

Default Reasoning from Conditional Knowledge Bases: Complexity and Tractable Cases*

Thomas Eiter[†]

Thomas Lukasiewicz[†]

Abstract

Conditional knowledge bases have been proposed as belief bases that include defeasible rules (also called defaults) of the form “ $\phi \rightarrow \psi$ ”, which informally read as “generally, if ϕ then ψ .” Such rules may have exceptions, which can be handled in different ways. A number of entailment semantics for conditional knowledge bases have been proposed in the literature. However, while the semantic properties and interrelationships of these formalisms are quite well understood, about their computational properties only partial results are known so far. In this paper, we fill these gaps and first draw a precise picture of the complexity of default reasoning from conditional knowledge bases: Given a conditional knowledge base KB and a default $\phi \rightarrow \psi$, does KB entail $\phi \rightarrow \psi$? We classify the complexity of this problem for a number of well-known approaches (including Goldszmidt et al.’s maximum entropy approach and Geffner’s conditional entailment), where we consider the general propositional case as well as natural syntactic restrictions (in particular, to Horn and literal-Horn conditional knowledge bases). As we show, the more sophisticated semantics for conditional knowledge bases are plagued with intractability in all these fragments. We thus explore cases in which these semantics are tractable, and find that most of them enjoy this property on feedback-free Horn conditional knowledge bases, which constitute a new, meaningful class of conditional knowledge bases. Furthermore, we generalize previous tractability results from Horn to q-Horn conditional knowledge bases, which allow for a limited use of disjunction. Our results complement and extend previous results, and contribute in refining the tractability/intractability frontier of default reasoning from conditional knowledge bases. They provide useful insight for developing efficient implementations.

Keywords: Reasoning under uncertainty, conditional knowledge bases, default reasoning, nonmonotonic inference, computational complexity, polynomial-time algorithms, Horn knowledge bases and variants.

1 Introduction

During the past decade, there has been extensive work on laying the foundations of inference systems for plausible reasoning in the presence of incomplete information. In particular, characterizing natural properties of such systems and their inference relations embodied was a major subject of study in nonmonotonic reasoning (cf. [37, 38]).

Conditional knowledge bases. A conditional knowledge base consists of a collection of strict statements in classical logic and a collection of defeasible rules (also called defaults). The former are statements that must always hold, while the latter are rules $\phi \rightarrow \psi$ that read as “generally, if ϕ then ψ .” Such rules may have exceptions, which can be handled in different ways. For example, the knowledge “penguins are birds”

*This paper is a substantially extended and revised version of preliminary work which appeared in: *Proc. 7th International Conference on Principles of Knowledge Representation and Reasoning (KR 2000)*, pp. 62–73, 2000, and in: *Proc. 7th European Workshop on Logics in Artificial Intelligence (JELIA 2000)*, 2000.

[†]Institut und Ludwig Wittgenstein Labor für Informationssysteme, Technische Universität Wien, Favoritenstraße 9-11, A-1040 Vienna, Austria. E-mail: eiter@kr.tuwien.ac.at, lukasiewicz@kr.tuwien.ac.at.

and “penguins don’t fly” can be represented by strict sentences, while the knowledge “birds fly” should be expressed by a defeasible rule (since penguins are birds that do not fly).

The semantics of a conditional knowledge base KB is given by the set of all defaults that are plausible consequences of KB . The literature contains several different proposals for plausible consequence relations and extensive work on their desired properties. The core of these properties are the rationality postulates proposed by Kraus, Lehmann, and Magidor [57]. It turned out that these rationality postulates constitute a sound and complete axiom system for several classical model-theoretic entailment relations under uncertainty measures on worlds. More precisely, they characterize classical model-theoretic entailment under preferential structures [85, 57], infinitesimal probabilities [1, 80], possibility measures [30], and world rankings [86, 48]. Moreover, they characterize an entailment relation based on conditional objects [31]. A survey of all these relationships is given in [8]. We will use the notion of ε -entailment to refer to these equivalent entailment relations. That their equivalence is not incidental is shown by Friedman and Halpern [35], who prove that many approaches are expressible as plausibility measures and thus they must, under some weak natural conditions, inevitably amount to the same notion of inference.

Mainly to solve problems with irrelevant information, the notion of rational closure as a more adventurous notion of entailment has been introduced by Lehmann [61, 63]. This notion of entailment is equivalent to entailment in system Z by Pearl [81] (which is generalized to variable strength defaults in system Z^+ by Goldszmidt and Pearl [47, 49]), to the least specific possibility entailment by Benferhat et al. [7], and to a conditional (modal) logic-based entailment by Lamarre [60]. Finally, mainly in order to solve problems with property inheritance from classes to exceptional subclasses, the maximum entropy approach to default entailment was proposed by Goldszmidt et al. [45] (and recently generalized to variable strength defaults by Bourne and Parsons [13]); the notion of lexicographic entailment was introduced by Lehmann [62] and Benferhat et al. [6]; the notion of conditional entailment was proposed by Geffner [41, 43]; and, an infinitesimal belief function approach was suggested by Benferhat et al. [9].

Motivation and goals of this work. While the semantic properties and interrelationships of the various formalisms are quite well understood, their computational properties are less explored. Algorithms for conditional knowledge bases have been described, for example, in [46, 63, 49, 22]. They are often used for a rough analysis of the computational complexity of the problems they solve. This way, in many cases, only rough upper bounds for the complexity of various computational problems have been established so far.

One of the goals of this paper is to fill these gaps and to draw a precise picture of the computational complexity of major formalisms for default reasoning from conditional knowledge bases. It thus complements and extends the previous work in [46, 63, 49, 22].

Our effort on characterizing the computational complexity of the various semantics serves several purposes. Firstly, precise computational relationships between various formalisms are established, that is, the feasibility of a polynomial time transformation of reasoning in one formalism into reasoning in another one can be assessed from our complexity results. Secondly, the results show that certain algorithms in the literature have optimal order under worst case complexity. Finally, the results give useful insight and background information when new algorithms for default reasoning are designed and practical implementations are developed; note that, to our knowledge, for the various semantics no or only prototype implementations are publicly available to date (see [12, 22]).¹

Another goal of this paper is to find, in the light of the results that emerge in the complexity characterization, meaningful cases in which default reasoning from conditional knowledge bases is tractable. In

¹The group of D. Dubois and H. Prade had developed experimental implementations of semantics equivalent to ε -semantics and system Z in the past for internal use, which have not been disseminated though (D. Dubois, personal communication).

particular, we aim at identifying nontrivial restrictions which, on the one hand, can be checked efficiently and, on the other hand, guarantee sufficient expressiveness such that relevant instances of the problem can be represented.

Main contributions and results. Our main contributions on the above issues are the following:

(1) First and foremost, we give a sharp characterization of the complexity of default reasoning from conditional knowledge bases under several semantics, improving on previous results. In particular, we address the following generic problem: Given a conditional knowledge base KB and a default $\phi \rightarrow \psi$, is it true that KB entails $\phi \rightarrow \psi$? Note that the precise formulation of this problem slightly varies in the different approaches and may involve further parameters. Our analysis includes formalisms for which only very rough or even no complexity results have been derived so far, namely proper ε -entailment [46], maximum entropy entailment [45] together with its variable-strength extension [13], and Geffner’s conditional entailment [41].

(2) We analyze the effect of compilation for ranking-based approaches, in terms of off-line computation of the ranking implicitly associated with the defaults in the knowledge base, such that it can be used on-line for default reasoning. Both the cost of computing the ranking and of its on-line use for default reasoning are examined.

(3) We analyze the impact of syntactical restrictions on the knowledge bases. In particular, we consider the restriction to the Horn case, where all strict statements are Horn clauses and all defeasible rules are of the form $\phi \rightarrow \psi$ with conjunctions of atoms ϕ and conjunctions of Horn clauses ψ , and the restriction to the literal-Horn case, where ψ is additionally a literal.

(4) We present new tractable cases for default reasoning from conditional knowledge bases. For this, we introduce two new classes of conditional knowledge bases, which generalize and restrict Horn conditional knowledge bases, respectively, and can be efficiently recognized. Our class of q-Horn conditional knowledge bases enriches in the spirit of [10] the expressiveness of a Horn KB by allowing limited use of disjunction in both classical statements and defeasible rules. For example, a default $saturday \rightarrow hiking \vee shopping$, which informally expresses that on Saturday, some person is normally out for either hiking or shopping, can be represented in q-Horn KB , while it can not be represented in a Horn KB . On the other hand, our class of feedback-free Horn conditional knowledge bases restricts the literal-Horn case by requesting that, roughly speaking, default consequents do not fire back into the classical knowledge of KB and that the defaults can be grouped into non-interfering clusters of bounded size. A number of examples in the class of feedback-free Horn KB ’s, taken from the literature, are given in Section 6.4. Note that, as shown by Example 6.6, this class allows for expressing taxonomic hierarchies that are augmented by default knowledge. A detailed picture of the hierarchy of all classes of conditional knowledge bases that we consider in this paper is given in Figure 8 on page 28.

Our main findings can be briefly summarized as follows.

- The approaches considered in this paper cover different complexity classes at the low end of the polynomial hierarchy, which range from co-NP (ε -entailment) to Π_2^P (Geffner’s conditional entailment). In general, they have lower complexity than well-known logical formalizations of nonmonotonic reasoning such as default logic, circumscription, or autoepistemic logic [88, 50, 33].
- The off-line computation of rankings does in general not pay off with respect to worst-case complexity, and in particular does not buy tractability. Furthermore, computing the ranking associated with a knowledge base is as difficult as solving the reasoning problem.

- Horn constraints have different effects on the various semantics. For some approaches, the restriction to the Horn case leads to tractability, while for the others, the complexity remains unchanged. Interestingly, for all semantics, Horn and literal-Horn knowledge bases have the same complexity. In particular, Geffner’s conditional entailment is Π_2^P -complete in the literal-Horn case, and thus harder than Reiter’s default logic in this case [56, 87].

- We show that previous tractability results for ε -entailment [46], proper ε -entailment [63, 49], z - and z^+ -entailment [49] in the Horn case can be extended to the q-Horn case. Thus, in all these approaches a limited use of disjunction is possible while tractability is retained. Furthermore, we show that in the feedback-free Horn case, default reasoning under z^* -entailment [45], z_s^* -entailment [13], *lex*-entailment [6], and *lex_p*-entailment [62] is tractable. To our knowledge, no or only limited tractable cases [6] for these notions of entailment from conditional knowledge bases have been identified so far.

- Our tractability results for the feedback-free Horn case are complemented by our proof that without a similar restriction on literal-Horn defaults, all the respective semantics remain intractable. In particular, this applies even for the case of a 1-literal Horn *KB*, in which each default is literal Horn and has at most one atom, and the classical knowledge in *KB* consists of Horn-clauses having at most two literals.

Structure of the paper. The rest of this paper is organized as follows. Section 2 contains some preliminaries on conditional knowledge bases and complexity classes that we need in this paper. In Section 3, we then review the various semantics for conditional knowledge bases that we consider in our study. In Section 4, we first formally define the inference problems to be analyzed, and then, after reviewing previous results, we overview and discuss our complexity results for these semantics. Section 5 is devoted to the proofs of our complexity results, and shows algorithms for some of the semantics. This section may be safely skipped by the reader less interested in technical details. In Section 6, we then explore the tractability/intractability frontier in more detail. We introduce q-Horn and feedback-free Horn default knowledge bases, for which we derive our tractability results and show the intractability results for the 1-literal Horn case. Section 7 considers related work, where we briefly address complexity results for conditional modal logics and discuss related complexity results in the fields of belief revision and nonmonotonic logics. The final Section 8 draws some conclusions and outlines issues for further research.

In order to distract not from the flow of reading, longer proofs and technical details have been moved to Appendices A–C.

2 Preliminaries

2.1 Conditional Knowledge Bases

We assume a set of *basic propositions* (or *atoms*) $At = \{p_1, p_2, \dots, p_n\}$ with $n \geq 1$. We use \perp and \top to denote the propositional constants *false* and *true*, respectively. The set of *classical formulas* is the closure of $At \cup \{\perp, \top\}$ under the Boolean operations \neg and \wedge . Classical formulas will be denoted by Greek lower letters α, β, \dots . We use $(\phi \Rightarrow \psi)$ and $(\phi \vee \psi)$ to abbreviate $\neg(\phi \wedge \neg\psi)$ and $\neg(\neg\phi \wedge \neg\psi)$, respectively, and adopt the usual conventions to eliminate parentheses. A *literal* is an atom p from At or its negation $\neg p$. A *Horn clause* is a classical formula $\phi \Rightarrow \psi$, where ϕ is either \top or a conjunction of atoms, and ψ is either \perp or an atom. A *definite Horn clause* is a Horn clause $\phi \Rightarrow \psi$, where ψ is an atom.

A *conditional rule* (or *default*) is an expression $\phi \rightarrow \psi$, where ϕ and ψ are classical formulas. A *conditional knowledge base* is a pair $KB = (L, D)$, where L is a finite set of classical formulas and D is

a finite set of defaults. Informally, L contains facts and rules that are certain, while D contains defeasible rules. In case $L = \emptyset$, we call KB a *default knowledge base*. A default $\phi \rightarrow \psi$ is *Horn* (resp., *literal-Horn*), if ϕ is either \top or a conjunction of atoms, and ψ is a conjunction of Horn clauses (resp., ψ is a literal). A *definite literal-Horn default* is a literal-Horn default $\phi \rightarrow \psi$, where ψ is an atom. Given a default d , we use $At(d)$ to denote the set of all atoms $a \in At$ that occur in d .

Given a conditional knowledge base $KB = (L, D)$, a *strength assignment* σ on KB is a mapping that assigns each default $d \in D$ a nonnegative integer $\sigma(d)$. A *priority assignment* on KB is a strength assignment π on KB such that $\{\pi(d) \mid d \in D\} = \{0, 1, \dots, k\}$ for some $k \geq 0$. Informally, a priority assignment is a strength assignment in which there are no “empty levels”.

An *interpretation* (or *world*) is a truth assignment $I: At \rightarrow \{\mathbf{true}, \mathbf{false}\}$, which is extended to classical formulas as usual. We use \mathcal{I}_{At} to denote the set of all worlds for At . The world I *satisfies* a classical formula ϕ , or I is a *model* of ϕ , denoted $I \models \phi$, iff $I(\phi) = \mathbf{true}$. I *satisfies* a default $\phi \rightarrow \psi$, or I is a *model* of $\phi \rightarrow \psi$, denoted $I \models \phi \rightarrow \psi$, iff $I \models \phi \Rightarrow \psi$. I *satisfies* a set K of classical formulas and defaults, or I is a *model* of K , denoted $I \models K$, iff I satisfies every member of K . A classical formula ϕ is a *logical consequence* of K , denoted $K \models \phi$, iff each model of K is also a model of ϕ . We write $K \not\models \phi$ iff it is not the case that $K \models \phi$. The world I *verifies* a default $\phi \rightarrow \psi$, denoted $I \models_v \phi \rightarrow \psi$, iff $I \models \phi \wedge \psi$. I *falsifies* a default $\phi \rightarrow \psi$, iff $I \models \phi \wedge \neg\psi$ (that is, $I \not\models \phi \rightarrow \psi$). A set of defaults D *tolerates* a default d under a set of classical formulas L iff $D \cup L$ has a model that verifies d . A set of defaults D is *under L in conflict* with a default $\phi \rightarrow \psi$ iff all models of $D \cup L \cup \{\phi\}$ satisfy $\neg\psi$.

A *world ranking* κ is a mapping $\kappa: \mathcal{I}_{At} \rightarrow \{0, 1, \dots\} \cup \{\infty\}$ such that $\kappa(I) = 0$ for at least one world I . It is extended to all classical formulas ϕ as follows. If ϕ is satisfiable, then $\kappa(\phi) = \min\{\kappa(I) \mid I \in \mathcal{I}_{At}, I \models \phi\}$; otherwise, $\kappa(\phi) = \infty$. A world ranking κ is *admissible* with a conditional knowledge base (L, D) iff $\kappa(\neg\phi) = \infty$ for all $\phi \in L$, and $\kappa(\phi) < \infty$ and $\kappa(\phi \wedge \psi) < \kappa(\phi \wedge \neg\psi)$ for all defaults $\phi \rightarrow \psi \in D$. A *default ranking* σ on D maps each $d \in D$ to a nonnegative integer.

We give an example that illustrates world rankings.

Example 2.1 The strict knowledge “all penguins are birds” and the defeasible knowledge “generally, birds fly”, “generally, penguins do not fly”, and “generally, birds have wings” can be represented by the following conditional knowledge base $KB = (L, D)$ over the set of atoms $At = \{penguin, bird, fly, wings\}$:

$$\begin{aligned} L &= \{penguin \Rightarrow bird\}, \\ D &= \{bird \rightarrow fly, penguin \rightarrow \neg fly, bird \rightarrow wings\}. \end{aligned}$$

It holds $At(bird \rightarrow fly) = \{bird, fly\}$ and $At(penguin \rightarrow \neg fly) = \{penguin, fly\}$.

Figure 1 shows three world rankings κ_0 , κ_1 , and κ_2 . It is easy to verify that κ_0 and κ_1 are admissible with KB (note that κ_0 and κ_1 are in fact the world rankings of KB in system Z and under maximum entropy, respectively). The world ranking κ_3 , however, is not admissible with KB , since L contains the classical formula $penguin \Rightarrow bird$, but $\kappa_3(penguin \wedge \neg bird) = \min(\kappa_3(I_2), \kappa_3(I_6), \kappa_3(I_{10}), \kappa_3(I_{14})) = 4 \neq \infty$. Moreover, D contains the default $bird \rightarrow wings$, but $\kappa_3(bird \wedge wings) = 0 = \kappa_3(bird \wedge \neg wings)$.

2.2 Complexity Classes

We assume some basic knowledge about complexity theory. In particular, we suppose familiarity with the classes P, NP, and co-NP. We now briefly introduce some other classes that we encounter in our analysis (see especially [79, 54, 84, 53] for further background).

The class P^{NP} (resp., NP^{NP}) contains all decision problems that can be solved in deterministic (resp., nondeterministic) polynomial time with an oracle for NP (informally, a subroutine for solving a problem in

	<i>penguin</i>	<i>bird</i>	<i>fly</i>	<i>wings</i>	κ_0	κ_1	κ_2
I_1	false	false	false	false	0	0	0
I_2	true	false	false	false	∞	∞	∞
I_3	false	true	false	false	1	2	2
I_4	true	true	false	false	1	2	2
I_5	false	false	true	false	0	0	0
I_6	true	false	true	false	∞	∞	∞
I_7	false	true	true	false	1	1	1
I_8	true	true	true	false	2	3	0

	<i>penguin</i>	<i>bird</i>	<i>fly</i>	<i>wings</i>	κ_0	κ_1	κ_2
I_9	false	false	false	true	0	0	0
I_{10}	true	false	false	true	∞	∞	4
I_{11}	false	true	false	true	1	1	1
I_{12}	true	true	false	true	1	1	1
I_{13}	false	false	true	true	0	0	0
I_{14}	true	false	true	true	∞	∞	∞
I_{15}	false	true	true	true	0	0	0
I_{16}	true	true	true	true	2	2	2

Figure 1: Some World Rankings

NP at unit cost). They are the classes Δ_2^P and Σ_2^P of the polynomial hierarchy, which has been introduced to capture the intrinsic complexity of problems that have complexity between NP and PSPACE. The class Π_2^P is the complementary class of Σ_2^P , which has Yes- and No-instances interchanged.

The class Δ_2^P has been refined to assess the number and quality of oracle calls for solving a problem:

The class D^P contains the problems that can be described as a logical conjunction of a problem P_1 in NP and a problem P_2 in co-NP. That is, given instances of I_1 and I_2 of P_1 and P_2 , respectively, the answer is “yes” if both I_1 and I_2 are Yes-instances, and “no” otherwise. Any problem in D^P can be solved with two NP oracle calls, and is intuitively easier than a problem complete for Δ_2^P .

The class $\Delta_2^P[O(\log n)]$ contains the problems in Δ_2^P that can be solved with $O(\log n)$ many oracle calls, where n is the size of the problem input. This class, also named Θ_2^P , is very robust and has many different equivalent characterizations [90]. In particular, it coincides with L^{NP} , logspace computability with an NP oracle, and with $P_{\parallel}^{\text{NP}}$, that is, polynomial time computability with an NP oracle where all oracle calls must be first prepared and then issued in parallel.

Qualitatively speaking, membership in $\Delta_2^P[O(\log n)]$ means that the problem can be solved efficiently by *parallelization* to the classical satisfiability problem (SAT), which may be solved by using one of the promising SAT-algorithms that have been developed (see e.g. [29]).

According to the current belief in complexity theory, the following is a strict hierarchy of inclusions:

$$P \subseteq \text{NP, co-NP} \subseteq D^P \subseteq \Delta_2^P[O(\log n)] = L^{\text{NP}} = P_{\parallel}^{\text{NP}} \subseteq \Delta_2^P = P^{\text{NP}} \subseteq \Sigma_2^P, \Pi_2^P.$$

For classifying problems that compute an output value (e.g., the set of atoms that are entailed by a classical formula ϕ), function classes similar to the classes above have been introduced (cf. [84, 53]). In particular, FP, $\text{FP}_{\parallel}^{\text{NP}}$ ($=\text{FL}^{\text{NP}}$), and FP^{NP} are the functional analogs of P, $P_{\parallel}^{\text{NP}}$ ($=L^{\text{NP}}$), and P^{NP} , respectively.

In this paper, unless stated otherwise, completeness for a decision class is with respect to standard polynomial time transformations. Furthermore, completeness for a function class is understood in terms of a natural generalization of polynomial time transformations: The problem P_1 reduces to P_2 , if there are polynomial time functions f and g such that for each instance I_1 of P_1 , the output for I_1 is given by $g(I_1, P_2(f(I_1)))$ ²; see [84, 53] for formal details. In case of P and FP, completeness is understood in terms of reductions that can be computed in logarithmic space.

In the sequel, unless stated otherwise, we consider presumably intractable problems (it has not been proved so far that $P \neq \text{NP}$) as intractable.

²Note that the first argument of g allows to access the original problem instance I_1 .

3 Semantics for Conditional Knowledge Bases

In this section, we recall some of the proposals for a semantics of conditional knowledge bases. To simplify the presentation, we shall adjust original definitions (without significant effects) to our framework, and use characterizations of semantics based on world rankings.

3.1 Examples

We now illustrate the different semantics of conditional knowledge bases along a classical example [49], which extends Example 2.1 by some more defaults.

Example 3.1 Consider the following conditional knowledge base $KB = (L, D)$, which represents the knowledge “all penguins are birds”, “generally, birds fly”, “generally, penguins do not fly”, “generally, birds have wings”, “generally, penguins live in the arctic”, and “generally, flying animals are mobile”.

$$L = \{penguin \Rightarrow bird\},$$

$$D = \{bird \rightarrow fly, penguin \rightarrow \neg fly, bird \rightarrow wings, penguin \rightarrow arctic, fly \rightarrow mobile\}.$$

We would like this conditional knowledge base to entail “generally, birds are mobile” (as birds generally fly, and flying animals are generally mobile) and “generally, red birds fly” (as the property “red” is not mentioned at all in KB and can thus be considered irrelevant to the flying ability of birds). Moreover, KB should entail “generally, penguins have wings” (as the set of all penguins is a subclass of the set of all birds, and thus penguins should inherit all properties of birds), and “generally, penguins do not fly” (as properties of more specific classes should override inherited properties of less specific classes).

The corresponding behavior of ε -entailment, z -entailment (that is, entailment in system Z), z^* -entailment (that is, entailment under maximum entropy), lex -entailment (that is, lexicographic entailment), and conditional entailment is shown in Table 1. In detail, $bird \rightarrow mobile$ is a plausible consequence of KB under all notions of entailment except for ε -entailment. Moreover, in this example, every notion of entailment except for ε -entailment ignores irrelevant information, while every notion of entailment except for ε - and z -entailment shows property inheritance from the class of all birds to the exceptional subclass of all penguins. Finally, the default $penguin \rightarrow \neg fly$ is entailed by KB under all notions of entailment.

Table 1: Plausible consequences of KB under different semantics

	$bird \rightarrow mobile$	$red \wedge bird \rightarrow fly$	$penguin \rightarrow wings$	$penguin \rightarrow \neg fly$
ε -entailment	–	–	–	+
z -entailment	+	+	–	+
z^* -entailment	+	+	+	+
lex -entailment	+	+	+	+
conditional entailment	+	+	+	+

The next example shows how ambiguities are handled under the different semantics.

Example 3.2 Let us now add the knowledge “generally, metal-winged objects fly” and “generally, light objects fly” to the conditional knowledge base $KB = (L, D)$ given in the previous example. That is, let us consider the conditional knowledge base $KB' = (L, D \cup \{metal-wings \rightarrow fly, light \rightarrow fly\})$.

What does KB' say about the ability to fly of light metal-winged penguins? Clearly, KB' is ambiguous on this point. That is, KB' should neither entail that light metal-winged penguins fly, nor that they do not fly.

It turns out that only ε - and conditional entailment show such a behavior. Under z - and lex -entailment, in contrast, KB' entails that light metal-winged penguins do not fly. Furthermore, under z^* -entailment, KB' entails that light metal-winged penguins fly.

Informally, the notions of z^+ -, z_s^* - and lex_p -entailment can be motivated as follows. Every notion of entailment in Table 1 is associated with a set of preference relations on \mathcal{I}_{At} (which is a singleton in case of z -, z^* -, and lex -entailment). These preference relations are implicitly encoded in the structure of KB . The notions of z^+ -, z_s^* - and lex_p -entailment are generalizations of z -, z^* -, and lex -entailment in which we can explicitly characterize these preference relations through additional strength and priority assignments.

3.2 ε -Semantics

We first describe the notions of ε -consistency, ε -entailment, and proper ε -entailment. These notions go back to Adams [1] and Pearl [80]. We define them in terms of world rankings (see especially Geffner's work [41, 42] for the equivalence to the original definitions).

A conditional knowledge base KB is ε -consistent iff there exists a world ranking that is admissible with KB . It is ε -inconsistent iff no such a world ranking exists.

A conditional knowledge base KB ε -entails a default $\phi \rightarrow \psi$ iff either $\kappa(\phi) = \infty$ (that is, ϕ is unsatisfiable) or $\kappa(\phi \wedge \psi) < \kappa(\phi \wedge \neg\psi)$ for all world rankings κ that are admissible with KB . Moreover, KB properly ε -entails $\phi \rightarrow \psi$ iff KB ε -entails $\phi \rightarrow \psi$ and KB does not ε -entail $\phi \rightarrow \perp$.

The next theorem is a simple generalization of a result by Adams [1], who stated it for $L = \emptyset$.

Theorem 3.1 (essentially [1]) *A conditional knowledge base (L, D) ε -entails a default $\phi \rightarrow \psi$ iff the conditional knowledge base $(L, D \cup \{\phi \rightarrow \neg\psi\})$ is ε -inconsistent.*

3.3 Systems Z and Z^+

Entailment in system Z (Pearl [81]) applies to ε -consistent conditional knowledge bases $KB = (L, D)$. It is linked to an ordered partition of D , a default ranking z , and a world ranking κ^z : Let (D_0, \dots, D_k) be the unique ordered partition of D such that, for $i = 0, \dots, k$, each D_i is the set of all defaults in $D - \bigcup\{D_j \mid 0 \leq j < i\}$ that are tolerated under L by $D - \bigcup\{D_j \mid 0 \leq j < i\}$. We call this (D_0, \dots, D_k) the z -partition of D . We next define the default ranking z as follows. For $j = 0, \dots, k$, each $d \in D_j$ is assigned the value j under z . Finally, the world ranking κ^z on all $I \in \mathcal{I}_{At}$ is defined as follows:

$$\kappa^z(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ 1 + \max_{d \in D: I \not\models d} z(d) & \text{otherwise.} \end{cases} \quad (1)$$

A default $\phi \rightarrow \psi$ is z -entailed by KB iff either $\kappa^z(\phi) = \infty$ or $\kappa^z(\phi \wedge \psi) < \kappa^z(\phi \wedge \neg\psi)$.

The notion of entailment in system Z^+ (Goldszmidt and Pearl [47, 49]) applies to ε -consistent conditional knowledge bases $KB = (L, D)$ with strength assignment σ . Entailment in system Z^+ is linked to a default ranking z^+ and a world ranking κ^+ , which are defined as the unique solution of the following system

of equations: For all $d = \phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$:

$$z^+(d) = \sigma(d) + \kappa^+(\phi \wedge \psi) \quad (2)$$

$$\kappa^+(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ 1 + \max_{d \in D: I \not\models d} z^+(d) & \text{otherwise.} \end{cases} \quad (3)$$

We are now ready to define z^+ -entailment as follows. A default $\phi \rightarrow \psi$ is z^+ -entailed by (KB, σ) at strength τ iff either $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \wedge \psi) + \tau < \kappa^+(\phi \wedge \neg\psi)$.

We note that for any ε -consistent conditional knowledge base (L, D) , the default ranking z and the world ranking κ^z coincide with z^+ and κ^+ for (L, D) under strength assignment $\sigma(d) = 0$, for all $d \in D$.

3.4 Maximum Entropy Semantics

The maximum entropy approach to default entailment has been introduced by Goldszmidt et al. [45]. Recently, it has been extended to variable strength defaults by Bourne and Parsons [13, 12].

In detail, z^* -entailment applies to ε -consistent minimal-core conditional knowledge bases $KB = (L, D)$, where KB is *minimal-core* iff for each default $d \in D$ there is a model I of $L \cup (D - \{d\})$ that falsifies d . This notion of entailment is linked to a default ranking z^* and a world ranking κ^* , which are defined as the unique solution of a system of equations similar to (2) and (3). For all $d = \phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$:

$$z^*(d) = 1 + \kappa^*(\phi \wedge \psi) \quad (4)$$

$$\kappa^*(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ \sum_{d \in D: I \not\models d} z^*(d) & \text{otherwise.} \end{cases} \quad (5)$$

A default $\phi \rightarrow \psi$ is z^* -entailed by KB iff either $\kappa^*(\phi) = \infty$ or $\kappa^*(\phi \wedge \psi) < \kappa^*(\phi \wedge \neg\psi)$.

The notion of z_s^* -entailment applies to ε -consistent conditional knowledge bases $KB = (L, D)$ with positive strength assignment σ . This notion of entailment is defined whenever the following system of equations has a unique solution z_s^*, κ_s^* with positive z_s^* .³ For all $\phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$:

$$\kappa_s^*(\phi \wedge \neg\psi) = \sigma(\phi \rightarrow \psi) + \kappa_s^*(\phi \wedge \psi) \quad (6)$$

$$\kappa_s^*(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ \sum_{d \in D: I \not\models d} z_s^*(d) & \text{otherwise.} \end{cases} \quad (7)$$

The uniqueness of z_s^* and κ_s^* is guaranteed by assuming that κ_s^* is *robust* [13], which is the following property: for all distinct defaults d_1 and d_2 in D , it holds that all models I_1 and I_2 of L having smallest

³Note that there may exist unique solutions z_s^*, κ_s^* to (6) and (7) in which some defaults are assigned a zero or negative rank. However, as argued in [12, p. 76], these defaults turn out to be redundant and should thus be removed.

ranks in κ_s^* such that $I_1 \not\models d_1$ and $I_2 \not\models d_2$, respectively, are different. That is, d_1 and d_2 do not have a common minimal falsifying model under L .

We say KB is *robust* iff the system of equations given by (6) and (7), for all $\phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$, has a unique solution z_s^*, κ_s^* such that z_s^* is positive and κ_s^* is robust.

We are now ready to define z_s^* -entailment as follows. A default $\phi \rightarrow \psi$ is z_s^* -entailed by (KB, σ) at strength τ iff either $\kappa_s^*(\phi) = \infty$ or $\kappa_s^*(\phi \wedge \psi) + \tau \leq \kappa_s^*(\phi \wedge \neg\psi)$.

The notion of z_s^* -entailment is a proper generalization of z^* -entailment:

Lemma 3.2 *Let $KB = (L, D)$ be a conditional knowledge base with strength assignment $\sigma(d) = 1$ for all $d \in D$. Suppose KB is ε -consistent and minimal-core. Then, the system of equations given by (6) and (7) for all $\phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$ has a unique solution z_s^*, κ_s^* , which coincides with z^*, κ^* . Moreover, κ_s^* is robust (and thus, also KB is robust).*

3.5 Lexicographic Entailment

The notion of lexicographic entailment goes back to Lehmann [62] and Benferhat et al. [6].

