I N F S Y S

R E S E A R C H

R E P O R T

INSTITUT FÜR INFORMATIONSSYSTEME

ABTEILUNG WISSENSBASIERTE SYSTEME

COMPLEXITY RESULTS FOR DEFAULT REASONING FROM CONDITIONAL KNOWLEDGE BASES

Thomas Eiter and Thomas Lukasiewicz

INFSYS RESEARCH REPORT 1843-99-10

NOVEMBER 1999

Institut für Informationssysteme
Abtg. Wissensbasierte Systeme
Technische Universität Wien
Favoritenstraße 11
A-1040 Wien, Austria
Tel: +43-1-58801-18405
Fax: +43-1-58801-18493
sek@kr.tuwien.ac.at
www.kr.tuwien.ac.at

TECHNISCHE UNIVERSITÄT WIEN

INFSYS RESEARCH REPORT
INFSYS RESEARCH REPORT 1843-99-10, NOVEMBER 1999

COMPLEXITY RESULTS FOR DEFAULT REASONING FROM CONDITIONAL KNOWLEDGE BASES

(PRELIMINARY REPORT)

Thomas Eiter [1] and Thomas Lukasiewicz [1]

Abstract. Conditional knowledge bases have been proposed as belief bases that include defeasible rules (also called defaults) of the form "$\phi \to \psi$", 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 algorithmic properties only partial results are known so far. In this paper, we fill these gaps and draw a precise picture of the complexity of default reasoning from conditional knowledge bases: Given a conditional knowledge base KB and a default $\phi \to \psi$, does KB entail $\phi \to \psi$? 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). We consider the general propositional case as well as syntactic restrictions (in particular, to Horn and literal-Horn conditional knowledge bases). Furthermore, we analyze the effect of precomputing rankings for the respective approaches. Our results complement and extend previous results, and contribute in exploring the tractability/intractability frontier of default reasoning from conditional knowledge bases.

[1]Institut für Informationssysteme, Technische Universität Wien, Favoritenstraße 11, A-1040 Wien, Austria. E-mail: eiter@kr.tuwien.ac.at, lukasiewicz@kr.tuwien.ac.at.

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. [23, 24]).

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" 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 [40].

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 [51, 40], infinitesimal probabilities [1, 48], possibility measures [18], and world rankings [52, 33]. Moreover, they characterize an entailment relation based on conditional objects [19]. A survey of all these relationships is given in [5]. 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 [22], 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 [44, 46]. This notion of entailment is equivalent to entailment in system Z by Pearl [49] (which is generalized to variable strength defaults in system Z^+ by Goldszmidt and Pearl [32, 34]), to the least specific possibility entailment by Benferhat et al. [4], and to a conditional (modal) logic-based entailment by Lamarre [43].

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. [30] (and recently generalized to variable strength defaults by Bourne and Parsons [7]); the notion of lexicographic entailment was introduced by Lehmann [45] and Benferhat et al. [3]; the notion of conditional entailment was proposed by Geffner [26, 28]; and, an infinitesimal belief function approach was suggested by Benferhat et al. [6].

While the semantic properties and interrelationships of the various formalisms are quite well understood, their algorithmic properties are less explored. Algorithms for conditional knowledge bases have been described, for example, in [31, 46, 34, 13]. 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.

The aim 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 [31, 46, 34, 13].

Our effort 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 are useful when constructing new algorithms for default reasoning.

The main contributions of this paper are as follows:

- 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 $\phi \to \psi$, is it true that KB entails $\phi \to \psi$? 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 [31], maximum entropy entailment [30] together with its variable-strength extension [7], and Geffner's conditional entailment [26].

- We reconsider Lehmann's lexicographic entailment [45], and correct an error of [13] in the derivation of the complexity of generalized lexicographic entailment [3].

- 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.

- 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 $\phi \to \psi$ 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.

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 $\mathrm{co\text{-}NP}$ (ε-entailment) to Π_2^{P} (Geffner's conditional entailment). In general, they have lower complexity than well-known formalisms of nonmonotonic reasoning such as default logic, circumscription, or autoepistemic logic [54, 35, 20].

- 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 [39, 53].

The rest of this paper is organized as follows. Section 2 contains some preliminaries on conditional knowledge bases and complexity classes. In Section 3, we review various semantics for conditional knowledge bases. In Section 4, we overview and discuss our complexity results. Section 5 contains the proofs of our results (this section may be safely skipped by the reader less interested in technical details). Finally, Section 6 considers related work and outlines future research.

2 Preliminaries

2.1 Conditional Knowledge Bases

We assume a set of *basic propositions* (or *atoms*) $At = \{p_1, p_2, \ldots, p_n\}$ 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, ... We use $(\phi \Rightarrow \psi)$ and $(\phi \vee \psi)$ to abbreviate $\neg(\phi \wedge \neg\psi)$ and $\neg(\neg\phi \wedge \neg\psi)$, respectively. We adopt the usual conventions to eliminate parentheses in combination with these abbreviations. 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).

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 $\pi(D) = \{\pi(d) \mid d \in D\} = \{0, 1, \ldots, k\}$ for some $k \geq 0$.

An *interpretation* (or *world*) is a truth assignment $I\colon At \rightarrow \{\textbf{true}, \textbf{false}\}$, 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 $I(\phi) = \textbf{true}$. 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. The world I *verifies* a default $\phi \rightarrow \psi$, denoted $I \models_{\mathrm{v}} \phi \rightarrow \psi$, 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 $\kappa\colon \mathcal{I}_{At} \rightarrow \{0, 1, \ldots\} \cup \{\infty\}$ 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 $\kappa(\phi) = \infty$ or $\kappa(\phi \wedge \psi) < \kappa(\phi \wedge \neg\psi)$ for all defaults $\phi \rightarrow \psi \in D$. A *default ranking* σ on D maps each $d \in D$ to an integer.

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 [47, 38, 50, 37] for further background).

The class P^{NP} (resp., $\mathrm{NP}^{\mathrm{NP}}$) 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 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 $\mathrm{D^P}$ 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 $\mathrm{D^P}$ can be solved with two NP oracle calls, and is intuitively easier than a problem complete for Δ_2^{P}.

The class $\Delta_2^{\mathrm{P}}[O(\log n)]$ 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 [55]. In particular, it coincides with $\mathrm{L^{NP}}$, logspace computability with an NP oracle, and with $\mathrm{P_{\|}^{NP}}$, 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 $\Delta_2^{\mathrm{P}}[O(\log n)]$ 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 heuristic SAT-algorithms that have been recently developed.

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

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

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. [50, 37]). In particular, FP, $\mathrm{FP_{\|}^{NP}}$ (=$\mathrm{FL^{NP}}$), and $\mathrm{FP^{NP}}$ are the functional analogs of P, $\mathrm{P_{\|}^{NP}}$ (=$\mathrm{L^{NP}}$), and $\mathrm{FP^{NP}}$, 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 $g(I_1, P_2(f(I_1)))$; see [50, 37] for formal details. In case of P and FP, completeness is understood in terms of reductions that can be computed in logarithmic space.

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 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 [48]. We define them in terms of world rankings (see especially Geffner's work [26, 27] 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.*

3.2 Systems Z and Z^+

The notions of entailment in systems Z and Z^+ (Pearl [49] and Goldszmidt and Pearl [32, 34]) apply to ε-consistent conditional knowledge bases $KB = (L, D)$. Moreover, entailment in system Z^+ additionally assumes a strength assignment σ on KB.

Entailment in system Z is linked to an ordered partition of D, a default ranking z, and a world ranking κ^z: Let $(D_0, \ldots, D_k)$ be the unique ordered partition of D such that, for $i = 0, \ldots, k$, 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 $(D_0, \ldots, D_k)$ the *z-partition* of D. We next define the default ranking z as follows. For $j = 0, \ldots, k$, 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:

$$\kappa^z(I) \;=\; \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ 1 + \max\limits_{d \in D:\, I \not\models d} Z(d) & \text{otherwise.} \end{cases} \tag{1}$$

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

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 $I \in \mathcal{I}_{At}$:

$$z^+(d) \;=\; \sigma(d) + \kappa^+(\phi \wedge \psi) \tag{2}$$

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

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 $d \in D$.

3.3 Maximum Entropy Semantics

The maximum entropy approach to default entailment has been introduced by Goldszmidt et al. [30]. Recently, it has been extended to variable strength defaults by Bourne and Parsons [7].

