Complexity Results for Default Reasoning from Conditional Knowledge Bases

Thomas Eiter and **Thomas Lukasiewicz**
Institut und Ludwig Wittgenstein Labor für Informationssysteme, TU Wien
Favoritenstraße 9-11, A-1040 Vienna, Austria
{eiter, lukasiewicz}@kr.tuwien.ac.at

Abstract

Conditional knowledge bases have been proposed as belief bases that include defeasible rules (also called defaults) of the form "$\phi \rightarrow \psi$", which informally read as "generally, if ϕ then ψ." Such rules may have exceptions, which can be handled in different ways. A number of entailment semantics for conditional knowledge bases have been proposed in the literature. However, while the semantic properties and interrelationships of these formalisms are quite well understood, about their 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 $\phi \rightarrow \psi$, does KB entail $\phi \rightarrow \psi$? We classify the complexity of this problem for a number of well-known approaches (including Goldszmidt et al.'s maximum entropy approach and Geffner's conditional entailment). 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 INTRODUCTION

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 [34].

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: They characterize classical model-theoretic entailment under preferential structures [45, 34], infinitesimal probabilities [1, 42], possibility measures [14], and world rankings [46, 27]. Moreover, they characterize an entailment relation based on conditional objects [15]. 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 [18], 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 [37, 39]. It is equivalent to entailment in system Z by Pearl [43] (which is generalized to variable strength defaults in system Z^+ by Goldszmidt and Pearl [26, 28]), to the least specific possibility entailment by Benferhat et al. [4], and to a conditional (modal) logic-based entailment by Lamarre [36].

Finally, mainly in order to solve problems with property inheritance from classes to exceptional subclasses, the max-

imum entropy approach to default entailment was proposed by Goldszmidt et al. [24] (and recently generalized to variable strength defaults by Bourne and Parsons [7]); the notion of lexicographic entailment was introduced by Lehmann [38] and Benferhat et al. [3]; the notion of conditional entailment was proposed by Geffner [20, 22]; 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 [25, 39, 28, 12]. 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 [25, 39, 28, 12].

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:

- We give a sharp characterization of the complexity of default reasoning from conditional knowledge bases under several semantics, improving on previous results. In particular, we address the following generic problem: Given a conditional knowledge base KB and a default $\phi \rightarrow \psi$, is it true that KB entails $\phi \rightarrow \psi$? Note that the precise formulation of this problem slightly varies in the different approaches and may involve further parameters.

- Our analysis includes formalisms for which only very rough or even no complexity results have been derived so far, namely proper ε-entailment [25], maximum entropy entailment [24] together with its variable-strength extension [7], and Geffner's conditional entailment [20].

- We reconsider Lehmann's lexicographic entailment [38], and correct an error of [12] 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 $\phi \rightarrow \psi$ with conjunctions of atoms ϕ and conjunctions of Horn clauses ψ, and the restriction to the literal-Horn case, where ψ is additionally a literal.

Our main findings can be briefly summarized as follows.

- The approaches considered in this paper cover different complexity classes at the low end of the polynomial hierarchy, which range from co-NP (ε-entailment) to Π_2^{P} (Geffner's conditional entailment). In general, they have lower complexity than well-known formalisms of nonmonotonic reasoning such as default logic, circumscription, or autoepistemic logic [48, 29, 16].

- As for the worst case, off-line computation of rankings does in general not pay off w.r.t. 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 [33, 47].

2 PRELIMINARIES

2.1 CONDITIONAL KNOWLEDGE BASES

We assume a set of *atoms* $At = \{p_1, \ldots, p_n\}$ with $n \geq 1$. By $\bot$ and $\top$ we denote the propositional constants *false* and *true*, respectively. The set of *classical formulas* is the closure of $At \cup \{\bot, \top\}$ under the Boolean operations $\neg$ and $\wedge$. We use $(\phi \Rightarrow \psi)$ and $(\phi \vee \psi)$ to abbreviate $\neg(\phi \wedge \neg\psi)$ and $\neg(\neg\phi \wedge \neg\psi)$, respectively, and adopt the usual conventions to eliminate parentheses. A *literal* is an atom p from At or its negation $\neg p$. A *Horn clause* is a classical formula $\phi \Rightarrow \psi$, where ϕ is either $\top$ or a conjunction of atoms, and ψ is either $\bot$ or 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 $\phi \to \psi$ 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 $\{\pi(d) \mid d \in D\} = \{0, 1, \ldots, k\}$ for some $k \geq 0$. Informally, a priority assignment is a strength assignment in which there are no "empty levels".

An *interpretation* (or *world*) is a truth assignment $I\colon At \to \{\textbf{true}, \textbf{false}\}$, 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 $\phi \to \psi$, or I is a *model* of $\phi \to \psi$, denoted $I \models \phi \to \psi$, 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 formula in K. The world I *verifies* a default $\phi \to \psi$, denoted $I \models_{\mathrm{v}} \phi \to \psi$, iff $I \models \phi \wedge \psi$. I *falsifies* a default $\phi \to \psi$, iff $I \models \phi \wedge \neg\psi$ (that is, $I \not\models \phi \to \psi$). 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 $\phi \to \psi$ iff all models of $D \cup L \cup \{\phi\}$ satisfy $\neg\psi$.