Lexicographic entailment as introduced in [6] applies to conditional knowledge bases $KB = (L, D)$ with priority assignment π , which defines an ordered partition (D_0, \dots, D_k) of D by $D_i = \{d \in D \mid \pi(d) = i\}$, for all $i \leq k$. It is used to define a *preference ordering* on worlds as follows. A world I is π -preferable to a world I' iff there exists some $i \in \{0, \dots, k\}$ such that $|\{d \in D_i \mid I \models d\}| > |\{d \in D_i \mid I' \models d\}|$ and $|\{d \in D_j \mid I \models d\}| = |\{d \in D_j \mid I' \models d\}|$ for all $i < j \leq k$. Note that this preference ordering can be expressed by a world ranking. A model I of a set of classical formulas \mathcal{F} is a π -preferred model of \mathcal{F} iff no model of \mathcal{F} is π -preferable to I .

A default $\phi \rightarrow \psi$ is lex_p -entailed by (KB, π) iff ψ is satisfied in every π -preferred model of $L \cup \{\phi\}$. We will omit π when it is clear from the context.

Note that lex_p -entailment is the only semantics for conditional knowledge bases among the ones examined in this paper in which a default $d \in D$ is not necessarily entailed by an ε -consistent $KB = (L, D)$.

The notion of lexicographic entailment in [62] is a special case of lexicographic entailment as above. It uses a particular priority assignment that is logically entrenched in KB , namely the default ranking z of KB (see Section 3.3). We then say that a default is *lex-entailed* by KB iff $\phi \rightarrow \psi$ is lex_p -entailed by (KB, z) . Note that this definition assumes that the default ranking z of KB exists, that is, that KB is ε -consistent.

It appears that, in a certain sense, *lex*-entailment is not less expressive than lex_p -entailment. That is, under a weak condition, priority assignments are expressible through logical entrenchment:

Theorem 3.3 *Let $KB = (L, D)$ be a conditional knowledge base such that every $d \in D$ has a verifying world, and let π be a priority assignment on KB . Then, there exists a conditional knowledge base $KB' = (L', D')$ and a formula ϕ' (depending only on KB and π) such that, for any default $\phi \rightarrow \psi$ over At , it holds that (KB, π) lex_p -entails $\phi \rightarrow \psi$ iff KB' *lex-entails* $\phi \wedge \phi' \rightarrow \psi$.*

Proof. The main idea behind the construction of KB' and ϕ' is to augment D by additional defaults such that the default ranking z of KB' coincides with the priority assignment π (see Appendix A). \square

The proof of the previous theorem shows in fact that the transformation of lex_p -entailment to *lex*-entailment is compliant with the Horn property. For later reference, we note the following.

Observation 3.1 *Let the conditional knowledge base KB' and the classical formula ϕ' be defined as in the proof of Theorem 3.3. Then, KB' is literal-Horn whenever KB is literal-Horn. Moreover, ϕ' is a conjunction of atoms. Finally, KB' and ϕ' can be constructed in polynomial time from KB and π .*

3.6 Conditional Entailment

The notion of conditional entailment has been introduced by Geffner [41, 43].

Given a conditional knowledge base $KB = (L, D)$, a *priority ordering* \prec on D is an irreflexive and transitive binary relation on D . We say \prec is *admissible* with KB iff each set of defaults $D' \subseteq D$ that is under L in conflict with some default $d \in D$ contains a default d' such that $d' \prec d$.

Based on \prec , we define a *preference ordering* on worlds as follows. A world I is \prec -*preferable* to a world I' , denoted $I \prec I'$, iff $\{d \in D \mid I \not\models d\} \neq \{d \in D \mid I' \not\models d\}$ and for each default $d \in D$ such that $I \not\models d$ and $I' \models d$, there exists a default $d' \in D$ such that $d \prec d'$, $I \models d'$, and $I' \not\models d'$. A model I of a set of classical formulas \mathcal{F} is a \prec -*preferred* model of \mathcal{F} iff no model of \mathcal{F} is \prec -preferable to I .

A default $\phi \rightarrow \psi$ is *conditionally entailed* by KB iff ψ is satisfied in every \prec -preferred model of $L \cup \{\phi\}$ of every priority ordering \prec that is admissible with KB .

A conditional knowledge base $KB = (L, D)$ is *conditionally consistent* iff there is a priority ordering \prec on D that is admissible with KB . The following lemma shows that in our framework of *finite* conditional knowledge bases, the notion of ε -consistency coincides with the notion of conditional consistency.

Lemma 3.4 *A conditional knowledge base KB is ε -consistent iff it is conditionally consistent.*

4 Complexity Characterization

In this section, we present and discuss our results on the complexity of the semantics described in the previous section. Prior to this, we need a formalization of the problems considered, which is given next.

4.1 Problem Statements

A *default reasoning problem* is a pair (KB, d) , where $KB = (L, D)$ is a conditional knowledge base and d is a default. It is *Horn* iff L is a finite set of Horn clauses, D is a finite set of Horn defaults, and d is a Horn default. It is *literal-Horn* iff L is a finite set of Horn clauses, D is a finite set of literal-Horn defaults, and d is a literal-Horn default. In case of z^+ - and z_s^* -entailment, we assume that KB and d have additionally a strength assignment $\sigma(KB)$ and a strength $\tau(d)$, respectively. In case of lex_p -entailment, we assume that KB has in addition a priority assignment $\pi(KB)$. The taxonomic hierarchy of default reasoning problems emerging from the definitions is shown in Figure 2, where the newly introduced literal-Horn class is emphasized in bold face.




Figure 2: Hierarchy of syntactic restrictions

Informally, a default reasoning problem represents the input for the entailment problem under a fixed semantics \mathcal{S} . We tacitly assume that KB satisfies any preconditions that the definition of \mathcal{S} -entailment in the previous section may request.

We analyze the computational complexity of the following problems:

- **ENTAILMENT:** Given a default reasoning problem (KB, d) , decide whether KB entails d under some fixed semantics \mathcal{S} . In case of z^+ - and z_s^* -entailment, decide whether d is z^+ - and z_s^* -entailed, respectively, by $(KB, \sigma(KB))$ at strength $\tau(d)$. In case of lex_p -entailment, we are asked whether d is lex_p -entailed by $(KB, \pi(KB))$.
- **RANKING:** Given a conditional knowledge base KB , compute the default ranking R of KB according to some fixed semantics \mathcal{S} (that is, the rank of each default in D).
- **RANK-ENTAILMENT:** Same as entailment, but the (unique) default ranking R of KB according to some fixed semantics \mathcal{S} is part of the problem input.

The problems RANKING and RANK-ENTAILMENT are relevant from a preprocessing perspective, in which the ranking R of a conditional knowledge base KB is computed in advance and then on-line available in the input for solving an entailment problem. The complexities of these problems give us some insight to the question of whether such preprocessing pays off in general.

4.2 Previous Results

As shown in Tables 3–5, complexity results for default reasoning from conditional knowledge bases have been obtained by several authors [46, 63, 49, 22]. Most of these results have been derived for default knowledge bases, though, and do not give a sharp complexity characterization.

Goldszmidt and Pearl [46] showed that deciding ε -consistency for general conditional knowledge bases (resp., Horn conditional knowledge bases) is in P^{NP} (resp., P). Lehmann and Magidor [63] proved that deciding preferential entailment (and thus also ε -entailment) for default knowledge bases is co-NP-complete. Furthermore, Lehmann and Magidor [63] and Goldszmidt and Pearl [49] showed that deciding ε -entailment for Horn default knowledge bases is in P. Finally, Goldszmidt and Pearl [46] proved that proper ε -entailment for general conditional knowledge bases (resp., Horn conditional knowledge bases) is in P^{NP} (resp., P). All these results easily carry over to our conditional knowledge bases. As the proofs are obtained by simple adjustments of the proofs for default knowledge bases, we omit them in this paper.

As for system Z^+ , the comprehensive work of Goldszmidt and Pearl [49] provides us with the following complexity results. As shown there, the problems ENTAILMENT, RANKING, and RANK-ENTAILMENT are all solvable in polynomial time for Horn default knowledge bases, while for general default knowledge bases, membership in the classes P^{NP} , FP^{NP} , and P_{\parallel}^{NP} , respectively, is an upper bound. A fortiori, since system Z is an instance of system Z^+ , all these upper bounds also hold for system Z . Again, it is straightforward that all these results carry over to our conditional knowledge bases.

Not much work has been done on determining the complexity of z^* - and z_s^* -entailment. Goldszmidt et al. [45] suspect that the complexity of z^* -entailment is high, and briefly note that, referring to results on Horn clause optimization [5], the problem should be NP-hard in the Horn case and thus intractable.

Cayrol et al. show in [22] that lex_p -entailment is P^{NP} -complete for default knowledge bases. Moreover, they state in [22] that lex_p -entailment is P^{NP} -hard for Horn default knowledge bases. However, the short proof sketch in [22] is inappropriate, since it mentions a reduction from a problem that is obviously in P_{\parallel}^{NP} ; this would only establish P_{\parallel}^{NP} -hardness for the Horn case. P^{NP} -hardness for the Horn case follows from proofs of related results by Nebel on the complexity of lexicographic belief revision [73] (see Section 7.2).

To our knowledge, no complexity results on Geffner’s conditional entailment have been derived so far.

4.3 Overview and Discussion

Our results on the complexity of default reasoning from conditional knowledge bases, together with results from the literature, are compactly summarized in Tables 3–5. They contain the three problems ENTAILMENT, RANKING, and RANK-ENTAILMENT from above, each of which is considered for the general case and the restrictions to the Horn and literal-Horn case.

It appears that a number of different complexity classes from P up to Π_2^P , the second level of the polynomial hierarchy, are covered. A first observation is that Geffner’s conditional entailment has the highest complexity (Π_2^P , Table 3) of all the formalisms considered in this paper. It is thus in the same league as a number of approaches to belief revision (see [55, 39, 73]) and major formalisms of nonmonotonic reasoning, such as circumscription [67, 65], Reiter’s default logic [82], McDermott and Doyle’s nonmonotonic logic [69, 68], and Moore’s autoepistemic logic [70], which are all Π_2^P -complete (see Sections 7.2 and 7.3 for further discussion). All other approaches in Table 3 have (considerably) lower complexity.

4.3.1 General case

At the low end of the complexity range, there are ε -entailment, which has the same complexity as classical logic, and proper ε -entailment, which has marginally higher complexity due to the additional ε -entailment requirement. At the high end, we have Geffner’s conditional entailment. Its high complexity is intuitively explained by an inherent pattern similar to reasoning under circumscription: To disprove that $KB = (L, D)$ entails $\phi \rightarrow \psi$, a \prec -preferred model I of $L \cup \{\phi\}$ under some admissible priority ordering \prec must be found such that ψ is false in I . As it turns out, such a guess can be verified in polynomial time with an NP oracle, where the oracle checks the \prec -preferredness of I (that is, minimality under \prec). This is similar to circumscription, that is, minimal model reasoning, where for disproving $CIRC(\phi) \models \psi$ the minimality of a guessed model M of ϕ in which ψ is false must be verified, which is a co-NP-complete problem [18] (see Section 7.3 for further discussion).

Also for the ranking-based approaches in Table 3 (that is, z -, z^+ -, z^* -, z_s^* -, lex -, and lex_p -entailment), the problem of verifying whether a model I of a formula ϕ is selected on the basis of $KB = (L, D)$ is (at least) co-NP-hard in general. However, there is a qualitative difference between them and Geffner’s approach. Each world ranking r induces a *modular partial ordering* on the models of $L \cup \{\phi\}$, in which *any two distinct models I_1 and I_2 of $L \cup \{\phi\}$ are comparable by their ranks $r(I_1)$ and $r(I_2)$* . The models with the same rank form a cluster, and the clusters are totally ordered by these ranks. The “preferred” models I of $L \cup \{\phi\}$ are those which have minimal rank $r(I)$. Using an NP oracle, it is possible to compute this minimal rank $r(I)$ in polynomial time, which is a polynomial-size certificate for recognizing preferred models efficiently. Intuitively, we have here a single well-connected search space, in which all preferred models of $L \cup \{\phi\}$ can be nailed down by this certificate.

On the other hand, in Geffner’s conditional entailment, two models $I_1 \neq I_2$ of $L \cup \{\phi\}$ may be incomparable, that is, neither $I_1 \prec I_2$ nor $I_2 \prec I_1$ may hold. In general, the search space for a preferred model of $L \cup \{\phi\}$ generally splits into an *exponential number* of completely disconnected search spaces, each of which is intractable and may contain a preferred model we are looking for. Moreover, there is no certificate computable in polynomial time with an NP oracle such that we can recognize preferred models efficiently from it (unless the polynomial hierarchy collapses).

More precisely, Geffner’s approach suffers from two sources of complexity: (i) the number of candidates for a preferred model I of $L \cup \{\phi\}$ in the possibly exponentially many disconnected search spaces, which are generated by incomparability of two models $I_1 \neq I_2$ of $L \cup \{\phi\}$ because of default violation; and (ii) the possibly exponential number of models I' of $L \cup \{\phi\}$ that are preferred to I , according to the admissibility

Table 2: Complexity of deciding ε -consistency

	general case	Horn case	literal-Horn case
ε -consistency	NP-complete	P-complete *	P-complete *

Table 3: Complexity of deciding entailment

	general case	Horn case	literal-Horn case
ε -entailment	co-NP-complete ⁺	P-complete **	P-complete **
proper ε -entailment	D^P -complete	P-complete *	P-complete *
z -entailment	P_{\parallel}^{NP} -complete	P-complete ***	P-complete ***
z^+ -entailment	P^{NP} -complete ***	P-complete ***	P-complete ***
z^* -entailment	P^{NP} -complete	P^{NP} -complete	P^{NP} -complete
z_s^* -entailment	P^{NP} -complete	P^{NP} -complete	P^{NP} -complete
lex -entailment	P^{NP} -complete	P^{NP} -complete	P^{NP} -complete
lex_p -entailment	P^{NP} -complete ++	P^{NP} -complete ++	P^{NP} -complete
conditional entailment	Π_2^P -complete	Π_2^P -complete	Π_2^P -complete

Table 4: Complexity of computing default rankings

	general case	Horn case	literal-Horn case
z	FP_{\parallel}^{NP} -complete	FP-complete ***	FP-complete ***
z^+	FP^{NP} -complete ***	FP-complete ***	FP-complete ***
z^*	FP^{NP} -complete	FP^{NP} -complete	FP^{NP} -complete
z_s^*	FP^{NP} -complete	FP^{NP} -complete	FP^{NP} -complete

Table 5: Complexity of deciding entailment given the default rankings

	general case	Horn case	literal-Horn case
z -entailment	P_{\parallel}^{NP} -complete ***	P-complete ***	P-complete ***
z^+ -entailment	P_{\parallel}^{NP} -complete ***	P-complete ***	P-complete ***
z^* -entailment	P^{NP} -complete	P^{NP} -complete	P^{NP} -complete
z_s^* -entailment	P^{NP} -complete	P^{NP} -complete	P^{NP} -complete

* Membership shown in [46].

** Membership shown in [63, 49] for default knowledge bases.

*** Membership was shown in [49] for default knowledge bases.

+ Shown in [63] for default knowledge bases.

++ Reported in [22] for default knowledge bases; the proof sketch for the Horn case in [22] shows merely P_{\parallel}^{NP} -hardness.

ordering \prec . Note that, contrary to expectation, classical inference is *not* listed as a principal source of complexity here, as results on the Horn restrictions (discussed below) show.

The mid-range of complexity is covered by the ranking-based approaches. Roughly speaking, for any classical formula α , the rank $r(\alpha)$ can be computed as follows:

1. Compute the default ranking R for KB ;
2. Compute $r(\alpha) = \min_{I \models \alpha} r(I)$, using R .

Algorithms for computing the default ranking R in Step 1 have been described in the literature. They can be reformulated to run in P^{NP} (see Section 5). Step 2 is feasible in P^{NP} by doing binary search on the range of the possible values for $r(I)$. This means that the condition “ $r(\phi) = \infty$ or $r(\phi \wedge \psi) < r(\phi \wedge \neg\psi)$ ” for the entailment of $\phi \rightarrow \psi$ from KB is decidable in P^{NP} , by simply checking the satisfiability of $L \cup \{\phi\}$, and if needed computing $r(\phi \wedge \psi)$, $r(\phi \wedge \neg\psi)$ and comparing them.

The complexity of Steps 1 and 2 is shown in Tables 4 and 5, respectively. As for Step 1, the rank $R(d)$ of a default d may range, except in case of z , over exponentially many possible values; in case of z , it ranges over $[0, \dots, n - 1]$ and thus over a linear number of values.⁴ Informally, the ranking R can be constructed bottom up, starting with defaults having lowest rank, and then computing the rank of the next default by doing a binary search on the range of its possible values. This resembles the FP^{NP} -complete problem of computing the lexicographic maximum model of a formula ϕ [58] and suggests that computing R has the same complexity. This intuition turns out to be correct in all cases except one. In case of z -entailment, it is possible to compute with parallel queries to an NP oracle in polynomial time a certificate, given by the sum of all ranks of all defaults, which allows to verify a proper guess for the ranking z (including further auxiliary data) in polynomial time. Given this certificate, recognizing the rank $z(d)$ of a default is in NP; since the number of possible values for $z(d)$ is bounded by the number of defaults, this means that the ranking z can then be determined with a polynomial number of parallel queries to an NP oracle. Since two rounds of parallel NP oracle queries can be replaced by a single round of NP queries (cf. [17]), this means that computing the z -ranking is in FP_{\parallel}^{NP} . We remark that a similar FP_{\parallel}^{NP} algorithm is also feasible for z^+ , if the strengths $\sigma(d)$ of defaults are bounded by a polynomial in the number of defaults.

Table 5 tells us that in all cases except z^+ , entailment does not become easier if the default ranking R is known. Thus, from a worst case perspective, precomputing the default ranking R does not pay off (but clearly saves time over repetitive computations). In case of z^+ , entailment becomes easier, because only the order of the defaults in z^+ is relevant, but not their actual ranks (which can thus be replaced by values from $[0, \dots, n - 1]$, and thus z^+ -entailment reduces to z -entailment).

4.3.2 Horn and literal-Horn case

In all tables, the results for the Horn and the literal-Horn case are the same. Thus, although it is in general not possible to simply split a Horn default $\phi \rightarrow \psi_1 \wedge \dots \wedge \psi_m$ into a semantically equivalent set of literal-Horn defaults $\phi \rightarrow \psi_1, \dots, \phi \rightarrow \psi_m$, it is possible to rewrite a Horn default reasoning problem to a literal-Horn one in polynomial time. Thus, the restriction of the general Horn to the literal-Horn case does not decrease the complexity.

⁴The difference between exponentially many possible values under z^* and a linear number of values under z is due to the difference between computing the world ranks by summation and maximization in equations (5) and (1), respectively.

At the low end of the complexity range, the Horn restriction gives tractability, as was (essentially) shown in [46, 63, 49]. Whereas, at the high end, Geffner’s conditional entailment has surprisingly its full complexity already in the literal-Horn case. This is exceptional, since related formalisms such as Reiter’s default logic and circumscription have lower complexity (more precisely, co-NP) in the Horn case [56, 87, 21].

Informally, in Geffner’s conditional entailment, the Horn property is not sufficient to eliminate the intractability of the preference check for a model of a formula, as this test is not simply reducible to a polynomial number of Horn satisfiability tests, as e.g. in Reiter’s default logic. Further syntactic restrictions are needed. One such restriction, which is efficiently checkable, is that the admissibility ordering \prec is empty.

At the mid-range of complexity, the syntactic restrictions have different effects. Tractability for z - and z^+ -entailment is gained since in the respective entailment algorithms, the NP oracle can be replaced by a polynomial time procedure. On the other hand, for z^* - and z_s^* -entailment, the Horn restriction does not decrease complexity. The intuitive reason is that here the rank of a default is given by the smallest *sum of ranks* of a set of violated defaults plus a value, while in the former cases this was the *maximum rank* of a single violated default plus a value. Computing the smallest sum is an intractable optimization problem, as there are exponentially many such sets; as pointed out in [45], it makes z^* (thus also z_s^*) entailment NP-hard in the Horn case. Contrary to this, the maximum violated rank can be computed by simply looping through the already ranked defaults.

Finally, in the Horn case, the precomputation of default rankings does not reduce the complexity of the entailment problem. This is explained by the fact that computing the rank of a classical formula ϕ amounts essentially to ranking a default.

4.3.3 Bottom line of the results

Our results and their discussion in the previous subsections lead us to the following conclusions.

- Among all the semantics that we analyze in this paper, Geffner’s conditional entailment is by far the computationally most expensive approach, and its computational nature is different from those of the other approaches. This complements the observation that the semantical relationships between Geffner’s and the other approaches are less established.
- For the other approaches, there is a trade-off between epistemic sophistication and computational complexity. More precisely, in the general case, the price for a more sophisticated semantics is rather modest and generally does not lead to another complexity level. Nonetheless, it affects properties such as efficient parallelization to SAT, which is only possible for ε -, proper ε -, and z -entailment. In the Horn case, however, an appealing semantics comes at the expense of computational intractability.
- Precomputing default rankings buys nothing or only little in the worst case (in particular, tractability cannot be gained this way).

4.4 Implications for Implementation

As far as implementations are concerned, our results provide a useful insight into the type of algorithms that may be feasible for solving default reasoning problems.

First of all, our results imply that polynomial time translations of the default reasoning problems into suitable other reasoning problems are feasible, such that existing algorithms and theorem provers might be used as an implementation platform. For example, deciding ε -consistency can be polynomially translated

to a classical SAT instance. Such a translation can be easily extracted from the proof of NP-membership (Theorem 5.2). Any of the sophisticated SAT packages (see e.g. [29]) can then be used for solving this instance. Moreover, ε -entailment and proper ε -entailment can be similarly polynomially reduced to one resp. two calls of a SAT procedure, which may be processed in parallel. As for the other semantics, theorem provers for logics with complexity up to Σ_2^P are needed as host for efficient translations. For example, DLV [78], DeRes [23], or a disjunctive extension of smodels [75], which all provide this expressiveness, might be used, as well as theorem provers based on quantified Boolean formulas [20, 83, 32]. However, efficient transformations of the problems to these logics remain to be designed.

In the case of problems with complexity P^{NP} or $P_{||}^{NP}$, such translations might not be very appealing, since the theorem provers mentioned above are tailored for solving problems whose complexity characteristics is given by the Σ_i^P (resp., Π_i^P) classes of the polynomial hierarchy, rather than the $P^{\Sigma_i^P}$ and $P_{||}^{\Sigma_i^P}$ classes. The definition and computational nature of these approaches suggests that reductions to optimization problems in integer programming might be a more suitable alternative (cf. [4, 89] for a similar approach in the areas of nonmonotonic reasoning and planning). Notice that e.g. computing a minimum nonnegative integer solution for a system of linear integer inequalities is FP^{NP} -hard, and thus all the problems in Tables 2–5 with complexity at most P^{NP} resp. FP^{NP} can be reduced to this problem.

For the development of genuine algorithms, the following can be learned from the complexity results in Tables 2–5 (see also [34] for similar considerations). The problems (resp., their complements) with complexity in NP (resp., co-NP) can be implemented by a standard backtracking algorithm. Such an algorithm is not expected to run in polynomial time, though. In case for the $P_{||}^{NP}$ - and $FP_{||}^{NP}$ -complete problems, parallelization of instances to problems in NP such as SAT might be implemented, along the lines of the algorithms exhibited in the proofs of Theorems 5.6 and 5.7. Alternatively, as done in [49] for problem RANK-ENTAILMENT in System Z^+ , algorithms can be designed for RANK-ENTAILMENT and ENTAILMENT which solve the problems with a logarithmic number of calls to an NP oracle. However, a similar algorithm for RANKING is unlikely to exist, since there is some evidence which suggests that FP^{NP} with logarithmically bounded oracle access is less powerful than $FP_{||}^{NP}$ [53].

For the problems with P^{NP} resp. FP^{NP} complexity, no efficient parallelization to NP problems seems feasible, and at least a linear number of calls to an NP oracle is mandatory. In principle, a backtracking strategy for finding an optimal solution could be pursued; it may use in a standard way the value of a best solution found so far for pruning the search space, but will have to explore an exponentially large tree.

Finally, in case of Geffner’s conditional entailment, the Π_2^P -completeness result means that a simple “flat” backtracking algorithm for disproving conditional entailment of $\phi \rightarrow \psi$ from KB which searches a tree for a polynomial solution path, where each step is efficiently possible (like, e.g., in the Davis-Putnam procedure), is infeasible. Rather, a tree must be searched whose terminal nodes represent instances of a co-NP-complete problem. This can be accomplished by a backtracking algorithm using calls to a classical propositional theorem prover (or, a SAT procedure and reversing the result); alternatively, a nested backtracking strategy is needed.

5 Derivation of Complexity Results

In this section, we formally derive the complexity results that have been summarized in the previous section. To simplify the treatment, we shall state known results at the beginning of each subsection, and prove the cases that remain for concluding the results in Tables 3 through 5. For the proofs of upper bounds, we shall describe algorithms that establish the results.

5.1 ε -Semantics

We start with the complexity of deciding ε -consistency, ε -entailment, and proper ε -entailment.

We now prove the new complexity results stated in Tables 2 and 3. We need the following lemma, which is essentially a reformulation of a similar result in [46] and [49].

Lemma 5.1 *A conditional knowledge base (L, D) is ε -consistent iff there exists an ordered partition (D_0, \dots, D_k) of D such that each default in D_i is tolerated under L by $\bigcup_{j=i}^k D_j$.*

We first show that the problem of deciding whether a conditional knowledge base is ε -consistent is NP-complete in the general case and P-hard in the literal-Horn case. Note that the former improves on the result in [46] that ε -consistency can be decided with a quadratic number of calls to a SAT oracle.

Theorem 5.2 *a) Deciding whether a given conditional knowledge base is ε -consistent is NP-complete. b) Deciding whether a given literal-Horn conditional knowledge base is ε -consistent is P-hard.*

Proof. a) We first show membership in NP. By Lemma 5.1, a conditional knowledge base (L, D) is ε -consistent iff (i) there exists an ordered partition (D_0, \dots, D_k) of D and (ii) for each set of defaults D_i and each default $d \in D_i$ there exists an interpretation I_d^i such that I_d^i is a model of L , that I_d^i verifies d , and that I_d^i satisfies $\bigcup_{j=i}^k D_j$. Such an ordered partition and such interpretations can be guessed and verified by a nondeterministic algorithm in polynomial time.

To show NP-hardness, we give a polynomial transformation from the NP-complete problem of deciding whether a propositional formula in CNF is satisfiable [40]. Let ϕ be a propositional formula in CNF. By Lemma 5.1, ϕ is satisfiable iff the default knowledge base $(\emptyset, \{\top \rightarrow \phi\})$ is ε -consistent.

b) See Appendix B. \square

Using ideas from the proof of Theorem 5.2 b), it can be shown that deciding whether a conditional knowledge base ε -entails a default is P-hard in the literal-Horn case:

Theorem 5.3 *Given a literal-Horn conditional knowledge base KB and a literal-Horn default d , deciding whether KB ε -entails d is P-hard.*

We finally show that deciding whether a conditional knowledge base properly ε -entails a default is D^P -complete in the general case and P-hard in the literal-Horn case. Note that this improves on the result in [46] that proper ε -entailment can be decided with a quadratic number of calls to a SAT oracle.

Theorem 5.4 *a) Given a conditional knowledge base KB and a default d , deciding whether KB properly ε -entails d is D^P -complete. b) Given a literal-Horn conditional knowledge base KB and a literal-Horn default d , deciding whether KB properly ε -entails d is P-hard.*

Proof. a) We now show membership in D^P . By Theorem 3.1, KB properly ε -entails $\phi \rightarrow \psi$ iff KB ε -entails $\phi \rightarrow \psi$ and $KB \cup \{\phi \rightarrow \neg\psi\}$ is ε -consistent. The problem of deciding whether KB ε -entails $\phi \rightarrow \psi$ is in co-NP [63]. By Theorem 5.2 a), the problem of deciding whether $KB \cup \{\phi \rightarrow \neg\psi\}$ is ε -consistent is in NP. Hence, the problem of deciding whether KB properly ε -entails $\phi \rightarrow \psi$ is in D^P .

The proofs of D^P -hardness and of b) are given in Appendix B. \square

5.2 Systems Z and Z^+

We next prove the new complexity results shown in Tables 3–5 for systems Z and Z^+ . In particular, we shall prove that each of the well-known upper bounds for system Z^+ is tight. We first, however, consider system Z for general conditional knowledge bases, for which we obtain slightly lower complexity.

Algorithm z -entailment**Input:** ε -consistent conditional knowledge base (L, D) and a default $\phi \rightarrow \psi$.**Output:** “Yes”, if (L, D) z -entails $\phi \rightarrow \psi$, otherwise “No”.

1. **if** $L \cup \{\phi\}$ is unsatisfiable **then return** “Yes”;
2. **for each** $i \in \{0, 1, \dots, n(n-1)/2\}$ **do**;
 if D has an admissible partition (D_0, \dots, D_k) of weight i
 then $w[i] := \text{true}$ **else** $w[i] := \text{false}$;
3. $w := \min\{i \mid w[i] = \text{true}, 0 \leq i \leq n(n-1)/2\}$;
4. $\alpha_1 := \phi \wedge \psi$; $\alpha_2 := \phi \wedge \neg\psi$;
5. **for each** $i \in \{0, 1, \dots, n\}$ and $j \in \{1, 2\}$ **do**
 if D has an admissible partition (D_0, \dots, D_k) of weight w and
 $L \cup \{\alpha_j\}$ has a model I s.t. $i = 1 + \max(\{\ell \mid I \not\models d, \text{ for some } d \in D_\ell\} \cup \{-1\})$
 then $k[i, j] := \text{true}$ **else** $k[i, j] := \text{false}$;
6. **for each** $j \in \{1, 2\}$ **do** $k_j := \min(\{i \mid k[i, j] = \text{true}, 0 \leq i \leq n\} \cup \{n+1\})$;
7. **if** $k_1 < k_2$ **then return** “Yes” **else return** “No”.

Figure 3: Algorithm z -entailment

We introduce some further notion. Given a conditional knowledge base $KB = (L, D)$, an ordered partition (D_0, \dots, D_k) of D is *admissible* with KB iff, for all $i = 0, \dots, k$, each default in D_i is tolerated under L by $\bigcup_{j=i}^k D_j$. The *weight* of an ordered partition (D_0, \dots, D_k) of D is defined as $\sum_{i=0}^k i \cdot |D_i|$.

Lemma 5.5 *Let $KB = (L, D)$ be a conditional knowledge base. The z -partition of D is the unique ordered partition (D_0, \dots, D_k) of D that is admissible with KB and that has the least weight.*

Armed with this result, we prove the following theorem.

Theorem 5.6 *a) Given a conditional knowledge base KB and a default d , deciding whether KB z -entails d is $P_{\parallel}^{\text{NP}}$ -complete. b) Given a literal-Horn conditional knowledge base KB and a literal-Horn default d , deciding whether KB z -entails d is P-hard.*

Proof. a) We now show membership in $P_{\parallel}^{\text{NP}}$. Towards this goal, consider algorithm z -entailment. It is not hard to see that it correctly decides whether KB z -entails the default $\phi \rightarrow \psi$: In Step 3, w is assigned the weight of the z -partition of D . Knowledge of w enables an easy check whether a given admissible partition (D_0, \dots, D_k) of D is the (unique) z -partition of D . This is exploited in Step 5, where it is checked for both $\phi \wedge \psi$ and $\phi \wedge \neg\psi$ whether they have a model I of rank at most i , for all $i \in \{0, 1, \dots, n\}$; note that for $i = 0$, the condition amounts to $I \models L \cup D$. In Step 6, the ranks of $\phi \wedge \psi$ and $\phi \wedge \neg\psi$ are determined from the results of these checks, after which entailment is decided.

Observe that all queries in Step 2 are in NP and can be decided in parallel; similar in Step 5. Taking Step 1 into account, the algorithm thus makes three “rounds” of parallel calls to an NP-oracle. It is well-known [17] that constantly many rounds of parallel NP oracle queries in a polynomial-time computation can be replaced by a single one. In algorithm z -entailment, the three rounds can be replaced by one round with $(1 + n(n-1)/2)(2n+2)$ parallel NP oracle queries (one for each triple (w, i, j) with $w \in \{0, 1, \dots, n(n-1)/2\}$, $i \in \{0, 1, \dots, n\}$, and $j \in \{1, 2\}$) from which the results of Steps 1, 3, and 6 can be concluded, and which we omit here for brevity. Thus the problem is in $P_{\parallel}^{\text{NP}}$. (Note that by applying binary search in Steps 2 and 5, w and k_1, k_2 can be computed in polynomial time with $O(\log n)$ successive

Algorithm z -ranking**Input:** ε -consistent conditional knowledge base (L, D) .**Output:** The z -ranking of (L, D) .

1. **for each** $i \in \{0, 1, \dots, n(n-1)/2\}$ **do**;
 if D has an admissible partition (D_0, \dots, D_k) of weight i
 then $w[i] := \text{true}$ **else** $w[i] := \text{false}$;
2. $w := \min\{i \mid w[i] = \text{true}, 0 \leq i \leq n(n-1)/2\}$;
3. **for each** $d \in D$ and $j \in \{0, 1, \dots, n-1\}$ **do**
 if D has an admissible partition (D_0, \dots, D_k) of weight w s.t. $d \in D_j$
 then $Z(d) := j$.

