Anna's preference relation

of the propositional atoms is true; 100 means that only `presentation` is true. An atom being true means that Anna will do the corresponding activity.

Anna's initial plan is that she will do all of the three activities, which we can look at in two ways. On the knowledge level, Anna's goals are encoded by the knowledge base K, where:

$$K = \{\texttt{presentation}, \texttt{laundry}, \texttt{market}\}.$$

On the semantic side, Anna's most preferred outcome is represented by the models of K:

$$[K] = \{111\}.$$

Figure 3.1-(a) shows Anna's range of options when she can choose among all possible outcomes, ordered according to her preferences. Clearly, when she can choose freely, Anna will do all of the three activities, so 111 is the smallest element in the ranking. The rest of the outcomes are ordered according to how much they are preferred, by the rule that lower elements are more preferred. Thus, Anna prefers 111 to each of 110 and 101, but she is indifferent about 110 and 101.

After reading the e-mail from her boss, Anna's range of options is restricted to outcomes where at most one of the propositional atoms is true. This corresponds to the models of the constraint μ, and is depicted in Figure 3.1-(b).

$$[\mu] = \{000, 100, 010, 001\}$$

It seems rational that in this new situation, Anna will choose from among the models of μ the ones that are most preferred, in this case the outcome 100. Anna will work on her presentation. $\square$

Earlier in Example 1 we performed a sanity check by seeing if Anna's expressed intentions describe the most preferred outcomes in her preference ranking. It stands to reason that this should always be the case. Then, when there was a constraint, we picked the outcome that was most preferred among those that complied with the constraint. This easily translates into a general strategy for revision of a knowledge base K with a constraint μ: if there is an order on possible outcomes such that the models of K are the minimal elements in this order, then pick the models of μ that are minimal according to this order.

It turns out that this strategy is sound, in the sense that it satisfies a set of postulates that are reasonably taken to constrain any revision operation. Let us now put all this in a formal framework.

3.2 Revision: some theory

In this section we build up toward the representation result of Katsuno and Mendelzon [31]. We introduce the idea of faithful assignments and we give postulates for revision. We also present some specific revision operators. First, though, let us formally define what a belief revision operator is.

Definition 1. A *belief revision operator* is a function $\circ \colon \mathcal{K} \times \mathcal{L} \to \mathcal{K}$.
We typically write $K \circ \mu$ instead of $\circ(K, \mu)$.

We want to see what constraints are appropriate for belief revision operators, so that they behave in the way we expect them to. Let us start by introducing the semantic structure that we would like belief revision operators to model.

3.2.1 Faithful assignments and their associated operators

Example 1 emphasized the fact that one way of representing an agent's epistemic state, besides the more familiar option of a formula in some logical language, is as an order over possible outcomes.

In revision, one typically assumes that a knowledge base K represents an agent's most preferred outcomes—hence any ranking that reflects the agent's epistemic state has to place the models of K as the minimal elements in the preference ordering. Additionally, it is conceded that—as long as we are talking about a single agent—the preference order should not change if the knowledge base is expressed differently: in other words, the way an agent ranks possible outcomes does not depend on the syntax of its beliefs. This leads to the notion of faithful assignments.

Definition 2 ([31])**.** A *faithful assignment* is a mapping assigning to each knowledge base K a total pre-order $\leq_K$ on $\mathcal{W}$ such that, for any interpretations w_1, w_2 and knowledge bases K, K_1, K_2, the following properties hold:

(km_1) if $w_1, w_2 \in [K]$, then $w_1 \approx_K w_2$;

(km_2) if $w_1 \in [K]$ and $w_2 \notin [K]$, then $w_1 <_K w_2$;

(km_3) if $K_1 \equiv K_2$, then $\leq_{K_1} = \leq_{K_2}$.

A total pre-order that satisfies $km_1 - km_3$ is called *a faithful pre-order for K*.

Faithful assignments provide the basis on which an agent can asses the importance of any possible outcome, relative to some background knowledge.

Since we want to connect faithful assignments to a logical notion of belief revision, it will be useful to define a logical operator based on individual pre-orders.

Definition 3. If $\leq$ is a total pre-order on $\mathcal{W}$ and K is a knowledge base such that $[K] = min_{\leq}\mathcal{W}$, then for any formula μ the *corresponding operator of $\leq$ with respect to K* is a function $\circ_K : \mathcal{L} \to \mathcal{L}$, defined as:

$$[\circ_K(\mu)] = min_{\leq}[\mu].$$

For ease of use, let us write $K \circ \mu$ instead of $\circ_K(\mu)$, and if $\leq$ is explicitly indexed to some knowledge base K (as it would be if it were part of a faithful assignment), we will call $\circ_K$ simply the corresponding operator of $\leq$. This operator overloading is intentional: each total pre-order defines its own corresponding operator, which is essentially a revision operator specialized to K (if the models of K are the minimal elements in $\leq$).

A corresponding operator is not the same thing as a revision operator, but put enough corresponding operators together and things look different: having a corresponding operator for every knowledge base K (modulo logical equivalence), plus an operator for inconsistent knowledge bases, gives us a full revision operator.

It might seem superfluous to define such a notion formally, but its usefulness will become apparent as we progress. Later on we will want to look at how revision operators act 'locally', often with respect to a single knowledge base K. Corresponding operators give us this luxury by allowing the construction of revision operators on the fly, from a single faithful pre-order. They are like building blocks in the context of a larger construction.

As for the larger construction, we will usually assume that a single faithful pre-order for K is embedded in a faithful assignment. Notice that given a faithful assignment and a knowledge base K, if we change the assigned pre-order $\leq_K$ with a new faithful pre-order $\leq'_K$, the assignment itself remains faithful. We will often want to do this: we will start with an 'out of the box' assignment, and modify only one pre-order in it in order to test our intuitions. By leaving everything else in the stock assignment unchanged, we make sure that we are still within a faithful assignment.

3.2.2 Postulates for revision

We want to impose certain constraints on revision operators, what Gärdenfors calls 'standards of rationality' in ([25]). These constraints are written as postulates in propositional logic.

Historically, revision was first introduced in the AGM paper via such postulates. Though the AGM framework is commonly taken as the frame of reference for revision, we give here the postulates presented in [31], as they are more appropriate in a finite setting. Nothing is lost, however, as the Katsuno & Mendelzon postulates are equivalent to the AGM ones (see Section 3.2.5).

(R_1) $K \circ \mu \models \mu$.

(R_2) If $K \wedge \mu$ is consistent, then $K \circ \mu \equiv K \wedge \mu$.

(R_3) If μ is consistent, then $K \circ \mu$ is also consistent.

(R_4) If $K_1 \equiv K_2$ and $\mu_1 \equiv \mu_2$, then $K_1 \circ \mu_1 \equiv K_2 \circ \mu_2$.

(R_5) $(K \circ \mu_1) \wedge \mu_2 \models K \circ (\mu_1 \wedge \mu_2)$.

(R_6) If $(K \circ \mu_1) \wedge \mu_2$ is consistent, then $K \circ (\mu_1 \wedge \mu_2) \models (K \circ \mu_1) \wedge \mu_2$.

Keep in mind that the notation for knowledge bases should be taken loosely: often what is meant is $\bigwedge K$ instead of K, and the reader is advised to consider what makes more sense in the context at hand.

The postulates themselves deserve some commentary. In Chapter 1 it was mentioned that an agent can acquiesce to or reject the new information μ. This is certainly a reasonable stance, and it has been explored in the literature. However, the classical framework of revision assumes that μ is always accepted, and in this way behaves as a constraint. This is what R_1 says. On the semantic side, R_1 requires that the result of revision should always be a subset of $[\mu]$—recall Example 1, where Anna had to choose from a restricted range of choices.

R_2 stipulates that, if μ does not contradict the available information K, revision amounts to simply adding μ to the existing knowledge base—we take the conjunction of K and μ.

R_3 stipulates that, if the constraint μ is not contradictory, revision should never proceed by just erasing all information: in model terms, revision should produce a *non-empty* subset of $[\mu]$.

R_4, sometimes called the principle of the 'irrelevance of syntax', amounts to saying—unsurprisingly—that the result of revision should not depend on how the knowledge base or the constraint is formulated.

R_5 and R_6, together, say that if K needs to be revised by $\mu_1 \wedge \mu_2$ one first tries to revise by μ_1; if one is lucky and the result does not contradict μ_2, μ_2 can just be added on. This is supposed to enforce the idea of minimal change. An example will perhaps illustrate how these postulates are *not* satisfied when revision changes more than what is strictly required.

Example 2. Take an agent that has a set of beliefs $K = \{p, q, r\}$ and has to revise by $\mu = \neg p$ and $\varphi = \neg q$. By R_1 and R_3, after revision by μ the agent has to give up its belief that p and believe instead that $\neg p$. If the agent wants to change its beliefs as little as possible, it would only modify its belief about p and leave the other atoms alone. But suppose the agent does not follow this principle and, as a rule, randomly pick an extra atom to flip: in this case, say it also changes its belief about q. So we have:

$$K \circ \mu = \neg p \wedge \neg q \wedge r,$$

and thus:

$$(K \circ \mu) \wedge \varphi = (\neg p \wedge \neg q \wedge r) \wedge \neg q = \neg p \wedge \neg q \wedge r.$$

Obviously, $(K \circ \mu) \wedge \varphi$ is consistent, so $R_5 - R_6$ would require that $(K \circ \mu) \wedge \varphi \equiv K \circ (\mu \wedge \varphi)$.

But now consider what happens when the agent revises independently by $\mu \wedge \varphi$: it has to give up its belief that p and that q and, because the agent follows this rule where it always revises more than is needed, it also gives up its belief that r. So we get:

$$K \circ (\mu \wedge \varphi) = \neg p \wedge \neg q \wedge \neg r.$$

Quite clearly, this violates R_5 and R_6—and, hopefully, it takes a step toward motivating their existence.

3.2.3 A representation result

Faithful assignments turn out to be intimately connected to revision operators, as the following theorem shows.

Proposition 2 ([31]). *A change operator $\circ\colon \mathcal{K} \times \mathcal{L} \to \mathcal{K}$ satisfies the revision postulates $R_1 - R_6$ if and only if there exists a faithful assignment mapping every knowledge base K to a total pre-order $\leq_K$ such that, for any formula μ:*

$$[K \circ \mu] = min_{\leq_K}[\mu].$$

This result is important and will provide a template for subsequent research. One can interpret it as showing that if an agent ranks its possible worlds in a faithful way and always picks the most preferred outcomes, then such an agent acts in accordance with postulates $R_1 - R_6$. Conversely, the result shows that an agent doing revision according to the postulates $R_1 - R_6$ behaves as if it always ranks possible outcomes according to the rules of faithful assignments.

Pragmatically, the Katsuno and Mendelzon representation result allows us to study revision operators from a model-theoretic perspective: finding operators that satisfy $R_1 - R_6$ becomes equivalent to finding faithful assignments. Some examples of such assignments are given below.

3.2.4 Revision operators

It turns out that a very general notion of distance between interpretations turns out to be enough to guarantee a faithful assignment.

Definition 4. A *non-symmetric pseudo-distance* between interpretations is a function $m\colon \mathcal{W} \times \mathcal{W} \to \mathbb{R}_+$ such that, for any $w_1, w_2 \in \mathcal{W}$:

$$m(w_1, w_2) = 0 \text{ if and only if } w_1 = w_2. \quad (minimality)$$

Both the drastic distance d_D and the Hamming distance d_H are non-symmetric pseudo-distances.

Given a non-symmetric pseudo-distance m one can define in a straightforward way $\leq_K$ for any knowledge base K. A revision operator follows immediately:

Definition 5. For a non-symmetric pseudo-distance m, knowledge base K and interpretation w, *the level of w with respect to K*, written as $l_K(w)$, is a positive number defined as follows:

$$l_K(w) = \begin{cases} min\{m(w, w') \mid w' \in [K]\}, & \text{if } K \text{ is consistent} \\ 1, & \text{otherwise.} \end{cases}$$

The relation $\leq_K^m$ is defined as:

$$w \leq_K^m w' \text{ iff } l_K(w) \leq l_K(w').$$

For a knowledge base K and formula μ, the operator $\circ^m$ is defined as:

$$[K \circ^m \mu] = min_{\leq_K^m}[\mu],$$

It is quite easy to see that $\leq_K^m$, thus defined, is a faithful pre-order.

Proposition 3. *If i is a non-symmetric pseudo-distance, the relation $\leq_K^i$ constructed as in Definition 5 is a total pre-order. It is also faithful.*

Proof. To show that $\leq_K$ is a total pre-order, we need to show that it is reflexive, transitive and total. This follows immediately from the fact that each $w \in \mathcal{W}$ is assigned a positive real, and the natural ordering of the reals is itself reflexive and transitive.

To show that $\leq_K$ is a also faithful, let us make a case distinction. First, suppose K is consistent. If $w_1, w_2 \in [K]$, then (by *minimality*):

$$l_K(w_1) = l_K(w_2) = 0,$$

so $w_1 \approx_K w_2$. This shows that km_1 holds. If $w_1 \in [K]$ and $w_2 \notin [K]$, we get that $l_K(w_1) = 0$ and $l_K(w_2) > 0$. To see why this latter statement holds, assume that $l_K(w_2) = 0$. Then there would have to be some $w' \in [K]$ such that:

$$m(w_2, w') = 0.$$

By *minimality*, it follows that $w_2 = w'$ and hence $w_2 \in [K]$. But this is a contradiction. This shows km_2.

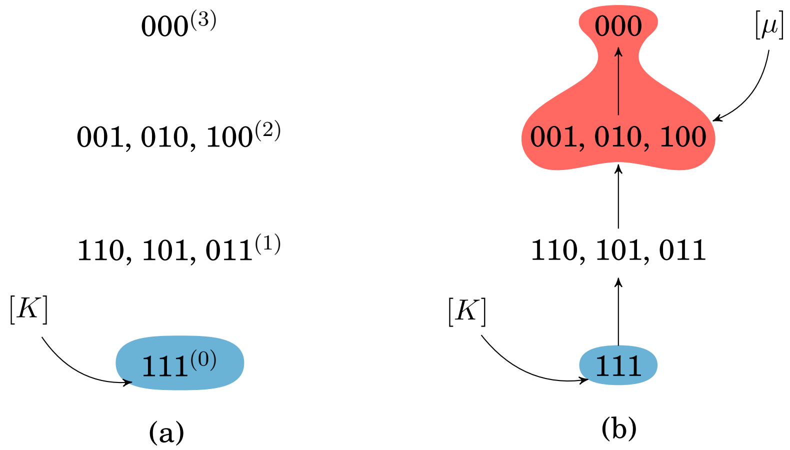

Figure 3.2: $\leq_K^{d_H}$ in action.

Since $\leq_K$ is constructed solely in terms of the models of K, the syntax of K does not matter, so km_3 holds.

Second, if K is inconsistent, all interpretations have the same level and they are all equal under $\approx_K$, hence all pre-orders under inconsistent knowledge bases are equal. This shows km_3. Since K has no models, km_1 and km_2 apply trivially. $\square$

Now that we know that $\leq_K^m$ is a faithful pre-order, the Katsuno and Mendelzon representation result (Proposition 2) guarantees that $\circ^m$ is a revision operator (i.e. satisfies postulates $R_1 - R_6$).

Let us look at a concrete case.

Example 3. Consider again Example 1, where we have the knowledge base K with $[K] = \{111\}$, and a formula μ with $[\mu] = \{000, 001, 010, 100\}$. In Example 1 the pre-order for K was hand-picked, under the assumption that it directly represented Anna's preferences. Now, by plugging in different distances in Definition 5, we get a different pre-order depending on the distance used.

When we use the Hamming distance d_H, then—since $[K]$ contains only 111—$l_K(w)$ will just be $d_H(w, 111)$, for any interpretation w. The levels assigned to each interpretation are shown in Figure 3.2-(a), and the pre-order generated is shown in Figure 3.2-(b).

Revising K by μ now gives us:

$$[K \circ^{d_H} \mu] = min_{\leq_K^{d_H}}[\mu] = \{001, 010, 100\},$$

in other words:

$$\begin{aligned} K \circ^{d_H} \mu \equiv \quad & (\neg\texttt{presentation} \land \neg\texttt{laundry} \land \texttt{market}) \lor \\ & (\neg\texttt{presentation} \land \texttt{laundry} \land \neg\texttt{market}) \lor \\ & (\texttt{presentation} \land \neg\texttt{laundry} \land \neg\texttt{market}). \end{aligned}$$

This means Anna will do exactly one of the activities, perhaps randomly choosing which one since she does not strictly prefer one over the others.

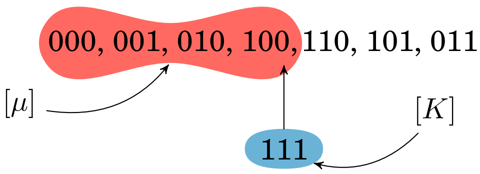

Figure 3.3: $\leq_K^{d_D}$ in action.

When we use the drastic distance d_D, we get that $l_K(111) = 0$ and $l_K(w) = 1$, for any interpretation other than 111. The order in this case is very simple (see Figure 3.3), and revision by μ yields:

$$K \circ^{d_D} \mu = min_{\leq_K^{d_D}}[\mu] = \{000, 001, 010, 100\}.$$

In this case, Anna will choose one of the activities or do nothing. □

It will be useful at this point to have a simple revision operator, one which we can invoke freely as an all-purpose tool. The operator associated with the assignment got using the drastic distance d_D seems simple enough, but it would be nice if we could express it in purely syntactic terms.

Definition 6. The *default assignment* is the assignment got by plugging in the drastic distance d_D in Definition 5.

By Definition 5, the default assignment gives a revision operator $\circ^{d_D}$. Let us introduce now another operator.

Definition 7. The *default operator* $\circ^{def} : \mathcal{K} \times \mathcal{L} \to \mathcal{K}$ is defined as:

$$K \circ^{def} \mu = \begin{cases} K \wedge \mu, \text{ if } K \wedge \mu \text{ is consistent,} \\ \mu, \text{ otherwise.} \end{cases}$$

It is straightforward to verify that the default operator satisfies postulates $R_1 - R_6$ and is, therefore, a revision operator. Intuitively, $\circ^{def}$ spells out a very simple revision strategy: if the new information μ is consistent with the knowledge base K, then simply add μ to K. If μ contradicts existing beliefs, then forget about them and just hold on to the new information.

Of course, the entire point of this is that $\circ^{d_D}$ and $\circ^{def}$ coincide.

Proposition 4. *The default revision operator $\circ^{def}$ is equivalent to $\circ^{d_D}$ (generated by the drastic distance), in the sense that:*

(i) *for any knowledge base K and formula μ, $K \circ^{def} \mu \equiv K \circ^{d_D} \mu$;*

(ii) *they correspond to the same faithful assignment (which we have called the default assignment).*

Proof. Let us start with (i). For a visual reminder of how $\circ^{d_D}$ assignment looks, see Example 3.

Take a knowledge base K and a formula μ. If $K \wedge \mu$ is consistent, then (by R_2), it follows that $K \circ^{d_D} \mu \equiv K \wedge \mu$. If $K \wedge \mu$ is inconsistent, then $[K \wedge \mu] = \emptyset$. Since $[K \circ^{d_D} \mu] = min_{\leq_K}[\mu]$ and the order $\leq_K$ has only two levels, it follows that $[K \circ^{d_D} \mu] = [\mu]$. So $K \circ^{d_D} \mu$ always behaves as the default operator.

Since $\circ^{def}$ is a revision operator, we know by the Katsuno and Mendelzon representation result (Proposition 2) that it corresponds to some faithful assignment. To show (ii), we need to show that $\circ^{def}$ can only correspond to an assignment that behaves as the one generated by the drastic distance. Consider, then, a knowledge base K and two interpretations w_1, w_2. As a reminder, φ_{w_1,w_2} is a propositional formula such that $[\varphi_{w_1,w_2}] = \{w_1, w_2\}$. In this setup $\leq_K$ is the faithful pre-order corresponding to $\circ^{def}$.

If $w_1, w_2 \in [K]$ or $w_1 \in [K], w_2 \notin [K]$, the situation is decided by $\leq_K$ being a faithful pre-order: we have that $w_1 \approx_K w_2$ or $w_1 <_K w_2$, respectively. The interesting case, then, is when $w_1, w_2 \notin [K]$. In this case $[K \wedge \varphi_{w_1,w_2}] = \emptyset$, so $K \wedge \varphi_{w_1,w_2}$ is inconsistent. We get that $K \circ^{def} \varphi_{w_1,w_2} = \varphi_{w_1,w_2}$, hence $min_{\leq_K}[\varphi_{w_1,w_2}] = [\varphi_{w_1,w_2}]$. This means that $w_1 \approx_K w_2$. Since w_1 and w_2 were arbitrarily chosen, it follows that all interpretations not in $[K]$ are equivalent with respect to $\leq_K$.[2] This shows that $\leq_K$ is the same pre-order as the one got using the drastic distance. $\square$

Other formulations

In the AGM framework presented in [1], the object of revision is a logical theory, and revision is defined using an additional belief change operator, called *expansion*.

Definition 8. The *expansion of a theory T by a formula μ*, written $T + \mu$, is defined as:

$$T + \mu := Cn(T \cup \{\mu\}).$$

Notice that the expansion of T by μ is by definition also a theory. Keeping this in mind, the revision of a theory T by μ, written $T * \mu$, is defined via the following postulates:

(K_1) $T * \mu = Cn(T * \mu)$.

(K_2) $\mu \in T * \mu$.

(K_3) $T * \mu \subseteq T + \mu$.

(K_4) If $\neg\mu \notin T$, then $T + \mu \subseteq T * \mu$.

(K_5) $T * \mu = \mathcal{L}$ only if $\nvdash \neg\mu$.

[2] The argument works in the same way when K is consistent or inconsistent.

(K_6) If $\vdash \mu_1 \leftrightarrow \mu_2$, then $T * \mu_1 = T * \mu_2$.

(K_7) $T * (\mu_1 \wedge \mu_2) \subseteq T * \mu_1 + \mu_2$.

(K_8) If $\neg\mu_2 \notin T * \mu_1$, then $T * \mu_1 + \mu_2 \subseteq T * (\mu_1 \wedge \mu_2)$.

Postulate K_1 ensures that the result of revision is also a theory. Other than this, there is a clear correspondence between the AGM postulates and the $R_1 - R_6$ postulates presented earlier. For instance, K_3 and K_4 taken together say that if T is consistent with μ (in the sense that $T \nvdash \neg\mu$), then $T * \mu = T + \mu$. The same idea is behind R_2: expansion, in a certain sense, is the set theoretic counterpart to conjunction.

The AGM paper goes on to explore more constructive ways of defining revision operators, the details of which should not concern us here. As a side note, one of their operators, defined from—what they call—the *full meet contraction function*, is essentially the same as the default operator introduced in Section 3.2.4.

The reader might have been struck by the very syntactic nature of AGM revision: there is virtually no semantics at play in it. This was a conscious decision: the original motivation for belief change, Gärdenfors writes, was to model transitions between epistemic states of an agent—what goes on, so to speak, in the agent's head when it changes its mind. From this perspective, the relationship between an agent's beliefs and reality (what semantics is purportedly about) does not matter, since an agent cannot go outside itself and inspect the relationship between its beliefs and the world:

> From the subject's point of view there is no way to tell whether she *accepts* something as knowledge, that is, has full belief in it, or whether her accepted knowledge is also *true*. When I speak of particular pieces of knowledge, I use the word "knowledge" in the sense of full belief, and I do not assume that this concept entails anything about the truth of the beliefs. (Gärdenfors [1988], p. 20)

This distrust for semantics, however, leads to tensions in other parts of the framework: to preserve the idea that agents are rational, the AGM authors force the agent to incorporate every consequence of its (explicit) beliefs. This leads to the object of revision being whole theories, which—as infinite sets—put the whole theory on a rather abstract footing. To make revision computationally feasible, epistemic states must in the least be scaled down to finite representations, which is why we go with the Katsuno and Mendelzon framework of a finite alphabet.

As for the incursion into semantics, this is perhaps unavoidable. What could be said to meet Gärdenfors' qualms is that, at least insofar as computational systems are concerned, semantics does not commit one to any definitive view about reality; it merely offers an alternative way of representing the part of reality we are interested in modelling.

Leaving the philosophical issues aside, it should be noted that there is a neat correspondence between revision operators $\circ$ and AGM revision operators $*$, as the following lemma shows.

Proposition 5 ([31]). *If T is a theory that is represented by a knowledge base K such that $T = \{\varphi \mid K \models \varphi\}$, μ is a propositional formula, $*$ is an AGM revision operator and $\circ$ its corresponding operator on knowledge bases, we have that $*$ satisfies $K_1 - K_8$ if and only if $\circ$ satisfies $R_1 - R_6$.*

Beside the formalisms presented here, there are other ways of thinking about revision, all of which turn out to be equivalent. One should mention Grove's System of Spheres and Epistemic Entrenchments, where the idea is to assign different priorities to formulas of a language: the more entrenched a formula is, the more resistant it is to revision. For nice reviews, see [43, 48]. As a side note, all these approaches end up assuming some sort of ranking over the objects of focus.

3.2.5 Other belief change operations

Section 3.2.4 already mentioned expansion as a belief change operator alongside revision. One need only add *contraction* to complete the picture.

Contraction, written $T-\varphi$, was conceived as the operation of removing φ from the theory T—recall the problem of derogating an item from a legal code mentioned at the beginning of the chapter. The AGM paper [1] featured postulates for contraction alongside those of revision, amd ever since then it has become customary to think of revision and contraction as inter-definable via the following set of identities:

$$T * \varphi := (T - \neg\varphi) + \varphi \qquad (\textbf{Levi identity})$$
$$T - \varphi := (T * \neg\varphi) \cap T \qquad (\textbf{Harper identity})$$

The AGM paper showed that if an operator $-$ satisfies the contraction postulates, then $*$ defined from $-$ with the Levi identity satisfies the revision postulates. Conversely, an operator $-$ defined from a revision operator $*$ with the Harper identity satisfies the contraction postulates.

This is not just a nice result, but another means of building revision operators: the AGM paper, in fact, always started from contraction and built its way up to revision—one reason behind this may be that it is in some sense easier to conceptualize contraction than revision. The semantic approach of Katsuno and Mendelzon opened up a more intuitive way of thinking directly about revision, by looking at how a revision operator acts on sets of interpretations (rather than on sets of formulas). Also, of all the belief change operators, revision has the strongest connections to merging, which is why it enjoys more attention here.

3.3 Merging: some history and motivation

Early work on merging ([3, 4]) came out of concerns about databases: we may want to combine several databases into a single, consistent database, so what are the general principles of doing this? What happens when we reach an inconsistency: which parts do we keep, which do we throw out? What if we also add an integrity constraint? In [3], Baral, Kraus and Minker considered four possible answers:

1. There is an Oracle that knows everything: if a contradiction arises, the Oracle decides what goes in the result of the merging.
2. There is some partial order on database items that allows us to choose when a conflict appears.
3. If there is a conflict between two databases over some atom p, define p's truth-value as *unknown*.
4. Choose a maximal amount of consistent information from the combined databases.

Their professional diagnosis regarding these four solutions was that 1 is impossible, 3 leads to loss of information, and 2 might be untenable because there is no access to priority information. They eventually settled for the 4^{th} alternative.[3]

In [46], Peter Revesz noted that there was a connection between combining several sources of information (which he called *arbitration*) and revision: Revesz thought of arbitration in the framework introduced by Katsuno and Mendelzon for revision. This meant coming up with postulates for an operator (in a finite propositional language), and describing the semantics of the operator through a representation result. Revesz managed to show that his arbitration operator $\triangleright$ corresponded to a *loyal assignment* on interpretations.

Though we will not go into the details of Revesz' work, one of his postulates deserves attention:

$$(M_7)\ \ (K_1 \triangleright \mu) \wedge (K_2 \triangleright \mu) \models (K_1 \cup K_2) \triangleright \mu.$$

M_7 is a postulate that tells us how $\triangleright$ should behave when two knowledge bases (K_1 and K_2) are combined into a single knowledge base ($K_1 \cup K_2$). We never saw this with revision.

In the end, though, the formalism that really caught on was the one proposed by Konieczny and Pino Pérez. An early version ([34]) featured postulates for a generic merging operator Δ, and introduced the notion of majority and arbitration operators. Constraints were added in a later version ([35]). Konieczny and Pino Pérez's idea was to encode a notion of *fairness* in the postulates, to ensure that merging does not favour one source over the others. The postulates were shown to correspond to syncretic assignments (see 3.4.1).

[3] For a very general overview of areas where multiple streams of information have to be combined into a single, coherent stream, see [5].

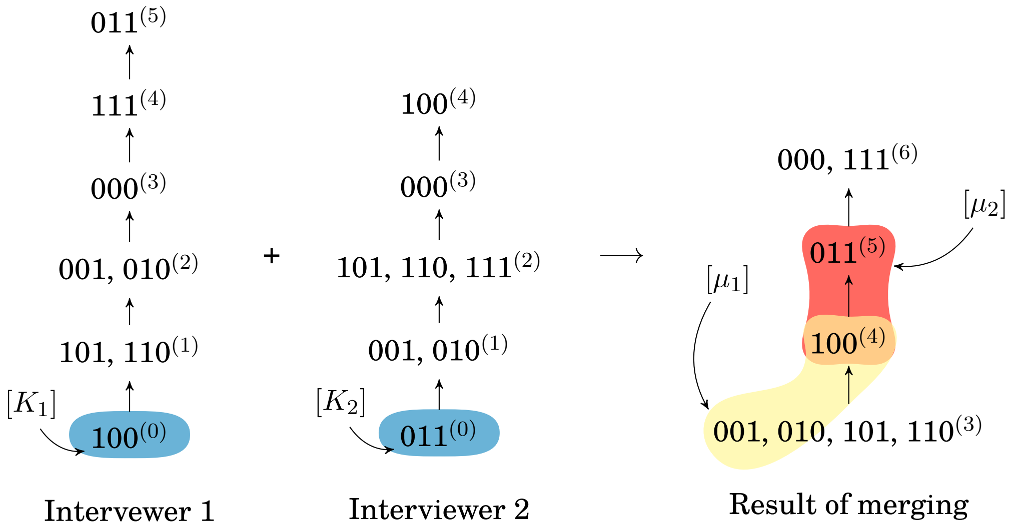

Figure 3.4: Preference relations for the two interviewers, and the ranking that results from a Borda count method of voting.

This work follows the Konieczny and Pino Pérez framework for merging. Let us motivate it further with an example.

Example 4. Two interviewers have to decide who to hire from a range of three applicants. After screening the candidates, each interviewer arrives at their own opinion about who should be hired. Unfortunately, their assessments differ.

We can represent the problem logically by having atoms p_1, p_2 and p_3 stand for the applicants, and knowledge bases K_1 and K_2 expressing the interviewers' conclusions:

$$K_1 = \{p_1, \neg p_2, \neg p_3\},$$
$$K_2 = \{\neg p_1, p_2, p_3\}.$$

In other words, Interviewer 1 wants to hire only applicant 1, while Interviewer 2 wants to hire applicants 2 and 3. Clearly, there is no common ground between the two interviewers: in logical terms, $\bigwedge K_1 \wedge \bigwedge K_2$ is inconsistent.

Here is, then, a situation where multiple agents have to settle on a common course of action, even though they may have competing views. As with revision, we will want to assume that possible worlds are ranked with respect to how much they are preferred. The additional twist is that, since we are now dealing with multiple agents, there needs to be a preference relation for each of them. In our example, we will assume that each interviewer ranks the possible options of which candidates to hire as in Figure 3.4.

Later on we will study general methods for ranking possible worlds, but in Figure 3.4 we assume that the rankings reflect directly the interviewers' preferences.

Notice, for instance, that Interviewer 1 has a strong preference for applicant 1 and an equally strong dislike for the other two applicants: she would rather see no one hired than have applicants 2 and 3 working together. The situation is reversed for Interviewer 2: she would like applicants 2 and 3 to be hired and would prefer leaving the positions unfilled rather than have candidate 1 on the job.

So who will get hired? This is the challenge that merging faces: as with revision, the aim is to describe a general process that ultimately results in a concrete recommendation, and which—at the same time—strives to be fair toward the parties concerned. More specifically, the individual preferences need to be aggregated into a single preference relation over possible worlds, while making sure that certain fairness constraints are met.

In our example, we can imagine that the interviewers settle the issue by voting on the possible outcomes: each interviewer assigns a score to every possible world according to its place in the preference ranking (see Figure 3.4), then they add the scores for each world and order the worlds according to their total score. The resulting ranking is shown in Figure 3.4. The method of voting is essentially a Borda count, with the only caveat that in our example worlds are more preferred if they have *lower* scores.

The advantage of representing the result as a total preference ranking on possible worlds is the fact that we can add constraints on the result of merging. Suppose the employer says that, because of budgetary reasons, exactly *one* candidate can be hired. This constrains the range of options to μ_1, where:

$$[\mu_1] = \{001, 010, 100\}.$$

Given the preference ranking, we can just choose the most preferred outcomes from the models of μ_1, that is:

$$min[\mu_1] = \{001, 010\}.$$

So either the second or third applicant should be hired.

Suppose, on the other hand, there is a different constraint: the final choice must be one of the initial first choices of the two interviewers (encoded by the constraint $[\mu_2] = \{100, 011\}$). Then the preferred candidate would be applicant 1:

$$min[\mu_2] = \{100\}.$$

Applicant 1 was interviewer 1's first choice, so it appears the final ranking favours Interviewer 1 over Interviewer 2. There is a sense in which this is unfair, since the fact that 100 has a lower score than 011 is just an artefact of the way in which Interviewer 1 structured her preferences, and not some impartial arbitration procedure. This leads to the possibility that, in the absence of any constraints on the merging process, the result might end up favouring one agent over the other, and that agents could ultimately game the system in their favour.