A *world ranking* κ is a mapping $\kappa\colon \mathcal{I}_{At} \to \{0, 1, \ldots\} \cup \{\infty\}$ 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 $\kappa(\phi) < \infty$ and $\kappa(\phi \wedge \psi) < \kappa(\phi \wedge \neg\psi)$ for all defaults $\phi \to \psi \in D$. A *default ranking* σ on D maps each $d \in D$ to an integer.

2.2 SEMANTICS FOR CONDITIONAL KNOWLEDGE BASES

We now recall some semantics for conditional knowledge bases. To simplify the presentation, we shall adjust original definitions (without significant effects) to our framework.

2.2.1 ε-Semantics (Adams [1] and Pearl [42])

We describe the notions of ε-consistency, ε-entailment, and proper ε-entailment in terms of world rankings (see especially Geffner's work [20, 21] 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 $\phi \to \psi$ iff either $\kappa(\phi) = \infty$ 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* the default $\phi \to \psi$ iff KB ε-entails $\phi \to \psi$ and KB does not ε-entail $\phi \to \bot$.

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

Theorem 2.1 (essentially [1]) *A conditional knowledge base (L, D) ε-entails a default $\phi \to \psi$ iff the conditional knowledge base $(L, D \cup \{\phi \to \neg\psi\})$ is ε-inconsistent.*

2.2.2 Systems Z and Z^+ (Pearl [43] and Goldszmidt and Pearl [26, 28])

The notions of entailment in systems Z and Z^+ apply to ε-consistent conditional knowledge bases $KB = (L, D)$. Entailment in system Z^+ additionally assumes a strength assignment σ on KB. It is associated with a default ranking z^+ and a world ranking κ^+, which are defined as the unique solution of the following system of equations: For all $d = \phi \to \psi \in D$ and all $I \in \mathcal{I}_{At}$:

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

$$\kappa^+(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ 1 + \max\limits_{d \in D\colon I \not\models d} z^+(d) & \text{otherwise.} \end{cases} \tag{2}$$

A default $\phi \to \psi$ is *z^+-entailed* by (KB, σ) *at strength* τ iff either $\kappa^+(\phi) = \infty$ or $\kappa^+(\phi \wedge \psi) + \tau < \kappa^+(\phi \wedge \neg\psi)$.

Finally, z-entailment is a special case of z^+-entailment: A default $\phi \to \psi$ is *z-entailed* by KB iff $\phi \to \psi$ is z^+-entailed by (KB, σ) at strength 0, where σ is given by $\sigma(d) = 0$ for all $d \in D$. In this special case, the rankings z^+ and κ^+ are denoted by z and κ^z, respectively.

2.2.3 Maximum Entropy Semantics (Goldszmidt et al. [24] and 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 the following system of equations: For all $d = \phi \to \psi \in D$ and all $I \in \mathcal{I}_{At}$:

$$z^\star(d) = 1 + \kappa^\star(\phi \wedge \psi) \tag{3}$$

$$\kappa^\star(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ \sum\limits_{d \in D\colon I \not\models d} z^\star(d) & \text{otherwise.} \end{cases} \tag{4}$$

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 $z^\star_{\rm s}$-entailment applies to ε-consistent conditional knowledge bases $KB = (L, D)$ with additional positive strength assignment σ. It is defined whenever the following system of equations (5) and (6) has a unique solution $z^\star_{\rm s}$ and $\kappa^\star_{\rm s}$. For all $\phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$:

$$\kappa^\star_{\rm s}(\phi \wedge \neg\psi) = \sigma(\phi \rightarrow \psi) + \kappa^\star_{\rm s}(\phi \wedge \psi) \quad (5)$$

$$\kappa^\star_{\rm s}(I) = \begin{cases} \infty & \text{if } I \not\models L \\ 0 & \text{if } I \models L \cup D \\ \sum\limits_{d \in D:\ I \not\models d} z^\star_{\rm s}(d) & \text{otherwise.} \end{cases} \quad (6)$$

The uniqueness of $z^\star_{\rm s}$ and $\kappa^\star_{\rm s}$ is guaranteed by assuming that $\kappa^\star_{\rm s}$ 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 $\kappa^\star_{\rm s}$ 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.

A default $\phi \rightarrow \psi$ is *$z^\star_{\rm s}$-entailed* by (KB, σ) *at strength* τ iff either $\kappa^\star_{\rm s}(\phi) = \infty$ or $\kappa^\star_{\rm s}(\phi \wedge \psi) + \tau \leq \kappa^\star_{\rm s}(\phi \wedge \neg\psi)$.

