Alternative foundations for Reiter's default logic

Thomas Linke and Torsten Schaub*
Institut für Informatik, Universität Potsdam
Postfach 60 15 53, D–14415 Potsdam, Germany
{linke,torsten}@cs.uni-potsdam.de

Abstract

We introduce an alternative conceptual basis for default reasoning in Reiter's default logic. In fact, most formal or computational treatments of default logic suffer from the necessity of exhaustive consistency checks with respect to the finally resulting set of conclusions; often this so-called extension is just about being constructed. On the theoretical side, this exhaustive approach is reflected by the usual fixed-point characterizations of extensions. Our goal is to reduce such global considerations to local and strictly necessary ones. For this purpose, we develop various techniques and instruments that draw on an analysis of interaction patterns between default rules, embodied by their mutual blocking behavior. These formal tools provide us with alternative means for addressing a variety of questions in default logic. We demonstrate the utility of our approach by applying it to three traditional problems. First, we obtain a range of criteria guaranteeing the existence and non-existence of extensions. Second, we get alternative characterizations of extensions that avoid fixed-point conditions. Finally, we furnish a formal account of default proofs that was up to now neglected in the literature.

1 Introduction

Default reasoning plays an important role in artificial intelligence systems since many underlying tasks involve reasoning from incomplete information. Reiter's *default logic* [40] is among the best known and most widely studied logical frameworks for addressing this form of reasoning. Apart from its natural and lucid language, it is also more expressive than competing formalisms, like Datalog with negation [9] or even circumscription [36]. That is, it can express succinctly knowledge situations where only exponential representations are available in the latter formalisms [8, 24].

Default logic augments classical logic by *default rules* that differ from standard inference rules in sanctioning inferences that rely upon given as well as absent information. Knowledge is represented in default logic by *default theories* (D, W) consisting of a set of formulas W and a set of default rules D. A default rule $\frac{\alpha:\beta}{\gamma}$ has two types of antecedents: A *prerequisite* α which is established if α can be derived and a *justification* β which is established if β can not be refuted. If both conditions hold, the *consequent* γ is concluded. In this way, a default theory may induce zero, one or multiple *extensions* of the facts in W.

The formal definition of extensions turns out to be more delicate than could have been expected. This is due to the context-sensitive nature of justifications. In fact, a default rule's justification could be refutable only just when all default conclusions contributing to an extension are known. This is why traditionally the non-refutability, or *consistency*, of a justification is verified with respect to the final extension. This leads to the common fixed-point characterizations of extensions. In such an approach, one is obliged to inspect the entire set of default rules in order to decide whether a particular default rule applies.

The goal of this paper is to furnish instruments that allow us to replace such exhaustive considerations by the strictly necessary ones. Our approach is based on the observation that the refutation of a justification is necessarily based on the existence of a proof for its negation. Although the default rules in these

*Affiliated with Simon Fraser University, Burnaby, Canada.

proofs are context-dependent, they delineate the set of rules that are relevant for establishing the consistency of a justification. The idea is thus to analyze default theories in order to extract all such interaction patterns among default rules in order to draw on this information for eliminating the reference to a final extension. In this way, the set of default rules must be inspect only once in its entirety, while subsequent considerations can be restricted to the extracted interaction patterns.

Our objective is nicely illustrated by applying it to the *extension-membership-problem,* which is concerned with deciding whether a default theory has an extension containing a given formula. Consider an example where birds fly, birds have wings, penguins are birds, and penguins don't fly along with a formalization through default theory

$$(D,W) = \left(\left\{ \frac{b\ :\ \neg ab_b}{f}, \frac{b\ :\ w}{w}, \frac{p\ :\ b}{b}, \frac{p\ :\ \neg ab_p}{\neg f} \right\}, \{\neg f \rightarrow ab_b, f \rightarrow ab_p, p\} \right). \tag{1}$$

We let δ_f, δ_w, δ_b, $\delta_{\neg f}$ abbreviate the previous default rules by appeal to their consequents. An analysis of these rules should provide us with the information that the application of the first rule depends on the *blockage* of the last one (and vice versa). The second and third rule can be applied no matter which of the other rules apply. We may thus infer $\neg f$ by fully ignoring the second and the third rule, while assuring that the first one is blocked. Notably, our initial analysis must be truly global and also extend to putatively unrelated parts of the theory. To see this, simply add the rule $\frac{:x}{\neg x}$, abbreviated $\zeta_{\neg x}$, destroying all previous extensions consistent with x. Now, the application of each rule depends additionally on the blockage of $\zeta_{\neg x}$. We can only apply $\delta_{\neg f}$ to derive $\neg f$, if both δ_f and $\zeta_{\neg x}$ are blocked.

We see that the application of rules depends on the blockage of other rules. The outcome of our analysis must thus provide information on which rules may block other rules (or even themselves). This is why we provide a formal account of rule interaction in terms of so-called *blocking sets.* They provide context-independent, candidate proofs for refuting a default rule's justification. But although we may allot somehow unlimited time to an analysis (seen as a compilation phase) we cannot provide unlimited space for its result. This is also crucial in our context, since the blocking information may be exponential in the size of the number of defaults. We address this problem by proposing a so-called *block graph.* This construct is obtained by abstracting from particular blocking situations, while keeping the essential interaction patterns between default rules. The size of the block graph is quadratic in the number of defaults. In the worst case, computation of the block graph has the same time complexity as the extension-membership-problem.

In fact, the block graph is our salient instrument for reducing extension-based issues to proof-oriented ones; its gathered information can be used in various ways.

- We may draw on the block graph for addressing the *existence-of-extension-problem,* aiming at deciding whether a default theory has an extension or not. As a result, we obtain a range of corresponding criteria that can be read off the block graph.

  In fact, this problem shows also up in our initial example. In addition to the rules needed for deriving a formula, we may furthermore have to consider rules menacing an encompassing extension, like $\zeta_{\neg x}$.

- Furthermore, we obtain novel characterizations of extensions by appeal to blockage-based concepts. The resulting characterizations do not only avoid global consistency checks, as needed in traditional characterizations, but they moreover eliminate the usual fixed-point condition.

- Analogously, we obtain characterizations of default proofs by means of blockage-based concepts that avoid referring to an encompassing extension.

  Moreover, the block graph can be used for focusing the formation of such default proofs on ultimately necessary defaults. In the above example, it will be the block graph of Theory (1) that tells us that the derivation of $\neg f$ through $\delta_{\neg f}$ requires that δ_f is inapplicable and that the two other rules may be completely ignored.

What makes our approach thus different from traditional, extension-oriented ones is that once we have extracted the blocking information subsequent problems can be addressed by considering default rules only by need.

Throughout this paper, we have adopted a rigorous blockage-based view. Although this notion was implicitly present in the literature, it has so far not been used as an explicit tool for addressing issues in default logic. For illustration, let us reconsider the applicability of default rules (c.f. Condition *3.* of Definition 2.1). The blockage-based view lets us distinguish three different classe of default theories:

1. If a theory contains *no interactions* among default rules, applicability depends on the rule in focus only. For instance, the applicability of $\delta_{\neg f}$ in the presence of δ_w and δ_b (and the absence of δ_f) depends only on $\delta_{\neg f}$ itself. On such non-conflicting theories, reasoning is reduced to deduction and can be accomplished by means of the constituent rules only.

2. A theory comprises *direct interactions* only, that is, the scope of each conflict affects the involved rules only. Then, applicability depends on the given rule plus the applicability of its blocking rules. For example, in the presence of δ_f, δ_w, and δ_b the applicability of $\delta_{\neg f}$ depends on $\delta_{\neg f}$ itself and δ_f. On such conflicting theories, reasoning must be done in a consistent way accounting for interactions between rules. Nonetheless, it can be accomplished by inspecting the constituent rules and (recursively) their blocking rules. Other rules that do not interact with these rules can be ignored.

3. A theory comprises *indirect interactions*, that is, conflicts that affect the application of rules not involved in the conflict. Now, applicability depends also on that of self-blocking yet unrelated rules and not only on the rules involved in the actual conflict. For instance, the applicability of $\delta_{\neg f}$ in the presence of δ_f, δ_w, δ_b, and $\zeta_{\neg x}$ depends on δ_f, $\delta_{\neg f}$, and $\zeta_{\neg x}$. As we have seen, such theories may have no extensions.

This three-fold classification is actually reflected by many notions that are introduced in the sequel. The first class is rather trivial, since it can be dealt with by standard logical means. As we will see these non-conflicting theories play nonetheless the role of a base case whenever conflicts are gradually removed from conflicting theories. Much more interesting default reasoning, involving conflicting information, is situated within the second class of theories. As a matter of fact, the profile of the underlying conflicting rules is captured by the block graph. The third class comprises the truly problematic theories. This is not only because the potential lack of extensions indicates a problem in the specification of the theory, but also because of ill-formed yet unrelated rules that necessitates the inspection of all default rules in the general case. In fact, the block graph provides a way out of this dilemma, since it allows us to decouple the existence of such ill-formed rules from the necessity to inspect the entire set of rules.

The rest of the paper is organized as follows. Section 2 gives a formal introduction to default logic. Section 3 introduces blocking sets from which we construct a default theory's block graph. This section elaborates also upon the underlying formal properties and gives further special-purpose concepts. The following three sections deal with applications of these concepts. Section 4 addresses the existence-of-extensions-problem. Section 5 provides alternative characterizations of extensions, while Section 6 is dedicated to the conception of default proofs. Section 7 grasps the presented approach in its entirety and discusses its impacts. Appendix A resumes some graph-theoretical background, needed for the proofs given in Appendix B. Sections 3, 4, and 6 heavily extend the work found in [29].

An extended version [31] of this paper contains further applications and refinements of the approach, including results on modular and skeptical reasoning as well as further insights into the relationship with variants of default logic. Also, it provides case studies on graph coloring, Hamiltonian circuits and taxonomies that aim at underpinning the utility of our approach.

Relationships to other work

The reduction of global concepts to local ones was already an issue in Reiter's seminal paper [40]. For capturing incremental constructions, Reiter isolated in [40] the property of semi-monotonicity and already showed that this property is only satisfied by the subclass of normal default theories (see Section 2 on formal details). This has led in the following years to the development of alternative default logics, among them justified [34], cumulative [6] and constrained default logic [16], all of which enjoy semi-monotonicity in full generality. Interestingly, Reiter and Criscuolo showed in [41] that normal default theories lack expressiveness and demonstrated that at least semi-normal default theories are needed for full expressive power. Since such theories lack semi-monotonicity in Reiter's default logic, the descendants of the original approach seem to be the right choice. But despite the fact that they allow for local constructions beyond normal default theories, they can also not account for the full expressive power of Reiter's default logic. In fact, as shown in [7], full "local constructibility", as embodied by semi-monotonicity, and full expressive power are incompatible. As a consequence, we cannot adopt one of the seemingly computational advantageous variants of default logic, if we want to keep the full expressive power of Reiter's original approach.

Our approach resides within Reiter's default logic, while aiming at shifting the emphasis from global, extension-based concepts to local, proof-oriented ones (whenever possible). For this purpose, we develop various techniques and instruments that draw on an analysis of interactions between default rules, manifested by their blocking behavior. While the basic intuition behind "blockage" is related to argumentation semantics [19, 5], our resulting instruments and their applications have clearly taken a more profound — since default logic specific — avenue of research.[1] The work closest to ours in the domain of default logic has been done in [28] and [37], where the notion of *conflict* is treated.

The first application of our concepts deals with the existence-of-extensions-problem. Interestingly, semi-monotonicity implies the existence of extensions, so that all of the above cited descendants of default logic guarantee extensions. Intuitively, this is due to the fact that "local constructibility" allows us to incrementally construct extensions without ever reconsidering any previous steps. The significance of the existence-of-extensions-problem has already led to several approaches, identifying subclasses of default theories always possessing extensions in Reiter's default logic. Among them, we find *normal* [40], *ordered* [21], *even*[2] [39], and *strongly stratified* [10] default theories. An algorithmic account of the existence-of-extensions-problem is proposed in [48]. We show in Section 4.4 that our conception around the block graph provides us with a range of criteria going beyond most of these proposals. The other advantage of our approach resides in its syntax-independence, which is absent in all of the previous approaches. To be fair, however, we note that our investment in constructing the block graph is also greater than that of the aforecited approaches. Lastly, we mention that notions like even- and oddness were also investigated in logic programming and truth-maintenance systems, which can be interpreted as restricted fragments of default logic.

The second application of our concepts results in a series of alternative characterizations of extensions that avoid the usual fixed-point condition. The first such non-fixed point characterizations was obtained by Etherington in [20] when defining a semantics for default logic. This was accomplished by imposing a strict partial order on the classes of models of the initial set of facts, whose maximum elements are put in correspondence with the extensions of the underlying theory. Alternative syntactic — yet fixed-point-based — characterizations were proposed in [35, 23, 47, 43]: A context-sensitive operator in the tradition of logic programming is used in [35]; extensions (of several default logics) are defined in terms of basic properties, like groundedness, regularity, etc. (see Section 2) in [23]; characterizations aiming at tableau-based implementations are developed in [47, 43]. An operational specification based on so-called processes is given in [1]. These processes amount to branches in the trees corresponding to the strict partial orders of Etherington-style semantics [20]; therefore they also avoid a fixed-point condition.

[1] We come back to these approaches in Section 3.5, 4.4, and 5.4, respectively.

[2] In the sequel, we use *ps-even* for referring to the notion of *evenness* proposed by Papadimitriou and Sideri in [39].

As with Etherington's semantics, the verification of consistency necessitates maximal sequences of rules. The principal difference between all of these approaches and ours rests on the necessity of exhaustive consistency checks. In our approach the block graph delineates the set of default rules that must be inspected for consistency checking. Also, our avoidance of fixed-point-conditions is different from that employed by Etherington and followers. While ours is accomplished by appeal to blocking relations, their approach relies on a post-filtering condition, verifying valid constructions posteriorly [20].

The final application of our concepts results in formal characterizations of default proofs. Although this question is closely related to the extension-membership-problem, it has so far been neglected in the literature. This is probably due to the fact that up to now default proofs were regarded as being extension-dependent. The extension-membership-problem is therefore usually approached by resorting to the following loop: Generate an extension, test if a formula in question is its member. If so, stop. Otherwise, repeat the loop. This procedure can be implemented by any of the extension-construction-procedures known from the literature, e.g. [26, 47, 38, 12]. In fact, Niemelä improves this approach in [38] by providing an extension-construction-procedure that allows to focus on extensions containing an initial query. We note that all of these approaches are primarily interested in the construction of extensions; the extension-membership-problem is only addressed indirectly. Proof-theoretic investigations of Reiter's default logic were done in [3, 4] on the basis of natural deduction and sequent calculi. In these calculi, however, consistency checking is also addressed in an exhaustive way by means of so-called "anti-calculi".

Extension-dependency and thus the need for exhaustive considerations vanish in the presence of semi-monotonicity. This was already exploited by Reiter in [40] for developing a query-answering procedure for normal default theories. Other local procedures were obtained in the aforecited variants of default logic. A local proof procedure for constrained default logic was given in [44]; [13] addresses the same task for Łukaszewicz' variant. Both approaches are local in the sense that they allow for deciding whether a set of default rules forms a default proof by looking at the constituent rules only. These approaches are thus centered around the concept of a default proof, which is missing in the former extension-oriented approaches. The extension-independent characterization of default proofs in Reiter's default logic is thus one of our major concerns.

2 Background

We start by completing our initial introduction to Reiter's default logic. A *default rule* $\frac{\alpha:\beta}{\gamma}$ is called *normal* if β is equivalent to γ; it is called *semi-normal* if β implies γ. We sometimes denote the *prerequisite* α of a default rule δ by $Pre(\delta)$, its *justification* β by $Just(\delta)$ and its *consequent* γ by $Cons(\delta)$.[3] A set of default rules D and a set of formulas W form a *default theory*[4] $\Delta = (D, W)$, that may induce one, multiple or even no *extensions* [40] in the following way.

Definition 2.1 *Let (D, W) be a default theory. For any set of formulas S, let $\Gamma(S)$ be the smallest set of formulas S' such that*

1. $W \subseteq S'$,
2. $Th(S') = S'$,
3. *For any* $\frac{\alpha:\beta}{\gamma} \in D$, *if* $\alpha \in S'$ *and* $\neg\beta \notin S$ *then* $\gamma \in S'$.

A set of formulas E is an extension of (D, W) if $\Gamma(E) = E$.

Observe that E is a fixed-point of Γ. Any such extension represents a possible set of beliefs about the world. For example, Default theory (1) has two extensions: $E_1 = Th(W \cup \{b, w, \neg f\})$ and $E_2 = Th(W \cup \{b, w, f\})$, while theory $(D \cup \{\frac{:x}{\neg x}\}, W)$ (where D and W are taken as in (1)) has no extension.

[3]This notation generalizes to sets of default rules in the obvious way, e.g. $Pre(D) = \{Pre(\delta) \mid \delta \in D\}$.

[4]If clear from the context, we sometimes refer to (D, W) as Δ and vice versa.

We call a default theory *coherent* if it has some extension.

For simplicity, we assume for the rest of the paper that default theories (D, W) comprise finite sets only. Additionally, we assume that for each default rule δ in D, we have that $W \cup \mathit{Just}(\delta)$ is consistent. This can be done without loss of generality because we can clearly eliminate all rules δ' from D for which $W \cup \mathit{Just}(\delta')$ is inconsistent without altering the set of extensions.

For a set of formulas S and a set of defaults D, define the *set of generating default rules* [40] as

$$\mathit{GD}(D, S) = \{ \delta \in D \mid S \models \mathit{Pre}(\delta) \text{ and } S \not\models \neg \mathit{Just}(\delta) \} \ . \tag{2}$$

By taking $\delta_f, \delta_w, \delta_b, \delta_{\neg f}$ to denote the default rules in (1), we see that the two extensions of (1) are generated by $\mathit{GD}(D, E_1) = \{\delta_w, \delta_b, \delta_{\neg f}\}$ and $\mathit{GD}(D, E_2) = \{\delta_f, \delta_w, \delta_b\}$, respectively.

Define a set of default rules D as *grounded* in a set of formulas S [46] iff there exists an enumeration $\langle \delta_i \rangle_{i \in I}$ of D such that for all $i \in I$ we have that

$$S \cup \mathit{Cons}(\{\delta_0, \ldots, \delta_{i-1}\}) \models \mathit{Pre}(\delta_i) \ . \tag{3}$$

Note that $\mathit{GD}(D, E)$ is grounded in W whenever E forms an extension of default theory (D, W). Conversely, the set $E = \mathit{Th}(W \cup \mathit{Cons}(\mathit{GD}(D, E)))$ forms an extension of (D, W) if $\mathit{GD}(D, E)$ is grounded in W.

Define a set of default rules D as *weakly regular* wrt a set of formulas S [28, 23] iff for each $\delta \in D$ we have that

$$S \cup \mathit{Cons}(D) \not\models \neg \mathit{Just}(\delta) \ . \tag{4}$$

Clearly, $\mathit{GD}(D, E)$ is weakly regular wrt W whenever E is an extension of some theory (D, W).

As shown in [43], maximal sets $D' \subseteq D$ of grounded and weakly regular default rules induce extensions of (D, W) in Łukaszewicz' variant of default logic [34]. That is, for each such D', $\mathit{Th}(W \cup \mathit{Cons}(D'))$ forms a, say, Łukaszewicz-extension. So, in contrast to what we have observed on Theory (1) for Reiter's default logic, $\mathit{Th}(W \cup \{b, w, \neg f\})$ and $\mathit{Th}(W \cup \{b, w, f\})$ do actually provide valid Łukaszewicz-extensions of Default theory $(D \cup \{\frac{:x}{\neg x}\}, W)$, where D and W are taken as in (1). This works because Łukaszewicz' variant enjoys the property of *semi-monotonicity*:

> For any sets of default rules $D' \subseteq D$, we have that if E' is an extension of (D', W) then there is an extension E of (D, W) where $E' \subseteq E$.

If a theory (D, W) enjoys this property, we can decide the extension-membership-problem for some formula φ by forming such a set E' with $\varphi \in E'$. Then, semi-monotonicity tells us that there is an extension E of (D, W) with $E' \subseteq E$. For forming a default proof for φ in Łukaszewicz' variant, it is hence sufficient to construct a grounded and weakly regular set of default rules $P \subseteq D$ with $W \cup \mathit{Cons}(P) \models \varphi$, while discarding all default rules in $D \setminus P$.[5]

Similar constructions are impossible in Reiter's default logic, due to the aforementioned reasons. On the other hand, we argue that there is no need for always inspecting all rules in $D \setminus P$. But then the question arises how to tell which rules must be considered and which rules can be ignored. An answer to this question is provided in the next section.

3 Representing interactions by block graphs

This section introduces the fundamental concepts on which we rely for analyzing possible interactions among default rules. We express these interactions through *blocking relations* that tell us which rules may block a rule in question. This information is then condensed in *block graphs* by abstracting from particular blocking situations, while keeping the essential interaction patterns between default rules. In the subsequent sections, we show how these instruments can be used for turning extension-based concepts into proof-oriented ones.

[5] We elaborate more upon the role of Łukaszewicz' variant as a "lower bound" for Reiter's default logic in [31].

3.1 Blocking sets

Our approach is founded on the concept of *blocking sets*. Given a default theory (D, W) and a default rule $\delta \in D$, intuitively, a blocking set for δ is a set of default rules $B \subseteq D$ whose joint application denies the application of δ. Such a blocking set provides a *candidate* for disabling the putatively applicable default rule δ. For this purpose, it is actually sufficient to refute a rule's justification, while ignoring its prerequisite. An existing derivation of a prerequisite can only be counterbalanced by refuting the justification of one of its default rules.

In order to become effective, however, a blocking set must be included in the set of generating default rules of an encompassing extension. That is, it must be grounded and the respective justifications must be consistent with the extension. In fact, groundedness can be effectively verified by looking at the candidate set only, while consistency is context-dependent since it refers to a final extensions. Our aim is however to capture *blockage* as an intrinsic feature of default theories rather than their resulting extensions. Moreover, we are often interested in showing that a critical blocking set *does not* apply, which rules out an extension-based characterization.

This leads us to the following definition of *blocking sets*.