It could be argued, then, that a fair merging procedure should ensure that none of the agents is given precedence over the others when the final ranking is computed. The general problem of merging ends up being tricky, precisely for this reason. □

We are left with the following questions: how should individual preferences be aggregated such that we get a unique preference relation? Under which conditions is the merging process fair?

The following subsections address these questions by looking at the problem from two perspectives: on the semantic side we look at ways of ranking interpretations, much like we did for revision with faithful assignments. For merging, however, we need to add certain conditions to ensure fairness. On the syntactic side we look at postulates that a merging operator should satisfy: again, the revision postulates need to be supplemented such that some notion of fairness is enforced. The two perspectives are then shown to coincide through the representation result obtained in [35].

3.4 Merging: some theory

We start by defining merging operators.

Definition 9. A *merging operator* is a function $\Delta\colon \mathcal{E} \times \mathcal{L} \to \mathcal{K}$.
We typically write $\Delta_\mu(E)$ instead of $\Delta(E, \mu)$.

In other words, merging operators map profiles and formulas (called *constraints*) onto knowledge bases. Intuitively, each knowledge base in a profile E represents the beliefs or goals of an agent.

3.4.1 Syncretic assignments

We give here the main semantic construction for merging: a syncretic assignment tells how the rankings for a profile should look like. The assignment should be fair when merging requires aggregating competing choices, and approximate a group consensus when something like that exists.

Definition 10. A *syncretic assignment* is a mapping assigning to each profile E a total pre-order $\leq_E$ on $\mathcal{W}$ such that, for any profiles E, E_1, E_2, knowledge bases K_1, K_2 and interpretations w_1, w_2:

- (kp_1) if $w_1 \in [E]$ and $w_2 \in [E]$, then $w_1 \simeq_E w_2$;
- (kp_2) if $w_1 \in [E]$ and $w_2 \notin [E]$, then $w_1 <_E w_2$;
- (kp_3) if $E_1 \equiv E_2$, then $\leq_{E_1} = \leq_{E_2}$;
- (kp_4) if $w_1 \in [K_1]$, there is $w_2 \in [K_2]$ such that $w_2 \leq_{\{K_1,K_2\}} w_1$;
- (kp_5) if $\begin{cases} w_1 \leq_{E_1} w_2, \\ w_1 \leq_{E_2} w_2, \end{cases}$ then $w_1 \leq_{E_1 \sqcup E_2} w_2$;
- (kp_6) if $\begin{cases} w_1 \leq_{E_1} w_2, \\ w_1 <_{E_2} w_2, \end{cases}$ then $w_1 <_{E_1 \sqcup E_2} w_2$.

$$\begin{array}{ccccc} w' & & w'' & & \\ \uparrow & & \uparrow & & w \\ w & + & w & \longrightarrow & \uparrow \\ \uparrow & & \uparrow & & w', w'' \\ w'' & & w' & & \\ & & & & \\ \leq_{\{K_1\}} & & \leq_{\{K_2\}} & & \leq_{\{K_1,K_2\}} \end{array}$$

Figure 3.5: Property kp_8 eliminates this outcome.

Additionally, we may want to consider the following properties:

(kp_7) If $w_1 <_{E_2} w_2$, then there is an $n \in \mathbb{N}$ such that $w_1 <_{E_1 \sqcup E_2^n} w_2$.

(kp_8) If $\begin{cases} w <_{\{K_1\}} w', \\ w <_{\{K_2\}} w'', \\ w' \approx_{\{K_1,K_2\}} w'', \end{cases}$ then $w <_{\{K_1,K_2\}} w'$.

What are these properties about? Notice that $kp_1 - kp_3$ are the conditions for faithful assignments: any pre-order $\leq_E$ that satisfies them is faithful, and hence has a corresponding revision operator.

Property kp_4 ensures fairness: if it would be false, then there would be a model of K_1 strictly smaller than any model of K_2 in $\leq_{\{K_1,K_2\}}$—some state of affairs consistent with K_1 would be preferred to anything K_2 thinks is desirable (think back to μ_2 in Example 4). We want to avoid this.

Property kp_5 says that if everyone prefers w_1 to w_2, this choice should be reflected in the final ranking. Property kp_6 sees to it that if someone has a strong preference for w_1 over w_2 and no one else disagrees with this, then the final ranking should make sure that w_1 is strictly preferred to w_2.

In kp_7, one wants to give some decision power to a majority: kp_7 says that if w_1 is strictly preferred to w_2 by E_2, then under an enough number of agents with the same preference structure, their preference prevails.

Property kp_8 only makes sense when $w'' <_{K_1} w <_{K_1} w'$ and $w' <_{K_2} w <_{K_2} w''$ (for all other arrangements, the property follows from kp_6). And in this situation, we want to eliminate the outcome of Figure 3.5. The reasoning behind this, if we look at Figure 3.5, is that in profile $\{K_1\}$ w'' is strongly preferred to w', with w somewhere in the middle. In the profile $\{K_2\}$ the situation is reversed—there is a real conflict here, in which two agents hold opposing views. So if in the merged order $\leq_{\{K_1,K_2\}}$ there is nothing to choose between w', w'' (they are equivalent), the order should favour the 'middle ground' between w' and w'', which in this case happens to be w.

Definition 11. An assignment that satisfies $kp_1 - kp_6 + kp_7$ is called a *majority syncretic assignment*. An assignment that satisfies $kp_1 - kp_6 + kp_8$ is called a *fair syncretic assignment*.

3.4.2 Postulates for merging

As with revision, we want to constrain the behaviour of a merging operator such that it conforms with certain standards of rationality. The following postulates have been proposed [35]:

(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 and K_2 are consistent, and $K_1 \models \mu$ 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 $\Delta_{\mu_1 \wedge \mu_2}(E) \models \Delta_{\mu_1}(E) \wedge \mu_2$.

Additionaly, we may want to consider the following postulates:

(Maj) There is $n \in \mathbb{N}$ such that $\Delta_\mu(E_1 \sqcup E_2^n) \models \Delta_\mu(E_2)$.

(Arb) If $\begin{cases} \Delta_{\mu_1}(K_1) \equiv \Delta_{\mu_2}(K_2), \\ \Delta_{\mu_1 \leftrightarrow \neg \mu_2}(\{K_1, K_2\}) \equiv (\mu_1 \leftrightarrow \neg \mu_2), \\ \mu_1 \not\models \mu_2, \mu_2 \not\models \mu_1, \end{cases}$ then $\Delta_{\mu_1 \vee \mu_2}(\{K_1, K_2\}) \equiv \Delta_{\mu_1}(K_1)$.

Definition 12. An operator that satisfies postulates $IC_0 - IC_8$ is called *an IC merging operator*. An operator that satisfies $IC_0 - IC_8 + Maj$ is called *a majority merging operator*. An operator that satisfies $IC_0 - IC_8 + Arb$ is called *an arbitration merging operator*.

Notice that if we replace $\Delta_\mu(E)$ with $E \circ \mu$, $IC_0 - IC_3 + IC_7 - IC_8$ become the $R_1 - R_6$ postulates for revision in Section 3.2.2. This is intentional, as merging is modelled to be an extension of the revision framework. The point is that if an operator satisfies $IC_0 - IC_3 + IC_7 - IC_8$, there automatically exists by the Katsuno and Mendelzon representation result a 'faithful' assignment on profiles, that is to say there exists an assignment that satisfies $kp_1 - kp_3$. We are fudging some notions here, since profiles (which merging is concerned about) should be distinguished from knowledge bases (which is what revision works with). However, when we switch to interpretations and how they are ranked this distinction does not really matter. Hence, any IC merging operator is also—from the semantic point of view—a revision operator

Postulates $IC_4, IC_5, IC_6, Maj, Arb$ are meant to capture properties $kp_4, kp_5, kp_6, kp_7, kp_8$ of syncretic assignments, and are best understood as their syntactic counterparts.

Observation 1. The condition that K_1 and K_2 need to be consistent in IC_4 is usually not explicitly included, though Konieczny and Pino Pérez mention separately in [35] that the postulate does not make sense unless it holds. I have chosen to write the postulate in this way because we want to be able to think of merging even when some of the knowledge bases are inconsistent (though that may not be of particular interest). Adding a separate clause saying that we are only going to deal with consistent knowledge bases might only create confusion, if it would be retracted later.

Observation 2. IC_8 is sometimes written in a more economical way as $\Delta_{\mu_1 \wedge \mu_2}(E) \models \Delta_{\mu_2}(E)$, since we already have that $\Delta_{\mu_1 \wedge \mu_2}(E) \models \mu_2$ by IC_0. I have chosen to write it like this for reasons of symmetry.

3.4.3 A representation result

As with revision, it can be shown that merging operators reflect the semantic structure of syncretic assignments, via the familiar strategy of always picking the minimal models of a constraint.

Proposition 6 ([35]). *An operator Δ satisfies the merging postulates $IC_0 - IC_8$ if and only if there exists a syncretic assignment mapping each profile E to a total pre-order $\leq_E$ such that, for any formula μ:*

$$[\Delta_\mu(E)] = min_{\leq_E}[\mu].$$

The pragmatic takeaway of this representation result is that we can think of merging as an operation on pre-orders, essentially one that aggregates different pre-orders into a single pre-order under some fairness constraints (this is spelled out more in subsequent sections). As such, merging begins to look very much like social choice theory (and voting, in particular), where the issue is to aggregate preference rankings as effectively as possible. The connections between merging and social choice have been touched upon and are subject to ongoing research (see [21, 23, 36]).[4]

It was mentioned in the previous sections that merging operators are also revision operators. What merging adds to the revision framework is a set of constraints about fairness and optimality: viewed from the semantic side, a pre-order $\leq_E$ is not independent from the other pre-orders in the assignment (as it is in revision), but depends on the pre-orders $\leq_{E_1}$, $\leq_{E_2}$, if $E_1 \sqcup E_2$. A particular consequence of this is that if an assignment is not fit for revision (i.e., it is not faithful), it will not be fit for merging also. This issue will come up in the case of Horn merging.

With regard to postulate kp_4, it is worth mentioning that IC_4 is relevant for enforcing it only when $\{K_1, K_2\}$ is inconsistent. If $\{K_1, K_2\}$ is consistent, then kp_4 comes 'for free', in the sense that we do not need a special postulate to make sure that it holds.

[4]The word 'profile' itself is taken from social choice theory. What is different in the case of merging is that agents are invited to rank possible worlds, or interpretations. In social choice theory, preference rankings are on entities unconnected to each other.

Lemma 1. *If $\{K_1, K_2\}$ is consistent, then in an assignment that assigns to each profile E a total pre-order $\leq_E$ such that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, if $kp_1 - kp_3$ are true, then kp_4 is also true.*

Proof. Take an assignment that satisfies the above conditions and $\leq_{\{K_1,K_2\}}$ the total pre-order it assigns to $\{K_1, K_2\}$. So, now, for any $w_1 \in [K_1]$ we can take a $w_2 \in [K_1 \wedge K_2]$ (which, by assumption, is non-empty). Obviously, $w_2 \in [K_2]$. By kp_1 and kp_2, $[K_1 \wedge K_2]$ contains the minimal elements of $\leq_{\{K_1,K_2\}}$, and this gives us that $w_2 \leq_{\{K_1,K_2\}} w_1$. □

Lemma 1 will come in handy in the context of Horn merging, where we will be forced to reconsider the kp_4.

3.4.4 Merging operators

By the Konieczny and Pino Pérez representation result (Proposition 6) we can produce look at merging as a logical operation and as a semantic operation that aggregates pre-orders on interpretations. In this section we define classes of merging operators from the semantic side, by defining a series of syncretic assignments.

As with faithful assignments for revision, a good starting point would be to use some notion of distance between interpretations in order to define the level of an interpretation with respect to a knowledge base K. This gives us a pre-order $\leq_K$. The next step is a strategy for aggregating different pre-orders. In doing this we need to keep in mind the constraints set by $kp_4 - kp_6$. To see how this adds to the revision framework, consider, for example, a profile $E = \{K_1, K_2, K_3\}$. We can split E in various ways:

$$E = \{K_1, K_2\} \sqcup \{K_3\} = \{K_1\} \sqcup \{K_2, K_3\} = \{K_1, K_3\} \sqcup \{K_2\}.$$

Consequently, $\leq_E$ depends on the interactions between $\leq_{\{K_1,K_2\}}$ and $\leq_{\{K_3\}}$, $\leq_{\{K_1\}}$ and $\leq_{\{K_2,K_3\}}$ and so on. This makes things trickier, as faithful assignments need to satisfy extra properties (i.e., $kp_4 - kp_6$).

In [35, 37] it was shown that, as long as the distance and aggregation functions satisfy some fairly general conditions, the assignment constructed from them will be syncretic.

Definition 13. A *pseudo-distance* between interpretations is a function $d\colon \mathcal{W}\times\mathcal{W} \to \mathbb{R}_+$ such that, for any $w_1, w_2 \in \mathcal{W}$:

$$d(w_1, w_2) = d(w_2, w_1) \qquad (symmetry)$$
$$d(w_1, w_2) = 0 \text{ if and only if } w_1 = w_2. \qquad (minimality)$$

A pseudo-distance extends the notion of a non-symmetric pseudo-distance that was required for defining faithful assignments (see Definition 4), by adding symmetry. The familiar distances d_H (the Hamming distance) and d_D (the drastic distance) are also pseudo-distances, in the sense defined here.

With pseudo-distances we can rank interpretations with respect to their degree of 'closeness', and create pre-orders $\leq_K$ for knowledge bases K. The next step is to aggregate these pre-orders.

Definition 14. An *aggregation function* is a function $f\colon \mathbb{R}_+^n \to \mathbb{R}_+$, for $n \in \mathbb{N}$, such that, for any $x_1, \ldots, x_n, x, y \in \mathbb{R}_+$ and any permutation σ, the following conditions hold:

· if $x \leq y$, then $f(x_1, \ldots, x, \ldots, x_n) \leq f(x_1, \ldots, y, \ldots, x_n)$; $(monotony)$

· $f(x_1, \ldots, x_n) = 0$ if and only if $x_1 = \cdots = x_n = 0$; $(minimality)$

· $f(x) = x$; $(identity)$

· $f(x_1, \ldots, x_n) = f(\sigma(x_1), \ldots, \sigma(x_n))$; $(symmetry)$

· if $f(x_1, \ldots, x_n) \leq f(y_1, \ldots, y_n)$, then $f(x_1, \ldots, x_n, z) \leq f(y_1, \ldots, y_n, z)$; $(composition)$

· if $f(x_1, \ldots, x_n, z) \leq f(y_1, \ldots, y_n, z)$, then $f(x_1, \ldots, x_n) \leq f(y_1, \ldots, y_n)$. $(decomposition)$

Putting pseudo-distances and aggregation functions together, we can define a general class of merging operators.

Definition 15. Take a pseudo-distance d between interpretations and an aggregation function f. If w is an interpretation and K is a consistent knowledge base, define *the distance between w and K* as:

$$d(w, K) = min\{d(w, w') \mid w' \in [K]\}.$$

If $E = \{K_1, \ldots, K_n\}$ is a profile where each $K \in E$ is consistent,[5] define *the level of w with respect to E* as:

$$l_E(w) = f(d(w, K_1), \ldots, d(w, K_n)).$$

Define $\leq_E$ as:

$$w \leq_E w' \text{ if and only if } l_E(w) \leq l_E(w').$$

Define an operator $\Delta^{d,f}$ as:

$$[\Delta^{d,f}(E)] = min_{\leq_E}[\mu].$$

The idea, therefore, is that the pre-order $\leq_E$ is computed from pre-orders for the component knowledge bases K by aggregating the levels of each interpretation, base-wise. It turns out that this is a sound merging strategy.

Proposition 7. *An operator $\Delta^{d,f}$ built using a pseudo-distance d and an aggregation function f as in Definition 15 satisfies $IC_0 - IC_8$.*

[5]There is no guarantee, however, that E will be consistent.

Observation 3. We have put this result in a form that will be easy to use later, but actually a stronger version holds. An aggregation function f that satisfies $monotony$, $minimality$ and $identity$ gives a merging operator $\Delta^{d,f}$ that satisfies $IC_0 - IC_2, IC_7 - IC_8$. $\Delta^{d,f}$ satisfies $IC_0 - IC_8$ if and only if f also satisfies $symmetry$, $composition$ and $decomposition$. See [33], or Theorems 8-9 in [37].

The problem now hinges on choosing an appropriate pseudo-distance d and an aggregation function f. It is easy to see, for instance, that the simple sum Σ satisfies all the properties of an aggregation function. We can also use the GMAX function, which–though strictly speaking not an aggregation function (it does not fit the definition)—can be used to define $\leq_E$ for a profile E. The details are given below.

Definition 16. For a profile $E = \{K_1, \ldots, K_n\}$ and an interpretation w, take $L_E(w)$ to be a list of all the distances $d(w, K_i)$, for $i \in \{1, \ldots, n\}$, written in descending order. In other words:

$$L_E(w) = (d(w, K_{i_1}), \ldots, d(w, K_{i_n})),$$

where $\{K_{i_1}, \ldots, K_{i_n}\} = \{K_1, \ldots, K_n\}$ and $d(w, K_{i_j}) \geq d(w, K_{i_{j+1}})$, for $j \in \{1, \ldots, n-1\}$.

Define the aggregation function GMAX as:

$$\mathrm{GMAX}(d(w, K_1), \ldots, d(w, K_n)) = L_E(w),$$

and use the lexicographic order $\leq_{lex}$ to compare interpretations. In other words:

$$w_1 \leq_E w_2 \text{ iff } L_E(w_1) \leq_{lex} L_E(w_2),$$

The following example shows how GMAX works to rank interpretations.

Example 5. Suppose $E = \{K_1, K_2, K_3\}$, d is some pseudo-distance and w, w' are interpretations such that:

$$\begin{array}{lll} d(w, K_1) = 2, & d(w, K_2) = 0, & d(w, K_3) = 3, \\ d(w', K_1) = 3, & d(w', K_2) = 1, & d(w', K_3) = 2. \end{array}$$

Then $L_E(w) = (3, 2, 0), L_E(w') = (3, 2, 1)$. We have that $L_E(w) \leq_{lex} L_E(w')$, hence $w \leq_E w'$. □

It turns out that if we plug into Definitions 15, 16 familiar distances, such as the Hamming distance d_H or the drastic distance d_D, we get operators that satisfy the IC merging postulates.

Proposition 8 ([35])**.** *Take a pseudo-distance d between interpretations. Then $\Delta^{d,\Sigma}$ and $\Delta^{d,GMAX}$ are IC-merging operators.*

Additionally:

- *$\Delta^{d,\Sigma}$ is a majority merging operator;*
- *$\Delta^{d,GMAX}$ is an arbitration merging operator.*

	$[K_i]$			
	100	011	Σ	GMAX
000	1	1	2	(1,1)
001	1	1	2	(1,1)
010	1	1	2	(1,1)
011	1	0	1	(1,0)
100	0	1	**1**	**(1,0)**
101	1	1	2	(1,1)
110	1	1	2	(1,1)
111	1	1	2	(1,1)

(a) Drastic distance d_D

	$[K_i]$			
	100	011	Σ	GMAX
000	1	2	3	(2,1)
001	2	1	**3**	**(2,1)**
010	2	1	**3**	**(2,1)**
011	3	0	3	(3,0)
100	0	3	**3**	(3,0)
101	1	2	3	(2,1)
110	1	2	3	(2,1)
111	2	1	3	(2,1)

(b) Hamming distance d_H

Figure 3.6: Merging K_1, K_2, K_3 under a constraint μ, different distances and different aggregation functions

Let us see how these operators work in a concrete example.

Example 6. Consider again the setup in Example 4, with knowledge bases K_1 and K_2 such that $[K_1] = \{100\}$ and $[K_2] = \{011\}$. Take the profile $E = \{K_1, K_2\}$ and the constraint μ with $[\mu] = \{001, 010, 100\}$.

There are various possibilities for how the merging should be computed: we can use the drastic distance d_D or the Hamming distance d_H as a pseudo-distance, in combination with either Σ or GMAX as an aggregation function. Figure 3.6 shows what we get in every case: the entries in the table under the models of K_i, $i \in \{1, 2\}$, give $d(w, K_i)$, for every interpretation $w \in \mathcal{W}$. We then aggregate these distances to get the level of each w in $\leq_E$, using an aggregation function $f \in \{\Sigma, \mathrm{GMAX}\}$. The level for each interpretation is shown under the corresponding function symbol.

The models of the constraint μ are shown in red—$\Delta_\mu^{d,f}(E)$ is then computed by choosing the minimal elements among them, written in bold font. Notice that we get different results, depending on the parameters chosen. For instance:

$$[\Delta_\mu^{d_D,\Sigma}(E)] = [\Delta_\mu^{d_D,\mathrm{GMAX}}(E)] = \{100\},$$

which means that $\Delta_\mu^{d_D,\Sigma}(E) \equiv \Delta_\mu^{d_D,\mathrm{GMAX}}(E) \equiv p_1 \wedge \neg p_2 \wedge \neg p_3$. However:

$$[\Delta_\mu^{d_H,\mathrm{GMAX}}(E)] = \{001, 010\},$$

which means that $\Delta_\mu^{d_H,\mathrm{GMAX}}(E) \equiv (\neg p_1 \wedge \neg p_2 \wedge p_3) \vee (\neg p_1 \wedge p_2 \wedge \neg p_3)$.

For an example in a larger alphabet, see [37], pp. 248-249. □

Observation 4. Definition 15 and the examples so far have all assumed that the knowledge bases making up a profile are all consistent. What happens if this is not the case? Though it may be debated whether it is of practical interest to perform merging when some of the knowledge bases are inconsistent, an approach for doing so has been outlined in [33].

The idea, roughly, is to introduce an extra aggregation step before computing the distance between an interpretation and a knowledge base. If $K = \{\varphi_1, \dots, \varphi_n\}$ is a knowledge base, all the φ's are propositional formulas , d is a pseudo-distance, w is an interpretation and $g\colon \mathbb{R}^n \to \mathbb{R}_+$ is the extra aggregation function, define $d(w, K)$ as:

$$d(w, K) = g(d(w, \varphi_1), \dots, d(w, \varphi_n)).$$

It is understood that $d(w, \varphi)$, where φ is a propositional formula, is defined as:

$$d(w, \varphi) = min\{d(w, w') \mid w' \in [\varphi]\}.$$

Once this is done, $d(w, E)$ and $\leq_E$ are constructed in Definition 15. An operator defined in this way is called a DA^2 ***merging operator***, and is denoted $\Delta^{d,f,g}$.[6]

The innovation brought by DA^2 operators is that $d(w, K)$ depends on the formulas that make up K rather than the models of K, which makes it possible—as mentioned in the beginning—to work with inconsistent knowledge bases.

The drawback is that many combinations of intuitive aggregation functions produce operators that do not satisfy all the IC postulates. As with the simpler distance based operators, the authors of [33] find some conditions on d, f and g under which $\Delta^{d,f,g}$ are IC merging operators.

One more thing deserves to be mentioned here. It can hardly escape notice that merging has deep connections with revision, indeed that it tightens the revision framework to cover interactions between knowledge bases. One way of seeing this is via the types of assignments captured by the revision and merging functions, respectively. Revision, by talking about faithful assignments, treats every pre-order $\leq_K$ individually: what happens in a pre-order $\leq_{K_1}$ does not depend on what happens in a neighbouring pre-order $\leq_{K_2}$ (unless K_1 and K_2 are logically equivalent). In merging, however, pre-orders $\leq_E$ are not independent of pre-orders $\leq_{E_1}, \leq_{E_2}$, if $E_1 \sqcup E_2 = E$. Merging, in this sense, places stricter conditions on assignments than revision, in order to assure that the assignments respect certain fairness conditions.

It would be natural, then, to ask if we can use revision operators to construct merging operators. A series of nice results regarding this appear already in [35], and they are presented below.

We know, by the Katsuno and Mendelzon representation result (Proposition 2), that for any revision operator $\circ$ we can find a faithful assignment which assigns to any knowledge base K a pre-order $\leq_K^\circ$ such that $[K \circ \mu] = min_{\leq_K^\circ}[\mu]$. The idea, then, is to start from faithful assignments and use them to build up syncretic assignments.

[6] According to [33], DA stands for ***D**istance-based **A**ggregation operator*, and 2 refers to the number of aggregation functions used.

Definition 17. Take a revision operator $\circ$. For any knowledge base K, $\leq_K^\circ$ is the pre-order assigned to K by the corresponding faithful assignment.

For any interpretation w, define the *the level of w with respect to K*, written $l_K^\circ(w)$, to be the height at which w appears in $\leq_K$. More formally, $l_K^\circ(w)$ is the length of the longest chain:

$$w_0 <_K \cdots <_K w_n,$$

where $w_0 \in [K]$ and $w_n = w$.

If $E = \{K_1, \ldots, K_n\}$ is a profile and f is an aggregation function, define *the level of w with respect to E* to be:

$$l_E^\circ(w) = f(l_{K_1}^\circ(w), \ldots, l_{K_2}^\circ(w)).$$

Define the relation $\leq_E$ as:

$$w_1 \leq_E w_2 \text{ iff } l_E^\circ(w_1) \leq l_E^\circ(w_2).$$

Define an operator Δ° as:

$$\Delta_\mu^\circ(E) = min_{\leq_E}[\mu].$$

The aggregation function f can be something familiar, like the sum Σ or the GMAX function. To use the GMAX function, the assignment is defined a bit differently but the intuition is the same: for a profile $E = \{K_1, \ldots, K_n\}$ and an interpretation w, create a vector $(l_{K_1}(w), \ldots, l_{K_n}(w))$ and rearrange it such that its elements appear in descending order, from left to right. Denote the ordered vector as $L_E(w)$. The lexicographic order of these vectors determines the order $\leq_E$, i.e.:[7]

$$w_1 \leq_E w_2 \text{ iff } L_E(w_1) \leq_{lex} L_E(w_2).$$

Definition 17 only says that we can aggregate pre-orders from faithful assignments to get pre-orders for profiles. The question is, however, if this always results in a syncretic assignment and, by extension, an IC merging operator. Unfortunately, this is not the case: a merging operator defined in this way does not generally satisfy all the IC merging postulates, in particular IC_4. The following example shows this.

Example 7. Take $K_1 = \{00\}, K_2 = \{10, 11\}$, and a faithful assignment that assigns to K_1 and K_2 the following pre-orders:

$$\leq_{K_1}: 00 <_{K_1} 01 \approx_{K_1} 10 \approx_{K_1} 11 \approx_{K_1},$$
$$\leq_{K_2}: 10 \approx_{K_2} 11 <_{K_2} 01 <_{K_2} 00.$$

Pre-orders $\leq_{K_1}$ and $\leq_{K_2}$ are clearly faithful, though it is not assumed that $\leq_{K_1}$ and $\leq_{K_2}$ were computed with any distance function.

Using Definition 17, interpretations get a level with respect to each knowledge base, and we can merge the profile $E = \{K_1, K_2\}$ using the sum as an aggregation function. The results are shown in Figure 3.7. Notice that in the merged pre-order

[7] This is practically the same strategy as in Definition 16.

	$[K_i]$		
	{00}	{10, 11}	Σ
00	0	2	2
01	1	1	2
10	1	0	1
11	1	0	1

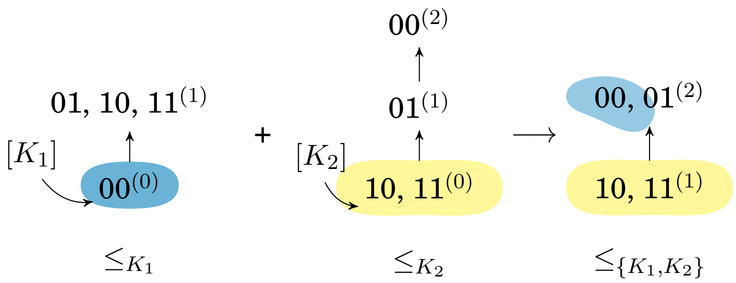

Figure 3.7: Merging two faithful pre-orders using the sum.

$\leq_{\{K_1,K_2\}}$ the models of K_2 are strictly lower than any of the models of K_1. This shows that IC_4 is not satisfied. The reader may convince herself of this by seeing that IC_4 is not true for $[\mu] = \{00, 10, 11\}$.

The problem with IC_4 can be fixed by placing an additional condition on assignments. First, let us extend the notion of distance between interpretations to distances between knowledge bases.

Definition 18. If $\circ$ is a revision operator and $l^{\circ}_{K}(w)$ is the level of w with respect to some knowledge base K as in Definition 17, then *the distance between two knowledge bases K_1 and K_2* is defined as:

$$d(K_1, K_2) = min\{l^{\circ}_{K_1}(w) \mid w \in [K_2]\}.$$

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

Definition 19. Knowledge bases K_1 and K_2 are *symmetric* if they satisfy the following condition:

$$d(K_1, K_2) = d(K_2, K_1) \qquad (symmetry)$$

Proposition 9 ([35])**.** *If $\circ$ is a revision operator and Δ° is an operator defined as in Definition 17 using either Σ or* GMAX *as an aggregation function, then Δ° is an IC merging operator iff any two knowledge bases are symmetric.*[8]

[8]Without *symmetry* on knowledge bases, Δ° satisfies all merging postulates except IC_4. With *symmetry*, Δ° defined with Σ is also a majority merging operator and Δ° defined with GMAX is also an arbitration merging operator.

Observation 5. Notice that in Example 7 *symmetry* does not hold for the knowledge bases K_1 and K_2:

$$d(K_1, K_2) = 1,$$
$$d(K_2, K_1) = 2.$$

One more interesting thing needs to be mentioned here. Proposition 3 shows that we can construct a revision operator $\circ$ from a very thin notion of distance m between interpretations, one that satisfies only *minimality*. If we now add *symmetry* on interpretations, we get *symmetry* for knowledge bases.

Lemma 2 ([35]). *If $\circ$ is a revision operator, then symmetry for knowledge bases holds if and only if $\circ$ is defined from a pseudo-distance.*

This shows that as long as we build our faithful assignments based on pseudo-distances we get *symmetry* for free. Since the familiar distances d_H and d_D are symmetric, any pre-orders got by using these distance measures can be aggregated (with Σ or GMAX), and we get a syncretic assignment.

CHAPTER 4

Revision in the Horn fragment

In this chapter we go over the work on revision in the Horn fragment by Delgrande and Peppas [16, 17]. We present the main difficulties in restricting ourselves to the Horn fragment, as well as what can be done to overcome them.

4.1 Restricting ourselves to the Horn fragment

Chapter 3 focused on revision and merging in the general framework of propositional logic. But for practical applications, it is often useful to limit the expressiveness of a language. A loss in expressiveness can be mitigated by the presence of more efficient reasoning procedures. A quick look at the success of Description Logics (formally, fragments of First Order Logic) shows this is a strategy that can pay off handsomely.

The choice of the propositional Horn fragment is motivated by several factors. For one thing, reasoning in the fragment can be done effectively: a notable result [20] shows that the satisfiability of a propositional Horn formula φ can be decided in linear time (linear in the number of occurrences of literals in φ).

Also, as is mentioned in [29], the rule-like nature of Horn formulas has a conceptual appeal. One can use (first-order) Horn clauses to build up sets inductively, using rules of the form:

$$\begin{aligned}&\mathsf{even}(x) \to \mathsf{odd}(\mathsf{successor}(x)).\\&\mathsf{even}(0).\end{aligned}$$

Expert knowledge is often expressed in terms of rules, and we may expect that reasoning procedures such as revision and merging produce results expressed in a similar way.

From a more theoretical point of view, a Horn clause theory always has a unique minimal model, and one can check if an atom is in the model by retracing the process through which the model was generated. As such, Horn clauses provide the basis of Prolog and related declarative programming frameworks, such as Datalog and Answer Set Programming (see also [18, 49]).

Our purpose is to characterize the class of revision operators that work in the restricted context of Horn propositional logic. The aim is to have a set of logical postulates and a representation result connecting them to a semantic structure. Preferably, the theory of Horn revision does not diverge too much from the theory of regular revision, as the same intuitions guide both inquiries. For now, let us define a Horn revision operator as follows.

Definition 20. A *Horn revision operator* is a function $\circ\colon \mathcal{K}_H \times \mathcal{L}_H \to \mathcal{K}_H$.
We write $K \circ \mu$ instead of $\circ(K, \mu)$.

That is to say, a Horn revision operator maps Horn knowledge bases and Horn formulas to Horn knowledge bases.

We will use the standard $R_1 - R_6$ postulates for revision (see Section 3.2.2) to constrain Horn revision operators. Quick inspection shows that they make sense when restricted to Horn knowledge bases and formulas. In consequence, we use them as such and we do not reiterate them here. The standard notion of faithful assignment also works when restricted to Horn knowledge bases, as it does not depend on the logical language used.

The question is whether something like the Katsuno and Mendelzon representation result which connects the revision postulates with faithful assignments (Proposition 2) still holds in the Horn fragment, and if any concrete operators exist for the Horn fragment. This is what we will address in the following sections.

4.2 Problems

Recall that, by Proposition 1, only sets of interpretations closed under intersection can be represented in the Horn fragment. Thus, a set like $\{01, 10\}$ is not representable by a Horn formula. In general, even if φ_1 and φ_2 are Horn formulas, $\varphi_1 \vee \varphi_2$ might not be expressible as a Horn formula. The problem this creates for belief revision is that the standard model-based operators introduced in Section 3.2.4 might produce results that are not in the fragment. This does, in fact, happen.