The notion of $z^\star_{\rm s}$-entailment is a proper generalization of $z^\star$-entailment:

Lemma 2.2 *Let $KB = (L, D)$ with strength assignment $\sigma(d) = 1$ for all $d \in D$ be ε-consistent and minimal-core. Then the system of equations given by (5) and (6) for all $\phi \rightarrow \psi \in D$ and all $I \in \mathcal{I}_{At}$ has a unique solution $z^\star_{\rm s}, \kappa^\star_{\rm s}$, which coincides with $z^\star, \kappa^\star$. Moreover, $\kappa^\star_{\rm s}$ is robust.*

2.2.4 Lexicographic Entailment (Lehmann [38] 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-preferable* 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 model I of a set of classical formulas $\mathcal{F}$ is a *π-preferred* model of $\mathcal{F}$ iff no model of $\mathcal{F}$ is π-preferable to I.

A default $\phi \rightarrow \psi$ is *lex$_{\rm p}$-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 [38] 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 2.2.2). We then say that a default $\phi \rightarrow \psi$ is *lex-entailed* by KB iff $\phi \rightarrow \psi$ is *lex$_{\rm p}$-entailed* by (KB, z). Note that this definition assumes that KB is ϵ-consistent.

It appears that, in a certain sense, lex-entailment is not less expressive than lex$_{\rm p}$-entailment. That is, under a weak condition, priority assignments are expressible through logical entrenchment:

Theorem 2.3 *Let $KB = (L, D)$ such that every $d \in D$ has a verifying world, and let π be a priority assignment on KB. Then, some $KB' = (L', D')$ and ϕ' exist such that, for any default $\phi \rightarrow \psi$ over At, it holds that (KB, π) lex$_{\rm p}$-entails $\phi \rightarrow \psi$ iff KB' lex-entails $\phi \wedge \phi' \rightarrow \psi$.*

2.2.5 Conditional Entailment (Geffner [20, 22])

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

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

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

2.3 COMPLEXITY CLASSES

We assume some familiarity with P, NP, and co-NP. We now briefly introduce some other classes that we encounter in our analysis (see [41, 32, 44, 31] for further background).

The class $\mathrm{P}^{\mathrm{NP}} = \Delta_2^{\mathrm{P}}$ (resp., $\mathrm{NP}^{\mathrm{NP}} = \Sigma_2^{\mathrm{P}}$) contains all decision problems that can be solved in deterministic (resp., nondeterministic) polynomial time with an oracle for NP. The class Π_2^{P} is the complementary class of Σ_2^{P}, which has Yes- and No-instances interchanged. The class D^{P} contains the problems that can be described as a logical conjunction of a problem in NP and a problem in co-NP. Any problem in D^{P} can be solved with two NP oracle calls, and is intuitively easier than a P^{NP}-complete problem. The class $\Delta_2^{\mathrm{P}}[O(\log n)]$ contains the problems in P^{NP} that can be solved with $O(\log n)$ many oracle calls, where n is the size of the problem input. This class coincides with $\mathrm{P}_{\parallel}^{\mathrm{NP}}$, that is, polynomial time computability with parallel NP oracle calls. According to the current belief in complexity theory, the following is a strict hierarchy of inclusions:

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

The classes FP, $\mathrm{FP}_{\|}^{\mathrm{NP}}$, and $\mathrm{FP}^{\mathrm{NP}}$ are the functional analogs of P, $\mathrm{P}_{\|}^{\mathrm{NP}}$, and P^{NP}, respectively.

Completeness for decision and function classes is w.r.t. standard polynomial time transformations (for the latter, see [35, 31]). In case of P and FP, completeness is w.r.t. logspace-reductions.

3 OVERVIEW OF COMPLEXITY RESULTS

3.1 PROBLEM STATEMENTS AND RESULTS

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 a strength assignment $\sigma(KB)$ and a strength $\tau(d)$, respectively. In case of lex_p-entailment, we assume that KB has a priority assignment $\pi(KB)$. We tacitly assume that KB satisfies any preconditions that the definition of entailment may request.

We analyze the 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 lex_p-entailment, we are asked whether d is 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 $\mathcal{S}$ (that is, the rank of each default in D).
- RANK-ENTAILMENT**:** Same as entailment, but the (unique) default ranking R of KB according to some fixed semantics $\mathcal{S}$ is part of the problem input.

The problems RANKING and RANK-ENTAILMENT are relevant from a preprocessing perspective, in which the ranking R of a conditional knowledge base is computed in advance and then on-line available in the input for solving an entailment problem.

Our results, together with results from the literature, are compactly summarized in Tables 1–3.

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 [29, 48, 16], which are all Π_2^{P}-complete. All other approaches in Table 1 have (considerably) lower complexity.

3.2 PREVIOUS RESULTS

Lehmann and Magidor [39] proved that preferential entailment (and thus also ε-entailment) for default knowledge bases is co-NP-complete. Moreover, Lehmann and Magidor [39] and Goldszmidt and Pearl [28] showed that ε-entailment for Horn default knowledge bases is in P. Finally, Goldszmidt and Pearl [25] proved that proper ε-entailment for general conditional knowledge bases (resp., Horn conditional knowledge bases) is in P^{NP} (resp., P).

As shown in [28], ENTAILMENT, RANKING, and RANK-ENTAILMENT in system Z^+ are in P, FP, and P, respectively, for Horn default knowledge bases, and in P^{NP}, $\mathrm{FP}^{\mathrm{NP}}$, and $\mathrm{P}_{\|}^{\mathrm{NP}}$, respectively, for general default knowledge bases. A fortiori, these upper bounds also hold for system Z.