Figure 4: Algorithm z -ranking

calls of an NP-oracle. This alternatively proves membership in $\Delta_2^P[O(\log n)] = P_{\parallel}^{\text{NP}}$.)

The proofs of $P_{\parallel}^{\text{NP}}$ -hardness and of b) are given in Appendix B. \square

Let us next consider the problem of computing the z -ranking of a conditional knowledge base. It appears that this problem has essentially the same complexity as z -entailment.

Theorem 5.7 *a) The problem of computing the default ranking z for a given conditional knowledge base KB is $\text{FP}_{\parallel}^{\text{NP}}$ -complete. b) The problem of computing the default ranking z for a given literal-Horn conditional knowledge base KB is FP-hard.*

Proof. a) We now show membership in $\text{FP}_{\parallel}^{\text{NP}}$. Consider algorithm z -ranking. Its Steps 1 and 2 are identical to Steps 2 and 3, respectively, of algorithm z -entailment, and compute the weight w of the z -partition of D . In Step 3 it is thus checked whether (D_0, \dots, D_k) is the z -partition, and thus $Z(d)$ is assigned the correct value. Notice that for each default $d \in D$, the query is true for exactly one value of i .

Let $n = |D|$. Steps 1 and 3, respectively, can be done by parallel NP-oracle calls. These two rounds of $n(n-1)/2 + 1$ and n^2 , respectively, parallel NP-oracle calls can be replaced by a single round of $n^2(n(n-1)/2 + 1) = n^4/2 - n^3/2 + n^2$ parallel NP-oracle calls as follows. For each triple (i, d, j) where $i \in \{0, 1, \dots, n(n-1)/2\}$, $d \in D$, and $j \in \{0, 1, \dots, n-1\}$, a query asks whether D has an ordered partition (D_0, \dots, D_k) that is admissible with KB , has weight i , and such that $d \in D_j$. It is now easy to see that $Z(d) = j$ iff the query for (i, d, j) is a query with smallest value for i that is answered “Yes”. It follows that computing the z -ranking is in $\text{FP}_{\parallel}^{\text{NP}}$.

The proofs of $\text{FP}_{\parallel}^{\text{NP}}$ -hardness and of b) are given in Appendix B. \square

Algorithm z -entailment can be modified to an algorithm z -rank-entailment that uses the z -ranking of the input, which saves on oracle queries, and can also be rewritten such that it makes one round of parallel NP oracle calls, whose number depends on the input. Our next result shows that it is unlikely that we can improve on this and find an algorithm that solves entailment with a fixed number of NP oracle calls (the proof is given in Appendix B).

Theorem 5.8 *a) Given a conditional knowledge base $KB = (L, D)$, a default d , and the ranking z of D , deciding whether KB z -entails d is $P_{\parallel}^{\text{NP}}$ -hard. b) Given a literal-Horn conditional knowledge base $KB = (L, D)$, a literal-Horn default d , and the ranking z of D , deciding whether KB z -entails d is P-hard.*

Let us now turn to z^+ -entailment. As already recalled at the beginning of this section, it is already known that for z^+ rankings the problems ENTAILMENT, RANKING, and RANK-ENTAILMENT are feasible in polynomial time with an oracle for NP. Thus, to verify the entries in Tables 3-5, it remains to show matching hardness results for ENTAILMENT and RANKING in the general case.

Theorem 5.9 *Given a conditional knowledge base KB with strength assignment σ and a default d , deciding whether (KB, σ) z^+ -entails d at given strength $\tau \geq 0$ is P^{NP} -hard. This remains true if $\tau = 0$ is fixed.*

Proof. We give a polynomial transformation from the following P^{NP} -complete problem [58]. Given a conjunction $\phi = \phi_1 \wedge \dots \wedge \phi_m$ of clauses ϕ_i on atoms x_1, \dots, x_n , which is asserted to be satisfiable, decide whether $I_{\text{lms}}(\phi) \models x_n$, where $I_{\text{lms}}(\phi)$ is the lexicographically maximum satisfying truth assignment of ϕ .

Let $k \geq 0$ be an integer such that $2^k \geq n$. For $i = 1, \dots, n$, let the defaults d_i^+ and d_i^- and their associated strengths $\sigma(d_i^+)$ and $\sigma(d_i^-)$, respectively, be defined as follows:

$$d_i^+ = \neg c_1 \wedge \dots \wedge \neg c_{i-1} \wedge \phi \rightarrow \neg x_i \vee c_i \quad \text{with strength } \sigma(d_i^+) = 0$$

$$d_i^- = \neg c_1 \wedge \dots \wedge \neg c_{i-1} \wedge \phi \rightarrow x_i \vee c_i \quad \text{with strength } \sigma(d_i^-) = 2^{k+n-i}.$$

Let $D = \{d_1^+, d_1^-, d_2^+, d_2^-, \dots, d_n^+, d_n^-\}$. Roughly speaking, we interpret each satisfying assignment for ϕ as the binary representation of a nonnegative integer, and the lexicographic order on satisfying assignments for ϕ as the usual order “ $<$ ” on the corresponding nonnegative integers. The defaults in D serve to add up the contribution of each bit (that is, each atom among x_1, \dots, x_n) in a satisfying assignment for ϕ . The default d_i^+ (resp., d_i^-) assesses a cost of 0 (resp., 2^{k+n-i}) for the truth value **true** (resp., **false**) of x_i in a satisfying assignment for ϕ . More precisely, the rank of every $d_1 \in \{d_1^+, d_1^-\}$ is given by $z^+(d_1) = \sigma(d_1)$, while the rank of every $d_i \in \{d_i^+, d_i^-\}$ with $i \in \{2, \dots, n\}$ is given by $z^+(d_i) = \sigma(d_i) + z^+(d_{i-1}) + 1$, where d_{i-1} is the unique default among d_{i-1}^+ and d_{i-1}^- that is falsified by $I_{\text{lms}}(\phi)$ (see Appendix B). That is, the ranks of d_i^+ and d_i^- with $i \in \{1, \dots, n\}$ reflect the restriction of $I_{\text{lms}}(\phi)$ to $\{x_1, \dots, x_i\}$. The factor 2^k in $\sigma(d_i^-)$ serves to account for the 1’s that are added at each ranking step according to equation (3) (in total, at most $n - 1$). Since the strengths of the defaults d_1^-, \dots, d_n^- decrease in a mathematical progression, minimization of the violation cost has the effect that, if possible, the x_i ’s are set to **true** rather than to **false**, for $i = 1, \dots, n$.

It can now be shown that x_n is true in $I_{\text{lms}}(\phi)$ iff $KB = (\emptyset, D)$ z^+ -entails the default $d = \neg c_1 \wedge \dots \wedge \neg c_n \wedge \phi \rightarrow x_n$ at strength $\tau = 0$ (see Appendix B). \square

Theorem 5.10 *The problem of computing the default ranking z^+ for a conditional knowledge base KB with strength assignment σ is FP^{NP} -hard.*

Proof. We give a polynomial transformation from the following FP^{NP} -complete problem [58]. Given a conjunction $\phi = \phi_1 \wedge \dots \wedge \phi_m$ of clauses on atoms x_1, \dots, x_n , which is asserted to be satisfiable, compute $I_{\text{lms}}(\phi)$, that is, the lexicographically maximum satisfying truth assignment of ϕ .

We slightly extend the construction in the proof of Theorem 5.9 as follows. We add to D there the default $d = \neg c_1 \wedge \dots \wedge \neg c_n \wedge \phi \rightarrow \top$ with $\sigma(d) = 0$. By a similar line of argumentation, it follows that (15) holds for all $i = 1, \dots, n$ and that $z^+(d) = a + n$ (recall that a is the integer having the complement of $I_{\text{lms}}(\phi)$ as binary representation). Since $I_{\text{lms}}(\phi)$ is easily computed from $z^+(d)$, the result follows. \square

5.3 Maximum Entropy Semantics

We now focus on the complexity of deciding z^* - and z_s^* -entailment. We first determine the precise complexity of z^* -entailment, and start with the lower complexity bound.

Algorithm z -rank-entailment**Input:** ε -consistent conditional knowledge base (L, D) , default $\phi \rightarrow \psi$, and z -ranking of D .**Output:** “Yes”, if (L, D) z -entails $\phi \rightarrow \psi$, otherwise “No”.

1. **if** $L \cup \{\phi\}$ is unsatisfiable **then return** “Yes”;
2. $\alpha_1 := \phi \wedge \psi$; $\alpha_2 := \phi \wedge \neg\psi$;
3. **for each** $i \in \{0, 1, \dots, n\}$ and $j \in \{1, 2\}$ **do**
 if $L \cup \{\alpha_j\}$ has a model I s.t. $i = 1 + \max(\{Z(d) \mid d \in D, I \not\models d\} \cup \{-1\})$
 then $k[i, j] := \text{true}$ **else** $k[i, j] := \text{false}$;
4. **for each** $i \in \{1, 2\}$ **do** $k_j := \min\{j \mid k[i, j] = \text{true}\}$;
5. **if** $k_1 < k_2$ **then return** “Yes” **else return** “No”.

Figure 5: Algorithm z -rank-entailment

Theorem 5.11 *Given a literal-Horn conditional KB, which is ε -consistent and minimal-core, and a literal-Horn default d , deciding whether KB z^* -entails d is P^{NP} -hard.*

Proof. We give a polynomial transformation from the following P^{NP} -complete problem. Given a set of weighted Horn clauses $C = \{\alpha_1 \Rightarrow \beta_1, \dots, \alpha_m \Rightarrow \beta_m\}$ on atoms x_1, \dots, x_n such that every $\alpha_i \Rightarrow \beta_i$ is satisfiable and has weight $w_i = 2^{c_i}$, where $c_i \geq 0$ is a nonnegative integer, and some $r \in \{1, \dots, m\}$, decide whether $I \models \alpha_r \Rightarrow \beta_r$ for every maximum weight world I under C , that is, world I such that $\sum_{I \models \alpha_i \Rightarrow \beta_i} w_i$ is maximum over all worlds in \mathcal{I}_{At} .

P^{NP} -hardness of this problem follows by a minor adaptation of the proof in [58] that computing the maximum weight assignment under a set C of weighted arbitrary clauses is FP^{NP} -complete; the proof in [58] implies that deciding whether $I \models \alpha_r \Rightarrow \beta_r$ holds for a particular clause $\alpha_r \Rightarrow \beta_r$ in every maximum weight assignment I under C is P^{NP} -complete.

We now construct a conditional knowledge base KB and a default d as requested such that $I \models \alpha_r \Rightarrow \beta_r$ holds for every maximum weight world I under C iff KB z^* -entails d .

The main idea behind this construction can be informally described as follows. For each Horn clause $\alpha_j \Rightarrow \beta_j$ with weight 2^{c_j} , we will introduce a default $d_{c_j, j}$. By additional Horn clauses and literal-Horn defaults in KB , we then ensure that $z^*(d_{c_j, j}) = 2^{c_j}$ and that $d_{c_j, j}$ is falsified by a model I of some conjunction of atoms ϕ iff I does not satisfy $\alpha_j \Rightarrow \beta_j$. Moreover, we will ensure that all the other defaults in KB are falsified by every model of ϕ . It will then follow that the minimal falsifying models of ϕ are exactly the maximum weight assignments under C . Hence, it finally just remains to choose an appropriate literal ψ , such that the default $\phi \rightarrow \psi$ is z^* -entailed by KB iff $\alpha_r \Rightarrow \beta_r$ holds in all minimal falsifying models of ϕ .

Let the set of atomic propositions be $At = \{x_1, \dots, x_n\} \cup A \cup B \cup T$, where $A = \{a_{i, j} \mid 1 \leq j \leq m, 0 \leq i \leq c_j\}$, $B = \{b_{i, j} \mid 1 \leq j \leq m, 0 \leq i \leq c_j\}$, and $T = \{t_{i, j} \mid 1 \leq j \leq m, 0 \leq i \leq c_j - 1\}$.

Then, $KB = (L, D)$ is defined as follows:

$$\begin{aligned} L &= \bigcup \{L_{i, j} \mid 1 \leq j \leq m, 0 \leq i \leq c_j\} \cup \{\phi_{i, j} \mid 1 \leq j \leq m, 0 \leq i \leq c_j\} \\ D &= \{d_{i, j} \mid 1 \leq j \leq m, 0 \leq i \leq c_j\}, \end{aligned}$$

where $L_{i,j}$, $\phi_{i,j}$, and $d_{i,j}$ are defined as follows:

$$\begin{aligned} L_{i,j} &= \{b_{i,j} \Rightarrow a_{0,j}, b_{i,j} \Rightarrow t_{0,j}, \dots, b_{i,j} \Rightarrow a_{i-1,j}, b_{i,j} \Rightarrow t_{i-1,j}\} \\ \phi_{i,j} &= \begin{cases} b_{i,j} \wedge t_{i,j} \Rightarrow \perp & \text{if } i < c_j \\ b_{i,j} \wedge \alpha_j \Rightarrow \beta_j & \text{if } i = c_j \end{cases} \\ d_{i,j} &= a_{i,j} \rightarrow b_{i,j}. \end{aligned}$$

Finally, let the literal-Horn default $d = \phi \rightarrow \psi$ be defined as follows:

$$\begin{aligned} \phi &= \bigwedge_{p \in AUT} p \\ \psi &= b_{c_k, k}. \end{aligned}$$

It can now be shown (see Appendix B) that (i) KB is ε -consistent, (ii) KB is minimal-core, (iii) the default ranking z^* is given by $z^*(d_{i,j}) = 2^i$ for all $d_{i,j} \in D$ (thus, $z^*(d_{c_j,j}) = 2^{c_j}$ for $j \leq m$), and (iv) $I \models \alpha_r \Rightarrow \beta_r$ holds in every I with maximum weight $\sum_{I \models \alpha_i \Rightarrow \beta_i} w_i$ iff KB z^* -entails d . \square

The next result follows from a slight extension of the construction in the proof of Theorem 5.11.

Theorem 5.12 *The problem of computing the default ranking z^* for an ε -consistent minimal-core literal-Horn conditional knowledge base is FP^{NP} -hard.*

Theorem 5.13 *Given a literal-Horn conditional knowledge base KB , which is ε -consistent and minimal-core, the ranking z^* of KB , and a literal-Horn default d , deciding whether KB z^* -entails d is P^{NP} -hard.*

Proof. Immediate by the proof of Theorem 5.11: The default ranking z^* for KB there is $Z(d_{i,j}) = 2^i$ for all $j \leq m$ and $i \leq c_j$. Hence, the problem there is easily reduced to the case where z^* is part of the input. \square

We next consider the extension of the maximum entropy approach z^* by variable strength defaults to z_s^* . Our goal is to show that z_s^* -entailment is in P^{NP} , and thus no harder than z^* -entailment. For this purpose, we use the algorithm z_s^* -**ranking** shown in Figure 6, which computes the default ranking z_s^* of an ε -consistent conditional knowledge base KB . If, during the computation, it is detected that KB is not robust, a special value *nil* is returned. This algorithm is essentially a reformulation of a similar algorithm in [13], and thus we do not analyze its correctness here (note that Step 4 in z_s^* -**ranking** is a technical trick to ensure that $\text{minf}(d)$ in Steps 5 and 6 has the rank of the current minimal falsifying model of d *excluding its own contribution*).

The next lemma shows that the rank $z_s^*(d)$ of each default d computed by this algorithm has an exponential upper bound, and thus can be represented by a polynomial number of bits.

Lemma 5.14 *Let $KB = (L, D)$ be an ε -consistent conditional knowledge base with positive strength assignment σ . Let $n = |D|$ be the cardinality of D and let $s = \max(\{\sigma(d) \mid d \in D\})$. Let $z_s^*(d_1), \dots, z_s^*(d_l)$, with $l \leq n$, denote the sequence of default ranks computed in algorithm z_s^* -ranking. Then $|z_s^*(d_i)| \leq s \cdot 2^{i-1}$ for all $i = 1, \dots, l$.*

By showing that algorithm z_s^* -ranking can be implemented to run in deterministic polynomial time if an NP-oracle is available, we prove the following result.

Theorem 5.15 *Given an ε -consistent conditional knowledge base $KB = (L, D)$ with positive strength assignment σ , the problem of computing the default ranking z_s^* for KB , if KB is robust, and returning *nil* otherwise, is in FP^{NP} .*

Algorithm z_s^* -ranking (essentially [13])**Input:** ε -consistent $KB = (L, D)$ with positive strength assignment σ .**Output:** Default ranking z_s^* , if the system of equations given by (6) and (7) for all $\phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$ has a unique solution z_s^* , κ_s^* such that z_s^* is positive and κ_s^* is robust; otherwise, *nil*.Notation: We use $\text{minv}(\phi \rightarrow \psi)$ and $\text{minf}(\phi \rightarrow \psi)$ to denote $\kappa_s^*(\phi \wedge \psi)$ and $\kappa_s^*(\phi \wedge \neg\psi)$, respectively, where $\kappa_s^*(\alpha) = \min_{I \in \mathcal{I}_{At}: I \models \alpha} \kappa_s^*(I)$ with $\kappa_s^*(I)$ as in (7).

1. **for each** $d \in D$ **do** $z_s^*(d) := \infty$;
2. **while** $\{d \in D \mid z_s^*(d) = \infty\} \neq \emptyset$ **do begin**
3. Take any $d \in D$ with $z_s^*(d) = \infty$ such that $\sigma(d) + \text{minv}(d)$ is minimal;
4. $z_s^*(d) := 0$;
5. **if** $\text{minf}(d) = \infty$ **then return nil**;
6. $z_s^*(d) := \sigma(d) + \text{minv}(d) - \text{minf}(d)$
7. **end**;
8. **if** κ_s^* satisfies (6) for all $\phi \rightarrow \psi \in D$ **and** z_s^* is positive and κ_s^* is robust
9. **then return** z_s^* **else return nil**.

Figure 6: Algorithm z_s^* -ranking**Proof.** Let $n = |D|$, let $s = \max(\{\sigma(d) \mid d \in D\})$, and let $|s|$ be the length of the binary representation of s .

It is sufficient to show the following properties of algorithm z_s^* -ranking: (i) computing $\text{minv}(d)$ and $\text{minf}(d)$ in Steps 3, 5, and 6 is in FP^{NP} ; (ii) deciding in Step 8 whether κ_s^* satisfies (6) for each $\phi \rightarrow \psi \in D$ is in P^{NP} ; and (iii) deciding whether z_s^* is positive and κ_s^* is robust in Step 8 is in P^{NP} . Lemma 5.14 implies that for every $\phi \rightarrow \psi \in D$, the values $\kappa_s^*(\phi \wedge \psi)$ and $\kappa_s^*(\phi \wedge \neg\psi)$ are from $\{0, \dots, s \cdot 2^n\} \cup \{\infty\}$. Thus, they can be computed by binary search in deterministic polynomial time with $O(|s| + n)$ call to an NP-oracle, since deciding whether some $I \in \mathcal{I}_{At}$ exists such that $\kappa_s^*(I) \leq w$ for a given value w is in NP. This already shows that (i) and (ii) are in FP^{NP} and P^{NP} , respectively. As for (iii), we check whether there are no two distinct defaults from D that have a common minimal falsifying model under L . If for every $d \in D$ the value $\text{minf}(d)$, which can be computed in FP^{NP} , is known, this is clearly in co-NP, and thus can be checked with a call to an NP-oracle. This proves the result. \square

The announced upper bound on z_s^* -entailment is now easily established.**Theorem 5.16** *Given an ε -consistent conditional knowledge base KB , asserted to be robust, with positive strength assignment σ , a default d , and an integer $\tau \geq 0$, deciding whether (KB, σ) z_s^* -entails d at strength τ , is in P^{NP} .***Proof.** Let $d = \phi \rightarrow \psi$. By Theorem 5.15, computing the z_s^* -ranking of KB is in FP^{NP} . Deciding whether $\kappa_s^*(\phi) = \infty$ is in NP, and by similar arguments as in the proof of Theorem 5.15, computing the ranks $\kappa_s^*(\phi \wedge \psi)$ and $\kappa_s^*(\phi \wedge \neg\psi)$ in a binary search using z_s^* is then in FP^{NP} ; testing $\kappa_s^*(\phi \wedge \psi) < \kappa_s^*(\phi \wedge \neg\psi)$ is simple. Overall, deciding whether (KB, σ) z_s^* -entails d is in P^{NP} . \square The matching hardness result is inherited from the P^{NP} -hardness of z^* -entailment in Theorem 5.11 and the fact that the minimal-core property implies robustness (Lemma 3.2).**Theorem 5.17** *Given an ε -consistent literal-Horn conditional knowledge base KB with positive strength assignment σ , where KB is asserted to be robust, a literal-Horn default d , and an integer $\tau \geq 0$, deciding*

whether $(KB, \sigma) z_s^*$ -entails d at strength τ , is P^{NP} -hard. This remains true if the ranking z_s^* is part of the input and $\tau = 0$ is fixed.

5.4 Lexicographic Entailment

We now concentrate on the complexity of lexicographic entailment. We remark that the upper complexity bound stated in Table 3 for lex_p -entailment follows immediately from Cayrol et al.'s work [22], which shows that lex_p -entailment is in P^{NP} for default knowledge bases (\emptyset, D) . This result can be easily extended to conditional knowledge bases (L, D) , by turning the formulas ϕ in the background knowledge L into defaults $d_\phi = \top \rightarrow \phi$, and assigning them priority $\pi(d_\phi) = \max_{d \in D} \pi(d) + 1$, provided L is consistent. This proves the result in Table 3.

We next prove that deciding lex_p -entailment is P^{NP} -hard for the Horn case. More precisely, we even show the stronger result that P^{NP} -hardness holds for the literal-Horn case.

Theorem 5.18 *Given a literal-Horn default knowledge base $KB = (\emptyset, D)$ with priority assignment π , and a literal-Horn default $d = \top \rightarrow \psi$, deciding whether (KB, π) lex_p -entails d is P^{NP} -hard.*

Proof. We give a polynomial transformation from the following P^{NP} -complete problem [58] (cf. Theorem 5.9). Given a conjunction $\alpha = \alpha_1 \wedge \dots \wedge \alpha_m$ of clauses α_i on atoms x_1, \dots, x_n , which is asserted to be satisfiable, decide whether $I_{lms}(\alpha) \models x_n$ where $I_{lms}(\alpha)$ is the lexicographically maximum model of α .

We construct $KB = (\emptyset, D)$, π , and $d = \top \rightarrow \psi$ as in the problem statement such that $I_{lms}(\alpha) \models x_n$ iff (KB, π) lex_p -entails d (see Appendix B). Roughly speaking, the main problem in this construction is to transform the clauses $\alpha_1, \dots, \alpha_m$ into Horn clauses $\alpha_1^*, \dots, \alpha_m^*$. We tackle this by replacing all positive literals x_i in $\alpha_1, \dots, \alpha_m$ by new negative literals $\neg y_i$, and express the relationships $x_i \wedge y_i \Rightarrow \perp$ and $\top \Rightarrow x_i \vee y_i$ in a suitable way. This is accomplished by introducing appropriate sets of defaults F_i and by exploiting the lexicographic preference ordering on worlds. \square

We next turn to the complexity of lex -entailment. Since this is a special case of lex_p -entailment, we can solve this problem by using an algorithm for lex_p -entailment, to which provide the z -partition of D from the conditional knowledge base $KB = (L, D)$ in the input. As this can be done without too much overhead, the complexity of lex_p -entailment gives us then an upper bound for lex -entailment. On the other hand, the transformation of general lexicographic default ranking to entrenched lexicographic default ranking from Theorem 3.3 allows us to infer that this complexity is also a lower bound.

Theorem 5.19 *a) Given an ε -consistent conditional knowledge base KB and a default d , deciding whether KB lex -entails d is in P^{NP} . b) Given an ε -consistent literal-Horn default knowledge base KB and a literal-Horn default d , deciding whether KB lex -entails d is P^{NP} -hard.*

5.5 Conditional Entailment

We finally concentrate on the complexity of Geffner's conditional entailment.

The following lemma is helpful for checking whether a priority ordering is admissible.

Lemma 5.20 *Let $KB = (L, D)$ be a conditional knowledge base. A priority ordering \prec on D is admissible with KB iff each $d \in D$ is tolerated by $D_d = D - \{d' \in D \mid d' \prec d\}$ under L .*

We are now ready to give an upper bound for the complexity of conditional entailment.

Theorem 5.21 *Given a conditional knowledge base $KB = (L, D)$ and a default $d = \phi \rightarrow \psi$, deciding whether KB conditionally entails d is in Π_2^P .*

Algorithm not-cond-entailment**Input:** Conditional knowledge base (L, D) and a default $\phi \rightarrow \psi$.**Output:** “Yes” iff (L, D) does not conditionally entail $\phi \rightarrow \psi$.

1. Guess a subset \prec of $D \times D$;
2. **if** \prec is not irreflexive **or** \prec is not transitive **then halt**; /* \prec is not a priority ordering */
3. **for each** $d \in D$ **do**
 - if** \neg (d is tolerated by $D - \{d' \in D \mid d' \prec d\}$ under L)
 - then halt**; /* \prec is not admissible */
4. Guess $I \in \mathcal{I}_{At}$;
5. **if** $I \not\models L \cup \{\phi\}$ **or** $I \models \psi$ **then halt**; /* wrong guess */
6. **if** some $J \in \mathcal{I}_{At}$ exists s.t. $J \prec I$ **and** $J \models L \cup \{\phi\}$ **then halt** /* I is not preferred */
else return “Yes”.

Figure 7: Algorithm **not-cond-entailment**

Proof. We show that the complement problem, that is, deciding whether KB does *not* conditionally entail d , is in Σ_2^P . Recall that d is not conditionally entailed by KB iff there exists a priority ordering \prec on D that is admissible with KB , and a \prec -preferred model I of $L \cup \{\phi\}$ such that $I \not\models \psi$. This is checked by the nondeterministic algorithm **not-cond-entailment**. In Step 3 there, Lemma 5.20 is applied for checking that the priority ordering \prec is admissible. It is easily seen that the (unnegated) queries in Step 3 and the query in Step 6 can be answered by an NP-oracle. Modulo these queries, each of the Steps 1–6 can be done in polynomial time. Hence, deciding whether KB does not conditionally entail d is in $\text{NP}^{\text{NP}} = \Sigma_2^P$. \square

We next show that this upper bound is tight, and that this holds even for the literal-Horn case.

Theorem 5.22 *Given a literal-Horn conditional knowledge base $KB = (L, D)$ and a literal-Horn default $d = \phi \rightarrow \psi$, deciding whether KB conditionally entails d is Π_2^P -hard.*

Proof. We give a polynomial transformation from the following canonical Π_2^P -complete problem [54]. Given a collection of clauses $\alpha_1, \dots, \alpha_l$ on the atoms $y_1, \dots, y_m, x_1, \dots, x_n$, where $m, n \geq 1$, decide whether the quantified Boolean formula $\Phi = \forall y_1 \dots \forall y_m \exists x_1 \dots \exists x_n (\alpha_1 \wedge \dots \wedge \alpha_l)$ evaluates to true. That is, is it true that for each truth assignment I_y to the variables y_1, \dots, y_m , there exists a truth assignment I_x to the variables x_1, \dots, x_n such that $I_y \cup I_x$ satisfies $\alpha_1 \wedge \dots \wedge \alpha_l$. Without loss of generality, we can assume that each clause α_i contains at least one variable from x_1, \dots, x_n .

We construct a literal-Horn conditional knowledge base $KB = (L, D)$ and a literal-Horn default $d = \phi \rightarrow \psi$ such that Φ evaluates to true iff KB conditionally entails d (see Appendix B). Roughly speaking, this construction involves two main problems. First, we must express somehow the validity of a quantified Boolean formula. This will be done by making use of the preference ordering on worlds. Second, we must transform the clauses $\alpha_1, \dots, \alpha_l$ into Horn clauses $\alpha_1^*, \dots, \alpha_l^*$. This will be done by replacing all positive literals y_i and x_j in $\alpha_1, \dots, \alpha_l$ by new negative literals $\neg y'_i$ and $\neg x'_j$, respectively. We will then use L to express the relationships $y_i \wedge y'_i \Rightarrow \perp$ and $x_j \wedge x'_j \Rightarrow \perp$, and the preference ordering on worlds to express the relationships $\top \Rightarrow y_i \vee y'_i$ and $\top \Rightarrow x_j \vee x'_j$; thus, $y_i \equiv \neg y'_i$ and $x_j \equiv \neg x'_j$. \square

6 New Tractable Cases

6.1 Overview

As the results in Section 4 show, a number of the more sophisticated semantics for conditional knowledge bases are intractable for all the classes of default reasoning problems that we have considered. Thus, the issue of meaningful, tractable classes of problems for these semantics naturally arises. On the other hand, it would be interesting to know whether the tractability results for the other semantics in the Horn case can be extended to more expressive classes of problems.

In this section, we tackle these issues and present new tractable cases for default reasoning from conditional knowledge bases. In response to the latter question, we introduce in Section 6.2 the class of *q-Horn* conditional knowledge bases. This class generalizes Horn conditional knowledge bases syntactically by allowing a restricted use of disjunction, and contains instances which cannot be represented in Horn conditional knowledge bases. As we show, the tractability results for Horn conditional knowledge bases in Tables 2–5 extend to q-Horn conditional knowledge bases, which can be regarded as a positive result regarding the tractability of default reasoning.

Finding meaningful tractable cases for the more sophisticated semantics for conditional knowledge bases turns out to be more challenging. A natural attempt is to show that a further restriction of the literal-Horn case leads to tractability. An obvious candidate restriction is bounding the size of the antecedents in the strict and classical rules to at most one atom. Unfortunately, as we show in Section 6.3, this does not buy tractability. An analysis of our proof reveals that the interaction of the defaults among each other and with the classical background knowledge must be controlled such that interferences have a local effect. This leads us to the class of *feedback-free Horn* (*ff-Horn*) default reasoning problems in Section 6.4. As we show, tractability is gained on this class for most of the intractable semantics in Tables 2–5. The refined hierarchy of default reasoning problem classes is shown in Figure 8, where the new classes introduced in this section are emphasized in bold face.

6.2 Q-Horn

6.2.1 Motivating example

Q-Horn conditional knowledge bases generalize Horn conditional knowledge bases by allowing a limited form of disjunction. The following example illustrates this kind of disjunctive knowledge.

Example 6.1 Assume that John is looking for Mary. Unfortunately, he did not find her at home. So, he is wondering where she might be. He knows that Mary might be having tea with her friends, that she might be in the library, or that she might be playing tennis. He also knows that these scenarios are pairwise exclusive and not exhaustive. Moreover, John knows that “generally, in the afternoon, Mary is having tea with her friends or she is in the library” and that “generally, on Friday afternoon, Mary is playing tennis”.

This knowledge can be expressed by the following conditional knowledge base $KB = (L, D)$:

$$\begin{aligned} L &= \{\neg tea \vee \neg library, \neg tea \vee \neg tennis, \neg library \vee \neg tennis\}, \\ D &= \{afternoon \rightarrow tea \vee library, Friday \wedge afternoon \rightarrow tennis\}. \end{aligned}$$

Suppose now that it is Friday afternoon and that John is wondering whether he should go to the library to look for Mary. That is, does KB entail $Friday \wedge afternoon \rightarrow library$?