Definition 3.1 *Let $\Delta = (D, W)$ be a default theory. For $\delta \in D$ and $B \subseteq D$, we define*

1. *B as a basic blocking set for δ, written $B \mapsto_\Delta \delta$, iff*

   BS1 $W \cup \mathit{Cons}(B) \models \neg \mathit{Just}(\delta)$ *and*

   BS2 *B is grounded in W.*

2. *B is an essential blocking set of δ, written $B \overset{\bullet}{\mapsto}_\Delta \delta$, iff $B \mapsto_\Delta \delta$ and*

   BS3 $(B \setminus \{\delta'\}) \not\mapsto_\Delta \delta''$ *for every $\delta' \in B$ and every $\delta'' \in B \cup \{\delta\}$.*

We define $\mathcal{B}_\Delta(\delta) = \{B \mid B \overset{\bullet}{\mapsto}_\Delta \delta\}$ as the set of all essential blocking sets of δ.

Observe that this definition treats default rules as monotonic inference rules by ignoring the rules' justifications. A blocking set for a rule δ amounts thus to a proof of $\neg \mathit{Just}(\delta)$ in a standard logical system augmented by inference rules obtained from D by ignoring default rules' justifications. Such systems are studied in [35]. Note that the consequents of a blocking set are not required to be consistent. This is needed, for instance, to detect groups of default rules whose joint application blocks any other default, like $\{\frac{:a}{c}, \frac{:b}{\neg c}\}$. Finally, note that for all $\delta \in D$ we have $\emptyset \not\mapsto_\Delta \delta$ because we require that $W \cup \mathit{Just}(\delta)$ is consistent.

For illustration, consider the following example.

$$\Delta = \left(\left\{ \frac{:\ a}{a}, \frac{:\ \neg a}{b}, \frac{a \wedge b\ :\ x}{c}, \frac{:\ \neg c}{d} \right\}, \emptyset \right) \tag{5}$$

Among the basic blocking sets, we have

$$\left\{\tfrac{:a}{a}\right\} \mapsto_\Delta \tfrac{:\neg a}{b} \qquad \text{and} \qquad \left\{\tfrac{:a}{a}, \tfrac{:\neg a}{b}, \tfrac{a \wedge b : x}{c}\right\} \mapsto_\Delta \tfrac{:\neg c}{d}\,. \tag{6}$$

In addition, all grounded supersets of $\{\frac{:\neg a}{b}\}$ and $\{\frac{:a}{a}, \frac{:\neg a}{b}, \frac{a \wedge b : x}{c}\}$ are basic blocking sets for $\frac{:\neg a}{b}$ and $\frac{:\neg c}{d}$, respectively. All these redundant supersets violate **BS3**, so that none of them is an essential blocking set. Moreover, the second basic blocking set in (6) is superfluous since it may never appear in an extension. To see this, observe that $\{\frac{:a}{a}, \frac{:\neg a}{b}, \frac{a \wedge b : x}{c}\}$ contains a blocking set for one of its constituent rules, given by the first blocking set in (6). In fact, such situations are also addressed by **BS3**. For instance, taking $\delta' = \frac{a \wedge b : x}{c}$ and $\delta'' = \frac{:\neg a}{b}$ in **BS3** shows that $\{\frac{:a}{a}, \frac{:\neg a}{b}, \frac{a \wedge b : x}{c}\}$ is no essential blocking set. Hence, the above theory has only a single essential blocking set:

$$\left\{\tfrac{:a}{a}\right\} \overset{\bullet}{\mapsto}_\Delta \tfrac{:\neg a}{b}\,.$$

We see that although **BS1** and **BS2** capture the basic characteristics of blocking sets, not all such sets are needed for blocking a given rule. This is addressed in the second part of Definition 3.1. In fact, **BS3** furnishes a concise specification imposing, first, that essential blocking sets are (set inclusion) minimal among all blocking sets of a rule and, second, that essential blocking sets do not contain blocking sets for their constituent rules. We have the following result.

Theorem 3.1 *Let $\Delta = (D, W)$ be a default theory and let $\delta \in D$ and $B \subseteq D$.*
We have that $B \overset{\bullet}{\mapsto}_\Delta \delta$ iff the following conditions hold.

1. $B \mapsto_\Delta \delta$,
2. $B' \mapsto_\Delta \delta$ *for no* $B' \subset B$*, and*
3. $B' \mapsto_\Delta \delta'$ *for no* $B' \subset B$ *and no* $\delta' \in B$.

Condition *2.* captures minimality, while Condition *3.* ensures a non-inclusion property. Note that the latter does not apply to entire blocking sets, where $B' = B$, since their set of consequents may be contradictory, as in $\{\frac{:a}{c}, \frac{:b}{\neg c}\}$.

The fact that **BS3** only eliminates superfluous blocking sets is guaranteed by the following result.

Theorem 3.2 *Let E be an extension of default theory $\Delta = (D, W)$ and let $\delta \in D$.*
Then the following are equivalent:

1. $E \models \neg Just(\delta)$
2. *there is some* $B \subseteq D$ *with* $B \mapsto_\Delta \delta$ *such that* $B \subseteq GD(D, E)$
3. *there is some* $B \subseteq D$ *with* $B \overset{\bullet}{\mapsto}_\Delta \delta$ *such that* $B \subseteq GD(D, E)$
4. *there is some* $B \in \mathcal{B}_\Delta(\delta)$ *such that* $B \subseteq GD(D, E)$

This theorem shows, that the existence of blocking sets provides necessary and sufficient conditions for refuting a default rule's justification according to the intuitions given in the introductory section. It also demonstrates that the blocking sets retained in $\mathcal{B}_\Delta(\delta)$ have the same effectiveness as their basic counterparts. Finally, it provides first evidence of how extension-oriented notions are expressible through blockage-based concepts. This is illustrated by the fact that condition $\mathcal{B}_\Delta(\delta) = \emptyset$ is sufficient for the consistency of $Just(\delta)$ with all extensions of Δ.

For further illustration, consider Theory (1) along with its blocking sets given in (7)–(10):

$$\mathcal{B}_\Delta(\delta_f) = \{\{\delta_{\neg f}\}\} \quad (7)$$
$$\mathcal{B}_\Delta(\delta_w) = \emptyset \quad (8)$$
$$\mathcal{B}_\Delta(\delta_b) = \emptyset \quad (9)$$
$$\mathcal{B}_\Delta(\delta_{\neg f}) = \{\{\delta_b, \delta_f\}\} \quad (10)$$

For example, $\{\delta_{\neg f}\}$ is the only blocking set for δ_f; it comprises a possible refutation of ab_b, the justification of δ_f. In general, a single default rule may have multiple blocking sets. For example, adding $\frac{:u}{v \wedge \neg v}$ to Theory (1) augments each set $\mathcal{B}_\Delta(\delta_i)$ by $\{\frac{:u}{v \wedge \neg v}\}$. The addition of $\frac{:x}{\neg x}$ to (1) leaves blocking sets (7)–(10) unaffected and yields $\mathcal{B}_\Delta\left(\frac{:x}{\neg x}\right) = \{\{\frac{:x}{\neg x}\}\}$, reflecting self blockage. Observe that Condition **BS3** allows us to discard blocking sets that block their own constituent rules. For instance, $\{\delta_f, \delta_b, \delta_{\neg f}\}$ is a putative blocking set of δ_w, but it is ruled out by **BS3** since it contains both δ_f and one of its blocking sets, $\{\delta_{\neg f}\}$.

Let us now look at further properties of the blocking sets kept in $\mathcal{B}_\Delta(\delta)$. In fact, they allow us to capture the conceptualization of consistency found in Reiter's default logic without any appeal to an encompassing extension, as shown next.

Theorem 3.3 *Let $\Delta = (D, W)$ be a default theory and let $D' \subseteq D$ be grounded in W.*
We have that D' is weakly regular wrt W iff we have for each $\delta' \in D'$ and each $B \subseteq D'$ that $B \notin \mathcal{B}_\Delta(\delta')$.

Seen as a mapping from default theories $\Delta = (D, W)$ to sets of sets of default rules $\mathcal{B}_\Delta(\delta)$ (taking δ as a parameter), we observe that $\lambda x.\mathcal{B}_x(\delta)$ is monotone with respect to the addition of default rules to D.

Theorem 3.4 *Let $\Delta = (D, W)$ and $\Delta' = (D', W)$ be default theories with $D \subseteq D'$.*
For $\delta \in D$, we have that

1. *$B \in \mathcal{B}_\Delta(\delta)$ implies $B \in \mathcal{B}_{\Delta'}(\delta)$ and*
2. *$B \in \mathcal{B}_{\Delta'}(\delta)$ and $B \subseteq D$ imply $B \in \mathcal{B}_\Delta(\delta)$.*

This result implies that blocking sets can be constructed in an incremental fashion (c.f. Corollary 3.9). Note that $\lambda x.\mathcal{B}_x(\delta)$ is not monotone with respect to W, since, for instance, adding a default rule's consequent to W eliminates this default from all blocking sets that contained it previously.

Let us now give some results establishing upper bounds for computational complexity. The following result for basic blocking sets draws on a similar result obtained in [45].

Theorem 3.5 *Let $\Delta = (D, W)$ be a default theory and let $\delta \in D$ and $B \subseteq D$.*
Deciding whether $B \mapsto_\Delta \delta$ holds is in co-NP.

For essential blocking sets, we may thus test for a given set B and a given δ whether $B \in \mathcal{B}_\Delta(\delta)$ with a polynomial number of calls to an NP-oracle. This gives the following result for determining whether there exists an essential blocking set for a given rule.

Theorem 3.6 *Let $\Delta = (D, W)$ be a default theory and let $\delta, \delta' \in D$.*
Deciding whether there exists an essential blocking set $B \in \mathcal{B}_\Delta(\delta)$ such that $\delta' \in B$ is in Σ_2^P.

As regards space complexity, we note that in the worst case, a theory with n rules may comprise $O(2^n)$ blocking sets. This is arguably an artifact of the problem in general rather than the specific approach at hand — there may simply be an exponential number of ways in which a set of defaults conflict. Consider for example the class of default theories discussed in [15], where we have

$$\begin{array}{rll} \frac{\alpha : \beta_{i,1}}{\beta_{i,1}} & \text{for} & i \in \{1,2\}\ , \\ \frac{\beta_{i,j} : \beta_{i',j+1}}{\beta_{i',j+1}} & \text{for} & i, i' \in \{1,2\} \text{ and } 1 \leq j < n, \text{ and} \\ \frac{\beta_{i,n} : \gamma}{\gamma} & \text{for} & i \in \{1,2\}\ . \end{array}$$

For a given n there are clearly 2^n "inferential paths" between α and γ. Given α, a default rule $\frac{: \neg\gamma}{\omega}$ is thus faced with $O(2^n)$ blocking sets. While this characterizes the worst case, in general we might expect the number of blocking sets to be more manageable. For example, in an inheritance hierarchy where a different "exception" type accounts for each level in the hierarchy, we would have a set of blocking sets that is linear in the size of the hierarchy. See [31] for detailed case-studies. On the other hand, the number of blocking sets is not related to the number of extensions of a given theory. To see this, observe that the default rules

$$\frac{: a_i}{\neg c_i} \quad \frac{: c_i}{\neg a_i} \quad \text{for} \quad i \in \{1, \ldots, n\}$$

induce 2^n extensions but only $2n$ blocking sets. That is, although we encounter an exponential number of extensions, we have only a linear number of blocking sets. The last two examples show that blockage- and extension-oriented approaches may work quite orthogonal to each other.

Finally, we show that blocking sets are independent from the representation of the underlying default theory. For this purpose, we define *syntactically equivalent* default theories as follows.

Definition 3.2 *Let $\Delta = (D, W)$ and $\Delta' = (D', W')$ be two default theories.*
We define Δ and Δ' to be syntactically equivalent if the following conditions hold.

1. *$\mathit{Th}(W) = \mathit{Th}(W')$ and*
2. *there is a bijective mapping $f : D \to D'$ such that for each $\delta \in D$, we have*

$$\begin{aligned} W &\models \mathit{Pre}(\delta) \equiv \mathit{Pre}(f(\delta)) , \\ W &\models \mathit{Just}(\delta) \equiv \mathit{Just}(f(\delta)) , \\ W &\models \mathit{Cons}(\delta) \equiv \mathit{Cons}(f(\delta)) . \end{aligned}$$

Observe that equivalence is actually defined modulo the set of premises W.

We generalize mapping f to sets and sets of sets of default rules by putting $f(B) = \{f(\delta) \mid \delta \in B\}$ for $B \subseteq D$ and $f(\mathcal{B}) = \{f(B) \mid B \in \mathcal{B}\}$ for $\mathcal{B} \subseteq 2^D$. Then, we have the following result showing that blocking sets are independent of the syntactical representation of the underlying default theory.

Theorem 3.7 *Let Δ and Δ' be syntactically equivalent default theories and $f : D \to D'$ some associated bijective mapping.*
We have for each $\delta \in D$ that $f(\mathcal{B}_\Delta(\delta)) = \mathcal{B}_{\Delta'}(f(\delta))$.

Two default theories are *extension equivalent,* if they have exactly the same extensions. Clearly, syntactic equivalence implies extension equivalence but not vice versa [35].

3.2 Block graph

Given the concept of blocking sets, we are ready to define our salient instrument: The *block graph* of a default theory.

Definition 3.3 *Let $\Delta = (D, W)$ be a default theory.*
The block graph $\Gamma_\Delta = (V_\Delta, A_\Delta)$ of Δ is a directed graph with vertices $V_\Delta = D$ and arcs

$$A_\Delta = \left\{(\delta', \delta) \mid \delta' \in B \text{ for some } B \in \mathcal{B}_\Delta(\delta)\right\} .$$

The block graph is an extract of the essential blocking information comprised in all blocking sets. This is done by abstracting from the membership of default rules in specific blocking sets. That is, there is an arc (δ', δ) between default rules δ' and δ in the block graph iff δ' belongs to *some* blocking set for δ.

For Default theory (1), we obtain the block graph given in Figure 1; it has arcs $(\delta_{\neg f}, \delta_f)$, $(\delta_f, \delta_{\neg f})$ and $(\delta_b, \delta_{\neg f})$.

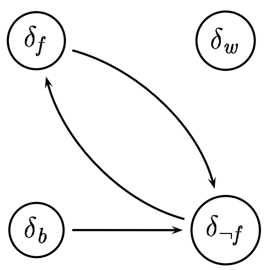

Figure 1: Block graph of Default theory (1).

We observe that the size of the block graph is always quadratic in the number of default rules, although there may be an exponential number of blocking sets in the worst case. The computational complexity associated with block graphs is directly related to that of blocking sets, as shown in the next result.

Theorem 3.8 *Let $\Delta = (D, W)$ be a default theory and let $G = (D, A)$ be a directed graph.*
The problem of deciding whether $G = \Gamma_\Delta$ is in Σ_2^P.

In view of the possibly exponential number of blocking sets, it is important to observe that we neither have to keep nor to recompute blocking sets during the construction of block graphs. This is expressed by the following corollary to the first part of Theorem 3.4.

Corollary 3.9 *Let $\Delta = (D, W)$ and $\Delta' = (D', W)$ be default theories with $D \subseteq D'$. Let $\Gamma_\Delta = (V_\Delta, A_\Delta)$ and $\Gamma_{\Delta'} = (V_{\Delta'}, A_{\Delta'})$ be the block graphs of Δ and Δ'.*
For $\delta, \delta' \in D$, we have that $(\delta, \delta') \in A_\Delta$ implies $(\delta, \delta') \in A_{\Delta'}$.

Observe that the second part of Theorem 3.4 is inapplicable, since it refers to blocking sets whose entity has disappeared in the block graph. To see this, consider a theory Δ' with rules $\frac{:a}{a}$, $\frac{a:b}{b}$, $\frac{:\neg b}{c}$ and no facts. Assume Δ is obtained from Δ' by deleting $\frac{:a}{a}$. Then, we have $\{\frac{:a}{a}, \frac{a:b}{b}\} \mapsto_{\Delta'} \frac{:\neg b}{c}$ while $\{\frac{a:b}{b}\} \not\mapsto_\Delta \frac{:\neg b}{c}$.

Another property carrying over from blocking sets to block graphs is that of syntax independence.

Corollary 3.10 *Let $\Delta = (D, W)$ and $\Delta' = (D', W')$ be syntactically equivalent default theories with block graphs $\Gamma_\Delta = (V_\Delta, A_\Delta)$ and $\Gamma_{\Delta'} = (V_{\Delta'}, A_{\Delta'})$. Let $f : D \to D'$ be some associated bijective mapping.*
We have for all $\delta, \delta \in D$ that $(\delta, \delta') \in A_\Delta$ iff $(f(\delta), f(\delta')) \in A_{\Delta'}$.

Although a block graph gives up full blocking information in return for better space complexity, it remains a powerful instrument for the purposes elaborated upon in sections 4 to 6. This is due to the fact that the block graph delineates the set of rules that comprise putative inconsistencies. Among others, this allows us to limit our attention to such sets when checking the consistency of δ's justification. The next section draws on this for providing a conceptual alternative to such consistency checks.

3.3 Supporting sets

From the perspective of blocking sets, a consistency check must guarantee that all blocking sets of a rule in focus are inapplicable. This leads us to the concept of *supporting sets*, which are intuitively simply blocking sets for blocking sets.

We first extend the notion of blocking sets to sets of rules: For a default theory $\Delta = (D, W)$ and sets $B, B' \subseteq D$, we call B' a *blocking set* for B, written $B' \overset{\bullet}{\mapsto}_\Delta B$, if there is some default rule $\delta \in B$ such that $B' \overset{\bullet}{\mapsto}_\Delta \delta$, or equivalently $B' \in \mathcal{B}_\Delta(\delta)$. Note that $B \not\overset{\bullet}{\mapsto}_\Delta \emptyset$ for each $B \subseteq D$.

With this, we may define the concept of supporting sets as follows.

Definition 3.4 *Let $\Delta = (D, W)$ be a default theory. For $\delta \in D$, we define the set $\mathcal{S}_\Delta(\delta)$ of all supporting sets for δ as*

$$\mathcal{S}_\Delta(\delta) = \{B'_1 \cup \cdots \cup B'_n \mid B'_i \subseteq D \text{ such that } B'_i \overset{\bullet}{\mapsto}_\Delta B_i \text{ for } i \in \{1, \ldots, n\} \text{ and } \mathcal{B}_\Delta(\delta) = \{B_1, \ldots, B_n\}\}$$

provided $\mathcal{B}_\Delta(\delta) \neq \emptyset$. Otherwise, we define it as $\mathcal{S}_\Delta(\delta) = \{\emptyset\}$.

Observe that $\mathcal{S}_\Delta(\delta) = \emptyset$ whenever $\mathcal{B}_\Delta(\delta') = \emptyset$ for all δ' in some $B_i \in \mathcal{B}_\Delta(\delta)$, because then for B_i there is no set of default rules B'_i such that $B'_i \overset{\bullet}{\mapsto}_\Delta B_i$, that is, $B'_1 \cup \cdots \cup B'_n$ is undefined.

The purpose of supporting sets is to rule out blocking sets as subsets of the generating default rules. Once a supporting set for δ has been applied (i.e. once it is a subset of the generating default rules), δ itself can be applied safely. Observe, however, that supporting sets may be inapplicable, as in the case of $\frac{:x}{\neg x}$, which forms both its own blocking and its own supporting set. The supporting sets in Theory (1) are given in (11)–(14):

$$\begin{aligned} \mathcal{S}_\Delta(\delta_f) &= \{\{\delta_b, \delta_f\}\} && (11)\\ \mathcal{S}_\Delta(\delta_w) &= \{\emptyset\} && (12)\\ \mathcal{S}_\Delta(\delta_b) &= \{\emptyset\} && (13)\\ \mathcal{S}_\Delta(\delta_{\neg f}) &= \{\{\delta_{\neg f}\}\} && (14)\end{aligned}$$

Consider the supporting set for δ_f in (11). We have to find *one* blocking set for each blocking set in $\mathcal{B}_\Delta(\delta_f) = \{\{\delta_{\neg f}\}\}$. In this easy case, we have to find some blocking set for $\delta_{\neg f}$, yielding $\{\delta_b, \delta_f\}$. Here, $\{\delta_b, \delta_f\}$ is the only supporting set for δ_f. Similarly, for $\delta_{\neg f}$, we have to find a blocking set B' for $\{\delta_b, \delta_f\}$ (see (10)). That is, we must have either $B' \in \mathcal{B}_\Delta(\delta_b)$ or $B' \in \mathcal{B}_\Delta(\delta_f)$. Because $\mathcal{B}_\Delta(\delta_b)$ is empty, we get $\{\delta_{\neg f}\} \in \mathcal{B}_\Delta(\delta_f)$ as the only supporting set for $\delta_{\neg f}$. The occurrence of $\delta_{\neg f}$ in its supporting set is due to the fact that there is a direct conflict between $\delta_{\neg f}$ and its blocking set. This says that $\delta_{\neg f}$ is safely applicable on its own. In general this need not be the case. For example, in default theory $(\{\frac{:b}{a}, \frac{:c}{\neg b}, \frac{:d}{\neg c}, \}, \emptyset)$ the last rule forms the single supporting set for the first one.

Since blocking sets are context-independent they represent merely candidate proofs for refuting a default rule's justification. That is, for a default rule there may be some extension containing such a blocking set that inhibits the application of the rule. Thus, given only a rule and one of its blocking sets, we cannot decide whether the rule applies without the final extension. The situation is different with supporting sets. Clearly, supporting sets are also context-independent. But unlike blocking sets, they are supposed to apply in the *same* extension as their supported rule. This can be made precise as follows.

Theorem 3.11 *Let $\Delta = (D, W)$ be a default theory. For every $\delta \in D$ and every extension E of Δ, we have that if $E \models Pre(\delta)$ and $S \subseteq GD(D, E)$ for some $S \in \mathcal{S}_\Delta(\delta)$, then $\delta \in GD(D, E)$.*

In fact, the joint application of a rule and one of its supporting sets can only be denied by an independent self-blocking part of the theory that destroys an encompassing extension. Thus, given a rule and one of its supporting sets, we can decide whether the rule (and its supporting set) applies whenever we can rule out sources of incoherence. This is one of the key ideas developed in Section 5 and 6.

3.4 Detecting and reconstructing blocking sets

Although the information gathered in the block graph is often sufficient for addressing problems like the existence-of-extension-problem, the detection and reconstruction of blocking sets remain important issues for deciding the applicability of default rules.