Not much work has been spent determining the complexity of $z^\star$- and $z_{\mathrm{s}}^\star$-entailment. Goldszmidt et al. [24] 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.

Cayrol et al. show in [12] that lex_p-entailment is P^{NP}-complete for default knowledge bases. Moreover, they state in [12] that lex_p-entailment is P^{NP}-hard for Horn default knowledge bases. However, the short proof sketch in [12] 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. P^{NP}-hardness for the Horn case — but not the literal-Horn case — can be concluded from proofs for related recent results on the complexity of lexicographic belief revision [40].

To our knowledge, there are no complexity results on Geffner's conditional entailment so far.

Note that under a number of semantics for conditional knowledge bases inherited from conditional modal logics, entailment is known to be co-NP-complete, cf. [8, 19].

3.3 DISCUSSION

3.3.1 General Case

At the low end of the complexity range, there are ε-entailment, which has the same complexity as classical logic, and proper ε-entailment, which has marginally higher complexity due to the additional ε-entailment requirement. At the high end, we have Geffner's conditional entailment. Its high complexity is intuitively explained by

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
$z_{\mathrm{s}}^{\star}$-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
$z_{\mathrm{s}}^{\star}$	$\mathrm{FP}^{\mathrm{NP}}$-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
$z_{\mathrm{s}}^{\star}$-entailment	P^{NP}-complete	P^{NP}-complete	P^{NP}-complete

* Membership shown in [25].
** Membership shown in [39, 28] for default knowledge bases.
*** Membership was shown in [28] for default knowledge bases.
$^{+}$ Shown in [39] for default knowledge bases. Note that deciding ε-consistency is NP-complete in the general case [17].
$^{++}$ Reported in [12] for default knowledge bases; the proof sketch for the Horn case in [12] shows merely $\mathrm{P}_{\parallel}^{\mathrm{NP}}$-hardness.

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 (i.e., minimality under $\prec$). This is similar to circumscription, i.e., 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 [9].

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, i.e., nei-

ther $I_1 \prec I_2$ 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 to 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, as results on the Horn restrictions (discussed below) show.

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

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

Algorithms for computing R in Step 1 have been described in the literature. They can be reformulated to run in P^{NP} (see [17]). 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 resembles the $\mathrm{FP}^{\mathrm{NP}}$-complete problem of computing the lexicographic maximum model of a formula ϕ [35] and suggests that computing R has the same complexity. This is true in all cases except one. For z-entailment, exploiting the fact that the sum of all default ranks is minimal at z, this can be done in polynomial time using parallel NP queries (see [17]).

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

3.3.2 Horn and Literal-Horn Case

In all tables, the results for the Horn and the literal-Horn case are the same. Thus, although it is in general not possible to simply split a Horn default $\phi \rightarrow \psi_1 \wedge \cdots \wedge \psi_m$ 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 [25, 39, 28]. 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 [33, 47, 11]. Informally, in Geffner's conditional entailment, the Horn property is not sufficient to eliminate the intractability of the preference check for a model of a formula, as this test is not simply reducible to a polynomial number of Horn satisfiability tests, as e.g. in Reiter's default logic. Further syntactic restrictions are needed. One such restriction, which is efficiently checkable, is that the admissibility ordering $\prec$ is 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 [24], 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.

4 DERIVATION OF RESULTS

Due to space reason, we focus here on maximum entropy semantics and conditional entailment. Detailed proofs of all results are given in the full paper [17].

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

Input: ε-consistent $KB = (L, D)$ with positive strength assignment σ.
Output: Default ranking $z^\star_{\rm s}$, if the system of equations (5) and (6), for all $\phi \to \psi \in D$ and $I \in \mathcal{I}_{At}$, 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 (6).

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 (5) for all $\phi \to \psi \in D$ **and** $\kappa^\star_{\rm s}$ is robust
9. **then return** $z^\star_{\rm s}$ **else return** *nil*.

Figure 1: Algorithm $z^\star_{\rm s}$**-ranking**

4.1 MAXIMUM ENTROPY SEMANTICS

The key for the P^{NP}-upper bound of $z^\star_{\rm s}$-entailment (hence, also $z^\star$-entailment) is algorithm $z^\star{}_{\rm s}$**-ranking** (see Fig. 1), adapted from [7]. It intializes all default ranks $z^\star_{\rm s}(d)$ to ∞ and then ranks one default at a time, using (6) and the current $z^\star_{\rm s}$. If the process fails, a special value *nil* is returned.

We 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 whether $\kappa^\star_{\rm s}$ satisfies (5) in Step 8 is in P^{NP} for each $\phi \to \psi \in D$; and (iii) deciding whether $\kappa^\star_{\rm s}$ is robust in Step 8 is in P^{NP}.