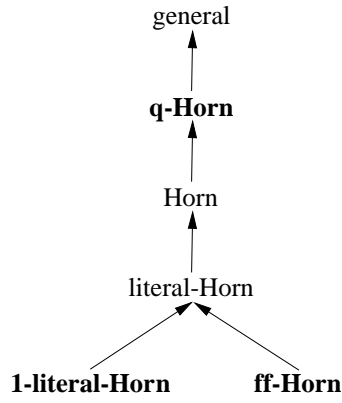


Figure 8: Refined hierarchy of syntactic restrictions

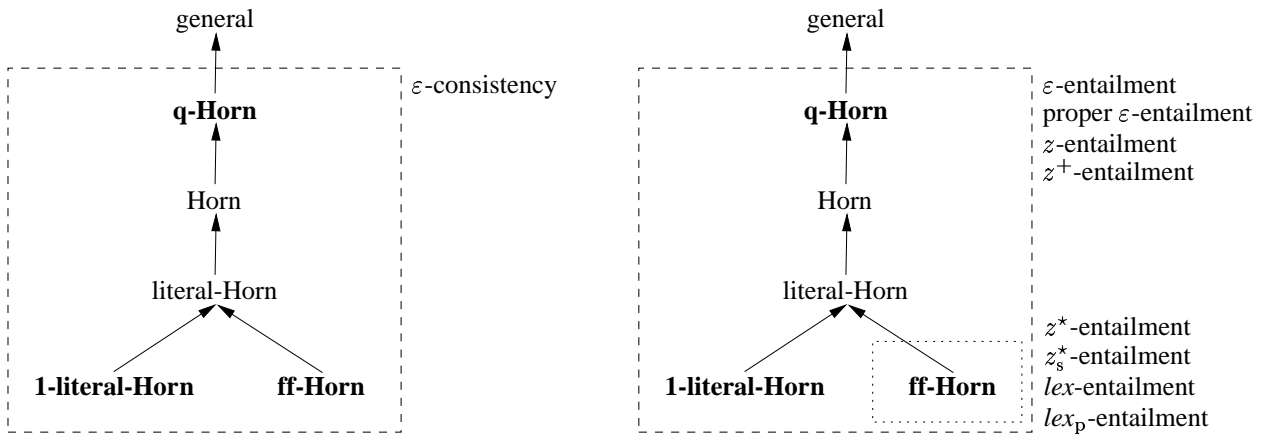


Figure 9: Tractability of CONSISTENCY and ENTAILMENT

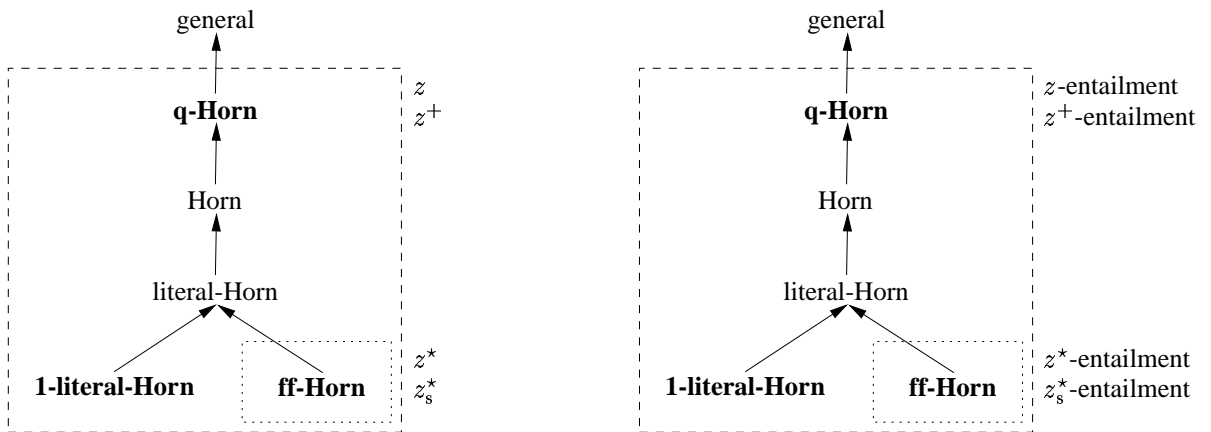


Figure 10: Tractability of RANKING and RANK-ENTAILMENT

6.2.2 Definitions

We now introduce q-Horn conditional knowledge bases. A *clause* is a disjunction of literals. A default $\phi \rightarrow \psi$ is *clausal* iff ϕ is either \top or a conjunction of literals, and ψ is a conjunction of clauses. A conditional knowledge base $KB = (L, D)$ is *clausal* iff L is a finite set of clauses and D is a finite set of clausal defaults. A default reasoning problem (KB, d) is *clausal* iff both KB and d are clausal.

A classical formula ϕ is in *conjunctive normal form* (or *CNF*) iff ϕ is either \top or a conjunction of clauses. We use the operator \sim to map each atom a to its negation $\neg a$, and each negated atom $\neg a$ to a . We define a mapping \mathcal{N} that associates each clausal default d with a classical formula in CNF as follows. If d is of the form $\top \rightarrow c_1 \wedge \cdots \wedge c_n$ with clauses c_1, \dots, c_n , then $\mathcal{N}(d) = c_1 \wedge \cdots \wedge c_n$. If d is of the form $l_1 \wedge \cdots \wedge l_m \rightarrow c_1 \wedge \cdots \wedge c_n$ with literals l_1, \dots, l_m and clauses c_1, \dots, c_n , then $\mathcal{N}(d)$ is defined as the conjunction of all $\sim l_1 \vee \cdots \vee \sim l_m \vee c_i$ with $i \in \{1, \dots, n\}$. We extend the mapping \mathcal{N} to classical formulas in CNF ϕ by defining $\mathcal{N}(\phi) = \phi$. We extend \mathcal{N} to finite sets K of classical formulas in CNF and clausal defaults as follows. Let K' denote the set of all $k \in K$ with $\mathcal{N}(k) \neq \top$. If $K' \neq \emptyset$, then we define $\mathcal{N}(K)$ as the conjunction of all $\mathcal{N}(k)$ with $k \in K'$. Otherwise, $\mathcal{N}(K)$ is defined as \top .

A *partial assignment* S is a set of literals such that for every atom $a \in At$ at most one of the literals a and $\neg a$ is in S . A classical formula in CNF ϕ is *q-Horn* iff there exists a partial assignment S such that (i) each clause in ϕ contains at most two literals outside of S , and (ii) if a clause in ϕ contains exactly two literals $u, v \notin S$, then neither $\sim u$ nor $\sim v$ belongs to S . Note that it is *not* assumed that the partial assignment S can be completed to a model of ϕ .

The class of q-Horn formulas generalizes the classes of quadratic, Horn, and disguised Horn formulas [11]. Recall that a classical formula in CNF $\phi = c_1 \wedge \cdots \wedge c_k$ over the atoms a_1, \dots, a_n is *quadratic* iff each clause c_i contains at most two literals. It is *Horn* iff each c_i is a Horn clause, and *disguised Horn* iff there exists a partial assignment S such that $|S| = n$ and that each clause c_i contains at most one literal not belonging to S . Informally, disguised Horn ϕ can be made Horn by “renaming” atoms. For quadratic (resp., Horn) ϕ , the partial assignment $S = \emptyset$ (resp., $S = \{\neg a_1, \dots, \neg a_n\}$) always satisfies (i) and (ii).

Example 6.2 The classical formulas $\neg a \wedge (b \vee c) \wedge (a \vee c)$, $\neg a \wedge (b \vee \neg c \vee \neg d) \wedge (\neg a \vee \neg b \vee \neg c \vee d)$, and $\neg a \wedge b \wedge (b \vee c)$ are quadratic, Horn, and disguised Horn, respectively, and thus q-Horn.

The classical formula $(\neg a \vee \neg b) \wedge (\neg a \vee \neg c) \wedge (\neg b \vee \neg c) \wedge (a \vee b) \wedge (a \vee c) \wedge (b \vee c \vee d)$ is q-Horn (as $S = \{d\}$ satisfies (i) and (ii)), but neither among Horn, quadratic, and disguised Horn.

The classical formula $(\neg a \vee \neg b \vee \neg c) \wedge (a \vee b \vee c)$ is not q-Horn.

A finite set K of classical formulas in CNF and clausal defaults is *q-Horn* iff $\mathcal{N}(K)$ is q-Horn. A conditional knowledge base $KB = (L, D)$ is *q-Horn* iff KB is clausal and $L \cup D$ is q-Horn. Clearly, every Horn conditional knowledge base is q-Horn, but not vice versa. A default reasoning problem (KB, d) is *q-Horn*, if KB is q-Horn and d is a clausal default.

Example 6.3 The conditional knowledge base $KB = (L, D)$ shown in Example 6.1 is q-Horn. More precisely, the classical formula $\mathcal{N}(L \cup D)$ associated with KB is given as follows:

$$\begin{aligned} \mathcal{N}(L \cup D) = & (\neg tea \vee \neg library) \wedge (\neg tea \vee \neg tennis) \wedge (\neg library \vee \neg tennis) \wedge \\ & (\neg afternoon \vee tea \vee library) \wedge (\neg Friday \vee \neg afternoon \vee tennis). \end{aligned}$$

It is now easy to verify that $\{\neg Friday, \neg afternoon\}$ is a partial assignment that satisfies (i) and (ii) as described above. That is, $\mathcal{N}(L \cup D)$ is q-Horn. Since KB is clearly clausal, it follows that KB is q-Horn.

The *size* of a classical formula in CNF ϕ , denoted $\|\phi\|$, is defined as the number of occurrences of literals in ϕ . We use $|\phi|$ to denote the number of clauses in ϕ . The *size* of a clausal default $d = \phi \rightarrow \psi$,

denoted $\|d\|$, is defined as $\|\phi\| + \|\psi\|$. The *size* of a finite set of clauses L , denoted $\|L\|$, is defined as the size of $\mathcal{N}(L)$. The *size* of a clausal conditional knowledge base $KB = (L, D)$, denoted $\|KB\|$, is defined as the size of $\mathcal{N}(L \cup D)$. We use $|D|$ to denote the cardinality of D .

6.2.3 Q-Horn formulas

We now give some preparative results. We first recall that the problem of deciding whether a q-Horn formula is satisfiable and the problem of recognizing q-Horn formulas are both tractable and can in fact be solved in linear time. These results go back to Boros et al. [10, 11].

Theorem 6.1 (see [10, 11])

- a) Given a q-Horn formula ϕ , deciding whether ϕ is satisfiable can be done in time $O(\|\phi\|)$.
- b) Given a classical formula in CNF ϕ , deciding whether ϕ is q-Horn can be done in time $O(\|\phi\|)$.

By these results, it is clear that q-Horn conditional knowledge bases can be efficiently recognized.

Proposition 6.2 *Given a clausal conditional knowledge base $KB = (L, D)$, deciding whether KB is q-Horn can be done in time $O(\|KB\|)$.*

The following lemma states the (immediate) observation that the set of all q-Horn formulas is closed under conjunction with literals and under decomposition of conjunctions into their components.

- Lemma 6.3** a) *If ϕ is a q-Horn formula and ψ is a conjunction of literals, then $\phi \wedge \psi$ is q-Horn.*
b) *If $\phi \wedge \psi$ is a q-Horn formula, then ϕ is q-Horn.*

6.2.4 ε -Semantics

We now prove the tractability results shown in Figures 9–10 for ε -semantics in the q-Horn case. The following theorem shows that deciding whether a q-Horn conditional knowledge base is ε -consistent is tractable.

Theorem 6.4 *Given a q-Horn conditional knowledge base $KB = (L, D)$, deciding whether KB is ε -consistent can be done in time $O(|D|^2 \|KB\|)$.*

Proof. A conditional knowledge base $KB = (L, D)$ is ε -consistent iff the z -partition of D exists [49]. That is, iff there exists an ordered partition (D_0, \dots, D_k) of D such that, for $i = 0, \dots, k$, each D_i is the set of all defaults in $D - \bigcup\{D_j \mid 0 \leq j < i\}$ that are tolerated under L by $D - \bigcup\{D_j \mid 0 \leq j < i\}$. Thus, deciding whether KB is ε -consistent can be reduced to $O(|D|^2)$ satisfiability tests on sets of classical formulas and defaults of the form $L \cup D' \cup \{\alpha \wedge \beta\}$ with $D' \subseteq D$ and $\alpha \rightarrow \beta \in D'$. Clearly, as $\alpha \rightarrow \beta \in D'$, such sets $L \cup D' \cup \{\alpha \wedge \beta\}$ are logically equivalent to $L \cup D' \cup \{\alpha\}$.

Assume now that KB is q-Horn. Then, by Lemma 6.3, every $L \cup D' \cup \{\alpha\}$ is q-Horn. Hence, by Theorem 6.1 a), each satisfiability test can be done in time $O(\|KB\|)$. Thus, in summary, deciding whether KB is ε -consistent can be done in time $O(|D|^2 \|KB\|)$. \square

The next result shows that deciding ε -entailment is tractable in the q-Horn case.

Theorem 6.5 *Given a q-Horn default reasoning problem (KB, d) , where $KB = (L, D)$ and $d = \phi \rightarrow \psi$, deciding whether KB ε -entails d can be done in time $O((\|\phi\| + \|\psi\|) |D|^2 (\|KB\| + \|d\|))$.*

Proof. By Theorem 3.1, a conditional knowledge base $KB = (L, D)$ ε -entails a default $d = \phi \rightarrow \psi$ iff $(L, D \cup \{\phi \rightarrow \neg\psi\})$ is ε -inconsistent. Hence, by the proof of Theorem 6.4, deciding whether KB ε -entails d can be reduced to $O(|D|^2)$ satisfiability tests on sets of classical formulas and defaults of the form

$L \cup D'' \cup \{\gamma\}$ with $D'' \subseteq D \cup \{\phi \rightarrow \neg\psi\}$ and $\gamma \rightarrow \delta \in D''$. We can now distinguish the three cases (i) $\phi \rightarrow \neg\psi \notin D''$, (ii) $\phi \rightarrow \neg\psi \in D''$ and $\gamma \rightarrow \delta \neq \phi \rightarrow \neg\psi$, and (iii) $\phi \rightarrow \neg\psi \in D''$ and $\gamma \rightarrow \delta = \phi \rightarrow \neg\psi$. That is, the satisfiability tests are done on sets of classical formulas and defaults of the form $L \cup D' \cup F$, where $D' \subseteq D$ and $F \in \{\{\alpha\}, \{\phi \rightarrow \neg\psi, \alpha\}, \{\neg\psi, \phi\}\}$ with $\alpha \rightarrow \beta \in D'$.

Suppose now that KB is q-Horn and d is clausal. Hence, ϕ is either \top or of the form $\phi_1 \wedge \dots \wedge \phi_k$ with literals ϕ_1, \dots, ϕ_k . Moreover, ψ is of the form $\psi_1 \wedge \dots \wedge \psi_n$ with clauses ψ_1, \dots, ψ_n . Thus, each satisfiability test on $L \cup D' \cup \{\phi \rightarrow \neg\psi, \alpha\}$ can be reduced to $\|\phi\|$ satisfiability tests on $L \cup D' \cup \{\neg\phi_i, \alpha\}$, $1 \leq i \leq k$, and $|\psi|$ satisfiability tests on $L \cup D' \cup \{\neg\psi_i, \alpha\}$, $1 \leq i \leq n$. Moreover, each satisfiability test on $L \cup D' \cup \{\neg\psi, \phi\}$ can be reduced to $|\psi|$ satisfiability tests on $L \cup D' \cup \{\neg\psi_i, \phi\}$, $1 \leq i \leq n$.

By Lemma 6.3, each such $L \cup D' \cup F$ with $F \in \{\{\alpha\}, \{\neg\phi_i, \alpha\}, \{\neg\psi_i, \alpha\}, \{\neg\psi_i, \phi\}\}$ is q-Horn. Hence, by Theorem 6.1 a), each satisfiability test can be done in time $O(\|KB\| + \|d\|)$. In summary, this shows that deciding whether KB ε -entails d can be done in time $O((\|\phi\| + |\psi|) |D|^2 (\|KB\| + \|d\|))$. \square

The following theorem shows that also deciding proper ε -entailment is tractable in the q-Horn case.

Theorem 6.6 *Given a q-Horn default reasoning problem (KB, d) , where $KB = (L, D)$ and $d = \phi \rightarrow \psi$, deciding whether KB properly ε -entails d can be done in time $O((\|\phi\| + |\psi|) |D|^2 (\|KB\| + \|d\|))$.*

Proof. Recall that KB properly ε -entails $\phi \rightarrow \psi$ iff KB ε -entails $\phi \rightarrow \psi$ and KB does not ε -entail $\phi \rightarrow \perp$ (that is, $\phi \rightarrow a \wedge \neg a$ for some $a \in At$). Hence, the result is immediate by Theorem 6.5. \square

6.2.5 Systems Z and Z^+

We next concentrate on entailment in systems Z and Z^+ . The following result shows that computing the default ranking z^+ is tractable in the q-Horn case. Since system Z^+ is a proper generalization of system Z , this result shows also that computing the default ranking z is tractable in the q-Horn case.

Theorem 6.7 *Given an ε -consistent q-Horn conditional knowledge base $KB = (L, D)$ with strength assignment σ , the default ranking z^+ can be computed in time polynomial in the input size.*

Proof. Given an ε -consistent conditional knowledge base $KB = (L, D)$ with strength assignment σ , the default ranking z^+ can be computed with $O(|D|^2 \log |D|)$ satisfiability tests on sets of classical formulas and defaults of the form $L \cup D' \cup \{\alpha\}$, where $D' \subseteq D$ and $\alpha \rightarrow \beta \in D'$ [49].

Assume now that KB is q-Horn. Then, by Lemma 6.3, each such $L \cup D' \cup \{\alpha\}$ is q-Horn. Hence, by Theorem 6.1 a), each satisfiability test can be done in time $O(\|KB\|)$. Thus, in summary, the default ranking z^+ can be computed in polynomial time. \square

We finally show that deciding z^+ -entailment is tractable in the q-Horn case. Again, since system Z^+ properly generalizes system Z , this result shows also that deciding z -entailment is tractable in the q-Horn case. Trivially, these tractability results remain true when z^+ and z , respectively, are part of the input.

Theorem 6.8 *Given a q-Horn default reasoning problem (KB, d) , where $KB = (L, D)$ is ε -consistent and has a strength assignment σ , deciding whether (KB, σ) z^+ -entails $d = \phi \rightarrow \psi$ at a given strength $\tau \geq 0$ can be done in time polynomial in the input size.*

Proof. Recall that (KB, σ) z^+ -entails d at strength τ iff $L \cup \{\phi\}$ is unsatisfiable or $\kappa^+(\phi \wedge \psi) + \tau < \kappa^+(\phi \wedge \neg\psi)$ (or, equivalently, $\kappa^+(\phi) + \tau < \kappa^+(\phi \wedge \neg\psi)$). Thus, we first have to check whether $L \cup \{\phi\}$ is unsatisfiable. If this is the case, then (KB, σ) z^+ -entails d at strength τ . Otherwise, we additionally have to decide whether $\kappa^+(\phi) + \tau < \kappa^+(\phi \wedge \neg\psi)$.

Suppose now that KB is q-Horn and d is clausal. Then, by Lemma 6.3, $L \cup \{\phi\}$ is q-Horn. Hence, by Theorem 6.1 a), deciding whether $L \cup \{\phi\}$ is unsatisfiable can be done in time $O(\|L\| + \|\phi\|)$.

By Theorem 6.7, the default ranking z^+ can be computed in polynomial time. Given the ranking z^+ , the values $\kappa^+(\phi)$ and $\kappa^+(\phi \wedge \neg\psi)$ can be computed with $O(\log |D|)$ satisfiability tests on sets of classical formulas and defaults of the form $L \cup D' \cup F$, where $F \in \{\{\phi\}, \{\phi, \neg\psi\}\}$ and $D' \subseteq D$ [49]. Since ψ is of the form $\psi_1 \wedge \dots \wedge \psi_n$ with clauses ψ_1, \dots, ψ_n , each satisfiability test on $L \cup D' \cup \{\phi, \neg\psi\}$ can be reduced to $|\psi|$ satisfiability tests on all $L \cup D' \cup \{\phi, \neg\psi_i\}$. By Lemma 6.3, each such $L \cup D' \cup F$ with $F \in \{\{\phi\}, \{\phi, \neg\psi_i\}\}$ is q-Horn. Hence, by Theorem 6.1 a), each satisfiability test can be done in time $O(\|KB\| + \|d\|)$. This shows that $\kappa^+(\phi)$ and $\kappa^+(\phi \wedge \neg\psi)$ can be computed in polynomial time.

In summary, deciding whether $(KB, \sigma) z^+$ -entails d at strength τ can be done in polynomial time. \square

6.3 Intractability Results for 1-Literal-Horn Case

How do we obtain tractability of deciding z^* -, z_s^* -, lex -, lex_p -, and conditional entailment? In particular, are there any syntactic restrictions on default reasoning problems that give tractability? We could, for example, further restrict literal-Horn defaults by limiting the number of atoms in the antecedent of each default as follows. A default $\phi \rightarrow \psi$ is *1-literal-Horn* iff ϕ is either \top or an atom, and ψ is a literal. A *1-Horn clause* is a classical formula $\phi \Rightarrow \psi$, where ϕ is either \top or an atom, and ψ is a literal. A conditional knowledge base $KB = (L, D)$ is *1-literal-Horn* iff L is a finite set of 1-Horn clauses and D is a finite set of 1-literal-Horn defaults. A default reasoning problem (KB, d) is 1-literal-Horn iff both KB and d are 1-literal-Horn.

Unfortunately, the following theorem shows that deciding z^* -entailment is still intractable even for this very restricted kind of default reasoning problems.

Theorem 6.9 *Given a 1-literal-Horn conditional knowledge base KB , which is ε -consistent and minimal-core, and a 1-literal-Horn default d , deciding whether $KB z^*$ -entails d is co-NP-hard.*

Informally, intractability is due to the fact that the default knowledge generally does not fix a unique instantiation of the atoms to truth values, in particular, when defaults “fire back” into the antecedents of other defaults, and when defaults are logically related through their consequents.

Since z_s^* -entailment is a proper generalization of z^* -entailment (see Lemma 3.2), it immediately follows that deciding z_s^* -entailment is intractable in the 1-literal-Horn case.

Corollary 6.10 *Given a 1-literal-Horn conditional knowledge base KB , which is ε -consistent and robust, a strength assignment σ on KB , a 1-literal-Horn default d , and a strength τ , deciding whether $(KB, \sigma) z_s^*$ -entails d at strength τ is co-NP-hard.*

The following theorem shows that also deciding lex -entailment, lex_p -entailment, and conditional entailment is intractable in the 1-literal-Horn case.

Theorem 6.11 *a) Given an ε -consistent 1-literal-Horn conditional knowledge base KB and a 1-literal-Horn default d , deciding whether KB lex -entails d is co-NP-hard.*

b) Given a 1-literal-Horn conditional knowledge base KB with priority assignment π and a 1-literal-Horn default d , deciding whether (KB, π) lex_p -entails d is co-NP-hard.

c) Given a 1-literal-Horn conditional knowledge base KB and a 1-literal-Horn default d , deciding whether KB conditionally entails d is co-NP-hard.

6.4 Feedback-free Horn

We will see that deciding s -entailment, where $s \in \{z^*, z_s^*, lex, lex_p\}$, becomes tractable, if we assume that the default reasoning problems can be decomposed into smaller problems of size bounded by a constant.

6.4.1 Motivating examples

We now give some examples to illustrate the main ideas behind this kind of decomposability. Roughly speaking, given a default reasoning problem $(KB, d) = ((L, D), \phi \rightarrow \psi)$, we solve one classical reasoning problem w.r.t. L and ϕ , and one reduced default reasoning problem (KB', d') , where KB' is obtained from KB by eliminating irrelevant defaults, and d' is obtained from d by adding literals to its antecedent.

In the sequel, we assume that conditional knowledge bases are implicitly associated with a strength assignment σ and a priority assignment π , when $s = z_s^*$ and $s = lex_p$, respectively.

Example 6.4 Consider again the conditional knowledge base $KB = (L, D)$ shown in Example 3.1. Assume now that we are wondering whether KB s -entails the defaults $penguin \rightarrow fly$, $red \wedge bird \rightarrow fly$, $bird \rightarrow mobile$, $penguin \rightarrow arctic$, or $penguin \rightarrow wings$, where $s \in \{z^*, z_s^*, lex, lex_p\}$. As it turns out, each of these problems can be reduced to one classical reasoning problem and one default reasoning problem.

For instance, deciding whether KB s -entails $red \wedge bird \rightarrow fly$ is reduced to a classical reasoning problem w.r.t. L and $red \wedge bird$, and to a default reasoning problem w.r.t. the conditional knowledge base $KB' = (L, \{bird \rightarrow fly, fly \rightarrow mobile\})$, which is obtained from KB by sensibly eliminating irrelevant defaults. More precisely, it is reduced to computing the least model of $L \cup \{red \wedge bird\}$ (that is, the set of all atoms that are logically entailed by $L \cup \{red \wedge bird\}$, which is given by $\{red, bird\}$), and the default reasoning problem whether KB' s -entails $red \wedge bird \wedge \neg penguin \wedge \neg arctic \rightarrow fly$ (which is true).

The next example considers the classical *Nixon diamond*.

Example 6.5 The defeasible knowledge “generally, quakers are pacifists” and “generally, republicans are not pacifists” can be expressed by the following conditional knowledge base $KB = (L, D)$:

$$\begin{aligned} L &= \emptyset, \\ D &= \{quaker \rightarrow pacifist, republican \rightarrow \neg pacifist\}. \end{aligned}$$

We are now asked whether Nixon, being a quaker and a republican, is also a pacifist. That is, we are wondering whether KB s -entails the default $quaker \wedge republican \rightarrow pacifist$, where $s \in \{z^*, z_s^*, lex, lex_p\}$. We will see that this default reasoning problem can be reduced to one classical reasoning problem w.r.t. the set of atoms $\{quaker, republican\}$ and one default reasoning problem w.r.t. the set of atoms $\{pacifist\}$.

The following example shows a taxonomic hierarchy adorned with some default knowledge [3].

Example 6.6 The strict knowledge “all birds and fishes are animals”, “all penguins and sparrows are birds”, “no bird is a fish”, “no penguin is a sparrow”, and the defeasible knowledge “generally, animals do not swim”, “generally, fishes swim”, and “generally, penguins swim” can be represented by the following conditional knowledge base $KB = (L, D)$:

$$\begin{aligned} L &= \{bird \Rightarrow animal, fish \Rightarrow animal, penguin \Rightarrow bird, \\ &\quad sparrow \Rightarrow bird, bird \Rightarrow \neg fish, penguin \Rightarrow \neg sparrow\}, \\ D &= \{animal \rightarrow \neg swims, fish \rightarrow swims, penguin \rightarrow swims\}. \end{aligned}$$

Do sparrows generally swim? That is, does KB s -entail $sparrow \rightarrow swims$, where $s \in \{z^*, z_s^*, lex, lex_p\}$? This default reasoning problem can be reduced to one classical reasoning problem w.r.t. the set of atoms $\{animal, bird, fish, sparrow, penguin\}$ and one default reasoning problem w.r.t. the set of atoms $\{swims\}$: We first compute the least model of $L \cup \{sparrow\}$ in which $sparrow$, $bird$, and $animal$ are true, and then decide whether $(L, \{animal \rightarrow \neg swims\})$ s -entails $sparrow \wedge bird \wedge animal \wedge \neg fish \wedge \neg penguin \rightarrow swims$.

6.4.2 Definitions

Suppose that for a literal-Horn conditional knowledge base $KB = (L, D)$, there exists a set of atoms $At_a \subseteq At$ such that L is defined over At_a and that all consequents of definite literal-Horn defaults in D are defined over $At - At_a$. The greatest such At_a , which clearly exists then, is called the *activation set* of KB . Intuitively, in any “context” given by L and ϕ , where ϕ is either \top or a conjunction of atoms from At , all those atoms in At_a that are logically entailed (resp., not logically entailed) by $L \cup \{\phi\}$ can be safely set to **true** (resp., **false**) in the preferred models of $L \cup \{\phi\}$.

For At_a , there exists a unique partition $\{At_1, \dots, At_n\}$ of $At - At_a$, where $n \geq 0$ (and each At_i is nonempty), such that (i) if $n > 0$, then every $d \in D$ is defined over some $At_a \cup At_i$ with $i \in \{1, \dots, n\}$, and (ii) n is maximal. We call this partition the *default partition* of KB . A conditional knowledge base $KB = (L, D)$ is *k-feedback-free Horn* (or *k-ff-Horn*) iff it is literal-Horn, it has an activation set At_a , and it has a default partition $\{At_1, \dots, At_n\}$ such that every At_i with $i \in \{1, \dots, n\}$ has cardinality at most k .

Example 6.7 The two conditional knowledge bases shown in Examples 6.5 and 6.6 are both 1-ff-Horn. Their activation sets are given by $\{quaker, republican\}$ and $\{animal, bird, fish, sparrow, penguin\}$, respectively. Moreover, their default partitions are given by $\{\{pacifist\}\}$ and $\{\{swims\}\}$, respectively.

The conditional knowledge bases shown in Example 6.4 is 2-ff-Horn. Its activation set and its default partition are given by $\{penguin, bird, red\}$ and $\{\{fly, mobile\}, \{arctic\}, \{wings\}\}$, respectively.

Given a literal-Horn default reasoning problem $(KB, d) = ((L, D), \phi \rightarrow \psi)$, we use L^+ (resp., D^+) to denote the set of all definite formulas in L (resp., D). An atom $b \in At$ is *active* w.r.t. KB and d iff $L^+ \cup D^+ \cup At(d) \models b$. A classical formula α (resp., default δ) is *active* w.r.t. KB and d iff all atoms in α (resp., δ) are active w.r.t. KB and d . An atom $b \in At$ (resp., classical formula α , default δ) is *inactive* w.r.t. KB and d iff it is not active w.r.t. KB and d . We often denote by \widehat{D} the set of all active defaults in D , and omit KB and d when they are clear from the context. Intuitively, to decide whether KB s -entails d , where $s \in \{z^*, z_s^*, lex, lex_p\}$, it is sufficient to consider all defaults in KB that are active w.r.t. KB and d .

This intuition is more formally expressed by the following lemma. Roughly speaking, this lemma implies that ψ is true in every preferred model of $L \cup \{\phi\}$ w.r.t. D iff ψ is true in every preferred model of $L \cup \{\phi\}$ w.r.t. the set of all active defaults in D .

Lemma 6.12 *Let $(KB, d) = ((L, D), \phi \rightarrow \psi)$ be a literal-Horn default reasoning problem, and let I be a model of L . Then, there exists a model I^* of L such that (i) $I^*(\gamma) = I(\gamma)$ for all active classical formulas γ , (ii) I^* satisfies all inactive defaults in D , and (iii) I^* satisfies the same active defaults in D as I .*

A default reasoning problem $(KB, d) = ((L, D), \phi \rightarrow \psi)$ is *k-ff-Horn*, where $k \geq 1$, iff (i) it is literal-Horn, and (ii) $(L, \widehat{D} \cup \{d\})$ has an activation set \widehat{At}_a and a default partition $\{At_1, \dots, At_n\}$ such that d is defined over some $\widehat{At}_a \cup At_j$ with $|At_j| \leq k$ and ψ being a literal over At_j , where \widehat{D} is the set of all active defaults in D w.r.t. KB and d . The class *k-ff-Horn* consists of all *k-ff-Horn* default reasoning problems. We define the class *feedback-free Horn* (or *ff-Horn*) by $\text{ff-Horn} = \bigcup_{k \geq 1} k\text{-ff-Horn}$.

Example 6.8 (Red birds) Consider the literal-Horn default reasoning problem (KB, d) , where $KB = (L, D)$ as in Examples 3.1 and 6.4, and $d = red \wedge bird \rightarrow fly$. The set \widehat{D} of all active defaults in D w.r.t. KB and d is given as follows:

$$\widehat{D} = \{bird \rightarrow fly, bird \rightarrow wings, fly \rightarrow mobile\}.$$

It is easy to verify that $(L, \widehat{D} \cup \{d\})$ has the activation set $\widehat{At}_a = \{penguin, bird, red, arctic\}$ (recall that $L = \{penguin \Rightarrow bird\}$) and the default partition $\{At_1, At_2\}$, where $At_1 = \{fly, mobile\}$ and $At_2 = \{wings\}$. Moreover, d is defined over $\widehat{At}_a \cup At_1$ with $|At_1| = 2$. That is, (KB, d) is 2-ff-Horn.

For Horn conditional knowledge bases (L, D) with activation set At_a , and classical formulas α that are either \top or conjunctions of atoms from At , we define the classical formula α^* as follows. If $L \cup \{\alpha\}$ is satisfiable, then α^* is the conjunction of all $b \in At$ with $L \cup \{\alpha\} \models b$ and all $\neg b$ with $b \in At_a$ and $L \cup \{\alpha\} \not\models b$. Otherwise, we define $\alpha^* = \perp$. Moreover, for satisfiable $L \cup \{\alpha\}$, we define the world I_α^* over the activation set At_a by $I_\alpha^*(b) = \mathbf{true}$ iff $L \cup \{\alpha\} \models b$, for all $b \in At_a$. Informally, if $L \cup \{\alpha\}$ is satisfiable, then α^* is the conjunction of all atoms $b \in At - At_a$ that occur in α , all atoms $b \in At_a$ that are logically entailed by $L \cup \{\alpha\}$, and all negations of atoms $b \in At_a$ that are not logically entailed by $L \cup \{\alpha\}$.

6.4.3 Recognizing feedback-free Horn

The following result shows that both recognizing k -ff-Horn conditional knowledge bases, and computing their activation set and default partition are tractable.