For expressing how a block graph Γ_Δ may facilitate addressing these issues, we let $\gamma^{-1}(\delta)$ denote the set of predecessors of δ in Γ_Δ and let $\gamma^{-2}(\delta)$ denote the set of predecessors of the nodes in $\gamma^{-1}(\delta)$. For $D' \subseteq D$, define $\gamma^{-1}(D') = \{\delta \in \gamma^{-1}(\delta') \mid \delta' \in D'\}$ and $\gamma^{-2}(D') = \{\delta \in \gamma^{-2}(\delta') \mid \delta' \in D'\}$.

The specific blocking and supporting sets are delineated by the block graph in the following way.

Theorem 3.12 *Let $\Delta = (D, W)$ be default theory and let Γ_Δ be its block graph.*
We have for all default rules $\delta \in D$ that

1. *$B \subseteq \gamma^{-1}(\delta)$ for all $B \in \mathcal{B}_\Delta(\delta)$, and*
2. *$S \subseteq \gamma^{-2}(\delta)$ for all $S \in \mathcal{S}_\Delta(\delta)$.*

In what follows, we give some sufficient, block graph based conditions for detecting blocking and supporting sets. Recall from Theorem 3.2 that the absence of blocking sets for a rule establishes the consistency of its justification. For illustration, consider the predecessor sets obtained from the block graph in Figure 1:

$$\begin{aligned} \gamma^{-1}(\delta_f) &= \{\delta_{\neg f}\} & \gamma^{-2}(\delta_f) &= \{\delta_f, \delta_b\}\\ \gamma^{-1}(\delta_w) &= \emptyset & \gamma^{-2}(\delta_w) &= \emptyset\\ \gamma^{-1}(\delta_b) &= \emptyset & \gamma^{-2}(\delta_b) &= \emptyset\\ \gamma^{-1}(\delta_{\neg f}) &= \{\delta_f, \delta_b\} & \gamma^{-2}(\delta_{\neg f}) &= \{\delta_{\neg f}\}\end{aligned} \qquad (15)$$

Our succinct example illustrates already some simple criteria that can be directly obtained from the block graph. First, whenever we have $\gamma^{-1}(\delta) = \emptyset$ a default rule δ is applicable without consistency check. Second, whenever we have $\gamma^{-2}(\delta) = \{\delta\}$ a default rule δ can support itself, [6] in an existing extension. This is thus inapplicable to rules like $\frac{:x}{\neg x}$.

In fact, the following conditions are sufficient for the absence of blocking and the presence of supporting sets, respectively, in a set of rules D':

$$\begin{aligned} D' \cap \gamma^{-1}(D') &= \emptyset & (16)\\ \gamma^{-2}(D') &\subseteq D', \quad \text{if for every } \delta' \in D' \text{ and } \delta'' \in \gamma^{-1}(\delta'), \text{ we have } \gamma^{-1}(\delta'') \neq \emptyset. & (17)\end{aligned}$$

The first condition tells us that D' contains no blocking set for any of its members. The second condition makes sure that all blocking sets of all members of D' are inhibited in D' because all supporting sets are present. The condition in (17) excludes cases where rules have blocking but no supporting sets.

More generally, we have the following sufficient criteria for the absence and presence, respectively, of blocking sets for a default rule δ in a grounded set of rules D'. There is no blocking set in D' for δ, if

$$W \cup \mathit{Cons}\big(D' \cap \gamma^{-1}(\delta)\big) \not\models \neg \mathit{Just}(\delta)\ . \qquad (18)$$

On the other hand, there is some blocking set in D' for δ, if

$$W \cup \mathit{Cons}\big(D' \cap \gamma^{-1}(\delta)\big) \models \neg \mathit{Just}(\delta)$$

and provided that $D' \cap \gamma^{-1}(\delta)$ is grounded.

Finally, let us consider the verification of weak regularity for a grounded set D' of default rules. According to Theorem 3.3, we must show that D' contains no blocking sets for its members. This is true if (18) holds for each $\delta \in D'$. However, given the underlying block graph $\Gamma_\Delta = (V_\Delta, A_\Delta)$, we may restrict our attention to rules δ that belong to the following subset of D' :

$$C' = \big\{\delta \in D' \mid (\delta', \delta) \in \big((D' \times D') \cap A_\Delta\big)\big\}\ .$$

For illustration, consider $D' = \{\delta_w, \delta_b, \delta_{\neg f}\}$ along with the block graph in Figure 1. We get $C' = \{\delta_{\neg f}\}$ and $D' \cap \gamma^{-1}(\delta_{\neg f}) = \{\delta_b\}$ for which (18) holds — while ignoring all rules in $D' \setminus C' = \{\delta_w, \delta_b\}$.

The last criteria illustrate the block graph's role in delineating sets of "critical rules". In the worst case, we are faced with a complete block graph, from which no gain is to be expected. Otherwise, it should be clear that the smaller the sets $\gamma^{-1}(\delta)$ and $\gamma^{-2}(\delta)$, the larger the pay-off obtained by means of the block graph. This question is further elaborated upon in [31], where coloring techniques are used to gather more information in block graphs.

3.5 Related approaches

As mentioned in the introductory section, our approach shares some of its basic intuitions with argumentation semantics. In the pioneering work of Dung in [19] an *argumentation framework* is a pair (A, R) where A is a set of arguments and $R \subseteq A \times A$ represents an "attack" relation between arguments. In this framework, a default theory can be (informally) interpreted via arguments of the form (J, ϕ), where $\mathit{Just}(P) \subseteq J$ for some "default proof" P of ϕ (c.f. Definition 6.1). And (J, ϕ) attacks (J', ϕ') iff $\neg\phi \in J'$. Observe that this abstract setting gives an infinite set of arguments, amounting to all possible default proofs drawable from an underlying default theory.

This framework in refined in [5] by considering *assumption-based frameworks* of form $(T, A, \overline{})$ where T and A are sets of formulas and $\overline{}$ is a mapping from A in the underlying language. A stands for a set of assumptions that can be used for extending a theory T; $\overline{}$ maps assumptions to their contrary. Among other nonmonotonic reasoning formalisms, default logic has been shown to be an instance of this abstract

[6] This is made precise in condition **PTD2** in Definition 5.1.

framework. Given a theory (D, W), default rules $\frac{\alpha : \beta}{\gamma} \in D$ are identified with monotonic inference rules of form $\frac{\alpha, M\beta}{\gamma}$. Adding these rules to classical logic, gives an inference relation $\vdash_D$. Further, let $T = W$ and $A = \{M\beta \mid \beta \in \mathit{Just}(D)\}$; the contrary $\overline{M\beta}$ of an assumption $M\beta \in A$ is $\neg\beta$. Given an assumption set $A' \subseteq A$ and an assumption $M\beta \in A$, A' is said to *attack* $M\beta$ iff $W \cup A' \vdash_D \neg\beta$.

Despite their different objects of discourse, blocking sets can be associated with redundancy-free attacks. In fact, we have that

- $B \mapsto_\Delta \delta$ implies that $\{M\beta \mid \beta \in \mathit{Just}(B)\}$ attacks $M\,\mathit{Just}(\delta)$, as well as
- $B \overset{\bullet}{\mapsto}_\Delta \delta$ implies that $\{M\beta \mid \beta \in \mathit{Just}(B)\}$ attacks $M\,\mathit{Just}(\delta)$.

To see that the converse does not hold in either case, consider theory (5). We get $A = \{Ma, M\neg a, Mx, M\neg c\}$. However, while A is an attack for $M\neg a$ and $M\neg c$ its counterpart $\left\{\frac{:a}{a}, \frac{:\neg a}{b}, \frac{a\wedge b : x}{c}, \frac{:\neg c}{d}\right\}$ is neither a basic nor an essential blocking set for any of its members. One can show that there is a correspondence between the minimal basic blocking sets of a rule δ and the minimal attacks of $\mathit{Just}(\delta)$. Moreover, in analogy to the discussion after (6), we have that $\{Ma, M\neg a, Mx\}$ attacks Mc, which has no essential blocking set as counterpart. This shows that attacks comprise much more redundancy than blocking sets. To be fair, however, one should bear in mind that the concept of an attack resides within an abstract framework, while our approach provides an infrastructure for reasoning, being specific to default logic. We return to this approach in Section 5.4.

As regards other work specific to default logic, we mention that the works in [28] and [37] treat the related notion of *conflict* dealing with minimal sets of default rules having inconsistent consequents rather than proof skeletons menacing particular justifications, as in our approach.

4 Existence of extensions

Determining whether a default theory has an extension is a fundamental problem in default logic. This question is also pertinent to the extension-membership-problem, since reasonable conclusions must reside in an existing extension. In previous works, broad subclasses of default theories always possessing extensions have been identified. Among them we find *normal* [40], *ordered* [21], *ps-even* [39][7], and *strongly stratified* [10] default theories.

We address this problem by exploiting blocking relations among default rules by means of the formal tools developed in the last section. This provides us with a range of sufficient conditions for the coherence and incoherence of default theories.

4.1 Block graph based criteria

To begin with, we call a default theory (D, W) *non-conflicting,* if it has no blocking sets, that is, if its block graph has no arcs; otherwise we call it *conflicting.* Non-conflicting default theories have unique extensions and trivially allow for inferences without consistency checks. [8]

Theorem 4.1 *Every non-conflicting default theory has a single extension.*

For instance, default theory $(\{\frac{:a}{a}, \frac{:b}{b}\}, \emptyset)$ is non-conflicting, yielding a block graph with no arcs. The same holds for Theory (1), when eliminating either $\frac{b : \neg ab_b}{f}$ or $\frac{p : \neg ab_p}{\neg f}$.

More interestingly, we call a default theory *well-ordered,* if its block graph is acyclic. Theorem 4.1 can be strengthened by the next result, which shows that well-ordered default theories have single extensions.

Theorem 4.2 *Every well-ordered default theory has a single extension.*

[7] We use ps-even for referring to the notion of *evenness* due to Papadimitriou and Sideri [39].

[8] Observe that although non-conflicting default theories ignore justifications, they are still non-monotonic because extensions may be invalidated after augmenting a non-conflicting theory.

For instance, default theory $(\{\frac{:a}{a}, \frac{:b\wedge\neg a}{b}\}, \emptyset)$ is well-ordered; its block graph contains a single arc, indicating that the first rule may block the second one (but *not* vice versa).

We call a default theory *even*, if its block graph contains cycles with even length only. Our main result of this section states that even default theories always have extensions.

Theorem 4.3 *Every even default theory has an extension.*

For instance, default theory $(\{\frac{:a\wedge\neg b}{a}, \frac{:b\wedge\neg a}{b}\}, \emptyset)$ is even; its block graph contains two arcs, indicating that the first rule may block the second one, *and* vice versa.

Evenness is also enjoyed by our initial default theory in (1), as can be easily verified by regarding its block graph in Figure 1. Unlike this, default theory $(\{\frac{:x}{\neg x}\}, \emptyset)$ is not even, since its block graph contains an odd cycle of length one.

The above criteria provide us with a strict hierarchy of default theories always possessing extensions. The advantage of these criteria is that they are easily verified, once a block graph has been computed. That is, they can be tested in polynomial time and they rely on a simple data structure. Moreover, they apply to general default theories and they are syntax-independent, unlike other approaches [21, 39] that apply to semi-normal default theories only and that give different results on equivalent yet syntactically different theories, as detailed in Section 4.4.

4.2 A blocking set based criterion for deciding coherence

The last criteria were based on the abstraction from specific blocking sets furnished by the block graph. Although this results in a much better space complexity, there is a price to pay. The criteria fail to capture the entire class of coherent theories. In fact, we can decide the existence-of-extension-problem for arbitrary default theories, when considering the underlying blocking sets instead of their block graph. This is to the best of our knowledge the first complete characterization for this problem.

For this, define a directed graph on all blocking sets of a default theory as follows.

Definition 4.1 *Let $\Delta = (D, W)$ be a default theory. The graph $\Omega_\Delta = (V_\Delta^\Omega, A_\Delta^\Omega)$ is a directed graph with vertices $V_\Delta^\Omega = \{B \in \mathcal{B}_\Delta(\delta) \mid \delta \in D\}$ and arcs*

$$A_\Delta^\Omega = \left\{(B, B') \mid B \overset{\bullet}{\mapsto}_\Delta B'\right\}.$$

The directed graph Ω_Δ represents the complete blocking information of Δ.

Furthermore, we need the following definition.

Definition 4.2 *Let $G = (V, A)$ be a directed graph and $K \subseteq V$ an independent subset of vertices. K is an inverse kernel of G iff for all nodes $u \in V \setminus K$ exists a $v \in K$ such that $(v, u) \in A$.*

Then, we have the following result.

Theorem 4.4 *Let $\Delta = (D, W)$ be a default theory.*
We have that Ω_Δ has an inverse kernel iff Δ has an extension.

In fact, we show in Proof 4.3 that every even theory Δ induces an inverse kernel in its graph Ω_Δ.

Moreover, there are non-even theories inducing such kernels. Let us illustrate this via the example, used by Etherington in [20] to show that semi-normal default theories may lack extensions:

$$\Delta = \left(\left\{\frac{:a\wedge\neg b}{a}, \frac{:b\wedge\neg c}{b}, \frac{:c\wedge\neg a}{c}\right\}, \emptyset\right), \tag{19}$$

for short $\delta_a, \delta_b, \delta_c$, respectively. This theory has no extension. From sets $\mathcal{B}_\Delta(\delta_a) = \{\{\delta_b\}\}$, $\mathcal{B}_\Delta(\delta_b) = \{\{\delta_c\}\}$, and $\mathcal{B}_\Delta(\delta_c) = \{\{\delta_a\}\}$, we get a block graph Γ_Δ with arcs $(\delta_c, \delta_b), (\delta_b, \delta_a), (\delta_a, \delta_c) \in A_\Delta$ whose arcs form an odd cycle. Adding formula $c \rightarrow b$ to (19) yields actually a theory Δ' whose only extension contains c. The block graph of this theory contains in addition to the odd cycle the arc (δ_c, δ_a). In fact, this arc counterbalances the self-blocking-behavior of the odd cycle. The last two situations cannot be distinguished via the criteria of Section 4.1. Thus, none of them is able to indicate an extension in the second case, since $\Gamma_{\Delta'}$ has still an odd cycle. This is different from the criterion expressed in Theorem 4.4 that indicates an extension in the second case. To see this, observe that $\{\delta_c\}$ forms an inverse kernel of $\Omega_{\Delta'}$.

4.3 Non-existence of extensions

Our exposition was so far dominated by tests guaranteeing the existence of extensions, although tests for their non-existence are also of interest. To see this, reconsider Theory (19). There, we observed how an odd cycle was counterbalanced by an arc from outside the odd cycle. In fact, we can show that an odd cycle destroys all extensions whenever there is no such arc (and no chords in the cycle).

Let us make this precise in the sequel. Given a cycle C in a directed graph, an arc between two nodes of C is called a *chord*, if it does not belong to the arcs of C.

Definition 4.3 *Let $\Delta = (D, W)$ be a default theory with block graph $\Gamma_\Delta = (D, A_\Delta)$.*
We define a cycle $C \subseteq D$ in Γ_Δ as harmful to Δ iff

1. *C has no chords and*
2. *there is no $\delta \in (D \setminus C)$ such that $(\delta, \delta') \in A_\Delta$ for some $\delta' \in C$.*

Now, we are ready to prove the following result on the non-existence of extensions.

Theorem 4.5 *Let $\Delta = (D, W)$ be default theory and $C \subseteq D$.*
If C is an harmful, odd cycle in Γ_Δ, then Δ has no extension.

For illustration, consider the odd cycle in the block graph of Theory (19). This cycle satisfies both conditions in Definition 4.3, indicating that (19) has no extension.

Although Theorem 4.5 does not furnish a complete characterization for default theories without extensions, it provides nonetheless an easy, block graph based test that allows us to shorten the gap towards those theories whose extensions are detectable by means of the criteria given in Section 4.1.

4.4 Related approaches to the existence-of-extension-problem

Historically, normal default theories were the first class for which the existence of extensions was demonstrated [40]. One may wonder why they have not played a special role so far. The reason is that normal default rules are involved in the reasoning process as any other rules. For instance, take a rule $\frac{:b}{c}$ and no facts; this gives an extension containing *c*. But adding normal rule $\frac{c:\neg b}{\neg b}$ destroys this extension and leaves us with an incoherent theory. This should illustrate that normal default rules deserve the same attention as any other rule. In fact, neither our approach nor any of the following ones is able to indicate — by its proper means[9] — the existence of extensions for normal default theories.

Etherington's ordered default theories. The pioneering work on the existence-of-extension-problem was done by Etherington in [20], although the problem was already discussed in [40]. Etherington's approach applies to semi-normal default theories in clausal form. The idea is to extract from such theories a relation on literals. Intuitively, this relation was meant to capture the inferential dependency among literals. Then, a semi-normal default theory was said to be *ordered* if the resulting relation was irreflexive, that is, no literal depended on itself.

Ordered theories were then supposed to possess at least one extension. Unfortunately this turned out to be wrong. To see this, consider default theory

$$\left(\left\{\frac{:c\wedge b}{c}, \frac{c:\neg b}{\neg b}\right\}, \emptyset\right) .$$

[9] That is, without providing a priori a special case handling normal default rules, as done in [10].

Although this semi-normal default theory is ordered according to [20, p. 86, Definition $\leqq$ and $\ll$], it has no extension. This counterexample applies also to the improvements made in [2].

We obtain from the previous theory a block graph having two arcs pointing from both default rules to the first one. This graph has thus an odd cycle that renders the underlying theory incoherent.

Papadimitriou and Sideri's even default theories. Papadimitriou and Sideri generalized Etherington's approach in [39]. Their approach is also restricted to semi-normal default theories in conjunctive normal form. In analogy to [20], a relation is extracted from these theories in order to capture the dependency among literals. This relation is used to define a directed graph with nodes D. Papadimitriou and Sideri show in [39] that any default theory, whose corresponding graph has only even cycles, possesses an extension. For clarity, we refer to such theories as being *ps-even* .

Consider the theories

$$\left(\left\{\frac{:c}{c}\right\},\emptyset\right) \qquad \text{and} \qquad \left(\left\{\frac{:c\wedge(c\vee\neg c)}{c}\right\},\emptyset\right)$$

both of which have the same extension $Th(\{c\})$. However, both theories yield a different graph in the approach of Papadimitriou and Sideri. While the first one satisfies ps-evenness, the second does not satisfy ps-evenness. This demonstrates that ps-evenness is syntax-dependent.

Our approach yields for both theories the same block graph, having a single node and no arcs. Both theories are thus recognized as being non-conflicting and as possessing a single extension. This shows that there are even default theories that are *not* ps-even. Conversely, we have the following result.

Theorem 4.6 *Let Δ be a semi-normal default theory in conjunctive normal form.*
If Δ is ps-even, then it is even.

We see that the block graph based criterion of evenness is more general than its counterpart in the approach of Papadimitriou and Sideri: (i) it is not restricted to a fragment of default logic, (ii) it is syntax-independent, and (iii) it is more expressive on the fragment dealt with by Papadimitriou and Sideri. To be fair, however, we note that our investment in constructing the block graph is also greater than that needed for constructing the graphs for ps-evenness [10].

Cholewiński's stratified default theories. Cholewiński adapts in [10] the notion of *stratification*, known from logic programming, for default logic. He then proves that so-called *strongly stratified default theories* always possess extensions. Intuitively, this criterion distinguishes default theories whose rules can be ordered by means of a *stratification function.*

In addition to this ordering condition, however, stratified default theories impose a rather severe restriction on the interplay between the premises W of a default theory and its default rules D. For stratified default theoriesit is required that the language of W and $Cons(D)$ must be disjoint. For instance, this prevents stratification techniques to recognize the existence of simple normal theories, like $\left(\left\{\frac{:a}{a},\frac{:b}{b}\right\},\{\neg a\vee\neg b\}\right)$.

As opposed to all aforementioned approaches, the definition of stratification provides a particular account for normal default rules, having *syntactically* equivalent justifications and consequents (see [10]). However, the approach fails to capture the existence of extensions for *semantically* normal default theories, due to a lack of syntax-independence. To see this, consider theories

$$\left(\left\{\frac{:a}{a},\frac{:b}{b}\right\},\emptyset\right) \qquad \text{and} \qquad \left(\left\{\frac{:a\wedge(x\vee\neg x)}{a\wedge(y\vee\neg y)},\frac{:b\wedge(y\vee\neg y)}{b\wedge(x\vee\neg x)}\right\},\emptyset\right)\ ,$$

both of which have the same extension $Th(\{a,b\})$. As detailed in [31], the first default theory is strongly stratified, which is not the case for the second one. This demonstrates that stratification is syntax-dependent. Of course, this is rectifiable by replacing the underlying concept of "syntactical equivalence" by "logical equivalence"; however, this means also passing from a subproblem in P to one in NP.

[10] Note that [17, 18] use similar graphs and kernels to ensure existence of extensions for the restricted class of propositional, prerequisite-free, conjunctive default theories.

Our approach yields for both theories the same arcless block graph, indicating the existence of a single extension. To be fair, we recall that the computation of the block graph is probably beyond NP.

Bondarenko et al.'s order-consistent assumption-based frameworks. In [5], a minimality condition is imposed on attacks in order to define an *attack relationship graph* for assumption-based frameworks. This graph is used to define *stratified* and *order-consistent* assumption-based frameworks (see [5] for details). We have the following result.

Theorem 4.7 *Let (D, W) be a default theory and let $\langle T, A, \overline{\ } \rangle$ be the corresponding assumption-based framework.*

1. *If $\langle T, A, \overline{\ } \rangle$ is stratified according to [5], then (D, W) is well-ordered.*
2. *If $\langle T, A, \overline{\ } \rangle$ is order-consistent according to [5], then (D, W) is even.*

To see that stratification and order-consistence are strictly weaker concepts than well-orderedness and evenness, respectively, consider the following extension of theory (5)

$$\left(\left\{\frac{:a}{a}, \frac{:\neg a}{b}, \frac{a \wedge b : x}{c}, \frac{:\neg c}{d}, \frac{:\neg d}{\neg a}\right\}, \emptyset\right)$$

This theory is neither stratified nor order-consistent, whilst it is well-ordered and thus also even. Let us explain this in terms of minimal basic blocking sets, since they correspond to minimal attacks. We have

$$\left\{\frac{:a}{a}, \frac{:\neg a}{b}, \frac{a \wedge b : x}{c}\right\} \mapsto_\Delta \frac{:\neg c}{d}, \quad \left\{\frac{:\neg c}{d}\right\} \mapsto_\Delta \frac{:\neg d}{\neg a} \quad \text{and} \quad \left\{\frac{:\neg d}{\neg a}\right\} \mapsto_\Delta \frac{:a}{a}.$$

This induces an odd cycle between $\frac{:\neg c}{d}$, $\frac{:\neg d}{\neg a}$, and $\frac{:a}{a}$ in the corresponding attack relationship graph. In contrast to this, the essential blocking sets of the above theory induce an acyclic block graph, which allows us to establish the existence of a single extension.