As for (i) and (ii), an inductive argument shows that $|z^\star_{\rm s}(d)| \leq s\, 2^{|D|}$, where $s = \max_{d \in D} \sigma(d)$, holds for each $d \in D$ that is assigned a value $< \infty$ (see Appendix). Thus, 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 $\{0, \ldots, s \cdot 2^n\} \cup \{\infty\}$. By binary search, they are polynomially computable with $O(n + \log s)$ calls to an NP-oracle, since given w, deciding whether some $I \in \mathcal{I}_{At}$ satisfies $\kappa^\star_{\rm s}(I) \leq w$ is in NP. Thus, (i) and (ii) are in $\mathrm{FP}^{\mathrm{NP}}$ and P^{NP}, respectively. For (iii), we test whether there are no $d_1 \neq d_2 \in D$ that have a common minimal falsifying model under L. If $\mathrm{minf}(d)$ is known for all $d \in D$—which can be computed in $\mathrm{FP}^{\mathrm{NP}}$—this is clearly in co-NP, and thus (iii) is in P^{NP}. Overall, computing the $z^\star_{\rm s}$-ranking of KB is thus in $\mathrm{FP}^{\mathrm{NP}}$.

Now consider deciding whether KB $z^\star_{\rm s}$-entails a default $\phi \to \psi$, once $z^\star_{\rm s}$ is known. Testing $\kappa^\star_{\rm s}(\phi) = \infty$ is in NP, and by similar arguments as above, computing $\kappa^\star_{\rm s}(\phi \wedge \psi)$ and $\kappa^\star_{\rm s}(\phi \wedge \neg\psi)$ in a binary search from $z^\star_{\rm s}$ is in $\mathrm{FP}^{\mathrm{NP}}$. Finally, testing $\kappa^\star_{\rm s}(\phi \wedge \psi) < \kappa^\star_{\rm s}(\phi \wedge \neg\psi)$ is simple. Put together, deciding whether KB $z^\star_{\rm s}$-entails d is in P^{NP}.

For the matching lower bound, we reduce the following P^{NP}-complete problem to $z^\star$-entailment. Given a set of weighted Horn clauses $C = \{\phi_i \mid 1 \leq i \leq m\}$ such that every ϕ_i is satisfiable and has weight $w_i = 2^{c_i}$, where $c_i \geq 0$ is an integer, and some $r \in \{1, \ldots, m\}$, decide whether $I \models \phi_r$ for every maximum weight world I for C, i.e., world $I \in \mathcal{I}_{At}$ such that $\sum_{I \models \phi_i} w_i$ is maximum. P^{NP}-hardness of this problem follows by a minor adaptation of the proof of Theorem 2.1 in [35].

We construct a literal-Horn conditional knowledge base KB and a literal-Horn default d such that $I \models \phi_r$ holds for every maximum weight world I of C iff KB $z^\star$-entails d (see Appendix). Informally, the idea behind our construction is as follows. For each Horn clause ϕ_i with weight 2^{c_i}, we introduce a default $d_{c_i,i}$. By additional Horn clauses and literal-Horn defaults in KB, we then ensure that $z^\star(d_{c_i,i}) = 2^{c_i}$ and that $d_{c_i,i}$ is falsified by a model I of some conjunction of atoms ϕ iff $I \not\models \phi_i$. Moreover, we ensure that all the other defaults in KB are falsified by every model of ϕ. This means 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 $d = \phi \to \psi$ is $z^\star$-entailed by KB iff ϕ_r holds in all minimal falsifying models of ϕ.

$\mathrm{FP}^{\mathrm{NP}}$-hardness of computing the $z^\star$-ranking for KB follows from an adaptation (see Appendix).

4.2 CONDITIONAL ENTAILMENT

As for the Π^{P}_2 upper bound, recall from the definitions that KB does not conditionally entail d iff there exist a priority ordering $\prec$ on D 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** (see Fig. 2). The crucial point there is Step 3, which exploits the following lemma (see Appendix):

Lemma 4.1 *A priority ordering $\prec$ on D is admissible with a conditional knowledge base $KB = (L, D)$ iff each $d \in D$ is tolerated by $D_d = D - \{d' \in D \mid d' \prec d\}$ under L.*

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".

Figure 2: Algorithm **not-cond-entailment**

It is easy to see that testing the condition in the lemma for a given $d \in D$ is in NP. Thus, all queries in Step 3 and the query in Step 6 (see Fig. 2) can be solved using an NP-oracle. Modulo these queries, each step in **not-cond-entailment** is polynomial. Hence, deciding whether KB does not conditionally entail d is in $\mathrm{NP}^{\mathrm{NP}} = \Sigma_2^{\mathrm{P}}$. Therefore, deciding conditional entailment is in $\Pi_2^{\mathrm{P}} = \text{co-}\Sigma_2^{\mathrm{P}}$.

To show that this upper bound is tight, we give a polynomial transformation from the following canonical Π_2^{P}-complete problem [32]. 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 (QBF) $\Phi = \forall y_1 \ldots \forall y_m \exists x_1 \ldots \exists x_n \, (\alpha_1 \wedge \cdots \wedge \alpha_l)$ evaluates to true.