Theorem 6.13 *a) Given a literal-Horn conditional knowledge base $KB = (L, D)$ and an integer $k \geq 1$, deciding whether KB is k -ff-Horn is possible in $O(\|L\| + \|D\|)$ time, that is, in time linear in the input size. b) Given a k -ff-Horn conditional knowledge base $KB = (L, D)$, computing the activation set At_a and the default partition $\{At_1, \dots, At_n\}$ is possible in $O(\|L\| + \|D\|)$ time, that is, in time linear in the input size.*

Proof. Let At_L be the set of all atoms $b \in At$ that occur in L , and let At_C be the set of all atoms $b \in At$ that occur in consequents of definite literal-Horn defaults $d \in D$. If $At_L \cap At_C \neq \emptyset$, then KB does not have any activation set, and thus KB is not k -ff-Horn. Otherwise, define $At_a = At - At_C$ and $\{At_1, \dots, At_n\}$ as the greatest partition of $At - At_a$ such that each $d \in D$ is defined over some $At_a \cup At_i$. If the cardinality of each At_i with $i \in \{1, \dots, n\}$ is at most k , then KB is k -ff-Horn. Otherwise, KB is not k -ff-Horn.

Computing the sets At_L , At_C , and At_a , computing the partition $\{At_1, \dots, At_n\}$ (that is, the connected components of the hypergraph $G = (V, E) = (At - At_a, \{At(d) - At_a \mid d \in D\})$), and deciding whether $At_L \cap At_C = \emptyset$ and whether each At_i has a cardinality of at most k can obviously be done in linear time using standard methods and data structures. The results follow from this. \square

The next result shows that also recognizing k -ff-Horn default reasoning problems is tractable.

Theorem 6.14 *a) Given a literal-Horn default reasoning problem (KB, d) with $KB = (L, D)$, and an integer $k \geq 1$, deciding whether (KB, d) is k -ff-Horn can be done in time linear in the input size. b) Given a k -ff-Horn default reasoning problem (KB, d) with $KB = (L, D)$, computing the set \widehat{D} of active defaults in D w.r.t. KB and d can be done in time linear in the input size.*

Proof. Since $L^+ \cup D^+ \cup At(d)$ is Horn, the set $At_R = \{b \in At \mid L^+ \cup D^+ \cup At(d) \models b\}$ can be computed in linear time. When At_R is given, the set \widehat{D} of all defaults that are active w.r.t. KB and d can be computed in linear time. By Theorem 6.13 b), the activation set At_a and the default partition $\{At_1, \dots, At_n\}$ of $(L, \widehat{D} \cup \{d\})$ can be computed in linear time. Clearly, determining $i \in \{1, \dots, n\}$ such that d is defined over $At_a \cup At_i$, and checking whether $|At_i| \leq k$ can be done in linear time. \square

6.4.4 Maximum entropy semantics: RANKING

In the sequel, let $KB = (L, D)$ be an ε -consistent k -ff-Horn conditional knowledge base with positive strength assignment σ . Let At_a denote the activation set of KB , and let $\{At_1, \dots, At_n\}$ be the default partition of KB . Let z_s^* be a ranking that maps each $d \in D$ to a positive integer, and let κ_s^* be defined by (7).

In order to compute the default ranking z_s^* , we have to compute ranks of the form $\kappa_s^*(\alpha \wedge \beta)$, where α is either \top or a conjunction of atoms from At_a , and β is either \top or a conjunction of literals over $At - At_a$. The following lemma shows that such $\kappa_s^*(\alpha \wedge \beta)$ coincide with $\kappa_s^*(\alpha^* \wedge \beta)$.

Lemma 6.15 *Let α be either \top or a conjunction of atoms from At_a . Let β be either \top or a conjunction of literals over $At - At_a$. Then, $\kappa_s^*(\alpha \wedge \beta) = \kappa_s^*(\alpha^* \wedge \beta)$.*

In the sequel, for each $i \in \{1, \dots, n\}$, let D_i be the set of all defaults in D that are defined over $At_a \cup At_i$. Let the function $\kappa_{s,i}^*$ on worlds I over At be defined as follows:

$$\kappa_{s,i}^*(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D_i \\ \sum_{d \in D_i: I \not\models d} z_s^*(d) & \text{otherwise.} \end{cases} \quad (8)$$

As D can be decomposed into the sets of defaults D_i over $At_a \cup At_i$, also κ_s^* can be decomposed:

Lemma 6.16 *Let α be either \top or a conjunction of atoms from At_a . Let $\beta = \beta_1 \wedge \dots \wedge \beta_n$, where each β_i is either \top or a conjunction of literals over At_i . Then, $\kappa_s^*(\alpha^* \wedge \beta) = \sum_{i \in \{1, \dots, n\}} \kappa_{s,i}^*(\alpha^* \wedge \beta_i)$.*

The following lemma will be useful to characterize the robustness of κ_s^* .

Lemma 6.17 *For each $i \in \{1, 2\}$, let α_i be either \top or a conjunction of atoms from At_a . Let β_1 and β_2 be conjunctions of literals over some At_j and At_k , respectively, with $j, k \in \{1, \dots, n\}$. Then, $L \cup \{\alpha_1 \wedge \beta_1\}$ and $L \cup \{\alpha_2 \wedge \beta_2\}$ have a common minimal model w.r.t. κ_s^* iff $\kappa_s^*(\alpha_1 \wedge \beta_1) = \kappa_s^*(\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2) = \kappa_s^*(\alpha_2 \wedge \beta_2)$ and both $L \cup \{\alpha_1 \wedge \alpha_2\}$ and $\beta_1 \wedge \beta_2$ are satisfiable.*

The following result shows that computing the default ranking z_s^* is tractable in the k -ff-Horn case. Since z_s^* is a proper generalization of z^* , this result shows also that computing the default ranking z^* is tractable in the k -ff-Horn case.

Theorem 6.18 *Let $k > 0$ a fixed integer. Given an ε -consistent k -ff-Horn conditional knowledge base $KB = (L, D)$ with positive strength assignment σ , computing the default ranking z_s^* for KB , if KB is robust, and returning nil otherwise, can be done in time polynomial in the input size.*

Proof. We now show that algorithm **z_s^* -ranking** (see Figure 6) can be done in polynomial time. Step 1 in **z_s^* -ranking** runs in $O(|D|)$. Steps 3–6 are performed $O(|D|)$ times. In particular, in Step 3, we first evaluate $O(|D|)$ expressions of the form $\sigma(d) + \text{minv}(d)$, determine the minimum of these values, and select one default d at which this minimum is attained. In Steps 5 and 6, we also evaluate $\text{minf}(d)$. In Step 8, we have to check that κ_s^* satisfies (6) for all $\phi \rightarrow \psi \in D$, that z_s^* is positive, and that κ_s^* is robust. For the former, we verify $O(|D|)$ expressions of the kind $\text{minf}(d) = \sigma(d) + \text{minv}(d)$, while for the latter, we apply Lemma 6.17 as follows. For every two distinct $\gamma_1 \rightarrow \delta_1, \gamma_2 \rightarrow \delta_2 \in D$, we verify that it is not the case that (a) $\kappa_s^*(\gamma_1 \wedge \neg \delta_1) = \kappa_s^*(\gamma_1 \wedge \gamma_2 \wedge \neg \delta_1 \wedge \neg \delta_2) = \kappa_s^*(\gamma_2 \wedge \neg \delta_2)$, and (b) both $L \cup \{\gamma_1 \wedge \gamma_2\}$ and $\neg \delta_1 \wedge \neg \delta_2$ are satisfiable. Thus, to prove that algorithm **z_s^* -ranking** can be done in polynomial time, it is now sufficient to show that the following tasks can be done in polynomial time:

- (i) computing $\text{minv}(d) = \kappa_s^*(\gamma \wedge \delta)$ for $d = \gamma \rightarrow \delta \in D$,
- (ii) computing $\text{minf}(d) = \kappa_s^*(\gamma \wedge \neg \delta)$ for $d = \gamma \rightarrow \delta \in D$,
- (iii) computing $\kappa_s^*(\gamma_1 \wedge \gamma_2 \wedge \neg \delta_1 \wedge \neg \delta_2)$ for distinct $\gamma_1 \rightarrow \delta_1, \gamma_2 \rightarrow \delta_2 \in D$,
- (iv) deciding whether $L \cup \{\gamma_1 \wedge \gamma_2\}$ and $\neg \delta_1 \wedge \neg \delta_2$ are satisfiable for distinct $\gamma_1 \rightarrow \delta_1, \gamma_2 \rightarrow \delta_2 \in D$.

Clearly, (iv) can be done in polynomial time, since deciding whether $L \cup \{\gamma_1 \wedge \gamma_2\}$ and $\neg \delta_1 \wedge \neg \delta_2$ are satisfiable can be done in linear time. To prove that (i)–(iii) can be done in polynomial time, it is now

sufficient to show that every $\kappa_s^*(\alpha \wedge \beta_1 \wedge \beta_2)$, where α is either \top or a conjunction of atoms from At_a , and β_1 and β_2 are either \top or conjunctions of literals from At_j and At_l , respectively, with $j \neq l$, can be computed in polynomial time. If $L \cup \{\alpha \wedge \beta_1 \wedge \beta_2\}$ is unsatisfiable, which can be checked in linear time as $L \cup \{\alpha \wedge \beta_1 \wedge \beta_2\}$ is Horn, then $\kappa_s^*(\alpha \wedge \beta_1 \wedge \beta_2) = \infty$. Otherwise, by Lemmata 6.15 and 6.16, it follows:

$$\kappa_s^*(\alpha \wedge \beta_1 \wedge \beta_2) = \kappa_{s,j}^*(\alpha^* \wedge \beta_1) + \kappa_{s,l}^*(\alpha^* \wedge \beta_2) + \sum_{i \in \{1, \dots, n\} - \{j, l\}} \kappa_{s,i}^*(\alpha^*).$$

Since $L \cup \{\alpha\}$ is Horn, α^* can be computed in linear time. Moreover, since $|At_i| \leq k$ for all $i \in \{1, \dots, n\}$, also $\kappa_{s,j}^*(\alpha^* \wedge \beta_1)$, $\kappa_{s,l}^*(\alpha^* \wedge \beta_2)$, and every $\kappa_{s,i}^*(\alpha^*)$ with $i \in \{1, \dots, n\} - \{j, l\}$ can be computed in polynomial time by simple exhaustive search: For $i = j$ (resp., $i = l$), we generate all worlds I over the set of atoms $At_a \cup At_i$ with $I \models \alpha^*$ and compute the minimum of $\kappa_{s,i}^*(I)$ subject to all such worlds I with $I \models \alpha^* \wedge \beta_1$ (resp., $I \models \alpha^* \wedge \beta_2$). Moreover, for every $i \in \{1, \dots, n\} - \{j, l\}$, we generate all worlds I over the set of atoms $At_a \cup At_i$ with $I \models \alpha^*$ and compute the minimum of $\kappa_{s,i}^*(I)$ subject to all such worlds I . Thus, $\kappa_s^*(\alpha \wedge \beta_1 \wedge \beta_2)$ can be computed in polynomial time. \square

6.4.5 Maximum entropy semantics: ENTAILMENT and RANK-ENTAILMENT

In the sequel, let $(KB, d) = ((L, D), \phi \rightarrow \psi)$ be a k -ff-Horn default reasoning problem with ε -consistent and robust KB . Let σ be a positive strength assignment on KB . Let \widehat{D} be the set of all defaults in D that are active w.r.t. KB and d , let \widehat{At}_a be the activation set of $(L, \widehat{D} \cup \{d\})$, and let $\{At_1, \dots, At_n\}$ be the default partition of $(L, \widehat{D} \cup \{d\})$. Let z_s^*, κ_s^* , where z_s^* is positive, be the unique solution of (6) and (7) for all $d = \phi \rightarrow \psi \in D$ and $I \in \mathcal{I}_{At}$. Let the function $\widehat{\kappa}_s^*$ on worlds I over At be defined as follows:

$$\widehat{\kappa}_s^*(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup \widehat{D} \\ \sum_{d \in \widehat{D}: I \not\models d} z_s^*(d) & \text{otherwise.} \end{cases} \quad (9)$$

The following lemma shows that κ_s^* coincides with $\widehat{\kappa}_s^*$ on all active classical formulas. That is, when we compute the rank of an active classical formula, we can restrict our attention to all active defaults in D .

Lemma 6.19 *Let γ be a classical formula that is active w.r.t. KB and d . Then, $\kappa_s^*(\gamma) = \widehat{\kappa}_s^*(\gamma)$.*

For every $i \in \{1, \dots, n\}$, let \widehat{D}_i denote the set of all defaults in \widehat{D} that are defined over $\widehat{At}_a \cup At_i$, and let $\widehat{\sigma}_i$ be the restriction of σ to \widehat{D}_i . Let $\widehat{z}_{s,i}^*$ be a default ranking that maps each default in \widehat{D}_i to a positive integer, and let the function $\widehat{\kappa}_{s,i}^*$ on worlds I over At be defined as follows:

$$\widehat{\kappa}_{s,i}^*(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup \widehat{D}_i \\ \sum_{d \in \widehat{D}_i: I \not\models d} \widehat{z}_{s,i}^*(d) & \text{otherwise.} \end{cases} \quad (10)$$

To decide whether (KB, σ) z_s^* -entails d at given strength $\tau > 0$, we will need all $z_s^*(\delta)$ with $\delta \in \widehat{D}_j$, where $j \in \{1, \dots, n\}$ such that δ is defined over $\widehat{At}_a \cup At_j$. The following lemma shows that the restriction of z_s^* to \widehat{D}_j coincides with the default ranking for (L, \widehat{D}_j) under strength assignment $\widehat{\sigma}_j$.

Lemma 6.20 Let $\widehat{z}_{s,j}^*(d) = z_s^*(d)$ for all $d \in \widehat{D}_j$, and let $\widehat{\kappa}_{s,j}^*$ be defined by (10) for all $I \in \mathcal{I}_{At}$ (with $i = j$). Then, $\widehat{\kappa}_{s,j}^*$ is robust, and $\widehat{z}_{s,j}^*, \widehat{\kappa}_{s,j}^*$ is the unique solution to the following system of equations:

$$\widehat{\kappa}_{s,j}^*(\gamma \wedge \neg\delta) = \widehat{\sigma}_j(\gamma \rightarrow \delta) + \widehat{\kappa}_{s,j}^*(\gamma \wedge \delta) \text{ for all } \gamma \rightarrow \delta \in \widehat{D}_j, \text{ and (10) for all } I \in \mathcal{I}_{At} \text{ (with } i = j). \quad (11)$$

We finally show that deciding z_s^* -entailment is tractable in the k -ff-Horn case. Again, since z_s^* properly generalizes z^* , this result shows also that deciding z^* -entailment is tractable in the k -ff-Horn case. Trivially, these tractability results remain true when z_s^* and z^* , respectively, are part of the input.

Theorem 6.21 Let $k > 0$ be fixed. Given a k -ff-Horn default reasoning problem $(KB, d) = ((L, D), \phi \rightarrow \psi)$, where KB is ε -consistent and robust, and a positive strength assignment σ on KB , deciding whether (KB, σ) z_s^* -entails d at given strength $\tau > 0$ can be done in time polynomial in the input size.

Proof. Recall that (KB, σ) z_s^* -entails d at strength τ iff $\kappa_s^*(\phi) = \infty$ or $\kappa_s^*(\phi \wedge \psi) + \tau \leq \kappa_s^*(\phi \wedge \neg\psi)$. That is, iff either (i) $L \cup \{\phi, \neg\psi\}$ is unsatisfiable, or (ii) both $L \cup \{\phi, \psi\}$ and $L \cup \{\phi, \neg\psi\}$ are satisfiable, and $\kappa_s^*(\phi \wedge \psi) + \tau \leq \kappa_s^*(\phi \wedge \neg\psi) (< \infty)$.

Since both $L \cup \{\phi, \psi\}$ and $L \cup \{\phi, \neg\psi\}$ are Horn, their satisfiability can be tested in linear time. Hence, it is now sufficient to show that $\kappa_s^*(\phi \wedge \psi) + \tau \leq \kappa_s^*(\phi \wedge \neg\psi)$ can be verified in polynomial time.

The classical formulas $\phi \wedge \psi$ and $\phi \wedge \neg\psi$ are of the form $\alpha \wedge \beta_1$ and $\alpha \wedge \beta_2$, respectively, where α is either \top or a conjunction of atoms from \widehat{At}_a , and both β_1 and β_2 are conjunctions of literals over some At_j with $j \in \{1, \dots, n\}$. Clearly, the default $\phi \rightarrow \psi$ is active. Thus, also $\alpha \wedge \beta_1$ and $\alpha \wedge \beta_2$ are active. By Lemmata 6.15, 6.16, and 6.19, it thus follows for $l \in \{1, 2\}$:

$$\begin{aligned} \kappa_s^*(\alpha \wedge \beta_l) &= \widehat{\kappa}_s^*(\alpha \wedge \beta_l) \\ &= \widehat{\kappa}_s^*(\alpha^* \wedge \beta_l) \\ &= \widehat{\kappa}_{s,j}^*(\alpha^* \wedge \beta_l) + \sum_{i \in \{1, \dots, n\} - \{j\}} \widehat{\kappa}_{s,i}^*(\alpha^*), \end{aligned}$$

where $\widehat{\kappa}_{s,j}^*$ is defined by (10) for all $I \in \mathcal{I}_{At}$, with $i = j$ and $\widehat{z}_{s,j}^*(d) = z_s^*(d)$ for all $d \in \widehat{D}_j$. This shows that $\kappa_s^*(\phi \wedge \psi) + \tau \leq \kappa_s^*(\phi \wedge \neg\psi)$ is equivalent to $\widehat{\kappa}_{s,j}^*(\phi^* \wedge \psi) + \tau \leq \widehat{\kappa}_{s,j}^*(\phi^* \wedge \neg\psi)$.

Since $L \cup \{\phi\}$ is Horn, ϕ^* can be computed in linear time. By Lemma 6.20, the default ranking $\widehat{z}_{s,j}^*$ can be computed by running z_s^* -ranking on the conditional knowledge base (L, \widehat{D}_j) under strength assignment $\widehat{\sigma}_j$. By Lemma 6.15, this can be done in polynomial time. Moreover, since $|At_j| \leq k$, both $\widehat{\kappa}_{s,j}^*(\phi^* \wedge \psi)$ and $\widehat{\kappa}_{s,j}^*(\phi^* \wedge \neg\psi)$ can be computed in polynomial time by simple exhaustive search. \square

Example 6.9 (Red birds continued) Let the 2-ff-Horn default reasoning problem (KB, d) be given by $KB = (L, D)$ of Examples 3.1 and 6.4, and $d = red \wedge bird \rightarrow fly$. Let $\sigma(\delta) = 1$ for all $\delta \in D$. Recall from Examples 3.1 and 6.8 that $L = \{penguin \Rightarrow bird\}$, $\widehat{D} = \{bird \rightarrow fly, bird \rightarrow wings, fly \rightarrow mobile\}$, $\widehat{At}_a = \{penguin, bird, red, arctic\}$, $At_1 = \{fly, mobile\}$, and $At_2 = \{wings\}$.

The default d is z_s^* -entailed by (KB, σ) at strength τ iff either (i) $L \cup \{red \wedge bird, \neg fly\}$ is unsatisfiable, or (ii) both $L \cup \{red \wedge bird, fly\}$ and $L \cup \{red \wedge bird, \neg fly\}$ are satisfiable, and

$$\kappa_s^*(red \wedge bird \wedge fly) + \tau \leq \kappa_s^*(red \wedge bird \wedge \neg fly).$$

Here, (ii) applies, and by Lemmata 6.15, 6.16, and 6.19, the latter inequality is equivalent to

$$\widehat{\kappa}_{s,1}^*((red \wedge bird)^* \wedge fly) + \tau \leq \widehat{\kappa}_{s,1}^*((red \wedge bird)^* \wedge \neg fly),$$

which is equivalent to

$$\widehat{\kappa}_{s,1}^*(\neg penguin \wedge bird \wedge red \wedge \neg arctic \wedge fly) + \tau \leq \widehat{\kappa}_{s,1}^*(\neg penguin \wedge bird \wedge red \wedge \neg arctic \wedge \neg fly),$$

where $\widehat{\kappa}_{s,1}^*$ is given through the ranking $\widehat{z}_{s,1}^*$ for $(L, \widehat{D}_1) = (L, \{bird \rightarrow fly, fly \rightarrow mobile\})$ under $\widehat{\sigma}_1 = \sigma|_{\widehat{D}_1}$.

It is now easy to verify that $\widehat{z}_{s,1}^*(d_1) = 1$ for all $d_1 \in \widehat{D}_1$, that both $L \cup \{red \wedge bird, fly\}$ and $L \cup \{red \wedge bird, \neg fly\}$ are satisfiable, that $\widehat{\kappa}_{s,1}^*(\neg penguin \wedge bird \wedge red \wedge \neg arctic \wedge fly) = 0$, and that $\widehat{\kappa}_{s,1}^*(\neg penguin \wedge bird \wedge red \wedge \neg arctic \wedge \neg fly) = 1$. This shows that $(KB, \sigma) z_s^*$ -entails $red \wedge bird \rightarrow fly$ at strength 1.

6.4.6 Lexicographic entailment

We now focus on lexicographic entailment. In what follows, let $(KB, d) = ((L, D), \phi \rightarrow \psi)$ be a k -ff-Horn default reasoning problem. Let π be a priority assignment on KB . Let \widehat{D} be the set of all defaults in D that are active w.r.t. KB and d , and let \widehat{At}_a (resp., $\{At_1, \dots, At_n\}$) be the activation set (resp., default partition) of $(L, \widehat{D} \cup \{d\})$. Let $\widehat{\pi}$ be the unique priority assignment on (L, \widehat{D}) that is consistent with π on KB (that is, $\widehat{\pi}(\delta) < \widehat{\pi}(\delta')$ iff $\pi(\delta) < \pi(\delta')$, for all $\delta, \delta' \in \widehat{D}$).

To decide whether $(KB, \pi) \text{lex}_p$ -entails d , we must check whether every π -preferred model of $L \cup \{\phi\}$ satisfies ψ . The following lemma shows that we can equivalently check whether every $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$ satisfies ψ . That is, we can restrict our attention to all active defaults in D .

Lemma 6.22 *Let γ be a classical formula that is active w.r.t. KB and d . Then,*

- a) *For every $\widehat{\pi}$ -preferred model I of $L \cup \{\phi\}$, there is a π -preferred model I^* of $L \cup \{\phi\}$ with $I^*(\gamma) = I(\gamma)$.*
- b) *Every π -preferred model of $L \cup \{\phi\}$ is also a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$.*

The next lemma shows that every $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$ satisfies ψ iff every $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$ satisfies ψ . That is, we can assume that every atom in \widehat{At}_a that is logically entailed (resp., not logically entailed) by $L \cup \{\phi\}$ is assigned the truth value **true** (resp., **false**).

Lemma 6.23 a) *For every $\widehat{\pi}$ -preferred model I of $L \cup \{\phi\}$, there exists a $\widehat{\pi}$ -preferred model J of $L \cup \{\phi^*\}$ such that $J|_{At-\widehat{At}_a} = I|_{At-\widehat{At}_a}$.*

- b) *Every $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$ is also a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$.*

In the sequel, for every $i \in \{1, \dots, n\}$, let \widehat{D}_i be the set of all defaults in \widehat{D} that are defined over $\widehat{At}_a \cup At_i$. Let $\widehat{\pi}_i$ be the unique priority assignment on (L, \widehat{D}_i) that is consistent with $\widehat{\pi}$ on KB (that is, $\widehat{\pi}_i(\delta) < \widehat{\pi}_i(\delta')$ iff $\widehat{\pi}(\delta) < \widehat{\pi}(\delta')$, for all $\delta, \delta' \in \widehat{D}_i$). Let $j \in \{1, \dots, n\}$ be such that d is defined over $\widehat{At}_a \cup At_j$.

The following lemma shows that every $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$ satisfies ψ iff every $\widehat{\pi}_j$ -preferred model of $L \cup \{\phi^*\}$ satisfies ψ . That is, we can restrict our attention to all defaults in the cluster \widehat{D}_j .

Lemma 6.24 a) *For every $\widehat{\pi}_j$ -preferred model I_j of $L \cup \{\phi^*\}$, there exists a $\widehat{\pi}$ -preferred model J of $L \cup \{\phi^*\}$ such that $J|_{At_j} = I_j|_{At_j}$.*

- b) *Every $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$ is a $\widehat{\pi}_j$ -preferred model of $L \cup \{\phi^*\}$.*

The following result shows that deciding lex_p -entailment is tractable in the k -ff-Horn case. Moreover, since computing the z -partition for ε -consistent conditional knowledge bases KB is tractable in the Horn case [49], this result shows also that deciding lex -entailment is tractable in the k -ff-Horn case.

Theorem 6.25 *Let $k > 0$ be fixed. Given a k -ff-Horn default reasoning problem $(KB, d) = ((L, D), \phi \rightarrow \psi)$ and a priority assignment π on KB , deciding whether $(KB, \pi) \text{lex}_p$ -entails d can be done in time linear in the input size.*

Proof. By Lemmata 6.22–6.24, (KB, π) lex_p -entails d iff either $L \cup \{\phi\}$ is unsatisfiable, or all $\hat{\pi}_j$ -preferred models of $L \cup \{\phi^*\}$ satisfy ψ . Since $L \cup \{\phi\}$ is Horn, deciding whether $L \cup \{\phi\}$ is unsatisfiable can be done in linear time. If $L \cup \{\phi\}$ is satisfiable, then we have to compute ϕ^* , which can be done in linear time. Moreover, as k is fixed, computing all worlds I over the set of atoms $\widehat{At}_a \cup At_j$ such that $I \models \phi^*$ can also be done in linear time. Finally, computing all such $\hat{\pi}_j$ -preferred worlds and verifying whether they all satisfy ψ can be done in linear time. \square

Example 6.10 (Red birds continued) Let the 2-ff-Horn default reasoning problem (KB, d) be given by $KB = (L, D)$ of Examples 3.1 and 6.4, and $d = red \wedge bird \rightarrow fly$. Let π map the defaults in $\{bird \rightarrow fly, bird \rightarrow wings, fly \rightarrow mobile\}$ and $\{penguin \rightarrow \neg fly, penguin \rightarrow arctic\}$ to 0 and 1, respectively.

By Lemmata 6.22–6.24, (KB, π) lex_p -entails $red \wedge bird \rightarrow fly$ iff either $L \cup \{red \wedge bird\}$ is unsatisfiable, or all $\hat{\pi}_1$ -preferred models of $L \cup \{(red \wedge bird)^*\} = L \cup \{\neg penguin \wedge bird \wedge red \wedge \neg arctic\}$ satisfy fly , where $\hat{\pi}_1$ is the priority assignment on $(L, \widehat{D}_1) = (L, \{bird \rightarrow fly, fly \rightarrow mobile\})$ that maps each element of \widehat{D}_1 to 0. It is now easy to verify that this is indeed the case. That is, (KB, π) lex_p -entails $red \wedge bird \rightarrow fly$.

Thus, we have established that s -entailment, where $s \in \{z^*, z_s^*, lex, lex_p\}$, is tractable in the ff-Horn case. The question is now whether a similar result also holds for Geffner’s conditional entailment, which is the remaining intractable notion of entailment in the 1-literal Horn case. It appears that the technique that we have successfully applied for the other approaches is not applicable to establish tractability for conditional entailment. There, the world rankings with respect to the full set D of defaults are equivalent to the sums of the world rankings with respect to the clusters $\widehat{D}_1, \dots, \widehat{D}_n$. In case of conditional entailment, however, a similar equivalence does not hold for preference orderings on worlds, since they are, in general, not total orderings. This requires the development of more complex techniques, which we leave for future work.

7 Related Work

In this section, we consider some work on complexity issues for related subjects.

7.1 Conditional Modal Logics

A stream of semantics for conditional knowledge bases, which we have not considered in this paper, is inherited from conditional modal logics, cf. [26, 27, 14, 15, 59, 36]. Roughly speaking, in these approaches a conditional statement $\phi \rightarrow \psi$ is true at a world w in a set of possible worlds W , if ψ is true in a set $f(w, \phi)$ of selected worlds in which ϕ is true. The worlds $f(w, \phi)$ may be the least exceptional, most normal, etc worlds from the view of w . To capture these notions, the possible worlds are related by a Lewis-style accessibility relation, which in general depends on the world w from which it is considered. An important note is that some conditional modal logics treat \rightarrow as a first class connective, and thus allow, in particular, nested use of \rightarrow , as well as Boolean combinations of defaults, which is not possible in our conditional knowledge bases. Thus, our complexity results for default reasoning from conditional knowledge bases have to be compared to complexity results for “flat” fragments of conditional modal logics, in which no nesting of \rightarrow is allowed, and no \rightarrow connective occurs inside the scope of another connective.

The work of Boutilier gives a deep study of modal conditional logics of normality [14, 16, 15], which goes beyond Delgrande’s early work on formalizing default reasoning through this approach [26]. Boutilier presented in [14] a conditional logic CT4D, which is equivalent to the modal logic S4.3 and whose flat fragment corresponds to a slight extension of rational consequence in [61]. In his later work [16], he introduced conditional logics CT4O and CO and showed that, in our terminology, entailment of a default $\phi \rightarrow \psi$ from

a default knowledge base KB under Lehmann’s preferential entailment (resp., rational entailment) [61] is equivalent to provability of $\phi \rightarrow \psi$ from KB under logic CT4O (resp., CO). Thus, for default reasoning in our setting, these two give the same semantics. Furthermore, Boutilier showed that ε -entailment of a default $\phi \rightarrow \psi$ from an ε -consistent default knowledge base KB is equivalent to provability of $\phi \rightarrow \psi$ from KB in logic CT4O. This and further observations on correspondences between different notions of consistency led Boutilier to suggest CT4O and CO as natural extensions of ε -semantics to the case of knowledge bases which contain nestings and Boolean combinations of defaults. Complexity results in [14, 15] show that under CT4D and CT4O semantics, entailment of a default $\phi \rightarrow \psi$ from a knowledge base KB , given by formulas from the flat fragment of CT4D and CT4O, respectively, is co-NP-complete. The same complexity applies to our more restrictive conditional knowledge bases.

An extensive analysis of the complexity of Lewis-style conditional modal logics has been carried out by Friedman and Halpern [36]. In their paper, they have analyzed the effect of semantic restrictions given by conditions on the set of worlds W_w which is considered possible at a world w such as *Normality* ($W_w \neq \emptyset$), *Reflexivity* ($w \in W_w$), and *Centering* (w is a minimal element in W_w with respect to w ’s accessibility relation \preceq_w). Halpern and Friedman gave axiomatizations of these conditions, and they determined the complexity of the logics emerging from (combinations of) these conditions, where they paid special attention to syntactical fragments of the full language. In particular, their results on bounded nestings of the conditional connective \rightarrow imply that entailment of a default statement $\phi \rightarrow \psi$ from a conditional knowledge base, which is given by a set of formulas KB in the flat fragment (i.e., no nesting of \rightarrow is allowed) of the language, is co-NP-complete for a wide range of conditional modal logics.

7.2 Belief Revision

In belief revision, one is concerned with the problem of incorporating a new belief, given through a sentence, into a current state of belief, given by a set of sentences. The new belief might contradict the current state of belief, though, and it is not immediately clear how this should be handled. Alchourrón, Gärdenfors, and Makinson (AGM) presented in the famous paper [2] several equivalent models for revision, which remove beliefs from the current state in order to reconcile it with the new piece of information, in a way such that a set of meaningful postulates is satisfied. As a salient feature, these postulates respect minimality of change. Since then, a number of different methods and operators for belief revision have been proposed, see e.g. [55, 39, 73]. Intuitively, default reasoning from conditional knowledge bases and belief revision are somehow related, since the derivation of plausible conclusions involves the retraction of statements which would lead to contradiction. The relationship has been considered more in detail in [16, 49]. Boutilier argues that default reasoning can be viewed as a special case of belief revision, and claims that “... *default reasoning can be thought of as the revision of a theory of expectations in order to incorporate what is known*” [16, p. 67]. In the same line, Goldszmidt and Pearl [16, 49] have shown that implementation and characterization issues in belief revision can be realized through default knowledge. We refer to [16, 49] for more details.

On the complexity side, a number of different revision approaches have been characterized, see e.g. [71, 76, 72, 64, 73]. In particular, the following reasoning problem has been considered there: Given a knowledge base, consisting of a set T of classical formulas, and classical formulas ϕ and ψ , is it the case that ψ is true in T after revision by ϕ ? This is also known as the Ramsey Test for conditional statements of the form “if ϕ were true, then ψ would be true.” It appeared that the computational complexity of this problem covers a whole range of complexity classes at the low end of the polynomial hierarchy up to its third level. In particular, it is Π_2^P -complete for a vast number of approaches, and thus has the same complexity as

Geffner’s conditional entailment. This implies that polynomial time translations between the Ramsey Test in these approaches and Geffner’s conditional entailment exist, which means that semantic relationships in terms of efficient (polynomial-time computable) embeddings among the formalisms (see [52] for various notions of embeddings) may exist.