The notion of $z^\star$-entailment applies to ε-consistent minimal-core conditional knowledge bases $KB = (L, D)$ without strength assignment, where KB is *minimal-core* iff for each default $d \in D$ there exists a model I of $L \cup (D - \{d\})$ that falsifies d. This notion of entailment is linked to a default ranking $z^\star$ and a world ranking $\kappa^\star$, 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 $I \in \mathcal{I}_{At}$:

$$z^\star(d) \;=\; 1 + \kappa^\star(\phi \wedge \psi) \tag{4}$$

$$\kappa^\star(I) \;=\; \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ \sum\limits_{d \in D:\, I \not\models d} z^\star(d) & \text{otherwise.} \end{cases} \tag{5}$$

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

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

$$\kappa^\star_{\mathrm{s}}(\phi \wedge \neg\psi) \;=\; \sigma(\phi \to \psi) + \kappa^\star_{\mathrm{s}}(\phi \wedge \psi) \tag{6}$$

$$\kappa^\star_{\mathrm{s}}(I) \;=\; \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ \sum\limits_{d \in D\,:\, I \not\models d} z^\star_{\mathrm{s}}(d) & \text{otherwise.} \end{cases} \tag{7}$$

The uniqueness of $z^\star_{\mathrm{s}}, \kappa^\star_{\mathrm{s}}$ is guaranteed by assuming that $\kappa^\star_{\mathrm{s}}$ is *robust* [7], 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 ranks in $\kappa^\star_{\mathrm{s}}$ 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 $\phi \to \psi \in D$ and all $I \in \mathcal{I}_{At}$, has a unique solution $z^\star_{\mathrm{s}}, \kappa^\star_{\mathrm{s}}$ such that $\kappa^\star_{\mathrm{s}}$ is robust.

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

The notion of $z^\star_{\mathrm{s}}$-entailment is a proper generalization of $z^\star$-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 $\phi \to \psi \in D$ and all $I \in \mathcal{I}_{At}$ has a unique solution $z^\star_{\mathrm{s}}, \kappa^\star_{\mathrm{s}}$, which coincides with $z^\star, \kappa^\star$. Moreover, $\kappa^\star_{\mathrm{s}}$ is robust (and thus, also KB is robust).*

Proof. 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 $\phi \to \psi \in D$ and all $I \in \mathcal{I}_{At}$ has a unique solution $z^\star_{\mathrm{s}}, \kappa^\star_{\mathrm{s}}$, which is given by the rankings $z^\star, \kappa^\star$ for KB.

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

3.4 Lexicographic Entailment

The notion of lexicographic entailment goes back to Lehmann [45] and Benferhat et al. [3].

Lexicographic entailment as introduced in [3] applies to conditional knowledge bases $KB = (L, D)$ with a priority assignment π on KB, which defines an ordered partition $(D_0, \ldots, D_k)$ 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-*preferred* to a world I' iff there exists some $i \in \{0, \ldots, k\}$ 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 default $\phi \to \psi$ is *lex*$_{\mathrm{p}}$-*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.

The notion of lexicographic entailment in [45] 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 3.2). We then say that a default is *lex-entailed* by KB iff $\phi \rightarrow \psi$ is *lex*$_\mathrm{p}$*-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$_\mathrm{p}$-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*$_\mathrm{p}$*-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.

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 $i \in \{0, \ldots, \pi(d)\}$, let $D_{d,i}$ contain the following defaults:

$$a_{d,i} \wedge \alpha \rightarrow \beta \qquad \text{if } i = \pi(d) \tag{8}$$

$$a_{d,i} \rightarrow \neg t_{d,i} \qquad \text{if } i < \pi(d) \tag{9}$$

$$a_{d,i} \rightarrow a_{d,i-1},\ a_{d,i} \rightarrow t_{d,i-1} \qquad \text{if } i > 0. \tag{10}$$

Informally, the default (8) corresponds to the default $\alpha \rightarrow \beta$ from D. By the auxiliary defaults (9) and (10), it is pulled to the level $\pi(d)$ in the z-partition of the augmented set of defaults. The defaults (10) recursively effect that ranking the default $a_{d,i} \wedge \alpha \rightarrow \beta$ for the z-partition (which requests that d has a verifying world) violates all defaults (9), which reside at levels $0, 1, \ldots, i-1$.

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

$$\begin{aligned} L' &= \bigcup \{\alpha \vee \neg a \mid \alpha \in L\} \\ D' &= \bigcup \{D_{d,i} \mid d \in D,\ i \in \{0, \ldots, \pi(d)\}\}\,. \end{aligned}$$

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 $(D'_0, \ldots, D'_k)$ of D' is given by $D'_i = \bigcup \{D_{d,i} \mid d \in D\colon\ \pi(d) \leq i\}$, for all $i = 0, \ldots, k$. 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

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

Satisfaction of a unmasks the background knowledge L, and satisfaction of $A \cup T$ effects that all defaults (9) are false while all defaults (10) are true. Furthermore, each model of ϕ' satisfies the default $a_{d,\pi(d)} \wedge \alpha \rightarrow \beta$ in (8) 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$_\mathrm{p}$-entails $\phi \rightarrow \psi$ iff KB' lex-entails $\phi \wedge \phi' \rightarrow \psi$. $\Box$

The proof of the previous theorem shows in fact that the transformation of lex$_\mathrm{p}$-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.*

3.5 Conditional Entailment

The notion of conditional entailment has been introduced by Geffner [26, 28].

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 *$\prec$-preferred* to a world I', denoted $I \prec I'$, iff $\{d \in D \,|\, I \not\models d\} \neq \{d \in D \,|\, I' \not\models d\}$ 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 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 $\prec$ admissible with KB.

4 Overview of Complexity Results

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 $z^{\star}$-entailment, we assume that KB and d have additionally a strength assignment $\sigma(KB)$ and a strength $\tau(d)$, respectively. In case of $\mathrm{lex_p}$-entailment, we assume that KB has additionally a priority assignment $\pi(KB)$.

Informally, a default reasoning problem represents the input for the entailment problem under a fixed semantics S. We tacitly assume that KB satisfies any preconditions that the definition of 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), we are asked whether KB entails d under some fixed semantics $\mathcal{S}$. In case of z^{+}- and $z^{\star}$-entailment, we are asked whether d is z^{+}- and $z^{\star}$-entailed, respectively, by $(KB, \sigma(KB))$ at strength $\tau(d)$. In case of $\mathrm{lex_p}$-entailment, we are asked whether d is $\mathrm{lex_p}$-entailed by $(KB, \pi(KB))$.

- RANKING**:** Given a conditional knowledge base KB, we are asked to compute the default ranking R of KB according to some fixed semantics $\mathcal{S}$.

- 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 is computed in advance and then on-line available in the input for solving an entailment problem. The complexities of these problems gives us some insight to the question of whether such preprocessing pays off in general.

Our results on the complexity of default reasoning from conditional knowledge bases, together with results from the literature, are compactly summarized in Tables 1–3. 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 1) of all the formalisms considered in this paper. It is thus in the same league as the major formalisms of nonmonotonic reasoning, such as Reiter's default logic, circumscription, and McDermott and Doyle's nonmonotonic logic [35, 54, 20], which are all Π_2^{P}-complete. All other approaches in Table 1 have (considerably) lower complexity.

4.2 Discussion

4.2.1 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 [10].

Also for the ranking-based approaches in Table 1, 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 ordering $\prec$. Note that, different from expectation, classical inference is *not* listed as a principal source of complexity here. This is justified by the results on the Horn restrictions discussed below.

Table 1: 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	$\mathrm{P}_{\parallel}^{\mathrm{NP}}$-complete	P-complete ***	P-complete ***
z^{+}-entailment	P^{NP}-complete ***	P-complete ***	P-complete ***
$z^{\star}$-entailment	P^{NP}-complete	P^{NP}-complete	P^{NP}-complete
$z^{\star}_{\mathrm{s}}$-entailment	P^{NP}-complete	P^{NP}-complete	P^{NP}-complete
lex-entailment	P^{NP}-complete	P^{NP}-complete	P^{NP}-complete
$\mathrm{lex}_{\mathrm{p}}$-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 2: Complexity of computing default rankings

	general case	Horn case	literal-Horn case
z	$\mathrm{FP}_{\parallel}^{\mathrm{NP}}$-complete	FP-complete ***	FP-complete ***
z^{+}	$\mathrm{FP}^{\mathrm{NP}}$-complete ***	FP-complete ***	FP-complete ***
$z^{\star}$	$\mathrm{FP}^{\mathrm{NP}}$-complete	$\mathrm{FP}^{\mathrm{NP}}$-complete	$\mathrm{FP}^{\mathrm{NP}}$-complete
$z^{\star}_{\mathrm{s}}$	$\mathrm{FP}^{\mathrm{NP}}$-complete	$\mathrm{FP}^{\mathrm{NP}}$-complete	$\mathrm{FP}^{\mathrm{NP}}$-complete

Table 3: Complexity of deciding entailment given the default rankings

	general case	Horn case	literal-Horn case
z-entailment	$\mathrm{P}_{\parallel}^{\mathrm{NP}}$-complete ***	P-complete ***	P-complete ***
z^{+}-entailment	$\mathrm{P}_{\parallel}^{\mathrm{NP}}$-complete ***	P-complete ***	P-complete ***
$z^{\star}$-entailment	P^{NP}-complete	P^{NP}-complete	P^{NP}-complete
$z^{\star}_{\mathrm{s}}$-entailment	P^{NP}-complete	P^{NP}-complete	P^{NP}-complete

* Membership shown in [31].
** Membership shown in [46, 34] for default knowledge bases.
*** Membership was shown in [34] for default knowledge bases.
$^{+}$ Shown in [46] for default knowledge bases.
$^{++}$ Reported in [13] for default knowledge bases; the proof sketch for the Horn case in [13] shows merely $\mathrm{P}_{\parallel}^{\mathrm{NP}}$-hardness.

