On the Computational Complexity of Assumption-based Argumentation for Default Reasoning

Yannis Dimopoulos
University of Cyprus
Dept. of Computer Science
yannis@cs.ucy.ac.cy

Bernhard Nebel
Universität Freiburg
Institut für Informatik
nebel@uni-freiburg.de

Francesca Toni
Imperial College
Dept. of Computing
ft@doc.ic.ac.uk

Abstract

Bondarenko *et al.* have recently proposed an abstract framework for default reasoning. Besides capturing most existing formalisms and proving that their standard semantics all coincide, the framework extends these formalisms by generalising the semantics of *admissible* and *preferred arguments*, originally proposed for logic programming only.

In this paper we analyse the computational complexity of *credulous* and *sceptical* reasoning under the semantics of admissible and preferred arguments for (the propositional variant of) the instances of the abstract framework capturing theorist, circumscription, logic programming, default logic, and autoepistemic logic. Although the new semantics have been tacitly assumed to mitigate the computational hardness of default reasoning under the standard semantics of *stable extensions*, we show that in many cases reasoning under the admissibility and preferability semantics is computationally harder than under the standard semantics. In particular, in the case of autoepistemic logic, sceptical reasoning under preferred arguments is located at the fourth level of the polynomial hierarchy, whereas the same form of reasoning under stable extensions is located at the second level.

1 Introduction

Bondarenko *et al.* [1] show that many logics for default reasoning, i.e., theorist [25], (many cases of) circumscription [20], Default Logic (DL) [26], Nonmonotonic Modal Logic [21], Autoepistemic Logic (AEL) [22], and Logic Programming (LP) can be understood as special cases of a single abstract framework. The standard semantics of all these logics can be understood as sanctioning a set of assumptions as a *stable extension* of a given theory, formulated in an underlying monotonic logic, iff the set of assumptions does not *attack* itself and it attacks every assumption not in the set. In abstract terms, an assumption can be attacked if its *contrary* can be proved, in the underlying monotonic logic, possibly with the aid of other conflicting assumptions.

Bondarenko *et al.* also propose two new semantics generalising, respectively, the admissibility semantics [8] and the semantics of preferred extensions [8] or partial stable models [27] for LP. In abstract terms, a set of assumptions is an *admissible argument* of a given theory, formulated in an underlying monotonic logic, iff it does not attack itself and it attacks all sets of assumptions which attack it. A set of assumptions is a *preferred argument* iff it is a maximal (with respect to set inclusion) admissible argument.

The new semantics are more general than the *stability semantics* since every stable extension is a preferred (and admissible) argument, but not every preferred argument is a stable extension. Moreover, the new semantics are more liberal because for most concrete logics for default reasoning, admissible and preferred arguments are always guaranteed to exist, whereas stable extensions are not. Finally, reasoning under the new semantics appears to be computationally easier than reasoning under the stability semantics. Intuitively, to show that a given sentence is justified by a stable extension, it is necessary to perform a global search amongst all the assumptions, to determine for each such assumption whether it or its contrary can be derived, independently of the sentence to be justified.[1] For the semantics of admissible and preferred arguments, however, a "local" search suffices. First, one has to construct a set of assumptions which, together with the given theory, (monotonically) derives the sentence to be justified, and then one has to augment the constructed set with further assumptions to defend it against all attacks [18, 6, 7].

[1] See [18, 7] for a more general discussion of the problems associated with computing the stability semantics.

However, from a complexity-theoretic point of view, it seems unlikely that the new semantics lead to better lower bounds than the standard semantics since all the "sources of complexity" one has in default reasoning are still present. There are potentially exponentially many assumption sets sanctioned by the semantics. Further, in order to test whether a sentence is entailed by a particular argument, one has to reason in the underlying monotonic logic. For this reason, one would expect that reasoning under the new semantics has the same complexity as under the stability semantics, i.e., it is complete for the first level of the polynomial hierarchy for LP and on the second level for logics with full propositional logic as the underlying logic [3]. However, previous results on the expressive power of DATALOG$^{\neg}$ queries by Saccà [28] suggest that this is not the case for LP. Indeed, Saccà's results imply that reasoning under the *preferability semantics* for LP is at the second level of the polynomial hierarchy.

In this paper we extend this analysis and provide complexity results for reasoning in the propositional variants of theorist, circumscription, LP, DL, and AEL under the new semantics. As it turns out, reasoning under the new semantics can be much harder than reasoning under the standard semantics. In particular, we show that sceptical reasoning in DL under the preferability semantics is on the third level of the polynomial hierarchy, that credulous reasoning in AEL under the *admissibility semantics* is on the third level of the polynomial hierarchy, and that sceptical reasoning in AEL under the preferability semantics is on the fourth level of the polynomial hierarchy.

The paper[2] is organised as follows. Section 2 summarises relevant features of the abstract framework of [1], its semantics and concrete instances. Section 3 gives complexity theory background and introduces the reasoning problems. Section 4 gives generic upper bounds for credulous and sceptical reasoning, parametric with respect to the complexity of the underlying monotonic logics. The generic results are instantiated to provide upper bounds for concrete instances of the abstract framework. Section 5 gives then completeness results for theorist and circumscription, Section 6 gives the completeness results for LP and DL, and Section 7 gives the completeness results for AEL. Section 8 discusses the results and concludes.

[2]This paper combines earlier papers [4, 5] by the same authors and it contains all formal proofs of the results in full.

2 Default Reasoning via Argumentation

Assume a **deductive system** ($\mathcal{L}$, $\mathcal{R}$), where $\mathcal{L}$ is some formal language with countably many sentences and $\mathcal{R}$ is a set of inference rules inducing a monotonic derivability notion $\vdash$. Given a theory $T \subseteq \mathcal{L}$ and a formula $\alpha \in \mathcal{L}$, $Th(T) = \{\alpha \in \mathcal{L} \mid T \vdash \alpha\}$ is the deductive closure of T. Then, an **abstract (assumption-based) framework** is a triple $\langle T, A, \overline{} \rangle$, where $T, A \subseteq \mathcal{L}$ and $\overline{}$ is a mapping from A into $\mathcal{L}$. T, the **theory**, is a set of beliefs, formulated in the underlying language, and can be extended by subsets of A, the set of **assumptions**. Indeed, an **extension** of an abstract framework $\langle T, A, \overline{} \rangle$ is a theory $Th(T \cup \Delta)$, with $\Delta \subseteq A$ (sometimes an extension is referred to simply as $T \cup \Delta$ or Δ). Finally, given an assumption $\alpha \in A$, $\overline{\alpha}$ denotes the **contrary** of α.

Theorist can be understood as a framework $\langle T, A, \overline{} \rangle$ where T and A are both arbitrary sets of sentences of classical (first-order or propositional) logic and the contrary $\overline{\alpha}$ of an assumption α is just its negation. $\vdash$ is ordinary classical provability.

Many cases of *circumscription*[3] can be understood similarly, except that the assumptions are negations of atomic sentences $\neg p(t)$, for all predicates p which are minimised, and atomic sentences $q(t)$ or their negations, for all predicates q which are fixed.

LP is the instance of the abstract framework $\langle T, A, \overline{} \rangle$ where T is a logic program, the assumptions in A are all negations $not\, p$ of atomic sentences p, and the contrary $\overline{not\, p}$ of an assumption is p. $\vdash$ is Horn logic provability, with assumptions, $not\, p$, understood as new atoms p^*, as in [12].

DL is the instance of the abstract framework $\langle T, A, \overline{} \rangle$ where the monotonic logic is first-order logic augmented with domain-specific inference rules of the form

$$\frac{\alpha_1, \ldots, \alpha_m, M\beta_1, \ldots, M\beta_n}{\gamma}$$

where α_i, β_j, γ are sentences in classical logic. T is a classical theory and A consists of all expressions of the form $M\beta$ where β is a sentence of classical logic. The contrary $\overline{M\beta}$ of an assumption $M\beta$ is $\neg\beta$.

AEL has, as the underlying language $\mathcal{L}$, a modal logic with a modal operator L, but the inference rules are those of classical logic. The assumptions have the form $\neg L\alpha$ or $L\alpha$. The contrary of $\neg L\alpha$ is α, and the contrary of $L\alpha$ is $\neg L\alpha$.