Analogous complexity correspondences can be noted between other revision methods and the semantics for conditional knowledge bases that we have studied in this paper. The comprehensive survey [73] lists several revision approaches which are P^{NP} -complete (in particular, linear revision and lexicographic revision) or P_{\parallel}^{NP} -complete (in particular, Dalal’s operator [24], cardinality maximal revision, and cut revision). Many of these correspondences seem to be more of a computational nature, as no immediate semantic relationship is apparent. However, for lexicographic revision, the Ramsey Test for ϕ, ψ on a classical knowledge base T amounts just to lex_p -entailment of $\phi \rightarrow \psi$ from a naturally corresponding default knowledge base; the proof in [73], showing P^{NP} -hardness of the Ramsey Test for the Horn case, thus establishes that the problem ENTAILMENT is P^{NP} -hard in the Horn case. A slight adaptation sharpens this to a proof for the literal-Horn case different from ours. However, this proof requests a Horn default of form $\phi \rightarrow \psi$ where both ϕ and ψ are atoms, while ours has $\phi = \top$ (cf. Theorem 5.18). Furthermore, by a suitable extension of the reduction in the proof of Theorem 6.11 b), which adds some more defaults (we do not carry this out here), we can show that P^{NP} -hardness holds also for the 1-literal Horn case. A similar result can not be concluded from [73].

7.3 Nonmonotonic Logics

Another area related to conditional knowledge bases—which is also related to belief revision—are non-monotonic logics. A number of nonmonotonic logics and formalisms have been proposed in the past decades for capturing common sense reasoning, including major formalisms such as circumscription [67, 65], default logic [82], Doyle and McDermott’s nonmonotonic logics [69, 68], and Moore’s autoepistemic logic [70]; see [66]. The computational complexity of nonmonotonic logics has been studied in many papers, e.g., [56, 87, 50, 74, 33, 21, 77] to mention a few comprehensive studies, and is quite well-understood. As in the case of belief revision, the complexity of most of these logics resides at the second level of the polynomial hierarchy. More precisely, the problem of deciding whether a given classical formula α is a consequence of a given knowledge base T under so called cautious reasoning, is a Π_2^P -complete problem. Entailment of a conditional $\phi \rightarrow \psi$ from a conditional knowledge base KB can be viewed as deciding logical consequence of ψ from $T \cup \{\phi\}$, where T is a theory in the underlying logic that is augmented by ϕ .

The most famous and influential among the nonmonotonic formalisms is perhaps Reiter’s default logic [82], in which a set of classical formulas is augmented by default rules of the form $\frac{\alpha : M\beta_1, \dots, M\beta_n}{\gamma}$ which read “if α is provable and each of β_1, \dots, β_n can be consistently assumed (i.e., does not lead to contradiction), then conclude that γ is provable.” Many variants and refinements of this approach have been developed, see e.g. [66]. In [56, 87], a rich taxonomy of classes of default rules $\frac{\alpha : M\beta}{\gamma}$ has been defined, by imposing syntactic conditions on their constituents α, β , and γ and on the structure of the set of defaults. Syntactically, our class of literal-Horn defaults corresponds to the class of Horn defaults in [56], and our class of 1-literal Horn defaults corresponds to the class of normal unary and prerequisite-free normal unary defaults in [56, 87]. However, our class of ff-Horn defaults has no corresponding class in [56, 87].

Semantically, Reiter’s default rules and our conditional rules are quite different. For example, from the conditional rules $a \rightarrow b$ and $\neg a \rightarrow b$ we can conclude that $\top \rightarrow b$ (that is, b) is true under all the semantics for conditional knowledge bases considered in this paper, while from the corresponding defaults $\frac{a : Mb}{b}$, $\frac{\neg a : Mb}{b}$ neither b nor $\neg b$ can be concluded. Furthermore, as argued in [43], a rule $\frac{a : Mb}{b}$ may be seen as a soft constraint for believing b when a is known, while a conditional rule $a \rightarrow b$ can be viewed as a hard

constraint to believe b in a limited context defined by a and possibly some background knowledge; see [43] for further discussion. Thus, because of these apparent differences, a comparison of complexity results for the syntactically corresponding classes of default knowledge bases and default theories as in [56, 87] is not much meaningful in general.

Furthermore, the results in [56, 87] are quite different from ours. As shown in [56], deciding whether ψ is a consequence of a default theory $T \cup \{\phi\}$, where the classical knowledge in T is Horn and both ϕ and ψ are atoms, is co-NP-complete in the case of Horn defaults in T , while it is polynomial in the case of normal unary defaults (with or without prerequisites) in T . On the other hand, for every semantics for conditional knowledge bases that we have considered in this paper, the corresponding entailment problem $\phi \rightarrow \psi$ is either tractable or intractable (co-NP-hard) in both cases (see Figure 9 and Section 6.3).

In [43], Geffner’s approach to conditional entailment has been considered outside the conditional camp as closest to prioritized circumscription [65], which is a refinement of circumscription by introducing groups of priorities $P_1 < P_2 < \dots < P_n$ for the different predicates $P = \bigcup_i P_i$ that should be minimized. Informally, circumscription selects “preferred” models of a set of classical formulas T , which are those having a smallest extension possible on the predicates in P ; a more sophisticated notion of circumscription allows for floating extensions of some of the remaining predicates, which is needed for deriving new positive conclusions. As noted in [43], in the propositional case, the difference between prioritized circumscription and conditional entailment is that in the latter approach the priorities are entrenched in the theory, while in the former, they are explicitly assigned. Furthermore, in conditional entailment strict and defeasible knowledge is separated, while there is no such distinction in circumscription; this can be accomplished by the use of appropriate abnormality predicates, though. For further discussion, we refer to [41, 43]. To our knowledge, no thorough formal study of the semantical relationship between conditional entailment and circumscription has been carried out so far.

Our results on the complexity of conditional entailment give some useful insights into this relationship. In fact, both Geffner’s conditional entailment of $\phi \rightarrow \psi$ from a conditional knowledge base and circumscriptive inference (with or without prioritization) $CIRC(T \cup \{\phi\}) \models \psi$ are Π_2^P -complete problems (cf. [33]). Thus, polynomial time mappings between these inference problems exist, which means that efficient (polynomial-time computable) semantical embeddings among the two formalisms might exist. On the other hand, the Π_2^P -hardness of conditional entailment also applies to the Horn case, for which circumscriptive inference has complexity lowered to co-NP (see [21]; for the case with priorities, this easily follows from [21] and results in [18]). Thus, no polynomial time translation of conditional entailment into circumscription is feasible in this case (unless the polynomial time hierarchy collapses), and thus also no efficient embedding is possible. Even in the case of 1-literal Horn theories, which corresponds to Horn-Krom theories considered in [21], conditional entailment is intractable, while circumscription without priorities is tractable if no propositional atoms are fixed or vary [21]. Thus, polynomial time translations from conditional entailment to circumscription on this fragment must use fixed as well as varying atoms.

8 Conclusion

In this paper, we have established a comprehensive picture of the complexity of major approaches to default reasoning from conditional knowledge bases, namely ϵ -semantics [1, 80], systems \mathcal{Z} and \mathcal{Z}^+ [81, 47, 49], maximum entropy semantics [45, 13], lexicographic entailment [62, 6], and conditional entailment [41, 43]. For most of these approaches, merely bounds for the complexity were known, but the precise complexity was unclear.

Our work contributes on two important issues. Firstly, it provides a complete and sharp characterization of the complexities of these approaches. As we have shown, they range from the first level (co-NP) up to the second level (Π_2^P) of the polynomial hierarchy, and populate several well-known complexity classes in between. Our analysis also covers the restriction of conditional knowledge bases to the Horn and literal-Horn case, which are important from a knowledge representation perspective. Our results may help in choosing for a particular application the “suitable” semantics, given the computational cost it has attached, if computational cost is an issue. Furthermore, the results give an answer to the issue of possible efficient translations of default reasoning in the approaches to other approaches. Moreover, they unveil the computational nature of the single problems and give as a clue about the feasibility of certain algorithms. This may be important for developing implementations of the various semantics for conditional knowledge bases, which are lacking to date. To our knowledge, only prototype implementations handling small examples have been developed so far, see [12]. Notice that in related areas such as nonmonotonic reasoning, knowledge about complexity results proved extremely useful for developing efficient implementations of reasoning systems such as DeReS [23], smodels [75], and DLV [78].

Secondly, our work contributes on a refinement of the tractability/intractability frontier of default reasoning from conditional knowledge bases, by establishing new tractable cases. In particular, we have introduced q -Horn (resp., ff-Horn) conditional knowledge bases, which are meaningful extensions (resp., restrictions) of Horn conditional knowledge bases. We have shown that previous tractability results can be extended to the q -Horn case, and that, on the other hand, intractable approaches become tractable for the ff-Horn case. Our results supply polynomial algorithms for these cases, or can be easily turned into such.

Several issues remains for further work. One issue is a more fine-grained picture of the complexity of the approaches. In the present paper, we did not pay attention to possible preprocessing or fixing parameters in the input. In the literature, two important approaches have been proposed in this respect. One approach is to measure the compilability of a knowledge representation formalism, according to the frameworks proposed in [19, 44], which roughly addresses the issue whether theories in one formalism can be mapped off-line to theories in another formalism such that on-line reasoning for varying queries becomes more efficient. The other approach is the concept of fixed-parameter tractability, which deals with the effect of fixing parameters in the problem input [28, 51]. Studying the amenability of the various semantics for conditional knowledge bases to these two approaches is an intriguing issue.

Another issue is to identify further tractable cases for the various approaches. For this purpose, it would be worthwhile to investigate new classes of conditional knowledge bases; some of them may be defined in the spirit of classes for nonmonotonic formalisms as in [87, 21]. Finally, further approaches to default reasoning (for example, the recent belief function approach by Benferhat et al. [9]) may be analyzed from a complexity point of view.

Acknowledgments. We are very grateful to the anonymous referees for carefully reading through the manuscript and for many useful comments and suggestions, which helped to improve this paper.

This work has been partially supported by the Austrian Science Fund under Project N Z29-INF and by a DFG grant.

A Appendix: Proofs for Section 3

Proof of Lemma 3.2. Since KB is minimal-core and $\sigma(d) = 1$ for all $d \in D$, equation (6) reduces to equation (4). That is, the system of equations given by (6) and (7) for all $\phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$ has a

unique solution z_s^*, κ_s^* with positive z_s^* , which is given by the rankings z^*, κ^* for KB .

We now show that κ^* is robust. Suppose it were not. Then, there would exist distinct $d_1, d_2 \in D$ that have a common minimal falsifying model I . By equation (5), it then follows $\kappa^*(I) \geq z^*(d_1) + z^*(d_2)$. By equation (4), we get $z^*(d_2) \geq 1$. Since KB is minimal-core, there exists a world I' that satisfies $L \cup (D - \{d_1\})$ and that falsifies d_1 . Hence, $\kappa^*(I') = z^*(d_1) < z^*(d_1) + z^*(d_2) \leq \kappa^*(I)$. But then I is not a minimal falsifying model of d_1 under L , which is a contradiction. It follows that κ^* is robust. \square

Proof of Theorem 3.3 (continued). Let $At' = At \cup A \cup T \cup \{a\}$ be the set of atoms for KB' , where $A = \{a_{d,i} \mid d \in D, 0 \leq i \leq \pi(d)\}$ and $T = \{t_{d,i} \mid d \in D, 0 \leq i < \pi(d)\}$. For each $d = \alpha \rightarrow \beta \in D$ and $i \in \{0, \dots, \pi(d)\}$, let $D_{d,i}$ contain the following defaults:

$$a_{d,i} \wedge \alpha \rightarrow \beta \quad \text{if } i = \pi(d) \quad (12)$$

$$a_{d,i} \rightarrow \neg t_{d,i} \quad \text{if } i < \pi(d) \quad (13)$$

$$a_{d,i} \rightarrow a_{d,i-1}, a_{d,i} \rightarrow t_{d,i-1} \quad \text{if } i > 0. \quad (14)$$

Informally, the default (12) corresponds to the default $\alpha \rightarrow \beta$ from D . By the auxiliary defaults (13) and (14), it is pulled to the level $\pi(d)$ in the z -partition of the augmented set of defaults. The defaults (14) imply that verifying the default $a_{d,i} \wedge \alpha \rightarrow \beta$ for the z -partition (which requests that d has a verifying world) violates all defaults (13), which reside at levels $0, 1, \dots, i-1$.

Define now $KB' = (L', D')$, where

$$L' = \bigcup \{\alpha \vee \neg a \mid \alpha \in L\}$$

$$D' = \bigcup \{D_{d,i} \mid d \in D, i \in \{0, \dots, \pi(d)\}\}.$$

The atom a serves in L' to mask the background knowledge L for the ranking process.

It is easy to see that KB' is ε -consistent, and that the z -partition (D'_0, \dots, D'_k) of D' is given by $D'_i = \bigcup \{D_{d,i} \mid d \in D: \pi(d) \leq i\}$, for all $i = 0, \dots, k$. Hence, each default $a_{d,\pi(d)} \wedge \alpha \rightarrow \beta$ with $d = \alpha \rightarrow \beta \in D$ is assigned the value $\pi(d)$ under the default ranking Z for KB' .

As for entailment of defaults, let now ϕ' be defined by

$$\phi' = a \wedge \left(\bigwedge_{p \in A \cup T} p \right).$$

Satisfaction of a unmaskes the background knowledge L , and satisfaction of $A \cup T$ implies that all defaults (13) are false while all defaults (14) are true. Furthermore, each model of ϕ' satisfies the default $a_{d,\pi(d)} \wedge \alpha \rightarrow \beta$ in (12) iff it satisfies $\alpha \rightarrow \beta$.

Thus, it is easily seen that for any formula ϕ over At , a world $I \in \mathcal{I}_{At'}$ is a z -preferred model of $\phi \wedge \phi'$ w.r.t. KB' iff the restriction of I to At is a π -preferred model of ϕ w.r.t. KB . This implies that for any default $\phi \rightarrow \psi$ over At , it holds that (KB, π) lex_p -entails $\phi \rightarrow \psi$ iff KB' lex -entails $\phi \wedge \phi' \rightarrow \psi$. \square

Proof of Lemma 3.4. Assume first that $KB = (L, D)$ is ε -consistent. By Lemma 5.1, there exists an ordered partition (D_0, \dots, D_k) of D such that each default in D_i is tolerated under L by $\bigcup_{j=i}^k D_j$. Let \prec be any total order on D such that $d \in D_i, d' \in D_j$, and $d \prec d'$ implies $i \leq j$. Clearly, \prec is irreflexive and transitive, and thus a priority ordering on D . Moreover, each $d \in D$ is tolerated under L by $D - \{d' \in D \mid d' \prec d\}$. Thus, by Lemma 5.20, \prec is admissible with KB . That is, KB is conditionally consistent.

Conversely, assume that $KB = (L, D)$ is conditionally consistent. That is, there exists a priority ordering \prec on D that is admissible with KB . Hence, there exists some $d \in D$ that is minimal w.r.t. \prec . Moreover, as \prec restricted to $(D - \{d\}) \times (D - \{d\})$ is a priority ordering on $D - \{d\}$ that is admissible with

$(L, D - \{d\})$, the conditional knowledge base $(L, D - \{d\})$ is conditionally consistent. Thus, we can define a sequence of defaults d_1, d_2, \dots, d_n such that $\{d_1, d_2, \dots, d_n\} = D$ and that each d_i is a minimal element in $\{d_i, d_{i+1}, \dots, d_n\}$ w.r.t. to a priority ordering \prec_i admissible with $(L, \{d_i, d_{i+1}, \dots, d_n\})$. It follows that each d_i is tolerated under L by $\{d_i, d_{i+1}, \dots, d_n\}$. Hence, by Lemma 5.1, KB is ε -consistent. \square

B Appendix: Proofs for Section 5

Proof of Theorem 5.2 (continued). b) We give a log-space reduction from the P-complete problem of deciding whether a given set $P = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_n \Rightarrow \psi_n\}$ of definite Horn clauses logically implies a given atom A (see e.g. [25]). Let D denote the set of literal-Horn defaults $\{\phi_1 \rightarrow \psi_1, \dots, \phi_n \rightarrow \psi_n, \top \rightarrow \neg A\}$. We now show that P logically entails A iff the literal-Horn default knowledge base (\emptyset, D) is ε -inconsistent. Assume first that (\emptyset, D) is ε -consistent. Hence, by Lemma 5.1, D contains at least one default d that is tolerated by D . Thus, there exists an interpretation I that verifies d and that satisfies D . Hence, I is a model of $P \cup \{\neg A\}$. That is, P does not logically entail A . Conversely, assume that P does not logically entail A . Let us consider the ordered partition $(D_1, D_2) = (\{\top \rightarrow \neg A\}, \{\phi_1 \rightarrow \psi_1, \dots, \phi_n \rightarrow \psi_n\})$ of D . Then, the default $\top \rightarrow \neg A$ is tolerated by $D_1 \cup D_2$, since there exists model of $P \cup \{\neg A\}$. Moreover, each default $d \in D_2$ is tolerated by D_2 , since the interpretation I that maps each ground atom to the truth value **true** always verifies each $d \in D_2$. That is, by Lemma 5.1, (\emptyset, D) is ε -consistent. \square

Proof of Theorem 5.3. We give a log-space reduction from the P-complete problem of deciding whether a given set $P = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_n \Rightarrow \psi_n\}$ of definite Horn clauses logically implies a given atom A (see, for example, [25]). Let D be the set of literal-Horn defaults $\{\phi_1 \rightarrow \psi_1, \dots, \phi_n \rightarrow \psi_n\}$. By Theorem 3.1 and the proof of Theorem 5.2 b), P logically entails A iff the literal-Horn default knowledge base (\emptyset, D) ε -entails the literal-Horn default $\top \rightarrow A$. \square

Proof of Theorem 5.4 (continued). a) We show D^P -hardness by a polynomial transformation from the D^P -complete problem SAT-UNSAT [79]: Given two propositional formulas α and β , decide whether α is satisfiable and β is unsatisfiable. Without loss of generality, we can assume that α and β are defined over disjoint sets of atoms. We now show that α is satisfiable and β is unsatisfiable iff the default knowledge base (\emptyset, \emptyset) properly ε -entails $\alpha \rightarrow \neg\beta$. Assume first that α is satisfiable and β is unsatisfiable. By Lemma 5.1, it follows that $(\emptyset, \{\alpha \rightarrow \beta\})$ is ε -inconsistent and that $(\emptyset, \{\alpha \rightarrow \top\})$ is ε -consistent. Hence, by Theorem 3.1, the default knowledge base (\emptyset, \emptyset) properly ε -entails $\alpha \rightarrow \neg\beta$. Conversely, assume that (\emptyset, \emptyset) properly ε -entails $\alpha \rightarrow \neg\beta$. Thus, by Theorem 3.1, it follows that $(\emptyset, \{\alpha \rightarrow \beta\})$ is ε -inconsistent and that $(\emptyset, \{\alpha \rightarrow \top\})$ is ε -consistent. Hence, by Lemma 5.1, it is immediate that α is satisfiable. Moreover, since α and β are defined over disjoint basic propositions, it also follows that β is unsatisfiable.

b) We give a log-space reduction from the following P-complete problem (see, e.g., [25]): Given a set $P = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_n \Rightarrow \psi_n\}$ of definite Horn clauses and an atom A , decide whether P entails A .

Let $D = \{\phi_1 \rightarrow \psi_1, \dots, \phi_n \rightarrow \psi_n\}$. We now show that P logically entails A iff the literal-Horn default knowledge base (\emptyset, D) properly ε -entails the literal-Horn default $\top \rightarrow A$. By Theorem 3.1 and the proof of Theorem 5.3, it is sufficient to show that $(\emptyset, D \cup \{\top \rightarrow \top\})$ is ε -consistent. By Lemma 5.1, this is indeed the case, as the world I that maps each ground atom to **true** always verifies each $d \in D \cup \{\top \Rightarrow \top\}$. \square

Proof of Lemma 5.5. Let $\mathcal{D} = (D_0, \dots, D_k)$ be the z -partition of D and let $\mathcal{D}' = (D'_0, \dots, D'_i)$ be an ordered partition of D that is admissible with KB and that has the least weight w . We now show by induction on i that $D_i = D'_i$ for all $i = 0, \dots, k$.

Basis: Let $i = 0$. Let us first assume that there is some $d \in D_0 - D'_0$. Hence, $\mathcal{D}'' = (D'_0 \cup \{d\}, D'_1 - \{d\}, \dots, D'_i - \{d\})$ is an ordered partition of D that is admissible with KB and that has a weight

smaller than w . But this contradicts the assumption that \mathcal{D}' has the least weight. Let us next assume that there is some $d \in D'_0 - D_0$. Hence, d is tolerated under L by D and not contained in D_0 . But this contradicts the assumption that \mathcal{D} is the z -partition of D . Hence, it holds $D_0 = D'_0$.

Induction: Let $i > 0$. By the induction hypothesis, we get $D_j = D'_j$ for all $j = 0, \dots, i-1$. Let us first assume that there is some $d \in D_i - D'_i$. Thus, $\mathcal{D}'' = (D'_0, \dots, D'_{i-1}, D'_i \cup \{d\}, D'_{i+1} - \{d\}, \dots, D'_l - \{d\})$ is an ordered partition of D that is admissible with KB and that has a weight smaller than w . But this contradicts the assumption that \mathcal{D}' has the least weight. Let us next assume that there is some $d \in D'_i - D_i$. Hence, d is tolerated under L by $\bigcup_{j=i}^k D_j$ and not contained in D_i . But this contradicts the assumption that \mathcal{D} is the z -partition of D . Hence, it holds $D_i = D'_i$.

That is, we get $D_i = D'_i$ for all $i = 0, \dots, k$ and thus also $k = l$. \square

Proof of Theorem 5.6 (continued). a) It remains to show $P_{\parallel}^{\text{NP}}$ -hardness. We give a polynomial transformation from the following $P_{\parallel}^{\text{NP}}$ -complete problem [90]. Given the propositional formulas in CNF $\alpha_1, \dots, \alpha_m$, we are asked whether the number of tautologies among $\alpha_1, \dots, \alpha_m$ is even. Without loss of generality, we can assume that $\alpha_1, \dots, \alpha_m$ are defined on pairwise disjoint sets of variables, that α_1 and α_2 are tautologies, that α_m is not a tautology, that m is odd, and that α_{i+1} is not a tautology iff α_i is not a tautology, cf. [90]. Let

$$\begin{aligned} d_1 &= b_1 \rightarrow \neg c_1, \\ d_i &= b_i \rightarrow b_{i-1} \wedge c_{i-1} \wedge \neg c_i, \quad \text{for } i = 2, \dots, m, \end{aligned}$$

and let $D = \{d_1, \dots, d_m\}$. Informally, verifying the default d_i , $i > 1$, for determining its rank requests that the default d_{i-1} is falsified, which means that d_{i-1} must have been already ranked.

It is thus easy to see that the default ranking z for the default knowledge base (\emptyset, D) is given by $Z(d_i) = i - 1$ for all $i = 1, \dots, m$. Let ϕ and ψ be defined as

$$\begin{aligned} \phi &= (\alpha_1 \Rightarrow b_1) \wedge \dots \wedge (\alpha_m \Rightarrow b_m) \wedge c_1 \wedge \dots \wedge c_m, \\ \psi &= (b_2 \wedge \neg b_3) \vee (b_4 \wedge \neg b_5) \vee \dots \vee (b_{m-1} \wedge \neg b_m), \end{aligned}$$

respectively. We now show that the number of tautologies among $\alpha_1, \dots, \alpha_m$ is even iff (\emptyset, D) z -entails $\phi \rightarrow \psi$. Let I be any interpretation that satisfies ϕ and that satisfies b_i iff α_i is a tautology. Hence, the number of tautologies among $\alpha_1, \dots, \alpha_m$ is even iff I satisfies ψ . Moreover, for any interpretation I' that satisfies ϕ and that satisfies some b_i with α_i not being a tautology, it holds $\kappa^z(I) < \kappa^z(I')$. Let us now first assume that $I \models \psi$. Hence, $I \models b_{2i} \wedge \neg b_{2i+1}$, and we get $\kappa^z(I) = \kappa^z(\phi \wedge \psi) < \kappa^z(\phi \wedge \neg\psi)$; thus, (\emptyset, D) z -entails $\phi \rightarrow \psi$. Let us next assume that $I \not\models \psi$. Hence, $I \models \phi \wedge \neg\psi$, and we get $\kappa^z(I) = \kappa^z(\phi \wedge \neg\psi) < \kappa^z(\phi \wedge \psi)$; thus, (\emptyset, D) does not z -entail $\phi \rightarrow \psi$.

b) We give a log-space reduction from the P-complete problem of deciding whether a given set $P = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_n \Rightarrow \psi_n\}$ of definite Horn clauses logically implies a given atom A (see e.g. [25]). Let D denote the set of literal-Horn defaults $\{\phi_1 \rightarrow \psi_1, \dots, \phi_n \rightarrow \psi_n\}$. We now show that P logically entails A iff the literal-Horn default knowledge base (\emptyset, D) z -entails the literal-Horn default $\top \rightarrow A$. Since the interpretation I that maps each ground atom to the truth value **true** always verifies each $d \in D$, it holds $Z(d) = 0$ for all $d \in D$. Moreover, since I also satisfies A , we get $\kappa^z(A) = 0$. Let us now first assume that P does not logically entail A . Thus, there exists a model of $P \cup \{\neg A\}$. That is, there exists a model of $D \cup \{\neg A\}$. Hence, we get $\kappa^z(\neg A) = 0$ and thus (\emptyset, D) does not z -entail $\top \rightarrow A$. Let us next assume that P logically entails A . Thus, all models of $\neg A$ falsify at least one default from $d \in D$. Hence, we get $\kappa^z(\neg A) = 1$ and thus (\emptyset, D) z -entails $\top \rightarrow A$. \square

Proof of Theorem 5.7 (continued). a) It remains to show $\text{FP}_{\parallel}^{\text{NP}}$ -hardness. We give a polynomial transformation from the following problem, which is easily seen to be $\text{FP}_{\parallel}^{\text{NP}}$ -complete. Given propositional

CNF formulas $\alpha_1, \dots, \alpha_m$ on disjoint sets of variables, $m \geq 1$, compute the truth values a_1, \dots, a_m such that $a_i = \mathbf{true}$ iff α_i is satisfiable. Let

$$\begin{aligned} d &= \top \rightarrow \neg c, \\ d_i &= c_i \rightarrow \alpha_i \vee c, \text{ for all } i = 1, \dots, m, \end{aligned}$$

and define $D = \{d, d_1, \dots, d_m\}$. Consider the default ranking z for the default knowledge base (\emptyset, D) . We now show that $Z(d_i) = 0$ iff α_i is satisfiable. Assume first that $Z(d_i) = 0$. Hence, d_i is tolerated by D . That is, there exists an interpretation I that verifies d_i and that satisfies D . Hence, I satisfies α_i and thus α_i is satisfiable. Assume next that $Z(d_i) > 0$. Suppose that α_i is satisfiable. Then there exists an interpretation I that satisfies $c_i \wedge \alpha_i \wedge \neg c$ and all $\neg c_j$ with $j \neq i$. Thus, I verifies d_i and satisfies D . That is, d_i is tolerated by D . But this contradicts the assumption $Z(d_i) > 0$. Hence, α_i must be unsatisfiable. This proves the $\text{FP}_{\parallel}^{\text{NP}}$ -hardness part.

We remark that it is unknown whether the z -ranking is computable with $O(\log n)$ many calls to an NP-oracle, where n is the input size. In fact, since $\text{FP}^{\text{NP}}[O(\log n)] \subseteq \text{FP}_{\parallel}^{\text{NP}}$ holds and the inclusion is believed to be strict [53], the $\text{FP}_{\parallel}^{\text{NP}}$ -hardness results suggests that computing the z -ranking is not in $\text{FP}^{\text{NP}}[O(\log n)]$.

b) We give a log-space reduction from the following problem.

Lemma B.1 *Given a set $P = \{\phi_1 \Rightarrow \psi_1, \dots, \phi_m \Rightarrow \psi_m\}$ of definite Horn clauses on the atoms $At = \{x_1, \dots, x_n\}$, computing $S = \{A \in At \mid P \models A\}$ is FP-complete.*

Proof (sketch). Obviously, the problem is in FP. Hardness for FP follows e.g. from the proof of P-hardness of deciding whether $P \models A$ holds for a given set P of definite Horn clauses and atom A in [25]. The log-space reduction there encodes the computation of a generic polynomial-time deterministic Turing machine M into this problem, such that $S = \{A \mid P \models A\}$ contains (among others) atoms reflecting the tape contents of M when it halts. From S , the output of M is easily extracted in log-space. \diamond

Let D be the set of literal-Horn defaults $\{\phi_1 \rightarrow \psi_1, \dots, \phi_m \rightarrow \psi_m, c_1 \rightarrow \neg x_1, \dots, c_n \rightarrow \neg x_n\}$, where c_1, \dots, c_n are pairwise distinct new variables. We now show that x_i is logically entailed by P iff the default ranking z for the literal-Horn default knowledge base (\emptyset, D) assigns $c_i \rightarrow \neg x_i$ the value 1. Let the world I be defined by $I(x_i) = \mathbf{true}$ and $I(c_i) = \mathbf{false}$ for all $i = 1, \dots, n$. It is easy to see that I satisfies D and verifies all $\phi_i \rightarrow \psi_i$ with $i = 1, \dots, n$. Hence, z maps all $\phi_i \rightarrow \psi_i$ with $i = 1, \dots, n$ to the value 0. To determine the value of the other defaults under z , let us now first assume that x_i is not logically entailed by P . That is, there is a model of $P \cup \{\neg x_i\}$. This model can be extended to a model I of D that verifies $c_i \rightarrow \neg x_i$ by defining $I(c_i) = \mathbf{true}$ and $I(c_j) = \mathbf{false}$ for all $i = 1, \dots, n$ with $i \neq j$. Hence, $c_i \rightarrow \neg x_i$ is tolerated by D and thus assigned the value 0 under z . Assume next that x_i is logically entailed by P . That is, there is no model of $P \cup \{\neg x_i\}$. Hence, $c_i \rightarrow \neg x_i$ is not tolerated by D . Let the world I be defined by $I(x_i) = \mathbf{false}$ and $I(c_i) = \mathbf{true}$ for all $i = 1, \dots, n$. It is now easy to see that I verifies all $c_j \rightarrow \neg x_j$ with $j = 1, \dots, n$. Thus, all $c_i \rightarrow \neg x_i$ such that x_i is logically entailed by P are assigned 1 under z . \square

Proof of Theorem 5.8. a) The result follows from the proof of Theorem 5.6 a). In detail, the reduction given in Theorem 5.6 a) also applies to the case where z is given in advance, as the ranking z for the constructed default knowledge base (\emptyset, D) is actually given in advance by $Z(d_i) = i - 1$ for all $i = 1, \dots, m$.

b) Similar to a), referring to the proof of Theorem 5.6 b) in place of Theorem 5.6 a). \square

Proof of Theorem 5.9 (continued). Let a denote the integer that has the complement of $I_{\text{lms}}(\phi)$ as binary representation. For example, $a = 10 = 2^3 + 2^1$ for $I_{\text{lms}}(\phi) = 0101$. We first show that the default ranking

z^+ is given as follows. For all $i = 1, \dots, n$:

$$\begin{aligned} z^+(d_i^+) &= 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1 \\ z^+(d_i^-) &= 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1 + 2^{k+n-i}. \end{aligned} \quad (15)$$

That is, the binary representation of $z^+(d_i^+)$ has the first $i - 1$ bits of $I_{\text{lms}}(\phi)$, padded with 0's to n bits, and k trailing bits added that account for cumulative extra costs $i - 1$. The binary representation of $z^+(d_i^-)$ is similar, but has the bit for x_i (i -th from left) set to 1.

Let us thus assume that $z^+(d_i^+)$ and $z^+(d_i^-)$ for all $i = 1, \dots, n$ are given by (15). For each world I , let $\kappa^+(I)$ be defined by (3). To show that (15) actually defines the solutions, we must show that (2) holds for all defaults $d = \phi \rightarrow \psi \in D$.