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 by the following steps:

1. Compute the default ranking R for KB;

2. Compute $r(\alpha) = \min_{I \models L \cup \{\alpha\}} r(I)$, 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 2 and 3, 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 $[0, \ldots, n-1]$ and thus over a linear number of values. 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 reminds of the $\mathrm{FP}^{\mathrm{NP}}$-complete problem of computing the lexicographic maximum model of a formula ϕ [41] 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. [9]), this means that computing the z-ranking is in $\mathrm{FP}^{\mathrm{NP}}_{\|}$. We remark that a similar $\mathrm{FP}^{\mathrm{NP}}_{\|}$ algorithm is also feasible for z^+, provided that the strengths $\sigma(d)$ of defaults are not too big (in particular, if each $\sigma(d)$ is bounded by a polynomial in the number of defaults).

Table 3 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 $[0, \ldots, n-1]$, and thus z^+-entailment reduces to z-entailment).

4.2.2 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 $\phi \rightarrow \psi_1 \wedge \cdots \wedge \psi_m$ into a semantically equivalent set of literal-Horn defaults $\phi \rightarrow \psi_1, \ldots, \phi \rightarrow \psi_m$, 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.

At the low end of the complexity range, the Horn restriction gives tractability, as was (essentially) shown in [31, 46, 34]. 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 [39, 53, 12].

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 bases are needed. One such restriction, which is efficiently checkable, is that the admissibility ordering $\prec$ is void.

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 $z^\star$- and $z^\star_\sigma$-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 [30], it makes $z^\star$ (thus also $z^\star_\sigma$) 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 Bottom Line of the Results

We have seen that Geffner's conditional entailment is computationally the most expensive approach among the ones analyzed. As far as the other approaches are concerned, there is a trade-off between an appealing semantics and the computational complexity. More precisely, in the general case, the price for a more appealing semantics is rather modest and generally does not lead to another complexity level. Nevertheless, 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 may be paid by computational intractability. Finally, precomputing default rankings buys nothing or only little in the worst case (in particular, tractability cannot be gained this way).

5 Derivation of 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 1 through 3. 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.

The following results on these problems are known in the literature. Goldszmidt and Pearl [31] showed that deciding ε-consistency for general conditional knowledge bases (resp., Horn conditional knowledge bases) is in P^{NP} (resp., P). Lehmann and Magidor [46] proved that deciding preferential entailment (and thus also ε-entailment) for default knowledge bases is co-NP-complete. Furthermore, Lehmann and Magidor [46] and Goldszmidt and Pearl [34] showed that deciding ε-entailment for Horn default knowledge bases is in P. Finally, Goldszmidt and Pearl [31] 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. Hence, it now just remains to prove the new complexity results stated in Table 1. We need the following lemma, which is essentially a reformulation of a similar result in [31] and [34].

Lemma 5.1 *A conditional knowledge base (L, D) is ε-consistent iff there exists an ordered partition $(D_0, \ldots, D_k)$ 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 [31] 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 $(D_0, \ldots, D_k)$ 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 $\bigcup_{j=i}^{k} D_i$. 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 [25]. Let ϕ be a propositional formula in CNF. By Lemma 5.1, ϕ is satisfiable iff the default knowledge base $(\emptyset, \{\top \to \phi\})$ is ε-consistent.

b) We give a log-space reduction from the P-complete problem of deciding whether a given set $P = \{\phi_1 \Rightarrow \psi_1, \ldots, \phi_n \Rightarrow \psi_n\}$ of definite Horn clauses logically implies a given atom A (see e.g. [14]). Let D denote the set of literal-Horn defaults $\{\phi_1 \to \psi_1, \ldots, \phi_n \to \psi_n, \top \to \neg A\}$. 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 $(D_1, D_2) = (\{\top \to \neg A\}, \{\phi_1 \to \psi_1, \ldots, \phi_n \to \psi_n\})$ of D. Then, the default $\top \to \neg A$ 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. $\Box$

We next prove 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.*

Proof. We give a log-space reduction from the P-complete problem of deciding whether a given set $P = \{\phi_1 \Rightarrow \psi_1, \ldots, \phi_n \Rightarrow \psi_n\}$ of definite Horn clauses logically implies a given atom A (see, for example, [14]). Let D be the set of literal-Horn defaults $\{\phi_1 \to \psi_1, \ldots, \phi_n \to \psi_n\}$. 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 $\top \to A$. $\Box$

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 [31] 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* $\mathrm{D^P}$*-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 first show membership in $\mathrm{D^P}$. By Theorem 3.1, KB properly ε-entails $\phi \to \psi$ iff KB ε-entails $\phi \to \psi$ and $KB \cup \{\phi \to \neg\psi\}$ is ε-consistent. The problem of deciding whether KB ε-entails $\phi \to \psi$ is in co-NP [46]. By Theorem 5.2 a), the problem of deciding whether $KB \cup \{\phi \to \neg\psi\}$ is ε-consistent is in NP. Hence, the problem of deciding whether KB properly ε-entails $\phi \to \psi$ is in $\mathrm{D^P}$.

We next show $\mathrm{D^P}$-hardness. We give a polynomial transformation from the $\mathrm{D^P}$-complete problem SAT-UNSAT [47]: Given two propositional formulas α and β in, 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 $\alpha \to \neg\beta$. Assume first that α is satisfiable and β is unsatisfiable. By Lemma 5.1, it follows that $(\emptyset, \{\alpha \to \beta\})$ is ε-inconsistent and that $(\emptyset, \{\alpha \to \top\})$ is ε-consistent. Hence, by Theorem 3.1, the default knowledge base $(\emptyset, \emptyset)$ properly ε-entails $\alpha \to \neg\beta$. Conversely, let us assume that $(\emptyset, \emptyset)$ properly ε-entails $\alpha \to \neg\beta$. Thus, by Theorem 3.1, it follows that $(\emptyset, \{\alpha \to \beta\})$ is ε-inconsistent and that $(\emptyset, \{\alpha \to \top\})$ 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, for example, [14]): Given a set $P = \{\phi_1 \Rightarrow \psi_1, \ldots, \phi_n \Rightarrow \psi_n\}$ of definite Horn clauses and an atom A, decide whether P logically implies A.

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

5.2 Systems Z and Z^+

The complexity results known in the literature for system Z^+ can be summarized as follows (cf. [34]): The problems ENTAILMENT, RANKING, and RANK-ENTAILMENT in system Z^+ are in P, FP, and P, respectively, for Horn default knowledge bases, and in $\mathrm{P^{NP}}$, $\mathrm{FP^{NP}}$, and $\mathrm{P^{NP}_{\|}}$, respectively, for general default knowledge bases. A fortiori, these upper bounds also hold for system Z. Note that these complexity results easily carry over to our conditional knowledge bases.

We shall prove that each of the 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.

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

Proof. Let $\mathcal{D} = (D_0, \ldots, D_k)$ be the z-partition of D and let $\mathcal{D}' = (D'_0, \ldots, D'_l)$ 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 $i = 0, \ldots, k$.

Algorithm *z*-entailment

Input: ε-consistent conditional knowledge base (L, D) and a default $\phi \to \psi$.
Output: "Yes", if (L, D) z-entails $\phi \to \psi$, otherwise "No."

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

Basis: Let $i = 0$. Let us first assume that there is some $d \in D_0 - D'_0$. Hence, $\mathcal{D}'' = (D'_0 \cup \{d\}, D'_1 - \{d\}, \ldots, D'_l - \{d\})$ 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 $d \in D'_0 - D_0$. 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 $j = 0, \ldots, i-1$. Let us first assume that there is some $d \in D_i - D'_i$. Thus, $\mathcal{D}'' = (D'_0, \ldots, D'_{i-1}, D'_i \cup \{d\}, D'_{i+1} - \{d\}, \ldots, D'_l - \{d\})$ 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 $i = 0, \ldots, k$ and thus also $k = l$. □

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 $\mathrm{P}^{\mathrm{NP}}_{\|}$-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 first show membership in $\mathrm{P}^{\mathrm{NP}}_{\|}$. Towards this goal, consider algorithm ***z*-entailment**. It is not hard to see that it correctly decides whether KB z-entails the default $\phi \to \psi$: 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 $(D_0, \ldots, D_k)$ 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 $i \in \{0, 1, \ldots, n\}$; 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 [9] 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 $1 + (1 + n(n-1)/2)(2n+2)$ parallel NP oracle queries, which we omit here. Thus the problem is in $\mathrm{P}_{\|}^{\mathrm{NP}}$. (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 calls of an NP-oracle. This alternatively proves membership in $\Delta_2^{\mathrm{P}}[O(\log n)] = \mathrm{P}_{\|}^{\mathrm{NP}}$.)