We construct a literal-Horn $KB = (L, D)$ and a literal-Horn default $d = \phi \to \psi$ such that Φ evaluates to true iff KB conditionally entails d (see Appendix). Roughly speaking, it solves two major problems. Firstly, expressing the validity of a QBF. This is done exploiting the preference ordering on worlds. Secondly, transforming the clauses $\alpha_1, \ldots, \alpha_l$ into Horn clauses $\alpha_1^\star, \ldots, \alpha_l^\star$. For that, all positive literals y_i and x_j in $\alpha_1, \ldots, \alpha_l$ are replaced by new negative literals $\neg y_i'$ and $\neg x_j'$, respectively. We use L to express $y_i \wedge y_i' \Rightarrow \bot$ and $x_j \wedge x_j' \Rightarrow \bot$, and the preference ordering on worlds to express $\top \Rightarrow y_i \vee y_i'$ and $\top \Rightarrow x_j \vee x_j'$; thus, $y_i \equiv \neg y_i'$ and $x_j \equiv \neg x_j'$.

5 SUMMARY AND CONCLUSION

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

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 [10, 23], and the effect of fixed parameters in the input [13, 30] 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 to gain tractability. Finally, other approaches to default reasoning (e.g., the recent one in [6]) may be analyzed from a complexity point of view.

APPENDIX: SELECTED PROOFS

Detailed proofs of all results are given in the full paper [17].

A MAXIMUM ENTROPY SEMANTICS

Lemma A.1 *Let $KB = (L, D)$ be ε-consistent with positive strength assignment σ. Let $s = \max(\{\sigma(d) \mid d \in D\})$. Let $z_{\mathrm{s}}^\star(d_1), \ldots, z_{\mathrm{s}}^\star(d_l)$, with $l \leq |D|$, be the sequence of default ranks computed in algorithm $z_{\mathrm{s}}^\star$-ranking. Then $|z_{\mathrm{s}}^\star(d_i)| \leq s \cdot 2^{i-1}$, for $1 \leq i \leq l$.*

Proof. The proof is by induction on $i = 1, \ldots, l$. Let $n = |D|$ be the cardinality of D.

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

$$|z_{\mathrm{s}}^\star(d_i)| \leq \sigma(d_i) + |\mathrm{minv}(d_i) - \mathrm{minf}(d_i)| \leq s + \textstyle\sum_{j=1}^{i-1} s \cdot 2^{j-1} = s + s \cdot (2^{i-1} - 1)) = s \cdot 2^{i-1}. \; \square$$

Theorem A.2 *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 reduction from the P^{NP}-complete problem on weighted Horn clauses $C = \{\phi_i \mid 1 \leq i \leq m\}$ in Section 4.1. Suppose $\phi_i = \alpha_i \Rightarrow \beta_i$ and recall it has weight $w_i = 2^{c_i}$.

The set of atoms is $At = \{x_1, \ldots, x_n\} \cup A \cup B \cup T$, 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:

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

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

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

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

$$\phi \;=\; \bigwedge_{p \in A \cup T} p, \qquad \psi \;=\; b_{c_k,\,k}\,.$$

We prove the following:

- KB is ε-consistent. This is shown using the fact (essentially proved in [25, 28]) that any $KB = (L, D)$ 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 $\bigcup_{j=i}^{k} D_j$.

  Let $(D_0, \ldots, D_k)$ 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 $d_{i,j} \in D_i$. Let I be a world such that $I \models \alpha_j \Rightarrow \beta_j$, $I \models a_{l,j}$, for all $l \leq i$, $I \models b_{i,j}$, $I \models t_{l,j}$, for all $l \leq i-1$, and I falsifies all remaining atoms $a_{i',j'}$, $b_{i',j'}$, and $t_{i',j'}$. Clearly such a world I exists. As easily seen, I verifies $d_{i,j}$ and $I \models L \cup \bigcup_{l=i}^{k} D_l$. Hence, $d_{i,j}$ is tolerated under L by $\bigcup_{l=i}^{k} D_l$. Thus, KB is ε-consistent.

- KB is minimal-core. Indeed, the world I such that $I \models a_{i,j}$ and $I \not\models p$ for any other atom p, falsifies the default $d_{i,j}$ while it satisfies $L \cup (D - \{d_{i,j}\})$.

- $z^\star(d_{i,j}) = 2^i$, for all $d_{i,j} \in D$ (thus, $z^\star(d_{c_j,j}) = 2^{c_j}$ for $j \leq m$). Straightforward by 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 show that $\kappa^\star(I) < \kappa^\star(I')$. As easily seen, 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 $j \leq n$. Then, $I'' \models L$, and no default $d_{i,j}$ 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 $j \leq m$. Hence,

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

  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 $\kappa^\star(I) < \kappa^\star(I')$.

  Assume first 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.

  Assume next 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. □

Theorem A.3 *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 A.2, which is $\mathrm{FP}^{\mathrm{NP}}$-complete: Given C as there, compute the weight w of a maximum weight world I of C.