Let $i \geq 1$ and let I be an interpretation that satisfies $\neg c_1 \wedge \dots \wedge \neg c_{i-1} \wedge c_i \wedge \phi$ and that coincides on the variables x_1, \dots, x_{i-1} with $I_{\text{lms}}(\phi)$. Hence, I verifies d_i^+ and d_i^- . Moreover, I satisfies all defaults d_j^+ and d_j^- with $j > i$. Finally, for all $j = 1, \dots, i - 1$, exactly one default among d_j^+ and d_j^- is falsified by I . More precisely, if $I \models x_j$, then d_j^+ is falsified by I . Otherwise, d_j^- is falsified by I . That is, the rank of the falsified default d_j among d_j^+ and d_j^- is given as follows:

$$z^+(d_j) = \begin{cases} 2^{k+n-j+1} \cdot (a \operatorname{div} 2^{n-j+1}) + j - 1 & \text{if } I \models x_j \\ 2^{k+n-j+1} \cdot (a \operatorname{div} 2^{n-j+1}) + j - 1 + 2^{k+n-j} & \text{if } I \not\models x_j. \end{cases}$$

For all $j = 1, \dots, i - 1$, we have $I \models x_j$ iff $I_{\text{lms}}(\phi) \models x_j$. The latter is equivalent to a 's bit for x_j (the j -th bit from left), denoted $a[j]$, being 0. Hence, we obtain the following:

$$\begin{aligned} z^+(d_j) &= \begin{cases} 2^{k+n-j+1} \cdot (a \operatorname{div} 2^{n-j+1}) + j - 1 + a[j] \cdot 2^{k+n-j} & \text{if } I_{\text{lms}}(\phi) \models x_j \\ 2^{k+n-j+1} \cdot (a \operatorname{div} 2^{n-j+1}) + j - 1 + a[j] \cdot 2^{k+n-j} & \text{if } I_{\text{lms}}(\phi) \not\models x_j \end{cases} \\ &= 2^{k+n-j} \cdot (a \operatorname{div} 2^{n-j}) + j - 1. \end{aligned}$$

Since $z^+(d_j)$ is maximal for $j = i - 1$, we get:

$$\kappa^+(I) = 1 + \max_{d \in D: I \not\models d} z^+(d) = 1 + \max_{j \in \{1, \dots, i-1\}} z^+(d_j) = 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1.$$

Hence, for every $d_i = \phi_i \rightarrow \psi_i \in \{d_i^+, d_i^-\}$, we get:

$$\kappa^+(\phi_i \wedge \psi_i) = \min_{J \in \mathcal{I}_{At}: J \models_{\forall} d_i} \kappa^+(J) \leq \kappa^+(I) = 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1.$$

Let I' be any other interpretation that satisfies $\neg c_1 \wedge \dots \wedge \neg c_{i-1} \wedge \phi$. Thus, for all $j = 1, \dots, i - 1$, exactly one default among d_j^+ and d_j^- is falsified by I' . Assume that I' does not coincide on the variables x_1, \dots, x_{i-1} with $I_{\text{lms}}(\phi)$. Hence, there must be some $j \leq i - 1$ such that x_j is true in I but false in I' , which means $I \not\models d_j^+$ and $I \models d_j^-$ while $I' \models d_j^+$ and $I' \not\models d_j^-$. Hence, we get:

$$\kappa^+(I') \geq 2^{k+n-j+1} \cdot (a \operatorname{div} 2^{n-j+1}) + j + 2^{k+n-j}.$$

Moreover, as $j < i$, we get:

$$2^{k+n-j+1} \cdot (a \operatorname{div} 2^{n-j+1}) \geq 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}).$$

Since $i - 1 - j < n \leq 2^k$, it holds $j + 2^{k+n-j} > i - 1$. Hence, it follows:

$$2^{k+n-j+1} \cdot (a \operatorname{div} 2^{n-j+1}) + j + 2^{k+n-j} > 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1.$$

That is, we get $\kappa^+(I) < \kappa^+(I')$ for any other such I' . For all $d_i = \phi_i \rightarrow \psi_i \in \{d_i^+, d_i^-\}$, it thus follows:

$$\kappa^+(\phi_i \wedge \psi_i) = \min_{J \in \mathcal{I}_{At}: J \models_v d_i} \kappa^+(J) = 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1.$$

But this shows that (2) holds for all $d = \phi \rightarrow \psi \in D$. Thus, (15) describes the actual z^+ -ranking of (\emptyset, D) .

Let I be an interpretation that satisfies $\neg c_1 \wedge \dots \wedge \neg c_n \wedge \phi$ and that coincides on the variables x_1, \dots, x_n with $I_{\text{ms}}(\phi)$. Let I' be any other interpretation that satisfies $\neg c_1 \wedge \dots \wedge \neg c_n \wedge \phi$ and that does not coincide on the variables x_1, \dots, x_n with $I_{\text{ms}}(\phi)$. By a line of argumentation similar to the one just pursued, it is easy to see that $\kappa^+(I) < \kappa^+(I')$.

Let us now assume that x_n is true in $I_{\text{ms}}(\phi)$. It then follows that

$$\kappa^+(\neg c_1 \wedge \dots \wedge \neg c_n \wedge \phi \wedge x_n) < \kappa^+(\neg c_1 \wedge \dots \wedge \neg c_n \wedge \phi \wedge \neg x_n).$$

That is, (KB, σ) z^+ -entails d at strength 0. Let us next assume that x_n is false in $I_{\text{ms}}(\phi)$. It then follows:

$$\kappa^+(\neg c_1 \wedge \dots \wedge \neg c_n \wedge \phi \wedge \neg x_n) < \kappa^+(\neg c_1 \wedge \dots \wedge \neg c_n \wedge \phi \wedge x_n).$$

That is, (KB, σ) does not z^+ -entail d at strength 0. \square

Proof of Theorem 5.11 (continued). (i) We first show that KB is ε -consistent. Let the ordered partition (D_0, \dots, D_k) of D be defined by $k = \max(c_1, \dots, c_m)$ and $D_i = \{d_{i,j} \mid d_{i,j} \in D\}$ for all $i = 1, \dots, k$. Consider any default $d_{i,j} \in D_i$. Let $I_{i,j}$ be a world such that $I_{i,j} \models \alpha_j \Rightarrow \beta_j$, $I_{i,j} \models a_{l,j}$, for all $l \leq i$, $I_{i,j} \models b_{i,j}$, $I_{i,j} \models t_{i,j}$, for all $l \leq i - 1$, and $I_{i,j}$ does not satisfy all remaining atoms $a_{i',j'}$, $b_{i',j'}$, and $t_{i',j'}$. Clearly such a world $I_{i,j}$ exists. It is now easy to see that $I_{i,j}$ verifies $d_{i,j}$ and that $I_{i,j} \models L \cup \bigcup_{l=i}^k D_l$. Hence, the default $d_{i,j}$ is tolerated under L by $\bigcup_{l=i}^k D_l$. Thus, by Lemma 5.1, KB is ε -consistent.

(ii) It is easy to see that KB is minimal-core. Indeed, the world I such that $I \models a_{i,j}$ and $I \not\models A$ for any other atomic proposition A , falsifies the default $d_{i,j}$ while it satisfies $L \cup (D - \{d_{i,j}\})$.

(iii) We now show by induction on $i = 0, \dots, k$ that the default ranking z^* is given by $z^*(d_{i,j}) = 2^i$ for all $d_{i,j} \in D$. Hence, in particular, $z^*(d_{c_j,j}) = 2^{c_j}$ for all $j = 1, \dots, m$.

Basis: Let $i = 0$. Since (D_0, \dots, D_k) is a partition of D that is admissible with KB , we get $z^*(d_{i,j}) = 1$ for all $d_{i,j} \in D_0$. *Induction:* Let $i > 0$. Let us consider any $d_{i,j} \in D_i$. Recall that the world $I_{i,j}$ described above satisfies $L \cup \bigcup_{l=i}^k D_l$ and, moreover, verifies $d_{i,j}$; furthermore, $I_{i,j}$ falsifies every $d_{l,j}$ such that $l \leq i - 1$. By the induction hypothesis, $z^*(d_{l,j}) = 2^l$ for every $l \leq i - 1$, and thus $\kappa^*(I_{i,j}) = \sum_{l=0}^{i-1} 2^l = 2^i - 1$. Hence, it follows $z^*(d_{i,j}) \leq 2^i$. On the other hand, every default $d_{l,j}$ where $l < i$ is falsified in every world I that satisfies L and verifies $d_{i,j}$. Hence, it follows $z^*(d_{i,j}) = 2^i$, which concludes the induction.

(iv) We finally show that $I \models \alpha_r \Rightarrow \beta_r$ holds in every I with maximum weight $\sum_{I \models \alpha_i \Rightarrow \beta_i} w_i$ iff KB z^* -entails d . We need some preparation as follows.

Let I be any world such that: (a) $I \models L \cup \{\phi\}$, (b) I is a maximum weight world under C , and (c) $I \models b_{c_j,j}$ iff $I \models \alpha_j \Rightarrow \beta_j$. Let I' be any world such that (a) holds but either (b) or (c) does not hold. We now show that $\kappa^*(I) < \kappa^*(I')$. It is easy to see that both I and I' falsify all defaults $d_{i,j}$, for $j \leq m$ and $i \leq c_j - 1$. Let I'' result from I' by redefining $b_{c_j,j}$ to $I'' \models b_{c_j,j}$ iff $I' \models \alpha_j \Rightarrow \beta_j$, for all $j \leq m$. Then,

$I'' \models L$, and no default $d_{i,j}$ that is satisfied by I'' is violated by $d_{i,j}$. Hence, it follows $\kappa^*(I'') \leq \kappa^*(I')$. Furthermore, I and I'' satisfy $d_{c_j,j}$ iff they satisfy $\alpha_j \Rightarrow \beta_j$, for all $j \leq m$. Hence, we get:

$$\begin{aligned}
\kappa^*(I) &= \sum_{d \in D} z^*(d) - \sum_{d \in D, I \models d} z^*(d) \\
&= \sum_{d \in D} z^*(d) - \sum_{j \in \{1, \dots, m\}, I \models \alpha_j \Rightarrow \beta_j} 2^{c_j} \\
&\leq \sum_{d \in D} z^*(d) - \sum_{j \in \{1, \dots, m\}, I'' \models \alpha_j \Rightarrow \beta_j} 2^{c_j} \\
&= \sum_{d \in D} z^*(d) - \sum_{d \in D, I'' \models d} z^*(d) \\
&= \kappa^*(I'')
\end{aligned}$$

It follows $\kappa^*(I) \leq \kappa^*(I'') \leq \kappa^*(I')$. Moreover, since (b) and (c) hold for I , while either (b) or (c) does not hold for I' , either $\kappa^*(I) < \kappa^*(I'')$ or $\kappa^*(I'') < \kappa^*(I')$, and thus $\kappa^*(I) < \kappa^*(I')$.

Let us now first assume that $I \models \alpha_r \Rightarrow \beta_r$ holds for every maximum weight world I under C . Consider any such I . Hence, there exists a world I' such that $I' \models x_i$ iff $I \models x_i$, $I' \models b_{c_j,j}$ iff $I \models \alpha_j \Rightarrow \beta_j$, and $I' \models L \cup \{\phi \wedge \psi\}$. Moreover, there is no maximum weight world I'' under C such that $I'' \models L \cup \{\phi \wedge \neg\psi\}$ and $I'' \models b_{c_j,j}$ iff $I'' \models \alpha_j \Rightarrow \beta_j$. It follows $\kappa^*(\phi \wedge \psi) < \kappa^*(\phi \wedge \neg\psi)$, and thus KB z^* -entails d .

Let us next assume that $I \not\models \alpha_r \Rightarrow \beta_r$ for some maximum weight world under C . Hence, there exists a world I' such that $I' \models x_i$ iff $I \models x_i$, $I' \models b_{c_j,j}$ iff $I \models \alpha_j \Rightarrow \beta_j$, and $I' \models L \cup \{\phi \wedge \neg\psi\}$. It follows $\kappa^*(\phi \wedge \psi) \geq \kappa^*(\phi \wedge \neg\psi)$, and thus KB does not z^* -entail d . \square

Proof of Theorem 5.12. We give a polynomial transformation from a suitable variant of the problem used in the reduction in the proof of Theorem 5.11, which is FP^{NP} -complete: Given a set of weighted Horn clauses $C = \{\alpha_1 \Rightarrow \beta_1, \dots, \alpha_m \Rightarrow \beta_m\}$ on n variables x_1, \dots, x_n , where each $\alpha_i \Rightarrow \beta_i$ is satisfiable and has weight $w_i = 2^{c_i}$, where $c_i \geq 0$ is a nonnegative integer, compute the weight w of a maximum weight world I under C , that is, $w = \max_{I \in \mathcal{I}_{At}} \sum_{I \models \alpha_i \Rightarrow \beta_i} w_i$. FP^{NP} -hardness of this problem can be established by a suitable adaptation of proofs in [58].

We slightly extend KB in the proof of Theorem 5.11 as follows. We introduce new atoms a^* and b^* , and the following set of literal-Horn clauses L^* and the literal-Horn default d^* :

$$\begin{aligned}
L^* &= \{b^* \Rightarrow a_{i,j} \mid 1 \leq j \leq m, 0 \leq i \leq c_j\} \cup \{b^* \Rightarrow t_{i,j} \mid 1 \leq j \leq m, 0 \leq i < c_j\} \\
d^* &= a^* \rightarrow b^*.
\end{aligned}$$

By a line of argumentation similar to the one in the proof of Theorem 5.11, it is easy to see that the extended conditional knowledge base $(L \cup L^*, D \cup \{d^*\})$ is ε -consistent and minimal-core. Moreover, its ranking z^* assigns all defaults $d_{i,j}$ the value 2^i and the default d^* the value $\sum_{j=1}^m \sum_{i=0}^{c_j} 2^i - w$. Consequently, the weight w of a maximum weight world under C is given by $w = 2 \cdot \sum_{j=1}^m 2^{c_j} - m - z^*(d^*)$, which can be easily computed from z^* . This proves the result. \square

Proof of Lemma 5.14. The claim is proved by induction on $i = 1, \dots, l$ as follows.

Basis: For $i = 1$, we get $z_s^*(d_1) = \sigma(d_1) \leq s$. *Induction:* Let $i > 1$. By the induction hypothesis, it holds that $|z_s^*(d_j)| \leq s \cdot 2^{j-1}$ for all $j = 1, \dots, i-1$. Hence, we get:

$$|z_s^*(d_i)| \leq \sigma(d_i) + |\text{minv}(d_i) - \text{minf}(d_i)| \leq s + \sum_{j=1}^{i-1} s \cdot 2^{j-1} = s + s \cdot (2^{i-1} - 1) = s \cdot 2^{i-1}. \quad \square$$

Proof of Theorem 5.18 (continued). Define $At = \{a, b, c, e\} \cup \{s_i, x_i, y_i \mid 1 \leq i \leq n\}$ and

$$F_i = \{a \rightarrow x_i, b \rightarrow y_i, c \wedge x_i \rightarrow \neg y_i, x_i \wedge y_i \rightarrow \neg s_i, \top \rightarrow s_i\} \quad \text{for } 1 \leq i \leq n.$$

It is now easy to verify that $I(x_i) = I(\neg y_i)$ holds for every world I that satisfies $a \wedge b \wedge c$ and a maximal number of defaults in F_i . That is, we can use F_i to express the relationships $x_i \wedge y_i \Rightarrow \perp$ and $\top \Rightarrow x_i \vee y_i$.

Let α_j^* be the Horn clause that is obtained from α_j by replacing each positive literal x_i by the new negative literal $\neg y_i$. Define $D = \bigcup_{i=0}^{n+2} D_i$, where the D_i are as follows:

$$\begin{aligned} D_i &= \{e \rightarrow x_{n-i}\} && \text{for all } i = 0, \dots, n-1 \\ D_n &= \{\alpha_1^*, \dots, \alpha_m^*\} \\ D_{n+1} &= \bigcup \{F_i \mid i = 1, \dots, n\} \\ D_{n+2} &= \{\top \rightarrow a, \top \rightarrow b, \top \rightarrow c, \top \rightarrow e\}. \end{aligned}$$

The priority assignment π on D is given by $\pi(d) = i$ for all $i = 0, \dots, n+2$ and $d \in D_i$. Finally, we define

$$d = \top \rightarrow x_n.$$

Observe that $KB = (\emptyset, D)$ and d are literal-Horn.

We now show that $I_{\text{lms}}(\alpha) \models x_n$ iff (KB, π) *lex*_P-entails d . It is sufficient to show that, for any preferred world I , its restriction to $X = \{x_1, \dots, x_n\}$, denoted by $I|X$, coincides with I_{lms} , and that, on the other hand, $I_{\text{lms}}(\alpha)$ can be extended to such an I .

Assume first that I is a preferred world. Hence, $I \models D_{n+2}$. Furthermore, I satisfies a maximal number of defaults in D_{n+1} ; this implies that $I(x_i) = I(\neg y_i)$, for all $i = 1, \dots, n$. Since α is satisfiable, preferredness of I then implies that $I \models D_n$ and $I|X \models \alpha$. The sets of defaults D_0, \dots, D_{n-1} then ensure that $I|X$ indeed coincides with $I_{\text{lms}}(\alpha)$.

Conversely, let I' be the world such that $I'(x_i) = I_{\text{lms}}(x_i)$, and $I'(y_i) = I'(\neg x_i)$ for all $i = 1, \dots, n$, and $I'(p) = \mathbf{true}$, for any other atom p . It is now easy to see that I' is a preferred world. \square

Proof of Theorem 5.19. a) By Theorem 5.7 a), computing the default ranking z for KB is in $\text{FP}_{\parallel}^{\text{NP}}$. Recall now that KB *lex*-entails d iff (KB, z) *lex*_P-entails d . As deciding the latter is in P^{NP} (see the discussion at the beginning of Section 5.4), deciding whether KB *lex*-entails d is also in P^{NP} .

b) We show that the P^{NP} -hard problem in Theorem 5.18, which is more general, is reducible to this problem. From the proof of Theorem 5.18, we may assume that every default $d \in D$ of $KB = (\emptyset, D)$ there has a verifying world. Thus, by Theorem 3.3 and Observation 3.1, we obtain such a reduction. \square

Proof of Lemma 5.20. Assume first that \prec is admissible with KB , and consider $d \in D$. Admissibility of \prec implies that D_d is under L not in conflict with d ; that is, d is tolerated under L by D_d .

Conversely, assume that every $d \in D$ is tolerated under L by D_d . Suppose that \prec is not admissible with KB . That is, some $D' \subseteq D$ is under L in conflict with some $d \in D$, and D' contains no default d' with $d' \prec d$. Hence, $D' \subseteq D_d$. Since D_d tolerates d under L , also D' tolerates d under L . But this contradicts the fact that D' is under L in conflict with d . Hence, \prec is admissible with KB . \square

The rest of the proof of Theorem 5.22 will make use of the following lemma shown by Geffner [41].

Lemma B.2 (Geffner [41]) *Let $KB = (L, D)$ be a conditional knowledge base. A default $\phi \rightarrow \psi$ is conditionally entailed by KB iff ψ is satisfied in every \prec -preferred model of $L \cup \{\phi\}$ of every minimal priority ordering \prec admissible with KB .*

Here, minimality is in terms of set inclusion, where \prec is viewed as set of pairs $\{(I, J) \mid I \prec J\}$.

We are now ready to complete the proof that conditional entailment is Π_2^P -hard for the literal-Horn case.

Proof of Theorem 5.22 (continued). Let $At = A_y \cup A_x \cup \{a\}$, where $A_y = \{a_i, b_i, y_i, y'_i \mid 1 \leq i \leq m\}$, $A_x = \{c_j, d_j, e_j, f_j, x_j, x'_j \mid 1 \leq j \leq n\}$. We define $KB = (L, D)$ as follows:

$$\begin{aligned} L &= L_1 \cup L_2 \cup L_3 \cup L_4, \\ D &= \bigcup_{i=1}^m D_{1,i} \cup \bigcup_{j=1}^n D_{2,j}, \end{aligned}$$

where the sets of Horn clauses $L_i, i = 1, 2, 3, 4$ and the default sets $D_{1,i}, D_{2,j}$ are defined as follows:

$$\begin{aligned} L_1 &= \{\alpha_1^* \vee \neg a, \dots, \alpha_l^* \vee \neg a\}, \\ L_2 &= \{\neg y_i \vee \neg y'_i \vee \neg a \mid i = 1, \dots, m\}, \\ L_3 &= \{\neg x_j \vee \neg x'_j \vee \neg a \mid j = 1, \dots, n\}, \\ L_4 &= \{x_j \Rightarrow \neg f_k, x'_j \Rightarrow \neg f_k \mid j, k = 1, \dots, n\}, \end{aligned}$$

where $\alpha_1^*, \dots, \alpha_l^*$ is obtained from $\alpha_1, \dots, \alpha_l$ by replacing the positive literals y_i and x_j by the negative literals $\neg y'_i$ and $\neg x'_j$, respectively; and

$$\begin{aligned} D_{1,i} &= \{a_i \rightarrow y_i, b_i \rightarrow y'_i\}, \\ D_{2,j} &= \{c_j \wedge d_j \rightarrow x_j, c_j \wedge e_j \rightarrow x'_j, c_j \rightarrow f_j\}. \end{aligned}$$

Finally, the default $d = \phi \rightarrow \psi$ is defined by

$$\begin{aligned} \phi &= a \wedge \left(\bigwedge_{i=1}^m (a_i \wedge b_i) \right) \wedge \left(\bigwedge_{j=1}^n (c_j \wedge d_j \wedge e_j) \right) \\ \psi &= \neg f_1. \end{aligned}$$

The set of all defaults that are conditionally entailed by KB is defined with respect to all priority orderings on D that are admissible with KB . By Lemma B.2, we can restrict our attention to all minimal priority orderings on D that are admissible with KB .

We note that every priority ordering \prec on D admissible with KB contains the following pairs:

$$c_j \rightarrow f_j \prec c_j \wedge d_j \rightarrow x_j, \tag{16}$$

$$c_j \rightarrow f_j \prec c_j \wedge e_j \rightarrow x'_j, \quad \text{for all } j = 1, \dots, n. \tag{17}$$

This is immediate from the observation that each set $\{c_j \rightarrow f_j\}$ tolerates under L neither $c_j \wedge d_j \rightarrow x_j$ nor $c_j \wedge e_j \rightarrow x'_j$.

Let \prec^* be the priority ordering on D that is given by exactly all pairs in (16) and (17). It is easy to see that each $d \in D$ is tolerated under L by $D - \{d' \in D \mid d' \prec^* d\}$; thus, by Lemma 5.20, \prec^* is admissible with KB . This means that \prec^* is the least (that is, unique minimal) priority ordering on D admissible with KB . Applying Lemma B.2, the default $d = \phi \rightarrow \psi$ is conditionally entailed by KB iff ψ is satisfied in every \prec^* -preferred model of $L \cup \{\phi\}$.

We are now ready to show that Φ evaluates to true iff KB conditionally entails d .

(\Leftarrow) Let us first assume that Φ evaluates to false. Hence, there exists a mapping $f : \{y_1, \dots, y_m\} \rightarrow \{\perp, \top\}$ such that the formula $\alpha = (\alpha_1 \wedge \dots \wedge \alpha_l) [y_1/f(y_1), \dots, y_m/f(y_m)]$ is unsatisfiable. Let I be the world such that (i) $I(y_i) = I(\neg y'_i) = \mathbf{true}$ iff $f(y_i) = \top$, for all $i = 1, \dots, m$; (ii) $I(x_j) = I(x'_j) = \mathbf{false}$, for all $j = 1, \dots, n$; and (iii) $I(p) = \mathbf{true}$ for any other atom p . Since each α_i^* contains at least one negative literal from $\neg x_1, \dots, \neg x_n, \neg x'_1, \dots, \neg x'_n$, we have $I \models L_1$. Clearly, $I \models L_2 \cup L_3$, and also $I \models L_4, I \models \phi$, and $I \models \neg\psi$. Hence, $I \models L \cup \{\phi, \neg\psi\}$. Suppose that $J \models L \cup \{\phi\}$ such that $J \prec^* I$. We now show that no such J exists, which proves that I is a \prec^* -preferred model of $L \cup \{\phi\}$. Since J is a model of L_2 and ϕ , it follows that J cannot satisfy both y_i and y'_i , for $i = 1, \dots, m$. As \prec^* does not define any preference between the defaults in $D_{1,i}$, for $i = 1, \dots, m$, it follows that J cannot falsify any default in $D_{1,i}$ that is not falsified by I . Hence, J and I falsify exactly the same defaults in $D_{1,i}$, for $i = 1, \dots, m$. That is, J must coincide on y_i and y'_i with I , for $i = 1, \dots, m$. Further, since I and J are different, J must either satisfy some x_j or x'_j , or falsify some f_j . The clauses in L_4 then have in the former case the effect that $J \models \neg f_j$ for all $j = 1, \dots, n$. Since in the latter case the default $c_j \rightarrow f_j$ from $D_{2,j}$ is violated by J but not by I , the fact that $J \prec^* I$ implies that either $c_j \wedge d_j \rightarrow x_j$ or $c_j \wedge e_j \rightarrow x'_j$, which are both violated by I , must be satisfied by J . Thus, either $J \models x_j$ or $J \models x'_j$ holds. The clauses in L_3 imply that only one can hold, and thus, $J(x_j) = J(\neg x'_j)$ holds, for all $j = 1, \dots, n$. Clearly, J is a model of $\alpha_1^* \wedge \dots \wedge \alpha_l^*$, as it satisfies L_1 and ϕ . Since $J(y_i) = J(\neg y'_i) = \mathbf{true}$ iff $f(y_i) = \top$, for all $i = 1, \dots, m$, it thus follows that J restricted to $\{x_1, \dots, x_n\}$ is a model of α . That is, α is satisfiable, which is a contradiction. Thus, J does not exist, and I is a \prec^* -preferred model of $L \cup \{\phi\}$. This shows that KB does not conditionally entail d .

(\Rightarrow) Conversely, let us assume that KB does not conditionally entail d . That is, there exists a \prec^* -preferred model I of $L \cup \{\phi\}$ such that $I \not\models \psi$, which means $I \models f_1$. The clauses in L_4 imply that $I(x_j) = I(x'_j) = \mathbf{false}$, for all $j = 1, \dots, n$; preferredness of I implies that $I(f_j) = \mathbf{true}$ (as the default $c_j \rightarrow f_j$ will be satisfied), for all $j = 1, \dots, n$. Let the mapping $f : \{y_1, \dots, y_m\} \rightarrow \{\perp, \top\}$ be defined by $f(y_i) = \top$ iff $I \models a_i \rightarrow y_i$. We now show that $\alpha = (\alpha_1 \wedge \dots \wedge \alpha_l) [y_1/f(y_1), \dots, y_m/f(y_m)]$ is unsatisfiable. Towards a contradiction, suppose there exists a truth assignment I' to x_1, \dots, x_n that satisfies α . Let I'' be the world that coincides on all y_i, y'_i with I , sets $I''(x_j) = I''(\neg x'_j) = I'(x_j)$, and sets $I''(f_j) = \mathbf{false}$, for all $j = 1, \dots, n$, and sets $I''(p) = \mathbf{true}$ for every other atom p . Then $I'' \models L \cup \{\phi\}$. The worlds I'' and I falsify exactly the same defaults in $D_{1,i}$, for all $i = 1, \dots, m$. Moreover, I satisfies $c_j \rightarrow f_j$ and falsifies both $c_j \wedge d_j \rightarrow x_j$ and $c_j \wedge e_j \rightarrow x'_j$, while I'' falsifies $c_j \rightarrow f_j$ and satisfies either $c_j \wedge d_j \rightarrow x_j$ or $c_j \wedge e_j \rightarrow x'_j$, for all $j = 1, \dots, n$. This shows $I'' \prec^* I$. Hence, I is not a \prec^* -preferred model of $L \cup \{\phi\}$, which is a contradiction. It follows that α is unsatisfiable, and thus Φ evaluates to false. \square

C Appendix: Proofs for Section 6

Proof of Theorem 6.9. We give a polynomial transformation from the complement of the NP-complete *one-in-three 3sat* problem for positive literals [40]: Given a set of variables $X = \{x_1, \dots, x_k\}$ and a set C of clauses $x_{1,1} \vee x_{1,2} \vee x_{1,3}, \dots, x_{n,1} \vee x_{n,2} \vee x_{n,3}$ such that $x_{i,j} \in X$ for all $i \in \{1, \dots, n\}$ and $j \in \{1, 2, 3\}$, decide whether there exists a truth assignment I_{oit} over X that satisfies exactly one variable in each clause.

We construct KB and d as in the statement of the theorem such that KB does not z^* -entail d iff such a truth assignment I_{oit} exists. Let the set of atoms be defined as $At = \{a, b, x_1, \dots, x_k\} \cup \{a_{i,j} \mid i \in \{1, \dots, n\}, j \in \{1, 2, 3\}\}$.

$j \in \{1, \dots, 4\}$. Let $KB = (L, D)$ and d be defined by:

$$\begin{aligned} L &= \{a \Rightarrow a_{i,j} \mid i \in \{1, \dots, n\}, j \in \{1, \dots, 4\}\} \cup \\ &\quad \{x_{i,j} \Rightarrow \neg x_{i,k} \mid i \in \{1, \dots, n\}, j, k \in \{1, \dots, 3\}, j < k\} \cup \\ &\quad \{x_{i,j} \Rightarrow \neg b \mid i \in \{1, \dots, n\}, j \in \{1, \dots, 3\}\}, \\ D &= \{a_{i,j} \rightarrow x_{i,j} \mid i \in \{1, \dots, n\}, j \in \{1, \dots, 3\}\} \cup \\ &\quad \{a_{i,4} \rightarrow b \mid i \in \{1, \dots, n\}\}, \\ d &= a \rightarrow b. \end{aligned}$$

It is easy to verify that KB is ε -consistent and minimal-core. Furthermore, it is easy to see that the z -partition of D is given by (D) . Hence, it follows $z^*(d) = 1$ for all $d \in D$.

We next show that every z^* -preferred model I of $L \cup \{a\}$ falsifies exactly $3n$ defaults in D . For every $i \in \{1, \dots, n\}$, the atoms $x_{i,1}, x_{i,2}, x_{i,3}, b$ are mutually exclusive under L . Hence, I falsifies at least three defaults among $a_{i,1} \rightarrow x_{i,1}, a_{i,2} \rightarrow x_{i,2}, a_{i,3} \rightarrow x_{i,3}, a_{i,4} \rightarrow b$, and thus at least $3n$ defaults in D . Moreover, a model J of $L \cup \{a\}$ that falsifies exactly $3n$ defaults in D is always given by $J(x) = \mathbf{false}$ for all $x \in X$, $J(a) = \mathbf{true}$, $J(b) = \mathbf{true}$, and $J(a_{i,j}) = \mathbf{true}$ for all $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, 4\}$.

We finally show that KB does not z^* -entail d iff there exists a truth assignment I_{oit} over X that satisfies exactly one variable in each clause from C . Assume first that such a truth assignment I_{oit} exists. Hence, the world I that is defined by $I|_X = I_{\text{oit}}$, $I(a) = \mathbf{true}$, $I(b) = \mathbf{false}$, and $I(a_{i,j}) = \mathbf{true}$ for all $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, 4\}$ is a model of $L \cup \{a\}$ that falsifies exactly $3n$ defaults in D . That is, I is a z^* -preferred model of $L \cup \{a\}$ with $I(b) = \mathbf{false}$. Moreover, the world J , as defined above, is a z^* -preferred model of $L \cup \{a\}$ with $I(b) = \mathbf{true}$. This shows that KB does not z^* -entail d .

Conversely, assume that KB does not z^* -entail d . That is, there exists a z^* -preferred model I of $L \cup \{a\}$ such that $I(b) = \mathbf{false}$. Hence, I falsifies exactly $3n$ defaults in D . More precisely, for every $i \in \{1, \dots, n\}$, it falsifies exactly two defaults among $a_{i,1} \rightarrow x_{i,1}, a_{i,2} \rightarrow x_{i,2}, a_{i,3} \rightarrow x_{i,3}$. This shows that $I|_X$ is a truth assignment over X that satisfies exactly one variable in each clause from C . \square

Proof of Theorem 6.11. The claims follow immediately from the proof of Theorem 6.9. Consider KB and d constructed in its proof. For the proof of b), let the priority assignment π on KB be defined by $\pi(d) = 1$ for all $d \in D$. For the proof of c), note that \emptyset is the least priority ordering on D admissible with KB . It now follows that KB does not lex -entail d (resp., (KB, π) does not lex_p -entail d), and KB does not conditionally entail d iff there is a truth assignment I_{oit} over X that satisfies exactly one variable in each clause from C . \square

Proof of Lemma 6.12. Let the world I^* be defined by $I^*(b) = I(b)$ for all active atoms $b \in At$ and by $I^*(b) = \mathbf{false}$ for all inactive atoms $b \in At$.

We first show that I^* is a model of L . Suppose not. That is, there exists some Horn clause $\alpha \Rightarrow \beta \in L$ such that $I^* \not\models \alpha \Rightarrow \beta$. That is, $I^*(\alpha) = \mathbf{true}$ and $I^*(\beta) = \mathbf{false}$. Hence, $I(\alpha) = \mathbf{true}$. Since $I \models \alpha \Rightarrow \beta$, we get $I(\beta) = \mathbf{true}$. Thus, all atoms in α are active, and β is an atom that is inactive. Since β is an atom, it holds that $\alpha \Rightarrow \beta \in L^+$. But this contradicts β being inactive. This shows that I^* is a model of L .

Clearly, as I^* coincides with I on all active atoms $b \in At$, it follows $I^*(\gamma) = I(\gamma)$ for all active classical formulas γ , and $I^* \models \delta$ iff $I \models \delta$ for all active defaults $\delta \in D$. This shows (i) and (iii).