[3] Namely, all cases where every model of the theory to be circumscribed is a Herbrand model of the theory, see [1] for more details.

Given an abstract framework $\langle T, A, ^{-}\rangle$ and an assumption set $\Delta \subseteq A$:

- Δ **attacks an assumption** $\alpha \in A$ iff $\overline{\alpha} \in Th(T \cup \Delta)$;
- Δ **attacks an assumption set** $\Delta' \subseteq A$ iff Δ attacks some assumption $\alpha \in \Delta'$.

Given that an assumption set $\Delta \subseteq A$ is **closed** iff $\Delta = A \cap Th(T \cup \Delta)$, the standard semantics of extensions of theorist [25], minimal models of circumscription [20], extensions of DL [26], stable expansions of AEL [22], and stable models of LP [14] correspond to the stability semantics of abstract frameworks, where an assumption set $\Delta \subseteq A$ is **stable** iff

1. Δ is closed,
2. Δ does not attack itself, and
3. Δ attacks each assumption $\alpha \notin \Delta$.

Assumption sets are always closed in the case of LP and DL. Frameworks with this property are referred to as **flat** [1]. Assumption sets might not be closed in the case of AEL. For example, given the theory $T = \{\neg Lp\}$ in AEL, the empty assumption set is not closed. Furthermore, assumption sets might not be closed in the case of theorist and circumscription. For example, if a formula is an assumption in theorist that is already derived by the theory, then the empty assumption set is not closed.

A **stable extension** is an extension $Th(T \cup \Delta)$ for some stable assumption set Δ. The standard semantics of circumscription [20] corresponds to the intersection of all stable extensions of the abstract framework corresponding to circumscription.

Bondarenko *et al.* argue that the stability semantics is unnecessarily restrictive, because it insists that an assumption set should take a stand on every issue (assumption). Thus, they define new semantics for the abstract framework, by generalising the argumentation-theoretic reformulation of [17] for the semantics originally proposed for LP by Dung [8]. The new semantics are defined in terms of "admissible" and "preferred" sets of assumptions/extensions. An assumption set $\Delta \subseteq A$ is **admissible** iff

1. Δ is closed,

2. Δ does not attack itself, and
3. for all closed sets of assumptions $\Delta' \subseteq A$, if Δ' attacks Δ then Δ attacks Δ'.

Maximal (with respect to set inclusion) admissible assumption sets are called **preferred**. In this paper we use the following terminology: an **admissible (preferred) argument** is an extension $Th(T \cup \Delta)$ for some admissible (preferred) assumption set Δ. Bondarenko *et al.* show that preferred arguments correspond to preferred extensions [8] and partial stable models [27] for LP.

Every stable assumption set/extension is preferred (and thus admissible) [1, Theorem 4.6], but not vice versa, in general. However, if the framework is **normal**, i.e., if every maximal closed assumption set that does not attack itself is a stable set, then the semantics of preferred and stable assumption sets coincide [1, Theorem 4.8]. Theorist and circumscription are normal frameworks, which implies that stability and preferability semantics are identical in these cases.

In the sequel we will use the following:

(Prop$_1$): Every preferred assumption set is (trivially) admissible and every admissible assumption set is a subset of some preferred assumption set;

(Prop$_2$): The empty assumption set is always admissible, trivially, for all flat frameworks;

(Prop$_3$): Every preferred extension is stable and every stable extension is preferred, for all normal frameworks.

Moreover, for any given semantics amongst the stability, admissibility and preferability semantics, we will use the terminology that "**a set of assumptions is sanctioned by a semantics**" to mean that the set of assumptions is stable/admissible/preferred, respectively.

3 Reasoning Problems and Computational Complexity

We will analyse the *computational complexity* of the following reasoning problems for the propositional variants of the frameworks for theorist, circumscription, LP, DL, and AEL under admissibility and preferability semantics:

- the **credulous reasoning problem**, i.e., the problem of deciding for any given sentence $\varphi \in \mathcal{L}$ whether $\varphi \in Th(T \cup \Delta)$ for *some* assumption set Δ sanctioned by the semantics;

- the **sceptical reasoning problem**, i.e., the problem of deciding for any given sentence $\varphi \in \mathcal{L}$ whether $\varphi \in Th(T \cup \Delta)$ for *all* assumption sets Δ sanctioned by the semantics.

Instead of the sceptical reasoning problem, we will often consider its complementary problem, i.e.

- the **co-sceptical reasoning problem**, i.e, the problem of deciding for any given sentence φ whether $\varphi \notin Th(T \cup \Delta)$ for *some* assumption set Δ sanctioned by the semantics.

Note that we are not advocating co-sceptical reasoning as interesting or useful epistemologically. Rather, we use it to support our complexity analysis.

In addition, we will consider a sub-problem of all these problems, namely:

- the **assumption set verification problem**, i.e., the problem of deciding whether a given set of assumptions Δ is sanctioned by the semantics.

We briefly revise fundamental notions from complexity theory.[4] We assume familiarity with the complexity classes P, NP, and co-NP, and with the notions of *many-one-reductions*, *Turing reductions*, and *hardness* and *completeness* with respect to these reductions.

The complexity of the above problems for all frameworks and semantics we consider is located at the lower end of the *polynomial hierarchy*. This is a (presumably) infinite hierarchy of complexity classes above NP defined by using *oracle machines*, i.e. Turing machines that are allowed to call a subroutine—the *oracle*—deciding some fixed problem in constant time. Let $\mathcal{C}$ be a class of decision problems. Then, $\mathsf{P}^{\mathcal{C}}$ denotes the class of problems that can be solved on a deterministic oracle machine in polynomial time with an oracle that decides a problem in $\mathcal{C}$. In general, for any class $\mathcal{X}$ defined by resource bounds, $\mathcal{X}^{\mathcal{C}}$ denotes the class of problems decidable on a Turing machine with a resource bound given

[4]Good textbooks covering the notions we introduce here have been written by Garey and Johnson [13] and Papadimitriou [24].

by $\mathcal{X}$ and an oracle for a problem in $\mathcal{C}$.[5] Based on these notions, the sets Δ_k^p, Σ_k^p, and Π_k^p are defined as follows:[6]

$$\Sigma_0^p = \Pi_0^p = \Delta_0^p = \mathsf{P}$$
$$\Sigma_{k+1}^p = \mathsf{NP}^{\Sigma_k^p},\ \Pi_{k+1}^p = \text{co-}\mathsf{NP}^{\Sigma_k^p},\ \Delta_{k+1}^p = \mathsf{P}^{\Sigma_k^p}$$

The "canonical" complete problems are SAT for Σ_1^p=NP and k-QBF for Σ_k^p ($k > 1$), where k-QBF is the problem of deciding whether the quantified boolean formula

$$\underbrace{\exists\vec{p}\,\forall\vec{q}\ldots}_{k \text{ alternating quantifiers starting with } \exists}\Phi(\vec{p},\vec{q},\ldots).$$

is true, for a formula Φ. The above problems remain complete for their respective classes when the innermost quantifier is $\exists$ and the formula Φ is in 3CNF, as well as when the innermost quantifier is $\forall$ and the formula Φ is in 3DNF [30]. The complementary of a k-QBF problem, denoted by co-k-QBF, is complete for Π_k^p.

All problems in the polynomial hierarchy can be solved in polynomial time iff $\mathsf{P} = \mathsf{NP}$. Further, all these problems can be solved by worst-case exponential time algorithms. Thus, the polynomial hierarchy might not seem too meaningful. However, different levels of the polynomial hierarchy differ considerably in practice, e.g. methods working for moderately sized instances of NP-complete problems do not work for Σ_2^p-complete problems.