We now show $\mathrm{P}_{\|}^{\mathrm{NP}}$-hardness. We give a polynomial transformation from the following $\mathrm{P}_{\|}^{\mathrm{NP}}$-complete problem [55]. Given the propositional formulas in CNF $\alpha_1, \ldots, \alpha_m$, we are asked whether the number of tautologies among $\alpha_1, \ldots, \alpha_m$ is even. Without loss of generality, we can assume that $\alpha_1, \ldots, \alpha_m$ 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 whenever α_i is not a tautology, cf. [55]. Let

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

and let $D = \{d_1, \ldots, d_m\}$. 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 $i = 1, \ldots, m$. Let ϕ and ψ be defined as

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

respectively. We now show that the number of tautologies among $\alpha_1, \ldots, \alpha_m$ 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 $\alpha_1, \ldots, \alpha_m$ 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, $I \models \phi \wedge \neg\psi$m, 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 $P = \{\phi_1 \Rightarrow \psi_1, \ldots, \phi_n \Rightarrow \psi_n\}$ of definite Horn clauses logically implies entails a given atom A (see e.g. [14]). Let D denote the set of literal-Horn defaults $\{\phi_1 \rightarrow \psi_1, \ldots, \phi_n \rightarrow \psi_n\}$. 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 $\top \rightarrow A$. □

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 conditional knowledge base is* $\mathrm{FP}_{\|}^{\mathrm{NP}}$*-complete. b) The problem of computing the default ranking z for a literal-Horne conditional knowledge base is* FP*-hard.*

Algorithm *z*-ranking

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

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

Proof. a) We first show membership in $\text{FP}_{\|}^{\text{NP}}$. 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 $(D_0, \ldots, D_k)$ 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 $n = |D|$. The queries in Steps 2 and 3, respectively, can be solved by parallel NP-oracle calls. These two rounds of $n(n-1)/2+1$ and n, 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 $i \in \{0, 1, \ldots, n(n-1)/2\}$, $d \in D$, and $j \in \{0, 1, \ldots, n-1\}$, a query asks whether D has an ordered partition $(D_0, \ldots, D_k)$ 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 $\text{FP}_{\|}^{\text{NP}}$.

We next show $\text{FP}_{\|}^{\text{NP}}$-hardness. We give a polynomial transformation from the following problem, which is easily seen to be $\text{FP}_{\|}^{\text{NP}}$-complete. Given propositional CNF formulas $\alpha_1, \ldots, \alpha_m$ on disjoint sets of variables, $m \geq 1$, compute the truth values $a_1, \ldots, a_m$ with $a_i =$ **true** iff α_i is satisfiable. Let

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

and define $D = \{d, d_1, \ldots, d_m\}$. 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 $\text{FP}_{\|}^{\text{NP}}$-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 $\text{FP}^{\text{NP}}[O(\log n)] \subseteq \text{FP}_{\|}^{\text{NP}}$ holds and the inclusion is believed to be strict [37], the $\text{FP}_{\|}^{\text{NP}}$-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 5.8 *Given a set* $P = \{\phi_1 \Rightarrow \psi_1, \ldots, \phi_m \Rightarrow \psi_m\}$ *of definite Horn clauses on the atoms* $At = \{x_1, \ldots, x_n\}$*, computing* $S = \{A \in At \mid P \models A\}$ *is* FP*-complete.*

Algorithm z-rank-entailment

Input: ε-consistent conditional knowledge base (L, D), default $\phi \to \psi$, and z-ranking of D.
Output: "Yes", if (L, D) z-entails $\phi \to \psi$, otherwise "No."

1. **if** ϕ is unsatisfiable **then return** "No";
2. $\alpha_1 := \phi \wedge \psi$; $\alpha_2 := \phi \wedge \neg\psi$;
3. **for each** $i \in \{0, 1, \ldots, n\}$ 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 $k[i, j]$:= true **else** $k[i, j]$:=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".

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

Algorithm ***z*-entailment** can be modified to an algorithm ***z*-rank-entailment** that uses the z-ranking provided in 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. As follows from our next result, it is unlikely that we can improve on this and find an algorithm that solves entailment with a fixed number of NP oracle calls.

Theorem 5.9 *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 $\mathrm{P}^{\mathrm{NP}}_{\|}$-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.*

Proof. a) The result follows from the proof of Theorem 5.6 a): Since the ranking z for the default knowledge base $(\emptyset, D)$ there is $Z(d_i) = i - 1$ for all $i = 1, \ldots, m$, the problem there is easily reduced to the case where z is given in advance. $\Box$

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

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 1-3, it remains to show matching hardness results for ENTAILMENT and RANKING in the general case.

Theorem 5.10 *Given a conditional knowledge base KB with strength assignment σ, a default d, and a strength τ, deciding whether (KB, σ) z^+-entails d at strength τ 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 [41]. Given a conjunction $\phi = \phi_1 \wedge \ldots \wedge \phi_m$ of clauses ϕ_i on atoms $x_1, \ldots, x_n$, which is asserted to be satisfiable, decide whether $I_{\mathrm{lsm}}(\phi) \models x_n$, where $I_{\mathrm{lsm}}(\phi)$ is the lexicographically maximum satisfying truth assignment of ϕ.

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

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

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

Let $D = \{d_1^+, d_1^-, d_2^+, d_2^-, \ldots, d_n^+, d_n^-\}$. Intuitively, the defaults d_i^+ and d_i^- assess the cost of the truth value for x_i in a satisfying assignment for ϕ. For determining the rank of the default d_i^+ (resp., d_i^-), exactly one of the defaults d_j^+ and d_j^- is violated, for each $j < i$. Since the strengths of the defaults d_i^-, $i \geq 1$, decrease in a mathematical progression, minimization of the violation cost effects that, if possible, x_i is set to **true** rather than to **false**. Thus, the ranks of d_i^+, d_i^- reflect the prefix of $I_{\mathrm{lsm}}(\phi)$. The factor 2^k serves to account for the 1's that are added at each ranking step according to equation (3).

We now show that x_n is true in $I_{\mathrm{lms}}(\phi)$ iff $KB = (\emptyset, D)$ z^+-entails the default $d = \neg c_1 \wedge \cdots \wedge \neg c_n \wedge \phi \rightarrow x_n$ with strength $\tau = 0$.

Let a denote the integer that has the complement of $I_{\mathrm{lms}}(\phi)$ as binary representation. For example, $a = 10 = 2^3 + 2^1$ for $I_{\mathrm{lms}}(\phi) = 0101$. We first show that the default ranking z^+ is given as follows. For all $i = 1, \ldots, n$:

$$\begin{aligned} z^+(d_i^+) &= 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1 \\ z^+(d_i^-) &= 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1 + 2^{k+n-i}\,. \end{aligned} \tag{11}$$

That is, the binary representation of $z^+(d_i^+)$ has the first $i-1$ bits of $I_{\mathrm{lms}}(\phi)$, 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 $i = 1, \ldots, n$ are given by (11). For each world I, let $\kappa^+(I)$ be defined by (3). To show that (11) actually defines the solutions, we must show that (2) holds for all defaults $d \in D$.

Let $i \geq 1$ and let I be an interpretation that satisfies $\neg c_1 \wedge \cdots \wedge \neg c_{i-1} \wedge c_i \wedge \phi$ and that coincides on the variables $x_1, \ldots, x_{i-1}$ with $I_{\mathrm{lms}}(\phi)$. Hence, I verifies d_i^+ and d_i^-. Moreover, I satisfies all defaults d_j^+ and d_j^- with $j > i$. Finally, for all $j = 1, \ldots, i-1$, exactly one default among d_j^+ and d_j^- is falsified by I. It is now easy to see that the falsified default among d_j^+ and d_j^- is assigned the value $2^{k+n-j} \cdot (a \operatorname{div} 2^{n-j}) + j - 1$ under z^+. Since this value is maximal for $j = i - 1$, we get

$$\min_{J \in \mathcal{I}_{At} \,:\, J \models_{\mathrm{v}} d_i} \kappa^+(J) \;\leq\; 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1$$

for all $d_i \in \{d_i^+, d_i^-\}$. Let I' be any other interpretation that satisfies $\neg c_1 \wedge \cdots \wedge \neg c_{i-1} \wedge \phi$. Thus, for all $j = 1, \ldots, i-1$, exactly one default among d_j^+ and d_j^- is falsified by I'. Assume that I' does not coincide on the variables $x_1, \ldots, x_{i-1}$ with $I_{\mathrm{lms}}(\phi)$. 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 $I' \models d_j^-$. 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, since $i - 1 - j \leq n < 2^k$, we get:

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

That is, we get $\kappa^+(I) < \kappa^+(I')$ for any other such I'. It thus follows:

$$\min_{J \in \mathcal{I}_{At} \colon J \models_{\mathrm{v}} d_i} \kappa^+(J) \;\; = \;\; 2^{k+n-i+1} \cdot (a \operatorname{div} 2^{n-i+1}) + i - 1$$

for all $d_i \in \{d_i^+, d_i^-\}$. But this shows that (2) holds for all defaults $d \in D$. Thus, (11) described the actual z^+-ranking of $(\emptyset, D)$.