We finally show (ii). Towards a contradiction, suppose that I^* falsifies some inactive literal-Horn default $\alpha \rightarrow \beta \in D$. That is, $I^*(\alpha) = \mathbf{true}$ and $I^*(\beta) = \mathbf{false}$. Thus, all atoms in α are active. Since $\alpha \rightarrow \beta$ is inactive, the literal β is inactive. Since $I^*(\beta) = \mathbf{false}$, it follows that β is an atom, and thus $\alpha \rightarrow \beta \in D^+$. But this contradicts β being inactive. Hence, I^* satisfies all inactive defaults in D . \square

Proof of Lemma 6.15. As $L \cup \{\alpha\}$ (resp., $L \cup \{\alpha^*\}$) and β are defined over disjoint sets of atoms, it follows that $L \cup \{\alpha \wedge \beta\}$ is satisfiable, iff both $L \cup \{\alpha\}$ and β are satisfiable, iff both $L \cup \{\alpha^*\}$ and β are satisfiable, iff $L \cup \{\alpha^* \wedge \beta\}$ is satisfiable. This shows that $\kappa_s^*(\alpha \wedge \beta) = \infty$ iff $\kappa_s^*(\alpha^* \wedge \beta) = \infty$.

Suppose now that $L \cup \{\alpha \wedge \beta\}$ is satisfiable. Since each model I of $\alpha^* \wedge \beta$ is also a model of $\alpha \wedge \beta$, we get $\kappa_s^*(\alpha \wedge \beta) \leq \kappa_s^*(\alpha^* \wedge \beta)$. Let I be any minimal model of $L \cup \{\alpha \wedge \beta\}$ w.r.t. κ_s^* . Let the world J be defined as $I_\alpha^* \cup I|_{At - At_\alpha}$. Clearly, J is a model of $\alpha^* \wedge \beta$. Since $I(b) = \mathbf{false}$ implies $J(b) = \mathbf{false}$, for all $b \in At_\alpha$, it follows $\{d \in D \mid I \not\models d\} \supseteq \{d \in D \mid J \not\models d\}$. As J is a model of $L \cup \{\alpha \wedge \beta\}$, and I is a minimal model of $L \cup \{\alpha \wedge \beta\}$ w.r.t. κ_s^* , we get $\{d \in D \mid I \not\models d\} = \{d \in D \mid J \not\models d\}$. Hence, we get $\kappa_s^*(I) = \kappa_s^*(J)$. It thus follows $\kappa_s^*(\alpha \wedge \beta) = \kappa_s^*(\alpha^* \wedge \beta)$. \square

Proof of Lemma 6.16. Obviously, $\kappa_s^*(\alpha^* \wedge \beta) = \infty$ iff $\sum_{i \in \{1, \dots, n\}} \kappa_{s,i}^*(\alpha^* \wedge \beta_i) = \infty$.

Now, let $L \cup \{\alpha^* \wedge \beta\}$ be satisfiable. As $\kappa_s^*(I) = \sum_{i \in \{1, \dots, n\}} \kappa_{s,i}^*(I)$ for all worlds I over At , we get:

$$\kappa_s^*(\alpha^* \wedge \beta) \geq \sum_{i \in \{1, \dots, n\}} \kappa_{s,i}^*(\alpha^* \wedge \beta). \quad (18)$$

For each $i \in \{1, \dots, n\}$, let I_i be a minimal model of $\alpha^* \wedge \beta$ w.r.t. $\kappa_{s,i}^*$. Since each D_i is defined over $At_\alpha \cup At_i$, we can define a model I of $\alpha^* \wedge \beta$ by $I = I_\alpha^* \cup I_1|_{At_1} \cup \dots \cup I_n|_{At_n}$. Thus, it follows:

$$\kappa_s^*(\alpha^* \wedge \beta) = \sum_{i \in \{1, \dots, n\}} \kappa_{s,i}^*(\alpha^* \wedge \beta). \quad (19)$$

Since β_1, \dots, β_n are defined over the pairwise disjoint sets of atoms At_1, \dots, At_n , it follows:

$$\kappa_{s,i}^*(\alpha^* \wedge \beta) = \kappa_{s,i}^*(\alpha^* \wedge \beta_i) \text{ for all } i \in \{1, \dots, n\}. \quad (20)$$

The claim now follows from (19) and (20). \square

Proof of Lemma 6.17. Assume first that $L \cup \{\alpha_1 \wedge \beta_1\}$ and $L \cup \{\alpha_2 \wedge \beta_2\}$ have a common minimal model I w.r.t. κ_s^* . Hence, both $L \cup \{\alpha_1 \wedge \alpha_2\}$ and $\beta_1 \wedge \beta_2$ are satisfiable. Moreover, it always holds $\kappa_s^*(\alpha_1 \wedge \beta_1) \leq \kappa_s^*(\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2)$ and $\kappa_s^*(\alpha_2 \wedge \beta_2) \leq \kappa_s^*(\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2)$. Since I is a model of $L \cup \{\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2\}$, it thus follows $\kappa_s^*(\alpha_1 \wedge \beta_1) = \kappa_s^*(\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2)$ and $\kappa_s^*(\alpha_2 \wedge \beta_2) = \kappa_s^*(\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2)$.

Conversely, as both $L \cup \{\alpha_1 \wedge \alpha_2\}$ and $\beta_1 \wedge \beta_2$ are satisfiable, also $L \cup \{\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2\}$ is satisfiable. Let I be any minimal model of $L \cup \{\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2\}$ w.r.t. κ_s^* . Clearly, I is a common model of $L \cup \{\alpha_1 \wedge \beta_1\}$ and $L \cup \{\alpha_2 \wedge \beta_2\}$. Moreover, $\kappa_s^*(\alpha_1 \wedge \beta_1) = \kappa_s^*(\alpha_1 \wedge \alpha_2 \wedge \beta_1 \wedge \beta_2) = \kappa_s^*(\alpha_2 \wedge \beta_2)$ implies that I is even a common *minimal* model of $L \cup \{\alpha_1 \wedge \beta_1\}$ and $L \cup \{\alpha_2 \wedge \beta_2\}$ w.r.t. κ_s^* . \square

Proof of Lemma 6.19. Clearly, $\kappa_s^*(\gamma) = \infty$ iff $L \cup \{\gamma\}$ is unsatisfiable iff $\widehat{\kappa}_s^*(\gamma) = \infty$.

Assume now that $L \cup \{\gamma\}$ is satisfiable. Since $\kappa_s^*(I) \geq \widehat{\kappa}_s^*(I)$ for all worlds I , it follows $\kappa_s^*(\gamma) \geq \widehat{\kappa}_s^*(\gamma)$. Let I be any model of $L \cup \{\gamma\}$ such that $\widehat{\kappa}_s^*(\gamma) = \widehat{\kappa}_s^*(I)$. By Lemma 6.12, there exists a model I^* of L such that $I^*(\gamma) = I(\gamma) = \mathbf{true}$, and I^* satisfies $D - \widehat{D}$ and the same defaults in \widehat{D} as I . Hence, we get $\kappa_s^*(I^*) \geq \kappa_s^*(\gamma)$ and $\kappa_s^*(I^*) = \widehat{\kappa}_s^*(I)$. It thus follows $\widehat{\kappa}_s^*(\gamma) \geq \kappa_s^*(\gamma)$. Hence, $\kappa_s^*(\gamma) = \widehat{\kappa}_s^*(\gamma)$. \square

Proof of Lemma 6.20. We first show that $\widehat{\kappa}_{s,j}^*, \widehat{\kappa}_{s,j}^*$ is a solution of (11). Consider any default $\gamma \rightarrow \delta \in \widehat{D}_j$. The classical formulas $\gamma \wedge \delta$ and $\gamma \wedge \neg \delta$ are of the form $\alpha \wedge \beta_1$ and $\alpha \wedge \beta_2$, respectively, where α is either \top or a conjunction of atoms from At_a , and both β_1 and β_2 are conjunctions of literals over At_j . As $\gamma \rightarrow \delta$ is

active, also $\alpha \wedge \beta_1$ and $\alpha \wedge \beta_2$ are active. By Lemmata 6.15, 6.16, and 6.19, it thus follows for $l \in \{1, 2\}$:

$$\begin{aligned}
\kappa_s^*(\alpha \wedge \beta_l) &= \widehat{\kappa}_s^*(\alpha \wedge \beta_l) \\
&= \widehat{\kappa}_s^*(\alpha^* \wedge \beta_l) \\
&= \widehat{\kappa}_{s,j}^*(\alpha^* \wedge \beta_l) + \sum_{i \in \{1, \dots, n\} - \{j\}} \widehat{\kappa}_{s,i}^*(\alpha^*) \\
&= \widehat{\kappa}_{s,j}^*(\alpha \wedge \beta_l) + \sum_{i \in \{1, \dots, n\} - \{j\}} \widehat{\kappa}_{s,i}^*(\alpha).
\end{aligned} \tag{21}$$

Hence, $\kappa_s^*(\gamma \wedge \neg \delta) = \sigma(\gamma \rightarrow \delta) + \kappa_s^*(\gamma \wedge \delta)$ iff $\widehat{\kappa}_{s,j}^*(\gamma \wedge \neg \delta) = \widehat{\sigma}_j(\gamma \rightarrow \delta) + \widehat{\kappa}_{s,j}^*(\gamma \wedge \delta)$. This shows that $\widehat{\kappa}_{s,j}^*, \widehat{\kappa}_{s,j}^*$ is a solution of (11). Its uniqueness follows from the robustness of $\widehat{\kappa}_{s,j}^*$, which we prove next.

Suppose that $\widehat{\kappa}_{s,j}^*$ is not robust. Then, by Lemma 6.17, there are two distinct $\gamma_1 \rightarrow \delta_1, \gamma_2 \rightarrow \delta_2 \in \widehat{D}_j$ such that $\widehat{\kappa}_{s,j}^*(\gamma_1 \wedge \neg \delta_1) = \widehat{\kappa}_{s,j}^*(\gamma_1 \wedge \gamma_2 \wedge \neg \delta_1 \wedge \neg \delta_2) = \widehat{\kappa}_{s,j}^*(\gamma_2 \wedge \neg \delta_2)$, and both $L \cup \{\gamma_1 \wedge \gamma_2\}$ and $\neg \delta_1 \wedge \neg \delta_2$ are satisfiable. We next show that for all atoms $c \in \widehat{At}_a$, it holds that $L \cup \{\gamma_1\} \models c$ iff $L \cup \{\gamma_2\} \models c$. Suppose the contrary. Without loss of generality, assume that γ_2 contains some atom $c \in \widehat{At}_a$ such that $L \cup \{\gamma_1\} \not\models c$. Let I be a minimal model of $L \cup \{\gamma_1 \wedge \gamma_2 \wedge \neg \delta_1 \wedge \neg \delta_2\}$, and thus also of $L \cup \{\gamma_1 \wedge \neg \delta_1\}$. Let the world I' be defined by $I' = I_{\gamma_1}^* \cup I|_{At - \widehat{At}_a}$. Then, I' is a model of $L \cup \{\gamma_1 \wedge \neg \delta_1\}$, it satisfies all the defaults in \widehat{D}_j that are satisfied by I , and it satisfies $\gamma_2 \rightarrow \delta_2$, which is falsified by I . It thus follows $\widehat{\kappa}_{s,j}^*(I') < \widehat{\kappa}_{s,j}^*(I)$. But this contradicts I being a minimal model of $L \cup \{\gamma_1 \wedge \neg \delta_1\}$. This shows that for all atoms $c \in \widehat{At}_a$, it holds that $L \cup \{\gamma_1\} \models c$ iff $L \cup \{\gamma_2\} \models c$. Hence, we can assume that $\gamma_1 \wedge \neg \delta_1, \gamma_1 \wedge \gamma_2 \wedge \neg \delta_1 \wedge \neg \delta_2$, and $\gamma_2 \wedge \neg \delta_2$ are of the form $\alpha \wedge \beta_1, \alpha \wedge \beta_1 \wedge \beta_2$, and $\alpha \wedge \beta_2$, respectively, where α is either \top or a conjunction of atoms from \widehat{At}_a , and both β_1 and β_2 are conjunctions of literals over At_j . By (21), it then follows $\kappa_s^*(\gamma_1 \wedge \neg \delta_1) = \kappa_s^*(\gamma_1 \wedge \gamma_2 \wedge \neg \delta_1 \wedge \neg \delta_2) = \kappa_s^*(\gamma_2 \wedge \neg \delta_2)$. But, by Lemma 6.17, this contradicts κ_s^* being robust. It thus follows that $\widehat{\kappa}_{s,j}^*$ is robust. \square

Proof of Lemma 6.22. We first show that every π -preferred model of $L \cup \{\phi\}$ satisfies $D - \widehat{D}$. Suppose the contrary. That is, there exists a π -preferred model I of $L \cup \{\phi\}$ such that $I \not\models d$ for some $d \in D - \widehat{D}$. Clearly, ϕ is active. Thus, by Lemma 6.12, there exists a model I^* of L such that $I^*(\phi) = I(\phi) = \mathbf{true}$, and I^* satisfies $D - \widehat{D}$ and the same defaults in \widehat{D} as I . This shows that I^* is a model of $L \cup \{\phi\}$ that is π -preferable to I . But this contradicts I being a π -preferred model of $L \cup \{\phi\}$. This shows that every π -preferred model of $L \cup \{\phi\}$ satisfies $D - \widehat{D}$.

a) Let I be a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$. Clearly, ϕ is active. Thus, by Lemma 6.12, there exists a model I^* of L such that $I^*(\phi) = I(\phi) = \mathbf{true}$, $I^*(\gamma) = I(\gamma)$, and I^* satisfies $D - \widehat{D}$ and the same defaults in \widehat{D} as I . We now prove that I^* is a π -preferred model of $L \cup \{\phi\}$. Suppose the contrary. That is, there exists a π -preferred model J of $L \cup \{\phi\}$ that is π -preferable to I^* . By the argumentation above, J satisfies $D - \widehat{D}$. But then J is $\widehat{\pi}$ -preferable to I , which contradicts I being a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$. This shows that I^* is a π -preferred model of $L \cup \{\phi\}$.

b) Let J be a π -preferred model of $L \cup \{\phi\}$. By the argumentation above, J satisfies $D - \widehat{D}$. Suppose that J is not a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$. That is, there is a model I of $L \cup \{\phi\}$ that is $\widehat{\pi}$ -preferable to J . Clearly, ϕ is active. Thus, by Lemma 6.12, there exists a model I^* of L such that $I^*(\phi) = I(\phi) = \mathbf{true}$, and I^* satisfies $D - \widehat{D}$ and the same defaults in \widehat{D} as I . But then I^* is π -preferable to J , which contradicts J being a π -preferred model of $L \cup \{\phi\}$. This shows that J is a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$. \square

Proof of Lemma 6.23. a) Let I be a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$. Let the world J be defined as $I_\phi^* \cup I|_{At - \widehat{At}_a}$. Clearly, J is a model of $L \cup \{\phi^*\}$ such that $J|_{At - \widehat{At}_a} = I|_{At - \widehat{At}_a}$. Since $I|_{\widehat{At}_a}$ is a superset of

$J|_{\widehat{At}_a}$, it follows $\{d \in D \mid I \not\models d\} \supseteq \{d \in D \mid J \not\models d\}$. Since J is a model of $L \cup \{\phi\}$, and I is a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$, we get $\{d \in D \mid I \not\models d\} = \{d \in D \mid J \not\models d\}$. Suppose now that there exists a model J' of $L \cup \{\phi^*\}$ that is $\widehat{\pi}$ -preferable to J . Hence, J' is a model of $L \cup \{\phi\}$ that is $\widehat{\pi}$ -preferable to I . But this contradicts I being a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$. Thus, J is a $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$.

b) Let I be a $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$. In particular, I is a model of $L \cup \{\phi\}$. Suppose now that I is not a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$. That is, there exists a model I' of $L \cup \{\phi\}$ that is $\widehat{\pi}$ -preferable to I . Let the world J be defined as $I_\phi^* \cup I'|_{At - \widehat{At}_a}$. Clearly, J is a model of $L \cup \{\phi^*\}$. Moreover, since $I'|_{\widehat{At}_a}$ is a superset of $J|_{\widehat{At}_a}$, it follows $\{d \in D \mid I' \not\models d\} \supseteq \{d \in D \mid J \not\models d\}$. Thus, J is a model of $L \cup \{\phi^*\}$ that is $\widehat{\pi}$ -preferable to I . But this contradicts I being a $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$. This shows that I is a $\widehat{\pi}$ -preferred model of $L \cup \{\phi\}$. \square

Proof of Lemma 6.24. a) For each $i \in \{1, \dots, n\}$ with $i \neq j$, let I_i be a $\widehat{\pi}_i$ -preferred model of $L \cup \{\phi^*\}$. Define the world J as $I_\phi^* \cup I_1|_{At_1} \cup \dots \cup I_n|_{At_n}$. Clearly, J is a model of $L \cup \{\phi^*\}$ with $J|_{At_j} = I_j|_{At_j}$. Assume now that J is not a $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$. That is, there exists a model J' of $L \cup \{\phi^*\}$ that is $\widehat{\pi}$ -preferable to J . Thus, there is some $i \in \{1, \dots, n\}$ such that J' is $\widehat{\pi}_i$ -preferable to J . That is, J' is $\widehat{\pi}_i$ -preferable to I_i . But this contradicts I_i being a $\widehat{\pi}_i$ -preferred model of $L \cup \{\phi^*\}$. Thus, J is a $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$.

b) Let I be a $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$. Suppose now that I is not a $\widehat{\pi}_j$ -preferred model of $L \cup \{\phi^*\}$. That is, there exists a model I' of $L \cup \{\phi^*\}$ that is $\widehat{\pi}_j$ -preferable to I . Let the world J be defined as $I_\phi^* \cup I'|_{At_j} \cup I|_{At - (\widehat{At}_a \cup At_j)}$. Clearly, J is a model of $L \cup \{\phi^*\}$. Moreover, J is $\widehat{\pi}$ -preferable to I . But this contradicts I being a $\widehat{\pi}$ -preferred model of $L \cup \{\phi^*\}$. Thus, I is a $\widehat{\pi}_j$ -preferred model of $L \cup \{\phi^*\}$. \square

References

- [1] E. W. Adams. *The Logic of Conditionals*, volume 86 of *Synthese Library*. D. Reidel, Dordrecht, Netherlands, 1975.
- [2] C. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [3] F. Bacchus, A. Grove, J. Y. Halpern, and D. Koller. From statistical knowledge bases to degrees of belief. *Artificial Intelligence*, 87:75–143, 1996.
- [4] C. Bell, A. Nerode, R. Ng, and V. S. Subrahmanian. Mixed integer programming methods for computing non-monotonic deductive databases. *Journal of the ACM*, 41(6):1178–1215, 1994.
- [5] R. Ben-Eliyahu. NP-complete problems in optimal Horn clauses satisfiability. Technical Report R-158, UCLA, Computer Science Dept., 1991.
- [6] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proc. 13th International Joint Conference on Artificial Intelligence (IJCAI-93)*, Chambéry, France, pp. 640–645. Morgan Kaufmann, 1993.
- [7] S. Benferhat, D. Dubois, and H. Prade. Representing default rules in possibilistic logic. In *Proc. 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, Cambridge, MA, pp. 673–684. Morgan Kaufmann, 1992.
- [8] S. Benferhat, D. Dubois, and H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. *Artificial Intelligence*, 92(1–2):259–276, 1997.
- [9] S. Benferhat, A. Saffiotti, and P. Smets. Belief functions and default reasoning. *Artificial Intelligence*. To appear.

- [10] E. Boros, Y. Crama, and P. L. Hammer. Polynomial-time inference of all valid implications for Horn and related formulae. *Annals of Mathematics and Artificial Intelligence*, 1:21–32, 1990.
- [11] E. Boros, P. L. Hammer, and X. Sun. Recognition of q-Horn formulae in linear time. *Discrete Applied Mathematics*, 55:1–13, 1994.
- [12] R. A. Bourne. *Default reasoning using maximum entropy and variable strength defaults*. PhD thesis, University of London, 1999.
- [13] R. A. Bourne and S. Parsons. Maximum entropy and variable strength defaults. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, Stockholm, Sweden, pp. 50–55. Morgan Kaufmann, 1999.
- [14] C. Boutilier. Conditional logics of normality as modal systems. In *Proc. 8th National Conference on Artificial Intelligence (AAAI-90)*, Boston, MA, pp. 594–599. AAAI Press/MIT Press, 1990.
- [15] C. Boutilier. Conditional logics of normality: A modal approach. *Artificial Intelligence*, 68:87–154, 1994.
- [16] C. Boutilier. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence*, 68:33–85, 1994.
- [17] S. Buss and L. Hay. On truth-table reducibility to SAT. *Information and Computation*, 91:86–102, 1991.
- [18] M. Cadoli. The complexity of model checking for circumscriptive formulae. *Information Processing Letters*, 44:113–118, 1992.
- [19] M. Cadoli, F. M. Donini, P. Liberatore, and M. Schaerf. Feasibility and unfeasibility of off-line processing. In *Proc. Fourth Israeli Symposium on Theory of Computing and Systems (ISTCS'96)*, Jerusalem, Israel, pp. 100–109. IEEE Computer Society, 1996.
- [20] M. Cadoli, A. Giovanardi, and M. Schaerf. An algorithm to evaluate quantified Boolean formulae. In *Proc. 16th National Conference on AI (AAAI-98)*, Madison, WI, pp. 262–267. AAAI Press/MIT Press, 1998.
- [21] M. Cadoli and M. Lenzerini. The complexity of propositional closed world reasoning and circumscription. *Journal of Computer and System Sciences*, 43:165–211, 1994.
- [22] C. Cayrol, M.-C. Lagasque-Schiex, and T. Schiex. Nonmonotonic reasoning: From complexity to algorithms. *Annals of Mathematics and Artificial Intelligence*, 22(3–4):207–236, 1998.
- [23] P. Cholewiński, V. W. Marek, A. Mikitiuk, and M. Truszczyński. Computing with default logic. *Artificial Intelligence*, 112(2–3):105–147, 1999.
- [24] M. Dalal. Investigations into a theory of knowledge base revision. In *Proc. 6th National Conference on Artificial Intelligence (AAAI-88)*, St. Paul, MN, pp. 475–479. AAAI Press/MIT Press, 1988.
- [25] E. Dantsin, T. Eiter, G. Gottlob, and A. Voronkov. Complexity and expressive power of logic programming. In *Proc. 12th IEEE International Conference on Computational Complexity (CCC-97)*, Ulm, Germany, pp. 82–101, 1997.
- [26] J. Delgrande. A first-order conditional logic for prototypical properties. *Artificial Intelligence*, 33(1):105–130, 1987.
- [27] J. Delgrande. On first-order conditional logics. *Artificial Intelligence*, 105(1–2):105–137, 1998.
- [28] R. G. Downey, M. F. Fellows, and M. R. Fellows. *Parameterized Complexity*. Springer, 1998.
- [29] D. Du, J. Gu, and P. M. Pardalos, editors. *Satisfiability Problem: Theory and Applications*, volume 35 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. American Mathematical Society, 1997.
- [30] D. Dubois and H. Prade. Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proc. 12th International Joint Conference on Artificial Intelligence (IJCAI-91)*, Sydney, Australia, pp. 419–424. Morgan Kaufmann, 1991.

- [31] D. Dubois and H. Prade. Conditional objects as non-monotonic consequence relationships. *IEEE Transactions on Systems, Man and Cybernetics*, 24:1724–1740, 1994.
- [32] U. Egly, T. Eiter, H. Tompits, and S. Woltran. Solving advanced reasoning tasks using quantified Boolean formulas. In *Proc. 18th National Conference on Artificial Intelligence (AAAI-00)*, Austin, TX, pp. 417–422. AAAI Press/MIT Press, 2000.
- [33] T. Eiter and G. Gottlob. Propositional circumscription and extended closed world reasoning are Π_2^P -complete. *Theoretical Computer Science*, 114(2):231–245, 1993. Addendum 118:315.
- [34] T. Eiter, G. Gottlob, and N. Leone. Abduction from logic programs: Semantics and complexity. *Theoretical Computer Science*, 189(1-2):129–177, 1997.
- [35] N. Friedman and J. Y. Halpern. Plausibility measures and default reasoning. *Journal of the ACM*. To appear.
- [36] N. Friedman and J. Y. Halpern. On the complexity of conditional logics. In *Proc. 4th International Conference on Principles of Knowledge Representation and Reasoning (KR-94)*, Bonn, Germany, pp. 202–213. Morgan Kaufmann, 1994.
- [37] D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors. *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume III: Nonmonotonic Reasoning and Uncertain Reasoning. Clarendon Press, Oxford, 1994.
- [38] D. M. Gabbay and P. Smets, editors. *Handbook on Defeasible Reasoning and Uncertainty Management Systems*. Kluwer Academic, Dordrecht, Netherlands, 1998.
- [39] P. Gärdenfors and H. Rott. Belief revision. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume IV: Epistemic and Temporal Reasoning, pp. 35–132. Clarendon Press, Oxford, 1995.
- [40] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, New York, 1979.
- [41] H. Geffner. *Default Reasoning: Causal and Conditional Theories*. MIT Press, Cambridge, MA, 1992.
- [42] H. Geffner. High probabilities, model preference and default arguments. *Mind and Machines*, 2:51–70, 1992.
- [43] H. Geffner and J. Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence*, 53(2–3):209–244, 1992.
- [44] G. Gogic, H. Kautz, C. H. Papadimitriou, and B. Selman. The comparative linguistics of knowledge representation. In *Proc. 14th International Joint Conference on Artificial Intelligence (IJCAI-95)*, Montréal, Canada, pp. 862–869. Morgan Kaufmann, 1995.
- [45] M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3):220–232, 1993.
- [46] M. Goldszmidt and J. Pearl. On the consistency of defeasible databases. *Artificial Intelligence*, 52(2):121–149, 1991.
- [47] M. Goldszmidt and J. Pearl. System Z^+ : A formalism for reasoning with variable strength defaults. In *Proc. 9th National Conference on Artificial Intelligence (AAAI-91)*, Anaheim, CA, pp. 399–404. AAAI Press/MIT Press, 1991.
- [48] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Proc. 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, Cambridge, MA, pp. 661–672. Morgan Kaufmann, 1992.
- [49] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84(1–2):57–112, 1996.
- [50] G. Gottlob. Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2(3):397–425, 1992.

- [51] G. Gottlob, F. Scarcello, and M. Sideri. Fixed parameter complexity in AI and nonmonotonic reasoning. In *Proc. 5th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR-99)*, El Paso, TX, LNCS 1730, pp. 1–18, Springer, 1999.
- [52] T. Janhunen. On the intertranslatability of autoepistemic, default and priority logics, and parallel circumscription. In *Proc. European Workshop on Logics in Artificial Intelligence (JELIA '98)*, Dagstuhl, Germany, LNCS 1489, pp. 216–232. Springer, 1998.
- [53] B. Jenner and J. Toran. Computing functions with parallel queries to NP. *Theoretical Computer Science*, 141:175–193, 1995.
- [54] D. S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 2, pp. 67–161. MIT Press, Cambridge, MA, 1990.
- [55] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:253–294, 1991.
- [56] H. Kautz and B. Selman. Hard problems for simple default logics. *Artificial Intelligence*, 49:243–279, 1991.
- [57] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 14(1):167–207, 1990.
- [58] M. W. Krentel. The complexity of optimization problems. *Journal of Computer and System Sciences*, 36(3):490–509, 1988.
- [59] P. Lamarre. S4 as the conditional logic of nonmonotonicity. In *Proc. 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR-91)*, Cambridge, MA, pp. 357–367. Morgan Kaufmann, 1991.
- [60] P. Lamarre. A promenade from monotonicity to non-monotonicity following a theorem prover. In *Proc. 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)*, Cambridge, MA, pp. 572–580. Morgan Kaufmann, 1992.
- [61] D. Lehmann. What does a conditional knowledge base entail? In *Proc. 1st International Conference on Principles of Knowledge Representation and Reasoning (KR-89)*, Toronto, Canada, pp. 212–222. Morgan Kaufmann, 1989.
- [62] D. Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15(1):61–82, 1995.
- [63] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, 1992.
- [64] P. Liberatore. The complexity of belief update. In *Proc. 15th International Joint Conference on Artificial Intelligence (IJCAI-97)*, Nagoya, Japan, pp. 68–73. Morgan Kaufmann, 1997.
- [65] V. Lifschitz. Computing circumscription. In *Proc. 9th International Joint Conference on Artificial Intelligence (IJCAI-85)*, Los Angeles, CA, pp. 121–127. Morgan Kaufmann, 1985.
- [66] W. Marek and M. Truszczyński. *Nonmonotonic Logics – Context-Dependent Reasoning*. Springer, 1993.
- [67] J. McCarthy. Circumscription – a form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.
- [68] D. McDermott. Non-monotonic logic II: Nonmonotonic modal theories. *Journal of the ACM*, 29:33–57, 1982.
- [69] D. McDermott and J. Doyle. Non-monotonic logic I. *Artificial Intelligence*, 13:41–72, 1980.
- [70] R. C. Moore. Semantical considerations on nonmonotonic logics. *Artificial Intelligence*, 25:75–94, 1985.
- [71] B. Nebel. Belief revision and default reasoning: Syntax-based approaches. In *Proc. 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR-91)*, Cambridge, MA, pp. 417–428. Morgan Kaufmann, 1991.

- [72] B. Nebel. Base revision operators and schemes: Semantics, representation, and complexity. In *Proc. 11th European Conference on Artificial Intelligence (ECAI-94)*, Amsterdam, Netherlands, pp. 341–345, 1994.
- [73] B. Nebel. How hard is it to revise a belief base? In D. Gabbay and P. Smets, editors, *Handbook on Defeasible Reasoning and Uncertainty Management Systems*, volume III: Belief Change, pp. 77–145. Kluwer Academic, Dordrecht, Netherlands, 1998.
- [74] I. Niemelä. On the decidability and complexity of autoepistemic reasoning. *Fundamenta Informaticae*, 17:117–155, 1992.
- [75] I. Niemelä and P. Simons. Smodels - an implementation of the stable model and well-founded semantics for normal logic programs. In *Proc. 4th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR-97)*, Dagstuhl, Germany, LNCS/LNAI 1265, pp. 420–429. Springer, 1997.
- [76] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence*, 57(2–3):227–270, 1992.
- [77] T. Eiter and G. Gottlob. The complexity of logic-based abduction. *Journal of the ACM*, 42(1):3–42, 1995.
- [78] T. Eiter, N. Leone, C. Mateis, G. Pfeifer, and F. Scarcello. The KR system dlv: Progress report, comparisons, and benchmarks. In *Proc. 6th International Conference on Principles of Knowledge Representation and Reasoning (KR-98)*, Trento, Italy, pp. 406–417. Morgan Kaufmann, 1998.
- [79] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, Reading, MA, 1994.
- [80] J. Pearl. Probabilistic semantics for nonmonotonic reasoning: A survey. In *Proc. 1st International Conference on Principles of Knowledge Representation and Reasoning (KR-89)*, Toronto, Canada, pp. 505–516. Morgan Kaufmann, 1989.
- [81] J. Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proc. 3rd Conference on Theoretical Aspects of Reasoning about Knowledge (TARK-90)*, Pacific Grove, CA, pp. 121–135. Morgan Kaufmann, 1990.
- [82] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [83] J. Rintanen. Improvements to the evaluation of quantified Boolean formulae. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, Stockholm, Sweden, pp. 1192–1197. Morgan Kaufmann, 1999.
- [84] A. Selman. A taxonomy of complexity classes of functions. *Journal of Computer and System Sciences*, 48:357–381, 1994.
- [85] Y. Shoham. A semantical approach to nonmonotonic logics. In *Proc. 2nd IEEE Symposium on Logic in Computer Science (LICS-87)*, Ithaca, New York, pp. 275–279, 1987.
- [86] W. Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In W. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume 2, pp. 105–134. Reidel, Dordrecht, Netherlands, 1988.
- [87] J. Stillman. It’s not my default: The complexity of membership problems in restricted propositional default logic. In *Proc. 8th National Conference on Artificial Intelligence (AAAI-90)*, Boston, MA, pp. 571–579. AAAI Press/MIT Press, 1990.
- [88] J. Stillman. The complexity of propositional default logic. In *Proc. 10th National Conference on Artificial Intelligence (AAAI-92)*, San Jose, CA, pp. 794–799. AAAI Press/MIT Press, 1992.
- [89] T. Vossen, M. Ball, A. Lotem, and D. S. Nau. On the use of integer programming models in AI planning. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, Stockholm, Sweden, pp. 304–309. Morgan Kaufmann, 1999.
- [90] K. Wagner. Bounded query classes. *SIAM Journal of Computing*, 19(5):833–846, 1990.