We slightly extend KB in the proof of Theorem A.2 as follows. We introduce new atoms $a^\star$ and $b^\star$, and the following set of literal-Horn clauses $L^\star$ and literal-Horn default $d^\star$:

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

By similar arguments as in the proof of Theorem A.2, it is easy to see that $KB' = (L \cup L^\star, D \cup \{d^\star\})$ 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 $\sum_{j=1}^{m} \sum_{i=0}^{c_j} 2^i - w$. Hence, 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 $z^\star$. □

B CONDITIONAL ENTAILMENT

Proof of Lemma 4.1. Assume that $\prec$ is admissible with KB, and let $d \in D$. Admissibility of $\prec$ implies that D_d is under L not in conflict with d (i.e., it tolerates d under L).

Conversely, assume every $d \in D$ is tolerated under L by D_d. Suppose that $\prec$ is not admissible with KB, i.e., 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 that D' is under L in conflict with d. Hence, $\prec$ is admissible with KB. □

Theorem B.4 *Given a literal-Horn conditional knowledge base $KB = (L, D)$ and a literal-Horn default $d = \phi \rightarrow \psi$, deciding whether KB conditionally entails d is Π_2^{P}-hard.*

Proof. (Sketch) We build a literal-Horn $KB = (L, D)$ and a literal-Horn default $d = \phi \rightarrow \psi$ such that the QBF Φ from Section 4.2 evaluates to true iff KB conditionally entails d.

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:

$$L \;=\; L_1 \cup L_2 \cup L_3 \cup L_4, \quad D \;=\; \bigcup_{i=1}^{m} D_{1,i} \cup \bigcup_{j=1}^{n} D_{2,j},$$

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:

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

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

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

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 the following lemma, we can restrict our attention to all minimal priority orderings $\prec$ on D admissible with KB, where $\prec$ is smaller than $\prec'$ iff $\{(I, J) \mid I \prec J\}$ is a proper subset of $\{(I, J) \mid I \prec' J\}$.

Lemma (Geffner [20]). *A default $\phi \rightarrow \psi$ is conditionally entailed by a conditional knowledge base $KB = (L, D)$ iff ψ is satisfied in every $\prec$-preferred model of $L \cup \{\phi\}$ of every minimal priority ordering $\prec$ admissible with KB.* ◇

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

$$\begin{aligned}
&c_j \rightarrow f_j \;\prec\; c_j \wedge d_j \rightarrow x_j\\
&\text{and } c_j \rightarrow f_j \;\prec\; c_j \wedge e_j \rightarrow x'_j, \text{ for } j = 1, \ldots, n.
\end{aligned} \tag{7}$$

Indeed, 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 $\prec^\star$ contain exactly all pairs in (7). By Lemma 4.1, it follows that $\prec^\star$ is admissible with KB. That is, $\prec^\star$ is the least priority ordering on D admissible with KB. Thus, by Geffner's lemma, KB conditionally entails $d = \phi \rightarrow \psi$ iff $I \models \psi$ for every $\prec^\star$-preferred model I of $L \cup \{\phi\}$.

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

($\Leftarrow$) Assume that Φ evaluates to false. Hence, a mapping $f : \{y_1, \ldots, y_m\} \rightarrow \{\bot, \top\}$ exists such that $\alpha = (\alpha_1 \wedge \cdots \wedge \alpha_l)\,[y_1/f(y_1), \ldots, y_m/f(y_m)]$ is unsatisfiable. Let I be such that (i) $I(y_i) = I(\neg y'_i) =$ **true** iff $f(y_i) = \top$, for all $i \leq m$; (ii) $I(x_j) = I(x'_j) =$ **false**, for all $j \leq n$; and (iii) $I(p) =$ **true** for any other atom p. Since each $\alpha_i^\star$ contains at least one negative literal from $\{\neg x_j,\ \neg x'_j \mid 1 \leq j \leq n\}$, 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, $I \models L \cup \{\phi, \neg\psi\}$. It can be shown [17] that no world J exists such that $J \models L \cup \{\phi\}$ and $J \prec^\star I$. Hence, I is a $\prec^\star$-preferred model of $L \cup \{\phi\}$. Thus, KB does not conditionally entail d.

($\Rightarrow$) Assume that KB does not conditionally entail d. Thus, a $\prec^\star$-preferred model I of $L \cup \{\phi\}$ exists 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 $j \leq n$; preferredness of I implies that $I(f_j) =$ **true** (as $c_j \rightarrow f_j$ will be satisfied), for all $j \leq n$. Let $f : \{y_1, \ldots, y_m\} \rightarrow \{\bot, \top\}$ be defined by $f(y_i) = \top$ iff $I \models a_i \rightarrow y_i$. We 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 some world I' on $x_1, \ldots, x_n$ exists such that $I' \models \alpha$. Let I'' coincide on all y_i, y'_i with I and such that $I''(x_j) = I''(\neg x'_j) = I'(x_j)$; $I''(f_j) =$ **false**, for all $j \leq n$; and $I''(p) =$ **true** for every other atom p. Then $I'' \models L \cup \{\phi\}$ and $I'' \prec^\star I$. Hence, I is not a $\prec^\star$-preferred model of $L \cup \{\phi\}$, which is a contradiction. Consequently, α is unsatisfiable, and thus Φ evaluates to false. □

Acknowledgements

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

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 IJCAI-93*, pages 640–645. 1993.