The complexity of the problems we are interested in has been extensively studied for existing logics for default reasoning under the standard, stability semantics [3, 15, 23, 19, 29, 2, 10]. Table 1 gives a partial summary of these results for the different logics. We note here that the semantics of circumscription has been originally proposed with respect to sceptical reasoning only. In this case, as shown in [1], reasoning in circumscription (restricted to Herbrand models) coincides with sceptical reasoning in theorist under the stability semantics. Moreover, we can naturally extend circumscription and define its credulous reasoning via a one-to-one correspondence with credulous reasoning in theorist under the stability semantics (see [1] for more details). Hence the complexity result for credulous reasoning in circumscription is a direct consequence of the respective result for

[5] Note that because using an oracle for a problem from $\mathcal{C}$ is identical to using an oracle for a problem from co-$\mathcal{C}$, we have $\mathcal{X}^{\mathcal{C}} = \mathcal{X}^{co-\mathcal{C}}$. For this reason, one usually does not use the notation $\mathcal{X}^{co-\mathcal{C}}$.

[6] The super-script p is only used to distinguish these classes from the analogous classes in the Kleene hierarchy.

theorist. The complexity results for reasoning in circumscription under the admissibility semantics, presented later in the paper, can be understood in a similar way.

	credulous reasoning	sceptical reasoning
Logic Programming	NP-complete	co-NP-complete
Theorist	Σ_2^p-complete	Π_2^p-complete
Circumscription	Σ_2^p-complete	Π_2^p-complete
Default Logic	Σ_2^p-complete	Π_2^p-complete
Autoepistemic Logic	Σ_2^p-complete	Π_2^p-complete

Table 1: Existing computational complexity results for the stability semantics

4 Generic Upper Bounds

In this section we give a number of *generic* upper bounds for reasoning under the admissibility and preferability semantics that are parametric on the complexity of the *derivability problem* in the underlying monotonic logic. This allows us to derive upper bounds for a wide range of concrete logics for default reasoning.

In the case of LP, the underlying logic is propositional Horn logic, hence the derivability problem is P-complete (under log-space reductions) [24, p. 176]. In the case of theorist, circumscription and AEL, the underlying logic is classical propositional logic, hence the derivability problem is co-NP-complete. Finally, in the case of DL, the underlying monotonic derivability is classical derivability extended with domain-specific inference rules. However, these extra inference rules do not increase the complexity of reasoning. Indeed, it is known (e.g. see [16, p.90]) that for any DL-like propositional monotonic rule system S, checking whether $S \not\models \varphi$ is NP-complete. Therefore, the following proposition follows immediately.

Proposition 1 *Given a DL framework $\langle T, A, ^{-} \rangle$, deciding for a sentence $\varphi \in \mathcal{L}$ and an assumption set $\Delta \subseteq A$ whether $\varphi \in Th(T \cup \Delta)$ is* co-NP-*complete.*

In order to decide the credulous and co-sceptical reasoning problems, one can apply the following non-deterministic algorithm:

Algorithm 2

1. *Guess an assumption set,*
2. *verify that it is sanctioned by the semantics, and*
3. *verify that the formula under consideration is derivable from the set of assumptions and the monotonic theory or not derivable from it, respectively.*

From this it follows that credulous reasoning and co-sceptical reasoning is in the complexity class $\mathsf{NP}^{\mathcal{C}}$, provided reasoning in the underlying logic is in $\mathcal{C}$ and the verification that an assumption set is sanctioned by the semantics can be done with polynomially many calls to a $\mathcal{C}$-oracle. For the stability semantics, we need indeed only polynomially many $\mathcal{C}$-oracle calls in order to verify that the assumption set Δ is not self-attacking and that it is closed and attacks all assumptions $\alpha \notin \Delta$. However, for the admissibility and preferability semantics the verification step does not seem to be so easy, as suggested by the following theorem.

Theorem 3 *For frameworks with an underlying monotonic logic with a derivability problem in $\mathcal{C}$, the* assumption set verification problem *is*

- *in* $\mathsf{P}^{\mathcal{C}}$ *under the stability semantics,*
- *in* co-$\mathsf{NP}^{\mathcal{C}}$ *under the admissibility semantics, and*
- *in* co-$\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}$ *under the preferability semantics.*

Proof: The first claim follows from the argument above that polynomially many $\mathcal{C}$-oracle calls are sufficient to verify that an assumption set is stable.

In order to prove the second claim, we give the following nondeterministic, polynomial-time algorithm that uses a $\mathcal{C}$-oracle and decides whether $\Delta \subseteq A$ is not admissible:

1. Check whether Δ is closed. If not, succeed, otherwise continue.
2. Guess an assumption set $\Delta' \subseteq A$.

3. Verify that Δ' is closed, using $|A - \Delta'|$ $\mathcal{C}$-oracle calls.

4. Verify that Δ' attacks Δ, using $|\Delta|$ $\mathcal{C}$-oracle calls.

5. Verify that Δ does not attack Δ', using $|\Delta'|$ $\mathcal{C}$-oracle calls.

Obviously, this algorithm succeeds iff Δ is not admissible, i.e., it decides the complement of the assumption set verification problem, thus proving the claim.

In order to prove the third claim, for any assumption set $\Delta \subseteq A$ that we want to verify, we give the following nondeterministic, polynomial-time algorithm that uses an $\mathsf{NP}^{\mathcal{C}}$-oracle:

1. Check whether Δ is admissible, using one $\mathsf{NP}^{\mathcal{C}}$-oracle call (by the second claim). If it is not, succeed. Otherwise continue.

2. Guess an assumption set $\Delta' \supset \Delta$.

3. Check whether Δ' is admissible, using one $\mathsf{NP}^{\mathcal{C}}$-oracle call (by the second claim). If it is, succeed. Otherwise fail.

Obviously, this algorithm succeeds iff Δ is not preferred. This means it decides the complement of the assumption set verification problem, thus proving the claim. ■

Furthermore, in the general case, there does not appear to be more efficient algorithms for the assumption set verification problems than the ones given in Theorem 3. For the special flat and normal frameworks, however, more efficient algorithms can be found, as demonstrated by the following two theorems.

Theorem 4 *For* flat *frameworks with an underlying monotonic logic with a derivability problem in $\mathcal{C}$, the* assumption set verification problem *is*

- *in* $\mathsf{P}^{\mathcal{C}}$ *under the admissibility semantics, and*
- *in* $\text{co-}\mathsf{NP}^{\mathcal{C}}$ *under the preferability semantics.*

Proof: We prove the first claim by giving the following deterministic, polynomial-time algorithm using a $\mathcal{C}$-oracle, for any assumption set $\Delta \subseteq A$ that we want to verify:

1. Check whether Δ attacks itself, using polynomially many $\mathcal{C}$-oracle calls. If it does, succeed. Otherwise continue.

2. Compute $A^* = \{\alpha \in A - \Delta \mid \Delta$ does not attack $\alpha\}$, using $|A - \Delta|$ calls to a $\mathcal{C}$-oracle.

3. Check whether $A^* \cup \Delta$ attacks Δ, using polynomially many $\mathcal{C}$-oracle calls. If it does, succeed. Otherwise fail.

It is easy to see that if this algorithm succeeds then Δ is not admissible, as $A^* \cup \Delta$ attacks Δ but, by (2), Δ does not attack A^* and, by (1), Δ does not attack itself.[7] Moreover, if the algorithm fails then Δ is admissible. Indeed, let Δ' be any attack against Δ. If $\Delta' \subseteq A^* \cup \Delta$, then, by monotonicity of the underlying logic, $A^* \cup \Delta$ attacks Δ, thus contradicting that the algorithm fails. Therefore, $\Delta' \not\subseteq A^* \cup \Delta$. Let $\alpha \in \Delta' - A^* - \Delta$. By (2), Δ attacks α. Thus, Δ attacks Δ', and, by (1), Δ is admissible.

The second claim of the theorem follows by reconsidering the algorithm used in the proof of Theorem 3 for the third claim, but using $\mathsf{P}^{\mathcal{C}}$-oracle calls at steps (1) and (3). ■

Due to (Prop$_3$), for normal frameworks the assumption set verification task under the preferability semantics is easier, as it can be reduced to that under the stability semantics. Therefore, the following result is a direct corollary of Theorem 3.

Proposition 5 *For* normal *frameworks with an underlying monotonic logic with a derivability problem in $\mathcal{C}$, the* assumption set verification problem *under the preferability semantics is in* $\mathsf{P}^{\mathcal{C}}$.