5 Alternative characterizations of extensions

This section furnishes alternative characterizations of extensions by appeal to blocking sets. It lays the formal foundations for our elaboration upon local, proof-oriented concepts for default logic. To this end, we shift the emphasis from extensions to their underlying sets of generating default rules. The application of a set of default rules depends on several issues. Apart from groundedness, it involves protecting the constituent default rules against blockage and assuring an encompassing extension. We start by giving a formal account of the first issue, while the second one can be addressed by the criteria developed in the last section.

5.1 Protectedness

The concept of a set of default rules being "protected against blockage" can be made precise as follows.

Definition 5.1 *Let $\Delta = (D, W)$ be a default theory.*
A set of default rules $D' \subseteq D$ is protected in Δ iff for each $\delta \in D'$ we have that

PTD1 $B \subseteq D'$ *for no* $B \in \mathcal{B}_\Delta(\delta)$ *and*

PTD2 $S \subseteq D'$ *for some* $S \in \mathcal{S}_\Delta(\delta)$.

In words, a set of defaults is protected if it contains no blocking set for any of its defaults and if it contains some supporting set for each constituent default. For example, $D' = \{\delta_b, \delta_f\}$ is protected in (1). A set like $\{\frac{:x}{\neg x}\}$ can not be protected. Although $\{\frac{:x}{\neg x}\}$ is its own supporting set, which establishes **PTD2**, it fails to satisfy **PTD1**.

We note that protectedness depends exclusively on the rules in D' and those connected to D' in the block graph. In fact, **PTD1** refers to rules in D' only so that it remains unaffected when increasing $D \setminus D'$. It is therefore monotonic wrt the addition of default rules to D. Semantically, **PTD1** is the blockage-oriented counterpart of weak regularity (c.f. Theorem 3.3). As opposed to the local character of **PTD1**, condition **PTD2** controls the interaction with rules external to D'. **PTD2** guarantees that there are no blocking sets outside of D'. The scope of this interaction is delineated by the pre-predecessors of D' in the block graph, among which we find the supporting sets needed for protecting D' against its blocking sets. In all, (grounded) protected sets can be regarded as fully independent components for generating default rules. This important fact is made precise in Theorem 5.3 below.

In fact, the generating default rules of an extension form themselves a protected set of default rules.

Theorem 5.1 *Let $\Delta = (D, W)$ be a default theory and let E be a set of formulas.*
If E is an extension of Δ, then $GD(D, E)$ is protected in Δ.

5.2 Characterizing extensions without fixed-points

By combining the notion of protectedness with a coherence condition, we obtain a series of alternative characterizations of extensions, all of which are based on Theorem 5.3 below.

For expressing this result, we first need the following definition.

Definition 5.2 *Let $\Delta = (D, W)$ be a default theory and $D' \subseteq D$.*
We define

$$\Delta|D' \qquad \text{as} \qquad \big(\ D \setminus (D' \cup \overline{D'})\ ,\ W \cup Cons\big(D'\big)\ \big)$$

where $\overline{D'} = \{\delta \in D \mid W \cup Cons(D') \models \neg Just(\delta)\}$.

The purpose of $\overline{D'}$ is to eliminate defaults whose justification is inconsistent with the facts of $\Delta|D'$. Intuitively, the operation $\Delta|D'$ results in a default theory, simulating the application of the rule set D' to theory Δ. This is made precise in the following theorem.

Theorem 5.2 *Let $\Delta = (D, W)$ be a default theory and let E be a set of formulas. Further, let $D' \subseteq GD(D, E)$ be grounded in W.*
We have that E is an extension of Δ iff E is an extension of $\Delta|D'$.

Using this concept, we can formulate the following major result.

Theorem 5.3 *Let $\Delta = (D, W)$ be a default theory and let E be a set of formulas.*
We have that E is an extension of Δ iff

$$E = Th\big(W \cup Cons\big(D'\big) \cup E'\big)$$

for some $D' \subseteq D$ such that

1. *D' is grounded in W,*
2. *D' is protected in Δ, and*
3. *$\Delta|D'$ has extension E'.*

The utility of this result stems from its decomposition of the definition of an extension into the formation of grounded, protected sets D' and a coherence condition on the theory simulating the application of the rules in D'. Notably, in a coherent context, the application of such a set of rules is fully independent of the rest of the theory. Observe also that verifying conditions *1.* and *2.* involves inspecting D' and predecessors of D' in the block graph only (see Theorem 5.9 below). The treatment of the remaining rules is (roughly) mapped onto an existence-of-extension-problem.

Taking D' in the "only-if"-direction of the last theorem as the generating default rules of E yields a non-conflicting default theory, as shown next.

Theorem 5.4 *Let $\Delta = (D, W)$ be a default theory.*
If E is an extension of Δ, then default theory $\Delta|GD(D, E)$ is non-conflicting.

That is, we get an arcless block graph $\Gamma_{\Delta|GD(D,E)} = (D \setminus (GD(D, E) \cup \overline{GD(D, E)}), \emptyset)$.

The "if"-direction of Theorem 5.3 is of great significance, since it furnishes construction principles for extensions, depending on the nature of the underlying default theory.

Theorem 5.5 *Let $\Delta = (D, W)$ be a non-conflicting default theory and let E be a set of formulas.*
We have that E is an extension of Δ iff $E = Th(W \cup Cons(D'))$ for a greatest set $D' \subseteq D$ being grounded in W.

Technically, this result is obtained as a by-product in the proof of Theorem 4.1.

The more interesting case is the conflicting yet coherent one.

Theorem 5.6 *Let $\Delta = (D, W)$ be an even default theory and let E be a set of formulas.*
We have that E is an extension of Δ iff $E = Th(W \cup Cons(D'))$ for some maximal $D' \subseteq D$ being grounded in W and protected in Δ.

That is, once the block graph indicates that a default theory iss even its extensions are induced by maximally grounded and protected sets of default rules. This definition does not only avoid a global consistency check, as needed in traditional ones, but it moreover gets rid of the usual fixed-point condition.

Even more surprisingly, this can also be achieved in the general case.

Theorem 5.7 *Let $\Delta = (D, W)$ be a default theory and let E be a set of formulas.*
We have that E is an extension of Δ iff $E = Th(W \cup Cons(D'))$ for some maximal $D' \subseteq D$ such that D' is grounded in W, D' is protected in Δ and $\Delta|D'$ is non-conflicting.

In contrast to even theories, we need in the general case an additional filter, stipulating that the resulting set D' induces an arcless block graph $\Gamma_{\Delta|D'}$. As a matter of fact, this is needed for covering the entire set of default rules D. While D' is conditioned by multiple constraints, $(D \setminus D')$ is taken care of through $\Gamma_{\Delta|D'}$. Thus, the rules in $(D \setminus D')$ are not necessarily (re)inspected due to the block graph.

We see that both Theorem 5.3 and 5.7 rely on block graph $\Gamma_{\Delta|D'}$. In fact, these block graphs can be obtained from Γ_Δ by arc-deletion only, as shown next.

Theorem 5.8 *Let $\Delta = (D, W)$ be a default theory such that $\Gamma_\Delta = (D, A_\Delta)$ contains no self loops.*
If $D' \subseteq D$ is grounded in W and protected in Δ, then $A_{\Delta|D'} \subseteq A_\Delta$.

Thus new arcs can only appear in $\Gamma_{\Delta|D'}$ in the presence of self loops. [11]

The obvious question is now: *Where have the global consistency check along with its underlying fixed-point construction gone?* The answer is: *They have been compiled away!* A fixed-point construction is usually needed for guessing the resulting extension. During the reconstruction of such an extension, all default rules are then already applied relative to the consistency requirements imposed by the final extension. In this way, it is impossible that the application of a default rule δ is subsequently invalidated by applying another default rule δ', whose consequent contradicts the justification of δ. That is, checking consistency against the pre-guessed extension makes it impossible to apply default rules under wrong consistency assumptions. Now, the block graph makes such kind of guesses obsolete, since it tells us which rules threaten the application of other rules. That is, when considering δ for application, the block graph indicates whether it is threatened by δ', and if this is the case which defaults are candidates for supporting the application of δ (by blocking δ'). Formally, this is accomplished by stipulating *protectedness*. In addition, we must account in the general case for default rules menacing the overall extension. This is addressed by requiring that the theory simulating the application of the generating default rules is non-conflicting, or equivalently, that its block graph has no arcs anymore.

[11] This is due to Condition **BS3** in Definition 3.1, which relies on the elimination of single rules; such a rule may constitute a self loop.

Consider our running example (1). The first extension E_1 is generated by $GD(D, E_1) = \{\delta_w, \delta_b, \delta_{\neg f}\}$. Clearly, $GD(D, E_1)$ is grounded; its protectedness is established by the sufficient conditions in (16) and (17). For maximality, we observe that the addition of δ_f would violate **PTD1**. Finally, we note that $\Delta|GD(D, E_1)$ leaves us with an empty set of rules, giving an empty block graph.

For a complement, consider the theory Δ' obtained by adding $\frac{:x}{\neg x}$ to (1). The set $GD(D, E_1) \cup \{\frac{:x}{\neg x}\}$ violates **PTD1**, so that we consider once more $GD(D, E_1)$, whose grounded- and protectedness are established as above. Now, however, $\Delta'|GD(D, E_1)$ yields a block graph, whose single arc is a loop at node $\frac{:x}{\neg x}$. Since all further subsets of $GD(D, E_1)$ bear even richer block graphs, Δ' has no extension.

5.3 Restricted semi-monotonicity

An important question is which rules in D must actually be inspected for deciding whether D' is a subset of the set of generating default rules of some extension. Recall that this question has a trivial answer in semi-monotonic default logics: *It is just D' and no other rules.* In Reiter's default logic, the answer can be read off the block graph: *It is D' along with its reachable predecessors.*

By letting $\gamma^\star(D')$ denote the set of all reachable predecessors[12] of rules from D' in Γ_Δ, we obtain the following result.

Theorem 5.9 (Restricted Semi-monotonicity) *Let $\Delta = (D, W)$ be a default theory and let $D' \subseteq D$ be a set of defaults.*

If $(\gamma^\star(D'), W)$ has an extension $E^\star$ and $\Delta|GD(\gamma^\star(D'), E^\star)$ is coherent, then Δ has an extension E with $E^\star \subseteq E$.

This result makes precise the block graph's role for limiting the search space by delineating the set of default rules that must be inspected for validating the application of a set of default rules.

As mentioned in the introductory section, semi-monotonicity was already isolated by Reiter in [40], where he showed that it is only satisfied by normal default theories in his default logic. This has led in the following years to the development of various alternative default logics, all of which enjoy semi-monotonicity in full generality. The result given in Theorem 5.9 is — to the best of our knowledge — the first result on semi-monotonicity capturing non-normal default theories in Reiter's default logic.

5.4 Related approaches for characterizing extensions

As mentioned above, the first non-fixed-point characterization of extensions was given in [20]. A rough syntactic characterization of Etherington's semantics amounts to constructing maximal sequences $\langle\delta_i\rangle_{i\in I}$ of default rules that are grounded and that satisfy

$$W \cup Cons(\{\delta_0, \ldots, \delta_{i-1}\}) \not\models \neg Just(\delta_i) \qquad \text{for all } i \in I. \tag{20}$$

Such a sequence is called *stable* if it is weakly regular. Stable sequences correspond to generating default rules of extensions and vice versa [20]. Condition (20) gives an approximation of weak regularity, which is then verified a posteriori.

Another interesting characterization of extension is given in [43]: E is an extension of (D, W) iff there is a grounded subset $D' \subseteq D$ such that $E = Th(W \cup Cons(D'))$ and for all $\delta \in D$ we have

1. If $\delta \in D'$ then $Pre(\delta) \in E$ and $\neg Just(\delta) \notin E$,
2. If $\delta \notin D'$ then $Pre(\delta) \notin E$ or $\neg Just(\delta) \in E$.

Unlike above, this characterization makes explicit reference to the rules in $D \setminus D'$. This reference is usually dealt with implicitly by appropriate maximality conditions. In fact, any valid D' is a *maximal* grounded set satisfying Condition 1.

[12] See Section A for a formal definition.

Let us now return to the argumentation frameworks discussed in Section 3.5. Dung gives in [19] a correspondence between extensions and his *stable sets*, defined as: $S = \{A \mid A \text{ is not attacked by } S\}$. This concept is refined in assumption-based frameworks [5], where *stable sets* are defined as sets of assumptions S satisfying the following conditions:

1. S does not attack itself, and
2. S attacks each assumption $\phi \notin S$.

Equivalently, given an assumption based framework $(T, A, \overline{})$ for some (D, W), a set S is stable iff

$$S = \{M\beta \mid \beta \in \mathit{Just}(D) \text{ and } W \cup S \not\vdash_D \neg\beta\} = \{A \mid A \text{ is not attacked by } S\}\ . \tag{21}$$

In [5], extensions are then put in correspondence with sets of form $\mathit{Th}(W \cup S)$. This gives a characterization of extensions in terms of sets of justifications, as opposed to sets of generating default rules. Condition 1. enforces weak regularity. Similar to Risch's second condition, Condition 2. stipulates that any rule (or justification) not contributing to the stable set can not be applied. As above, this induces maximal sets satisfying Condition 1.

Our family of characterizations differs in several respects from those listed above. First, our basic characterization in Theorem 5.3 does not rely on maximal sets of rules; it is applicable to arbitrary subsets D' of D. Second, it deals with rules in $D \setminus D'$ by appeal to coherence, which can be addressed in several ways. For instance, by using the block graph, it allows us to avoid the common fixed-point characterization of extensions (c.f. Theorem 5.7). Third, we use supporting sets for protecting D' instead of (meta)conditions forbidding blocking sets outside of D'.

Among the computationally motivated characterizations of extensions, Niemelä describes in [38] sophisticated conflict-resolution techniques for an extension-construction-procedure. Interestingly, as assumption-based frameworks, it relies on characterizing extensions by sets of justifications, called *full sets*. Full sets contain those justifications that are consistent with the set, obtained by closing the initial set of facts under classical inferences and those default rules (used as monotonic inference rules) whose justifications belong to the full set. By definition, this is equivalent to the notion of stable sets given in (21). Computationally, full sets are determined by techniques borrowed from the Davis-Putnam-Procedure [14]. In contrast to this, Marek et al. advocate in [11] stratification techniques as the primary tool of their extension-construction-procedure.

Finally, let us briefly return to argumentation frameworks in order to investigate the relationship between so-called *admissible* assumption sets and protected sets. A set of assumptions S is *admissible* if

1. S does not attack itself, and
2. for all sets of assumptions S', if S' attacks some $\phi \in S$, then S attacks some $\phi' \in S'$.

Maximal admissible sets S are used in [5] for defining extensions of form $\mathit{Th}(W \cup S)$ that differ from Reiter's extensions. Consider theory $\left(\left\{\frac{:\neg a}{a}, \frac{a:b}{c}, \frac{:\neg c}{\neg c}\right\}, \emptyset\right)$. While the only admissible set is $\emptyset$, we get an additional protected set, namely $\left\{\frac{:\neg c}{\neg c}\right\}$. The concept of extensions defined by maximal protected sets is elaborated upon in [30].

6 Characterizations of default proofs

This section addresses the extension-membership-problem. This problem has actually three dimensions: First, the decision-oriented one that is merely concerned with the abstract question whether a default theory has an extension containing a given formula. Second, the proof-oriented one that aims at providing an adequate notion of a default proof. And finally the algorithmic dimension that deals with query-answering procedures whose aim is then to find the aforementioned default proofs. In what follows, our emphasis lies on the characterization of default proofs rather than the algorithmic aspects dealing with the search for these proofs.

6.1 Formal foundations for default proofs

The discussion in the last section was dominated by the consideration of maximal sets of default rules D' in Theorem 5.3. It is however equally important to notice that D' needs not to be maximal. This is because classical logic allows us to sanction viable parts of an extension E, given by $Th(W \cup Cons(D'))$, without constructing the remaining part, namely E'. In fact, taking D' as a grounded subset of some generating default rules attributes to it the character of an *extension-dependent default proof*: Given the set of generating default rules $GD(D,E)$ for an extension E, a default proof of some formula φ is simply a grounded set $D' \subseteq GD(D,E)$ such that $W \cup Cons(D') \models \varphi$. In the context furnished by E, we do neither have to care about the consistent application of the rules in D' (this is assured by $D' \subseteq GD(D,E)$) nor (trivially) about the existence of an encompassing extension. Both issues, addressed by conditions *2.* and *3.* in Theorem 5.3, are however of crucial importance, whenever no such extension is provided.

For example, in Default theory (1), the set $D' = \{\delta_b, \delta_f\}$ may serve as a default proof for f. It satisfies conditions *1.* and *2.* in Theorem 5.3 because it is grounded and protected wrt (1). For showing that f belongs to an existing extension, we must show that $\Delta|D'$ is coherent. This can be accomplished without computing such an extension. We get a non-conflicting theory

$$\Delta|D' \;=\; (D \setminus (\{\delta_b, \delta_f\} \cup \{\delta_{\neg f}\}), W \cup \{f, b\}) \;=\; (\{\delta_w\}, \{p, ab_p, f, b\})$$

which must have a single extension due to its arcless block graph. Hence, we have shown that f is a default conclusion of (1) without computing the corresponding extension.

The following corollary to Theorem 5.3 makes the previous ideas precise.

Corollary 6.1 *Let $\Delta = (D, W)$ be a default theory and let φ be a formula.*
We have that $\varphi \in E$ for some extension E of Δ iff

$$W \cup Cons\big(D'\big) \models \varphi$$

for some $D' \subseteq D$ such that

1. *D' is grounded in W,*
2. *D' is protected in Δ, and*
3. *$\Delta|D'$ is coherent.*

The formation of default proofs thus boils down to finding a grounded and protected set of default rules that allows for deriving a query, provided that it is applicable within an existing extension.

Whenever Condition *3.* can be addressed by one of the criteria given in Section 4, Corollary 6.1 represents a characterization of default proofs expressed entirely in terms of blocking and supporting sets. In case $\Delta|D'$ agrees with the syntactic formats stipulated in either of [21, 39, 10], these approaches work just as fine. The test is trivial if $\Delta|D'$ is normal or non-conflicting.

6.2 Default proofs from non-conflicting default theories

For conceptual clarity, we start with default proofs for the simple case of non-conflicting default theories.

Definition 6.1 (Pure default proof[13]) *Let $\Delta = (D, W)$ be a default theory and φ a formula.*
A set of default rules $P \subseteq D$ is a pure default proof for φ from Δ iff

P1 $W \cup Cons(P) \models \varphi$,

P2 *P is grounded in W.*

[13] Observe that although pure default proofs ignore justifications, they are still non-monotonic because they may be invalidated after augmenting the underlying non-conflicting theory.

We see that basic blocking sets are pure default proofs for negated justifications. Observe that the simple nature of non-conflicting theories makes consistency checks obsolete. For example, let Δ'' be the default theory obtained from (1) by leaving out $\frac{b:\neg ab_b}{f}$. Then, Δ'' is non-conflicting and $P = \{\frac{p:\neg ab_p}{\neg f}\}$ is a pure default proof for $\neg f$. This proof can be found without consistency checking nor any measures guaranteeing the existence of an encompassing extension.

We observe the following result for non-conflicting default theories.

Theorem 6.2 *Let Δ be a non-conflicting default theory and φ a formula.*
We have that $\varphi \in E$ for an extension E of Δ iff there is a pure default proof for φ from Δ.

6.3 Default proofs from conflicting yet coherent default theories

Whenever we have a conflicting yet coherent default theory, we have to take supporting sets into account because then it is necessary to protect the constituent default rules of a default proof. This leads us to the concept of *protected default proofs.*

Definition 6.2 (Protected default proof) *Let $\Delta = (D, W)$ be a default theory and φ a formula.*
A set of default rules $P \subseteq D$ is a protected default proof for φ from Δ iff

PP1 *P is a pure default proof for φ from Δ,*

PP2 *P is protected in Δ.*

Note that this characterization substitutes global consistency checks by the determination of supporting sets. In fact, any protected default proof consists of rules needed for deriving φ and supplementary rules needed for protecting the derivation against blocking sets. Hence, we have to make sure that there is some supporting set for each default in the proof, as expressed in the following definition.

The next result shows that it is sufficient to inspect the set $\gamma^\star(P)$ of all reachable predecessors (in the block graph) of a pure default proof P when checking its protectedness.

Theorem 6.3 *Let $\Delta = (D, W)$ be a default theory and let $P \subseteq D$ be a pure default proof for φ from Δ.*
If default theory $(\gamma^\star(P), W)$ has extension E with $\varphi \in E$, then there is a protected default proof for φ from (D, W).

Observe that without any restriction on the theory, the existence of a protected default proof for φ does not guarantee an extension containing φ. That is, soundness and completeness necessitate coherent theories (c.f. Theorem 6.4).

For illustration, let us return to the even default theory in (1). We have seen in Section 6.2 that

$$P = \{\delta_{\neg f}\} = \left\{\frac{p:\neg ab_p}{\neg f}\right\} \tag{22}$$

is a pure default proof for $\neg f$, that is, **PP1** holds. For verifying **PP2**, we must address **PTD1** and **PTD2**. For **PTD1**, it is sufficient to observe that the members of P are not connected in the block graph. The fact that $\mathcal{S}_\Delta(\delta_{\neg f}) = \{\{\delta_{\neg f}\}\}$ establishes **PTD2**. Hence P is a protected default proof for $\neg f$ from (1). Provided that one uses approximation condition (17) for establishing **PTD2**, this proof is found without any consistency checks and no measures guaranteeing the existence of an encompassing extension.

Similar arguments show that

$$P' = \left\{\frac{p:b}{b}, \frac{b:\neg ab_b}{f}\right\} \qquad \text{and} \qquad P'' = \left\{\frac{p:b}{b}, \frac{b:w}{w}\right\} \tag{23}$$

are protected default proofs for f and w from Theory (1), respectively. Note that we only have to warrant a supporting set for $\frac{b:\neg ab_b}{f}$ since it is the only default rule in $P' \cup P''$ having a predecessor in the block graph. As with $\frac{p:\neg ab_p}{\neg f}$ above, however, $\frac{b:\neg ab_b}{f}$ supports itself, so that no other defaults have to be taken into account. The second default proof is formed without any supporting sets.