Let I be an interpretation that satisfies $\neg c_1 \wedge \cdots \wedge \neg c_n \wedge \phi$ and that coincides on the variables $x_1, \ldots, x_n$ with $I_{\mathrm{lms}}(\phi)$. Let I' be any other interpretation that satisfies $\neg c_1 \wedge \cdots \wedge \neg c_n \wedge \phi$ and that does not coincide on the variables $x_1, \ldots, x_n$ with $I_{\mathrm{lms}}(\phi)$. 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 $I_{\mathrm{lms}}(\phi)$. It then follows that

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

That is, *KB* z^+-entails d with strength 0. Let us next assume that x_n is false in $I_{\mathrm{lms}}(\phi)$. It then follows $\kappa^+(\neg c_1 \wedge \cdots \wedge \neg c_n \wedge \phi \wedge \neg x_n) < \kappa^+(\neg c_1 \wedge \cdots \wedge \neg c_n \wedge \phi \wedge x_n)$. That is, *KB* does not z^+-entail d with strength 0. $\Box$

Theorem 5.11 *The problem of computing the default ranking z^+ for a conditional knowledge base KB with strength assignment σ is* $\mathrm{FP}^{\mathrm{NP}}$*-hard.*

Proof. We give a polynomial transformation from the following $\mathrm{FP}^{\mathrm{NP}}$-complete problem [41]. Given a conjunction $\phi = \phi_1 \wedge \ldots \wedge \phi_m$ of clauses on atoms $x_1, \ldots, x_n$, which is asserted to be satisfiable, compute $I_{\mathrm{lsm}}(\phi)$, that is, the lexicographically maximum satisfying truth assignment of ϕ.

We slightly extend the construction in the proof of Theorem 5.10 as follows. Add to the set D of defaults there the default $d = \neg c_1 \wedge \cdots \wedge \neg c_n \wedge \phi \to \top$ with $\sigma(d) = 0$. By a similar line of argumentation, it follows that (11) holds for all $i = 1, \ldots, n$ and that $z^+(d) = a + n$ (recall that a is the integer having the complement of $I_{\mathrm{lms}}(\phi)$ as binary representation). Since $I_{\mathrm{lms}}(\phi)$ is easily computed from $z^+(d)$, the result follows. $\Box$

5.3 Maximum Entropy Semantics

We now focus on the complexity of deciding $z^\star$- and $z_{\mathrm{s}}^\star$-entailment.

Not much work has been spent determining the complexity of $z^\star$- and $z_{\mathrm{s}}^\star$-entailment. Goldszmidt et al. [30] suspect that the complexity of $z^\star$-entailment is high, and briefly note that, referring to results on Horn optimization problems, [2], the problem should be NP-hard in the Horn case. We now determine the precise complexity of $z^\star$-entailment, and start with the lower complexity bound.

Theorem 5.12 *Given a literal-Horn conditional KB that is ε-consistent and minimal-core, and a literal-Horn default d, deciding whether KB $z^\star$-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 $C = \{\alpha_1 \Rightarrow \beta_1, \ldots, \alpha_m \Rightarrow \beta_m\}$ on atoms $x_1, \ldots, x_n$ 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 $r \in \{1, \ldots, m\}$, decide whether $I \models \alpha_r \Rightarrow \beta_r$ for every maximum weight world I for 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 [41] that computing the maximum weight assignment for a set C of weighted arbitrary clauses is $\mathrm{FP}^{\mathrm{NP}}$-complete; the proof in [41] 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 for 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 of C iff KB $z^\star$-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 $z^\star(d_{c_j,j}) = 2^{c_j}$ 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 for C. Hence, it finally just remains to choose an appropriate literal ψ, such that the default $\phi \rightarrow \psi$ is $z^\star$-entailed by KB iff $\alpha_r \Rightarrow \beta_r$ holds in all minimal falsifying models of ϕ.

Let the set of atomic propositions be $At = \{x_1, \ldots, x_n\} \cup A \cup B \cup T$, 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:

$$
\begin{aligned}
L_{i,j} &= \{b_{i,j} \Rightarrow a_{0,j},\ b_{i,j} \Rightarrow t_{0,j},\ \ldots,\ b_{i,j} \Rightarrow a_{i-1,j},\ b_{i,j} \Rightarrow t_{i-1,j}\} \\
\phi_{i,j} &= \begin{cases} b_{i,j} \wedge t_{i,j} \Rightarrow \bot & \text{if } i < c_j \\ b_{i,j} \wedge \alpha_i \Rightarrow \beta_i & \text{if } i = c_j \end{cases} \\
d_{i,j} &= a_{i,j} \rightarrow b_{i,j}\,.
\end{aligned}
$$

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

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

We first show that KB is ε-consistent. Let the ordered partition $(D_0, \ldots, D_k)$ of D be defined by $k = \max(c_1, \ldots, c_m)$ and $D_i = \{d_{i,j} \mid d_{i,j} \in D\}$ for all $i = 1, \ldots, k$. 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 $l \leq i$, $I_{i,j} \models b_{i,j}$, $I_{i,j} \models t_{l,j}$, 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.

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 $L \cup (D - \{d_{i,j}\})$.

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

Basis: Let $i = 0$. Since $(D_0, \ldots, D_k)$ is a partition of D that is admissible with KB, we get $z^\star(d_{i,j}) = 1$ 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, $z^\star(d_{l,j}) = 2^l$ for every $l \leq i - 1$, and thus $\kappa^\star(I_{i,j}) = \sum_{l=0}^{i-1} 2^l = 2^i - 1$. Hence, it follows $z^\star(d_{i,j}) \leq 2^i$. 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 $z^\star(d_{i,j}) = 2^i$, which concludes the induction.

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 $z^\star$-entails d. We need some preparation as follows.

Let I be any world such that: (i) $I \models L \cup \{\phi\}$, (ii) I is a maximum weight world of C, and (iii) $I \models b_{c_j,j}$ iff $I \models \alpha_j \Rightarrow \beta_j$. Let I' be any world such that (i) holds but either (ii) or (iii) does not hold. We now show that $\kappa^\star(I) < \kappa^\star(I')$. 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 $j \leq n$. Then, $I'' \models L$, and no default $d_{i,j}$ that is satisfied by I'' is violated by $d_{i,j}$. Hence, it follows $\kappa^\star(I'') \leq \kappa^\star(I')$. 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:

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

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

Let us now first assume that $I \models \alpha_r \Rightarrow \beta_r$ holds for every maximum weight world I of 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'' of 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 $\kappa^\star(\phi \wedge \psi) < \kappa^\star(\phi \wedge \neg\psi)$, and thus KB $z^\star$-entails d.

Let us next assume that $I \not\models \alpha_r \Rightarrow \beta_r$ for some maximum weight world of 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 $\kappa^\star(\phi \wedge \psi) \geq \kappa^\star(\phi \wedge \neg\psi)$, and thus KB does not $z^\star$-entail d. $\Box$

Theorem 5.13 *The problem of computing the default ranking $z^\star$ for an ε-consistent minimal-core literal-Horn conditional knowledge base is* $\mathrm{FP}^{\mathrm{NP}}$*-hard.*

Proof. We give a polynomial transformation from a suitable variant of the problem used in the reduction in the proof of Theorem 5.12, which is $\mathrm{FP}^{\mathrm{NP}}$-complete: Given a set of weighted Horn clauses $C = \{\alpha_1 \Rightarrow$

$\beta_1, \ldots, \alpha_m \Rightarrow \beta_m\}$ on n variables $x_1, \ldots, x_n$, 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 of C, that is, $w = \max_{I \in \mathcal{I}_{At}} \sum_{I \models \alpha_i \Rightarrow \beta_i} w_i$. $\mathrm{FP}^{\mathrm{NP}}$-hardness of this problem can be established by a suitable adaptation of proofs in [41].