We could now apply directly algorithm 2 described above in combination with the above results for deriving upper bounds for the credulous and sceptical reasoning problems. However, some of the upper bounds thus obtained can be reduced, as follows.

Directly from (Prop$_1$), we have the following result.

[7] Note that if the framework is not flat, then the assumption set $A^* \cup \Delta$ might not be closed. Therefore, even if (3) succeeds, Δ can still be admissible, as it may attack an assumption that is *derivable* from $A^* \cup \Delta$.

Proposition 6 *Credulous reasoning under the admissibility semantics is equivalent to credulous reasoning under the preferability semantics.*

Thus, it follows directly that credulous reasoning under the admissibility semantics has the same upper bound as credulous reasoning under the preferability semantics. In particular, for normal frameworks we get the same upper bound for credulous reasoning under the admissibility semantics as for the stability semantics.

In addition, co-sceptical and sceptical reasoning under the admissibility semantics is often much simpler than suggested by the upper bounds of the respective assumption set verification problem combined with algorithm 2. For example, in flat frameworks $\langle T, A, ^- \rangle$ the sceptical reasoning problem reduces to the classical derivability from the theory T, because of (Prop$_2$). This might be the case even for non-flat frameworks. We call an assumption-based framework $\langle T, A, ^- \rangle$ **simple** iff there is no admissible assumption set whenever T is inconsistent in the underlying monotonic logic [8], and otherwise there exists a minimal (with respect to set inclusion) admissible set $\Delta_m = A \cap Th(T)$.

Proposition 7 *For* flat *frameworks and for* simple *frameworks with an underlying monotonic logic with a derivability problem in $\mathcal{C}$, the* sceptical reasoning problem *under the admissibility semantics is in $\mathcal{C}$.*

All the results in this section combined with algorithm 2 give the next theorem, specifying upper bounds for the reasoning problems for all the types of frameworks considered so far.

Theorem 8 *Upper bounds for the different reasoning problems, types of frameworks, and semantics are as specified in the following table:*

Frameworks	*Stability*		*Admissibility*		*Preferability*	
	cred.	*scept.*	*cred.*	*scept.*	*cred.*	*scept.*
general	$\mathsf{NP}^{\mathcal{C}}$	co-$\mathsf{NP}^{\mathcal{C}}$	$\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}$	co-$\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}$	$\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}$	co-$\mathsf{NP}^{\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}}$
normal	$\mathsf{NP}^{\mathcal{C}}$	co-$\mathsf{NP}^{\mathcal{C}}$	$\mathsf{NP}^{\mathcal{C}}$	co-$\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}$	$\mathsf{NP}^{\mathcal{C}}$	co-$\mathsf{NP}^{\mathcal{C}}$
flat	$\mathsf{NP}^{\mathcal{C}}$	co-$\mathsf{NP}^{\mathcal{C}}$	$\mathsf{NP}^{\mathcal{C}}$	$\mathcal{C}$	$\mathsf{NP}^{\mathcal{C}}$	co-$\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}$
simple	$\mathsf{NP}^{\mathcal{C}}$	co-$\mathsf{NP}^{\mathcal{C}}$	$\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}$	$\mathcal{C}$	$\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}$	co-$\mathsf{NP}^{\mathsf{NP}^{\mathsf{NP}^{\mathcal{C}}}}$

[8]Note that not all deductive systems underlying an abstract framework are equipped with a notion of inconsistency. For example, the instance of the framework for LP is not. Moreover, note that the notion of inconsistency is a separate notion from that of contrary.

Proof: The results for the stability semantics follow from applying algorithm 2, with step (2) of the algorithm solvable by a call to a $\mathsf{P}^{\mathcal{C}}$ oracle (by Theorem 3), and step (3) solvable by a call to a $\mathcal{C}$-oracle. This gives an upper bound of $\mathsf{NP}^{\mathsf{P}^{\mathcal{C}}}$, which coincides with $\mathsf{NP}^{\mathcal{C}}$, for both the credulous and co-sceptical reasoning problems.

The results for the admissibility semantics in the first row and for sceptical reasoning under the preferability semantics in the first row follow by the same argument.

The result for credulous reasoning under the preferability semantics in the first row follows from Proposition 6 and the corresponding result for the admissibility semantics.

The results for admissibility and preferability semantics in the second row are justified as follows. Credulous reasoning under admissibility and preferability semantics as well as co-sceptical reasoning under the preferability semantics can be shown to be in $\mathsf{NP}^{\mathsf{P}^{\mathcal{C}}}$, which equals $\mathsf{NP}^{\mathcal{C}}$, by using algorithm 2 and applying Propositions 6 and 5. Further, the upper bound for sceptical reasoning under the admissibility semantics is the general upper bound given in the first row.

The results for admissibility and preferability semantics in the third row follow by applying Proposition 7 for sceptical reasoning under admissibility and, for the other columns, algorithm 2 and Theorem 4.

Finally, the results for admissibility and preferability semantics in the fourth row are the general results in the first row, with the exception of the result for sceptical reasoning under the admissibility semantics given by Proposition 7. ■

As shown in the table, the upper bounds derived for sceptical reasoning under the admissibility semantics are sometimes lower than those derived for sceptical reasoning under the stability semantics. However, in these cases it amounts to deriving monotonic conclusions from the theory T and ignoring the assumptions completely. In other words, in these cases, default reasoning is trivialised.

5 Simple, Normal Frameworks: Theorist and Circumscription

The concrete frameworks for theorist and circumscription are *normal* [1] and *simple*, as shown below.

Lemma 9 *The frameworks for theorist and circumscription are simple.*

Proof: Circumscription is a special instance of theorist. Thus, we only need to prove the theorem for theorist.

If the given theorist theory T is inconsistent then the corresponding framework admits no admissible argument, as any closed assumption set attacks itself.

Assume that T is consistent. Then, we only need to prove that $\Delta = Th(T) \cap A$ attacks every closed assumption set Δ' which attacks Δ. Now, if $\Delta = \emptyset$, then there is no set Δ' that attacks Δ. If $\Delta \neq \emptyset$, then Δ' attacks Δ iff $T \cup \Delta'$ is inconsistent and, as Δ' is closed, $\Delta' = A$. Thus, necessarily Δ attacks Δ'. ■

For both frameworks, the credulous and sceptical reasoning problems reach the respective upper bounds specified in Theorem 8. Indeed, due to Proposition 6 and (Prop_3), credulous reasoning under admissibility and preferability semantics is identical to credulous reasoning under the standard, stability semantics, leading to the result that the complexity is also identical. Thus, the next proposition follows immediately from the results in Table 1.

Proposition 10 *Credulous reasoning in theorist and circumscription under the admissibility and preferability semantics is Σ_2^p-complete.*

Directly by (Prop_3), sceptical reasoning under the preferability semantics is identical to sceptical reasoning under the stability semantics. Thus, the next proposition follows immediately from the results in Table 1.

Proposition 11 *Sceptical reasoning in theorist and circumscription under the preferability semantics is Π_2^p-complete.*

Finally, sceptical reasoning under the admissibility semantics is trivial because the frameworks are simple and sceptical reasoning reduces to monotonic derivability from the theory.

Proposition 12 *Sceptical reasoning in theorist and circumscription frameworks under the admissibility semantics is* co-NP*-complete.*

In other words, for the concrete frameworks for theorist and circumscription, we either get the same results as under the stability semantics or we get trivial results.

6 Flat Frameworks: Logic Programming and Default Reasoning

As in the case of theorist and circumscription, in the case of LP and DL the upper bounds specified in Theorem 8 are tight.

Since the concrete framework for LP is flat, sceptical reasoning under the admissibility semantics reduces to reasoning in the underlying monotonic logic, i.e., derivability in propositional Horn theories, which is P-complete.

Proposition 13 *Sceptical reasoning in LP under the admissibility semantics is* P*-complete.*

From Theorem 8, again because the LP framework is flat, credulous reasoning under the admissibility and preferability semantics is in NP^{P}, which equals NP. NP-completeness can be obtained as a direct corollary of an earlier result by Saccà [28], that the expressive power of DATALOG$^{\neg}$ queries under the "possible M-stable semantics" (corresponding to credulous reasoning under the admissibility and preferability semantics) coincides with DB-NP, i.e. the class of all databases that are recognisable in NP. From this result the following theorem follows immediately.