As a result, we get that protected default proofs furnish a sound and complete concept for addressing the extension-membership-problem on coherent, or in our case even default theories.

Theorem 6.4 *Let $\Delta = (D, W)$ be a even default theory and φ a formula.*
We have that $\varphi \in E$ for an extension E of Δ iff there is a protected default proof for φ from Δ.

The stipulation of evenness stems from the fact that it relies, as protectedness, on the underlying block graph. Thus the interplay of both concepts can be characterized in a direct way. Therefore, evenness should not be seen as a restriction on coherent theories. There should be more general classes of coherent default theories, guaranteeing correctness and completeness of protected default proofs that remain to be discovered.

As a corollary to Theorem 6.6, we obtain that deciding whether there is a protected default proof for formula φ from a default theory is in Σ_2^P.

6.4 Default proofs from general default theories

When dealing with arbitrary default theories, we must guarantee that a default proof resides in an encompassing extension. Clearly, this should be done without computing such an extension. Given a protected default proof P from a theory Δ, this can be accomplished by checking whether $\Delta|P$ is coherent.

Definition 6.3 (General default proof) *Let $\Delta = (D, W)$ be a default theory and φ a formula.*
A general default proof for φ from Δ is a set of default rules $P \subseteq D$ such that

DP1 *P is a protected default proof for φ from Δ and*

DP2 *$\Delta|P$ is coherent.*

For illustration, let us add self-circular default rule $\zeta_{\neg f} = \frac{:f}{\neg f}$ to the theory in (1). The resulting theory Δ' has a single extension: $Th(W \cup \{b, w, \neg f\})$. The block graph of Δ' is given in Figure 2. It is

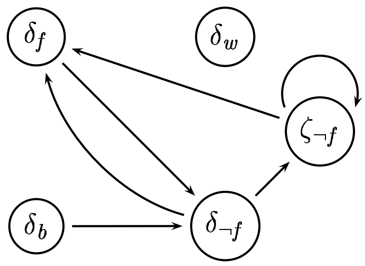

Figure 2: Block graph of Δ'.

obtained from the one in Figure 1 by adding arcs $(\zeta_{\neg f}, \zeta_{\neg f})$, $(\delta_{\neg f}, \zeta_{\neg f})$, and $(\zeta_{\neg f}, \delta_f)$. In what follows, we reconsider the previous protected default proofs in the light of this change.

We start with the protected default proof given in (22), viz.

$$P = \{\delta_{\neg f}\} = \left\{ \frac{p:\neg ab_p}{\neg f} \right\} .$$

For establishing **DP1**, we observe that the set of predecessors of $\delta_{\neg f}$ remains unchanged. For verifying **DP2**, we consider the block graph of $\Delta'|P$, which is obtained from $\Gamma_{\Delta'}$ by deleting arcs and vertices. First, we delete in $\Gamma_{\Delta'}$ all defaults in $P = \{\delta_{\neg f}\}$ along with all adjacent arcs. The same is done with δ_f and $\zeta_{\neg f}$ because $W \cup Cons(\{\delta_{\neg f}\}) \models \neg Just(\delta)$ for $\delta \in \{\delta_f, \zeta_{\neg f}\}$ (c.f. Definition 5.2). As a result, we obtain a graph with vertex set $\{\delta_w, \delta_b\}$ and no arcs. This implies that $\Delta'|P$ is coherent and hence that P is a general default proof for $\neg f$ from Δ'.

Next, consider the first protected proof in (23), viz.

$$P' = \{\delta_b, \delta_f\} = \left\{ \frac{p : b}{b}, \frac{b : \neg ab_b}{f} \right\} . \tag{24}$$

Unlike above, we encounter an additional predecessor of P' in the augmented block graph, namely $\zeta_{\neg f}$. The predecessors of $\zeta_{\neg f}$ yield two candidates for forming a supporting set: $\delta_{\neg f}$ and $\zeta_{\neg f}$. We see that their common consequent $\neg f$ is inconsistent with those in P'. In other words, augmenting P' by either of them would violate **PP2**, that is, **PTD1**. Hence there is no way to form a supporting set for δ_f and so P' is no protected default proof from Δ'. This is reflected by the fact that Δ' has no extension containing f.

Finally, let us consider the second protected proof in (23), viz.

$$P'' = \{\delta_b, \delta_w\} = \left\{ \frac{p : b}{b}, \frac{b : w}{w} \right\} .$$

Although the absence of predecessors of P'' in $\Gamma_{\Delta'|P''}$ establishes **DP1**, the rules in P'' are insufficient to guarantee an encompassing extension since there is an odd cycle in $\Gamma_{\Delta'|P''}$, as shown in Figure 3. In contrast to the proof in (24), we may strengthen P'' in order to establish **DP2**. To see this, observe that there is an arc entering the odd cycle, namely $(\delta_{\neg f}, \zeta_{\neg f})$. This indicates that $\delta_{\neg f}$ may block $\zeta_{\neg f}$. Hence, consider $P'' \cup \{\delta_{\neg f}\}$. The fact that $\{\delta_b, \delta_w\} \cup \{\delta_{\neg f}\}$ satisfies **DP1** is established as shown above. For **DP2**, we must inspect the block graph of $\Delta'|(P'' \cup \{\delta_{\neg f}\})$, which turns out to be the empty graph. In all, $P'' \cup \{\delta_{\neg f}\}$ is thus a valid default proof for *w* from Δ'.

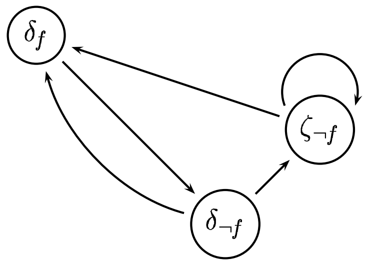

Figure 3: Block graph of $\Delta'|P''$.

We note how the block graph indicated a way to eliminate the odd cycle in the resulting block graph. Condition **DP2** may thus be verified by first eliminating the defaults involved in the default proof from the block graph and then by checking whether the resulting block graph is even, acyclic, or arcless. In order to apply these criteria, however, we might have to add additional rules eliminating dangerous odd cycles. These additional rules can be identified by means of the results in Section 4.3.

We have the following result, establishing soundness and completeness in the general case.

Theorem 6.5 *Let $\Delta = (D, W)$ be a (general) default theory and φ a formula.*
We have that $\varphi \in E$ for an extension E of Δ iff there is a general default proof for φ from Δ.

We have seen that the formation of default proofs benefits considerably from the usage of block graphs. That is, we may restrict our attention to ultimately necessary default rules and so avoid constructing entire extensions.

Finally, it follows from Theorem 6.5 that the existence of general default proofs is decidable within the complexity class of the extension-membership-problem, as made precise in the following theorem.

Theorem 6.6 *Let $\Delta = (D, W)$ be a default theory and φ a formula.*
The problem of deciding whether there is a general default proof for φ from Δ is in Σ_2^P.

7 Conclusion

We have introduced new theoretical foundations for default logic. The fundamental idea was to avoid global concepts like fixed-point conditions and exhaustive consistency checks by replacing them by rather local concepts that allow for considering only strictly necessary default information. This was accomplished by putting forward the notion of blockage. Our formal account of blockage is given by the concept of an (essential) blocking set; this can be regarded as a specific, redundancy-free instance of attacking arguments, used in the abstract frameworks of argumentation semantics [19, 5]. In contrast to these unifying frameworks, however, our elaboration is specific to default logic. This has led to much stronger and many additional results than obtained at the abstract level in [19, 5].

In our specific setting, blocking sets represent context-independent proof skeletons that may be used for refuting a default rule's justification. That is, for a default rule there may be some extension containing such a blocking set that inhibits the application of the rule. Thus, given only a rule and one of its blocking sets, we cannot decide whether the rule applies without any information about the extension at hand. The situation is different with supporting sets. Clearly, supporting sets are also context-independent. But unlike blocking sets, they are supposed to apply in the *same* extension as their supported rule. In fact, the joint application of a rule and one of its supporting sets can only be denied by a self-blocking part of the theory that destroys an encompassing extension. Thus, given a rule and one of its supporting sets, we can decide whether the rule (and its supporting set) applies whenever we can rule out sources of incoherence. From a general complexity-theoretic point of view, solving a coherence problem is as hard as constructing an encompassing extension. In our case, however, the block graph makes the difference since it indicates sources of incoherence.

The block graph furnishes our instrument for abstracting from particular blocking situations, while keeping the essential interaction patterns between default rules. In view of the possibly exponential number of blocking sets, it is thus a trade-off between full blockage information and feasible space complexity. In fact, for improving the quality of the resulting block graphs, we tried to eliminate as many redundant blocking sets as possible. This has led us to the concept of essential blocking sets that are provably as effective as basic blocking sets.

In fact, the block graph tells us exactly which default rules must be considered for applying a rule in question. This can be made more precise by returning to the classification of default theories, given in the introductory section. Consider the applicability of a set of rules P:

1. The fact that a theory contains no interactions corresponds to an arcless block graph. For establishing the applicability of P, it is thus sufficient to verify groundedness of P; no other rules must be inspected.

2. The fact that a theory contains merely direct interactions is mirrored by an even (or acyclic) block graph. For establishing the applicability of P, it is sufficient to verify groundedness and protectedness of P. This necessitates the inspection of $\gamma^\star(P)$, namely P and all reachable predecessors of rules in P in the block graph.

3. The fact that a theory comprises both direct and indirect interactions amounts in our framework to a non-even block graph. For establishing the applicability of P, we are obliged to verify in addition to groundedness and protectedness of P, also the coherent interplay of P with the rest of the theory. This necessities the inspection of all default rules D: While $\gamma^\star(P)$ is considered for establishing groundedness and protectedness of P, the coherent interplay of P and $(D \setminus P)$ is taken care of through the block graph of $\Delta|P$.

The block graph's role for limiting the search space by delineating the scope of default rules that must be inspected for validating the application of a set of default rules is made precise in a restricted semi-monotonicity result. This result is to the best of our knowledge the first one capturing semi-monotonicity beyond normal default theories in Reiter's default logic.

Practically, the block graph is obtained from an initial analysis of the default theory. Notably it is of quadratic space complexity, unlike other approaches like [26, 47] that face an exponential blow-up in the worst case. The computational complexity of the block graph amounts in the worst case to that of the extension-membership-problem. The initial effort put into the block graph pays off the more, the sparser the block graph.

In all, the block graph (along with its underlying blocking sets) provide a powerful structural tool for analyzing default theories, as demonstrated by a variety of applications:

Existence-of-extension-problem. We address this problem by furnishing a range of criteria guaranteeing the existence of extensions, each of which can be read off the block graph in polynomial time. We show that these criteria are simpler and go beyond existing approaches. Our criteria are fully syntax-independent and allow for treating general default theories.

For a complement, we give also a criterion indicating non-existence of extensions. Although all these criteria provide no complete characterization of default theories with and without extensions, they furnish nonetheless easy, block graph based tests that allow us to shorten the gap between both classes of theories.

Notably, we can decide the existence-of-extension-problem, when considering blocking sets. Although this faces exponential space complexity in practice, this result is to the best of our knowledge the first complete characterization for coherent default theories.

Characterizations of extensions. By appeal to the block graph, we obtain novel characterizations of extensions that do not only avoid global consistency checks, as needed in traditional characterizations. Moreover, the usual fixed-point condition is fully eliminated.

These characterizations are not only of theoretical importance, but they are also of practical relevance, since they lead to a basic procedure for constructing extensions. In fact, we show in [31] such a construction can be decomposed into an incremental construction of a Łukaszewicz-extension [34] and a block graph based condition.

Extension-membership-problem. We provide a series of characterizations of default proofs along the classification of default theories described above. In each case, we give a soundness and completeness result.

To the best of our knowledge, these characterizations furnish the first definitions of default proofs in Reiter's default logic that do not refer to an outer extension. This is reflected by the fact that up to now all computational approaches to the extension-membership-problem are extension-oriented [26, 47, 38, 12]. In contrast to this, our characterizations of default proofs avoid exhaustive consistency checks with respect to an outer extension by focusing on the strictly necessary default rules by means of the block graph.

Apart from these applications, our blockage-based concepts allow for addressing various other issues. First, the block graph furnishes sufficient conditions for restricting or even omitting consistency-checks. Second, as detailed in [32], it provides means for supporting skeptical and modular reasoning. And finally its resulting concepts, like protectedness, allow us to make the relationship to Łukaszewicz's interpretation of default theories much more precise, as shown in [30].

Given this fundamental framework, one may now divide a computational problem like query-answering in an off-line and an on-line process: One may start with an analytic compilation phase resulting in the block graph Γ_Δ of a default theory Δ. The subsequent query-answering phase aims at finding a default proof P such that $\Delta|P$ possesses an extension. The unavoidable examination of the entire set of default rules is then done only once in the compilation phase; this allows for inspecting only the ultimately necessary default rules during the actual query answering phase. To be more precise, a default rule belongs to P only (i) if it contributes to the derivation of the query or if it is needed (ii) for supporting a constituent rule of the proof or (iii) for supporting an encompassing extension. While (i) is fixed by the standard inferential relation, (ii) and (iii) are determined by the block graph. For delineating such a proof, we can draw on Γ_Δ for detecting and eventually recomputing the blocking and supporting sets of its constituent rules. Blocking sets are found among the direct predecessors of a rule, while the search for its supporting sets can be restricted to its pre-predecessors. Analogous decompositions of other computational problems like extension-construction are straightforward and should arguably be mutually beneficial for existing computational approaches. Of particular interest are also computational approaches to argumentation, e.g. [27].

To sum up, the salient contribution of our paper was to shift global, extension-based concepts in default logic towards local, proof-oriented ones. For instance, we have shown how to replace global consistency checks by rather local proof-based constructions that are guided by the underlying block graph. This provides a formal account of "local constructibility" that was up to now always associated with semi-monotonic default theories.

The next major steps on this research avenue are manifold. One such avenue will deal with algorithmic and implementation-oriented issues that exploit the theoretical framework proposed in this paper. Another one has to address more fine-grained complexity issues. For instance, is it possible to characterize restricted classes of theories having particular block graphs that lead to reduced complexity? Dually, one may also consider well-known fragments, like logic programming, and investigate their particular blockage structure.

Acknowledgements

We would like to thank Philippe Besnard, Gerhard Brewka, Robert Mercer, Ilkka Niemelä, Hans Tompits, and Mirek Truszczyński for commenting on previous versions of this paper.

This work was partially supported by the German Science Foundation (DFG) within Project "Nichtmonotone Inferenzsysteme zur Verarbeitung konfligierender Regeln".

A Graph-theoretical background

Definition A.1 *A directed graph G is a pair $G = (V, A)$ such that V is a finite, non-empty set of vertices and $A \subseteq V \times V$ is a set of arcs. If G is a directed graph $G^{-1} = (V, A')$ denotes the directed graph where the orientation for all arcs in G is switched, formally, $(u, v) \in A'$ iff $(v, u) \in A$.*

A *directed cycle* in $G = (V, A)$ is a finite subset $C \subseteq V$ such that $C = \{v_1, \cdots, v_n\}$ and $(v_i, v_{i+1}) \in A$ for each $1 \leq i < n$ and $(v_n, v_1) \in A$. The arcs $A(C)$ of a cycle C are defined as $A(C) = \{(v_i, v_{i+1}) \mid 1 \leq i < n\} \cup \{(v_n, v_1)\}$. The *length* of a directed cycle in a graph is the total number of arcs occurring in the cycle. Additionally, we call a cycle *even* (*odd*) if its length is even (odd).

Definition A.2 *Let $G = (V, A)$ be a directed graph and $U \subseteq V$ a subset of vertices. U is independent wrt G iff for all $u, v \in U$ we have $(u, v) \notin A$.*

Thus a subset $U \subseteq V$ is independent wrt a graph (V, A) if there is no arc between nodes in U.

Definition A.3 *Let $G = (V, A)$ be a directed graph and $K \subseteq V$ an independent subset of vertices. K is a kernel of G iff for all nodes $u \in V \setminus K$ exists a $v \in K$ such that $(u, v) \in A$.*

A kernel K is an independent set of vertices such that for every vertex not in K there is an arc to some vertex in K.

Theorem A.1 **[42]** *Let $G = (V, A)$ be a directed graph without odd cycles, then G has a kernel.*

For a directed graph $G = (V, A)$ and a vertex $v \in V$, we define the set of all *predecessors* of v as $\gamma^{-1}(v) = \{u \mid (u, v) \in A\}$. All *reachable predecessors* $\gamma^{\star}(v)$ of v are defined as follows:

$$\gamma^{\star}(v) = \bigcup_{i \in I} \gamma^{-i}(v)$$

where $\gamma^{-0}(v) = \{v\}$ and for all $i > 1$ we have $\gamma^{-i}(v) = \{u \mid (u, w) \in A \text{ and } w \in \gamma^{-(i-1)}(v)\}$. For a set of vertices $U \subseteq V$, we define

$$\gamma^{\star}(U) = \bigcup_{v \in U} \gamma^{\star}(v).$$

B Proofs of results

B.1 Auxiliary technical results

First, we recall the following specification of extensions [40].

Theorem B.1 *Let (D, W) be a default theory and let E be a set of formulas.*
Define $E_0 = W$ and for $i \geq 0$

$$E_{i+1} = \mathit{Th}(E_i) \cup \left\{ \gamma \;\middle|\; \frac{\alpha : \beta}{\gamma} \in D, \alpha \in E_i, \neg\beta \notin E \right\}.$$

Then, E is an extension of (D, W) iff $E = \bigcup_{i \geq 0} E_i$.

The above procedure is not strictly iterative since E appears in the specification of E_{i+1}.

We need the following lemmas and theorems when proving the results of this paper. All proofs for the results in this subsection are straight forward and can be found in [33]. The next lemma gives a further property of grounded sets of default rules, namely that the union of grounded sets is also grounded under certain conditions.

Lemma B.2 *Let D_1 and D_2 be enumerable sets of default rules, $D = D_1 \cup D_2$ and W be a set of formulas. If D_1 is grounded in W and D_2 is grounded in $W \cup \mathit{Cons}(D_1)$, then D is grounded in W.*

Now we are ready to formulate an alternative characterization of the extensions of a default theory which serves as a theoretical basis for proving the further results.

Theorem B.3 *Let (D, W) be a default theory and let E be a set of formulas. Then, E is an extension of (D, W) iff*

1. *$E = \mathit{Th}(W \cup \mathit{Cons}(\mathit{GD}(D, E)))$ and*
2. *$\mathit{GD}(D, E)$ is grounded in W.*

Lemma B.4 *Let $\Delta = (D, W)$ be a default theory and $D_1 \subseteq D$ a subset of default rules such that default theory (D_1, W) has extension E_1 and default theory (D, E_1) has extension E_2. If $\mathit{GD}(D_1, E_1) \subseteq \mathit{GD}(D, E_2)$, then E_2 is an extension of Δ.*

Lemma B.5 *Let $\Delta = (D, W)$ be a default theory. Then Γ_Δ has an odd cycle iff Ω_Δ has an odd cycle.*

Lemma B.6 *Let $\Delta = (D, W)$ be a default theory, $\delta \in D$ a default rule and C a subset of D. If C is a complete support of δ then C is grounded in W.*

B.2 Proofs of results occurring in the text

Proof 3.1 Let $B \subseteq D$ and $\delta \in D$. If $B = \emptyset$ the proposition trivially holds. Thus let $B \neq \emptyset$ for the rest of this proof.

By Definition 3.1 $B \in \mathcal{B}_\Delta(\delta)$ is equivalent to $B \mapsto_\Delta \delta$ and **BS3**. That is, $B \mapsto_\Delta \delta$ and

$$(B \setminus \{\delta'\}) \not\mapsto_\Delta \delta'' \qquad \text{for each } \delta' \in B \text{ and each } \delta'' \in B \cup \{\delta\}.$$

This is equivalent to $B \mapsto_\Delta \delta$ and

$$\begin{array}{ll} (B \setminus \{\delta'\}) \not\mapsto_\Delta \delta & \text{for each } \delta' \in B \text{ and} \\ (B \setminus \{\delta'\}) \not\mapsto_\Delta \delta'' & \text{for each } \delta' \in B \text{ and each } \delta'' \in B. \end{array}$$

According to Definition 3.1, this is equivalent to $B \mapsto_\Delta \delta$ and

$$\begin{array}{ll} B' \not\mapsto_\Delta \delta & \text{for each } B' \subset B \text{ and} \\ B' \not\mapsto_\Delta \delta'' & \text{for each } B' \subset B \text{ and each } \delta'' \in B. \end{array}$$

Finally, this is equivalent to

$$\begin{array}{ll} B \mapsto_\Delta \delta & \text{and} \\ B' \mapsto_\Delta \delta & \text{for no } B' \subset B \text{ and} \\ B' \mapsto_\Delta \delta'' & \text{for no } B' \subset B \text{ and no } \delta'' \in B. \end{array}$$

■

Proof 3.2 Let E be an extension of default theory $\Delta = (D, W)$ and $\delta \in D$. According to Theorem B.3 we have $E = Th(W \cup Cons(GD(D,E)))$ and $GD(D,E)$ is grounded in W.

First we prove that *1.* and *3.* are equivalent. If $E \models \neg Just(\delta)$ then $W \cup Cons(GD(D,E)) \models \neg Just(\delta)$. Then there is a minimal set $B \subseteq GD(D,E)$ such that $B \mapsto_\Delta \delta$. That is, $B' \mapsto_\Delta \delta$ for no $B' \subset B$.

Now let $\delta', \delta'' \in B$. If $B \setminus \{\delta'\}$ is not grounded in W then by Definition 3.1 we have $B \setminus \{\delta'\} \not\mapsto_\Delta \delta''$. Now let $B \setminus \{\delta'\}$ be grounded in W. Because $B \subseteq GD(D,E)$, it follows $E \not\models \neg Just(\delta'')$ for each $\delta'' \in B$. By monotonicity of $\models$ we obtain that $W \cup Cons(B \setminus \{\delta'\}) \not\models \neg Just(\delta'')$ for each $\delta'' \in B$. Hence $B \setminus \{\delta'\} \mapsto_\Delta \delta''$ for no $\delta', \delta'' \in B$, which implies $B' \mapsto_\Delta \delta''$ for no $B' \subset B$ and no $\delta'' \in B$. By Theorem 3.1 we conclude $B \in \mathcal{B}_\Delta(\delta)$. The other direction trivialy holds.