4.2.1 Standard operators do not work in the Horn fragment

Consider the faithful pre-order for $[K] = \{111\}$ computed with the Hamming distance d_H, and assume it is embedded in a faithful assignment. Example 3 in Section 3.2.4 illustrates this pre-order. For $[\mu] = \{000, 100, 010, 001\}$ and the operator corresponding to this pre-order, we get:

$$[K \circ \mu] = \{100, 010, 001\}.$$

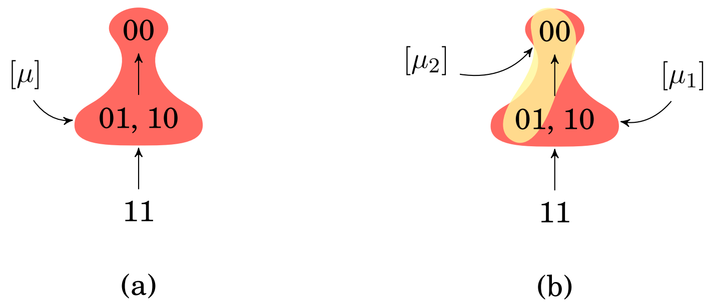

Figure 4.1: Standard revision operators might not work in the Horn fragment; if we try to repair them they might not satisfy the revision postulates.

Since $[K \circ \mu]$ is not closed under intersection, the result of revision is not in the Horn fragment.

Or, for a simpler case, consider the two letter alphabet $\mathcal{U} = \{p_1, p_2\}$, the knowledge base $K = \{p_1, p_2\}$ and the pre-order $\leq_K$ constructed using the Hamming distance (see Figure 4.1-(a)). If we revise by $\mu = (p_1 \wedge p_2) \to \bot$, we have that $[\mu] = \{00, 01, 10\}$ and:

$$min_{\leq_K}[\mu] = \{01, 10\}.$$

Since $\{01, 10\}$ is not closed under intersection, there is no Horn formula that represents it. Though it is easy to see what the result of revision would look like in full propositional logic (the formula $(p_1 \wedge \neg p_2) \vee (\neg p_1 \wedge p_2)$, or something equivalent) there is no way to express this in the Horn fragment. The problem goes beyond distance based operators: indeed, *any* pre-order where there exist sets of models whose minimal elements are not representable by Horn formulas will not work for Horn revision (the drastic distance, however, does not suffer from this problem).

We can try to fix this by redefining how revision by μ is computed on the semantic side: if classical revision produces a set of interpretations that is not representable by a Horn formula, take its closure under intersection. In other words:

$$[K \circ \mu] := Cl_{\cap}(min_{\leq_K}[\mu]).$$

This strategy ensures that revision stays in the Horn fragment and has been pursued in [12]. However, it leads to problems elsewhere.

4.2.2 Revision axioms may not be satisfied

If we take up the suggestion that we always close the set of minimal elements, we get an operator that stays within the Horn fragment. Unfortunately, this operator does not necessarily satisfy the revision postulates, as Examples 8, 9 show.

Example 8. Consider again the pre-order generated via Hamming distance for the two letter alphabet and $K = \{p_1, p_2\}$. Take $\mu_1 = p_1 \wedge p_2 \to \bot$, $\mu_2 = p_1 \to \bot$. The models of μ_1 and μ_2 are $[\mu_1] = \{00, 01, 10\}$ and $[\mu_2] = \{00, 01\}$. They are represented

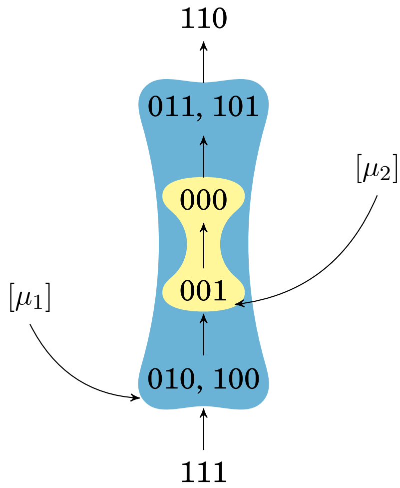

Figure 4.2: A faithful pre-order whose associated operator does not satisfy all the revision postulates

in Figure 4.1-(b). The understanding is that we close any set of minimal elements under intersection if it is not representable by a Horn formula.

However, notice that R_5 does not hold here for the corresponding operator. On the one hand, we have:

$$\begin{aligned}
[K \circ \mu_1] \quad &= Cl_{\cap}(min_{\leq_K}[\mu_1]) = Cl_{\cap}(\{01, 10\}) = \{00, 01, 10\}. \\
[(K \circ \mu_1) \wedge \mu_2] \quad &= \{00, 01\}.
\end{aligned}$$

On the other hand:

$$[K \circ (\mu_1 \wedge \mu_2)] = Cl_{\cap}(min_{\leq_K}[\mu_1 \wedge \mu_2]) = Cl_{\cap}(\{01\}) = \{01\}.$$

If R_5 held it should be the case that $[(K \circ \mu_1) \wedge \mu_2] \subseteq [K \circ (\mu_1 \wedge \mu_2)]$, which is clearly not true. R_6 does hold, though this is merely an artefact of the particular setup we have chosen, as the following example demonstrates. $\square$

Example 9. For the three letter alphabet, take a faithful assignment that assigns to $K = \{p_1, p_2, p_3\}$ the pre-order in Figure 4.2, and otherwise behaves as the default assignment.

Let us test R_5 and R_6 with $\mu_1 = p_1 \wedge p_2 \rightarrow \perp$ and $\mu_2 = (p_1 \rightarrow \perp) \wedge (p_2 \rightarrow \perp)$. We have $[K] = \{111\}$, $[\mu_1] = \{000, 001, 010, 011, 100, 101\}$ and $[\mu_2] = \{000, 001\}$. Revision gives us:

$$\begin{aligned}
[K \circ \mu_1] \quad &= Cl_{\cap}(min_{\leq_K}[\mu_1]) = Cl_{\cap}(\{010, 100\}) = \{000, 010, 100\}, \\
[(K \circ \mu_1) \wedge \mu_2] \quad &= \{000\}.
\end{aligned}$$

At the same time:

$$[K \circ (\mu_1 \wedge \mu_2)] = Cl_{\cap}(min_{\leq_K}[\mu_1 \wedge \mu_2]) = Cl_{\cap}(\{001\}) = \{001\}.$$

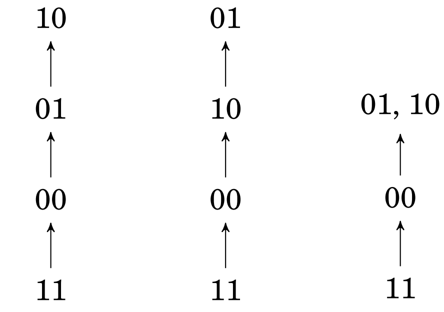

(a) $\leq^1_K$ (b) $\leq^2_K$ (c) $\leq^3_K$

Figure 4.3: Different pre-orders can yield the same corresponding operator

Since $(K \circ \mu_1) \wedge \mu_2$ is consistent, if R_5 and R_6 held it should follow that $[(K \circ \mu_1) \wedge \mu_2] = [K \circ (\mu_1 \wedge \mu_2)]$, but this is clearly not the case. Neither R_5 nor R_6 hold here. □

It seems that as long as there is no restriction on what the minimal models of a formula can be, candidate operators either do not stay in the Horn fragment or fail to satisfy all the revision postulates. The proposal of Delgrande and Peppas [17] is to restrict the pre-orders.

Definition 21. A total pre-order $\leq_K$ is *Horn compliant* if and only if for any Horn formula φ, $min_{\leq_K}[\varphi]$ is representable by a Horn formula.

We say that a faithful assignment is Horn compliant if for any knowledge base K, the pre-order it assigns to K is Horn compliant.

The idea, now is that as long as we work with Horn compliant pre-orders the above problems fade away: the minimal models of any Horn formula will be representable by a Horn formula, so we do not need to take their closure to guarantee that they are in the Horn fragment.

4.2.3 Different rankings yield the same revision operator

A related issue is that different pre-orders (even when Horn compliant) might yield the same corresponding operator. Consider the following example.

Example 10. In the two letter alphabet, take $K = \{p_1, p_2\}$ and three faithful assignments that assign to K the pre-orders $\leq^1_K, \leq^2_K, \leq^3_K$ in Figure 4.3, while otherwise behaving as the default assignment.

It is easy to see that all of these pre-orders yield the same corresponding operator. We just have to show that for any Horn formula μ, it is the case that:

$$min_{\leq^1_K}[\mu] = min_{\leq^2_K}[\mu] = min_{\leq^3_K}[\mu].$$

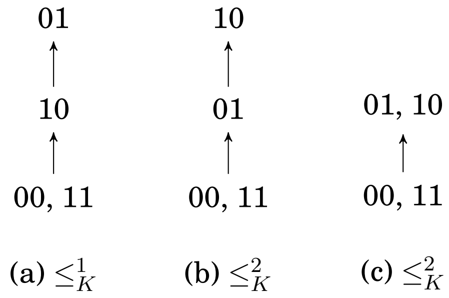

Figure 4.4: The default operator is associated to several pre-orders.

We can convince ourselves of this by inspecting the three pre-orders for different Horn formulas μ. If $[\mu]$ contains only one element, the conclusion is immediate. Next, if $11 \in [\mu]$, then $min_{\leq^i_K}[\mu] = \{11\}$, for $i \in \{1, 2, 3\}$. If $11 \notin [\mu]$ but $00 \in [\mu]$, then $min_{\leq^i_K}[\mu] = \{00\}$. This exhausts all the possibilities: if $00, 11$ would not be in $[\mu]$, then $[\mu]$ would have to be $\{01, 10\}$, which is not be closed under intersection. This would be a contradiction, since μ is assumed to be a Horn formula. $\square$

One can look at Example 10 from a different perspective: a Horn revision operator cannot distinguish between pre-orders $\leq^1_K, \leq^2_K$ and $\leq^3_K$. More to the point, it cannot distinguish between 10 and 01 when neither is among the minimal models of some formula φ.

Though this is not in itself an obstacle to having a Horn revision operator, it does point to a potential weakness of such an operator, since it might lose some control over what the pre-orders look like.

Observation 6. As a side-effect of this issue, consider the operator defined as:

$$K \circ \mu = \begin{cases} K \wedge \mu, & \text{if } K \wedge \mu \text{ is consistent,} \\ \mu, & \text{otherwise.} \end{cases}$$

In Section 3.2.4 we called this the default revision operator. Though it is a revision operator even in the Horn case, it does not enforce the default assignment (induced by the drastic distance) any more. To see this, take $[K] = \{11, 00\}$, and three assignments that assign to K the pre-orders $\leq^1_K, \leq^2_K, \leq^3_K$ in Figure 4.4, while otherwise behaving as the default assignment. Reasoning as in Example 10, we conclude that all of these three assignments yield the same revision operator (the default one), though only $\leq_{K_3}$ belongs to the default assignment.

4.2.4 The revision postulates allow undesirable structures

What makes the representation result of Katsuno and Mendelzon useful is that it works in both directions: revision based on faithful assignments satisfies the revision postulates, and any operator satisfying the revision postulates induces some

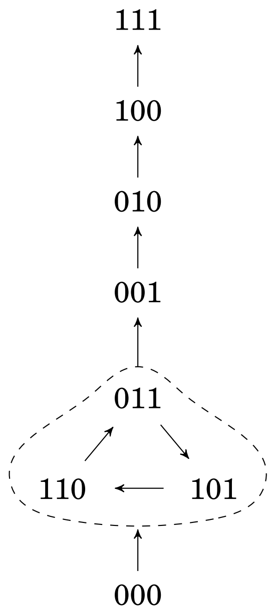

Figure 4.5: A pseudo-preorder whose associated operator satisfies the revision postulates.

pre-order $\leq_K$, for any K. However, this turns out not to be the case in the Horn fragment.

Example 11. In the three letter alphabet, take $K = \{p_1 \to \bot, p_2 \to \bot, p_3 \to \bot\}$ and an assignment that assigns to K the pseudo-preorder $\leq_K$ in Figure 4.5, while otherwise behaving as the default assignment. If $\circ$ is the corresponding operator of $\leq$, Delgrande and Peppas show in [16, 17] that, for K fixed as above, it satisfies the revision postulates (the reader may convince herself of this by directly checking that the postulates are true). $\square$

Since $\leq_K$ is embedded in a faithful assignment, $\circ$ can be extended to arbitrary knowledge bases. This means that we have a full-scale operator, defined on top of a pseudo-preorder, that satisfies the revision postulates. To make matters worse, there is no faithful assignment that produces the same revision operator.

Proposition 10 ([16, 17])**.** *There is no pre-order $\leq^\star$ that has the same corresponding operator as $\leq$ and which does not contain a non-transitive cycle between* $110, 011, 101$.

Proof. Suppose there is such a pre-order $\leq^\star$. Take $[\mu] = \{110, 011, 010\}$. We know that $[K \circ \mu] = \{110\}$, because we have $\leq_K$ to tell us this. Since $\leq^\star$ is equivalent to $\leq_K$, it follows that $min_{\leq^\star_K}[\mu] = \{110\}$. So $110 <^\star 011$.

By the same reasoning, we get that $011 <^\star 101$ and $101 <^\star 110$. This creates the cycle. $\square$

Non-transitive cycles are clearly a problem for Horn revision, as a representation result in the style of Katsuno and Mendelzon (see Prop. 2) does not hold any more. More precisely, one direction does not imply the other: doing revision in the Horn fragment according to the revision postulates does not guarantee that there is a ranking on possible worlds that has the desired properties.

How do we move on from here? One way is to see what sort of structure the Horn revision postulates *do* enforce on the set of possible worlds. Following this thread we would get a representation result for the Horn case, which would in all likelihood be weaker than the Katsuno-Mendelzon one. It could be argued that this is how the situation stands and we should leave things at that.

Another way would be to put the question differently: how would we have to modify the revision postulates such that they force the assignment to be faithful even in the Horn case? The answer to this question takes us in a different direction: rather than studying the sort of semantic structure described by the standard revision postulates, we stick to one particular type of structure (faithful pre-orders) and strengthen the postulates so as to characterize it and nothing else.

There is a philosophical reason to go along the second route: as mentioned in Chapter 3, faithful assignments represent agents' preferences, and the postulates are seen as a tool for characterizing them in a logical language. Allowing the postulates to range over non-faithful pre-orders would mean that we are allowing different types of rankings. And, though it might be useful to study alternative ways of representing the structure of an agent's preferences, we should probably not allow rankings that are unrealistic. In particular, transitivity seems like a core feature of any preference ranking. Non-transitive cycles (as in Example 11) seem to violate our usual intuitions about how people rank their choices.

4.3 Modifying the revision framework

We would like, therefore, to introduce some constraints that eliminate cycles. Delgrande and Peppas [16, 17] have proposed adding the following postulate-schema to the standard postulates $R_1 - R_6$:

($Acyc$) If, for any $n \geq 1$ and $\mu_0, \ldots, \mu_n \in \mathcal{L}_H$, all the following formulas are consistent:

$$\begin{cases} K \circ \mu_0 \wedge \mu_n, \\ K \circ \mu_1 \wedge \mu_0, \\ \ldots \\ K \circ \mu_n \wedge \mu_{n-1}, \end{cases}$$

then $K \circ \mu_n \wedge \mu_0$ is also consistent.

The rationale behind $Acyc$ is that it renders non-transitive cycles in the assignment harmless. To see how this is done, let us see the effect $Acyc$ has on the pseudo-preorder in Example 11 (see Figure 4.5).

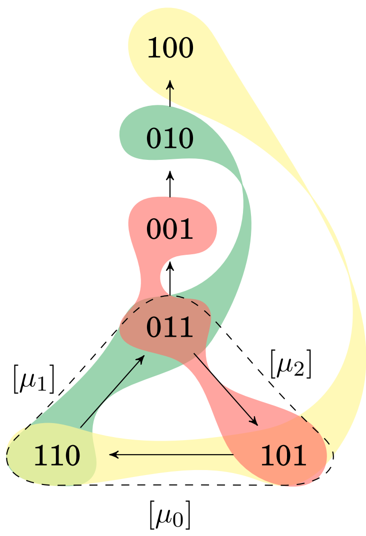

Figure 4.6: This non-transitive cycle is not possible if $Acyc$ is true.

Take $\mu_0 = \varphi_{101,110}, \mu_1 = \varphi_{110,011}, \mu_2 = \varphi_{011,101}$. Notice that:

$$\begin{aligned}
[\mu_0] &= \{101, 110, 100\},\\
[\mu_1] &= \{110, 011, 010\},\\
[\mu_2] &= \{011, 101, 001\}.
\end{aligned}$$

The part of the pre-order that is relevant here is shown in Figure 4.6. We get that:

$$\begin{aligned}
&[K \circ \mu_0 \wedge \mu_2] = min_{\leq_K}[\mu_0] \cap [\mu_2] = \{101\},\\
&[K \circ \mu_1 \wedge \mu_0] = \{110\},\\
&[K \circ \mu_2 \wedge \mu_1] = \{011\}.
\end{aligned}$$

This makes the antecedent of $Acyc$ true. However, its consequent is false since $[K \circ \mu_2 \wedge \mu_0] = \emptyset$. So a non-transitive cycle like the one in Figures 4.5 and 4.6 would not be possible if $Acyc$ were true.

It should be mentioned that, though a pre-order without non-transitive cycles will satisfy $Acyc$, $Acyc$ does not, in general *eliminate* cycles. To see this, consider the pseudo-preorder in Figure 4.7 and its corresponding operator $\circ$. It turns out that $\circ$ also satisfies the revision postulates—not only that, but it also satisfies $Acyc$.

The argument that $\circ$, associated to the pseudo-preorder in Figure 4.7, satisfies the revision postulates is essentially the same as for Example 11. Too see that it also satisfies $Acyc$, let us focus on the cycle and take $\mu_0 = \varphi_{001,100}, \mu_1 = \varphi_{010,100}, \mu_2 =$

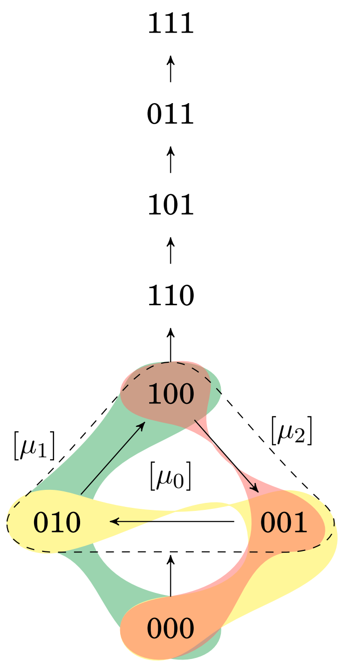

Figure 4.7: A pseudo-preorder with a non-transitive cycle that satisfies the revision postulates, as well as $Acyc$.

$\varphi_{100,001}$. We get:

$$[K \circ \mu_0 \wedge \mu_2] = \{000\},$$
$$[K \circ \mu_1 \wedge \mu_0] = \{000\},$$
$$[K \circ \mu_2 \wedge \mu_1] = \{000\}.$$

Also:

$$[K \circ \mu_2 \wedge \mu_0] = \{000\}.$$

So $K \circ \mu_2 \wedge \mu_0$ is consistent, and $Acyc$ is satisfied for these particular formulas and $n = 2$. It is easy to see that $Acyc$ is also satisfied in general.

The moral is that $Acyc$ has a subtle effect on what structures are allowed: Example 11 was harmful because there did not exist any equivalent pre-order that did the same job as the one with the non-transitive cycle. For the pre-order in Figure 4.7 it turns out that we *can* find an equivalent pre-order. Observation 7, coming later, makes this point.

Though it eliminates potentially unwanted cycles, we have to make sure that $Acyc$ does not eliminate too much. Delgrande and Peppas have shown that:

Proposition 11 ([16, 17])**.** *In the case of regular propositional logic, $Acyc$ is logically implied by postulates $R_1 - R_6$.*

The question is whether $Acyc$ is enough to guarantee that Horn compliant faithful assignments match with Horn revision functions. It turns out this is the case.

4.4 A representation result

Proposition 12 ([16, 17]). *An operator* $\circ\colon \mathcal{K}_H \times \mathcal{L}_H \to \mathcal{K}_H$ *satisfies* $R_1 - R_6 + Acyc$ *if and only if there exists a Horn compliant faithful assignment mapping every Horn knowledge base* $K \in \mathcal{L}_H$ *to a Horn compliant faithful pre-order* $\leq_K$ *such that, for any Horn formula* $\mu \in \mathcal{L}_H$*,* $\circ$ *is the corresponding operator for* $\leq_K$*. In other words:*

$$[K \circ \mu] = min_{\leq}[\mu].$$

We do not give the full proof here, just sketch the general idea. For one direction, we assume that for a Horn compliant faithful assignment and an operator $\circ$ defined on top of it, $\circ$ satisfies the revision postulates $R_1 - R_6$ and $Acyc$. This is a matter of straightforward checking.

For the other direction, the task is to define a Horn compliant faithful assignment assuming that a Horn revision function is given. In [17], James Delgrande and Pavlos Peppas propose a construction that proceeds in three steps, explained below and illustrated in Example 12.

Example 12. We will show how the Delgrande and Peppas definition of a pre-order $\leq_K$ works, given a revision operator $\circ$ and a knowledge base K.

Let us take $K = \{p_1 \to \bot, p_2 \to \bot, p_3 \to \bot\}$, $[K] = \{000\}$. What we need, first of all, is to specify a revision operator $\circ$. How do we do that? One option would be to give a table of the values for $K \circ \mu$, for Horn formulas μ (modulo logical equivalence). Since there is no reason why the reader should trust us when we say that $\circ$ thus introduced is a revision operator, an additional argument would then have to be given that $\circ$ satisfies the revision postulates. This would end up being a very tedious example.

A simpler method is to take $\circ$ as the corresponding operator of an already existing pre-order $\leq$, for instance the one in Figure 4.8. Since $\leq$ is obviously faithful, and at this point it has already been argued that one direction of Proposition 12 holds, we can safely assert that $\circ$ satisfies the revision postulates.

Now that we have a revision operator $\circ$ (given by $\leq$ in Figure 4.8) and we know how it acts on K, let us use the construction suggested by Delgrande and Peppas [17] to obtain a Horn compliant faithful pre-order $\leq_K$.

Step 1. We start by defining a relation $\leq'_K$ on interpretations, in the following way. For any $w_1, w_2 \in \mathcal{W}$:

$$w_1 \leq'_K w_2 \text{ iff } w_1 \in [K \circ \varphi_{w_1,w_2}].$$

As a reminder, φ_{w_1,w_2} is a Horn formula such that $[\varphi_{w_1,w_2}] = Cl_\cap(\{w_1, w_2\})$.

It is easy to see that $\leq'_K$ is reflexive, though in general it is neither transitive nor total. For instance in our running example, by bootstrapping on $\leq$ we get that

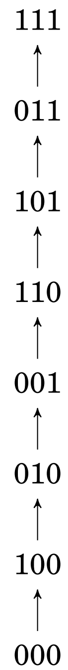

Figure 4.8: A pre-order $\leq$ which we can use to define a revision operator.

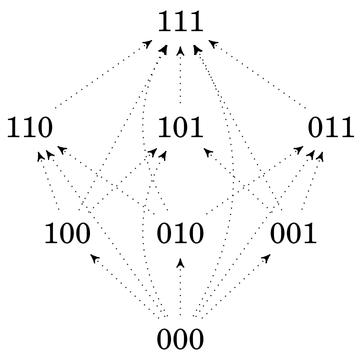

Figure 4.9: Step 1. $\leq'_K$

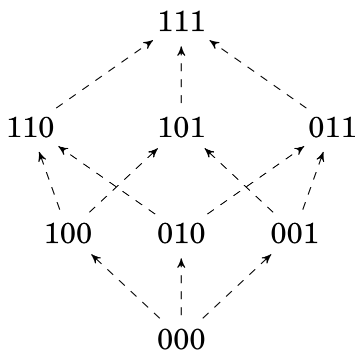

Figure 4.10: Step 2. $\leq^*_K$

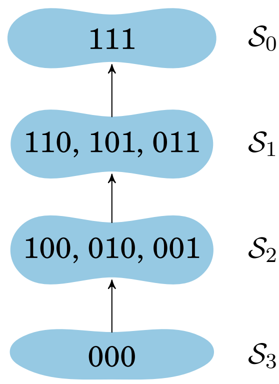

Figure 4.11: Step 3. $\leq_K$

$[K \circ \varphi_{100,010}] = min_{\leq}\{000, 100, 010\} = \{000\}$. Thus, $000 <'_K 100$ and $000 <'_K 010$, but 100 and 010 are not in $\leq'_K$. The full graphic of $\leq'_K$ is shown in Figure 4.9.

Step 2. We take the transitive closure of $\leq'_K$, and denote it by $\leq^*_K$. By design, $\leq^*_K$ is a pre-order on interpretations, but it is still not total; $\leq^*_K$ for our running example is shown in Figure 4.10.

Step 3. We partition the set of interpretations in a series of levels, defined on the basis of $\leq^*_K$: the first level is made up of the maximal elements of $\leq^*_K$, and denoted $\mathcal{S}_0$. After this is done, we put $\mathcal{S}_0$ aside. Then we take the maximal elements of the remaining interpretations and put them in level $\mathcal{S}_1$, and so on. Since there is a finite number of interpretations, this process eventually reaches an end.

The order $\leq_K$ is then defined by taking all the elements in a level $\mathcal{S}_i$ to be equivalent, and strictly smaller than elements in all the preceding levels $\mathcal{S}_j$, $j < i$. In our running example, we get:

$$000 <_K 100 \approx_K 010 \approx_K 001 <_K 110 \approx_K \ldots .$$

The full picture of $\leq_K$ together with the levels is shown in Figure 4.11. Delgrande and Peppas [17] show that all the above notions are well-defined: in particular, any subset of elements in $\leq^*_K$ always has maximal elements, so the definition of the levels makes sense; they also show, through a number of lemmas, that $\leq_K$ is the pre-order we were looking for. □

Observation 7. We have seen that Horn revision operators that give rise to non-transitive cycles like the one in Figure 4.5 are made impossible by adding the postulate $Acyc$. What made that example particularly problematic was that there was no equivalent pre-order (Proposition 10). What about, then, the pseudo-preorder in Figure 4.7? It also contains a non-transitive cycle, and it is *not* rendered impossible by $Acyc$.

The key thing, in the case of Figure 4.7, is that in this case there exists an equivalent pre-order (i.e., it has the same associated revision operator and is transitive). We can find an equivalent pre-order by following through the construction suggested by Delgrande and Peppas and illustrated in the previous example: it turns out that we obtain the same pre-order as that in Figure 4.11.

In support of this claim, it is enough to look at the interpretations involved in the cycle, two at a time. Notice that $[K \circ \varphi_{010,100}] = \{000\}$. This will mean that in the initial step 1, we have $000 <'_K 010$, $000 <'_K 100$ and $010, 100$ are unconnected by $\leq'_K$. The same goes for the pairs $100, 001$ and $001, 010$. For the remaining interpretations, notice that they are ranked in the same way as in Figure 4.8. All this is to say that $\leq'_K$ will end up looking exactly as in the example we just studied, with the rest of the construction following suit.

In fact, the proof of Proposition 12 that we have just sketched offers a constructive method for building a Horn compliant faithful pre-order from a revision operator that also satisfies $Acyc$.

A remaining question is whether we can point to any specific Horn revision operators. By Proposition 12, it is sufficient to find a Horn compliant syncretic assignment that satisfies $Acyc$. In [17], Delgrande and Peppas given an example of such an assignment. The details are given below.

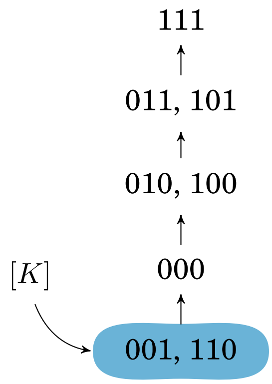

Figure 4.12: Definition 21 gives us this Horn compliant pre-order for $[K] = \{001, 110\}$.

Definition 22. If K is a knowledge base and w_1, w_2 are interpretations, define $\leq_K$ as:

$$w_1 \leq_K w_2 \text{ iff either } w_1 \in [K], \text{ or } w_1, w_2 \notin [K] \text{ and } |w_1| \leq |w_2|.$$

The point of this definition is that it always places $w_1 \cap w_2$ below w_1 and w_2 in a pre-order, and this is sufficient to ensure Horn compliance. As an example, take $[K] = \{001, 110\}$. By Definition 21, $\leq_K$ will look as in Figure 4.12.

CHAPTER 5

Merging in the Horn fragment: challenges and difficulties

In this chapter we present the main difficulties related to Horn merging. Ideas for how to address them are left for the next chapter. Let us start by defining Horn merging operators.

Definition 23. A *Horn merging operator* is a function $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$.
We write $\Delta_\mu(E)$ instead of $\Delta(E, \mu)$.

In other words, a Horn merging operator maps a Horn profile and a Horn formula (the constraint) to a Horn knowledge base.

The problem of Horn merging is much the same as for classical merging: we want to characterize the class of Horn merging operators through a representation result linking a set of logical constraints to a semantic structure. We will use the IC-merging postulates, restricted to the Horn fragment, as the constraints for our Horn merging operator. Quick inspection shows that none of the IC postulates is problematic in the Horn case.

Of the two additional postulates that were presented in Section 3.4.2, Maj can also be used without modification. However, the arbitration postulate Arb does not make sense in the Horn fragment, since it uses disjunction. Disjunction is not defined in the Horn fragment, and neither is it always expressible in terms of other notions. Hence, in order to make Arb functional in the Horn case we would have to replace it with a different postulate, one that preserves the original intuition.

Much as with revision, restricting our merging operators turns out not to be straightforward: we lose all the standard operators, and we open the door to a slew of unintended structures to take shelter under the standard postulates. Correcting this will require patching up both the notion of a syncretic assignment and the postulates we use to model them. Some of the problems that arise for merging are inherited from revision (in a sense to be made more precise below), and for those we will adapt the solutions suggested by Delgrande and Peppas [17] in the context of Horn revision. Other problems, however, turn out to be specific to merging. We will look at these in turn.

5.1 Problems inherited from revision

It was remarked in Section 3.4.2 that a merging operator is also a revision operator. When we identify E with $\bigwedge E$, postulates $IC_0 - IC_3 + IC_7 - IC_8$ are equivalent to postulates $R_1 - R_6$ for revision. What's more, properties $kp_1 - kp_3$ of syncretic assignments are exactly the properties of faithful assignments. That is to say, on the semantic side a merging operator has all the properties of a revision operator.

The analogy goes further. Though we must formally distinguish between profiles and knowledge bases, on the semantic side this distinction is often blurred. For the specific merging operators we have considered, the pre-order for a profile E is often computed by aggregating the pre-order for the knowledge bases in K, using some familiar aggregation function and notion of distance between interpretations. Now, the formal requirements fall on $\leq_E$, not on the pre-orders for particular knowledge bases. However, consider a profile E that contains a single knowledge base, $E = \{K\}$. In this case any standard aggregation function guarantees that $\leq_E$ is exactly $\leq_K$, since by Definition 14 the aggregation function satisfies *identity*—or, to put it differently, aggregating $\leq_K$ alone outputs $\leq$.

The point here is that if a syncretic assignment has pre-orders that are not Horn compliant this will reflect on the merging operator, which—for this reason—will not stay in the Horn fragment. Remember, this was a prominent issue with Horn revision (see Section 4.2.1). For instance, the Hamming distance is off the table from the start, at least in combination with any of the usual aggregation functions. As was already remarked in Section 4.2.1, pre-orders built using the Hamming distance d_H are not always Horn compliant: if $[K] = \{111\}$ and $E = \{K\}$, then $\leq_E$ ends up equating 100, 010 and 001, while placing 000 above them. That is to say, if we take a profile $E = \{K\}$, then any aggregation function will assign to E the same pre-order as the one for K. So for a Horn formula $[\mu] = \{000, 100, 010, 001\}$, we get:

$$min_{\leq_E}[\mu] = \{100, 010, 001\},$$

which is not closed under intersection.

Let us define a Horn compliant assignment for profiles. It is clear what is meant by this, but it does no harm to be precise.

Definition 24. An assignment for profiles is *Horn compliant* if and only if for any profile E, the pre-order it assigns to E is Horn compliant.

It is apparent, then, that the Hamming distance d_H in combination with standard aggregation functions does not produce Horn compliant assignments.

A similar thing applies for non-standard structures, like the non-transitive pre-orders of Section 4.2.4. The point is, if a pseudo-preorder $\leq_K$ satisfies the revision postulates, then the pre-order $\leq_E$ where $E = \{K\}$ will turn out to satisfy the corresponding merging postulates. That is to say, non-transitive cycles need to be enforced with a version of $Acyc$ adapted to the merging scenario:

($Acyc$) If for every $n \geq 1$, all of the following are consistent:

$$\begin{cases} \Delta_{\mu_0}(E) \wedge \mu_n, \\ \Delta_{\mu_1}(E) \wedge \mu_0, \\ \dots \\ \Delta_{\mu_n}(E) \wedge \mu_{n-1}, \end{cases}$$

then $\Delta_{\mu_n}(E) \wedge \mu_0$ is consistent.

$Acyc$ makes a difference only in the Horn fragment. In classical propositional logic, it follows from postulates $IC_0 - IC_3 + IC_7 - IC_8$.

Proposition 13. *$Acyc$ follows from postulates $IC_0 - IC_3 + IC_7 - IC_8$ in full propositional logic.*

Proof. We can show this by induction. For $n = 1$, $Acyc$ follows immediately. Let us suppose now that if $Acyc$ is true for every $i \leq n$, then it is also true for $n + 1$.