Theorem 14 *Credulous reasoning in LP under the admissibility and preferability semantics is* NP*-complete.*

Again from Theorem 8, sceptical reasoning in LP under the preferability semantics is in co-$\mathsf{NP}^{\mathsf{NP}^{\mathsf{P}}}$, which coincides with Π_2^p. Π_2^p-completeness can be obtained again as a direct corollary of the result proven again by Saccà [28], that the expressive power of DATALOG$^{\neg}$ queries under the "definite M-stable semantics" (corresponding to sceptical preferability semantics) coincides with the class DB-Π_2^p, i.e. the class of all databases that are recognisable in DB-Π_2^p. From this result the following theorem follows immediately.

Theorem 15 *Sceptical reasoning in LP under the preferability semantics is* Π_2^p*-complete.*

Therefore, for LP, credulous reasoning under admissibility and preferability semantics has the same complexity as under the stability semantics (see Table 1), whereas sceptical reasoning is either one level lower but trivial, under the admissibility semantics, or one level higher, under the preferability semantics, than under the stability semantics.

Since the instance of the framework for DL is flat, sceptical reasoning under the admissibility semantics reduces to reasoning in the underlying monotonic logic, i.e., derivability in propositional classical logic, which is co-NP-complete.

Proposition 16 *Sceptical reasoning in DL under the admissibility semantics is* co-NP*-complete.*

By Proposition 6, credulous reasoning under the preferability semantics coincides with credulous reasoning under the admissibility semantics. From Theorem 8, credulous reasoning under the admissibility and preferability semantics is in $\mathsf{NP}^{\mathsf{NP}}$, which coincides with Σ_2^p. Σ_2^p-hardness, and therefore Σ_2^p-completeness, can be proven by a reduction from 2-QBF.

Theorem 17 *Credulous reasoning in DL under the admissibility and preferability semantics is Σ_2^p-complete.*

Proof: By Proposition 6, it suffices to prove the theorem for the admissibility semantics. Membership follows from Theorem 8. To prove hardness, we use a straightforward reduction from 2-QBF to the credulous reasoning problem under the admissibility semantics.

Assume the quantified boolean formula $\exists p_1, \ldots, p_n \, \forall q_1, \ldots, q_m \Phi$, with Φ a formula in 3DNF over the propositional variables $p_1, \ldots, p_n, q_1, \ldots, q_m$. We construct a DL theory $(\emptyset, D)$ such that the given quantified boolean formula is true iff some admissible argument for the framework corresponding to $(\emptyset, D)$ contains Φ.

Let D consists of the default rules

$$\frac{Mp_i}{p_i}; \quad \frac{M\neg p_i}{\neg p_i}$$

for each $i = 1, \ldots, n$, simulating the choice of a truth value for each propositional variable p_i in Φ. Obviously, $(\emptyset, D)$ can be constructed in log-space. Moreover, it is easy to see that the given 2-QBF is true iff there exists an admissible extension of the framework corresponding to $(\emptyset, D)$ containing Φ. ■

Again from Theorem 8, in DL, sceptical reasoning under the preferability semantics is in co-$\mathsf{NP}^{\mathsf{NP}^{\mathsf{NP}}}$, which coincides with Π_3^p. Π_3^p-hardness, and therefore Π_3^p-completeness, can be proven by a reduction from 3-QBF.

Theorem 18 *Sceptical reasoning in DL under the preferability semantics is* Π_3^p-*complete.*

Proof: Membership follows from Theorem 8. To prove hardness, we use a reduction from 3-QBF to the co-sceptical reasoning problem under the preferability semantics.

Assume the quantified boolean formula $\Psi = \exists p_1, \ldots, p_n \, \forall q_1, \ldots, q_m \, \exists r_1, \ldots, r_k \Phi$, with Φ a formula in 3CNF over the propositional variables $p_1, \ldots, p_n$, $q_1, \ldots, q_m$, $r_1, \ldots, r_k$. We construct a DL theory $(\emptyset, D)$ such that Ψ is true iff some sentence F is not contained in some preferred argument for the framework corresponding to $(\emptyset, D)$.

The language of $(\emptyset, D)$ contains atoms $p_1,\ldots, p_n$, $q_1,\ldots, q_m$, and $r_1,\ldots, r_k$ as well as atoms $t_1,\ldots, t_n$, $s_1,\ldots, s_m$, intuitively holding true iff a truth value for the variables $p_1,\ldots, p_n$, $q_1,\ldots, q_m$, respectively, has been chosen. D consists of the default rules

$$\frac{M(p_i \wedge t_i)}{p_i \wedge t_i}; \; \frac{M(\neg p_i \wedge t_i)}{\neg p_i \wedge t_i}; \; \frac{M(q_j \wedge s_j)}{q_j \wedge s_j}; \; \frac{M(\neg q_j \wedge s_j)}{\neg q_j \wedge s_j}$$

for each $i = 1,\ldots, n$, $j = 1,\ldots, m$, simulating the choice of a truth value for each p_i and q_j in Φ,

$$\frac{M\Phi}{\bigwedge_{j=1,\ldots,m} \neg s_j}$$

to prohibit choices of truth values for all the q_j's that render Φ satisfiable,

$$\frac{M\neg s_j}{\bigwedge_{h=1,\ldots,m} \neg s_h}; \; \frac{M\neg t_i}{\bigwedge_{h=1,\ldots,n} \neg t_h \wedge \bigwedge_{h=1,\ldots,m} \neg s_h}$$

for each $i = 1, \ldots, n$, $j = 1, \ldots, m$, to enforce that truth value choices are made either for all q_j's or for no q_j and that truth value choices are made either for all p_i's or for none of the p_i's and q_j's, and finally

$$\frac{M\Phi}{\neg\Phi}; \; \frac{M\neg t_i}{t_i}; \; \frac{M\neg s_j}{s_j}.$$

for each $i = 1, \ldots, n$, $j = 1, \ldots, m$ to guarantee that no admissible assumption set contains $M\Phi$ or any of $M\neg t_i$ and $M\neg s_j$.

Obviously, $(\emptyset, D)$ can be constructed in log-space. Moreover, we prove that Ψ is true iff there is a preferred argument not containing $F = \wedge_{j=1,..,m} s_j$. In other words, the 3-QBF Ψ can be reduced to co-sceptical reasoning in DL under the preferability semantics.

In the sequel we will use the following terminology. If v is a truth assignment to the p_i's, we denote by Δ_v^p the assumption set

$$\{M(p_i \wedge t_i) | v(p_i) = true, i = 1, \ldots, n\} \cup \{M(\neg p_i \wedge t_i) | v(p_i) = false, i = 1, \ldots, n\}.$$

Similarly, if u is a truth assignment to the q_j's, we denote by Δ_u^q the assumption set

$$\{M(q_j \wedge s_j) | v(q_j) = true, j = 1, \ldots, m\} \cup \{M(\neg q_j \wedge s_j) | v(q_j) = false, j = 1, \ldots, m\}.$$

First of all, it is obvious that no admissible assumption set can contain any of the assumptions $M\Phi$, $M\neg s_i$, $M\neg t_i$ (as, if it did, it would attack itself). Furthermore, it is easy to see that for any truth assignment v to the p_i's, the set Δ_v^p is an admissible set. Moreover, every preferred assumption set must contain a set Δ_v^p for some truth assignment v to the p_i's. Finally, if Δ_v^p is not preferred, then there exists a truth assignment u to the q_i's such that $\Delta_v^p \cup \Delta_u^q$ is preferred.

Assume that Ψ is true under a particular truth assignment v to the p_i's. Obviously, Δ_v^p does not derive $F = \wedge_{j=1,..,m} s_j$. We show that the set Δ_v^p is a preferred assumption set.

Suppose that it is not, and that we can extend Δ_v^p by the set Δ_u^q, for some truth assignment u to the q_i's, thus obtaining an admissible set. Then, $\Delta_v^p \cup \Delta_u^q$ counter attacks the attack $\{M\Phi\}$, i.e. $\neg\Phi$ belongs to the extension given by $\Delta_v^p \cup \Delta_u^q$. As a consequence, Ψ is not true under the truth assignment v: contradiction.