We slightly extend KB in the proof of Theorem 5.12 as follows. We introduce new atoms $a^\star$ and $b^\star$, and following set of literal-Horn clauses $L^\star$ and literal-Horn default $d^\star$:

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

By a line of argumentation similar to the one in the proof of Theorem 5.12, it is easy to see that the extended conditional knowledge base $(L \cup L^\star, D \cup \{d^\star\})$ is ε-consistent and minimal-core. Moreover, its ranking $z^\star$ assigns all defaults $d_{i,j}$ the value 2^i and the default $d^\star$ the value $\sum_{j=1}^{m} \sum_{i=0}^{c_j} 2^i - w$. Consequently, the weight w of a maximum weight world of C is given by $w = 2 \cdot \sum_{j=1}^{m} 2^{c_j} - m - z^\star(d^\star)$, which can be easily computed from $z^\star$. This proves the result. □

Theorem 5.14 *Given a literal-Horn conditional knowledge base KB that is ε-consistent and minimal-core, the ranking $z^\star$ of KB, and a literal-Horn default d, deciding whether KB $z^\star$-entails d is P^{NP}-hard.*

Proof. The result follows from the proof of Theorem 5.12: The default ranking $z^\star$ 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 $z^\star$ is part of the input. □

We next consider the extension of the maximum entropy approach $z^\star$ by variable strength defaults to $z^\star_{\mathrm{s}}$. Our goal is to show that $z^\star_{\mathrm{s}}$-entailment is in P^{NP}, and thus no harder than $z^\star$-entailment. For this purpose, we use the algorithm **$z^\star_{\mathrm{s}}$-ranking**, which computes the default ranking $z^\star_{\mathrm{s}}$ of a conditional knowledge base KB that is ε-consistent. 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 [7], and thus we do not analyze its correctness here.

The next lemma is shows that the rank $z^\star_{\mathrm{s}}(d)$ 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 5.15 *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 $z^\star_{\mathrm{s}}(d_1), \ldots, z^\star_{\mathrm{s}}(d_l)$, with $l \leq n$, denote the sequence of default ranks computed in algorithm $z^\star_{\mathrm{s}}$-ranking. Then $|z^\star_{\mathrm{s}}(d_i)| \leq s \cdot 2^{i-1}$ for all $i = 1, \ldots, l$.*

Proof. The claim is proved by induction on $i = 1, \ldots, l$ as follows.

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

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

By showing that algorithm $z^\star_{\mathrm{s}}$-ranking can be implemented to run in deterministic polynomial time if an NP-oracle is available, we prove the following result.

Algorithm $z^\star_{\rm s}$-ranking (essentially [7])

Input: ε-consistent conditional knowledge base $KB = (L, D)$ with positive strength assignment σ.

Output: Default ranking $z^\star_{\rm s}$, if the system of equations given by (6) and (7) for all $\phi \to \psi \in D$ and all $I \in \mathcal{I}_{At}$ has a unique solution $z^\star_{\rm s}$, $\kappa^\star_{\rm s}$ such that $\kappa^\star_{\rm s}$ is robust; otherwise, *nil*.

Notation: We use $\mathrm{minv}(\phi \to \psi)$ and $\mathrm{minf}(\phi \to \psi)$ to denote $\kappa^\star_{\rm s}(\phi \wedge \psi)$ and $\kappa^\star_{\rm s}(\phi \wedge \neg\psi)$, respectively, where $\kappa^\star_{\rm s}(\alpha) = \min_{I \in \mathcal{I}_{At}: I \models \alpha} \kappa^\star_{\rm s}(I)$ with $\kappa^\star_{\rm s}(I)$ as in (7).

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

Theorem 5.16 *Given an ε-consistent conditional knowledge base $KB = (L, D)$ with positive strength assignment σ, the problem of computing the default ranking $z^\star_{\rm s}$ for KB, if KB is robust, and returning nil otherwise, is in* $\mathrm{FP}^{\mathrm{NP}}$.

Proof. Let $n = |D|$ and let $s = \max(\{\sigma(d) \mid d \in D\})$. Furthermore, denote by $|s|$ the length of the binary representation of s.

It is sufficient to show the following properties of algorithm **$z^\star_{\rm s}$-ranking**: (i) computing $\mathrm{minv}(d)$ and $\mathrm{minf}(d)$ in Steps 3,5, and 6 is in $\mathrm{FP}^{\mathrm{NP}}$; (ii) deciding in Step 8 whether $\kappa^\star_{\rm s}$ satisfies (6) for each $\phi \to \psi \in D$ is in P^{NP}; and (iii) deciding whether $\kappa^\star_{\rm s}$ is robust in Step 8 is in P^{NP}. Lemma 5.15 implies that for every $\phi \to \psi \in D$, the values $\kappa^\star_{\rm s}(\phi \wedge \psi)$ and $\kappa^\star_{\rm s}(\phi \wedge \neg\psi)$ are from $\{0, \ldots, s \cdot 2^n\} \cup \{\infty\}$. 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 $\kappa^\star_{\rm s}(I) \leq w$ for a given value w is in NP. This already shows that (i) and (ii) are in $\mathrm{FP}^{\mathrm{NP}}$ 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 each value $\mathrm{minf}(d)$, for $d \in D$, is known, which can be computed in $\mathrm{FP}^{\mathrm{NP}}$, this is clearly in co-NP, and thus can be checked with a call to an NP-oracle. This proves the result. $\Box$

The announced upper bound on $z^\star_{\rm s}$-entailment is now easily established.

Theorem 5.17 *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) $z^\star_{\rm s}$-entails d at strength τ, is in* P^{NP}.

Proof. Let $d = \phi \to \psi$. By Theorem 5.16, computing the $z^\star_{\rm s}$-ranking of KB is in $\mathrm{FP}^{\mathrm{NP}}$. Deciding whether $\kappa^\star_{\rm s}(\phi) = \infty$ is in NP, and by similar arguments as in the proof of Theorem 5.16, computing the ranks $\kappa^\star_{\rm s}(\phi \wedge \psi)$ and $\kappa^\star_{\rm s}(\phi \wedge \neg\psi)$ in a binary search using $z^\star_{\rm s}$ is then in $\mathrm{FP}^{\mathrm{NP}}$; testing $\kappa^\star_{\rm s}(\phi \wedge \psi) < \kappa^\star_{\rm s}(\phi \wedge \neg\psi)$ is simple. Overall, deciding whether KB $z^\star_{\rm s}$-entails d is in P^{NP}. $\Box$

The matching hardness result is inherited from the P^{NP}-hardness of $z^{\star}$-entailment in Theorem 5.12 and the fact that the minimal-core property implies robustness (Lemma 3.2).

Theorem 5.18 *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, σ) $z^{\star}_{\mathrm{s}}$-entails d at strength τ, is P^{NP}-hard. This remains true if the ranking $z^{\star}_{\mathrm{s}}$ is part of the input and $\tau = 0$ is fixed.*

5.4 Lexicographic Entailment

Cayrol et al. show in [13] that $\mathrm{lex_p}$-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 1.

As for the lower bound, Cayrol et al. show in [13] that $\mathrm{lex_p}$-entailment is P^{NP}-hard for general default knowledge bases. Moreover, they state in [13] that $\mathrm{lex_p}$-entailment is P^{NP}-hard for Horn default knowledge bases. However, the short proof sketch in [13] is inappropriate, since it mentions a reduction from a problem that is obviously in $\mathrm{P}^{\mathrm{NP}}_{\|}$; this would only establish $\mathrm{P}^{\mathrm{NP}}_{\|}$-hardness for the Horn case.

For this reason, we next prove that the problem 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 5.19 *Given a literal-Horn default knowledge base $KB = (\emptyset, D)$ with priority assignment π, and a literal-Horn default d, deciding whether (KB, π) $\mathrm{lex_p}$-entails d is P^{NP}-hard.*

Proof. We give a polynomial transformation from the following P^{NP}-complete problem [41] (cf. Theorem 5.10). Given a conjunction $\alpha = \alpha_1 \wedge \ldots \wedge \alpha_m$ of clauses α_i on atoms $x_1, \ldots, x_n$, which is asserted to be satisfiable, decide whether $I_{\mathrm{lsm}}(\alpha) \models x_n$ where $I_{\mathrm{lsm}}(\alpha)$ is the lexicographically maximum model of α,.

We now construct $KB = (\emptyset, D)$, π, and $d = \phi \rightarrow \psi$ as in the problem statement such that $I_{\mathrm{lsm}}(\alpha) \models x_n$ iff (KB, π) $\mathrm{lex_p}$-entails d.

Roughly speaking, the main problem in this construction is to transform the clauses $\alpha_1, \ldots, \alpha_m$ into Horn clauses $\alpha^{\star}_1, \ldots, \alpha^{\star}_m$. We tackle this by replacing all positive literals x_i in $\alpha_1, \ldots, \alpha_m$ 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.

In detail, let the set of atoms be $At = \{a, b, c, e\} \cup \{s_i, x_i, y_i \mid 1 \leq i \leq n\}$. Define

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

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 $\alpha^{\star}_j$ 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:

$$\begin{aligned}
D_i &= \{e \rightarrow x_{n-i}\} \qquad \text{for all } i = 0, \ldots, n-1 \\
D_n &= \{\alpha^{\star}_1, \ldots, \alpha^{\star}_m\} \\
D_{n+1} &= \bigcup \{F_i \mid i = 1, \ldots, n\} \\
D_{n+2} &= \{\top \rightarrow a, \top \rightarrow b, \top \rightarrow c, \top \rightarrow e\}\,.
\end{aligned}$$

The priority assignment π on D is given by $\pi(d) = i$ for all $i = 0, \ldots, n+2$ 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 $I_{\mathrm{lsm}}(\alpha) \models x_n$ iff KB $\mathrm{lex_p}$-entails d. It is sufficient to show that, for any preferred world I, its restriction to $X = \{x_1, \ldots, x_n\}$, denoted by $I|X$, coincides with I_{lsm}, and that, on the other hand, $I_{\mathrm{lsm}}(\alpha)$ 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 $i = 1, \ldots, n$. Since α is satisfiable, preferredness of I then implies that $I \models D_n$ and $I|X \models \alpha$. The sets of defaults $D_0, \ldots, D_{n-1}$ then ensure that $I|X$ indeed coincides with $I_{\mathrm{lsm}}(\alpha)$.