Since $\Delta_{\mu_1}(E) \wedge \mu_0$ is consistent, there exists an interpretation $w \in [\Delta_{\mu_1}(E) \wedge \mu_0]$. We know that $(\mu_1 \vee \mu_2) \wedge \mu_1 \equiv \mu_1$, so $\Delta_{(\mu_1 \vee \mu_2) \wedge \mu_1}(E) \equiv \Delta_{\mu_1}(E)$. We infer that $w \in [\Delta_{(\mu_1 \vee \mu_2) \wedge \mu_1}(E)]$. Now, applying IC_7, we get that $[\Delta_{(\mu_1 \vee \mu_2) \wedge \mu_1}(E)] \subseteq [\Delta_{\mu_1 \vee \mu_2}(E) \wedge \mu_1]$, which implies that $w \in [\Delta_{\mu_1 \vee \mu_2}(E)]$. Since we also know that $w \in [\mu_0]$, it follows that $\Delta_{\mu_1 \vee \mu_2}(E) \wedge \mu_0$ is consistent.

Similarly, we get that $\Delta_{\mu_3}(E) \wedge (\mu_1 \vee \mu_2)$ is consistent. Applying the induction hypothesis, to μ_0, $\mu_1 \vee \mu_2$, μ_3, $\dots$, μ_{n+1}, we get the conclusion. $\square$

These problems, inherited from revision, can be fixed by adapting the ideas of Delgrande and Peppas in [17]. Let us look now at problems specific to merging.

5.2 Standard operators do not stay in the Horn fragment

We cannot, in general, rely on the standard operators to work in the Horn case. Let us look at a couple of examples.

We have eliminated assignments built with the Hamming distance d_H on the ground that they are not Horn compliant. The drastic distance d_D fares no better, though for a different reason: even though a pre-order $\leq_K$ built with d_D is Horn compliant, aggregating several of these pre-order does not always produce a Horn compliant pre-order. Thus, in this case, we identify $\leq_E$ with $\leq_K$. Consider knowledge bases $[K_1] = \{01\}$, $[K_2] = \{10\}$ and the profile $E = \{K_1, K_2\}$. Table 5.1 shows

	$[K_i]$			
	{01}	{10}	Σ	GMAX
00	1	1	2	(1,1)
01	0	1	**1**	**(1,0)**
10	1	0	**1**	**(1,0)**
11	1	1	2	(1,1)

Table 5.1: $\leq_E$ aggregated using Σ and GMAX, under levels assigned with the drastic distance d_D.

	$[K_i]$			
	{000}	{111}	Σ	GMAX
000	0	1	**1**	**(1,0)**
001	1	2	3	(2,1)
010	1	2	3	(2,1)
100	1	2	3	(2,1)
011	2	3	5	(3,2)
101	2	3	5	(3,2)
110	2	3	5	(3,2)
111	3	0	3	(3,0)

Table 5.2: $\leq_E$ aggregated using Σ and GMAX, under levels assigned by $\circ^{DP}$.

the pre-order $\leq_E$ got using Σ and GMAX as aggregation functions. Notice, under both aggregation functions the pre-order that results for $\leq_E$ is not Horn compliant: in both cases, $01 \approx_E 10 <_E 00$, which means that the set of minimal elements of $[\mu] = \{00, 01, 10\}$ is not representable by a Horn formula.

5.3 Postulates may not be satisfied

What about the assignment that Delgrande and Peppas define in [17] for their Horn revision operator (see Definition 22)? We already know, from Definition 17, that we can aggregate pre-orders from a faithful assignment. We also know that this assignment is Horn compliant. As a reminder, this assignment orders interpretations w_1 and w_2 according to the number of bits in them equal to 1.

Though Horn compliance is guaranteed, problems arise when we aggregate pre-orders, both with Σ and GMAX. As an example, take $[K_1] = \{000\}$, $[K_2] = \{111\}$ and the profile $E = \{K_1, K_2\}$. The pre-orders $\leq_{K_1}$ and $\leq_{K_2}$ are used to assign levels to each interpretation w, which we then proceed to aggregate using the usual functions. Table 5.2 shows the results. Notice that both aggregation functions produce the same pre-order $\leq_E$, and that (in both cases) $000 <_E 111$. Off the bat, this means

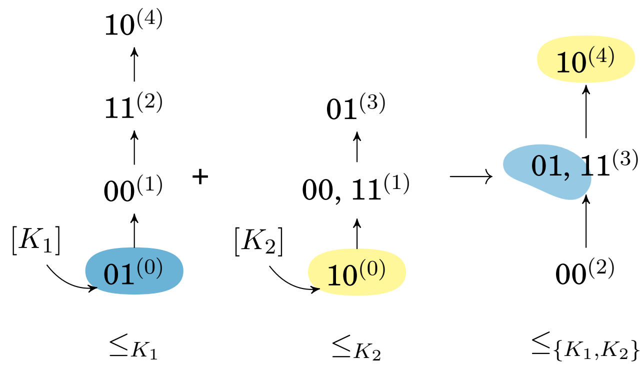

Figure 5.1: $\leq_{\{K_1,K_2\}}$ satisfies the merging axioms, but not kp_4

that as an assignment for profiles, this does not respect kp_4 and is therefore not syncretic. As a reminder, in the setting of regular propositional logic, kp_4 was used to ensure that the merging operator satisfies postulate IC_4. And, though it is not clear at this point that there is a correspondence between merging operators and syncretic assignments in the Horn case, our fears turn out to be true.

Supposing that a merging operator is defined in the usual manner, as $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for a Horn formula μ and a Horn profile E, IC_4 is not satisfied in this assignment. To see this, consider a Horn formula μ such that $[\mu] = \{000, 111\}$. Obviously, $K_1 \models \mu$ and $K_2 \models \mu$. But we have $[\Delta_\mu(\{K_1, K_2\}) \wedge K_1] = \{000\}$, whereas $[\Delta_\mu(\{K_1, K_2\}) \wedge K_2] = \emptyset$.

5.4 Undesirable structures

We will show that there exist assignments that are not syncretic, yet satisfy the merging postulates. It turns out that the merging postulates may be satisfied, even though properties $kp_4 - kp_6$ are individually violated. We show this with counter-examples pertaining to each of the properties.

5.4.1 Property kp_4

In Example 13, we consider an assignment where the merging postulates are satisfied, yet property kp_4 is false. As a reminder, kp_4 says that for any $w_1 \in [K_1]$ there is $w_2 \in [K_2]$ such that $w_2 \leq_{\{K_1,K_2\}} w_1$. The postulate we have to look out for is IC_4, which says that if K_1 and K_2 are consistent and $K_1 \models \mu$, $K_2 \models \mu$, then $\Delta_\mu(\{K_1, K_2\}) \wedge K_1$ is consistent if and only if $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is consistent.

Example 13. Consider belief bases K_1, K_2 such that $[K_1] = \{01\}$, $[K_2] = \{10\}$, and an assignment that works as in Figure 5.1, when restricted to K_1 and K_2.

Figure 5.1 shows the level of each interpretation with respect to the respective knowledge base in parentheses, and the pre-orders that are thus formed. It is worth noting, $\leq_{K_1}$ and $\leq_{K_2}$ are *not* generated using any familiar notion of distance—the levels were assigned by hand. However, merging of $\leq_{K_1}$ and $\leq_{K_2}$ to get $\leq_{\{K_1,K_2\}}$ is done through the Σ aggregation function (i.e., by adding up the levels of each interpretation). $\square$

A (Horn) merging operator Δ is defined on top of this assignment in the usual way, by taking $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any $\mu \in \mathcal{L}_H$. It makes sense to do this, as we argue below.

Proposition 14. *There exists a Horn merging operator Δ that behaves as in Figure 5.1 with respect to K_1, K_2 and $\{K_1, K_2\}$.*

Proof. The full assignment (of which Figure 5.1 is only a fragment) is given in Section 7.4, and there it becomes clearer that one can define a merging operator on top of it. In this section let us concentrate on $\leq_{K_1}$, $\leq_{K_2}$ and $\leq_{\{K_1,K_2\}}$.

Direct inspection shows that $\leq_{K_1}$, $\leq_{K_2}$ and $\leq_{\{K_1,K_2\}}$ are faithful pre-orders and that they are Horn compliant. This alone makes it possible to associate with each of the profiles $\{K_1\}, \{K_2\}, \{K_1, K_2\}$ an operator Δ, defined as $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, if E is one of these three profiles and μ is a Horn formula. By the Katsuno and Mendelzon representation result (see Proposition 2), Δ satisfies Postulates $IC_0 - IC_3 + IC_7 - IC_8$, since these are essentially the postulates for revision.

Because we aggregate $\leq_{K_1}$ and $\leq_{K_2}$ with Σ, $kp_5 - kp_6$ are also satisfied. It is easy to see why this holds, as kp_5 and kp_6 reduce to a couple of inequalities between positive integers. This guarantees that $IC_5 - IC_6$ are also satisfied, since they depend only on kp_5 and kp_6.

The only postulate left to be checked is IC_4. Even here, the only problematic case concerns the models of K_1 and K_2 in $\leq_{\{K_1,K_2\}}$. Take, then, a Horn formula μ such that $[\mu] = Cl_\cap([K_1] \cup [K_2]) = \{00, 01, 10\}$. Obviously, $K_1 \vDash \mu$ and $K_2 \vDash \mu$, so we are in the range of application of IC_4 (in fact, this is exactly the kind of situation IC_4 is employed for).

We have that $[\Delta_\mu(\{K_1, K_2\})] = \{00\}$, 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. $\square$

At the same time, though, kp_4 is not true for this assignment, because $01 <_{\{K_1,K_2\}} 10$.

Could it be the case, however, that there is an alternative pre-order where kp_4 is satisfied, and which yields the same operator as $\leq_{\{K_1,K_2\}}$?

Proposition 15. *There is no pre-order satisfying kp_4 which yields the same operator as $\leq_{\{K_1,K_2\}}$.*

Proof. Suppose there is a pre-order $\leq^\star_{\{K_1,K_2\}}$ where kp_4 is satisfied, and which yields the same merging operator as $\leq_{\{K_1,K_2\}}$.

If $\leq^{\star}_{\{K_1,K_2\}}$ satisfies kp_4, then it should be the case that $10 \leq^{\star}_{\{K_1,K_2\}} 01$. However, we know that:

$$\begin{aligned} [\Delta_{\varphi_{10,11}}(\{K_1, K_2\})] &= \{11\}, \\ [\Delta_{\varphi_{01,11}}(\{K_1, K_2\})] &= \{01, 11\}. \end{aligned}$$

This implies that $01 \approx^{\star}_{\{K_1,K_2\}} 11 <^{\star}_{\{K_1,K_2\}} 10$, which creates a contradiction. $\square$

Proposition 15 shows that there is no way of emulating the same merging operator while making kp_4 true.

5.4.2 Properties kp_5 and kp_6

For full propositional logic postulates IC_5 and IC_6 make sure that an assignment corresponding to a merging operator satisfies properties kp_5 and kp_6, respectively. In this subsection we make the case that this becomes problematic in the Horn case: even though there might exist an assignment that satisfies all the merging postulates, the examples show it is doubtful whether it has to also satisfy kp_5 and/or kp_6. We remind the reader of postulates $IC_5 - IC_6$ and properties $kp_5 - kp_6$:

(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 $\Delta_\mu(E_1 \sqcup E_2) \models \Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$;

(kp_5) if $\begin{cases} w_1 \leq_{E_1} w_2, \\ w_1 \leq_{E_2} w_2, \end{cases}$ then $w_1 \leq_{E_1 \sqcup E_2} w_2$;

(kp_6) if $\begin{cases} w_1 \leq_{E_1} w_2, \\ w_1 <_{E_2} w_2, \end{cases}$ then $w_1 <_{E_1 \sqcup E_2} w_2$.

The following examples show how $IC_5 - IC_6$ can fail to enforce $kp_5 - kp_6$ in the case of Horn logic. The first one targets property kp_5.

Example 14. Assume there exists a faithful assignment which for two profiles E_1 and E_2 behaves as in Figure 5.2 with respect to interpretations 000, 010, 100 and 110, and is otherwise Horn compliant and syncretic.

Notice that kp_5 does not hold in this setup: we have $010 \approx_{E_1} 100, 010 \approx_{E_2} 100$, but $010 <_{E_1 \sqcup E_2} 100$. This raises the question of whether IC_5 could be satisfied for a merging operator defined on top of this assignment.

First of all, let us show that under our assumptions such an operator could exist. We have said that the assignment is Horn compliant and syncretic everywhere, except on the fragment shown in Figure 5.2. It turns out the assignment is Horn compliant here as well: we have that $\min_{\leq_E}\{000, 010, 100\} = \{000\}$ for any pre-order $\leq_E \in \{\leq_{E_1}, \leq_{E_2}, \leq_{E_1 \sqcup E_2}\}$. Any other subset of interpretations is linearly ordered, hence its set of minimal elements is representable by a Horn formula.

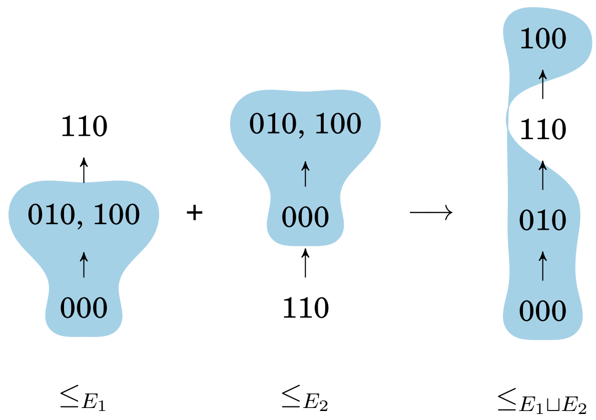

Figure 5.2: Counter-example for kp_5.

This implies that the assignment is Horn compliant everywhere. It makes sense, then, to define a Horn operator $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ such that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any Horn formula μ and Horn profile E. We want to claim that Δ satisfies all the Horn merging postulates. Indeed, this is true anywhere Δ is under-girded by a Horn compliant syncretic assignment, which—we have assumed—is true everywhere except in Figure 5.2. All we are left to check is Figure 5.2.

Since the assignment is faithful, we can assert, by the left-to-right direction of the Delgrande and Peppas representation result for Horn revision (Proposition 12), that the postulates pertaining to Horn revision ($IC_0 - IC_3 + IC_7 - IC_8 + Acyc$) are true. Since the assignment is, by assumption, syncretic everywhere except the part displayed in Figure 5.2, we can also assert that IC_4 is satisfied.

All that remain are postulates $IC_5 - IC_6$ and, to cut the suspense, the only case that could be problematic is the one posed by the interpretations for which kp_5 does not hold, which are interpretations 010 and 100.

Notice, however, that there is no Horn formula that represents exactly the set $\{010, 100\}$. The best we can do is $[\varphi_{010,100}] = \{000, 010, 100\}$, and in this case we have that $\Delta_{\varphi_{010,100}}(E_1) \wedge \Delta_{\varphi_{010,100}}(E_2)$ is consistent, and:

$$[\Delta_{\varphi_{010,100}}(E_1) \wedge \Delta_{\varphi_{010,100}}(E_2)] = [\Delta_{\varphi_{010,100}}(E_1 \sqcup E_2)] = \{000\}.$$

We have just shown that for $[\mu] = \{000, 100, 010\}$, IC_5 and IC_6 are true. It is trivial to show that $IC_5 - IC_6$ are also satisfied for any other formulas whose models are among the featured interpretations. $\square$

It is perhaps surprising to see that IC_5 can be satisfied in an assignment where kp_5 does not hold, but closer thought shows this is to be expected: since in the Horn fragment we cannot capture 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 kp_5.

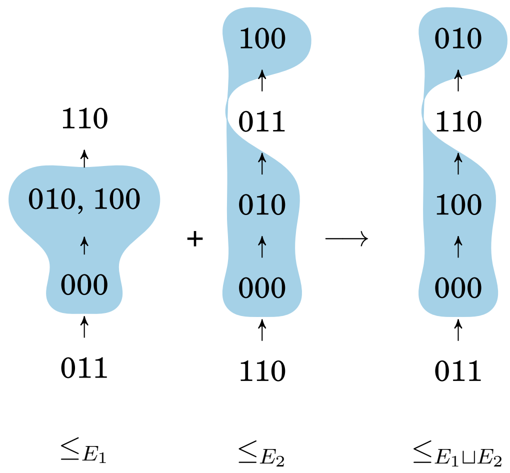

Figure 5.3: Counter-example 1 for kp_6.

This might not be a problem if there existed an equivalent assignment described by Δ, where kp_5 holds: in our case, it would mean that there exists a pre-order $\leq^\star_{E_1 \sqcup E_2}$ such that $100 \approx^\star_{E_1 \sqcup E_2} 010$. However, as the following proposition shows, this is not possible.

Proposition 16. *There is no Horn compliant syncretic assignment based on the operator Δ from Example 14 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. Suppose there is such an assignment, and keep in mind that (by assumption) it is described by the same operator Δ that describes Figure 5.2.

Looking at $\leq_{E_1 \sqcup E_2}$, we get that:

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

This leads us to conclude that $010 <^\star_{E_1 \sqcup E_2} 110 <^\star_{E_1 \sqcup E_2} 100$, and by transitivity $010 <^\star_{E_1 \sqcup E_2} 100$. This contradicts the assumption that $100 \approx^\star_{E_1 \sqcup E_2} 010$. □

Example 14 will serve as a template for subsequent ones showing the different ways in which postulates $IC_5 - IC_6$ can fail to enforce properties $kp_5 - kp_6$ of syncretic assignments. The following examples are directed at kp_6, and the reasoning is analogous to Example 14. Hence we skip the ceremony and only point out how IC_6 is true, even though kp_6 for the underlying assignment is not.

Example 15. Assume there exists an assignment which behaves locally as in Figure 5.3 and is syncretic otherwise. Notice that $010 \approx_{E_1} 100$ and $010 <_{E_2} 100$, so kp_6 would

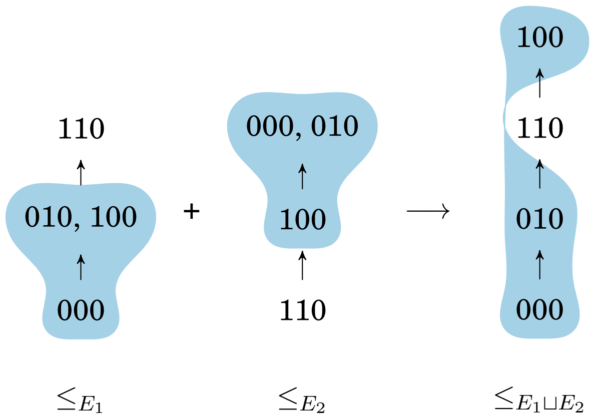

Figure 5.4: Counter-example 2 for kp_6.

require that $010 <_{E_1 \sqcup E_2} 100$. Instead, we have $100 <_{E_1 \sqcup E_2} 010$. At the same time, $\Delta_{\varphi_{010,100}}(E_1) \wedge \Delta_{\varphi_{010,100}}(E_2)$ is consistent, and:

$$[\Delta_{\varphi_{010,100}}(E_1) \wedge \Delta_{\varphi_{010,100}}(E_2)] = [\Delta_{\varphi_{010,100}}(E_1 \sqcup E_2)] = \{000\}.$$

$\square$

Could we 'fix' the assignment in Figure 5.3 by putting $010 <_{E_1 \sqcup E_2} 100$, so that kp_6 holds? The answer is no.

Proposition 17. *There is no Horn compliant syncretic assignment based on the same operator as that in Figure 5.3 which assigns to* $E_1 \sqcup E_2$ *a pre-order* $\leq^\star_{E_1 \sqcup E_2}$ *where* $010 <^\star_{E_1 \sqcup E_2} 100$.

Proof. We already have:

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

So $100 <^\star_{E_1 \sqcup E_2} 110 <^\star_{E_1 \sqcup E_2} 010$, and by transitivity $100 <^\star_{E_1 \sqcup E_2} 010$. So it cannot be the case that $010 \approx^\star_{E_1 \sqcup E_2} 100$. $\square$

For similar reasons, we cannot fix the assignment by putting $100 <_{E_2} 010$. This shows that if a Horn merging operator Δ behaves as in Figure 5.3 there is no escaping the fact that kp_6 does not hold in the assignment Δ describes.

Let us now look at another type of counter-example for kp_6.

Example 16. Assume there exists an assignment which behaves locally as in Figure 5.4 and is syncretic otherwise. Notice that $100 \approx_{E_1} 010$ and $100 <_{E_2} 010$, so kp_6 would require that $100 <_{E_1 \sqcup E_2} 010$. Instead, we have $010 <_{E_1 \sqcup E_2} 100$. At the same time, we have that:

$$[\Delta_{\varphi_{010,100}}(E_1) \wedge \Delta_{\varphi_{010,100}}(E_2)] = \{000\} \cap \{100\} = \emptyset,$$

Thus, $\Delta_{\varphi_{010,100}}(E_1) \wedge \Delta_{\varphi_{010,100}}(E_2)$ is inconsistent and $IC_5 - IC_6$ are satisfied (IC_6 is trivially satisfied, since the condition for its application does not apply). As before, this cannot be fixed by finding an equivalent pre-order $\leq_{E_1 \sqcup E_2}$ where 100 is placed strictly lower than 010 (the argument is the same as for Example 14). $\square$

5.4.3 Discussion

In the regular case, a faithful assignment (that satisfies$kp_1 - kp_3$) is coerced tp satisfy $kp_4 - kp_6$ by postulates $IC_4 - IC_6$. This is achieved because we can apply IC_4 to the formula $[\mu] = [K_1] \cup [K_2]$ and $IC_5 - IC_6$ to $[\mu] = [\varphi_{w_1,w_2}]$, for any two interpretations w_1, w_2. The postulates thus control which interpretations can appear in $min_{\leq_E}[\mu]$, for a Horn profile E. However, in the Horn case we have less control on what can appear in $min_{\leq_E}[\mu]$, especially if $[\mu] = \{w_1, w_2, w_1 \cap w_2\}$.

With respect to $kp_5 - kp_6$, problematic cases could appear for two reasons. On the one hand we can have:

$$\Delta_\mu(E_1) \wedge \Delta_\mu(E_2) \text{ is inconsistent,}$$

and μ is the formula characterizing the interpretations that contradict kp_5 (or kp_6). In this case IC_5 is trivially true, and IC_6 just fails to apply. Hence, if there are problematic cases where this is the case, axioms IC_5, IC_6 simply are of no use.

The other case is where:

$$\begin{aligned} &\Delta_\mu(E_1) \wedge \Delta_\mu(E_2) \text{ is consistent, and} \\ &\Delta_\mu(E_1) \wedge \Delta_\mu(E_2) \equiv \Delta_\mu(E_1 \sqcup E_2). \end{aligned}$$

Such a situation, if possible, would show that IC_5, IC_6 can apply meaningfully, but fail to detect problematic cases.

It is apparent by now that the 'extra' interpretations 011 and 110 appear in these examples as wedges between interpretations, to guarantee that their order cannot be switched without changing the operator on which it is based: for instance, the presence of 110 between 010 and 100 in $\leq_{E_1 \sqcup E_2}$ (Example 14) forces any Horn operator modelling this arrangement to replicate it as such and, implicitly, place 100 strictly lower than 010. This occurs in a context where we cannot say anything about the pair $\{010, 100\}$ directly, as it is not representable by a Horn formula.

It does, however, suggest the possibility that if we could somehow guarantee that such wedges do not exist, then we could rearrange 100 and 010 in ways that conform to kp_5 and/or kp_6. We will be exploring this possibility in the next chapter.

CHAPTER 6

Merging in the Horn fragment: a representation result

In this chapter we explore the correspondence between syncretic assignments and Horn merging operators. As we have seen from the previous chapter, the standard framework for merging does not work in the Horn case. We explore viable solutions by adapting the Delgrande and Peppas [17] work for Horn revision, and extending it to Horn merging.

Section 6.1 presents one half of a representation theorem: it shows that one can define a Horn merging operator that satisfies $IC_0 - IC_8 + Acyc$ on top of a syncretic assignment, if the assignment satisfies the additional condition of Horn compliance. The following sections tackle the reverse problem: what conditions does a Horn merging operator Δ need to satisfy in order to describe a syncretic assignment?

6.1 From assignments to operators

In this section we present a first half of a representation result, which takes us from horn compliant syncretic assignments to Horn merging operators. As a reminder, a pre-order $\leq$ on interpretations is Horn compliant if for any Horn formula μ, the set $min_{\leq}[\mu]$ is representable by a Horn formula (see Definition 21).

Theorem 1. *If there exists a Horn compliant syncretic assignment mapping each profile E to a total pre-order $\leq_E$, then an operator $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ defined as:*

$$[\Delta_\mu(E)] = min_{\leq_E}[\mu],$$

satisfies postulates $IC_0 - IC_8 + Acyc$.

Observation 8. Note that under the condition of Horn compliance, $min_{\leq_E}[\mu]$ is representable in the Horn fragment, for any Horn formula μ. This makes the definition formally correct, since it ensures that the result of merging is in the Horn fragment.

Proof. Assume that there exists a Horn compliant syncretic assignment mapping each profile E to a total pre-order $\leq_E$, and define the operator $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ as:

$$[\Delta_\mu(E)] = min_{\leq_E}[\mu],$$

Let us take each of the postulates $IC_1 - IC_8 + Acyc$ in turn and show they are satisfied.

(IC_0) $\Delta_\mu(E) \models \mu$.

This is equivalent to $min_{\leq_E}[\mu] \subseteq [\mu]$, which holds in virtue of how the set of minimal elements is defined (see Section 2.3).

(IC_1) If μ is consistent, then $\Delta_\mu(E)$ is consistent.

Suppose, on the contrary, that μ is consistent and $\Delta_\mu(E)$ is inconsistent. Then $[\mu] \neq \emptyset$ and $min_{\leq_E}[\mu] = \emptyset$. Take, now, $w_0 \in [\mu]$. Since $min_{\leq_E}[\mu] = \emptyset$, w_0 cannot be a minimal model of μ. Hence there exists $w_1 \in [\mu]$ such that $w_1 <_E w_0$.

By the same reasoning w_1 is not minimal, so there exists $w_2 \in [\mu]$ such that $w_2 <_E w_1$. Iterating this process, we can generate a sequence $w_0, w_1, \ldots, w_n, \ldots$ of models of μ, such that:

$$\cdots <_E w_n <_E \cdots \leq_E w_1 <_E w_0.$$

Since $[\mu]$ is finite, some element w_i in this chain must appear more than once, and we get:

$$w_i <_E w_{i+k} <_E \cdots <_E w_{i+1} <_E w_i.$$

By transitivity we get that $w_i \leq_E w_{i+1} <_E w_i$, which is a contradiction.

(IC_2) If $\bigwedge E$ is consistent with μ, then $\Delta_\mu(E) \equiv \bigwedge E \wedge \mu$.

As a first thing to keep in mind we have, by assumption, that $[E] \cap [\mu] \neq \emptyset$. This tells us that there exists $w_0 \in [E] \cap [\mu]$ and, consequently, that $[\mu] \neq \emptyset$. By IC_1, it follows that $min_{\leq_E}[\mu] \neq \emptyset$. Now, let us show that $min_{\leq_E}[\mu] = [E] \cap [\mu]$ by double inclusion.

First, take $w_1 \in min_{\leq_E}[\mu]$. It is immediate that $w_1 \in [\mu]$. Suppose, however, that $w_1 \notin [E]$. Consider the w_0 from above: since $w_0 \in [E]$ and the assignment is syncretic, we have (by kp_2) that $w_0 <_E w_1$. At the same time, since $w_0, w_1 \in [\mu]$ and, moreover, $w_1 \in min_{\leq_E}[\mu]$, we get that $w_1 \leq_E w_0$. This leads to a contradiction. So $w_1 \in [E]$, and therefore $min_{\leq_E}[\mu] \subseteq [E] \cap [\mu]$.

For the reverse inclusion, take $w_1 \in [E] \cap [\mu]$ and suppose $w_1 \notin min_{\leq_E}[\mu]$. Since $min_{\leq_E}[\mu] \neq \emptyset$, there exists $w_2 \in min_{\leq_E}[\mu]$. In virtue of this we get that $w_2 <_E w_1$. But we also have that $w_1 \in [E]$, and by kp_1 and kp_2, $w_1 \leq_E w_2$. This leads to a contradiction, so $w_1 \in min_{\leq_E}[\mu]$ and therefore $[E] \cap [\mu] \subseteq min_{\leq_E}[\mu]$, which concludes the proof.

(IC_3) If $E_1 \equiv E_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}(E_1) \equiv \Delta_{\mu_2}(E_2)$.

From $E_1 \equiv E_2$ and kp_3 we have that $\leq_{E_1} = \leq_{E_2}$. Since $\mu_1 \equiv \mu_2$, we get that $[\mu_1] = [\mu_2]$. Putting these two results together, we have that $min_{\leq_{E_1}}[\mu_1] = min_{\leq_{E_2}}[\mu_2]$, and hence $\Delta_{\mu_1}(E_1) \equiv \Delta_{\mu_2}(E_2)$.

(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.

Take K_1, K_2 and μ such that $K_1 \models \mu$, $K_2 \models \mu$, and suppose $min_{\leq_{\{K_1,K_2\}}}[\mu] \cap [K_1] \neq \emptyset$. Take $w_1 \in min_{\leq_{\{K_1,K_2\}}}[\mu] \cap [K_1]$. Then $w_1 \in [K_1]$, and by kp_4 there exists $w_2 \in [K_2]$ such that $w_2 \leq_{\{K_1,K_2\}} w_1$. Because $[K_2] \subseteq [\mu]$, it also follows that $w_2 \in [\mu]$.

Putting this together with the fact that $w_1 \in min_{\leq_{\{K_1,K_2\}}}[\mu]$, we get that $w_2 \in min_{\leq_{\{K_1,K_2\}}}[\mu]$ and hence (because $w_2 \in [K_2]$), $w_2 \in min_{\leq_{\{K_1,K_2\}}}[\mu] \cap [K_2]$. This shows that $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is consistent. The converse is completely analogous.

(IC_5) $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2) \models \Delta_\mu(E_1 \sqcup E_2)$.

If $min_{\leq_{E_1}}[\mu] \cap min_{\leq_{E_2}}[\mu] = \emptyset$, the conclusion is immediate. Otherwise, take $w_0 \in min_{\leq_{E_1}}[\mu] \cap min_{\leq_{E_2}}[\mu]$. We want to show that $w_0 \in min_{\leq_{E_1 \sqcup E_2}}[\mu]$.

Take any $w \in [\mu]$ and suppose $w \leq_{E_1 \sqcup E_2} w_0$. We have that $w_0 \leq_{E_1} w$, and $w_0 \leq_{E_2} w$, hence (by kp_5), $w_0 \leq_{E_1 \sqcup E_2} w$. This shows that $w_0 \in min_{\leq_{E_1 \sqcup E_2}}[\mu]$.

(IC_6) If $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$ is consistent, then $\Delta_\mu(E_1 \sqcup E_2) \models \Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$.

From the fact that $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$ is consistent we infer that:

$$min_{\leq_{E_1}}[\mu] \cap min_{E_2}[\mu] \neq \emptyset,$$

so there exists $w_0 \in min_{\leq_{E_1}}[\mu] \cap min_{E_2}[\mu]$. This shows that $w_0 \in [\mu]$ and $[\mu] \neq \emptyset$, so (by IC_1) $min_{\leq_{E_1 \sqcup E_2}}[\mu] \neq \emptyset$.

Take, now, $w \in min_{\leq_{E_1 \sqcup E_2}}[\mu]$ and suppose $w \notin min_{\leq_{E_1}}[\mu] \cap min_{\leq_{E_1}}[\mu]$. Then, either $w \notin min_{\leq_{E_1}}[\mu]$ or $w \notin min_{\leq_{E_2}}[\mu]$.

If $w \notin min_{\leq_{E_1}}[\mu]$, recall the w_0 from above. We get that $w_0 <_{E_1} w$ and $w_0 \leq_{E_2} w$. By kp_6, we have $w_0 <_{E_1 \sqcup E_2} w$. But at the same time, since $w \in min_{\leq_{E_1 \sqcup E_2}}[\mu]$, we have that $w \leq_{E_1 \sqcup E_2} w_0$. Together, these conclusions lead to a contradiction.

The case $w \notin min_{\leq_{E_2}}[\mu]$ is completely analogous and also also leads to a contradiction. So the original assumption must be false. In other words, $w \in min_{\leq_{E_1}}[\mu] \cap min_{\leq_{E_2}}[\mu]$, and so $min_{\leq_{E_1 \sqcup E_2}}[\mu] \subseteq min_{\leq_{E_1}}[\mu] \cap min_{\leq_{E_2}}[\mu]$.

(IC_7) $\Delta_{\mu_1}(E) \wedge \mu_2 \models \Delta_{\mu_1 \wedge \mu_2}(E)$.

IC_7 is equivalent to $min_{\leq_E}[\mu_1] \cap [\mu_2] \subseteq min_{\leq_E}[\mu_1 \wedge \mu_2]$. The case when $min_{\leq_E}[\mu_1] \cap [\mu_2] = \emptyset$ is immediate, hence assume $min_{\leq_E}[\mu_1] \cap [\mu_2] \neq \emptyset$ and take $w \in min_{\leq_E}[\mu_1] \cap [\mu_2]$. Observe that this implies $w \in [\mu_1]$, and since $w \in [\mu_2]$, we have $w \in [\mu_1] \cap [\mu_2]$

which gives $w \in [\mu_1 \wedge \mu_2]$. We now want to show that w is among the minimal elements of $[\mu_1 \wedge \mu_2]$.