[4] S. Benferhat, D. Dubois, and H. Prade. Representing default rules in possibilistic logic. In *Proceedings KR'92*, pages 673–684. 1992.

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

[6] S. Benferhat, A. Saffiotti, and P. Smets. Belief functions and default reasoning. In *Proceedings UAI'95*, pages 19–26. 1995.

[7] R. A. Bourne and S. Parsons. Maximum entropy and variable strength defaults. In *Proceedings IJCAI-99*, pages 50–55. 1999.

[8] C. Boutilier. Conditional logics of normality as modal systems. In *Proceedings AAAI-90*, pages 594–599, 1990.

[9] M. Cadoli. The complexity of model checking for circumscriptive formulae. *Inf. Process. Lett.*, 44:113–118, 1992.

[10] M. Cadoli, F. M. Donini, P. Liberatore, and M. Schaerf. Feasibility and unfeasibility of off-line processing. In *Proceedings ISTCS'96*, pages 100–109. 1996.

[11] M. Cadoli and M. Lenzerini. The complexity of propositional closed world reasoning and circumscription. *J. Comput. Syst. Sci.*, 43:165–211, 1994.

[12] C. Cayrol, M.-C. Lagasquie-Schiex, and T. Schiex. Nonmonotonic reasoning: From complexity to algorithms. *Ann. Math. Artif. Intell.*, 22(3–4):207–236, 1998.

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

[14] D. Dubois and H. Prade. Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proceedings IJCAI-91*, pages 419–424. 1991.

[15] D. Dubois and H. Prade. Conditional objects as nonmonotonic consequence relationships. *IEEE Trans. Syst. Man Cybern.*, 24:1724–1740, 1994.

[16] T. Eiter and G. Gottlob. Propositional circumscription and extended closed world reasoning are Π_2^P-complete. *Theor. Comput. Sci.*, 114(2):231–245, 1993. Addendum 118:315.

[17] T. Eiter and T. Lukasiewicz. Complexity results for default reasoning from conditional knowledge bases. Technical Report INFSYS RR-1843-99-10, Institut für Informationssysteme, TU Wien, 1999.

[18] N. Friedman and J. Y. Halpern. Plausibility measures and default reasoning. *J. ACM*. To appear.

[19] N. Friedman and J. Y. Halpern. On the complexity of conditional logics. In *Proceedings KR'94*, pages 202–213. 1994.

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

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

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

[23] G. Gogic, H. Kautz, C. H. Papadimitriou, and B. Selman. The comparative linguistics of knowledge representation. In *Proceedings IJCAI-95*, pages 862–869. 1995.

[24] M. Goldszmidt, P. Morris, and J. Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 15(3):220–232, 1993.

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

[26] M. Goldszmidt and J. Pearl. System Z^+: A formalism for reasoning with variable strength defaults. In *Proceedings AAAI-91*, pages 399–404. 1991.

[27] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Proceedings KR'92*, pages 661–672. 1992.

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

[29] G. Gottlob. Complexity results for nonmonotonic logics. *J. Log. Comput.*, 2(3):397–425, June 1992.

[30] G. Gottlob, F. Scarcello, and M. Sideri. Fixed parameter complexity in AI and nonmonotonic reasoning. In *Proceedings LPNMR'99* , pages 1–18, 1999.

[31] B. Jenner and J. Toran. Computing functions with parallel queries to NP. *Theor. Comput. Sci.*, 141:175–193, 1995.

[32] 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.

[33] H. Kautz and B. Selman. Hard problems for simple default logics. *Artif. Intell.*, 49:243–279, 1991.

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

[35] M. W. Krentel. The complexity of optimization problems. *J. Comput. Syst. Sci.*, 36(3):490–509, 1988.

[36] P. Lamarre. A promenade from monotonicity to nonmonotonicity following a theorem prover. In *Proceedings KR'92*, pages 572–580. 1992.

[37] D. Lehmann. What does a conditional knowledge base entail? In *Proceedings KR'89*, pages 212–222. 1989.

[38] D. Lehmann. Another perspective on default reasoning. *Ann. Math. Artif. Intell.*, 15(1):61–82, 1995.

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

[40] B. Nebel. How hard is it to revise a belief base? In D. M. Gabbay and P. Smets, editors, *Handbook on Defeasible Reasoning and Uncertainty Management Systems*, volume 3: Belief Change, pages 77–145. Kluwer Academic, Dordrecht, Netherlands, 1998.

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

[42] J. Pearl. Probabilistic semantics for nonmontonic reasoning: A survey. In *Proceedings KR'89*, pages 505–516. 1989.

[43] J. Pearl. System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings TARK'90* , pages 121–135. 1990.

[44] A. Selman. A taxonomy of complexity classes of functions. *J. Comput. Syst. Sci.*, 48:357–381, 1994.

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

[46] 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. D. Reidel, Dordrecht, 1988.

[47] J. Stillman. It's not my default: The complexity of membership problems in restricted propositional default logic. In *Proceedings AAAI-90*, pages 571–579, 1990.

[48] J. Stillman. The complexity of propositional default logic. In *Proceedings AAAI-92*, pages 794–799, 1992.