Next, we prove that *1.* and *2.* are equivalent. If E is an extension such that $E \models \neg Just(\delta)$, then $W \cup Cons(GD(D,E)) \models \neg Just(\delta)$. Hence there is a minimal subset $B' \subseteq GD(D,E)$ such that $B' \mapsto_\Delta \delta$. According to Definition 3.1 the other direction also holds.

Trivially, *2.* and *4.* are equivalent by Definition 3.1. ■

Proof 3.3 Let $D' \subseteq D$ be a grounded set of default rules.

If D' is weakly regular, then for each $\delta \in D'$ we have $W \cup Cons(D') \not\models \neg Just(\delta)$. Assume that there is a $\delta \in D'$ and there is a $B \in \mathcal{B}_\Delta(\delta)$ such that $B \subseteq D'$. Then by Definitions 3.1 and 3.1 we have $W \cup Cons(B) \models \neg Just(\delta)$. Because $B \subseteq D'$ it follows $W \cup Cons(D') \models \neg Just(\delta)$ which is a contradiction to the fact that D' is weakly regular. Hence for each $\delta \in D'$ and each $B \in \mathcal{B}_\Delta(\delta)$ we have $B \not\subseteq D'$.

For the other direction let $B \notin \mathcal{B}_\Delta(\delta)$ for each $\delta \in D'$ and for each $B \subseteq D'$. By definition, this implies $B \not\mapsto_\Delta \delta$ and thus B is not grounded in W or $W \cup Cons(B) \not\models \neg Just(\delta)$ for each $\delta \in D'$ and for each $B \subseteq D'$. Assuming that D' would not be weakly regular means that there is some $\delta \in D'$ s.t. $W \cup Cons(D') \models \neg Just(\delta)$. It follows that there is a minimal subset $B \subseteq D'$ s.t. B is grounded in W (since D' is grounded in W) and $W \cup Cons(B) \models \neg Just(\delta)$. Since this contradicts the fact that according to our premise we have B is not grounded in W or $W \cup Cons(B) \not\models \neg Just(\delta)$ for each $\delta \in D'$ and for each $B \subseteq D'$ it follows that D' is weakly regular.

■

Proof 3.4 Let $\Delta = (D, W)$ and $\Delta' = (D', W)$ be default theories with $D \subseteq D'$ and $\delta \in D$. Let $B \in \mathcal{B}_\Delta(\delta)$. Then $B \subseteq D'$ and $B \mapsto_\Delta \delta$ and **BS3** by Definition 3.1. Because Conditions $B \not\mapsto_\Delta \delta$ and **BS3** are independent from $D \setminus B$ it follows that they are also independent from $D' \setminus B$ and thus B is a set of default rules fulfilling these conditions wrt Δ'. That is, $B \in \mathcal{B}_{\Delta'}(\delta)$.

Under the premise that $B \subseteq D$ the same argument shows that $B \in \mathcal{B}_{\Delta'}(\delta)$ implies $B \in \mathcal{B}_\Delta(\delta)$. ■

Proof 3.5 For this proof we need the following proposition.

> **Proposition B.7** [22, Propsition 1.1] *Let R be a set of monotonic inference rules of the form $\frac{\varphi_1, \cdots, \varphi_n}{\psi}$. Let T be the modal theory obtained by replacing each rule in R by $L\varphi_1 \wedge \cdots \wedge L\varphi_n \to \psi$. Then for every set of propositional formulas I and every propositional formula ω, we have that ω belongs to the least set of formulas containing I and closed under propositional consequence and the rule from R iff $I \cup T \models_N \omega$.*

For a default rule $\delta = \frac{\alpha : \beta}{\gamma}$ define

$$t_m(\delta) = L\alpha \to \gamma.$$

Observe that $t_m(\delta)$ is a monotonic inference rule. For a set of defaults D' define

$$t_m(D') = \{t_m(\delta) \mid \delta \in D'\}.$$

Let $\Delta = (D, W)$ be a default theory and let $\delta \in D$ and $B \subseteq D$.

According to [45, Theorem 6], we know that reasoning in N and NT is computationally equivalent to reasoning in propositional logic. That is, the problem of deciding if $I \models_{\mathcal{S}} \varphi$ for a given finite set I of modal formulas and a given modal formula φ is co-NP-complete (for $\mathcal{S}$ being modal logics N or NT).

Since in Definition 3.1 defaults are treated as monotonic inference rules (justification are not considered) it follows from Proposition B.7 that $B \mapsto_\Delta \delta$ iff $t_n(B) \cup W \models_{\mathcal{S}} \neg Just(\delta)$. Thus the problem of deciding whether relation $\mapsto_\Delta$holds for a given set of defaults and a given default is reducible to the problem of deciding whether relation $\models_{\mathcal{S}}$ holds for a given set of formulas and a given formula. Therefore the first problem is at least as hard as the second one, which is co-NP-complete. Thus deciding whether $B \mapsto_\Delta \delta$ holds is in co-NP[14]. ■

Proof 3.6 Let $B \subseteq D$ be a subset of default rules and let $\delta \in D$ be a default rule. According to Definition 3.1 we have that B is a blocking set for δ iff

$$B \mapsto_\Delta \delta \text{ and} \tag{25}$$
$$\forall \delta' \in B : \forall \delta^\star \in B \cup \{\delta\} : B \setminus \{\delta'\} \not\mapsto_\Delta \delta^\star. \tag{26}$$

Therefore, to verify that a given set B is a blocking set for a given default δ can be done by $(n(n+1)+1) = n^2 + n + 1$ times testing the relation $\mapsto_\Delta$. According to Theorem 3.5, testing whether for a given set B and a given default δ we have $B \mapsto_\Delta \delta$ is in co-NP. Thus testing whether $B \in \mathcal{B}_\Delta(\delta)$ can be done with a polynomial number of calls to an NP-oracle.

In order to show that there is a blocking set for a given default δ we may guess a subset $B \subseteq D$ nondeterministically and check that $\delta'' \in B$ and that $B \in \mathcal{B}_\Delta(\delta)$. Thus the problem of deciding whether there is a blocking set for δ is nondeterministically Turing-reducible to a co-NP-complete problem (propositional SAT) and hence in Σ_2^P. ■

Proof 3.7 Let $\delta \in D$. $B' \in f(\mathcal{B}_\Delta(\delta))$ iff $B' = f(B)$ for some $B \in \mathcal{B}_\Delta(\delta)$. This is according to Definition 3.1 equivalent to the following conditions:

$W \cup Cons(B) \models \neg Just(\delta)$ and
B is grounded in W and
$\forall \delta' \in B : \forall \delta'' \in B \cup \{\delta\} : B \setminus \{\delta'\} \not\mapsto_\Delta \delta''$.

[14] Since there is a one-to-one correspondence between modal logic N and reasoning with monotonic inference rules and propositional logic the problem is co-NP-complete.

By definition this is equivalent to

$$\begin{array}{ll} W \cup \mathit{Cons}(B) \models \neg \mathit{Just}(\delta) & \text{and} \\ B \text{ is grounded in } W & \text{and} \\ \forall \delta' \in B : \forall \delta'' \in B \cup \{\delta\} : & \\ \quad (W \cup \mathit{Cons}(B \setminus \{\delta'\}) \not\models \neg \mathit{Just}(\delta'') & \text{or} \\ \quad B \setminus \{\delta'\} \text{ is not grounded in } W). & \end{array}$$

Since $\mathit{Th}(W) = \mathit{Th}(W')$ this is equivalent to

$$\begin{array}{ll} W' \cup \mathit{Cons}(B) \models \neg \mathit{Just}(\delta) & \text{and} \\ B \text{ is grounded in } W' & \text{and} \\ \forall \delta' \in B : \forall \delta'' \in B \cup \{\delta\} : & \\ \quad (W' \cup \mathit{Cons}(B \setminus \{\delta'\}) \not\models \neg \mathit{Just}(\delta'') & \text{or} \\ \quad B \setminus \{\delta'\} \text{ is not grounded in } W'). & \end{array}$$

According to Definition 3.2 this is equivalent to

$$\begin{array}{ll} W' \cup \mathit{Cons}(B') \models \neg \mathit{Just}(f(\delta)) & \text{and} \\ B' \text{ is grounded in } W' & \text{and} \\ \forall f(\delta') \in B' : \forall f(\delta'') \in B' \cup \{f(\delta)\} : & \\ \quad (W' \cup \mathit{Cons}(B' \setminus \{f(\delta')\}) \not\models \neg \mathit{Just}(f(\delta'')) & \text{or} \\ \quad B' \setminus \{f(\delta')\} \text{ is not grounded in } W'). & \end{array}$$

By definition this is equivalent to $B' \in \mathcal{B}_{\Delta'}(f(\delta))$. In all for each $\delta \in D$ we get

$$f(\mathcal{B}_{\Delta}(\delta)) = \mathcal{B}_{\Delta'}(f(\delta))\,.$$

■

Proof 3.8 Let $\Delta = (D, W)$ be a default theory s.t. $n = |D|$ and let $G = (D, A)$ be a directed graph with nodes D.

In order to decide whether $G = \Gamma_{\Delta}$ we may guess a set $B_{(\delta,\delta')} \subseteq D$ for each possible arc $(\delta, \delta') \in D \times D$ in G nondeterministically. That is, we guess n^2 subsets of default rules nondeterministically. Next, for each of this subsets $B_{(\delta,\delta')}$ we check that $\delta \in B_{(\delta,\delta')}$ and that $(\delta, \delta') \in A$ and that $B_{(\delta,\delta')} \in \mathcal{B}_{\Delta}(\delta')$. Observe that the first and second of this checks can be done directly (without calling an oracle) and the last check can be done with a polynomial number of calls to an NP-oracle (see proof of Theorem 3.6). Altogether we need a polynomial number of calls to an NP-oracle, because we have to process n^2 subsets $B_{(\delta,\delta')}$ and the multiplication of polynomials leads to polynomials. Thus the problem of deciding whether $G = \Gamma_{\Delta}$ can be done with a polynomial number of calls to an NP-oracle and hence is in Σ_2^P. ■

Proof 3.9 This is an immediate consequence of Theorem 3.4. ■

Proof 3.10 This is an immediate consequence of Theorem 3.7. ■

Proof 3.11 Let $\Delta = (D, W)$ be a default theory, $\delta \in D$ and let E be an extension of Δ. Furthermore let $S \in \mathcal{S}_{\Delta}(\delta)$ be a supporting set of δ. If $S \subseteq \mathit{GD}(D, E)$ then by Definition 3.4 there is no blocking set for δ in $\mathit{GD}(D, E)$. Therefore from $E \models \mathit{Pre}(\delta)$ it follows by definition of the generating defaults of an extension that $\delta \in \mathit{GD}(D, E)$. ■

Proof 3.12 *1.* and *2.* follow immediately from Definition 3.1 and 3.4, respectively. ■

Proof 4.1 If $\Delta = (D, W)$ is non-conflicting then we have $A_{\Delta} = \emptyset$.

Let D' be a maximal subset of D, which is grounded in W. Observe that D' is unique. Define

$$E = \mathit{Th}\big(W \cup \mathit{Cons}(D')\big)\,. \tag{27}$$

We now prove that E is an extension of $\Delta = (D, W)$ by showing the following two statements

$$E = Th(Cons(GD(D,E)) \cup W) \text{ and} \tag{28}$$

$$GD(D,E) \text{ is grounded in } W. \tag{29}$$

Then by Theorem B.3 E is an extension of Δ. It is sufficient to show that $D' = GD(D,E)$, because this implies (28) and (29). Because of (27), D' is grounded in W and $E = Th(W \cup Cons(D'))$, for all $\delta' \in D'$ we have

$$E \models Pre(\delta'). \tag{30}$$

Assume there is a $\delta' \in D'$ such that $E \models \neg Just(\delta')$. Then D' is not weakly regular wrt W and Theorem 3.3 implies that there is $\delta \in D'$ and $B \subseteq D'$ such that $B \in \mathcal{B}_\Delta(\delta)$. Thus, according to Definition 3.3 and 3.1, for each $\delta \in B$ we have $(\delta, \delta') \in A_\Delta$, which is a contradiction to the fact that Δ is non-conflicting. Thus, there is no $\delta' \in D'$ with $E \models \neg Just(\delta')$ and it follows $E \not\models \neg Just(\delta')$ for each $\delta' \in D'$. Hence, with (30) we have $D' \subseteq GD(D,E)$.

By definition, if $\delta \in GD(D,E)$ then $E = Th(W \cup Cons(D')) \models Pre(\delta)$. Thus $\delta \in D'$, because otherwise D' would not be a maximal subset of D, which is grounded in W. Finally, we obtain $D' = GD(D,E)$ and are done.

■

Proof 4.2 That every well-ordered default theory has at least one extension is a direct consequence of Theorem 4.3, since the proof of Theorem 4.3 does not rely on Theorem 4.2.

It remains to show that every well-ordered theory has exactly one extension. Assume that $\Delta = (D, W)$ is well-ordered and has two different extensions E_1 and E_2. Then the corresponding sets of generating defaults $GD(D,E_1)$ and $GD(D,E_2)$ are different and both grounded in W. Clearly, there are enumerations $\langle\delta_i\rangle_{i\in I}$ of $GD(D,E_1)$ and $\langle\delta'_i\rangle_{j\in J}$ of $GD(D,E_2)$ such that (3) holds for both of them and there is a maximal k for which the following condition holds:

$$\delta_i = \delta'_i \text{ for all } i,j \leq k \text{ and } \delta_{k+1} \neq \delta'_{k+1}. \tag{31}$$

That is, there are no other enumerations of $GD(D,E_1)$ and $GD(D,E_2)$ with a greater k fulfilling Condition (31).

Observe that $\{\delta_1, \cdots, \delta_{k+1}\}$ and $\{\delta'_1, \cdots, \delta'_{k+1}\}$ are grounded in W. Furthermore, we have that $W \cup Cons(\{\delta_1, \cdots, \delta_{k+1}\}) \models \neg Just(\delta'_{k+1})$ and $W \cup Cons(\{\delta'_1, \cdots, \delta'_{k+1}\}) \models \neg Just(\delta_{k+1})$, because otherwise k would not be maximal. Since $\{\delta_1, \cdots, \delta_{k+1}\} \subseteq GD(D,E_1)$ and $\{\delta'_1, \cdots, \delta'_{k+1}\} \subseteq GD(D,E_2)$, according to Theorem 3.2, there are minimal subsets $B \subseteq \{\delta_1, \cdots, \delta_{k+1}\}$ and $B' \subseteq \{\delta'_1, \cdots, \delta'_{k+1}\}$ such that $B \in \mathcal{B}_\Delta(\delta'_{k+1})$ and $B' \in \mathcal{B}_\Delta(\delta_{k+1})$. By Definition 3.3 this implies that Δ is not well-ordered, which is a contradiction. Hence, every well-ordered default theory has a single extension.

■

Proof 4.3 If Γ_Δ has no odd cycles, then according to Lemma B.5 Ω_Δ has no odd cycles. Trivially, Ω_Δ^{-1} has no odd cycles. With Theorem A.1 there is a kernel K for Ω_Δ^{-1}. According to Definition A.3 of a kernel, the following two conditions hold for K:

$$K \text{ is independent wrt } \Omega_\Delta \qquad \text{and} \tag{32}$$

$$\forall B \in V_\Delta^\Omega \setminus K : \exists B' \in K \text{ such that } (B', B) \in A_\Delta^\Omega. \tag{33}$$

That is, K is an inverse kernel of Ω_Δ. According to Theorem 4.4, Δ has an extension.

■

Proof 4.4 Let $\Delta = (D, W)$ be a default theory such that Ω_Δ has an inverse kernel K. Then we have

$$K \text{ is independent wrt } \Omega_\Delta \qquad \text{and} \tag{34}$$

$$\forall B \in V_\Delta^\Omega \setminus K : \exists B' \in K \text{ such that } (B', B) \in A_\Delta^\Omega. \tag{35}$$

Define $K^\star = \cup_{B \in K} B$, the union of all blocking sets in K. Because each blocking set is grounded in W, with Lemma B.2 it follows that $K^\star$ is grounded in W. Additionally, we have that for each $\delta \in K^\star$ there is a blocking set $B \in K$ with $\delta \in B$.

Let $\delta \in K^\star$ and assume $W \cup \mathit{Cons}(K^\star) \models \neg \mathit{Just}(\delta)$. Then $K^\star$ is not weakly regular wrt W. According to Theorem 3.3, it follows that there is $\delta' \in K^\star$ and there is $B' \subseteq K^\star$ such that $B' \in \mathcal{B}_\Delta(\delta')$. If $B' \in K$ we obtain a contradiction to (34), because otherwise there would be a $B \in K$ with $\delta' \in B$ and thus $(B', B) \in A^\Omega_\Delta$. If $B' \notin K$ there is a $B \in K$ such that $(B, B') \in A^\Omega_\Delta$ (see (35)). Then, by definition there is a $\delta'' \in B'$ with $B \in \mathcal{B}_\Delta(\delta'')$. Because $B' \subseteq K^\star$, it follows $\delta'' \in K^\star$. Hence, there is $B'' \in K$ with $\delta'' \in B''$. Putting all together, we obtain $B, B'' \in K$, $\delta'' \in B''$ and $B \in \mathcal{B}_\Delta(\delta'')$. This implies $(B, B'') \in A^\Omega_\Delta$, which again is a contradiction to (34). Therefore we have that the assumption was false and $K^\star$ is weakly regular wrt W, that is, for all $\delta \in K^\star$ we obtain

$$W \cup \mathit{Cons}(K^\star) \not\models \neg \mathit{Just}(\delta). \tag{36}$$

Now define

$$E^\star = \mathit{Th}(W \cup \mathit{Cons}(K^\star))\,. \tag{37}$$

With the fact that $K^\star$ is grounded in W and with (36) we conclude $\mathit{GD}(K^\star, E^\star) = K^\star$. Thus, with Theorem B.3

$$E^\star \text{ is an extension of default theory } (K^\star, W). \tag{38}$$

Next we show that default theory $\Delta^\star = (D^\star, E^\star)$ has an extension E. Without loss of generality, let $D^\star = D \setminus \{\delta \mid E^\star \models \neg \mathit{Just}(\delta)\}$. That is, for each $\delta \in D^\star$ the following condition holds:

$$E^\star \not\models \neg \mathit{Just}(\delta). \tag{39}$$

Assume there are defaults $\delta, \delta' \in D^\star$ such that $(\delta, \delta') \in A_{\Delta^\star}$ for the block graph $\Gamma_{\Delta^\star} = (V_{\Delta^\star}, A_{\Delta^\star})$. Then, by Definition 3.3 there is $B^\star \subseteq D^\star$ with $\delta \in B^\star$ such that $B^\star \in \mathcal{B}_{\Delta^\star}(\delta')$. Because $K^\star$ is grounded in W and because of (37) and (39) there is a superset B of $B^\star$ such that $B \in \mathcal{B}_\Delta(\delta')$. We distinguish the following three cases.

1. Let $B \in K$ and $\delta' \in K^\star$. We immediately get a contradiction to (34), because then there would be a $B' \in K$ with $\delta' \in B'$ such that $(B, B') \in A^\Omega_\Delta$.

2. Let $B \in K$ and $\delta' \notin K^\star$. Then $B \subseteq K^\star$. According to Definition 3.1 $B \in \mathcal{B}_\Delta(\delta')$ implies $W \cup \mathit{Cons}(B) \models \neg \mathit{Just}(\delta')$ and it follows $E^\star \models \neg \mathit{Just}(\delta')$. But this is a contradiction to (39) since $\delta' \in D^\star$.

3. For the last case let $B \notin K$ (for $\delta' \in K^\star$ or $\delta' \notin K^\star$). Then (35) implies that there is a $B' \in K$ such that $(B', B) \in A^\Omega_\Delta$, which implies that there is a $\overline{\delta} \in B$ with $B' \in \mathcal{B}_\Delta\left(\overline{\delta}\right)$. By Definition 3.1 it follows that $W \cup \mathit{Cons}(B') \models \neg \mathit{Just}\left(\overline{\delta}\right)$. Since $B' \in K$ we have that $W \cup \mathit{Cons}(K^\star) \models \neg \mathit{Just}\left(\overline{\delta}\right)$ and with (37) this is a contradiction to $\overline{\delta} \in B \subseteq D^\star$ and (39).

In any case we obtain a contradiction, which shows that the assumption was false, hence, there are no defaults $\delta, \delta' \in D^\star$ such that $(\delta, \delta') \in A_{\Delta^\star}$ for the block graph $\Gamma_{\Delta^\star}$. In other words $\Delta^\star$ is non-conflicting and thus has an extension E, according to Theorem 4.1.

Next we prove that $K^\star \subseteq \mathit{GD}(D, E)$. Let $\delta \in K^\star$, then $E \models \mathit{Pre}(\delta)$, because $K^\star$ is grounded in W and $E^\star = \mathit{Th}(W \cup \mathit{Cons}(K^\star))$. According to Theorem 3.2 the following condition is true

$$E \not\models \neg \mathit{Just}(\delta) \text{ iff there is no } B' \subseteq \mathit{GD}(D, E) \text{ such that } B' \in \mathcal{B}_{\Delta^\star}(\delta)\,. \tag{40}$$

Assume there is a $B' \subseteq \mathit{GD}(D, E)$ such that $B' \in \mathcal{B}_{\Delta^\star}(\delta)$. Then we distinguish the following two cases.

1. Let $B' \in K$. Because $\delta \in K^\star$, there is a $B \in K$ with $\delta \in B$ and according to Definition 4.1 we have $(B', B) \in A^\Omega_\Delta$ and $B, B' \in K$. But this is impossible according to (34).

2. Let $B' \notin K$. Then there is a $B \in K$ such that $(B, B') \in A^\Omega_\Delta$. Therefore, there is a $\overline{\delta} \in B'$ such that $E \models \neg Just(\overline{\delta})$, because $B \subseteq K^\star \subseteq E^\star \subseteq E$ and thus $\overline{\delta} \notin GD(D, E)$. But this is a contradiction to $\overline{\delta} \in B' \subseteq GD(D, E)$.