For that, assume $w \notin min_{\leq_E}[\mu_1 \wedge \mu_2]$. We already have that $[\mu_1 \wedge \mu_2] \neq \emptyset$, which (by IC_1) gives us $min_{\leq_E}[\mu_1 \wedge \mu_2] \neq \emptyset$. Take, then, $w_0 \in min_{\leq_E}[\mu_1 \wedge \mu_2]$. We conclude that $w_0 <_E w$. Regarding w_0, we can also conclude that $w_0 \in [\mu_1]$. However, we also have that $w \in min_{\leq_E}[\mu_1]$, which implies $w \leq_E w_0$. We have arrived at a contradiction.

(IC_8) If $\Delta_{\mu_1}(E) \wedge \mu_2$ is consistent, then $\Delta_{\mu_1 \wedge \mu_2}(E) \models \Delta_{\mu_1}(E) \wedge \mu_2$.

By hypothesis, $min_{\leq_E}[\mu_1] \cap [\mu_2] \neq \emptyset$, hence there exists $w_0 \in min_{\leq_E}[\mu_1] \cap [\mu_2]$. It follows that $w_0 \in [\mu_1]$, $w_0 \in [\mu_2]$ and, from here, that $w_0 \in [\mu_1 \wedge \mu_2]$.

We want to show that $min_{\leq_E}[\mu_1 \wedge \mu_2] \subseteq min_{\leq_E}[\mu_1] \cap [\mu_2]$, so take $w \in min_{\leq_E}[\mu_1 \wedge \mu_2]$. Obviously, $w \in [\mu_1 \wedge \mu_2]$, hence $w \in [\mu_1]$ and $w \in [\mu_2]$. Suppose, however, that $w \notin min_{\leq_E}[\mu_1]$. This, together with $w_0 \in min_{\leq_E}[\mu_1]$, implies that $w_0 < w$. On the other hand, $w \in min_{\leq_E}[\mu_1 \wedge \mu_2]$ and $w_0 \in [\mu_1 \wedge \mu_2]$, so $w \leq w_0$. We have arrived at a contradiction.

$(Acyc)$ If for every $n \in \mathbb{N}$, all of the following are consistent:

$$\begin{cases} \Delta_{\mu_0}(E) \wedge \mu_n, \\ \Delta_{\mu_1}(E) \wedge \mu_0, \\ \dots \\ \Delta_{\mu_n}(E) \wedge \mu_{n-1}, \end{cases}$$

then $\Delta_{\mu_n}(E) \wedge \mu_0$ is consistent.

By hypothesis, we have that $min_{\leq_E}[\mu_0] \cap [\mu_n] \neq \emptyset$, and $min_{\leq_E}[\mu_{i+1}] \cap [\mu_i] \neq \emptyset$, for $0 \leq i < n$. This means that there exist $w_0, w_1, \dots, w_n$ such that:

$$\begin{gathered} w_0 \in min_{\leq_E}[\mu_1] \cap [\mu_0], \\ w_1 \in min_{\leq_E}[\mu_2] \cap [\mu_1], \\ \dots \\ w_{n-1} \in min_{\leq_E}[\mu_n] \cap [\mu_{n-1}], \\ w_n \in min_{\leq_E}[\mu_0] \cap [\mu_n]. \end{gathered}$$

From the first two statements we can infer that $w_0 \in min_{\leq_E}[\mu_1]$ and $w_1 \in [\mu_1]$, which implies that $w_0 \leq_E w_1$. Following up, $w_1 \leq_E w_2$, ..., $w_{n-1} \leq_E w_n$ and $w_n \leq_E w_0$. By transitivity, it follows that $w_n \leq_E w_{n-1}$. Putting this together with the fact that $w_{n-1} \in min_{\leq_E}[\mu_n]$ and $w_n \in [\mu_n]$, it follows that $w_n \in min_{\leq_E}[\mu_n]$. We also have that $w_n \in [\mu_0]$, so $w_n \in min_{\leq_E}[\mu_n] \cap [\mu_0]$. Clearly, then, $min_{\leq_E}[\mu_n] \cap [\mu_0] \neq \emptyset$.

(Maj) There exists $n > 0$ such that $\Delta_\mu(E_1 \sqcup E_2^n) \models \Delta_\mu(E_2)$.

We want to show that there exists an $n > 0$ such that $min_{\leq_{E_1 \sqcup E_2^n}}[\mu] \subseteq min_{\leq_{E_2}}[\mu]$. Suppose, on the contrary, that for all $n > 0$ there exists $w_n \in [\mu]$ such that $w_n \in min_{\leq_{E_1 \sqcup E_2^n}}[\mu]$ and $w_n \notin min_{\leq_{E_2}}[\mu]$.

Since $[\mu]$ is finite, one of the w_i's will have to appear infinitely often. In other words, there is an infinite set $M = \{n_1, n_2, \dots\} \subseteq \mathbf{N}^\star$ and a $w \in [\mu]$ such that for every $n_i \in M$ we have $w \in min_{\leq_{E_1 \sqcup E_2^{n_i}}}[\mu]$ and $w \notin min_{\leq_{E_2}}[\mu]$. Without loss of generality assume M is ordered, i.e. $n_1 \leq n_2 \leq \dots$.

Obviously $[\mu] \neq \emptyset$, and (by IC_1) $min_{\leq_{E_2}}[\mu] \neq \emptyset$. Take $w_0 \in min_{\leq_{E_2}}[\mu]$. Two things are of interest here. First, we have that $w \leq_{E_1 \sqcup E_2^{n_i}} w_0$, for every $n_i \in M$. Second, we get that $w_0 <_{E_2} w$.

Therefore, by kp_7, there exists a $k \in \mathbf{N}$ such that $w_0 <_{E_1 \sqcup E_2^k} w$. Now, since M is infinite, there will always exist an $n_i \in M$ such that $n_i > k$—in fact, there will be an infinity. Take the smallest such n_i.

We know that $w_0 <_{E_2} w$ and $w_0 <_{E_1 \sqcup E_2^k} w$, hence by kp_6 we get that:

$$w_0 <_{(E_1 \sqcup E_2^k) \sqcup E_2} w.$$

Effectively, this means that $w_0 <_{E_1 \sqcup E_2^{k+1}} w$. We iterate this until we get $w_0 <_{E_1 \sqcup E_2^{n_i}} w$. This leads to a contradiction, because we have already established that $w \leq_{E_1 \sqcup E_2^{n_i}} w_0$. $\square$

This is already a useful result, since it provides a sufficient condition for having a Horn merging operator. We will use it in Chapter 7.

6.2 From operators to assignments

The purpose now is to see under what conditions an operator $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ describes a Horn compliant syncretic assignment mapping every profile E to a total pre-order $\leq_E$, such that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any Horn formula μ. The previous chapter has already made clear that in the Horn case the standard merging postulates are not fit to the task.

We start by adapting the Delgrande and Peppas result for Horn revision [17] to merging. We will show that for a Horn operator satisfying the revision postulates (in our case $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$), there exists a Horn compliant assignment that satisfies $kp_1 - kp_3$. The question, then, will to enforce the rest of the properties for syncretic assignments

6.2.1 Defining a pre-order $\leq_E$ based on a merging operator

In the following we will assume we are given a Horn merging operator $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ that satisfies (at least) $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$.

Take a Horn profile E. We begin by defining a partial pre-order $\leq'_E$ based on Δ.

Definition 25. Define the relation $\leq'_E$ on interpretations as follows:

$$w_1 \leq'_E w_2 \text{ iff } w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E)].$$

If $[\varphi_{w_1,w_2}] = \{w_0, w_1, w_2\}$, where $w_0 = w_1 \cap w_2$, the postulates allow for the situation when $[\Delta_{\varphi_{w_1,w_2}}(E)] = \{w_0\}$, so $w_1 \not\leq'_E w_2$ and $w_2 \not\leq'_E w_1$. In other words, $\leq'_E$ is not necessarily total. We will be looking to extend this pre-order to a total pre-order later on. The very next step, however, is to take the transitive closure.

Definition 26. Define $\leq^*_E$ to be the transitive closure of $\leq'_E$. In other words:

$$w_1 \leq^*_E w_2 \text{ iff there exist interpretations } u_0, \ldots, u_n \text{ such that}$$
$$w_1 = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n = w_2.$$

By construction $\leq^*_E$ is reflexive and transitive, in other words $\leq^*_E$ is a partial pre-order.

Definition 27. Define the following chain:

$$\begin{aligned} \mathcal{S}_0 &= max_{\leq^*_E}\mathcal{W}, \\ \mathcal{S}_1 &= max_{\leq^*_E}(\mathcal{W}\backslash\mathcal{S}_0), \\ &\ldots \\ \mathcal{S}_{i+1} &= max_{\leq^*_E}(\mathcal{W}\backslash \textstyle\bigcup_{j=0}^{i}\mathcal{S}_j). \end{aligned}$$

A couple of observations are in order here.

Observation 9. The chain is constructed by progressively taking elements out of $\mathcal{W}$. Since $\mathcal{W}$ is finite, the chain eventually reaches a point after which all subsequent $\mathcal{S}_i$'s are equal to the empty set. Take m to be the first index where this happens, i.e. $\mathcal{S}_m \neq \emptyset$, $\mathcal{S}_{m+i} = \emptyset$, for any $i \geq 1$.

Now, any non-empty set of interpretations $\mathcal{S}$ has at least one maximal element (this is proved later, see Lemma 7), thus all the $\mathcal{S}_i$'s from $\mathcal{S}_0$ up to $\mathcal{S}_m$ are non-empty. Clearly, the $\mathcal{S}_i$'s up to $\mathcal{S}_m$ are disjoint and their union is $\mathcal{W}$, so it is appropriate to conclude that $\{\mathcal{S}_i \mid 0 \leq i \leq m\}$ is a partition of $\mathcal{W}$.

Observation 10. Since $\{\mathcal{S}_i \mid 0 \leq i \leq m\}$ is a partition of $\mathcal{W}$, any interpretation $w \in \mathcal{W}$ belongs to exactly one $\mathcal{S}_i$. We will call an $\mathcal{S}_i$ to which w belongs a *level of* w. Since $\mathcal{S}_i$, in this sense, is unique, it is appropriate to talk of *the* level of w in this partition.

Definition 28. Define the relation $\leq_E$ as follows:

$$w_1 \leq_E w_2 \text{ iff there are } \mathcal{S}_i \text{ and } \mathcal{S}_j \text{ such that } w_1 \in \mathcal{S}_i, w_2 \in \mathcal{S}_j \text{ and } i \geq j.$$

The claim is that $\leq_E$, thus defined, is Horn compliant and satisfies $kp_1 - kp_3$.

6.2.2 Proving that $\leq_E$ is Horn compliant and faithful

We show in this section, following [17], that $\leq_E$ is Horn compliant and satisfies $kp_1 - kp_3$.

Lemma 3. *Take two interpretations w_1, w_2 such that $w_1 \leq'_E w_2$. Then for any Horn formula μ, if $w_1 \in [\mu]$ and $w_2 \in [\Delta_\mu(E)]$, then $w_1 \in [\Delta_\mu(E)]$.*

Proof. Take $\mu \in \mathcal{L}_H$ and two interpretations w_1, w_2 such that $w_1 \leq'_E w_2$, $w_1 \in [\mu]$ and $w_2 \in [\Delta_\mu(E)]$. We want to show that $w_1 \in [\Delta_\mu(E)]$.

By assumption, $w_2 \in [\Delta_\mu(E)]$ and obviously $w_2 \in [\varphi_{w_1,w_2}]$. Thus, $w_2 \in [\Delta_\mu(E) \wedge \varphi_{w_1,w_2}]$, so $\Delta_\mu(E) \wedge \varphi_{w_1,w_2}$ is consistent. By applying postulates IC_7 and IC_8 it follows that:

$$[\Delta_\mu(E) \wedge \varphi_{w_1,w_2}] = [\Delta_{\mu\wedge\varphi_{w_1,w_2}}(E)].$$

Since $w_2 \in [\Delta_\mu(E) \wedge \varphi_{w_1,w_2}]$, it follows that $w_2 \in [\Delta_{\mu\wedge\varphi_{w_1,w_2}}(E)]$. By IC_0, $[\Delta_{\mu\wedge\varphi_{w_1,w_2}}(E)] \subseteq [\mu \wedge \varphi_{w_1,w_2}]$. We conclude that $w_2 \in [\mu \wedge \varphi_{w_1,w_2}]$, and consequently that $w_2 \in [\mu]$. $\square$

Lemma 4. *For any Horn formula μ, $min_{\leq'_E}[\mu] = [\Delta_\mu(E)]$.*

Proof. We prove the lemma by double inclusion.

(i) $min_{\leq'_E}[\mu] \subseteq [\Delta_\mu(E)]$, for any Horn formula μ.

Take $\mu \in \mathcal{L}_H$ and assume there is an interpretation w_1 such that $w_1 \in min_{\leq'_E}[\mu]$ and $w_1 \notin [\Delta_\mu(E)]$. Since $w_1 \in min_{\leq'_E}[\mu]$, it follows that $[\mu] \neq \emptyset$, so (by IC_1) $[\Delta_\mu(E)] \neq \emptyset$. Take $w_2 \in [\Delta_\mu(E)]$ and let us show, first, that $w_1 \not\leq'_E w_2$ and $w_2 \not\leq'_E w_1$.

> Assume, on the contrary, that $w_1 \leq'_E w_2$. Since $w_1 \in [\mu]$ and $w_2 \in [\Delta_\mu(E)]$, then (by Lemma 3) we get that $w_1 \in [\Delta_\mu(E)]$, which contradicts our starting assumption. So $w_1 \not\leq'_E w_2$.
>
> Assume, now, that $w_2 \leq'_E w_1$. We have just shown that $w_1 \not\leq'_E w_2$, so this means that $w_2 <'_E w_1$. We know that $w_2 \in [\Delta_\mu(E)]$, so (by IC_0) $w_2 \in [\mu]$. Since $w_1 \in [\mu]$, $w_2 <'_E w_1$ contradicts the fact that $w_1 \in min_{\leq'_E}[\mu]$.

Denote $w_0 = w_1 \cap w_2$. We will next show that $[\Delta_{\varphi_{w_1,w_2}}(E)] = \{w_0\}$.

> By IC_0, we have that $[\Delta_{\varphi_{w_1,w_2}}(E)] \subseteq [\varphi_{w_1,w_2}]$. Since $[\varphi_{w_1,w_2}] \neq \emptyset$, by IC_1 we get $[\Delta_{\varphi_{w_1,w_2}}(E)] \neq \emptyset$. Also, we have that $w_1 \notin [\Delta_{\varphi_{w_1,w_2}}(E)]$, because otherwise we would have that $w_1 \leq'_E w_2$, and we have just shown this cannot happen. Similarly, $w_2 \notin [\Delta_{\varphi_{w_1,w_2}}(E)]$. So the only possibility left is that $[\Delta_{\varphi_{w_1,w_2}}(E)] = \{w_0\}$.

As a direct consequence, we have that:

$$[\Delta_{\varphi_{w_1,w_2}}(E)] \cap [\varphi_{w_0,w_1}] = [\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_0,w_1}] = \{w_0\}. \tag{6.1}$$

From this it follows that $\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_0,w_1}$ is consistent, so by IC_7 and IC_8:

$$[\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_0,w_1}] = [\Delta_{\varphi_{w_1,w_2} \wedge \varphi_{w_0,w_1}}(E)]. \tag{6.2}$$

Since $w_0 \subseteq w_1$ we can say that, while $[\varphi_{w_0,w_1}]$ surely contains w_0 and w_1, it does not contain more than that—it is already closed under intersection. On the other hand, $[\varphi_{w_1,w_2}]$ contains w_0, w_1 and w_2. So $[\varphi_{w_0,w_1}] \subseteq [\varphi_{w_1,w_2}]$, and it follows that:

$$[\varphi_{w_1,w_2}] \cap [\varphi_{w_0,w_1}] = [\varphi_{w_0,w_1}].$$

So $[\varphi_{w_1,w_2} \wedge \varphi_{w_0,w_1}] = [\varphi_{w_0,w_1}]$, which is to say that $\varphi_{w_1,w_2} \wedge \varphi_{w_0,w_1} \equiv \varphi_{w_0,w_1}$. By IC_3, we get:

$$[\Delta_{\varphi_{w_1,w_2} \wedge \varphi_{w_0,w_1}}(E)] = [\Delta_{\varphi_{w_0,w_1}}(E)]. \tag{6.3}$$

Connecting results 6.1, 6.2 and 6.3, we get:

$$[\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_0,w_1}] = [\Delta_{\varphi_{w_0,w_1}}(E)] = \{w_0\}. \tag{6.4}$$

From 6.4 and the definition of $\leq'_E$ (see Definition 25) we can infer that $w_0 \leq'_E w_1$. Also, we get that $w_1 \not\leq'_E w_0$ (because otherwise, by the same definition, we would get that $w_1 \in [\Delta_{\varphi_{w_0,w_1}}(E)]$). Putting these two things together, we have just shown that:

$$w_0 <'_E w_1. \tag{6.5}$$

Let us now focus our attention on something slightly different: we have that $w_1 \in [\mu]$ (because $w_1 \in min_{\leq'_E}[\mu]$ by assumption) and $w_2 \in [\mu]$ (because $w_2 \in [\Delta_\mu(E)]$ by assumption, and $[\Delta_\mu(E)] \subseteq [\mu]$ by IC_0). Also, μ is a Horn formula, so $w_1 \cap w_2 = w_0 \in [\mu]$. And because $w_1 \in min_{\leq'_E}[\mu]$, it follows that $w_1 \leq'_E w_0$. With result 6.5, this leads to a contradiction.

(ii) $[\Delta_\mu(E)] \subseteq min_{\leq'_E}[\mu]$, for any Horn formula μ.

Take $\mu \in \mathcal{L}_H$ and $w_1 \in [\Delta_\mu(E)]$. We will show that $w_1 \leq'_E w_2$, for any $w_2 \in [\mu]$. So take a $w_2 \in [\mu]$.

Obviously $w_1 \in [\varphi_{w_1,w_2}]$, and by assumption $w_1 \in [\Delta_\mu(E)]$. So $[\Delta_\mu(E) \wedge \varphi_{w_1,w_2}] \neq \emptyset$, which is to say that $\Delta_\mu(E) \wedge \varphi_{w_1,w_2}$ is consistent. By IC_7 and IC_8, it follows that:

$$[\Delta_\mu(E) \wedge \varphi_{w_1,w_2}] = [\Delta_{\mu \wedge \varphi_{w_1,w_2}}(E)]. \tag{6.6}$$

Take, again, $w_0 = w_1 \cap w_2$. Since $w_1, w_2 \in [\mu]$ and $[\mu]$ is a Horn formula, we get that $w_0 \in [\mu]$. So $[\varphi_{w_1,w_2}] \subseteq [\mu]$ and hence $[\mu \wedge \varphi_{w_1,w_2}] = [\varphi_{w_1,w_2}]$. By IC_3, now, we get:

$$[\Delta_{\mu \wedge \varphi_{w_1,w_2}}(E)] = [\Delta_{\varphi_{w_1,w_2}}(E)]. \tag{6.7}$$

Connecting 6.6 and 6.7, we get that:

$$[\Delta_\mu(E) \wedge \varphi_{w_1,w_2}] = [\Delta_{\varphi_{w_1,w_2}}(E)]. \tag{6.8}$$

By assumption $w_1 \in [\Delta_\mu(E)]$ and $w_1 \in [\varphi_{w_1,w_2}]$, so by 6.8 we get that $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E)]$, so (by Definition 25 of $\leq'_E$) $w_1 \leq'_E w_2$. This shows that $w_1 \in min_{\leq'_E}[\mu]$. $\square$

Lemma 5. *If* $w_1 \leq'_E w_2 \leq'_E \cdots \leq'_E w_n \leq'_E w_1$*, then* $w_1 \leq'_E w_n$.

Proof. Take interpretations $w_1, \ldots, w_n$ such that $w_1 \leq'_E w_2 \leq'_E \cdots \leq'_E w_n \leq'_E w_1$. From Definition 25 of $\leq'_E$ and the given inequalities, we get:

$$\begin{gathered} w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E)], \\ w_2 \in [\Delta_{\varphi_{w_2,w_3}}(E)], \\ \ldots \\ w_n \in [\Delta_{\varphi_{w_n,w_1}}(E)]. \end{gathered}$$

It follows that:

$$\begin{gathered} w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_n,w_1}], \\ w_2 \in [\Delta_{\varphi_{w_2,w_3}}(E) \wedge \varphi_{w_1,w_2}], \\ \ldots \\ w_{n-1} \in [\Delta_{\varphi_{w_{n-1},w_n}}(E) \wedge \varphi_{w_{n-2},w_{n-1}}], \\ w_n \in [\Delta_{\varphi_{w_n,w_1}}(E) \wedge \varphi_{w_{n-1},w_n}]. \end{gathered}$$

Applying $Acyc$, we get that:

$$[\Delta_{\varphi_{w_n,w_1}}(E) \wedge \varphi_{w_1,w_2}] \neq \emptyset$$

This allows us to apply IC_8 which, together with IC_8, gives us:

$$[\Delta_{\varphi_{w_n,w_1}}(E) \wedge \varphi_{w_1,w_2}] = [\Delta_{\varphi_{w_n,w_1} \wedge \varphi_{w_1,w_2}}(E)] \tag{6.9}$$

At the same time we have $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_n,w_1}]$, which implies that $\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_n,w_1}$ is consistent. Applying IC_7 and IC_8 again, we infer:

$$[\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_n,w_1}] = [\Delta_{\varphi_{w_1,w_2} \wedge \varphi_{w_n,w_1}}(E)] \tag{6.10}$$

Obviously $\varphi_{w_n,w_1} \wedge \varphi_{w_1,w_2} \equiv \varphi_{w_1,w_2} \wedge \varphi_{w_n,w_1}$, so we can apply IC_3 and connect 6.9 and 6.10 to get:

$$[\Delta_{\varphi_{w_n,w_1}}(E) \wedge \varphi_{w_1,w_2}] = [\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_n,w_1}] \tag{6.11}$$

We know that $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_n,w_1}]$, so from result 6.10 we get that $w_1 \in [\Delta_{\varphi_{w_n,w_1}}(E) \wedge \varphi_{w_1,w_2}]$. *A fortiori*, $w_1 \in [\Delta_{\varphi_{w_n,w_1}}(E)]$ and $w_1 \leq'_E w_n$, which is we wanted to show. $\square$

Lemma 6. *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 $u_0, \ldots, u_n$ 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:

$$w_2 \leq'_E u_1 \leq'_E \cdots \leq'_E u_n \leq'_E w_1 \leq'_E w_2 \tag{6.12}$$

From Lemma 5 and result 6.12 we get that $w_2 \leq'_E w_1$, which together with the assumption that $w_1 <'_E w_2$ leads to a contradiction. $\square$

Lemma 7. *For every set of interpretations $\mathcal{S} \neq \emptyset$, $max_{\leq_E^*}\mathcal{S} \neq \emptyset$.*

Proof. Take a set of interpretations $\mathcal{S} \neq \emptyset$ and assume that $max_{\leq_E^*}\mathcal{S} = \emptyset$. Then for any $w \in \mathcal{S}$ there is $w' \in \mathcal{S}$ such that $w <_E^* w'$.

So, starting with $w_0 \in \mathcal{S}$ one can build up a sequence:

$$w_0 <_E^* w_1 <_E^* w_2 <_E^* \ldots .$$

Since $\mathcal{S}$ is finite, some of the interpretations in the sequence will have to repeat themselves. In other words, there exist interpretations w_i, w_j such that:

$$w_i <_E^* w_{i+1} <_E^* \cdots <_E^* w_{j-1} <_E^* w_j <_E^* w_i.$$

Applying Lemma 6 it follows that $w_i \leq_E^* w_j$, which contradicts the fact that $w_j <_E^* w_i$. $\square$

Lemma 8. *The pre-order $\leq_E$ is total and extends $\leq_E^*$, in the sense that if $w_1 \leq_E^* w_2$, then $w_1 \leq_E w_2$.*

Proof. It was argued in Observation 9 that the set of levels $\{\mathcal{S}_i \mid 0 \leq i \leq m\}$ forms a partition of $\mathcal{W}$. Thus, any two interpretations w_1, w_2 belong to some (unique) level $\mathcal{S}_i$, $\mathcal{S}_j$, respectively. Since $i \geq j$ or $j \geq i$, it follows that $w_1 \leq_E w_2$ or $w_2 \leq_E w_1$. In other words, $\leq_E$ is total.

Suppose $w_1 \leq_E^* w_2$ and $w_1 \in \mathcal{S}_i$, $w_2 \in \mathcal{S}_j$. We want to show that $w_1 \leq_E w_2$, or in other words that $i \geq j$.

Assume, on the contrary, that $j > i$. We have that $w_1 \in \mathcal{S}_i$, which means (by Definition 27 of $\mathcal{S}_i$) that $w_1 \in \max_{\leq_E^*}(W \setminus \bigcup_{k=0}^{i-1} \mathcal{S}_k)$. By assumption, $w_2 \in \mathcal{S}_j$, with $j > i$. The fact that $j > i$ tells us that $\mathcal{S}_j$ comes later in the sequence $\mathcal{S}_0, \mathcal{S}_1, \ldots$ than $\mathcal{S}_i$, which means that $w_2 \in \mathcal{W} \setminus \bigcup_{k=0}^{i-1} \mathcal{S}_k$. Since w_1 is a maximal element of this set and (by assumption) $w_1 \leq_E^* w_2$, it follows that $w_2 \leq_E^* w_1$, and w_2 is also a maximal element in $\mathcal{W} \setminus \bigcup_{k=0}^{i-1} \mathcal{S}_k$.

To see why this last statement holds, take $w_3 \in \mathcal{W} \setminus \bigcup_{k=0}^{i-1} \mathcal{S}_k$ such that $w_2 \leq_E^* w_3$. Since $w_1 \leq_E^* w_2$ and $\leq_E^*$ is transitive (by definition), we have $w_1 \leq_E^* w_3$. We know that w_1 is maximal in $\mathcal{W} \setminus \bigcup_{k=0}^{i-1} \mathcal{S}_k$, so $w_3 \leq_E^* w_1$. Using transitivity again, we get that $w_3 \leq_E^* w_2$, so w_2 is maximal in $\mathcal{W} \setminus \bigcup_{k=0}^{i-1} \mathcal{S}_k$.

This means, to recap, that w_1 and w_2 are both maximal $\mathcal{W} \setminus \bigcup_{k=0}^{i-1} \mathcal{S}_k$. So they both belong to the same level. But this is a contradiction. $\square$

Lemma 9. *If w_1 and w_2 are interpretations, it holds that $w_1 \approx_E w_2$ iff $w_1, w_2 \in \mathcal{S}_i$.*

Proof. We know that $w_1 \approx_E w_2$ iff $w_1 \leq_E w_2$ and $w_2 \leq_E w_1$, which—according to Definition 28 of $\leq_E$—is equivalent to there being some levels $\mathcal{S}_i$ and $\mathcal{S}_j$ such that $w_1 \in \mathcal{S}_i, w_2 \in \mathcal{S}_j$, and $i \geq j$ and $j \geq i$. This, in turn, is equivalent to $i = j$, which means that w_1 and w_2 are on the same level. $\square$

Lemma 10. *If w_1 and w_2 are interpretations, it holds that $w_1 <_E w_2$ iff $w_1 \in \mathcal{S}_i, w_2 \in \mathcal{S}_j$ and $i > j$.*

Proof. Though it is not difficult to grasp why this is true, for the sake of being rigorous let us prove the result one direction at a time.

(i) If $w_1 <_E w_2$, then $w_1 \in \mathcal{S}_i, w_2 \in \mathcal{S}_j$ and $i > j$.

Unpacking the hypothesis, we have that $w_1 \leq_E w_2$ and $w_2 \not\leq_E w_1$. The first of these statements guarantees (by Definition 28 of $\leq_E$) that there are $\mathcal{S}_i, \mathcal{S}_j$, such that $w_1 \in \mathcal{S}_i, w_2 \in \mathcal{S}_j$[1] and $i \geq j$. Suppose $i = j$. This would imply that $w_2 \leq_E w_1$, which contradicts the second statement. Hence, $i > j$.

(ii) If $w_1 \in \mathcal{S}_i, w_2 \in \mathcal{S}_j$ and $i > j$, then $w_1 <_E w_2$.

Using the hypothesis and Definition 28 of $\leq_E$, we get that $w_1 \leq_E w_2$. Suppose, also, that $w_2 \leq_E w_1$. Then $w_1 \approx_E w_2$, and (by Lemma 10), it follows that $i = j$. But this is a contradiction, so $w_2 \not\leq_E w_1$. $\square$

Lemma 11. *If $\mathcal{S}_k$ is the last element in the sequence $\mathcal{S}_0, \mathcal{S}_1, \ldots$ that intersects $[\mu]$, then $min_{\leq_E}[\mu] = \mathcal{S}_k \cap [\mu]$.*

Proof. We will show both directions.

(i) $min_{\leq_E}[\mu] \subseteq \mathcal{S}_k \cap [\mu]$.

Take $w \in min_{\leq_E}[\mu]$ and suppose $w \notin \mathcal{S}_k \cap [\mu]$. Since obviously $w \in [\mu]$, it can only be the case that $w \notin \mathcal{S}_k$. Since $\mathcal{S}_k$ is by hypothesis the last element in the sequence $\mathcal{S}_0, \mathcal{S}_1, \ldots$ that intersects $[\mu]$, it follows that $w \in \mathcal{S}_i$, for some $i < k$. If we take now some $w' \in \mathcal{S}_k \cap [\mu]$, we get (by Lemma 10) that $w' <_E w$. But this contradicts the fact that w is $\leq_E$-minimal in $[\mu]$.

(ii) $\mathcal{S}_k \cap [\mu] \subseteq min_{\leq_E}[\mu]$.

Take $w \in \mathcal{S}_k \cap [\mu]$ and some $w' \in [\mu]$ such that $w' \leq_E w$. To show that w is also $\leq_E$-minimal in $[\mu]$, we need to show that $w \leq_E w'$.

Suppose, on the contrary, that $w' <_E w$. This means (by Lemma 10) that $w' \in \mathcal{S}_l$, for some $l > k$. But this contradicts the fact that $\mathcal{S}_k$ is the last element in the sequence $\mathcal{S}_0, \mathcal{S}_1, \ldots$ that intersects $[\mu]$. $\square$

[1] As mentioned in Observation 10, each interpretation belongs to exactly one level.

Lemma 12. *For any Horn formula μ, $min_{\leq_E}[\mu] = min_{\leq'_E}[\mu]$.*

Proof. If μ is inconsistent, then $[\mu] = \emptyset$ and the equality obviously holds. Assume, then, that $[\mu] \neq \emptyset$ and take $\mathcal{S}_k$ to be the last set in the sequence $\mathcal{S}_0, \ldots, \mathcal{S}_m$ that intersects $[\mu]$. By Lemma 9, we have that $min_{\leq_E}[\mu] = \mathcal{S}_k \cap [\mu]$, so what we need to prove is that $\mathcal{S}_k \cap [\mu] = min_{\leq'_E}[\mu]$. We will do this by showing the double inclusion.

(i) $\mathcal{S}_k \cap [\mu] \subseteq min_{\leq'_E}[\mu]$.

Assume, by contradiction, that there is $w_1 \in \mathcal{S}_k \cap [\mu]$ such that $w_1 \notin min_{\leq'_E}[\mu]$. Then there exists $w_2 \in [\mu]$ such that $w_2 <'_E w_1$. From this and Lemma 6 it follows that $w_2 <^*_E w_1$.

Because $w_2 <^*_E w_1$, w_2 cannot be in $\mathcal{S}_k$. Since $\mathcal{S}_k$ is the last set in the sequence that contains models of $[\mu]$, it follows that w_2 must come earlier in the sequence, i.e. $w_2 \in \mathcal{S}_j$, for some $j < k$ (see Figure [...]). But if $w_2 \in \mathcal{S}_j$, then for any $l > j$ and for any $w \in \mathcal{S}_l$, we have that $w <^*_E w_2$. In particular, this implies that $w_1 <'_E w_2$, because $k > j$. This is a contradiction.

(ii) $min_{\leq'_E}[\mu] \subseteq \mathcal{S}_k \cap [\mu]$.

Assume, by contradiction, that there is $w_1 \in min_{\leq'_E}[\mu]$ such that $w_1 \notin \mathcal{S}_k \cap [\mu]$. Since, obviously, $w_1 \in [\mu]$, it has to be that $w_1 \notin \mathcal{S}_k$. But $\mathcal{S}_k$ is the last element in the sequence $\mathcal{S}_0, \mathcal{S}_1, \ldots$ intersecting $[\mu]$, so $w_1 \in \mathcal{S}_j$, for some $j < k$.

Take $w_2 \in \mathcal{S}_k \cap [\mu]$. Since $\mathcal{S}_k$ comes after $\mathcal{S}_j$ in the sequence $\mathcal{S}_0, \mathcal{S}_1, \ldots$, then (by Lemma 10) it is the case that $w_2 <_E w_1$. This implies that $w_1 \not\leq'_E w_2$. Why? Assume, on the contrary, that $w_1 \leq'_E w_2$. Then $w_1 \leq^*_E w_2$ and $w_1 \leq_E w_2$ (since, by Lemma 8, $\leq_E$ extends $\leq^*_E$). But this contradicts the fact that $w_2 <_E w_1$.

In order to generate the final contradiction, let us show that $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E)]$.