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

We next turn to the complexity of lex-entailment. Since this is a special case of $\mathrm{lex_p}$-entailment, we can solve this problem by using an algorithm for $\mathrm{lex_p}$-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 $\mathrm{lex_p}$-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 5.20 *a) Given an ε-consistent conditional knowledge base KB and a default d, deciding whether KB lex-entails d is in* $\mathrm{P^{NP}}$. *b) Given an ε-consistent literal-Horn default knowledge base KB and a literal-Horn default d, deciding whether KB lex-entails d is* $\mathrm{P^{NP}}$*-hard.*

Proof. a) By Theorem 5.7 a), computing the default ranking z for KB is in $\mathrm{FP}^{\mathrm{NP}}_{\|}$. Let the priority assignment π on KB be given by z. Then, KB lex-entails d iff (KB, π) $\mathrm{lex_p}$-entails d. As deciding the latter is in $\mathrm{P^{NP}}$ (see the discussion at the beginning of Section 5.4), deciding whether KB lex-entails d is also in $\mathrm{P^{NP}}$.

b) We show that the $\mathrm{P^{NP}}$-hard problem in Theorem 5.19, which is more general, is reducible to this problem. From the proof of Theorem 5.19, 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. $\Box$

5.5 Conditional Entailment

We finally concentrate on the complexity of Geffner's conditional entailment. To our knowledge, no complexity results have been derived for this notion of entailment so far.

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

Lemma 5.21 *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 $D_d = D - \{d' \in D \,|\, d' \prec d\}$ under L.*

Proof. 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. $\Box$

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

Algorithm Not-Cond-Entailment

Input: Conditional knowledge base (L, D) and a default $\phi \to \psi$.
Output: "Yes" iff (L, D) does not conditionally entail $\phi \to \psi$.

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".

Theorem 5.22 *Given a conditional knowledge base $KB = (L, D)$ and a default $d = \phi \to \psi$, deciding whether KB conditionally entails d is in Π_2^{P}.*

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 5.21 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 $\mathrm{NP}^{\mathrm{NP}} = \Sigma_2^{\mathrm{P}}$. $\Box$

We next show that this upper bound is tight, and that this holds even for the literal-Horn case. The proof of this result will make use of the following lemma shown by Geffner [26].

Lemma 5.23 (Geffner [26]) *Let $KB = (L, D)$ be a conditional knowledge base. A default $\phi \to \psi$ 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 prove that conditional entailment is Π_2^{P}-hard for the literal-Horn case.

Theorem 5.24 *Given a literal-Horn conditional knowledge base $KB = (L, D)$ and a literal-Horn default $d = \phi \to \psi$, 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 [38]. Given a collection of clauses $\alpha_1, \ldots, \alpha_l$ on the atoms $y_1, \ldots, y_m, x_1, \ldots, x_n$, where $m, n \geq 1$, decide whether the quantified Boolean formula $\Phi = \forall y_1 \ldots \forall y_m \exists x_1 \ldots \exists x_n\, (\alpha_1 \wedge \cdots \wedge \alpha_l)$ evaluates to true. That is, is it true that for each truth assignment I_y to the variables $y_1, \ldots, y_m$, there exists a truth assignment I_x to the variables $x_1, \ldots, x_n$ such that $I_y \cup I_x$ satisfies $\alpha_1 \wedge \cdots \wedge \alpha_l$. Without loss of generality, we can assume that each clause α_i contains at least one variable from $x_1, \ldots, x_n$.

We now construct a literal-Horn conditional knowledge base $KB = (L, D)$ and a literal-Horn default $d = \phi \to \psi$ such that Φ evaluates to true iff KB conditionally entails d.

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 $\alpha_1, \ldots, \alpha_l$ into Horn clauses $\alpha_1^\star, \ldots, \alpha_l^\star$. This will be done by replacing all positive literals y_i and x_j in $\alpha_1, \ldots, \alpha_l$ by new negative literals $\neg y_i'$ and $\neg x_j'$, respectively. We will then use L to express the relationships $y_i \wedge y_i' \Rightarrow \bot$ and $x_j \wedge x_j' \Rightarrow \bot$, 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 $x_j \equiv \neg x_j'$.

Let $At = A_y \cup A_x \cup \{a\}$, where $A_x = \{a_i,\, b_i,\, y_i,\, y_i' \mid 1 \leq i \leq m\}$, $A_x = \{c_j,\, d_j,\, e_j,\, f_j,\, x_j,\, x_j' \mid 1 \leq j \leq n\}$. 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_i^1, D_j^2 are defined as follows:

$$\begin{aligned} L_1 &= \{\alpha_1^\star \vee \neg a, \ldots, \alpha_l^\star \vee \neg a\}, \\ L_2 &= \{\neg y_i \vee \neg y_i' \vee \neg a \mid i = 1, \ldots, m\}, \\ L_3 &= \{\neg x_j \vee \neg x_j' \vee \neg a \mid j = 1, \ldots, n\}, \\ L_4 &= \{x_j \Rightarrow \neg f_k,\; x_j' \Rightarrow \neg f_k \,|\, j, k = 1, \ldots, n\}, \end{aligned}$$

where $\alpha_1^\star, \ldots, \alpha_l^\star$ is obtained from $\alpha_1, \ldots, \alpha_l$ 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

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

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 5.23, 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:

$$c_j \rightarrow f_j \;\prec\; c_j \wedge d_j \rightarrow x_j, \tag{12}$$

$$c_j \rightarrow f_j \;\prec\; c_j \wedge e_j \rightarrow x_j', \qquad \text{for all } j = 1, \ldots, n. \tag{13}$$

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 $\prec^\star$ be the priority ordering on D that is given by exactly all pairs in (12) and (13). It is easy to see that each $d \in D$ is tolerated under L by $D - \{d' \in D \mid d' \prec^\star d\}$; thus, by Lemma 5.21, $\prec^\star$ is admissible with KB. This means that $\prec^\star$ is the least (that is, unique minimal) priority ordering on D admissible with KB. Applying Lemma 5.23, the default $d = \phi \to \psi$ is conditionally entailed by KB iff ψ is satisfied in every $\prec^\star$-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 $f : \{y_1, \ldots, y_m\} \to \{\bot, \top\}$ such that the formula $\alpha = (\alpha_1 \wedge \cdots \wedge \alpha_l)\,[y_1/f(y_1), \ldots, y_m/f(y_m)]$ 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 $i \leq m$; (ii) $I(x_j) = I(x'_j) =$ **false**, for all $j = 1, \ldots, n$; and (iii) $I(p) = \mathtt{true}$ for any other atom p. Since each $\alpha_i^\star$ contains at least one negative literal from $\neg x_1, \ldots, \neg x_n, \neg x'_1, \ldots, \neg x'_n$, 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 $J \prec^\star I$. We show that no such J exists, which proves that I is a $\prec^\star$-preferred model of $L \cup \{\phi\}$. Because of the defaults in $D_{1,i}$, J must coincide on y_i and y'_i with I, for $i = 1, \ldots, m$. 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 effect that $J \models \neg f_j$ for all $j = 1, \ldots, n$. Since the default $c_j \to f_j$ from $D_{2,j}$ is violated by J but not by I, the fact that $J \prec^\star I$ implies that either $c_j \wedge d_j \to x_j$ or $c_j \wedge e_j \to x'_j$, 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 $j = 1, \ldots, n$. However, this means that α is satisfiable, which is a contradiction. Thus, J does not exist, and I is a $\prec^\star$-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 $\prec^\star$-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 $I(x_j) = I(x'_j) =$ **false**, for all $j = 1, \ldots, n$; preferredness of I implies that $I(f_j) =$ **true** (as the default $c_j \to f_j$ will be satisfied), for all $j = 1, \ldots, n$. Let the mapping $f\colon \{y_1, \ldots, y_m\} \to \{\bot, \top\}$ be defined by $f(y_i) = \top$ iff $I \models a_i \to y_i$. We now show that $\alpha = (\alpha_1 \wedge \cdots \wedge \alpha_l)\,[y_1/f(y_1), \ldots, y_m/f(y_m)]$ is unsatisfiable. Towards a contradiction, suppose there exists a truth assignment I' to $x_1, \ldots, x_n$ 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 $I''(f_j) =$ **false**, for all $j = 1, \ldots, n$, and sets $I''(p) = \mathtt{true}$ for every other atom p. Then $I'' \models L \cup \{\phi\}$ and $I'' \prec^\star I$. This means that I is not a $\prec^\star$-preferred model of $L \cup \{\phi\}$, which is a contradiction. It follows that α is unsatisfiable, and thus Φ evaluates to false. $\Box$

6 Related Work and Conclusion

In this paper, we have established a sharp picture of the complexity of major approaches to default reasoning from conditional knowledge bases.