With (40) it follows that $E \not\models \neg Just(\delta)$ and we get that $\delta \in GD(D, E)$. Hence, $K^\star \subseteq GD(D, E)$. To sum up, we have seen that $E^\star$ is an extension of $(K^\star, W)$, $(D, E^\star)$ has extension E (because trivially $\Delta^\star = (D^\star, E^\star)$ and $(D, E^\star)$ have the same extensions) and $GD(K^\star, E^\star) \subseteq GD(D, E)$. According to Lemma B.4 E is an extension of (D, W).

Let E be an extension of $\Delta = (D, W)$ and define $K = \{B \mid B \in \mathcal{B}_\Delta(\delta)$ for some $\delta \in D$ and $B \subseteq GD(D, E)\}$. We show that K is an inverse kernel wrt Ω_Δ by proving (34) and (35) for K.

Since $B \subseteq GD(D, E)$ for each $B \in K$ it follows by Definition 4.1 that K is independent wrt Ω_Δ, that is (34) holds for K. Otherwise $GD(D, E)$ would not be weakly regular (see Theorem 3.3).

Now let $B \in V^\Omega_\Delta \setminus K$. Then there exists a $\delta \in B \setminus GD(D, E)$ s.t. $W \cup Cons(GD(D, E)) \models \neg Just(\delta)$. If there would be no such δ in B then B would be a subset of $GD(D, E)$ because B is grounded in W. But this would imply that B is in K, which is a contradiction to $B \in V^\Omega_\Delta \setminus K$. From $W \cup Cons(GD(D, E)) \models \neg Just(\delta)$ we conclude with Theorem 3.3 that there is some $B' \subseteq GD(D, E)$ s.t. $B' \in \mathcal{B}_\Delta(\delta)$. This means that for each $B \in V^\Omega_\Delta \setminus K$ there is some $B' \in K$ s.t. $(B', B) \in A^\Omega_\Delta$, that is (35). Hence K is an inverse kernel wrt Ω_Δ. ■

Proof 4.5 If C is an (directed) harmful, odd cycle in Γ_Δ then there are only arcs between defaults in C which belong to the cycle. Let $C = \{\delta_1, \cdots, \delta_{2n+1}\}$ and for each $1 \leq i \leq 2n+1$ let B_i the blocking set of δ_{i+1} such that $\delta_i \in B_i$ (see Definition 3.3). Since C is harmful we have $B_i = \{\delta_i\}$ for each $1 \leq i \leq 2n+1$.

Now assume Δ has an extension E. Then we have $E = Th(W \cup Cons(GD(D, E)))$ and trivially not all of the above B_i can be subsets of $GD(D, E)$. Thus without loss of generality, assume $B_1 \not\subseteq GD(D, E)$. Since B_1 is the only blocking set of B_2 it follows that $B_2 \subseteq GD(D, E)$. Hence, $B_3 \not\subseteq GD(D, E)$, because $B_2 \in \mathcal{B}_\Delta(\delta_3)$ and $\delta_3 \in B_3$. After repeating this argument n times it follows that $B_{2n+1} \not\subseteq GD(D, E)$. But know we are able to conclude that $B_1 \subseteq GD(D, E)$ with the same argumentation as above. This is, because C is harmful. Therefore we obtain a contradiction, which shows that our assumption was false, and thus Δ has no extension. ■

Proof 4.6 We prove the proposition by showing that Δ is not ps-even under the assumption that Δ is not even. For definition of ps-even, $G(\Delta)$ and odd cycles of $G(\Delta)$ used in this proof see [39]. So let Δ be not even, that is, there is an odd cycle $C = \{\delta_1, \cdots, \delta_{2n+1}\}$ in Γ_Δ. Then we have $(\delta_i, \delta_{i+1}) \in A_\Delta$ for each $1 \leq i < 2n+1$ and $(\delta_{2n+1}, \delta_1) \in A_\Delta$. According to Definition 3.3, for each i we have $\delta_i \in B_i$ for some $B_i \in \mathcal{B}_\Delta(\delta_{i+1})$. Thus by the definition of blocking sets

$$W \cup Cons(B_i) \models \neg Just(\delta_{i+1}) \text{ and} \tag{41}$$

$$B_i \text{ is grounded in } W. \tag{42}$$

Before we continue, we give a translation of the second proposition of Lemma 1 in [39] into our terminology. Let $A, B \subseteq D$ be subsets of defaults and $\delta, \delta' \in D$ defaults. Then the second proposition in Lemma 1 in [39] states

$$\begin{aligned} &(W \cup Cons(A) \models Pre(\delta') \text{ and } W \cup Cons(B) \not\models \neg Just(\delta')) \qquad \text{and} \\ &(W \cup Cons(A) \not\models Pre(\delta') \text{ or } W \cup Cons(B \cup \{\delta\}) \models \neg Just(\delta')) \\ &\text{implies } (\delta, \delta') \in E_1. \end{aligned} \tag{43}$$

Now let $\delta_{i+1} \in B_{i+1}$, $A = B_{i+1}$, $B = B_i \setminus \{\delta_i\}$ and $\delta = \delta_i$ and $\delta' = \delta_{i+1}$. Then (41), (42) and (43) imply $(\delta_i, \delta_{i+1}) \in E_1$ for each $1 \leq i < 2n+1$ and with the same argumentation $(\delta_{2n+1}, \delta_1) \in E_1$.

Hence there is an odd cycle (as defined in [39]) in $G(\Delta)$. That is, Δ is not ps-even and the proof is finished. ■

Proof 4.7 Let $\Delta = (D, W)$ be a default theory and let $\mathcal{A}_\Delta = \langle T, A, \overline{} \rangle$ be the corresponding assumption-based framework.

Before staring with the actual proof, observe that we have the following lemma:

> **Lemma B.8** *If (δ, δ') is an arc in the block graph of Δ then $(Just(\delta), Just(\delta'))$ is an arc in the attack relationship graph of $\mathcal{A}$.*
>
> This lemma is an immediate consequence from the definition of an attack and the definition of the attack relationship graph in [5].

For the first part of the theorem let $\mathcal{A}_\Delta$ be stratified according to [5]. Then by Definition 7.2 in [5] there is no cycle in the attack relationship graph of $\mathcal{A}_\Delta$. Thus, according to Lemma B.8 there is no cycle in the block graph of Δ. Hence Δ is well-ordered.

Now let $\mathcal{A}_\Delta$ be order-consistent according to [5] and let $G = (V, A)$ be the attack relationship graph corresponding to $\mathcal{A}_\Delta$. Bondarenko et al. define a relation $\prec$ between two nodes v and v' ($v \prec v'$) iff there exists both a path with an even number of edges and a path with an odd number of edges from v to v' in G (see [5, Definition 7.6]). $\mathcal{A}_\Delta$ is order-consistent iff the graph $(V, \{(v, v') \mid v \prec v'\})$ has no cycles. Now assume that there is an odd cycle C in G. Then for any two nodes v, v' in C there is both a path with an even number of edges and a path with an odd number of edges from v to v'. If the shortest path through the cycle C from v to v' is even (odd) the path from v to v' over v' (adding one run through the cycle to the original path from v to v') is odd (even). Hence for any two nodes v, v' in C we have $v \prec v'$ and graph $(V, \{(v, v') \mid v \prec v'\})$ has a cycles. That is, $\mathcal{A}_\Delta$ is not order-consistent.

Since $\mathcal{A}_\Delta$ is order-consistent the corresponding attack relationship graph has no odd cycles. Then, according to Lemma B.8 there is no odd cycle in the block graph of Δ and Δ is even. ■

Proof 5.1 Let $\Delta = (D, W)$ be a default theory and E an extension of Δ. Then with Theorem B.3 we have

$$E = Th(W \cup Cons(GD(D, E))) \quad \text{and} \tag{44}$$

$$GD(D, E) \text{ is grounded in } W. \tag{45}$$

For proving that $GD(D, E)$ is protected in Δ let $\delta \in GD(D, E)$. By definition of the generating default rules for E, we know that $E \models Pre(\delta)$ and $E \not\models \neg Just(\delta)$. Therefore $GD(D, E)$ contains no blocking set of δ. Thus for each $B \in \mathcal{B}_\Delta(\delta)$ there exists a $\delta' \in B$ such that $E \models \neg Just(\delta')$, because by Definition 3.1 each blocking set B of δ is grounded in W. Otherwise from (45) and an easy induction, it would follow that $B \subseteq GD(D, E)$. More formally we have

$$\forall B \in \mathcal{B}_\Delta(\delta) : \exists \delta' \in B \text{ such that } W \cup Cons(GD(D, E)) \models \neg Just(\delta'). \tag{46}$$

According to Theorem 3.2 this implies

$$\forall B \in \mathcal{B}_\Delta(\delta) : \exists \delta' \in B : \exists B' \in \mathcal{B}_\Delta(\delta') \text{ such that } B' \subseteq GD(D, E). \tag{47}$$

Thus, by Definition 3.4 for each $\delta \in GD(D, E)$ exists a $S \in \mathcal{S}_\Delta(\delta)$ such that $S \subseteq GD(D, E)$ and that is $GD(D, E)$ is protected in Δ. ■

Proof 5.2 Let $D' \subseteq GD(D, E)$ be a set of default rules s.t. D' is grounded in W. First, recall the following definition. For a set D' of default rules we define $\overline{D'} := \{\delta \in D \mid W \cup Cons(D') \models \neg Just(\delta)\}$. Let $D'' := D \setminus (D' \cup \overline{D'})$ for the rest of this proof.

Let E be an extension of Δ. According to Theorem B.3, we have

$$\begin{array}{l} E = Th(W \cup Cons(GD(D, E))) \quad \text{and} \\ GD(D, E) \text{ is grounded in } W. \end{array} \tag{48}$$

Next we show $\mathit{GD}(D, E) = \mathit{GD}(D'', E) \cup D'$. Let $\delta \in \mathit{GD}(D, E)$. According to definition, it follows $E \models \mathit{Pre}(\delta)$ and $E \not\models \neg\mathit{Just}(\delta)$. If $\delta \in D'$ we are ready. Let $\delta \notin D'$. Because $D' \subseteq \mathit{GD}(D, E)$ from $E \not\models \neg\mathit{Just}(\delta)$ with (48) it follows $W \cup \mathit{Cons}(D') \not\models \neg\mathit{Just}(\delta)$. That is, $\delta \notin \overline{D'}$ and we conclude $\delta \in D''$. Since $E \models \mathit{Pre}(\delta)$ and $E \not\models \neg\mathit{Just}(\delta)$ it follows that $\delta \in \mathit{GD}(D'', E) \cup D'$. Now let $\delta \in \mathit{GD}(D'', E) \cup D'$. If $\delta \in D'$ we are ready, because $D' \subseteq \mathit{GD}(D, E)$. Let $\delta \notin D'$, that is, $\delta \in \mathit{GD}(D'', E)$. By definition, we have $E \models \mathit{Pre}(\delta)$ and $E \not\models \neg\mathit{Just}(\delta)$, that is, $\delta \in \mathit{GD}(D, E)$. All in all we obtain

$$\mathit{GD}(D, E) = \mathit{GD}(D'', E) \cup D'. \tag{49}$$

From (48), (49) and the fact that D' is grounded in W we conclude (see Lemma B.2)

$$\begin{aligned} &E = \mathit{Th}(W \cup \mathit{Cons}(D' \cup \mathit{GD}(D'', E))) \qquad \text{and} \\ &\mathit{GD}(D'', E) \text{ is grounded in } W \cup \mathit{Cons}(D'). \end{aligned} \tag{50}$$

According to Theorem B.3, E is an extension of $\Delta|D'$.

For the other direction let E be an extension of $\Delta|D'$. According to Theorem B.3, we have (50). Since D' is grounded in W, according to Theorem B.2, it follows that $D' \cup \mathit{GD}(D'', E)$ is grounded in W. By using (50) instead of (48), a similar argumentation as above shows that (49) and thus (48) is also true under the current premises. As above, it follows that E is an extension of Δ. ■

Proof 5.3 Let $\Delta = (D, W)$ be a default theory and E an extension of Δ. Then with Theorem B.3 we have

$$E = \mathit{Th}(W \cup \mathit{Cons}(\mathit{GD}(D, E))) \qquad \text{and} \tag{51}$$

$$\mathit{GD}(D, E) \text{ is grounded in } W. \tag{52}$$

If we set $E' = E$ and $D' = \mathit{GD}(D, E)$ then $\Delta|D' = (D \setminus \mathit{GD}(D, E), W \cup \mathit{Cons}(\mathit{GD}(D, E)))$. With the abbreviations $D'' = D \setminus \mathit{GD}(D, E)$ and $W' = W \cup \mathit{Cons}(\mathit{GD}(D, E))$ and Theorem B.3 E' is an extension of $\Delta|D'$ iff

$$E = \mathit{Th}\big(W' \cup \mathit{Cons}\big(\mathit{GD}(D'', E')\big)\big) \qquad \text{and} \tag{53}$$

$$\mathit{GD}(D'', E') \text{ is grounded in } W'. \tag{54}$$

Assume $\delta \in \mathit{GD}(D'', E')$. Then δ is not in $\mathit{GD}(D, E)$ because $D'' = D \setminus \mathit{GD}(D, E)$, that is $E \not\models \mathit{Pre}(\delta)$ or $E \models \neg\mathit{Just}(\delta)$. Because $E = E'$, this implies that δ is not in $\mathit{GD}(D'', E')$, which is a contradiction. Thus $\mathit{GD}(D'', E')$ is empty. Now we have seen (54) because the empty set trivially is grounded in W' and (53) follows from (51). That is E' is an extension of $\Delta|D'$.

It remains to show that D' is grounded in W and protected in Δ. That D' is grounded in W follows immediately from (52) and D' is protected in Δ according to Theorem 5.1.

For the other direction let

$$E = \mathit{Th}(W \cup \mathit{Cons}(D') \cup E') \tag{55}$$

for a $D' \subseteq D$ such that D' is grounded in W, D' is protected in Δ and E' is an extension of default theory $\Delta|D'$. If we set $D'' = D \setminus D'$ then for $\Delta|D' = (D'', W \cup \mathit{Cons}(D'))$ we have

$$E' = \mathit{Th}\big(W \cup \mathit{Cons}\big(D' \cup \mathit{GD}(D'', E')\big)\big) \qquad \text{and} \tag{56}$$

$$\mathit{GD}(D'', E') \text{ is grounded in } W \cup \mathit{Cons}\big(D'\big) \tag{57}$$

according to Theorem B.3. Therefore we have

$$E = E'. \tag{58}$$

Next we show $D' \subseteq \mathit{GD}(D,E)$. According to Definition 5.1 for each $\delta \in D'$ we have the following two conditions:

$$B \subseteq D' \text{ for no } B \in \mathcal{B}_\Delta(\delta) \tag{59}$$
$$S \subseteq D' \text{ for some } S \in \mathcal{S}_\Delta(\delta)\,. \tag{60}$$

Let $\delta \in D'$. With the groundedness of D' and (56) we have $E \models \mathit{Pre}(\delta)$. Assume $E \models \neg \mathit{Just}(\delta)$. Then with (56) and (58) we have that $D' \cup \mathit{GD}(D'',E')$ is not weakly regular. According to Theorem 3.3 there is some $\delta' \in D' \cup \mathit{GD}(D'',E')$ and some $B' \subseteq D' \cup \mathit{GD}(D'',E')$ s.t. $B' \in \mathcal{B}_\Delta(\delta')$. With (55) and (56) it follows that $E' \models \neg \mathit{Just}(\delta')$. Therefore

$$\delta' \notin \mathit{GD}(D'',E') \tag{61}$$

We distinguish the following cases:

1. $B' \subseteq D'$ and $\delta' \in D'$:
   In this case we obtain a contradiction to (60), because D' contains a default and one of its blocking sets.

2. $B' \subseteq D'$ and $\delta' \notin D'$:
   In this case we have $\delta' \in \mathit{GD}(D'',E')$ which is a contradiction to (61).

3. $B' \not\subseteq D'$ and $\delta' \in D'$:
   In this case we have $D' \cap \mathit{GD}(D'',E') \neq \emptyset$.

4. $B' \not\subseteq D'$ and $\delta' \notin D'$:
   Again, we have $\delta' \in \mathit{GD}(D'',E')$ which is a contradiction to (61).

Therefore the assumption is false and we have $E \not\models \neg \mathit{Just}(\delta)$ which implies

$$D' \subseteq \mathit{GD}(D,E). \tag{62}$$

Finally, we show that $\mathit{GD}(D,E) = D' \cup \mathit{GD}(D'',E')$.

Let $\delta \in \mathit{GD}(D,E)$, then by definition $E \models \mathit{Pre}(\delta)$ and $E \not\models \neg \mathit{Just}(\delta)$. If $\delta \in D'$ we are ready, because of (62). Let $\delta \notin D'$. From $E \models \mathit{Pre}(\delta)$ and (58) we conclude $E' \models \mathit{Pre}(\delta)$. From $E \not\models \neg \mathit{Just}(\delta)$ we get $E' \not\models \neg \mathit{Just}(\delta)$. Therefore $\delta \in \mathit{GD}(D'',E')$ which implies $\mathit{GD}(D,E) \subseteq D' \cup \mathit{GD}(D'',E')$.

Now let $\delta \in D' \cup \mathit{GD}(D'',E')$. If $\delta \in \mathit{GD}(D'',E')$ then by definition we have $E' \models \mathit{Pre}(\delta)$ which implies $E \models \mathit{Pre}(\delta)$ and $E' \not\models \neg \mathit{Just}(\delta)$. From (58) it follows $E \not\models \neg \mathit{Just}(\delta)$ and also $\delta \in \mathit{GD}(D,E)$. That is, for the first case we get $D' \cup \mathit{GD}(D'',E') \subseteq \mathit{GD}(D,E)$. For the second case let $\delta \notin \mathit{GD}(D'',E')$, that is $\delta \in D'$. With (62) it follows $\delta \in \mathit{GD}(D,E)$.

Finally we obtain $\mathit{GD}(D,E) = D' \cup \mathit{GD}(D'',E')$. According to (55) and (56) it follows $E = \mathit{Th}(W \cup \mathit{GD}(D,E))$. Because D' is grounded in W with (57) we conclude that $\mathit{GD}(D,E)$ is grounded in W. Therefore Theorem B.3 implies that E is an extension of Δ.

■

Proof 5.4 If E is an extension of (D,W) then according to Theorem B.3 we have

$$E = \mathit{Th}(W \cup \mathit{Cons}(\mathit{GD}(D,E))) \tag{63}$$

According to definition of generating defaults (see (2)) for each default rule δ **not** in $\mathit{GD}(D,E)$ we have

$$E \not\models \mathit{Pre}(\delta) \qquad \text{or} \qquad E \models \neg \mathit{Just}(\delta). \tag{64}$$

Set $\Delta' = (D',W') = \Delta|\mathit{GD}(D,E)$ then by definition we have $D' = D \setminus (\mathit{GD}(D,E) \cup \overline{\mathit{GD}(D,E)})$ and $W' = W \cup \mathit{Cons}(\mathit{GD}(D,E))$. It remains to show that $A_{\Delta'} = \emptyset$, because then Δ' is non-conflicting.

Assume $(\delta_1, \delta_2) \in A_{\Delta'}$ for $\delta_1, \delta_2 \in D'$, that is δ_1 and δ_2 are not in the generating defaults of extension E, that is (64) holds for δ_1 and δ_2. According to Definitions 3.3 and 3.1, exists a $B \subseteq D'$ such that $\delta_1 \in B$ and $B \in \mathcal{B}_\Delta(\delta_2)$, because $(\delta_1, \delta_2) \in A_{\Delta'}$. That is in particular B is grounded in W'. According to the definition of $|$, we have $W' \not\models \neg Just(\delta)$ for each δ in B. For each δ in B we have $W' \not\models Pre(\delta)$ because of (63) and (64). This implies that B is not grounded in W' which is a contradiction to (63). Thus our assumption was false and for all $\delta_1, \delta_2 \in B$ we have $(\delta_1, \delta_2) \notin A_{\Delta'}$ That is, the block graph $\Gamma_{\Delta|GD(D,E)}$ has no arcs and by definition $\Delta' = \Delta|GD(D, E)$ is non-conflicting. ■

Proof 5.5 This result is obtained as a direct consequence of the proof of Theorem 4.1. ■

Proof 5.6 Let E be an extension of Δ. Theorems 5.1 and 5.4 imply that $GD(D, E)$ is protected in Δ and $\Gamma_{\Delta|GD(D,E)}$ is arcless, respectively.

Assume that $GD(D, E)$ is not maximal s.t. $GD(D, E)$ is grounded in W and protected in Δ. Then there is a superset D'' of $GD(D, E)$ for which the above conditions hold. That is, there is a $\delta \in D''$ and $\delta \notin GD(D, E)$ s.t. $E \models Pre(\delta)$. Since D'' is protected in Δ, according to Theorem 3.3, D'' is weakly regular. Therefore we have $E \not\models \neg Just(\delta)$ and it follows that $\delta \in GD(D, E)$ which is a contradiction. Hence $GD(D, E)$ is a maximal set with the desired properties.

For the other direction let $E = Th(W \cup Cons(D'))$ for some maximal $D' \subseteq D$ s.t. D' is grounded in W and protected in Δ. Set $(D'', W'') = \Delta|D'$, then $D'' = D \setminus (D' \cup \overline{D'})$ and $W'' = W \cup Cons(D')$.

According to Theorem 5.8, we have $A_{\Delta|D'} \subseteq A_\Delta$. Since Δ is even $\Delta|D'$ is even and thus has an extension E'. According to Theorem 5.3, we have that $E' = Th(W \cup Cons(D' \cup GD(D'', E')))$ is an extension of Δ. Theorems B.3 and 5.1 imply that $D' \cup GD(D, E)$ is grounded in W and protected in Δ.

Since D' was a maximal set being grounded in W and protected in Δ we have

$$GD(D'', E') = \emptyset.$$

Finally, this implies $E' = E$ and we obtain E is an extension of Δ. ■

Proof 5.7 Let E be an extension of Δ. According to Theorem B.3, we have

$$E = Th(W \cup Cons(GD(D, E))) \quad \text{and}$$
$$GD(D, E) \text{ is grounded in } W.$$

Theorems 5.1 and 5.4 imply that $GD(D, E)$ is protected in Δ and $\Gamma_{\Delta|GD(D,E)}$ is arcless, respectively.