> Assume, on the contrary, that $w_1 \notin [\Delta_{\varphi_{w_1,w_2}}(E)]$. Then $w_1 \not\leq'_E w_2$. Notice, we also have that $w_2 \not\leq'_E w_1$: if, on the contrary, it would hold that $w_2 \leq'_E w_1$, then from this and the fact that $w_1 \in min_{\leq'_E}[\mu]$ and $w_2 \in [\mu]$ it would follow that $w_1 \leq'_E w_2$, which we already know does not hold.
>
> So $w_2 \not\leq'_E w_1$, and hence $w_2 \notin [\Delta_{\varphi_{w_1,w_2}}(E)]$. At the same time $[\varphi_{w_1,w_2}] \neq \emptyset$, so (by IC_1) $[\Delta_{\varphi_{w_1,w_2}}(E)] \neq \emptyset$, and (by IC_0) $[\Delta_{\varphi_{w_1,w_2}}(E)] \subseteq [\varphi_{w_0,w_1}]$. It follows that $[\Delta_{\varphi_{w_1,w_2}}(E)] = \{w_0\}$, where $w_0 = w_1 \cap w_2$.
>
> This shows that $\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_0,w_1}$ is consistent, and by IC_7 and IC_8:
>
> $$[\Delta_{\varphi_{w_1,w_2}}(E) \wedge \varphi_{w_0,w_1}] = [\Delta_{\varphi_{w_1,w_2} \wedge \varphi_{w_0,w_1}}(E)] = \{w_0\} \tag{6.13}$$
>
> At the same time, $[\varphi_{w_1,w_2} \wedge \varphi_{w_0,w_1}] = [\varphi_{w_0,w_1}]$. So by IC_3:
>
> $$[\Delta_{\varphi_{w_1,w_2} \wedge \varphi_{w_0,w_1}}(E)] = [\Delta_{\varphi_{w_0,w_1}}(E)] \tag{6.14}$$

Connecting equalities 6.13 and 6.14, we get that:

$$[\Delta_{\varphi_{w_0,w_1}}(E)] = \{w_0\},$$

which implies that $w_0 <'_E w_1$. Because $\mu \in \mathcal{L}_H$ and $w_1, w_2 \in [\mu]$, it is also the case that $w_0 \in [\mu]$. In the end this leads to a contradiction, because $w_1 \in min_{\leq'_E}[\mu]$.

We have shown, then, that $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E)]$. From this it follows that $w_1 \leq'_E w_2$. But we have already argued that $w_1 \not\leq'_E w_2$, so we have derived a contradiction. □

Proposition 18. *The pre-order $\leq_E$ is Horn compliant and it corresponds to the operator Δ, in the sense that for any Horn formula μ, $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$.*

Proof. By Lemma 4, $[\Delta_\mu(E)] = min_{\leq'_E}[\mu]$. By Lemma 12, $min_{\leq'_E}[\mu] = min_{\leq_E}[\mu]$. Putting these two things together, we have that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, so $\leq_E$ corresponds to the operator Δ.

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. □

6.2.3 Is $\leq_E$ syncretic?

We started by assuming we have a Horn merging operator Δ that satisfies $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$. Based on Δ, we defined in Section 6.2.1 an assignment that assigns to each profile E a total pre-order $\leq_E$. We showed in Section 6.2.2 that $\leq_E$ is Horn compliant, for any Horn profile E, and that it satisfies:

$$[\Delta_\mu(E)] = min_{\leq_E}[\mu].$$

In doing this, we have been following the lead of Delgrande and Peppas in [17]. But their work was tailored for Horn revision: in so doing, they were looking for Horn compliant *faithful* assignments, whereas we want the assignments to be syncretic. Let us convince ourselves, then, if the assignment defined in Section 6.2.1 is suitable for merging and satisfies properties $kp_1 - kp_6$.

To anticipate, the answer is no: we will show that $kp_1 - kp_3$ hold, but that—reconsidering the examples from the previous chapter—$kp_4 - kp_6$ are problematic.

First, let us show that for any profile E and interpretations w_1, w_2, the following holds:

(kp_1) if $w_1 \in [E]$ and $w_2 \in [E]$, then $w_1 \simeq w_2$;

(kp_2) if $w_1 \in [E]$ and $w_2 \notin [E]$, then $w_1 < w_2$.

If E is inconsistent, then $[E] = \emptyset$ and $kp_1 - kp_2$ hold trivially. If E is consistent, take a tautology $[\mu]$, for instance $\mu = p \to p$. Then $[\mu] = \mathcal{W}$, so obviously $[E \wedge \mu] \neq \emptyset$. Then, by IC_2:

$$[\Delta_\mu(E)] = [\bigwedge E \wedge \mu] = [E] \cap \mathcal{W} = [E].$$

This means that $min_{\leq_E}[\mu] = [E]$, and—since $[\mu] = \mathcal{W}$—it follows that $min_{\leq_E}\mathcal{W} = [E]$. Thus, the minimal elements of the entire set of interpretations $\mathcal{W}$ with respect to the pre-order $\leq_E$ are exactly the models of E, which is what kp_1 and kp_2 amount to.

(kp_3) If $E_1 \equiv E_2$, then $\leq_{E_1} = \leq_{E_2}$.

Let us observe, first, that if $E_1 \equiv E_2$, then $\leq'_{E_1} = \leq'_{E_2}$. For this, take two interpretations w_1 and w_2 and suppose that $w_1 \leq'_{E_1} w_2$. Then, by Definition 25 of $\leq'_{E_1}$, we get that $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E_1)]$. Since $E_1 \equiv E_2$, by IC_3 it follows that $[\Delta_{\varphi_{w_1,w_2}}(E_1)] = [\Delta_{\varphi_{w_1,w_2}}(E_2)]$ and $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E_2)]$, hence $w_1 \leq'_{E_2} w_2$. Applying this reasoning in reverse, we get that $w_1 \leq'_{E_1} w_2$ if and only if $w_1 \leq'_{E_2} w_2$. This shows that $\leq'_{E_1} = \leq'_{E_2}$.

Quite clearly, then, the transitive closures $\leq^*_{E_1}$, $\leq^*_{E_2}$ of $\leq'_{E_1}$ and $\leq'_{E_2}$, respectively, are equal. And, since they depend only on $\leq^*_{E_1}$ and $\leq^*_{E_2}$, it must be the case that $\leq_{E_1}$ and $\leq_{E_2}$ are equal as well. This shows that kp_3 holds.

Up to now we have been assuming that Δ satisfies only $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$. This ensures that $kp_1 - kp_3$ are true. With regard to $kp_4 - kp_6$, even in the normal case, we need more constraints on Δ. In the normal case, $IC_4 - IC_6$ enforce $kp_4 - kp_6$, but does this hold here? Unfortunately, no. Recall that Section 5.4 of the previous chapter we presented a series of counter-examples: assignments that otherwise are Horn compliant and syncretic, except for not satisfying kp_4, kp_5 and kp_6, respectively. At the same time, it was shown that the merging operator based on those assignments satisfies IC_4, IC_5 and IC_6, respectively. We have also shown no assignments equivalent to the problematic one. This shows that $IC_4 - IC_6$ are not enough to enforce $kp_4 - kp_6$.

6.2.4 An alternative for IC_4 and kp_4

Example 13 in the previous chapter illustrates why it is difficult to force models of K_1 and K_2 to be on an equal footing in $\leq_{\{K_1,K_2\}}$ with our logical postulates: the operator Δ might end up talking about the elements in the closure of $[K_1] \cup [K_2]$, while the elements of $[K_1]$ and $[K_2]$ themselves remain inaccessible. We can try to negotiate the intuition of fairness that IC_4 and kp_4 encode by looking at modified versions of these statements.

As Lemma 1 shows, the case that needs special attention is the one where $\{K_1, K_2\}$ is inconsistent: if $\{K_1, K_2\}$ is consistent, then an assignment which satisfies $kp_1 - kp_3$ automatically also satisfies kp_4.

In the normal case (of full propositional logic), IC_4 takes care of kp_4 when $\{K_1, K_2\}$ is inconsistent by applying to $\mu = K_1 \vee K_2$: IC_4, in this case, makes sure that the minimal elements of $[K_1] \cup [K_2]$ do not consist of models belonging to one knowledge base alone. In the Horn case, however, $[K_1] \cup [K_2]$ is not necessarily representable by a Horn formula. We *can*, for instance, consider $Cl_\cap([K_1] \cup [K_2])$, but if the models of $[K_1] \cup [K_2]$ are inaccessible to the merging operator (as it happens in Example 13), then a merging operator Δ may satisfy IC_4 while having no control over the inter-

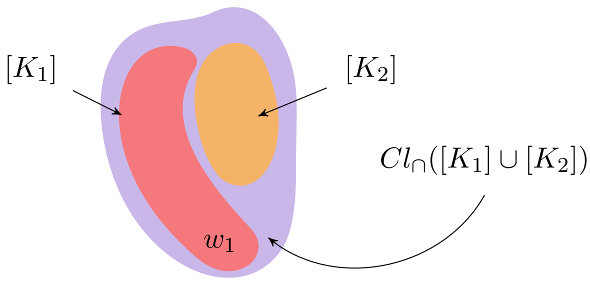

Figure 6.1: $Cl_\cap([K_1] \cup [K_2])$

pretations in $[K_1] \cup [K_2]$, which will then be permitted to vary in ways that contradict kp_4.

It turns out, though, that in the case when $\{K_1, K_2\}$ is inconsistent we can prove a slightly amended property:

(kp_4') If $\{K_1, K_2\}$ is inconsistent, then for any $w_1 \in [K_1]$ there exists an interpretation $w_2 \in Cl_\cap([K_1] \cup [K_2]) \backslash [K_1]$ such that $w_2 \leq_{\{K_1,K_2\}} w_1$.

Proposition 19. *IC_4 implies kp_4'.*

Proof. Take $w_1 \in [K_1]$ and suppose, on the contrary, that:

$$\text{for any } w \in Cl_\cap([K_1] \cup [K_2]) \backslash [K_1], w_1 <_{\{K_1,K_2\}} w. \tag{6.15}$$

The situation is described visually in Figure 6.1. Consider now a Horn formula μ such that $[\mu] = Cl_\cap([K_1] \cup [K_2])$. One thing we can immediately say about $[\mu]$ is that it is consistent, since $[K_1]$ is non-empty ($w_1 \in [K_1]$ by hypothesis) and $[K_1] \subseteq [\mu]$. It follows from IC_1 that $min_{\leq_{\{K_1,K_2\}}}[\mu] \neq \emptyset$.

We can use this to show that $min_{\leq_{\{K_1,K_2\}}}[\mu] \cap [K_1] \neq \emptyset$. Suppose, on the contrary, that we have:

$$min_{\leq_{\{K_1,K_2\}}}[\mu] \cap [K_1] = \emptyset. \tag{6.16}$$

Because $min_{\leq_{\{K_1,K_2\}}}[\mu] \neq \emptyset$, there exists $w' \in min_{\leq_{\{K_1,K_2\}}}[\mu]$. Since $w_1 \in [\mu]$, it follows that:

$$w' \leq_{\{K_1,K_2\}} w_1 \tag{6.17}$$

Notice that w' cannot be in $[K_1]$ because of assumption 6.16. But $w' \in [\mu]$, and hence $w' \in [\mu] \backslash [K_1]$. This, together with our original assumption 6.15, implies that:

$$w_1 <_{\{K_1,K_2\}} w'. \tag{6.18}$$

Now, inequalities 6.17 and 6.18 lead to a contradiction, hence assumption 6.16 is false.

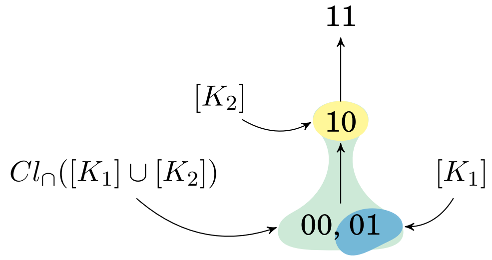

Figure 6.2: $\leq_{\{K_1,K_2\}}$. kp'_4 is true, but IC_4 is not.

Next, notice that $K_1 \models \mu$, $K_2 \models \mu$. Also, we have just proven that $\Delta_\mu(\{K_1, K_2\}) \wedge K_1$ is consistent. Applying IC_4, we get that $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is consistent. In other words, there exists $w_2 \in min_{\leq_{\{K_1,K_2\}}}[\mu] \cap [K_2]$.

From $w_2 \in min_{\leq_{\{K_1,K_2\}}}[\mu]$ and $w_1 \in [\mu]$ it follows that $w_2 \leq_{\{K_1,K_2\}} w_1$. It is also the case that $w_2 \in [K_2]$, and since we have assumed that $\{K_1, K_2\}$ is inconsistent, we have that $[K_1] \cap [K_2] = \emptyset$. This implies that $w_2 \notin [K_1]$. Consequently, $w_2 \in [\mu] \backslash [K_1]$. Our original assumption 6.15 yields $w_1 <_{\{K_1,K_2\}} w_2$ and we have derived a contradiction, showing that assumption 6.15 is false. □

The sad truth, however, is that in the other direction kp'_4 does not guarantee IC_4: if we replace kp_4 with kp'_4 in the definition of a Horn compliant syncretic assignment, then a merging operator defined in the usual way will not necessarily satisfy IC_4. The following example illustrates this fact.

Example 17. Take two knowledge bases $[K_1] = \{01\}$ and $[K_2] = \{10\}$ and a Horn compliant assignment that assigns to $\{K_1, K_2\}$ the pre-order in Figure 6.2. We assume the assignment satisfies properties $kp_1 - kp_3, kp'_4$ and $kp_5 - kp_6$. Notice that $\leq_{\{K_1,K_2\}}$ in Figure 6.2 does, as a matter of fact, satisfy kp'_4: in the interesting case, we have that for $w_1 = 01 \in [K_1]$, there exists $w_2 = 00 \in Cl_\cap([K_1] \cup [K_2]) \backslash [K_1]$ such that $w_2 \leq_{\{K_1,K_2\}} w_1$.

If we define an operator Δ as $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any Horn profile E and Horn formula μ, we notice that IC_4 is not true. Take a Horn formula $[\mu] = \{00, 01, 10\}$, Clearly, $K_1 \models \mu$ and $K_2 \models \mu$, and $[\Delta_\mu(\{K_1, K_2\}) \wedge K_1] = \{01\}$, so $\Delta_\mu(\{K_1, K_2\}) \wedge K_1$ is consistent. Now IC_4 would require that $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is also consistent. However, $[\Delta_\mu(\{K_1, K_2\}) \wedge K_2] = \emptyset$. □

This could be fixed by adopting the weaker version of IC_4 for the Horn case below. Recalling a notion introduced in Section 2.2, if K_1 and K_2 are Horn knowledge bases, $K_{1,2}$ denotes a Horn knowledge base such that $[K_{1,2}] = Cl_\cap([K_1] \cup [K_2])$. Then, for any Horn formula μ, we require that:

(IC'_4) If K_1 and K_2 are consistent, and if $K_1 \models \mu$, $K_2 \models \mu$, then the following holds: if $\Delta_\mu(\{K_1, K_2\}) \wedge K_1$ is consistent, then $\Delta_\mu(\{K_1, K_2\}) \wedge K_{1,2}$ is consistent.

Observation 11. It is easy to see that IC'_4 and kp'_4 are both weaker versions of IC_4 and kp_4, respectively: that is to say, in the case of normal propositional logic $IC_{4'}$ is implied by IC_4; also, kp_4 implies kp'_4.

It can be shown now that kp'_4 and IC'_4 correspond in the Horn case.

Proposition 20. *If Δ is a Horn operator that satisfies $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$ and there is a Horn compliant assignment that assigns to any Horn profile E a total pre-order $\leq_E$ such that for any Horn formula μ, $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, and if the assignment satisfies $kp_1 - kp_3$, then Δ satisfies IC'_4 if and only if the assignment satisfies kp'_4.*

Proof. First, notice that if $\{K_1, K_2\}$ is consistent, the result follows immediately. Hence, we concentrate on the case when $\{K_1, K_2\}$ is inconsistent. We prove this result one direction at a time.

(i) If Δ satisfies IC'_4, then the assignment satisfies kp'_4.

The proof for this is essentially similar to the one for Proposition 19, except for the fact that we should use IC'_4 instead of IC_4 and $K_{1,2}$ instead of K_2. We do not reiterate the proof here.

(ii) If the assignment satisfies kp'_4, then Δ satisfies IC'_4.

Take a Horn formula μ such that $K_1 \models \mu$, $K_2 \models \mu$ and suppose $\Delta_\mu(\{K_1, K_2\}) \wedge K_1$ is consistent. Then there exists $w_1 \in [\Delta_\mu(\{K_1, K_2\}) \wedge K_1]$, which means that $w_1 \in [K_1]$ and (by kp'_4) there is $w_2 \in Cl_\cap([K_1] \cup [K_2]) \backslash [K_1]$ such that $w_2 \leq_{\{K_1,K_2\}} w_1$. Since $[K_1] \subseteq [\mu]$ and $[K_2] \subseteq [\mu]$, it follows that $w_1, w_2 \in [\mu]$, and because $w_1 \in min_{\leq_{\{K_1,K_2\}}}[\mu]$, then also $w_2 \in min_{\leq_{\{K_1,K_2\}}}[\mu]$. Because $w_2 \in Cl_\cap([K_1] \cup [K_2]) \backslash [K_1]$ and $Cl_\cap([K_1] \cup [K_2]) \backslash [K_1] \subseteq Cl_\cap([K_1] \cup [K_2])$, it follows that $w_2 \in [K_{1,2}]$. The conclusion, now, is easily derived. □

6.3 Capturing elusive orderings

The stumbling block in the case of Horn merging are pairs of interpretations w_1, w_2 not closed under intersection, which—when not in $min_{\leq_E}[\varphi_{w_1,w_2}]$—escape the notice of a Δ under the current formulation of the postulates. We would like, then, to characterize the case when $w_1 <_E w_2$, for some profile E. Note that in the normal case this is easy, since we have:

$$w_1 <_E w_2 \text{ iff } w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E)] \text{ and } w_2 \notin [\Delta_{\varphi_{w_1,w_2}}(E)].$$

In the Horn case, this is still true if $w_1 \subseteq w_2$ or $w_2 \subseteq w_1$. However, if $\{w_1, w_2\}$ is not closed under intersection, the equivalence breaks down.

To be more precise, one direction still holds: if $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E)]$ then $w_1 \leq_E w_2$. Why? Because $w_1 \leq'_E w_2$ by definition, and this carries over to $\leq_E$ by Lemma 8.

But the converse does not hold: if $w_1 \leq_E w_2$, it is not necessarily the case that $w_1 \in [\Delta_{\varphi_{w_1,w_2}}(E)]$. We can have any sort of arrangement between w_1 and w_2 with neither of them being in $[\Delta_{\varphi_{w_1,w_2}}(E)]$ (this happens repeatedly throughout the counterexamples).

The moral is that in the Horn case we cannot always judge the order between w_1 and w_2 simply by looking at $[\Delta_{\varphi_{w_1,w_2}}(E)]$, which is particularly worrying for properties $kp_5 - kp_6$. So is the situation hopeless? Fortunately this is not the case. It turns out we can still control the order of w_1 and w_2, albeit in a more roundabout way.

We will rely on the construction of $\leq_E$ on the basis of postulates $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$ from Section 6.2.1. Let us start by fixing a few facts about elements in $\leq_E$.

6.3.1 Some facts about $\leq_E$

By Lemma 10, $w_1 <_E w_2$ if and only if w_1 comes later than w_2 in the sequence $\mathcal{S}_0, \mathcal{S}_1, \ldots$, that is to say if and only if $w_1 \in \mathcal{S}_i$, $w_2 \in \mathcal{S}_j$ and $i > j$. If we picture the levels with $\mathcal{S}_0$ on the top and the rest of the levels following underneath it, then $w_1 <_E w_2$ is equivalent to w_1 being on a lower level than w_2. The problem of figuring out when $w_1 <_E w_2$ now becomes the problem of seeing when w_1 is on a lower level than w_2.

One suggestion is that w_1 is lower than w_2 if there is a chain of interpretations going from w_1 to w_2 in a $\leq^*_E$-chain. A moment's reflection shows this is not a necessary condition, since w_1 might end up lower than w_2 even though they are not connected directly in $\leq^*_E$. The following example shows this.

Example 18. Take a profile $[E] = \{000\}$ and a Horn merging operator satisfying $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$ that generates (through Definitions 25 and 26 the pre-order $\leq^*_E$) in Figure 6.3.

Notice that in the final pre-order $\leq_E$ 001 and 100 end up on the same level and 110 ends up on the level below them. Thus, $110 <_E 001$, even though in 110 and 001 are not connected in $\leq^*_E$.

What we can do, however, is connect both w_1 and w_2 to some element of the top level $\mathcal{S}_0$: in Figure 6.3, 001 and 110 are both connected to 011 on top level through a $\leq^*_E$-chain. The strategy, then, will be to characterize the elements of the top level $\mathcal{S}_0$, and show that any interpretation is connected to the top level through some continuous $\leq^*_E$-chain. We will then judge the position of an interpretation w in the $\leq_E$ hierarchy by judging the distance from w to the top level.

Lemma 13. *For any $i \geq 1$ and any $w \in \mathcal{S}_i$, there is $w' \in \mathcal{S}_{i-1}$ such that $w <^*_E w'$.*

Proof. Suppose this is not the case, i.e. there is a level $i \geq 1$ and there is a $w \in \mathcal{S}_i$ for which there is no $w' \in \mathcal{S}_{i-1}$ such that $w <^*_E w'$ (see Figure 6.4).

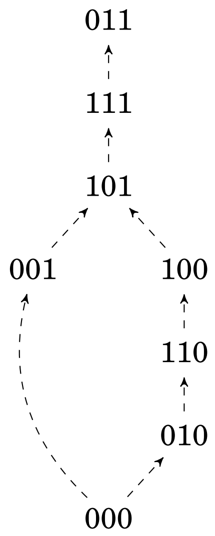

Figure 6.3: 110 is lower than 001, even though they are not connected in $\leq_K^*$

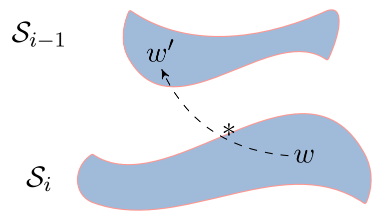

Figure 6.4: There needs to be $w' \in \mathcal{S}_{i-1}$ such that $w <_E^* w'$

It turns out that $w \in \mathcal{S}_{i-1}$, as the following argument shows. We know, by definition, that:

$$\mathcal{S}_{i-1} = max_{\leq_E^*}(\mathcal{W} \setminus \bigcup_{j=0}^{i-2} \mathcal{S}_j).$$

So take $w^\star \in \mathcal{M} \setminus \bigcup_{j=0}^{i-2} \mathcal{S}_j$ such that $w \leq_E^* w^\star$. Because of our assumption, $w^\star$ cannot be in $\mathcal{S}_{i-1}$, so $w^\star$ must be in one of the lower levels, i.e. $w^\star \in \mathcal{S}_k$, for some $k \geq i$.

Now, it cannot be the case that $k > i$. This would mean that $w^\star$ is on a strictly lower level than w, which (by Lemma 10) implies that $w^\star <_E^* w$.

The only possibility, therefore, is that $k = i$ and $w^\star \in \mathcal{S}_i$. Then, by the definition of $\mathcal{S}_i$, we get that $w^\star \leq_E^* w$. But, because $w^\star$ was chosen as an arbitrary element of $\mathcal{M} \setminus \bigcup_{j=0}^{i-2} \mathcal{S}_j$, this effectively shows that w is a maximal element in $\mathcal{M} \setminus \bigcup_{j=0}^{i-2} \mathcal{S}_j$. Therefore $w \in \mathcal{S}_{i-1}$, which is a contradiction. $\square$

Corollary 1. For any $w_i \in \mathcal{S}_i, i \geq 1$, there is a sequence of interpretations $w_{i-1}, \ldots, w_0$, such that $w_{i-1} \in \mathcal{S}_{i-1}, \ldots, w_0 \in \mathcal{S}_0$, and $w_i <_E^* w_{i-1} <_E^* \cdots <_E^* w_0$ (see Figure 6.5).

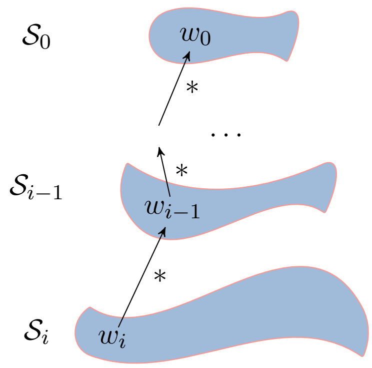

Figure 6.5: Lemma 13 gives us a chain $w_i <_E^* w_{i-1} <_E^* \cdots <_E^* w_0$

Let us now characterize the elements of the top level S_0.

Definition 29. An interpretation w_0 is *a top-level interpretation in* $\leq_E$ if $w_0 \in \mathcal{S}_0$.

Lemma 14. *For $w \in \mathcal{W}$, we have that w is a top-level interpretation in $\leq_E$ iff the following holds: for any $u_0, u_1, \ldots, u_n$ such that $w = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n$, we have that $u_n \leq'_E u_{n-1} \leq'_E \cdots \leq'_E u_1 \leq'_E w = u_0$.*

Proof. For one direction, take $w \in \mathcal{S}_0$, $u_0, u_1, \ldots, u_n$ such that $w = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n$, and suppose one of the reverse inequalities does not hold.

This means that there are two interpretations u_i, u_{i+1} somewhere in the chain such that $u_i <'_E u_{i+1}$. It follows that, in the transitive closure of $\leq'_E$, we get:

$$w = u_0 \leq_E^* u_1 \leq_E^* \cdots \leq_E^* u_i <_E^* u_{i+1} \leq_E^* \cdots \leq_E^* u_n,$$

and in particular $w <_E^* u_{i+1}$. But this is a contradiction, since w is assumed to be on the top level.

For the other direction, take w for which the condition holds, and suppose it is not on the top level $\mathcal{S}_0$. This means that $w \in \mathcal{S}_i$, $i \geq 1$. But now, by applying Lemma 13, we get that there is a $w' \in \mathcal{S}_{i-1}$ such that $w <_E^* w'$. Let us unpack this.

The fact that $w <_E^* w'$ implies that there are $u_0, u_1, \ldots, u_{n-1}, u_n$ such that:

$$w = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n = w'.$$

This follows from the definition of $\leq_E^*$. Now we apply our assumption, and get that:

$$w = u_0 \approx'_E u_1 \approx'_E \cdots \approx'_E u_{n-1} \approx'_E u_n = w'.$$

When we take again the transitive closure of $\leq'_E$, we get that $w \approx_E^* w'$. This contradicts our previous conclusion that $w <_E^* w'$. □

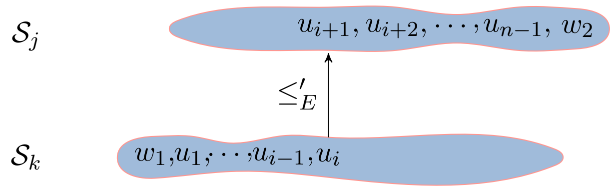

Figure 6.6: A jump from w_1 on $\mathcal{S}_k$ to w_2 on $\mathcal{S}_j$, $j < k$

Top-level interpretations provide a good sky-hook on which to anchor other interpretations. One thing that needs to be established, though, is how this anchoring is to be done. To this end, the notion of a jump will prove useful.

Definition 30. For two interpretations w_1, w_2, we say there is *a jump from* w_1 *to* w_2 *in* $\leq_E$ if there are $u_0, u_1, \ldots, u_n$ such that $w_1 = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n = w_2$, and exactly one of those inequalities is strict.

Jumps are meant to capture the move from one level to a higher one, *via* the $\leq'_E$ relation (see Figure 6.6).

Lemma 15. *If there is a jump from an interpretation* w_1 *to an interpretation* w_2 *in* $\leq_E$*, then* w_2 *is on a higher level than* w_1*. In other words, if* $w_1 \in \mathcal{S}_k$*, then* $w_2 \in \mathcal{S}_j, j < k$*.*

Proof. Take $w_1 \in \mathcal{S}_k$, such that there is a jump from w_1 to w_2. Then there are $u_0, u_1, \ldots, u_n$ such that:

$$w_1 = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n = w_2,$$

with exactly one of the inequalities being strict. Take the strict inequality to be between u_i and u_{i+1}, that is to say $u_i <'_E u_{i+1}$. Now, when we consider the transitive closure of $\leq'_E$, we get that $w_1 <^*_E w_2$, so (by Lemma 10), w_2 is on a higher level than w_1. □

Lemma 16. *If* $w_1 <^*_E w_2$ *and* w_2 *is on a level immediately before the level of* w_1 *in* $\leq_E$ *(i.e.,* $w_1 \in \mathcal{S}_i$ *and* $w_2 \in \mathcal{S}_{i-1}$*, for some* $i \geq 1$*), then there is a jump from* w_1 *to* w_2*.*

Proof. The fact that $w_1 <^*_E w_2$ means that there are $u_0, \ldots, u_n$ such that:

$$w_1 = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n = w_2.$$

We will show that exactly one of those inequalities is strict. Suppose, on the contrary, that this is not the case. In other words, suppose that either none of the inequalities is strict, or that there is more than one strict inequality.

If none of the inequalities is strict, we have:

$$w_1 = u_0 \approx'_E u_1 \approx'_E \cdots \approx'_E u_{n-1} \approx'_E u_n = w_2,$$

which in turn leads to $w_1 \approx^*_E w_2$, and w_1, w_2 end up on the same level.

Suppose that more than one inequality is strict, i.e. $u_i <'_E u_{i+1}$ and $u_j <'_E u_{j+1}$, for some $i < j$, with possibly even more strict inequalities. Then we get:

$$w_1 \leq^*_E \cdots \leq^*_E u_i <^*_E u_{i+1} \leq^*_E \cdots \leq^*_E u_j < u_{j+1} \leq^*_E \cdots \leq^*_E w_2.$$

In particular, we have that:

$$w_1 <^*_E u_{i+1} <^*_E u_{j+1} \leq^*_E w_2.$$

This puts w_2 at least two levels higher up than w_1.

In both cases, we get a contradiction with the fact that w_2 is on the level immediately before the level of w_1. □

The converse to Lemma 16 does not hold. That is, if there is a jump from w_1 to w_2, it does not follow that w_2 is on a level immediately above the level of w_1—there can be jumps that, so to say, skip levels [perhaps an example].

Fortunately, we know from Lemma 13 that for any interpretation w, we can find an interpretation w' on the immediately preceding level, such that $w <^*_E w'$, so there is a $\leq^*_E$-chain that visits every level and ends at the top (Corollary 1). Let us put these results together.

Corollary 2. For any interpretation $w_i \in \mathcal{S}_i$, there is a sequence of interpretations $w_{i-1}, \ldots, w_0$, such that $w_{i-1} \in \mathcal{S}_{i-1}, \ldots, w_0 \in \mathcal{S}_0$, and there is a jump from w_k to w_{k-1} in $\leq_E$, for any $k \in \{0, \ldots, i\}$.

Proof. By Corollary 1, there is a sequence of interpretations $w_{i-1}, \ldots, w_0$, such that $w_{i-1} \in \mathcal{S}_i, \ldots, w_0 \in \mathcal{S}_0$, and:

$$w_i <^*_E w_{i-1} <^*_E \cdots <^*_E w_0.$$

From Lemma 16 it follows that there is a jump from w_k to w_{k-1}, for any $k \in \{0, \ldots, i\}$. □

Corollary 2 shows that we can get from any interpretation w to the top level through a series of successive jumps that visit every level. This result will prove to be crucial in what is to come.

6.3.2 Doing things with formulas

Let us recall the notions we will use in this section. It is no mystery that single interpretations can be characterized by propositional formulas, even in the Horn fragment. For instance, if $\mathcal{U} = \{p_1, p_2, p_3\}$, then the interpretation 111 is characterized by the formula $p_1 \wedge p_2 \wedge p_3$, 110 is characterized by $p_1 \wedge p_2 \wedge (p_3 \rightarrow \bot)$, and so on. A general procedure is easily discernible.

In Section 2.1 we have introduced the notion of complete formulas: propositional formulas whose set of models contains only one element. In Section 2.2 we have named $\mathcal{L}^c_H$ the set of complete Horn propositional formulas. The formula $p_1 \wedge p_2 \wedge p_3$

from above is a complete formula. So is $p_1 \wedge p_1 \wedge p_2 \wedge p_3$. The formula $p_1 \wedge p_2$, however, is not complete, since $[p_1 \wedge p_2] = \{110, 111\}$. In general, we will denote by σ_w a formula such that $[\sigma_w] = \{w\}$. When the context is clear, we write σ_i for σ_{w_i}.

Recall also from Section 2.2 that if μ_i, μ_j are Horn formulas, $\mu_{i,j}$ is the Horn formula such that $[\mu_{i,j}] = Cl_\cap([\mu_i] \cup [\mu_j])$. We cannot express classical disjunction in the Horn fragment, but $\mu_{i,j}$ is as close an approximation of it as we get. If σ_i and σ_j are complete formulas with $[\sigma_i] = \{w_i\}, [\sigma_j] = \{w_j\}$, then $\sigma_{i,j}$ is basically φ_{w_i,w_j}, also introduced in Section 2.2.

The purpose now is to define some syntactic condition (something that could figure in a postulate) on complete formulas, which corresponds to their models being on different levels in the pre-order $\leq_E$ defined from a merging operator Δ that satisfies $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$. In the following, we will always assume that such a Horn merging operator Δ is given, and that $\leq_E$ is the total pre-order built from it as in Section 6.2.1.