Conversely, assume that the framework corresponding to $(\emptyset, D)$ admits a preferred argument Δ that does not derive $F = \wedge_{j=1,..,m} s_j$. We prove that Ψ is true.

Clearly there exists some truth assignment v to the p_i's such that $\Delta_v^p \subseteq \Delta$. Since Δ is preferred and it does not contain F, none of the sets $\Delta_v^p \cup \Delta_u^q$, for every possible truth assignment u to the q_i's, is admissible. This means that none of these sets of assumptions can counter attack the attack $\{M\Phi\}$ and derive $\neg\Phi$. Therefore, Ψ is true. ■

Therefore, as in the LP case, in the DL case credulous reasoning under the admissibility and preferability semantics has the same complexity as under the stability semantics (see Table 1), whereas sceptical reasoning is either one level lower but trivial, under the admissibility semantics, or one level higher, under the preferability semantics, than under the stability semantics.

Note that similar results to the one obtained above for DL have been recently obtained for disjunctive logic programming [11].

7 General Frameworks: Autoepistemic Logic

AEL is neither flat, simple, nor normal. This means that we cannot expect any simplifications when reasoning in AEL frameworks. As a matter of fact, the upper bounds for *general frameworks*, which apply of course, are also tight for AEL.

By Proposition 6, credulous reasoning under the preferability semantics coincides with credulous reasoning under the admissibility semantics. From Theorem 8, since the reasoning problem in the underlying monotonic logic for AEL is classical reasoning in propositional logic (coNP-complete), credulous reasoning under the admissibility and preferability semantics is in $\mathsf{NP}^{\mathsf{NP}^{\mathsf{NP}}}$, which coincides with Σ_3^p. Σ_3^p-hardness, and therefore Σ_3^p-completeness, can be proven by a reduction from 3-QBF.

Theorem 19 *Credulous reasoning in AEL under the admissibility and preferability semantics is Σ_3^p-complete.*

Proof: By Proposition 6, it suffices to prove the theorem for the admissibility semantics. Membership follows from Theorem 8. To prove hardness, we use a reduction from 3-QBF to the credulous reasoning problem under the admissibility semantics.

Assume the following quantified boolean formula $\Psi = \exists p_1, \ldots, p_n \, \forall q_1, \ldots, q_m \, \exists r_1, \ldots, r_k \Phi$, with Φ a formula in 3CNF over the propositional variables $p_1, \ldots, p_n$, $q_1, \ldots, q_m$, $r_1, \ldots, r_k$. We construct an AEL theory T such that Ψ is true iff some sentence F is contained in some admissible argument for the framework corresponding to T.

The language of T contains atoms $p_1, \ldots, p_n$, $q_1, \ldots, q_m$, and $r_1, \ldots, r_k$ as well as atoms $t_1, \ldots, t_n$, intuitively holding if a truth value for the variables $p_1, \ldots, p_n$ has been chosen, and an atom s used to prevent that any truth value for the q_j's can be chosen. T consists of the sentences:

$$\neg L \neg p_i \;\rightarrow\; p_i \wedge t_i,$$

$$\begin{aligned} \neg L p_i &\rightarrow \neg p_i \wedge t_i, \\ \neg L \neg \Phi & \\ \neg L \neg q_j &\rightarrow q_j \wedge s \wedge \neg L s, \\ \neg L q_j &\rightarrow \neg q_j \wedge s \wedge \neg L s, \end{aligned}$$

for each $i = 1,\ldots,n$, $j = 1,\ldots,m$.

Obviously, T can be constructed in log-space. Now we prove that the framework corresponding to T admits an admissible extension containing $F = \bigwedge_{i=1,..,n} t_i$ iff Ψ is true. This means that the given 3-QBF can be reduced to credulous reasoning under the admissibility semantics.

Assume that the framework corresponding to T admits an admissible extension Δ deriving F. Then, for each $i = 1,\ldots,n$, either $\neg L\neg p_i$ or $\neg Lp_i$ is part of Δ, for F to be derived by it. Further, $\neg L\neg\Phi$ must be part of Δ, for Δ to be closed, and thus admissible. Finally, none of the assumptions $\neg L\neg q_i$, $\neg Lq_i$ can be part of Δ, for otherwise Δ, if closed, would attack itself and thus be non-admissible.

Consider any assumption set A that attacks Δ. A must attack one of the assumptions $\neg L\neg p_i$, $\neg Lp_i$, or $\neg L\neg\Phi$. However, if A attacked any of $\neg L\neg p_i$, $\neg Lp_i$, then Δ would immediately counter-attack A. Therefore, for A to be an assumption set that can possibly render Δ non-admissible, it must make the same choices on the p_i's as Δ, and attack $\neg L\neg\Phi$. For A to attack $\neg L\neg\Phi$, then A must derive $\neg\Phi$, by including, in addition to the assumptions from $\{\neg L\neg p_i, \neg Lp_i\}_{i=1,\ldots,n}$ already chosen by Δ, assumptions from the set $\{\neg L\neg q_j, \neg Lq_j\}_{j=1,\ldots,m}$. Such choices cannot be counter-attacked by Δ without making it self-attacking. Therefore, since Δ is admissible, no such A exists. This means that, for the given choices on the p_i's in Δ, no choices for the q_j's exist that make $\neg\Phi$ true. In other words, for the given choice of the p_i in Δ, and for all choices of the truth values for the q_j's, there exists an assignment of truth values to the r_l's that makes Φ true, which implies that Ψ is necessarily true.

Conversely, assume that there is no admissible extension of the framework corresponding to T deriving F above. Then, regardless of the choices for the p_i's, there is always an attack on $\neg L\neg\Phi$, deriving $\neg\Phi$, that cannot be counter-attacked while keeping the candidate set of assumptions non-self-attacking. Then, by the arguments presented above, the given 3-QBF formula Ψ cannot be true. ■

Again from Theorem 8, sceptical reasoning under the admissibility and preferability semantics in AEL is in co-$\mathsf{NP}^{\mathsf{NP}^{\mathsf{NP}}}$, which coincides with Π_3^p. Π_3^p-hardness, and therefore Π_3^p-completeness, can be proven by a reduction from 3-QBF.

Theorem 20 *Sceptical reasoning in AEL under the admissibility semantics is* Π_3^p-*complete.*

Proof: Membership follows from Theorem 8. To prove hardness, we use a reduction from 3-QBF to the co-sceptical reasoning problem under the admissibility semantics.

We use the reduction in the proof of the previous Theorem 19, but extend the theory T constructed there to the theory $T' = T \cup \{L \bigwedge_i t_i\}$.

Any admissible set Δ must contain the assumptions $\neg L \neg \Phi$ and $L \bigwedge_i t_i$ in order for Δ to be closed. Furthermore, any admissible extension of T' must contain $\bigwedge_i t_i$ because otherwise it is attacked by $\neg L \bigwedge_i t_i$ without having a counter-attack. From this fact and the above observations it follows that T' has an admissible extension iff the given 3-QBF formula Ψ is true. Given that if no admissible extension exists all co-sceptical queries will be answered negatively, the above is equivalent to the fact that $\neg \bigwedge_i t_i$ is not a sceptical consequence of T' iff Ψ is true, i.e., the construction is a log-space reduction from 3-QBF to co-sceptical reasoning under the admissibility semantics. ■

Again from Theorem 8, in AEL, sceptical reasoning under the preferability semantics is in co-$\mathsf{NP}^{\mathsf{NP}^{\mathsf{NP}^{\mathsf{NP}}}}$, which coincides with Π_4^p. Π_4^p-hardness, and therefore Π_4^p-completeness, can be proven by a reduction from 4-QBF.

Theorem 21 *Sceptical reasoning in AEL under the preferability semantics is* Π_4^p-*complete.*

Proof: Membership follows from Theorem 8. To prove hardness, we use a reduction from 4-QBF to the co-sceptical reasoning problem under the preferability semantics.