A stream of semantics for conditional knowledge bases that we have not considered in this paper is inherited from conditional modal logics, cf. [15, 16, 8, 42, 21]. Roughly speaking, $\phi \to \psi$ evaluates in a Lewis-style system at a world w to true, if in a set of selected worlds in which ϕ is true (the most plausible, least exceptional ones etc) also ψ is true. Such logics allow nested use of $\to$, which is not the case in conditional knowledge bases. For example, Boutilier introduced 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 [40]; by his results, entailment from a conditional knowledge base is co-NP-complete under CT4D semantics. Friedman and Halpern analyzed the complexity of conditional modal logics based on different sets of modal axioms; their results imply that entailment of a default from a conditional knowledge base is co-NP-complete under a wide range of modal semantics [21].

Several issues remains for further work. One issue is a fine-grained analysis of the effect of preprocessing and fixing parameters in the input. In particular, the amenability of the various semantics to compilability according to the frameworks proposed in [11, 29], and the effect of fixed parameters in the input [17, 36] seem to be worthwhile to study.

Another issue is to identify tractable classes for the various approaches. Our results on literal-Horn knowledge bases suggest that for the "hard" semantics, severe further restrictions have to be imposed in order to gain tractability. Finally, further approaches to default reasoning (for example, the recent approach in [6]) may be analyzed from a complexity point of view.

References

[1] E. W. Adams. *The Logic of Conditionals*, volume 86 of *Synthese Library*. D. Reidel, Dordrecht, Netherlands, 1975.

[2] R. Ben-Eliyahu. NP-complete problems in optimal horn clauses satisfiability. Technical Report R-158, UCLA, Computer Science Dept., 1991.

[3] S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, pages 640–645. Morgan Kaufmann, 1993.

[4] S. Benferhat, D. Dubois, and H. Prade. Representing default rules in possibilistic logic. In *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, pages 673–684. Morgan Kaufmann, 1992.

[5] S. Benferhat, D. Dubois, and H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. *Artificial Intelligence*, 92(1–2):259–276, 1997.

[6] S. Benferhat, A. Saffiotti, and P. Smets. Belief functions and default reasoning. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, pages 19–26, 1995. (See also Technical Report TR/IRIDIA/95-05, Université Libre de Bruxelles, Belgium, 1998).

[7] R. A. Bourne and S. Parsons. Maximum entropy and variable strength defaults. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, 1999.

[8] C. Boutilier. Conditional logics of normality as modal systems. In *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 594–599, 1990.

[9] S. Buss and L. Hay. On truth-table reducibility to SAT. *Information and Computation*, 91:86–102, 1991.

[10] M. Cadoli. The complexity of model checking for circumscriptive formulae. *Information Processing Letters*, 44:113–118, 1992.

[11] M. Cadoli, F. M. Donini, P. Liberatore, and M. Schaerf. Feasibility and unfeasibility of off-line processing. In *Proceedings of the Fourth Israeli Symposium on Theory of Computing and Systems (ISTCS'96)*, pages 100–109. IEEE Computer Society, 1996.

[12] M. Cadoli and M. Lenzerini. The complexity of propositional closed world reasoning and circumscription. *Journal of Computer and System Sciences*, 43:165–211, 1994.

[13] C. Cayrol, M.-C. Lagasquie-Schiex, and T. Schiex. Nonmonotonic reasoning: From complexity to algorithms. *Annals of Mathematics and Artificial Intelligence*, 22(3–4):207–236, 1998.

[14] E. Dantsin, T. Eiter, G. Gottlob, and A. Voronkov. Complexity and expressive power of logic programming. In *Proceedings of the Twelth IEEE International Conference on Computational Complexity (CCC '97)*, pages 82–101, 1997.

[15] J. Delgrande. A first-order conditional logic for prototypical properties. *Artificial intelligence*, 33(1):105–130, 1987.

[16] J. Delgrande. On first-order conditional logics. *Artificial intelligence*, 105(1–2):105–137, 1998.

[17] R. G. Downey, M. F. Fellows, and M. R. Fellows. *Parameterized Complexity*. Springer, 1998.

[18] D. Dubois and H. Prade. Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, pages 419–424. Morgan Kaufmann, 1991.

[19] D. Dubois and H. Prade. Conditional objects as non-monotonic consequence relationships. *IEEE Transactions on Systems, Man and Cybernetics*, 24:1724–1740, 1994.

[20] 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.

[21] N. Friedman and J. Y. Halpern. On the complexity of conditional logics. In *Proceedings of the 4th International Conference on Knowledge Representation and Reasoning*, pages 202–213. Morgan Kaufmann, 1994.

[22] N. Friedman and J. Y. Halpern. Plausibility measures and default reasoning. *Journal of the ACM*, 1999.

[23] 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.

[24] D. M. Gabbay and P. Smets, editors. *Handbook on Defeasible Reasoning and Uncertainty Management Systems*. Kluwer Academic, Dordrecht, Netherlands, 1998.

[25] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, New York, 1979.

[26] H. Geffner. *Default Reasoning: Causal and Conditional Theories*. MIT Press, Cambridge, MA, 1992.

[27] H. Geffner. High probabilities, model preference and default arguments. *Mind and Machines*, 2:51–70, 1992.

[28] H. Geffner and J. Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence*, 53(2–3):209–244, 1992.

[29] G. Gogic, H. Kautz, C. H. Papadimitriou, and B. Selman. The comparative linguistics of knowledge representation. In C. Mellish, editor, *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 862–869. Morgan Kaufmann, 1995.

[30] 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.

[31] M. Goldszmidt and J. Pearl. On the consistency of defeasible databases. *Artificial Intelligence*, 52(2):121–149, 1991.

[32] M. Goldszmidt and J. Pearl. System Z^+: A formalism for reasoning with variable strength defaults. In *Proceedings of the 8th National Conference on Artificial Intelligence*, pages 399–404. Morgan Kaufmann, 1991.

[33] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, pages 661–672. Morgan Kaufmann, 1992.

[34] M. Goldszmidt and J. Pearl. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*, 84(1–2):57–112, 1996.

[35] G. Gottlob. Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2(3):397–425, June 1992.

[36] G. Gottlob, F. Scarcello, and M. Sideri. Fixed parameter complexity in AI and nonmonotonic reasoning. In *Proceedings of the 4th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR-99)*, 1999. To appear.

[37] B. Jenner and J. Toran. Computing functions with parallel queries to NP. *Theoretical Computer Science*, 141:175–193, 1995.

[38] D. S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 2, pages 67–161. MIT Press, Cambridge, MA, 1990.

[39] H. Kautz and B. Selman. Hard problems for simple default logics. *Artificial Intelligence*, 49:243–279, 1991.

[40] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 14(1):167–207, 1990.

[41] M. W. Krentel. The complexity of optimization problems. *Journal of Computer and System Sciences*, 36(3):490–509, 1988.

[42] P. Lamarre. S4 as the conditional logic of nonmonotonicity. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR-91)*, pages 357–367, 1991.

[43] P. Lamarre. A promenade from monotonicity to non-monotonicity following a theorem prover. In *Proceedings of the 3rd International Conference on Knowledge Representation and Reasoning*, pages 572–580. Morgan Kaufmann, 1992.

[44] D. Lehmann. What does a conditional knowledge base entail? In *Proceedings of the 1st International Conference on Knowledge Representation and Reasoning*, pages 212–222. Morgan Kaufmann, 1989.

[45] D. Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15(1):61–82, 1995.

[46] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, 1992.

[47] C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, Reading, MA, 1994.

[48] J. Pearl. Probabilistic semantics for nonmontonic reasoning: A survey. In *Proc. of the 1st International Conference on Knowledge Representation and Reasoning*, pages 505–516. Morgan Kaufmann, 1989.

[49] J. Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 121–135. Morgan Kaufmann, 1990.

[50] A. Selman. A taxonomy of complexity classes of functions. *Journal of Computer and System Sciences*, 48:357–381, 1994.

[51] Y. Shoham. A semantical approach to nonmonotonic logics. In *Proceedings of the 2nd IEEE Symposium on Logic in Computer Science*, pages 275–279, 1987.

[52] 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, pages 105–134. Reidel, Dordrecht, Netherlands, 1988.

[53] J. Stillman. It's not my default: The complexity of membership problems in restricted propositional default logic. In *Proc. of the 7th National Conference on Artificial Intelligence*, pages 571–579, 1990.

[54] J. Stillman. The complexity of propositional default logic. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 794–799, 1992.

[55] K. Wagner. Bounded query classes. *SIAM Journal of Computing*, 19(5):833–846, 1990.