Definition 31. A complete Horn formula $\sigma_0 \in \mathcal{L}^c_H$ is a *top-level formula with respect to* E if for any $\sigma_1, \ldots, \sigma_n \in \mathcal{L}^c_H$, it is the case that:

$$\text{if} \quad \begin{cases} \sigma_0 \wedge \Delta_{\sigma_{0,1}}(E), \\ \sigma_1 \wedge \Delta_{\sigma_{1,2}}(E), \\ \ldots \\ \sigma_{n-1} \wedge \Delta_{\sigma_{n-1,n}}(E) \end{cases} \quad \text{are all consistent,}$$

$$\text{then} \quad \begin{cases} \sigma_1 \wedge \Delta_{\sigma_{0,1}}(E), \\ \sigma_2 \wedge \Delta_{\sigma_{1,2}}(E) \\ \ldots \\ \sigma_n \wedge \Delta_{\sigma_{n-1,n}}(E), \end{cases} \quad \text{are also consistent.}$$

We will call $\mathcal{L}^E_{Top} \subseteq \mathcal{L}^c_H$ the set of top-level formulas with respect to E.

Top-level formulas are meant to characterize interpretations on the top level $\mathcal{S}_0$ of $\leq_E$.

Lemma 17. *A complete formula σ_0 is a top-level Horn formula with respect to E iff w_0 is a top-level interpretation in $\leq_E$, where* $[\sigma_0] = \{w_0\}$.

Proof. We will prove the lemma one direction at a time.

(i) If σ_0 is a top-level Horn formula with respect to E, then $w_0 \in [\sigma_0]$ is a top-level interpretation in $\leq_E$.

Take a top-level formula σ_0 with $[\sigma_0] = \{w_0\}$, and interpretations $u_0, \ldots, u_n$ such that:

$$w_0 = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n.$$

We will show that all of these inequalities also hold in reverse, which (by Lemma 14) implies that w_0 is a top level interpretation.

Let us denote by σ_i a complete formula such that $[\sigma_i] = u_i$. Since $u_0 \leq'_E u_1$, we get that $u_0 \in [\Delta_{\sigma_{0,1}}(E)]$, which allows us to conclude that $u_0 \in [\sigma_0 \wedge \Delta_{\sigma_{0,1}}(E)]$. Replicating this reasoning, we get that:

$$\begin{cases} u_0 \in [\sigma_0 \wedge \Delta_{\sigma_{0,1}}(E)], \\ u_1 \in [\sigma_1 \wedge \Delta_{\sigma_{1,2}}(E)], \\ \dots \\ u_{n-1} \in [\sigma_{n-1} \wedge \Delta_{\sigma_{n-1,n}}(E)] \end{cases}$$

Our assumption here is that σ_0 is a top-level formula, so by applying the definition we get that:

$$\begin{cases} \sigma_1 \wedge \Delta_{\sigma_{0,1}}(E), \\ \sigma_2 \wedge \Delta_{\sigma_{1,2}}(E), \\ \dots \\ \sigma_n \wedge \Delta_{\sigma_{n-1,n}}(E), \end{cases}$$

are all consistent. Let us show this implies that $u_i \leq'_E u_{i-1}$, for any $i \in \{1, \dots, n\}$.

Let us look at these consistency statements one at a time. If $\sigma_1 \wedge \Delta_{\sigma_{0,1}}(E)$ is consistent, then $[\sigma_1 \wedge \Delta_{\sigma_{0,1}}(E)] \neq \emptyset$. But we know that $[\sigma_1] = \{u_1\}$, so $u_1 \in [\Delta_{\sigma_{0,1}}(E)] = min_{\leq_E}[\sigma_{0,1}]$. We also know that $u_0 \in [\sigma_{0,1}]$, so we can conclude that $u_1 \leq'_E u_0$. Iterating this reasoning for each of the consistency statements, we get that:

$$u_n \leq'_E u_{n-1} \leq'_E \cdots \leq'_E u_1 \leq'_E u_0 = w_0.$$

By Lemma 14, this implies that w_0 is a top interpretation.

(ii) If w_0 is a top-level interpretation in $\leq_E$, then σ_0 with $[\sigma_0] = \{w_0\}$ is a top-level Horn formula with respect to E.

Take a top-level interpretation w_0, and $\sigma_0 \in \mathcal{L}_H^c$ a complete formula such that $[\sigma_0] = \{w_0\}$. Also, take $\sigma_1, \dots, \sigma_n \in \mathcal{L}_H^c$ such that:

$$\begin{cases} \sigma_0 \wedge \Delta_{\sigma_{0,1}}(E), \\ \sigma_1 \wedge \Delta_{\sigma_{1,2}}(E), \\ \dots \\ \sigma_{n-1} \wedge \Delta_{\sigma_{n-1,n}}(E) \end{cases}$$

are consistent. We want to show that the condition for σ_0 being a top-level Horn formula holds.

All the previous consistency statements allow us to take:

$$\begin{cases} u_0 \in [\sigma_0 \wedge \Delta_{\sigma_{0,1}}(E)], \\ u_1 \in [\sigma_1 \wedge \Delta_{\sigma_{1,2}}(E)], \\ \dots \\ u_{n-1} \in [\sigma_{n-1} \wedge \Delta_{\sigma_{n-1,n}}(E)] \end{cases}$$

Based on all this, what can we conclude about $u_0, \dots, u_n$? First, since $u_0 \in [\sigma_0]$ and $[\sigma_0] = \{w_0\}$, it follows that $w_0 = u_0$. Second, since $u_0 \in [\Delta_{\sigma_{0,1}}(E)]$ and $u_1 \in [\sigma_{0,1}]$, we get that $u_0 \leq'_E u_1$. All in all:

$$w_0 = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n.$$

Since our assumption is that w_0 is a top-level interpretation, we can apply Lemma 14 and get:

$$u_n \leq'_E u_{n-1} \leq'_E \cdots \leq'_E u_1 \leq'_E u_0 = w_0.$$

From $u_n \leq'_E u_{n-1}$ it follows that $u_n \in [\Delta_{\sigma_{n-1,n}}(E)]$, so $u_n \in [\sigma_n \wedge \Delta_{\sigma_{n-1,n}}(E)]$, which shows that $\sigma_n \wedge \Delta_{\sigma_{n-1,n}}(E)$ is consistent. In general, this gives us that $\sigma_i \wedge \Delta_{\sigma_{i-1,i}}(E)$ is consistent, for $i \in \{1, \dots, n\}$, which is significant because it means that σ_0 is a top-level Horn formula with respect to E. $\square$

Let us now characterize jumps from one formula to another.

Definition 32. If $\sigma, \sigma' \in \mathcal{L}_H^c$ are complete Horn formulas, we say there is *a jump from σ to σ' with respect to E* if there are $\sigma_0, \dots, \sigma_n$ such that $\sigma_0 = \sigma, \sigma_n = \sigma'$, and the following statements hold:

$$\begin{cases} \sigma_0 \wedge \Delta_{\sigma_{0,1}}(E), \\ \sigma_1 \wedge \Delta_{\sigma_{1,2}}(E), \\ \dots \\ \sigma_{n-1} \wedge \Delta_{\sigma_{n-1,n}}(E), \end{cases}$$

are all consistent, and exactly one of the following is inconsistent:

$$\begin{cases} \sigma_1 \wedge \Delta_{\sigma_{0,1}}(E), \\ \sigma_2 \wedge \Delta_{\sigma_{1,2}}(E), \\ \dots \\ \sigma_n \wedge \Delta_{\sigma_{n-1,n}}(E). \end{cases}$$

We can show now that jumps between complete Horn formulas correspond to jumps between their interpretations.

Lemma 18. *If $\sigma, \sigma' \in \mathcal{L}_H^c$ are complete Horn formulas and $[\sigma] = \{w_1\}, [\sigma'] = \{w_2\}$, then there is a jump from σ to σ' with respect to E iff there is a jump from w_1 to w_2 in $\leq_E$.*

Proof. We will prove the lemma on direction at a time.

(i) If there is a jump from σ to σ' with respect to E, with $[\sigma] = \{w_1\}$ and $[\sigma'] = \{w_2\}$, then there is a jump from w_1 to w_2 in $\leq_E$.

Take two complete Horn formulas σ, σ', with $[\sigma] = \{w_1\}, [\sigma'] = \{w_2\}$, and assume there is a jump from σ to σ'. This means that there are $\sigma_0, \ldots, \sigma_n$ such that $\sigma_0 = \sigma, \sigma_n = \sigma'$, and:

$$\begin{cases} \sigma_0 \wedge \Delta_{\sigma_{0,1}}(E), \\ \sigma_1 \wedge \Delta_{\sigma_{1,2}}(E), \\ \ldots \\ \sigma_{n-1} \wedge \Delta_{\sigma_{n-1,n}}(E), \end{cases}$$

are all consistent.

Take, then:

$$\begin{cases} u_0 \in [\sigma_0 \wedge \Delta_{\sigma_{0,1}}(E)], \\ u_1 \in [\sigma_1 \wedge \Delta_{\sigma_{1,2}}(E)], \\ \ldots \\ u_{n-1} \in [\sigma_{n-1} \wedge \Delta_{\sigma_{n-1,n}}(E)]. \end{cases}$$

It follows immediately that $u_0 = w_1, u_n = w_2$. More importantly, since $u_{i-1} \in [\Delta_{\sigma_{i-1,i}}(E)]$ and $u_i \in [\sigma_{i-1,i}]$ for $i \in 1, \ldots, n$, we get that $u_{i-1} \leq'_E u_i$, or:

$$u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n.$$

We need to show that exactly one of these inequalities is strict. Suppose, on the contrary, that this is not the case. This means either none of the inequalities is strict, or more than one is strict.

If none of the inequalities is strict, we have that:

$$u_0 \approx'_E u_1 \approx'_E \cdots \approx'_E u_{n-1} \approx'_E u_n.$$

But then neither of the formulas $\sigma_i \wedge \Delta_{\sigma_{i-1,i}}(E)$ is inconsistent, since $u_i \in [\sigma_i \wedge \Delta_{\sigma_{i-1,i}}(E)]$. This contradicts what we know about these formulas from the definition of a jump from σ to σ'.

If more than one inequality is strict, we also derive a contradiction. Suppose $u_{i-1} <'_E u_i$ and $u_{j-1} <'_E u_j$, for some $1 \leq i < j \leq n$.

Now if $u_{i-1} <'_E u_i$, then $u_i \notin [\Delta_{\sigma_{i-1,i}}(E)]$. Since $[\sigma_i] = \{u_i\}$, it follows that $[\sigma_i \wedge \Delta_{\sigma_{i-1,i}}(E)] = \emptyset$, or $\sigma_i \wedge \Delta_{\sigma_{i-1,i}}(E)$ is inconsistent. The same holds for $\sigma_j \wedge \Delta_{\sigma_{j-1,j}}(E)$. Together, these contradict what we know about the jump from σ to σ' from its definition.

To conclude, exactly one of the inequalities in the $\leq'_E$-chain is strict, hence there is a jump from w_1 to w_2.

(*ii*) If there is a jump from w_1 to w_2 in $\leq_E$, then there is a jump from σ to σ' with respect to E, with $[\sigma] = \{w_1\}$ and $[\sigma'] = \{w_2\}$.

Suppose there is a jump from w_1 to w_2. This means there are $u_0, \ldots, u_n$ such that:

$$w_1 = u_0 \leq'_E u_1 \leq'_E \cdots \leq'_E u_{n-1} \leq'_E u_n = w_2,$$

and exactly one of the inequalities is strict. Take $\sigma_0, \ldots, \sigma_n$ such that $[\sigma_k] = \{u_k\}, k \in \{0, \ldots, n\}$, and suppose that $u_{i-1} <'_E u_i$, is the one strict inequality, for some $i \in \{1, \ldots, n\}$.

Quite clearly, if we consider the statements in the following group:

$$\begin{cases} \sigma_0 \wedge \Delta_{\sigma_{0,1}}(E), \\ \sigma_1 \wedge \Delta_{\sigma_{1,2}}(E), \\ \ldots \\ \sigma_{n-1} \wedge \Delta_{\sigma_{n-1,n}}(E), \end{cases}$$

all of them are consistent, since $u_i \in [\sigma_i \wedge \Delta_{\sigma_{i,i+1}}(E)]$, for $i \in \{1, \ldots, n\}$. Half of the definition for a jump from σ to σ', then, is satisfied.

For the other half, let us show that $\sigma_i \wedge \Delta_{\sigma_{i-1,i}}(E)$ is the only inconsistent formula among the formulas of its kind. First, because $u_{i-1} <'_E u_i$, it follows that $u_i \notin [\Delta_{\sigma_{i-1,i}}(E)]$, so $\sigma_i \wedge \Delta_{\sigma_{i-1,i}}(E)$ is inconsistent. Second, suppose $\sigma_j \wedge \Delta_{\sigma_{j-1,j}}(E)$ is also inconsistent, for some $j \neq i$. Then we would get that $u_j \notin [\Delta_{\sigma_{j-1,j}}(E)]$, and hence $u_{j-1} <'_E u_j$. But this contradicts the fact that $u_{i-1} <'_E u_i$ is the only strict inequality in the $\leq'_E$-chain. This concludes the proof. □

The notion of a jump will be used towards defining the level of a complete formula, a notion analogous to the level of an interpretation in the order $\leq_E$.

Definition 33. If $\sigma \in \mathcal{L}^c_H$ is a complete Horn formula, then *an ascending ladder for σ with respect to E* is a sequence $\sigma_0, \sigma_1, \ldots, \sigma_n$ of formulas such that $\sigma_0 = \sigma, \sigma_n \in \mathcal{L}^E_{Top}$, and there is a jump from σ_i to σ_{i+1}, for any $i \in \{0, n-1\}$.

Proposition 21. *For any $\sigma \in \mathcal{L}^c_H$, there exists an ascending ladder for σ with respect to E.*

Proof. Suppose $[\sigma] = \{w\}$. By Corollary 2, there are $u_0, \ldots, u_n$ such that $w = u_0, u_n \in \mathcal{S}_0$, u_i is on a level immediately above the level of u_{i-1}, and there is a jump from u_{i-1} to u_i, for every $i \in \{0, \ldots, n\}$.

Take $\sigma_0 = \sigma$, and σ_i to be complete formulas such that $[\sigma_i] = \{u_i\}$, for $i \geq 1$. Since u_n is a top-level interpretation, then by Lemma 17, σ_n is a top-level formula. By Lemma 18, to every jump from u_{i-1} to u_i there corresponds a jump from σ_{i-1} to σ_i, for every $i \in \{0, \ldots, n\}$. This shows that there is an ascending ladder for σ. □

The purpose is to use ascending ladders as a means of characterizing a formula's 'distance from the top': the number of steps it takes to reach the top level starting from the formula. Proposition 21 guarantees that there exists an ascending ladder for any $\sigma \in \mathcal{L}^c_H$, in other words that there exist paths to the top from every complete formula.

In general there can exist many ascending ladders for any σ, depending on how the jumps are done. We would like an ascending ladder that captures where the model of σ is in terms of the hierarchy of levels $\mathcal{S}_0, \mathcal{S}_1, \ldots$ in $\leq_E$.

Definition 34. If $\sigma_0, \dots, \sigma_n$ is an ascending ladder for σ with respect to E, then we take n to be its *length*.

Definition 35. For $\sigma \in \mathcal{L}_H^c$, the *level of σ with respect to E* is the length of its longest ascending ladder with respect to E.

We denote by $l_E(\sigma)$ the level of σ with respect to E.

Proposition 22. *If $[\sigma] = \{w\}$, then the level of σ with respect to E is the same as the level of w in $\leq_E$.*

Proof. Suppose this is not the case, i.e. take m to be the level of σ with respect to E and n the level of w in $\leq_E$, with $m \neq n$.

We have that $w \in \mathcal{S}_n$, so (by Corollary 2) there is a sequence of interpretations $w_n, w_{n-1}, \dots, w_0$, such that $w_n = w \in \mathcal{S}_n, w_{n-1} \in \mathcal{S}_{n-1}, \dots, w_0 \in \mathcal{S}_0$, and there is a jump from w_k to w_{k-1}, for any $k \in \{0, \dots, n\}$. Let us take σ_i to be complete formulas such that $\sigma_0 = \sigma$ and $[\sigma_i] = \{w_i\}$, for $i \geq 0$. By Lemma 17, σ_0 is a top-level formula; by Lemma 18, there is a jump between σ_{k-1} and σ_k, for any $k \in \{1, \dots, n\}$. Thus, $\sigma_0, \dots, \sigma_n$ is an ascending ladder for σ with respect to E, and its length is n.

Now, we cannot have $m < n$, since the level of σ is supposed to be the length of the longest ascending ladder of σ with respect to E, and we have just found an ascending ladder of length n. It follows that $m > n$.

But then (by the definition of ascending ladders) there is a sequence $\sigma_0, \sigma_1, \dots, \sigma_m$ of complete Horn formulas (no connection with the σ's above) such that $\sigma_0 = \sigma, \sigma_m \in \mathcal{L}_{Top}^E$, and there is a jump from σ_{i-1} to σ_i, for any $i \in \{1, m\}$.

If we take $[\sigma_i] = \{w_i\}$ (again, no connection to the interpretations considered above), we get (by Lemma 18) that there is a jump from w_{i-1} to w_i, for any $i \in \{1, m\}$. Since a jump translates into a difference of at least one level (from Lemma 15) and we have m jumps, we get that the level of w_0 is grater than m and (as we have assumed that $m > n$), definitely greater than n. But $w_0 = w$, and the level of w in $\leq_E$ is n. We have arrived at a contradiction.

Thus, the only possibility is that $m = n$. □

This now practically solves our problem, since we have found a characterization of where a given interpretation stands in a particular order $\leq_E$ in terms of formulas of $\mathcal{L}_H$ and the Δ operator.

Corollary 3. For any interpretations w_1, w_2 and $\sigma_1, \sigma_2 \in \mathcal{L}_H^c$ such that $[\sigma_1] = \{w_1\}, [\sigma_2] = \{w_2\}$, it is the case that $w_1 <_E w_2$ iff $l_E(\sigma_1) > l_E(\sigma_2)$.

Proof. By Lemma 10, $w_1 <_E w_2$ iff the level of w_1 is greater than the level of w_2 in the hierarchy of levels $\mathcal{S}_0, \mathcal{S}_1, \dots$ in $\leq_E$. By Proposition 22, the level of w_1 is greater than the level of w_2 in $\leq_E$ iff $l_E(\sigma_1) > l_E(\sigma_2)$. □

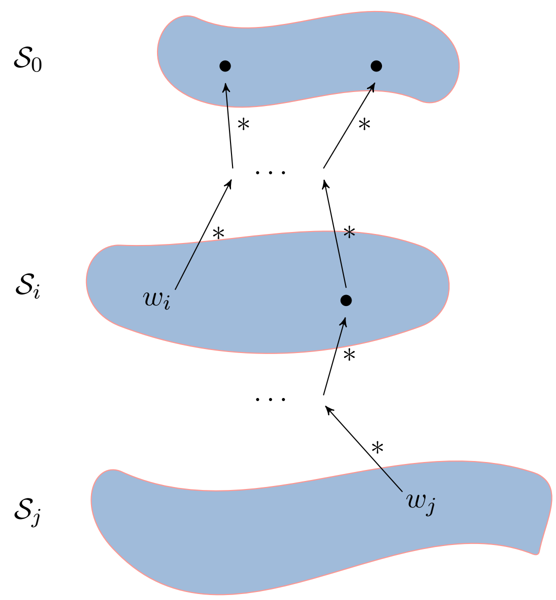

Figure 6.7: Distance from the top for interpretations on different levels

6.3.3 Conclusion

To spell out the intuition behind all this: we start from the fact that the relative positioning of w_1 and w_2 in $\leq_E$ cannot be represented with $\Delta_{\varphi_{w_1,w_2}}(E)$ only. We show this does not imply that the order between w_1 and w_2 is invisible to a Horn merging operator—the order between w_1 and w_2 can be captured in $\mathcal{L}_H$ by looking at the interaction of φ_{w_1,w_2} with the other formulas in $\mathcal{L}_H$.

More, precisely, the order between w_1 and w_2 in $\leq_E$ depends on the distance between w_1 and w_2 and the top level $\mathcal{S}_0$: the one that is 'farther' from the top is lower in the hierarchy (see Figure 6.7).

This translates into a relation on formulas: it turns out that w_1 is farther from the top than w_2 in $\leq_E$ exactly when σ_1 is farther from the top than σ_2 with respect to E, given suitably defined notions of 'top' and 'distance from the top' for formulas. To put it more shortly, we have recoded the construction of $\leq_E$ in the actual language of Horn logic, using the operator Δ.

6.4 kp_5, kp_6 (and kp_4) revisited

In these conditions it becomes possible to capture kp_5 and kp_6. We can replace postulates IC_5, IC_6 with:

(IC_5') For any $\sigma_1, \sigma_2 \in \mathcal{L}_H^c$, if $l_{E_1}(\sigma_1) \geq l_{E_1}(\sigma_2)$ and $l_{E_2}(\sigma_1) \geq l_{E_2}(\sigma_2)$, then $l_{E_1 \sqcup E_2}(\sigma_1) \geq$

$l_{E_1 \sqcup E_2}(\sigma_2)$.

(IC'_6) For any $\sigma_1, \sigma_2 \in \mathcal{L}^c_H$, if $l_{E_1}(\sigma_1) \geq l_{E_1}(\sigma_2)$ and $l_{E_2}(\sigma_1) > l_{E_2}(\sigma_2)$, then $l_{E_1 \sqcup E_2}(\sigma_1) > l_{E_1 \sqcup E_2}(\sigma_2)$.

We can even get an equivalent for IC_4 which targets kp_4 specifically:

(IC''_4) For any $\sigma_1 \in \mathcal{L}^c_H$ such that $\sigma_1 \models K_1$, there is a $\sigma_2 \in \mathcal{L}^c_H$ such that $\sigma_2 \models K_2$ and $l_{\{K_1,K_2\}}(\sigma_2) \geq l_{\{K_1,K_2\}}(\sigma_1)$.

There is still the question of the relationship between these postulates and the standard postulates for merging in the normal case: ideally, in the normal case they are logically implied by the IC postulates. We leave this matter to further reflection. In the absence of definitive information about how these postulates relate to the standard ones, we can formulate the theorem that follows as below:

Theorem 2. *Take a Horn merging operator Δ that satisfies $IC_0 - IC_3 + IC_7 - IC_8 + Acyc$ and the Horn compliant assignment defined from it as in Section 6.2.1, which we know satisfies $kp_1 - kp_3$. Then the assignment satisfies $kp_4 - kp_6$ if and only if Δ satisfies IC''_4, IC'_5 and IC'_6.*

Proof. This follows immediately, using Corollary 3. □

We have thus obtained a qualified representation result, which enforces $kp_4 - kp_6$ on an assignment as long as it is built using the construction outlined in Section 6.2.1 and (the rather unwieldy) postulates IC''_4, IC'_5 and IC'_6 hold.

CHAPTER 7

Specific Horn merging operators

In this chapter we introduce a couple of specific Horn merging operators.

7.1 Some more notation

Before proceeding, let us gather some notation that will be used in the rest of this chapter.

We have seen that a pre-order $\leq$ on interpretations can be represented as a graph, by placing elements lower in the order lower on the graph, or through levels assigned to each interpretation. A level is a positive integer, used to fix the position of an interpretation in a particular pre-order. We will be making use of this feature heavily in the rest of this chapter. The convention will usually be that if K is a knowledge base and w is an interpretation, then $l_K(w)$ denotes the level of w in $\leq_K$. The pre-order $\leq_K$ is completely determined if the levels of every interpretation are given, by the rule:

$$w_1 \leq_K w_2 \text{ iff } l_K(w_1) \leq l_K(w_2).$$

We will usually specify a pre-order by giving the levels of an interpretation in it. If $w \in [K]$, then $l_K(w) = 0$, always. This is to ensure that the pre-order is at least faithful. If K is a complete knowledge base and $[K] = \{w\}$, then for any interpretation w' we write $l_w(w')$ instead of $l_K(w')$.

If $\leq_1$ and $\leq_2$ are two total pre-orders specified through their levels, $\leq_{1+2}$ denotes the pre-order got by aggregating $\leq_1$ and $\leq_2$ through the sum aggregation function: that is to say, we add the levels interpretation-wise.

The graph representation is useful to visualize the relative positions of interpretations in a pre-order. In his chapter, though, we will mostly represent a pre-order

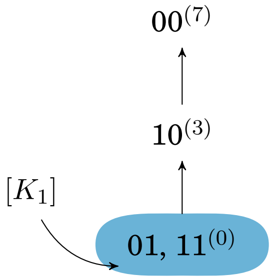

	$[K_1]$	$[K_2]$	$\leq_{K_1} + \leq_{K_2}$
00	7	0	7
01	0	2	2
10	3	2	5
11	0	5	5

Figure 7.1: Different representations of pre-orders

$\leq_K$ through a table, where each column represents the vector of levels corresponding to each interpretation. The head of the column will usually be a list of models of K, or the way we obtained the pre-order.

As an example, Figure 7.1 shows on the left the graph representation for a pre-order $\leq_{K_1}$ with $[K_1] = \{01, 11\}$. On the right, it shows the same pre-order in a table, alongside another pre-order $\leq_{K_2}$ and $\leq_{K_1} + \leq_{K_2}$.

The same things apply to pre-orders $\leq_E$, where E is a profile.

Since our purpose will be to construct complete assignments for Horn profiles, it will be useful to break assignments down into components.

Definition 36. A *basic assignment* is a faithful assignment for consistent knowledge bases.

Observation 12. Since $\leq_{K_1}=\leq_{K_2}$ if $K_1 \equiv K_2$, a basic assignment can be thought of as a finite list of pre-orders, each corresponding to a non-empty subset of $\mathcal{W}$ (i.e., we can pick one pre-order $\leq_K$ as a representative of the entire equivalence class of pre-orders $\leq_{K'}$, where $K' \equiv K$).

For an alphabet with n propositional atoms the set of interpretations $\mathcal{W}$ has 2^n elements. Each non-empty subset of $\mathcal{W}$ represents one possible consistent knowledge base, modulo logical equivalence. Thus, for a basic assignment we need $2^{2^n} - 1$ pre-orders. The number is smaller if we focus only on the sets of interpretations that are representable by a Horn formula.

An even more basic component of an assignment will turn out to be the assignment for complete knowledge bases K, that is to say knowledge bases that have exactly one model.

Definition 37. An *initial assignment* is an assignment for complete knowledge bases.

We will represent an initial assignment as a $2^n \times 2^n$ square matrix, where n is the size of the alphabet. For instance, Table 7.2 shows the initial assignment for the assignment defined by Delgrande and Peppas for a Horn revision operator (see

	00	01	10	11
00	0	1	1	1
01	1	0	2	2
10	1	2	0	2
11	2	3	3	0

Figure 7.2: Basic assignment for the assignment defined by Delgrande and Peppas for Horn revision (Definition 22).

Definition 22). The idea will be to take a suitable basic assignment and convert it into a Horn compliant, syncretic assignment for profiles. For how this is done, once we have a suitable basic assignment, see Section 7.3. The following section clarifies what we mean by 'suitable'.

7.2 Some general properties for Horn compliance

Theorem 1 tells us that a Horn compliant syncretic assignment can be used to define a Horn merging operator that satisfies the standard postulates $IC_0 - IC_8$ and $Acyc$. It is useful, because it gives us a guide to building specific Horn merging operators.

The standard methods for defining syncretic assignments (familiar notions of distance between interpretations and intuitive aggregation functions) fail in the Horn case because they are typically not Horn compliant. The up-side is that if we can make sure that an assignment is Horn compliant, even through multiple aggregations, then we can use it to define a Horn merging operator. So if we decide to use familiar aggregation functions, the problem narrows down to making sure that the pre-orders always remain Horn compliant.

In the following, we will mostly work with the sum aggregation function. For Horn compliance, we prove some general properties that characterize Horn compliant pre-orders.

The first thing to notice is that to check that a pre-order $\leq$ is Horn compliant one need not go through *all* possible subsets $[\mu] \subseteq \mathcal{W}$, where $\mu \in \mathcal{L}_H$, as a naive reading of the definition would suggest. In fact, it is sufficient to check just triples of interpretations $\{w_0, w_1, w_2\}$, where $w_1 \not\subseteq w_2$, $w_2 \not\subseteq w_1$ and $w_0 = w_1 \cap w_2$.

Lemma 19. *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 shows is that Horn compliance means that this arrangement never occurs in $\leq$.

Now, 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 $\{w_0, w_1, w_2\}$.

Conversely, 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. □

We can break down our problem into two parts. The first part consists in defining a set of pre-orders $\leq_K$ for consistent knowledge bases K, and is often simple enough. The next step is to find some condition which preserves Horn compliance through sum aggregation: that is, if $\leq_1$ and $\leq_2$ are Horn compliant, what needs to happen in order for $\leq_{1+2}$ to be Horn compliant? The following property turns out to be sufficient.

Definition 38. A pre-order $\leq$ is *well-behaved* if and only for any interpretations w_0, w_1, w_2 such that $w_1 \not\subseteq w_2$, $w_2 \not\subseteq w_1$ and $w_0 = w_1 \cap w_2$, it is the case that:

$$w_0 \leq w_1 \text{ or } w_0 \leq w_2, \text{ and}$$
$$|min\{l(w_1), l(w_2)\} - l(w_0)| \leq |max\{l(w_1), l(w_2)\} - l(w_0)|$$

That is to say, $\leq$ is well-behaved if w_0 is smaller than at least one of w_1 and w_2 and, in case w_0 is between w_1 and w_2, then w_0 is closer to the smaller of the two. Notice that if w_0 is smaller than both w_1 and w_2, then $\leq$ is automatically well behaved. Also notice that a well-behaved pre-order $\leq$ is also Horn compliant. We can now show that well-behavedness is transmitted through sum aggregation.

Proposition 23. *If $\leq_1$ and $\leq_2$ are well-behaved, then $\leq_{1+2}$ is well-behaved.*

Proof. We will show that $\leq_{1+2}$ is well-behaved by case analysis. Take a triple of interpretations $\{w_0, w_1, w_2\}$ such that $w_1 \not\subseteq w_2$, $w_2 \not\subseteq w_1$ and $w_0 = w_1 \cap w_2$.

(i) Suppose w_0 is smaller than one of the other two interpretations in *both* pre-orders. In other words, there exists $w_i \in \{w_1, w_2\}$ such that $l_1(w_0) \leq l_1(w_i)$ and $l_2(w_0) \leq l_2(w_i)$.

Obviously, then, $l_{1+2}(w_0) \leq l_{1+2}(w_i)$, so the first property of well-behaved pre-orders is satisfied. To see that the other property also holds, let us look at the possible cases. Suppose $l_1(w_0) \leq l_1(w_1)$ and $l_2(w_0) \leq l_2(w_1)$. If w_0 is smaller than both w_1 and w_2 in both pre-orders, then w_0 is smaller than both w_1 and w_2 in the aggregated pre-order, and the second property holds trivially.

Suppose, then, that $l_1(w_2) < l_1(w_0) \le l_1(w_1)$ and $l_2(w_0) \le l_2(w_1)$. Here we also have two cases. First, suppose $l_2(w_2) \le l_2(w_0)$. So take $m > 0$, $n, p, q \ge 0$ such that:

$$l_1(w_1) = l_1(w_0) + n,$$
$$l_1(w_2) = l_1(w_0) - m,$$
$$l_2(w_1) = l_2(w_0) + q,$$
$$l_2(w_2) = l_2(w_0) - p.$$

Because $\le_1$ and $\le_2$ are well-behaved, we have that $m \le n$ and $p \le q$. The second condition for well-behavedness translates as:

$$|min\{n + q, -m - p\}| \le |max\{n + q, -m - p\}|.$$

Since $-m - p < 0$ and $n + q \ge 0$, it has to be the case that $-m - p$ is the minimal element of $\{n + q, -m - p\}$. Then the inequality we have to prove becomes:

$$| - m - p| \le |n + q|,$$

or $m + p \le n + q$. Because $m \le n$ and $p \le q$, this follows immediately.

Second, suppose $l_2(w_0) < l_2(w_2)$. Take $m, p > 0$ and $n, q \ge 0$ such that:

$$l_1(w_1) = l_1(w_0) + n,$$
$$l_1(w_2) = l_1(w_0) - m,$$
$$l_2(w_1) = l_2(w_0) + q,$$
$$l_2(w_2) = l_2(w_0) + p.$$

The condition of well-behavedness translates as:

$$|min\{n + q, p - m\}| \le |max\{n + q, p - m\}|.$$

If both $n + q$ and $p - m$ are positive, this follows immediately. If $p - m < 0$, then it is clearly the minimal element in $\{n + q, p - m\}$, and the inequality to prove becomes:

$$|p - m| \le |n + q|,$$

or $m - p \le n + q$. Since $m \le n$, the conclusion follows immediately.

(ii) Since it cannot be the case that w_0 is greater than both w_1 and w_2 in either $\le_1$ or $\le_2$, the only left possibility is that $w_1 <_1 w_0 \le_1 w_2$ and $w_2 <_2 w_0 \le_2 w_1$, or other some such combination obtained through symmetry. Let us focus on this case.

The fact that $w_1 <_1 w_0 \le_1 w_2$, means that $l_1(w_1) < l_1(w_0) \le l_1(w_2)$. Take, therefore, $m > 0$ and $n \ge 0$ such that:

$$l_1(w_1) = l_1(w_0) - m,$$
$$l_1(w_2) = l_1(w_0) + n.$$

Because $\leq_1$ is, by assumption, well-behaved, it follows that $m \leq n$. Analogously, take $p > 0$ and $q \geq 0$ such that:

$$l_2(w_2) = l_2(w_0) - p,$$
$$l_2(w_1) = l_2(w_0) + q.$$

We also have that $p \leq q$.