Assume that $GD(D, E)$ is not maximal s.t. $GD(D, E)$ is grounded in W, protected in Δ and $\Gamma_{\Delta|D'}$ is arcless. Then there is a superset D'' of $GD(D, E)$ for which the above conditions hold. That is, there is a $\delta \in D''$ and $\delta \notin GD(D, E)$ s.t. $E \models Pre(\delta)$. Since D'' is protected in Δ, according to Theorem 3.3, D'' is weakly regular. Therefore we have $E \not\models \neg Just(\delta)$ and it follows that $\delta \in GD(D, E)$ which is a contradiction. Hence $GD(D, E)$ is a maximal set with the desired properties.

For the other direction let $E = Th(W \cup Cons(D'))$ for some maximal $D' \subseteq D$ s.t. D' is grounded in W and protected in Δ and $\Gamma_{\Delta|D'}$ is arcless (that is, $\Delta|D'$ is non-conflicting). Set $(D'', W'') = \Delta|D'$, then $D'' = D \setminus (D' \cup \overline{D'})$ and $W'' = W \cup Cons(D')$.

Next, we show $D' = GD(D, E)$. Since D' is grounded and protected, by definition, we know that $D' \subseteq GD(D, E)$. Let $\delta \in GD(D, E)$ then $E \models Pre(\delta)$ and $E \not\models \neg Just(\delta)$. Assume $\delta \notin D'$. Then $\delta \notin D' \cup \overline{D'}$, that is, $\delta \in D''$. Since D' is maximal set with the above properties, for $D' \cup \{\delta\}$ not all properties (grounded in W and protected in Δ and $\Gamma_{\Delta|(D' \cup \{\delta\})}$ arcless) hold. Because $D' \cup \{\delta\}$ is grounded in W, we know that $D' \cup \{\delta\}$ is not protected in Δ or for $D' \cup \{\delta\}$ we have that $\Gamma_{\Delta|(D' \cup \{\delta\})}$ is not arcless. But $\Gamma_{\Delta|(D' \cup \{\delta\})}$ is arcless since $\Gamma_{\Delta|D'}$ is arcless and $D' \subseteq D' \cup \{\delta\}$. That is, $D' \cup \{\delta\}$ is not protected in Δ. Then $D' \cup \{\delta\}$ is not weakly regular or it contains no supporting set for δ. Assume $D' \cup \{\delta\}$ is not weakly regular. According to Theorem 3.3, we have that there is a $\delta' \in D' \cup \{\delta\}$ and a $B' \in \mathcal{B}_\Delta(\delta')$ s.t. $B' \subseteq D' \cup \{\delta\}$. We have the following two cases.

1. If $B' \subseteq D'$ then $W \cup Cons(D') \models \neg Just(\delta')$. But this is a contradiction to D' protected in Δ or $\delta \in GD(D, E)$, depending on whether $\delta' \in D'$ or $\delta' = \delta$, respectively.

2. If $B' \not\subseteq D'$ then we know that $\delta \in B'$. Again we have two cases:

   (a) $\delta' \in D'$. Since D' is protected there is a supporting set S of δ' s.t. $S \subseteq D'$ and $S \in \mathcal{B}_\Delta(\delta)$ because otherwise D' would not be protected. Hence $E \models \neg Just(\delta)$ which is a contradiction to $E \not\models \neg Just(\delta)$ (see above).

   (b) $\delta' \notin D'$ implies $\delta' = \delta$. We have that $\Gamma_{\Delta|D'}$ is not arcless, because $\{\delta\} \in \mathcal{B}_{\Delta|D'}(\delta)$. This follows form the fact that for each single default δ and each default theory Δ form $\{\delta\} \mapsto_\Delta \delta'$ it follows that $\{\delta\} \in \mathcal{B}_\Delta(\delta')$. Again this is a contradiction.

Therefore the assumption was false and $D' \cup \{\delta\}$ is weakly regular wrt W. Now assume that $D' \cup \{\delta\}$ contains no supporting set for δ. That is, δ has supporting sets different from empty set. Therefore, according to Definition 3.4, $\mathcal{B}_\Delta(\delta) \neq \emptyset$. And for all blocking sets $B \in \mathcal{B}_\Delta(\delta)$ we have that $B \not\subseteq D' \cup \{\delta\}$, because otherwise $D' \cup \{\delta\}$ would not be weakly regular wrt W. Since D' contains no supporting set for δ it follows that there is a $B \in \mathcal{B}_\Delta(\delta)$ s.t. $B \subseteq D' \cup D''$. From $\delta \in D''$ we conclude $W'' = W \cup Cons(D') \not\models \neg Just(\delta)$. Therefore there exists a minimal $B'' \subseteq B \cap D''$ s.t. $B'' \mapsto_{\Delta|D'} \delta$. That is,

$$B'' \subseteq B. \tag{65}$$

If we assume $B'' \notin \mathcal{B}_{\Delta|D'}(\delta)$ there is a $B' \subset B''$ and a $\delta' \in B''$ s.t. $B' \mapsto_{\Delta|D'} \delta'$ (see Theorem 3.1). Now (65) implies $B' \subset B$ and $\delta' \in B$. Then $B' \cup (B \setminus B'') \subset B$ and $B' \cup (B \setminus B'') \mapsto_\Delta \delta'$, that is, $B \notin \mathcal{B}_\Delta(\delta)$. But this is a contradiction and we obtain $B'' \in \mathcal{B}_{\Delta|D'}(\delta)$. Now this again is a contradiction to the fact that $\Gamma_{\Delta|D'}$ is arcless. Thus our initial assumption was false, that is there is a supporting set for δ in $D' \cup \{\delta\}$ and therefore $D' \cup \{\delta\}$ is protected in Δ. This is a contradiction to the maximality of D' and hence the assumption that $\delta \notin D'$ was false.

We have seen that $D' = GD(D, E)$ and according to Theorem B.3, E is an extension of Δ. ■

Proof 5.8 Let $D' \subseteq D$ a set of default rules which is is grounded in W and protected in Δ. Set $\Delta|D' = (D'', W'')$. Then by Definition 5.2 of $|$ we have

$$\begin{array}{l} W'' = W \cup Cons(D') \\ D'' = D \setminus (D' \cup \overline{D'}) \text{ where } \overline{D'} = \{\delta \in D \mid W \cup Cons(D') \models \neg Just(\delta)\}. \end{array} \tag{66}$$

If $(\delta^\star, \delta) \in A_{\Delta|D'}$ then by Definition 3.3 there is some $B^\star \in \mathcal{B}_{\Delta|D'}(\delta)$ such that $\delta^\star \in B^\star$. By Definition 3.1 we have that $B^\star \mapsto_{\Delta|D'} \delta$ and **BS3** hold for $B^\star$. From $B^\star \mapsto_{\Delta|D'} \delta$ it follows that there is a minimal $B \subset D$ with $B = B^\star \cup \overline{B}$ for a $\overline{B} \subseteq D'$ and $B \mapsto_\Delta \delta$. According to Theorem 3.1, $B \in \mathcal{B}_\Delta(\delta)$ iff for each $\delta' \in B$ and each $B' \subset B$ we have $B' \not\mapsto_\Delta \delta'$. Assume $B \notin \mathcal{B}_\Delta(\delta)$, that is, there is a $\delta' \in B$ and a $B' \subset B$ such that $B' \mapsto_\Delta \delta'$. By definition it follows $W \cup Cons(B') \models \neg Just(\delta')$. Furthermore, by monotonicity of $\models$, we obtain $W \cup Cons(B) \models \neg Just(\delta')$ and $\delta' \in B$. That is, B is not weakly regular. Theorem 3.3 implies the existence of $\delta'' \in B$ and $B'' \subset B$ such that $B'' \in \mathcal{B}_\Delta(\delta'')$. We distinguish the following cases:

1. $B'' \subseteq D'$ and $\delta'' \in D'$:
   In this case we obtain a contradiction to the fact that D' is a protected default proof, because it contains a default and one of its blocking sets (see Definition 6.2).

2. $B'' \subseteq D'$ and $\delta'' \notin D'$:
   Since $\delta'' \in B$ and $B = B^\star \cup \overline{B}$ for $\overline{B} \subseteq D'$ and $B^\star \subseteq D''$ we conclude $\delta'' \in B^\star \subseteq D''$. Since B'' is a blocking set of δ'' with (66) and $B'' \subseteq D'$ we have $\delta'' \notin D''$, which is a contradiction.

3. $B'' \not\subseteq D'$ and $\delta'' \in D'$:
   In this case there exists a $\overline{\delta} \in B''$ such that $\overline{\delta} \notin D'$ (that is $\overline{\delta} \in D''$). Furthermore, D' contains a blocking set of $\overline{\delta}$, because D' is protected in Δ and $\delta'' \in D'$. Thus by (66) $\overline{\delta} \notin D''$ which is a contradiction.

4. $B'' \not\subseteq D'$ and $\delta'' \notin D'$:
   Then $\delta'' \in B^\star$. If $\delta'' \in B''$ we obtain a contradiction to the fact that Γ_Δ has no self loops. If $\delta'' \notin B''$ we have that $B^\star \neq B''$. Therefore $B^\star \cap B'' \subset B^\star$ and $\delta'' \in B^\star$ such that $(B'' \cap B^\star) \mapsto_\Delta \delta''$. But this is a contradiction to $B^\star \in \mathcal{B}_{\Delta|D'}(\delta)$.

Hence the above assumption was false and $B \in \mathcal{B}_\Delta(\delta)$. Finally, we conclude that $A_{\Delta|D'} \subseteq A_\Delta$. ∎

Proof 5.9 Let $\Delta = (D, W)$, $D^\star = \gamma^\star(D')$ for some $D' \subseteq D$ and $\Delta^\star = (D^\star, W)$. Furthermore let $E^\star$ be an extension of $\Delta^\star$. Then $GD(D^\star, E^\star)$ is protected in $\Delta^\star$. Clearly, $GD(D^\star, E^\star)$ is also protected in Δ since $D^\star \subseteq D$.

Let E' be an extension of $\Delta|GD(D^\star, E^\star)$, which exists because $\Delta|GD(D^\star, E^\star)$ is coherent. Because $GD(D^\star, E^\star)$ is grounded in W, Theorem 5.3 implies that $E = Th(W \cup Cons(GD(D^\star, E^\star)) \cup E')$ is an extension of Δ. Since $E^\star = Th(W \cup Cons(GD(D^\star, E^\star)))$ we conclude $E^\star \subseteq E$. ∎

Proof 6.1 This Corollary follows directly from Theorem 5.3. ∎

Proof 6.2 Let $\Delta = (D, W)$ be a non-conflicting default theory and E an extension of Δ with $\varphi \in E$. Then with Theorem B.3 we have $E = Th(W \cup Cons(GD(D, E)))$ and $GD(D, E)$ is grounded in W. Let $D_p = GD(D, E)$, then **P2** and **P1** are trivially true. Thus D_p is a pure default proof for φ from Δ.

For the other direction let D_p be a pure default proof for φ from Δ. Then by Definition 6.1 **P2** and **P1** hold for D_p. Since Δ is non-conflicting it has an unique extension $E = Th(W \cup Cons(D'))$ where $D' \subseteq D$ is a maximal subset of defaults which is grounded in W (see proof of Theorem 4.1). Because according to **P1** D_p is grounded in W it follows that $D_p \subseteq D'$. Therefore from **P2** it follows that E is an extension of Δ with $\varphi \in E$. ∎

Proof 6.3 Without loss of generality, let $P \subseteq D$ be minimal, pure default proof for φ from Δ. Let $\Delta^\star = (\gamma^\star(P), W)$. If default theory $\Delta^\star$ has an extension E with $\varphi \in E$ then, according to Theorem 5.1 the corresponding set of generating defaults $GD(\gamma^\star(P), E)$ is protected in $\Delta^\star$ and grounded in W. According to the definition of $\gamma^\star(P)$ there are no arcs (δ', δ) in Γ_Δ such that $\delta' \in D \setminus \gamma^\star(P)$ and $\delta \in \gamma^\star(P)$. Thus $B \in \mathcal{B}_\Delta(\delta)$ implies $B \subseteq \gamma^\star(P)$ for each $\delta \in \gamma^\star(P)$. Since $\gamma^\star(P) \subseteq D$ with Theorem 3.4 we have $\mathcal{B}_{\Delta^\star}(\delta) \subseteq \mathcal{B}_\Delta(\delta)$ for each $\delta \in \gamma^\star(P)$. $B \in \mathcal{B}_\Delta(\delta)$ implies $B \subseteq \gamma^\star(P)$. Hence, according to Theorem 3.4 we have $B \in \mathcal{B}_{\Delta^\star}(\delta)$. In all we get $\mathcal{B}_{\Delta^\star}(\delta) = \mathcal{B}_\Delta(\delta)$ for each $\delta \in \gamma^\star(P)$. Therefore $GD(\gamma^\star(P), E)$ is also protected in Δ (see Definition 5.1 and Theorem 3.4). Hence $GD(\gamma^\star(P), E)$ is a protected default proof for φ from Δ. ∎

Proof 6.4 Let $\Delta = (D, W)$ be a even default theory and E an extension of Δ with $\varphi \in E$. Then with Theorem B.3 we have

$$E = Th(W \cup Cons(GD(D, E))) \quad \text{and} \tag{67}$$

$$GD(D, E) \text{ is grounded in } W. \tag{68}$$

We set $D_\varphi = GD(D, E)$ and prove that D_φ is a protected default proof for φ from Δ. (67) and (68) imply that conditions **P2** and **P1** hold for D_φ. Thus **PP1** holds. Theorem 5.1 implies that D_φ is protected in Δ (that is **PP2**) and thus a protected default proof fpr φ.

For the other direction let D_φ be a protected default proof for φ from Δ. Then by Definition 6.2 we have **PP1** and **PP2** for D_φ. According to Theorem 5.8 we have that $A_{\Delta|D_\varphi} \subseteq A_\Delta$, which implies that $\Delta|D_\varphi$ is even, because Δ is even. Now let E' be an extension of $\Delta|D_\varphi$ (E' exists according to Theorem 4.2) and set $E = Th(D_\varphi \cup E')$. Then, according to Corollary 6.1, E is an extension of Δ with $\varphi \in E$. ∎

Proof 6.5 Let $\Delta = (D, W)$ be a default theory and E an extension of Δ with $\varphi \in E$. Set $D_\varphi = GD(D, E)$. Observe that in the first half of the proof of Theorem 6.4 we do not have used that Δ is even. Therefore here the same argumentation applies to show **DP1**, that is D_φ is a protected default proof for φ from Δ. According to Theorem 5.4 $\Delta|D_\varphi$ is non-conflicting, that is **DP2**. Thus D_φ is a (general) default proof for φ.

For the other direction let D_φ be a (general) default proof for φ from Δ, that is **DP1** and **DP2** hold for D_φ. Therefore D_φ is a protected default proof, that is **PP1** and **PP2** hold for D_φ, according to Definition 6.2. According to Definition 6.1 **P2** and **P1** follow for D_p because D_p is also a pure default proof. By **DP2** $\Delta|D_\varphi$ is has an extension E'. Set $E = \mathit{Th}(D_\varphi \cup E')$, then by Corollary 6.1 E is a classical extension of Δ with $\varphi \in E$.

■

Proof 6.6 This theorem follows immedialely from Theorem 6.5 because the extension-membership-problem is in Σ_2^P [25]. ■

References

[1] G. Antoniou. *Non-monotonic Reasoning.* MIT Press, 1998.

[2] P. Besnard. *An Introduction to Default Logic*. Symbolic Computation. Springer Verlag, 1989.

[3] P. Bonatti. Proof systems for default and autoepistemic logic. In P. Miglioli, U. Moscato, D. Mundici, M. Ornaghi, editors, *Proceedings Tableaux'96*, pages 127–142. Springer Verlag, 1996.

[4] P. Bonatti and N. Olivetti. A sequent calculus for skeptical default logic. In *Proceedings Tableaux'97*, pages 107–121. Springer Verlag, 1997.

[5] A. Bondarenko, P. Dung, R. Kowalski, and F. Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93(1-2):63–101, 1997.

[6] G. Brewka. Cumulative default logic: In defense of nonmonotonic inference rules. *Artificial Intelligence*, 50(2):183–205, 1991.

[7] G. Brewka. *Nonmonotonic Reasoning: Logical Foundations of Commonsense.* Cambridge University Press, Cambridge, 1991.

[8] M. Cadoli, F. Donini, and M. Schaerf. Is intractability of nonmonotonic reasoning a real drawback. In *Proceedings of the National Conference on Artificial Intelligence*, pages 946–951. The AAAI Press/The MIT Press, 1994.

[9] S. Ceri, G. Gottlob, and L. Tanca. *Logic Programming and Databases.* Springer Verlag, 1990.

[10] P. Cholewiński. Reasoning with stratified default theories. In V. Marek and A. Nerode, editors, *Proceedings of the Third International Conference on Logic Programming and Nonmonotonic Reasoning*, pages 273–286. Springer Verlag, 1995.

[11] P. Cholewiński, V. Marek, A. Mikitiuk, and M. Truszczyński. Experimenting with nonmonotonic reasoning. In L. Sterling, editor, *Proceedings of the International Conference on Logic Programming*, pages 267–281. MIT Press, 1995.

[12] P. Cholewiński, V. Marek, and M. Truszczyński. Default reasoning system DeReS. In *Proceedings of the Fifth International Conference on the Principles of Knowledge Representation and Reasoning*, pages 518–528. Morgan Kaufmann Publishers, 1996.

[13] V. Ciorba. A query answering algorithm for łukaszewicz' general open default theory. In J. Alferes, L. Pereira, and E. Orlowska, editors, *Fifth European Workshop on Logics in Artificial Intelligence*, pages 208–223. Springer Verlag, 1996.

[14] M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7:201–215, 1960.

[15] J. Delgrande and T. Schaub. Compiling specificity into approaches to nonmonotonic reasoning. *Artificial Intelligence*, 90(1-2):301–348, 1997.

[16] J. Delgrande, T. Schaub, and W. Jackson. Alternative approaches to default logic. *Artificial Intelligence*, 70(1-2):167–237, 1994.

[17] Y. Dimopoulos and V. Magirou. A graph-theoretic approach to default logic. *Information and Computation*, 112:239–256, 1994.

[18] Y. Dimopoulos and A. Torres. Graph theoretical structures in logic programs and default theories. *Theoretical Computer Science*, 170:209–244, 1996.

[19] P. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):222–238, 1995.

[20] D. Etherington. *Reasoning with Incomplete Information: Investigations of Non-Monotonic Reasoning.* Research Notes in AI, Pitman, 1987.

[21] D. Etherington. Formalizing nonmonotonic reasoning. *Artificial Intelligence*, 31:41–85, 1987.

[22] M. Fitting, V. Marek, and M. Truszczyński. The pure logic of necessitation. *Journal of Logic and Computation*, 2:349–373, 1992.

[23] C. Froidevaux and J. Mengin. Default logic: A unified view. *Computational Intelligence*, 10(3):331–369, 1994.

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

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

[26] U. Junker and K. Konolige. Computing the extensions of autoepistemic and default logic with a TMS. In *Proceedings of the National Conference on Artificial Intelligence*, pages 278–283. The AAAI Press/The MIT Press, 1990.

[27] A. Kakas and F. Toni. Computing argumentation in logic programming. *Journal of Logic and Computation*, 9:515–562, 1999.

[28] F. Lévy. Computing extensions of default theories. In R. Kruse and P. Siegel, editors, *Proceedings of the European Conference on Symbolic and Quantitative Approaches for Uncertainty*, pages 219–226. Springer Verlag, 1991.

[29] T. Linke and T. Schaub. An approach to query-answering in reiter's default logic and the underlying existence of extensions problem. In J. Dix, L. Fariñas del Cerro, and U. Furbach, editors, *Proceedings of the Sixth European Workshop on Logics in Artificial Intelligence*, pages 233–247. Springer Verlag, 1998.

[30] T. Linke and T. Schaub. Default reasoning via blocking sets. In M. Gelfond, N. Leone, and G. Pfeifer, editors, *Proceedings of the Fifth International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'99)*, pages 247–261. Springer Verlag, 1999.

[31] T. Linke and T. Schaub. New foundations for reasoning with reiter's default logic. Technical report, Institute of Informatics, University of Potsdam, 1999.

[32] T. Linke and T. Schaub. On bottom-up pre-processing techniques for automated default reasoning. In A. Hunter and S. Parsons, editors, *Proceedings of the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 268–278. Springer Verlag, 1999.

[33] Th. Linke. *New Foundations for Automation of Default Reasoning.* Akad. Verl.-Ges. Aka 2000, Berlin, 2000.

[34] W. Łukaszewicz. Considerations on default logic — an alternative approach. *Computational Intelligence*, 4:1–16, 1988.

[35] V. Marek and M. Truszczyński. *Nonmonotonic logic: context-dependent reasoning.* Artifical Intelligence. Springer Verlag, 1993.

[36] J. McCarthy. Circumscription — a form of nonmonotonic reasoning. *Artificial Intelligence*, 13(1-2):27–39, 1980.

[37] J. Mengin. Prioritized conflict resolution for default reasoning. In *Proceedings of the European Conference on Artificial Intelligence*, pages 376–380, 1994.

[38] I. Niemelä. Towards efficient default reasoning. In C. Mellish, editor, *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 312–318. Morgan Kaufmann Publishers, 1995.

[39] C. Papadimitriou and M. Sideri. Default theories that always have extensions. *Artificial Intelligence*, 69:347–357, 1994.

[40] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–132, 1980.

[41] R. Reiter and G. Criscuolo. On interacting defaults. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 270–276, 1981.

[42] M. Richardson. Solutions of irreflexive relations. *Annals of Mathematics*, 3:573–590, 1953.

[43] V. Risch. Analytic tableaux for default logics. *Journal of Applied Non-Classical Logics*, 6(1):71–88, 1996.

[44] T. Schaub. A new methodology for query-answering in default logics via structure-oriented theorem proving. *Journal of Automated Reasoning*, 15(1):95–165, 1995.

[45] G. Schwarz and M. Truszczyński. Subnormal modal logics for knowledge representation. In *Proceedings of the National Conference on Artificial Intelligence*, pages 438–443. The AAAI Press/The MIT Press,, 1993.

[46] C. Schwind. A tableaux-based theorem prover for a decidable subset of default logic. In M. Stickel, editor, *Proceedings of the Conference on Automated Deduction*, pages 528–542. Springer Verlag, 1990.

[47] C. Schwind and V. Risch. Tableau-based characterization and theorem proving for default logic. *Journal of Automated Reasoning*, 13:223–242, 1994.

[48] A. Zhang and V. Marek. On the classification and existence of structures in default logic. *Fundamenta Informaticae*, 8(4):485–499, 1990.