Assume the following quantified boolean formula $\Psi = \exists p_1,\ldots,p_n\, \forall q_1,\ldots,q_m\, \exists r_1,\ldots,r_k\, \forall s_1,\ldots,s_o\, \Phi$, with Φ a formula in 3DNF over the propositional variables $p_1,\ldots,p_n$, $q_1,\ldots,q_m$, $r_1,\ldots,r_k$, and $s_1,\ldots,s_o$. We construct an AEL theory T such that Ψ is true iff a particular sentence F is not contained in some preferred argument of T.

The language of T contains atoms $p_1,\ldots,p_n$, $q_1,\ldots,q_m$, and $r_1,\ldots,r_k$, $s_1,\ldots,s_o$ as well as atoms $t_1,\ldots,t_m$, the latter intuitively holding iff a truth value for the variables $q_1,\ldots,q_m$ has been chosen. Finally, we have atoms v and w. The atom v is used to block the truth assignment to the q_j's and w is used to prohibit any choices on assumptions $\{\neg L \neg r_h, \neg L r_h\}$ in the preferred argument.

T consists of the following sentences:

$$
\begin{aligned}
\neg L \neg p_i &\rightarrow p_i, & (1)\\
\neg L p_i &\rightarrow \neg p_i, & (2)\\
\neg L \neg q_j \wedge \neg L v &\rightarrow q_j \wedge t_j, & (3)\\
\neg L q_j \wedge \neg L v &\rightarrow \neg q_j \wedge t_j, & (4)\\
\neg L t_j &\rightarrow v, & (5)\\
\neg L \neg r_h &\rightarrow r_h \wedge w \wedge \neg L w, & (6)\\
\neg L r_h &\rightarrow \neg r_h \wedge w \wedge \neg L w, & (7)\\
\Phi &\rightarrow v, & (8)
\end{aligned}
$$

for each $i = 1,\ldots,n$, $j = 1,\ldots,m$, $h = 1,\ldots,k$.

Now we claim that there exists a preferred extension not containing $F = \bigwedge t_i$ iff Ψ is true.

First, one notes that an assumption set containing non-conflicting assumptions from the set $\{\neg L\neg p_i, \neg L p_i\}$ is an admissible set. Let Δ be a maximal such set.

Secondly, it is obvious that Δ can be expanded (in a non-trivial way) only by adding the assumption $\neg Lv$ and assumptions from the set $\{\neg L\neg q_j, \neg L q_j\}$. Let us call this expanded set Δ'. Such a set Δ' is only admissible if we make choices for all q_j's because otherwise Δ' can be attacked by $\neg L t_j$ (using $\neg L t_j \rightarrow v$) for which there is no counter-attack from Δ'.

Thirdly, the set Δ' cannot be further expanded using assumptions from $\{\neg L\neg r_h, \neg L r_h\}$, because these assumptions lead to an immediate self-attack.

Fourthly, an assumption set Δ' containing assumptions from $\{\neg L\neg q_j, \neg L q_j\}$ together with $\neg Lv$ can only be admissible if $\neg Lv$ cannot be attacked by any assumption set.

The only way to construct an attack A against $\neg Lv$ in Δ', which is not immediately counter-attacked by Δ', would be to use all assumptions in Δ' and assumptions from $\{\neg L\neg r_h, \neg L r_h\}$. Note that such assumptions cannot be counter-attacked by Δ'. Now the only way to attack $\neg Lv$ would be to make v true, and in order to do so, one has to make Φ true.

Assuming now that Δ' is admissible means that for all possible choices for the r_h's, Φ is not derivable, i.e., there is always a truth assignment to the s_k's that makes $\neg\Phi$ true. This means that Δ cannot be expanded by assumptions from $\{\neg L\neg q_j, \neg L q_j\}$ together with $\neg Lv$, if under the truth assignment to the p_i's corresponding to the assumptions in Δ, for all truth assignments to the q_j's, there is always an truth assignment to the r_h's that makes Φ true. In other words, if there exists a preferred assumption set that does not render $\bigwedge t_j$ true under T, Ψ is true.

Conversely, let us assume that Ψ is true. Let Δ be an assumption set containing assumptions from $\{\neg L\neg p_i, \neg Lp_i\}$ corresponding to a truth assignment to the p_i's that makes $\forall q_1,\ldots, q_m \exists r_1,\ldots, r_k \forall s_1,\ldots, s_o \Phi$ true. This assumption set cannot be expanded to Δ' by assumptions from $\{\neg L\neg q_j, \neg Lq_j\}$ together with $\neg Lv$, because for any such expansion there exists a value assignment to the r_h's which makes Φ true, corresponding to a set of choices from $\{\neg L\neg r_h, \neg Lr_h\}$ which together with Δ' is an assumptions set that leads together with T to the derivation of Φ and v, hence attacking Δ'. For this reason, there exists a preferred extension not containing both $\neg Lv$ and choices from $\{\neg L\neg q_j, \neg Lq_j\}$, and hence this preferred extension does not contain $\bigwedge t_j$. ■

Therefore, all reasoning problems are harder in AEL under the admissibility and preferability semantics than under the stability semantics (see Table 2). Indeed, credulous reasoning under admissibility and preferability semantics is one level higher than under the stability semantics (Theorem 19); sceptical reasoning under the admissibility semantics is one level higher than under the stability semantics (Theorem 20); sceptical reasoning under the preferability semantics is two levels higher than under the stability semantics (Theorem 21).

Moreover, whereas reasoning under the stability semantics has the same complexity in AEL as in DL, reasoning under the admissibility and preferability semantics is harder in AEL than in DL. Indeed, sceptical reasoning under the preferability semantics and credulous reasoning under the admissibility and preferability semantics are one level harder for AEL than for DL, and sceptical reasoning under the admissibility semantics is two levels harder.

We note that various complexity results for the *parsimonious* and *moderately grounded* semantics for AEL are presented in [9]. It would be interesting to see how the semantics for AEL provided by preferred/admissible arguments (for the instance of the abstract framework for AEL) relate to the semantics of [9]. This is however outside the scope of the present paper.

8 Conclusion and Discussion

We have studied the computational complexity of the credulous and sceptical reasoning problems under the new admissibility and preferability semantics for the abstract framework for default reasoning proposed in [1], for a number of concrete instances of the abstract framework, namely theorist, circumscription, logic programming (LP), default logic (DL) and autoepistemic logic (AEL). These new semantics are presented in [1] as "simpler" alternatives to the conventional stability semantics for all instances of the framework (see Section 1 for a discussion of this issue).

Table 2 summarises the results we have proven (for the admissibility and preferability semantics) as well as existing results in the literature (for the stability semantics). In the table, "$\mathcal{X}$-c." stands for "$\mathcal{X}$-complete." We have proven the results by appealing to properties of the frameworks, whenever possible. In particular, we have used the properties (proven in [1]) that default logic and logic programming are flat frameworks and that theorist and circumscription are normal frameworks. In addition, we have introduced the new property that frameworks are *simple*, and proven that theorist and circumscription satisfy such property. Autoepistemic logic is a general framework in that it does not satisfy any special property amongst the ones considered.

Framework	Property	Admissibility		Preferability		Stability	
		cred.	scept.	cred.	scept.	cred.	scept.
AEL	general	Σ_3^p-c.	Π_3^p-c.	Σ_3^p-c.	Π_4^p-c.	Σ_2^p-c.	Π_2^p-c.
DL	flat	Σ_2^p-c.	co-NP-c.	Σ_2^p-c.	Π_3^p-c.	Σ_2^p-c.	Π_2^p-c.
LP		NP-c.	P–c.	NP-c.	Π_2^p-c.	NP-c.	co-NP–c.
Theorist	simple &	Σ_2^p-c.	co-NP-c.	Σ_2^p-c.	Π_2^p-c.	Σ_2^p-c.	Π_2^p-c.
Circumscription	normal	Σ_2^p-c.	co-NP-c.	Σ_2^p-c.	Π_2^p-c.	Σ_2^p-c.	Π_2^p-c.