Now assume, on the contrary, that $w_1 <_{1+2} w_0$ and $w_2 <_{1+2} w_0$. This means that $l_1(w_1) + l_2(w_1) < l_1(w_0) + l_2(w_0)$ and $l_1(w_2) + l_2(w_2) < l_1(w_0) + l_2(w_0)$. Or, by substituting the above equalities:

$$l_1(w_0) - m + l_2(w_0) + q < l_1(w_0) + l_2(w_0),$$
$$l_1(w_0) + n + l_2(w_0) - p < l_1(w_0) + l_2(w_0).$$

Simplifying, we get:

$$q - m < 0,$$
$$n - p < 0.$$

This means that $q < m$ and $n < p$. Because $m \leq n$ and $p \leq q$, we get that $q < m \leq n < p \leq q$, which leads to a contradiction.

We also need to show:

$$|min\{l_{1+2}(w_1), l_{1+2}(w_2)\} - l_{1+2}(w_0)| \leq |max\{l_{1+2}(w_1), l_{1+2}(w_2)\} - l_{1+2}(w_0)|.$$

When we unpack this, it is equivalent to:

$$|min\{l_1(w_1) + l_2(w_1), l_1(w_2) + l_2(w_2)\} - (l_1(w_0) + l_2(w_0))| \leq$$
$$|max\{l_1(w_1) + l_2(w_1), l_1(w_2) + l_2(w_2)\} - (l_1(w_0) + l_2(w_0))|.$$

Or, to put it differently:

$$|min\{(l_1(w_1) + l_2(w_1)) - (l_1(w_0) + l_2(w_0)), (l_1(w_2) + l_2(w_2)) - (l_1(w_0) + l_2(w_0))\}| \leq$$
$$|max\{(l_1(w_1) + l_2(w_1)) - (l_1(w_0) + l_2(w_0)), (l_1(w_2) + l_2(w_2)) - (l_1(w_0) + l_2(w_0))\}|.$$

Plugging in the earlier equalities, this is equivalent to:

$$|min\{q - m, n - p\}| \leq |max\{q - m, n - p\}|.$$

Suppose it does not hold, or $|min\{q - m, n - p\}| > |max\{q - m, n - p\}|$. A quick case analysis shows that this leads to a contradiction. If $min\{q - m, n - p\} \geq 0$, the contradiction is immediate. Therefore, assume $min\{q - m, n - p\} < 0$, and that $min\{q-m, n-p\} = q-m$. Since $q-m < 0$, we get that $q < m$. Adding this up to what we already know about m, n, p, q from the fact that $\leq_1$ and $\leq_2$ are well-behaved, we get that:

$$p \leq q < m \leq n$$

Since $|min\{q-m, n-p\}| > |max\{q-m, n-p\}|$, it follows that $|q-m| > |n-p|$, or $m-q > n-p$. This is a contradiction.

If $min\{q-m, n-p\} = n-p$, then $n-p < 0$ and $n < p$. Hence $m \leq n < p \leq q$. From $|min\{q-m, n-p\}| > |max\{q-m, n-p\}|$, we derive that $|n-p| > |q-m|$, or $p-n > q-m$. This is, again, a contradiction. □

In Chapter 2 it was mentioned that if we want to aggregate pre-orders using Σ or GMAX, the assignment should be symmetric (otherwise the merging operator does not satisfy IC_4). The condition of symmetry turns out to play an important role: if symmetry holds, then the basic assignment is completely determined by the initial assignment.

Lemma 20. *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 complete knowledge base, the pre-order $\leq_K$ is assumed to be given. Let us suppose, now, that $[K] = \{w_1, \ldots, w_n\}$, for $n > 1$. We want to show that if w is an interpretation, then the level of w in $\leq_K$ is determined by the assignment for complete knowledge bases.

Take, then, an interpretation $w_i \in \mathcal{W}$. We denote by $l_K(w_i)$ the level of w_i in $\leq_K$, and 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 $min\{l_K(w) \mid w \in K_i\} = min\{l_K(w_i)\} = l_K(w_i)$. We want to show that this is determined by the assignment for complete knowledge base. For this, remember that K_i is a complete knowledge base, therefore $l_{K_i}(w)$ is determined, for any interpretation w, therefore $min\{l_{K_i}(w) \mid w \in [K]\}$ is completely determined. □

Observation 13. Notice that symmetry requires the matrix representing the initial assignment to be a symmetric matrix, with 0 on the main diagonal. For instance, the assignment defined by Delgrande and Peppas for Horn revision (Definition 22) is not symmetric, and this can be read off of its initial assignment matrix (see Table 7.2).

Let us look at an example of how symmetry allows us to determine the pre-order $\leq_K$, where K is not complete, from the initial assignment.

Example 19. Consider the initial assignment in Table 7.1. Notice that it is symmetric. Take, now, a knowledge base $[K] = \{10, 11\}$. Let us apply symmetry to compute the vector levels for $\leq_K$. Take interpretation 00. By symmetry, we have

	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 7.1: Applying symmetry to determine pre-orders from the initial assignment.

that $l_K(00) = l_{00}(K)$. Now:

$$l_{00}(K) = min\{l_{00}(10), l_{00}(11)\} = min\{2, 3\} = 2.$$

Similarly, we can find the levels of every other interpretation in $\leq_K$. □

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 24. *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 fall out easily. If either of K_1 or K_2 is complete, then $l_{K_1}(K_2) = l_{K_2}(K_1)$ by definition. If $[K_1] \cap [K_2] \neq \emptyset$, then $l_{K_1}(K_2) = l_{K_2}(K_1) = 0$.

The only case left to analyse is when K_1 and K_2 are consistent, non-complete knowledge bases and they share no models. Suppose, then, that:

$$[K_1] = \{w_1, \ldots, w_m\},$$
$$[K_2] = \{w'_1, \ldots, w'_n\}.$$

Then:

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

The last step there was taken by applying symmetry. To visualize what this statement says, consult Table 7.3 and focus on the square of dots in the upper right corner: $min\{l_{w'_1}(w_1), \ldots, l_{w'_1}(w_m)\}$ takes the minimum of the dotted elements in the w'_1-column, while $min\{l_{w'_n}(w_1), \ldots, l_{w'_n}(w_m)\}$ 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.

	w_1	...	w_m	w'_1	...	w'_n
w_1	0			•	...	•
...				•	...	•
w_m			0	•	...	•
w'_1	•	...	•	0		
...	•	...	•			
w'_n	•	...	•			0

Figure 7.3: Symmetry

Completely analogously, we get that $l_{K_2}(K_1)$ takes the minimum element in the lower left dotted square of Table 7.3. Crucially, remember that the initial assignment matrix is symmetric—hence, the sub-matrix we have selected in Table 7.3 is also symmetric. It follows that the dotted squares contain the same elements, and so they have the same minima. This proves that $l_{K_2}(K_1) = l_{K_1}(K_2)$. $\square$

Symmetry and well-behavedness combined allow us to take a ground assignment and turn into a Horn compliant syncretic assignment.

Theorem 3. *If a basic assignment determined through symmetry from a symmetric initial assignment is also well-behaved, we can define a Horn compliant syncretic assignment from it.*

Proof. The basic assignment ends up being symmetric, by Proposition 24. Construct now an assignment for profiles as follows:

- if $E = \{K\}$ is a profile that contains only one (consistent) knowledge base, then assign to E the pre-order $\leq$ from the initial assignment that corresponds to the models of K;
- if $E = \{K_1, \ldots, K_n\}$ contains more than one consistent knowledge bases, then assign to E the pre-order $\leq_{\Sigma_1^n}$, where $\leq_i$ is the pre-order in the initial assignment corresponding to the models of K_i.

Since the basic assignment is well-behaved (by assumption), it follows by Proposition 23 that aggregating any two pre-orders from the basic assignment produces a well-behaved pre-order. By an induction argument, well-behavedness is preserved through arbitrary aggregations. This means that $\leq_E$ is well-behaved, for any profile E. As mentioned earlier, well-behavedness implies Horn compliance.

We have argued above that the resulting assignment is Horn compliant. The assignment is by design faithful, so it satisfies properties $kp_1 - kp_3$. We have also argued above that the assignment is symmetric, so it satisfies kp_4. Since aggregation is done with the sum function, it also satisfies $kp_5 - kp_6$. This means that the assignment is Horn compliant and syncretic $\square$

	00	01	10	11	00, 01	00, 10	00, 11	01, 11	10, 11	00, 01, 10	00, 01, 11	00, 10, 11
00	0	1	1	2	0	0	0	1	1	0	0	0
01	1	0	2	3	0	1	1	0	2	0	0	1
10	1	2	0	3	1	0	1	2	0	0	1	0
11	2	3	3	0	2	2	0	0	0	2	0	0

Table 7.2: Ground assignment for the 2 letter alphabet

Observation 14. If we have an initial assignment that is symmetric, and the basic assignment that results from it through symmetry is well-behaved, then by Theorem 3 we can define a syncretic assignment from it. By Theorem 1 this gives us a Horn merging operator.

We can now use all the knowledge we gathered to define some Horn compliant syncretic assignments.

7.3 A Horn compliant syncretic assignment for the two letter alphabet

We first define the initial assignment. As a reminder, $|w|$ is the number of bits in w equal to 1. Now, if $[K] = \{w^*\}$ is a complete knowledge base, define the level of an interpretation w with respect to K as:

$$l_K(w) = \begin{cases} 0, \text{ if } w \in [K], \\ |w^*| + |w|, \text{ otherwise.} \end{cases}$$

For any other knowledge base, the levels are computed by symmetry. Aggregation is done through the sum aggregation function.

For the two letter alphabet, space allows us to present the whole basic assignment (see Table 7.2). Notice that the initial assignment is symmetric, and the rest of the basic assignment is symmetric by design. By Proposition 24, the entire basic assignment is therefore symmetric.

The basic assignment is also well-behaved—this can be checked by inspecting Table 7.2. It follows by Theorem 3 that this gives rise to a Horn compliant syncretic assignment.

7.4 The counter-example for kp_4

From Section 5.4.1, we owe the reader a faithful, Horn compliant assignment that does not satisfy kp_4 but still satisfies IC_4. In Section 5.4.1 we only show what $\leq_{K_1}$ and $\leq_{K_2}$, when $[K_1] = \{01\}$, $[K_2] = \{10\}$, have to be. In Table 7.3 we give the full basic assignment. We aggregate by adding up the levels.

	00	01	10	11	00, 01	00, 10	00, 11	01, 11	10, 11	00, 01, 10	00, 01, 11	00, 10, 11
00	0	1	1	1	0	0	0	1	1	0	0	0
01	1	0	3	1	0	1	1	0	2	0	0	1
10	1	4	0	1	1	0	1	1	0	0	1	0
11	1	2	1	0	1	1	0	0	0	1	0	0

Table 7.3: Counter-example for kp_4

Remember that when we sum aggregate $\leq_{K_1}$ and $\leq_{K_2}$ we get a pre-order that does not satisfy kp_4, but its associated operator satisfies IC_4 (see Section 5.4.1).

Notice that the initial assignment is not completely symmetric. However, the basic assignment for the non-basic part is computed using symmetry from the basic part. It can be checked by hand that except for the non-symmetric part we introduced explicitly, the rest of the basic assignment is symmetric (we only need to additionally check symmetry for pre-orders $\leq_{K_i}$ and $\leq_{K_j}$, where K_i and K_j are non-complete and do not share any models).

Direct inspection of the basic assignment in Table 7.3 shows that it is well-behaved. Hence, by Proposition 23, the basic assignment is Horn compliant and stays Horn compliant through arbitrary aggregations.

By Theorem 3, this translates into a Horn compliant, almost-syncretic assignment.

7.5 A Horn compliant syncretic assignment for an n letter alphabet

We define here an assignment for the general case of an alphabet of size n. As before, we define the initial assignment and extend it to the basic assignment through symmetry. 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 the entries above the main diagonal are fixed by symmetry. We aggregate using the sum.

An additional provision is that in the initial assignment matrix we rearrange the rows and columns thus: we put first the interpretations with 0 bits that equal to 1, then the interpretations with exactly 1 bits equal to 1, and so on. As an example, in the 3 letter alphabet we would have the following order: 000, 001, 010, 100, 011, 101, 110, 111. This is slightly different than the order we had been putting interpretations in up to now.

For just this definition, let us denote by $(w_0, \ldots, w_{2^n})$ the vector of interpretations in $\mathcal{W}$, where n is the size of the alphabet. This vector is ordered as we have described in the previous paragraph.

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

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

	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 7.4: The recursive relation for levels

	000	001	010	100	011	101	110	111
000	0	1	2	3	4	5	6	7
001	1	0	3	5	7	9	11	13
010	2	3	0	8	12	16	20	24
100	3	5	8	0	20	28	36	44
011	4	7	12	20	0	48	64	80
101	5	9	16	28	48	0	112	144
110	6	11	20	36	64	112	0	256
111	7	13	24	44	80	144	256	0

Table 7.5: Operator for the three letter alphabet

This means that the levels on the first column are $0, 1, 2, \dots 2^n - 1$.

Second, for $1 \leq i \leq 2^n - 1$, 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. See, for instance, Table 7.4: 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 7.5 shows the matrix that we get for the three letter alphabet. Notice the order of interpretations in the heads of rows and columns.

As mentioned above, the numbers above the main diagonal are fixed by symmetry from the number below the main diagonal.

A couple of things should be noted here, such that the proof that follows goes smoothly.

Observation 15. 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.

	w_a	w_b
$\dots$	$\dots$	$\dots$
w_{i_m}	$a+$	$b+$
$\dots$	$a+$	$b+$
w_{i_n}	$a+$	$b+$
$\dots$	$\dots$	$\dots$

Table 7.6: Extracting a sub-matrix from the initial assignment matrix

Observation 16. 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 a subset two interpretations $w_a, \dots, w_b$, which appear in this order in the matrix, and we extract their columns from the initial assignment matrix (see Figure 7.6).

We get a sub-matrix of the original one as in Figure 7.6. Suppose, now, that we select some subset of interpretations $\{w_{i_m}, \dots, w_{i_n}\}$, which also appear in this order. Take:

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

that is to say: the smallest level on the column for w_a, except 0, is a. 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_b, is b. We represent this by writing $b+$.

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 observation, let us show that the assignment we have defined is well-behaved.

Proposition 25. *The basic assignment just defined is well-behaved.*

Proof. Take a triple of interpretations $\{w_0, w_1, w_2\}$ such that $w_1 \not\subseteq w_2$, $w_2 \not\subseteq w_1$ and $w_0 = w_1 \cap w_2$. We want to show that in any of the pre-orders $\leq_K$ of the basic assignment, it is the case that $w_0 \leq_K w_1$ or $w_0 \leq_K w_2$, and that:

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

First, notice that the assignment for complete knowledge bases 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.

Let us see look now at knowledge bases K that have more than one model. We will do a case distinction.

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

Table 7.7: The level of w_0 in $\leq_K$ has to be smaller than the levels of w_1 and w_2.

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

Table 7.8: The level of w_0 in $\leq_K$ has to be smaller than the level of w_2.

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 $[K] = \{w_{i_1}, \ldots, w_{i_k}\}$. By symmetry, we have:

$$l_K(w_0) = l_{w_0}(K),$$

which means that:

$$l_K(w_0) = min\{l_{w_0}(w_{i_1}), \ldots, l_{w_0}(w_{i_k})\}.$$

Similarly, we get that:

$$l_K(w_1) = min\{l_{w_1}(w_{i_1}), \ldots, l_{w_1}(w_{i_k})\}.$$

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 $[K] = \{w_{i_1}, \ldots, w_{i_k}\}$ (see Table 7.7, 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:

$$min\{l_{w_0}(w_{i_1}), \ldots, l_{w_0}(w_{i_k})\} = a.$$
$$min\{l_{w_1}(w_{i_1}), \ldots, l_{w_1}(w_{i_k})\} = b.$$

Using Observations 15 and 16, it follows that $a \leq b$. Similarly, it follows that $a \leq c$, which implies the conclusion. Second, suppose $w_1 \in [K]$, and that $[K] = \{w_1, w_{i_1}, \ldots, w_{i_k}\}$. 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 7.8, which was also completed using symmetry, and by Observations 15 and 16, 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 non-complete knowledge base. $\square$

Since the assignment is symmetric (by design) and, as we have just shown, it is well-behaved, then Theorem 3 tells us that we can construct a Horn compliant syncretic assignment from it.

CHAPTER 8

Discussion

We conclude by discussing related work and offering a summary of what has been achieved, together with indications of future work.

8.1 Related work

The work closest to our own is by Delgrande and Peppas on Horn revision [16, 17], presented at length in Chapter 4. In fact, our own work began as an attempt to extend Delgrande and Peppas' results to merging. The hope was that a representation result for merging would fall easily from their construction of a faithful preorder from a Horn revision operator. This did not happen, and the Horn merging landscape turned out to be more complex than initially suspected.

Delgrande and Peppas [16, 17] approach revision directly, by exploring connections between Horn revision operators and model-based constructions. This itself adapts previous work of Katsuno and Mendelzon [31]. Other approaches to Horn revision take a different route, by defining revision operators from contraction operators with the Levi identity (see Section 3.2.5):

$$T * \varphi := (T - \neg\varphi) + \varphi \qquad \textbf{(Levi identity)},$$

where T is a theory and φ a formula. This is a tactic for defining revision operators that goes back to at least the original AGM paper [1]. Using the Levi identity in the Horn becomes less straightforward, since contraction ends up having unwanted properties, and there is still no universal agreement over what a contraction operator for Horn logic should look like [6, 15, 19]. In fact, a Horn revision operator defined from a contraction operator has been presented only recently [57]. It is defined from a variant of the Levi identity and a model-based Horn contraction function [56]. The model-based Horn contraction of [56] is an interesting alternative to the approach of Delgrande and Peppas [16, 17] (which we follow here), as it does not rely on the pre-orders to be Horn compliant and total.

Another important approach to Horn revision, which has the added advantage of working for merging as well, is to attempt to 'repair' the results of standard operators [12, 13]. As was mentioned in Sections 4.2.1 and 5.2, standard revision and merging operators are unsuitable for use in the Horn fragment because they may produce results that are not within the fragment—in model terms, the set of models selected as the result of revision/merging might not be closed under intersection. One can attempt to repair these results, either by closing them under intersection, or by selecting subsets that are already closed. Unfortunately, this usually results in operators that typically do not satisfy all the postulates.

We see the strategy of repairing as complementary to our own, as it answers the same basic question: how should revision/merging proceed when the underlying language is restricted to the Horn fragment? We have chosen to focus on 'pure' operators, that (by definition) always stay within the Horn fragment and satisfy all the stated postulates—but by doing so we have had to abandon all the standard operators and come up with ways of constructing new ones (see Chapter 7). But one can just as well start from the standard operators and, by repairing them in suitable ways, obtain approximate operators—operators that satisfy only a subset of the postulates—for the Horn fragment.

The same type of concerns that motivate our inquiry are also present in works on ontology revision and merging. Ontologies are normally expressed in some description logic (DL), which is usually some fragment of first order logic tailored so as to be computationally efficient. Studying belief change for DL languages requires, therefore, adapting the operators to the particular fragment of interest. This is similar to the problem we tackled in this work.

Using belief change in DL languages is appropriate, since maintenance tasks that have to be performed on ontologies are powerfully reminiscent of the belief change operations that have been touched upon here. Medical ontologies often contain axioms that on later reflection turn out to be counter-intuitive and have to be removed: this can be seen as an instance of contraction. A related problem arises when new axioms are added and consistency has to be maintained: this can be seen as an instance of revision. There is a growing literature modelling these tasks as problems of belief change [24, 44, 45, 47, 54, 55]. We also mention here research on updating knowledge bases not directly inspired by belief change, but which still bears connection to it [50, 51].

Model-based approaches to ontology merging are largely inspired by Konieczny and Pino Pérez [35, 37]. We mention here [28], where a family of merging operators called *dilation operators* are presented, and [10], where an existing framework for revision is used to define operations for ontologies.

8.2 Summary

Our purpose was to obtain a characterization of merging operators in the context of Horn propositional logic. In Chapter 5 we showed that this presents some difficulties. Standard merging operators produce results that are not in the Horn fragment, hence they cannot be used as out-of-the-box tools. Furthermore, the limited expressiveness of the Horn fragment means that standard merging postulates fail to characterize the same semantic structure (syncretic assignments) as in the case of full propositional logic. The consequence is that a straightforward generalization of the Konieczny and Pino Pérez representation result [35] to the Horn case is not possible.

Drawing on Delgrande and Peppas' work on Horn revision [16, 17], we introduced in Chapter 6 a restricted notion of syncretic assignments, which required that every ranking on interpretations be Horn compliant (see Definition 21). We then proved, as one half of a representation result, that a Horn merging operator defined on top of this restricted notion satisfies all the standard postulates (Theorem 1).

For the other half of a representation result, we were faced with the problem that there were unwanted structures which, notwithstanding, fit in with the postulates (see Chapter 5). This included pseudo-preorders with non-transitive cycles, as well as pre-orders which did not satisfy the conditions of fairness expected from merging operators. As in Horn revision, pseudo-preorders were eliminated by introducing a special postulate called $Acyc$. Fairness was enforced by introducing a set of specialized postulates. These postulates work on the condition that pre-orders are built with a procedure inspired by Delgrande and Peppas [16, 17], but their relation to the standard postulates is still open. We thus obtained a qualified version of a representation result (see Theorem 2).

The first half of the representation result (Theorem 1) tells us that a syncretic assignment can be used to construct a Horn merging operator, as long as the assignment is Horn compliant. We used this insight in Chapter 7, where we explored the conditions under which a set of pre-orders is Horn compliant and continues to be Horn compliant under successive sum aggregations. We defined the notion of a *well-behaved* pre-order, which guarantees this, and we coupled it with other insights (e.g., symmetry for knowledge bases) to construct a series of Horn compliant syncretic assignments.

8.3 Conclusions and future work

There is plenty of space for improvement. In the short term, we would like to obtain a better characterization of merging operators for the Horn fragment. This means solving the open issues left with Theorem 2, or finding better ways to enforce the fairness properties on an assignment built from a Horn merging operator. Other issues left to be addressed are how to model the arbitration postulate Arb for the Horn fragment (recall that in its standard formulation it could not be expressed in Horn logic), and how to characterize the postulates Maj and Arb semantically.

The issue of specific merging operators for the Horn fragment is still very much open. In Chapter 7 we presented a few samples, but a systematic treatment and classification should be attempted.

In the long run, the notion of merging should be extended to other types of formulas, not just Horn ones. These include dual-Horn, Krom and affine formulas. A unified framework for merging in arbitrary fragments would require a more abstract framework, along the lines of [12].

Bibliography

[1] Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985). On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *J. Symb. Log.*, 50(2):510–530.

[2] Alchourrón, C. E. and Makinson, D. (1981). Hierarchies of Regulations and Their Logic. In Hilpinen, R., editor, *New Studies in Deontic Logic*, pages 125–148. D. Reidel Publishing Company.

[3] Baral, C., Kraus, S., and Minker, J. (1991). Combining Multiple Knowledge Bases. *IEEE Trans. on Knowl. and Data Eng.*, 3(2):208–220.

[4] Baral, C., Kraus, S., Minker, J., and Subrahmanian, V. S. (1992). Combining Knowledge Bases Consisting Of First-Order Theories. *Computational Intelligence*, 8(1):45–71.

[5] Bloch, I., Hunter, A., Appriou, A., Ayoun, A., Benferhat, S., Besnard, P., Cholvy, L., Cooke, R., Cuppens, F., Dubois, D., Fargier, H., Grabisch, M., Kruse, R., Lang, J., Moral, S., Prade, H., Saffiotti, A., Smets, P., and Sossai, C. (2001). Fusion: General concepts and characteristics. *International Journal of Intelligent Systems*, 16(10):1107–1134.

[6] Booth, R., Meyer, T., Varzinczak, I. J., and Wassermann, R. (2014). On the Link between Partial Meet, Kernel, and Infra Contraction and its Application to Horn Logic. *CoRR*, abs/1401.3902.

[7] Brewka, G., Eiter, T., and McIlraith, S. A., editors (2012). *Principles of Knowledge Representation and Reasoning: Proceedings of the Thirteenth International Conference, KR 2012, Rome, Italy, June 10-14, 2012*. AAAI Press.

[8] Cabalar, P. and Son, T. C., editors (2013). *Logic Programming and Nonmonotonic Reasoning, 12th International Conference, LPNMR 2013, Corunna, Spain, September 15-19, 2013. Proceedings*, volume 8148 of *Lecture Notes in Computer Science*. Springer.

[9] Cawsey, A., Galliers, J., Reece, S., and Sparck Jones, K. (1992). Automating the Librarian: Belief Revision As a Base for System Action and Communication with the User. *Comput. J.*, 35(3):221–232.

[10] Cóbe, R., Resina, F., and Wassermann, R. (2013). Merging Ontologies via Kernel Contraction. In Bax, M. P., Almeida, M. B., and Wassermann, R., editors, *ONTOBRAS*, volume 1041 of *CEUR Workshop Proceedings*, pages 94–105. CEUR-WS.org.

[11] Creignou, N., Kolaitis, P., and Zanuttini, B. (2008). Structure identification of boolean relations and plain bases for co-clones. *J. Comput. Syst. Sci.*, 74(7):1103 – 1115.

[12] Creignou, N., Papini, O., Pichler, R., and Woltran, S. (2014a). Belief revision within fragments of propositional logic. *J. Comput. Syst. Sci.*, 80(2):427–449.

[13] Creignou, N., Papini, O., Rümmele, S., and Woltran, S. (2014b). Belief merging within fragments of propositional logic. *CoRR*, abs/1404.6445.

[14] Dechter, R. and Pearl, J. (1992). Structure Identification in Relational Data. *Artif. Intell.*, 58(1-3):237–270.

[15] Delgrande, J. P. (2008). Horn Clause Belief Change: Contraction Functions. In Brewka, G. and Lang, J., editors, *KR*, pages 156–165. AAAI Press.

[16] Delgrande, J. P. and Peppas, P. (2011). Revising Horn Theories. In Walsh, T., editor, *IJCAI*, pages 839–844. IJCAI/AAAI.

[17] Delgrande, J. P. and Peppas, P. (forthcoming). Revising Horn Theories.

[18] Delgrande, J. P., Peppas, P., and Woltran, S. (2013). AGM-Style Belief Revision of Logic Programs under Answer Set Semantics. In [8], pages 264–276.

[19] Delgrande, J. P. and Wassermann, R. (2010). Horn Clause Contraction Functions: Belief Set and Belief Base Approaches. In Lin, F., Sattler, U., and Truszczynski, M., editors, *KR*, pages 143–152. AAAI Press.

[20] Dowling, W. F. and Gallier, J. H. (1984). Linear-time algorithms for testing the satisfiability of propositional horn formulae. *J. Log. Prog.*, 1(3):267 – 284.

[21] Eckert, D. and Pigozzi, G. (2005). Belief merging, judgment aggregation and some links with social choice theory. In Delgrande, J. P., Lang, J., Rott, H., and Tallon, J.-M., editors, *Belief Change in Rational Agents*, volume 05321 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany.

[22] Eiter, T. and Gottlob, G. (1992). On the Complexity of Propositional Knowledge Base Revision, Updates, and Counterfactuals. *Artif. Intell.*, 57(2-3):227–270.

[23] Everaere, P., Konieczny, S., and Marquis, P. (2014). On Egalitarian Belief Merging. In Baral, C., De Giacomo, G., and Eiter, T., editors, *KR*. AAAI Press.

[24] Flouris, G., Plexousakis, D., and Antoniou, G. (2006). On Generalizing the AGM Postulates. In Penserini, L., Peppas, P., and Perini, A., editors, *STAIRS*, volume 142 of *Frontiers in Artificial Intelligence and Applications*, pages 132–143. IOS Press.

[25] Gärdenfors, P. (1988). *Knowledge in Flux. Modeling the Dynamics of Epistemic States*. Cambridge: MIT Press.

[26] Gärdenfors, P. (2011). Notes on the History of Ideas Behind AGM. *J. Philos. Logic*, 40:115–120.

[27] Gärdenfors, P. and Makinson, D. (1988). Revisions of Knowledge Systems Using Epistemic Entrenchment. In Vardi, M. Y., editor, *TARK*, pages 83–95. Morgan Kaufmann.

[28] Gorogiannis, N. and Hunter, A. (2008). Merging First-Order Knowledge Using Dilation Operators. In Hartmann, S. and Kern-Isberner, G., editors, *FoIKS*, volume 4932 of *Lecture Notes in Computer Science*, pages 132–150. Springer.

[29] Hodges, W. (1993). Logical Features of Horn Clauses. In Gabbay, D., Hogger, C., and Robinson, J., editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 1. Logical Foundations*, pages 449–503. Clarendon Press, Oxford.

[30] Katsuno, H. and Mendelzon, A. O. (1991). On the Difference between Updating a Knowledge Base and Revising It. In Allen, J. F., Fikes, R., and Sandewall, E., editors, *KR*, pages 387–394. Morgan Kaufmann.

[31] Katsuno, H. and Mendelzon, A. O. (1992). Propositional Knowledge Base Revision and Minimal Change. *Artif. Intell.*, 52(3):263–294.

[32] Kavvadias, D. and Sideri, M. (1999). The Inverse Satisfiability Problem. *SIAM J. Comput.*, 28(1):152–163.

[33] Konieczny, S., Lang, J., and Marquis, P. (2004). DA^2 merging operators. *Artif. Intell.*, 157(1-2):49–79.

[34] Konieczny, S. and Pino Pérez, R. (1998). On the Logic of Merging. In Cohn, A. G., Schubert, L. K., and Shapiro, S. C., editors, *KR*, pages 488–498. Morgan Kaufmann.

[35] Konieczny, S. and Pino Pérez, R. (2002). Merging Information Under Constraints: A Logical Framework. *J. Log. Comput.*, 12(5):773–808.

[36] Konieczny, S. and Pino Pérez, R. (2005). Propositional belief base merging or how to merge beliefs/goals coming from several sources and some links with social choice theory. *European Journal of Operational Research*, 160(3):785–802.

[37] Konieczny, S. and Pino Pérez, R. (2011). Logic Based Merging. *J. Philosophical Logic*, 40(2):239–270.

[38] Lau, R. Y., Bruza, P. D., and Song, D. (2004). Belief Revision for Adaptive Information Retrieval. In *Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '04, pages 130–137, New York, NY, USA. ACM.

[39] Lau, R. Y. K., Bruza, P. D., and Song, D. (2008). Towards a Belief-revision-based Adaptive and Context-sensitive Information Retrieval System. *ACM Trans. Inf. Syst.*, 26(2):8:1–8:38.

[40] Lewis, D. (1973). *Counterfactuals*. Basil Blackwell Ltd.

[41] Logan, B., Reece, S., and Jones, K. S. (1994). Modelling Information Retrieval Agents with Belief Revision. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '94, pages 91–100, New York, NY, USA. Springer-Verlag New York, Inc.

[42] Makinson, D. (2003). Ways of Doing Logic: What was Different about AGM 1985? *J. Log. Comput.*, 13(1):5–15.

[43] Peppas, P. (2008). Belief Revision. In van Harmelen, F., Lifschitz, V., and Porter, B., editors, *Handbook of Knowledge Representation*, pages 317–360. Elsevier B.V.

[44] Qi, G. and Du, J. (2009). Model-based Revision Operators for Terminologies in Description Logics. In Boutilier, C., editor, *IJCAI*, pages 891–897.

[45] Qi, G., Liu, W., and Bell, D. (2006). A Revision-based Approach to Handling Inconsistency in Description Logics. *Artif. Intell. Rev.*, 26(1-2):115–128.

[46] Revesz, P. Z. (1997). On the Semantics of Arbitration. *International Journal of Algebra and Computation*, 07(02):133–160.

[47] Ribeiro, M. M. and Wassermann, R. (2009). Base Revision for Ontology Debugging. *J. Log. Comput.*, 19(5):721–743.

[48] Rodrigues, O., Gabbay, D., and Russo, A. (2011). Belief revision. In Gabbay, D. and Guenthner, F., editors, *Handbook of Philosophical Logic, Vol. 16*, pages 1–114. Springer Science+Business Media B.V.

[49] Schwind, N. and Inoue, K. (2013). Characterization Theorems for Revision of Logic Programs. In [8], pages 485–498.

[50] Slota, M. and Leite, J. (2012a). A Unifying Perspective on Knowledge Updates. In del Cerro, L. F., Herzig, A., and Mengin, J., editors, *JELIA*, volume 7519 of *Lecture Notes in Computer Science*, pages 372–384. Springer.

[51] Slota, M. and Leite, J. (2012b). Robust Equivalence Models for Semantic Updates of Answer-Set Programs. In [7], pages 158–168.

[52] Williams, M.-A. (1998). Applications of belief revision. In Freitag, B., Decker, H., Kifer, M., and Voronkov, A., editors, *Transactions and Change in Logic Databases*, volume 1472 of *Lecture Notes in Computer Science*, pages 287–316. Springer Berlin Heidelberg.

[53] Zanuttini, B. and Hébrard, J.-J. (2002). A unified framework for structure identification. pages 335–339.

[54] Zhuang, Z., Wang, Z., Wang, K., and Qi, G. (2014). Contraction and Revision over DL-Lite TBoxes. In Brodley, C. E. and Stone, P., editors, *AAAI*, pages 1149–1156. AAAI Press.

[55] Zhuang, Z. Q. and Pagnucco, M. (2009). Belief Contraction in the Description Logic EL. In Grau, B. C., Horrocks, I., Motik, B., and Sattler, U., editors, *Description Logics*, volume 477 of *CEUR Workshop Proceedings*. CEUR-WS.org.

[56] Zhuang, Z. Q. and Pagnucco, M. (2012). Model Based Horn Contraction. In [7], pages 169–178.

[57] Zhuang, Z. Q., Pagnucco, M., and Zhang, Y. (2013). Definability of Horn Revision from Horn Contraction. In *IJCAI*, pages 1205–1211.