Table 2: Overview of complexity results

The table shows that reasoning under the new semantics can be much harder than reasoning under the conventional stability semantics. In particular, for AEL, sceptical reasoning under the admissibility and preferability semantics is one and two level harder, respectively, than under the stability semantics, whereas credulous reasoning under both new semantics is one level harder than under the stability semantics. Also, for DL and LP, sceptical reasoning under the preferability semantics is one level harder than under the stability semantics, whereas sceptical reasoning under the admissibility semantics is one level simpler than under the stability semantics, but it reduces to monotonic reasoning in the logic underlying the framework, thus becoming a trivial form of non-monotonic reasoning.

There appears to be a clash between these results and the intuition spelled out in Section 1, that admissibility and preferability arguments are seemingly easier to compute than stable extensions. However, our results are not as surprising as they might at first appear. Since the admissibility and preferability semantics *do not* restrict the number of extensions, one would expect that default reasoning under these semantics is as hard as under the stability semantics. The higher complexity of the sceptical reasoning problem under the preferability semantics is due to the fact that in order to verify that an assumption set is preferred, one needs to check that *none* of its supersets is admissible.

Of course, our results do not contradict the expectation that in practice constructing admissible arguments is often easier than constructing stable extensions. For example, given the propositional logic program $P \cup \{p\}$, with P any set of clauses not defining the atom p, the empty set for the query p that can be constructed "locally", without accessing P. Moreover, if $P \cup \{p\}$ is locally stratified or order-consistent [1], p is guaranteed to be a credulous consequence of the program under the stability semantics. Indeed, in all cases where the stability semantics coincides with the preferability semantics (e.g. for stratified and order-consistent abstract frameworks) any sound (and complete) computational mechanism for the admissibility semantics is sound (and complete) for the stability semantics.

The "locality" feature of the admissibility semantics renders it a feasible alternative to the stability semantics in the first-order case, when the propositional version of the given abstract framework is infinite. For example, given the (negation-free) logic program: $\{q(f(X));\ p(0)\}$, the empty set of assumptions is an admissible argument for the query $p(0)$ that can be constructed "locally", even though the propositional version of the corresponding abstract framework is infinite.

The complexity results in this paper show that sceptical reasoning under admissibility and preferability semantics is trivial and highly complex, respectively. However, this does not seem to matter for the envisioned applications of this semantics, because credulous reasoning only is required for these applications [18]. For example, in argumentation in practical reasoning in general and legal reasoning in particular, unilateral arguments are put forwards and defended against all counterarguments, in a credulous manner. Indeed, these domains appear to be particularly well suited for credulous reasoning under the admissibility semantics. In general, the results presented in this paper indicate that reasoning under the new semantics is harder. On the positive side, they indicate that the new semantics allows us to encode more complex reasoning patterns than when reasoning with the stability semantics.

Acknowledgments

The first author has been partially supported by the DFG as part of the graduate school on *Human and Machine Intelligence* at the University of Freiburg and the second author has been partially supported by the University of Cyprus.

References

[1] Andrei Bondarenko, Phan Minh Dung, Robert A. Kowalski, and Francesca Toni. An abstract, argumentation-theoretic framework for default reasoning. *Artificial Intelligence*, 93(1–2):63–101, 1997.

[2] Marco Cadoli, Maurizio Lenzerini. The Complexity of Propositional Closed World Reasoning and Circumscription. *Journal of Computer and System Sciences* 48(2): 255–310, 1994.

[3] Marco Cadoli and Marco Schaerf. A survey of complexity results for non-monotonic logics. *The Journal of Logic Programming*, 17:127–160, 1993.

[4] Yannis Dimopoulos, Bernhard Nebel and Francesca Toni. Preferred Arguments are Harder to Compute than Stable Extensions, *Proc. IJCAI'99, 16th International Joint Conference on Artificial Intelligence*, pages 36–41, T. Dean ed., Morgan Kaufmann, 1999.

[5] Yannis Dimopoulos, Bernhard Nebel and Francesca Toni. Finding Admissible and Preferred Arguments Can Be Very Hard, *Proc. KR'2000, 7th International Conference on Principles of Knowledge Representation and Reasoning* pages 53–61, T. Cohn, F. Giunchiglia, B. Selman eds, Morgan Kaufmann, 2000.

[6] Phan Minh Dung, Robert A. Kowalski, and Francesca Toni. Synthesis of Proof Procedures for Default Reasoning, *Proc. LOPSTR'96, International Workshop on Logic Program Synthesis and Transformation*, pages 313-324 J. Gallagher ed., LNCS 1207, Springer Verlag, 1996.

[7] Phan Minh Dung, Robert A. Kowalski, and Francesca Toni. Argumentation-theoretic proof procedures for default reasoning. Technical report, Imperial College, London, UK, 1998.

[8] Phan Minh Dung. Negation as hypothesis: an abductive foundation for logic programming. In K. Furukawa, editor, *Proceedings of the 8th International Conference on Logic Programming*, pages 3–17, Paris, France, 1991. MIT Press.

[9] Thomas Eiter and Georg Gottlob. Reasoning with parsimonious and moderately grounded expansions. Fundamenta Informaticae 17(1,2):31–54, 1992.

[10] Thomas Eiter and Georg Gottlob. Propositional Circumscription and Extended Closed-World Reasoning are Π_2^p-Complete. *Theoretical Computer Science* 114(2): 231–245, 1993.

[11] Thomas Eiter, Nicola Leone, and Domenico Saccà. Expressive power and complexity of partial models for disjunctive deductive databases. *Theoretical Computer Science*, 206:181–218, 1998.

[12] Kave Eshghi and Robert A. Kowalski. Abduction compared with negation as failure. In G. Levi and M. Martelli, editors, *Proceedings of the 6th International Conference on Logic Programming*, pages 234–254, Lisbon, Portugal, 1989. MIT Press.

[13] Michael R. Garey and David S. Johnson. *Computers and Intractability—A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, CA, 1979.

[14] Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In K. Bowen and R. A. Kowalski, editors, *Proceedings of the 5th International Conference on Logic Programming*, pages 1070–1080, Seattle, WA, 1988. MIT Press.

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

[16] Georg Gottlob. The complexity of default reasoning under the stationary fixed point semantics. *Information and Computation*, 121:81–92, 1995.

[17] Antonis C. Kakas, Robert A. Kowalski and Francesca Toni. The role of abduction in logic programming. *Handbook of Logic in Artificial Intelligence and Logic Programming*, 5: 235–324, D.M. Gabbay, C.J. Hogger and J.A. Robinson eds., Oxford University Press, 1998.

[18] Robert A. Kowalski and Francesca Toni. Abstract argumentation. *Journal of Artificial Intelligence and Law*, 4(3–4):275–296, 1996.

[19] Victor W. Marek and Miroslaw Truszczynski. *Nonmonotonic logic: context-dependent reasoning*. Springer-Verlag, Berlin, Heidelberg, New York, 1993.

[20] John McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13(1–2):27–39, 1980.

[21] Drew McDermott. Nonmonotonic logic II: Nonmonotonic modal theories. *Journal of the Association for Computing Machinery*, 29(1):33–57, 1982.

[22] Robert C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence,* 25:75–94, 1985.

[23] Ilkka Niemelä. Towards automatic autoepistemic reasoning. In *Logics in AI*, volume 478 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin, Heidelberg, New York, 1990.

[24] Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, Reading, MA, 1994.

[25] David Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.

[26] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1):81–132, April 1980.

[27] Domenico Saccà and Carlo Zaniolo. Stable models and non-determinism in logic programs with negation. In *Proceedings of the 9th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database-Systems (PODS-90)*, pages 205–217, 1990.

[28] Domenico Saccà. The expressive power of stable models for bound and unbound DATALOG queries. *Journal of Computer and System Sciences*, 54:441–464, 1997.

[29] J. Stillman. The complexity of propositional default logics. In *Proceedings of the 10th National Conference of the American Association for Artificial Intelligence (AAAI-92)*, pages 794–799, San Jose, CA, July 1992. MIT Press.

[30] L. Stockmeyer and A. Meyer. Word Problems Requiring Exponential Time: Preliminary Report. In *Proceedings of the 5th ACM Symposium on the Theory of Computing, (STOC 1973)*, pages 1–9, 1973.