ELSEVIER
Artificial Intelligence 109 (1999) 71–109
Artificial Intelligence
www.elsevier.com/locate/artint

A directly cautious theory of defeasible consequence for default logic via the notion of general extension [☆]

G. Aldo Antonelli [1]

Logic & Philosophy of Science, School of Social Sciences, University of California, Irvine, CA 92697-5100, USA

Received 17 April 1997; received in revised form 3 March 1998

Abstract

This paper introduces a generalization of Reiter's notion of "extension" for default logic. The main difference from the original version mainly lies in the way conflicts among defaults are handled: in particular, this notion of "general extension" allows defaults not explicitly triggered to pre-empt other defaults. A consequence of the adoption of such a notion of extension is that the collection of all the general extensions of a default theory turns out to have a nontrivial algebraic structure. This fact has two major technical fall-outs: first, it turns out that every default theory has a general extension; second, general extensions allow one to define a well-behaved, skeptical relation of defeasible consequence for default theories, satisfying the principles of Reflexivity, Cut, and Cautious Monotonicity formulated by D. Gabbay.

Alternative developments are also considered, namely: (i) the important case of semi-normal default theories (which turn out to have a unique minimal extension); (ii) defeasible consequence as based on extensions that are nonminimal ("optimal" in the sense of Manna and Shamir); and (iii) a variant of general extensions that, at the cost of a slight complication, avoids certain somewhat counterintuitive results.

This approach, inspired by Kripke's theory of truth, is parallel to, and generalizes over, the treatment of defeasible inheritance over cyclic networks obtained in the author's previous work; vindicates the skeptical approach to defeasible inheritance of Horty, Thomason, and Touretzky; and has similarities to the results of Fitting and Przymusiński in logic programming. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Default logic; Defeasible consequence; Truth theory; Logic programming; Nonmonotonic inheritance

---

[☆] Partial, informal expositions of the contents of this paper can be found in [2,3].

[1] Email: aldo@uci.edu.

PII: S0004-3702(99)00014-4

Introduction

The emergence of so-called *nonmonotonic logics* through the work of, among others, McCarthy [23,24], McDermott and Doyle [25], Moore [26], and Reiter [30] is undoubtedly one of the most significant developments both in logic and in artificial intelligence. Among the several formalisms proposed in the literature, Reiter's *default logic* in particular seems to provide a very flexible general-purpose framework. This paper introduces a generalization of Reiter's notion of "extension" for default logic, which can be motivated by the following considerations.

All nonmonotonic logics aim to provide a formal account of the fact that reasoners can reach conclusions tentatively, reserving the right to retract them in the light of further information. As a consequence, a certain amount of "jumping to the conclusion" is built right into any nonmonotonic framework, in the form (particularly explicit in the case of default logic) of *defeasible inference rules*. Nonmonotonicity amounts precisely to the proviso that if conflicts were to arise, one would "retract" some of the conclusions in order to restore consistency.

However, there are in principle two different kinds of conflict that can arise: (i) conflicts between tentatively endorsed conclusions and newly learned facts (strict information) about the world; and (ii) conflicts among the conclusions of the defeasible rules. All nonmonotonic frameworks handle the first kind of conflict in the same way: rules whose conclusions are inconsistent with the given information are not triggered; hence, if new facts were to be learned, old conclusions would be retracted (nonmonotonicity). On the other hand, there are in principle two different ways to handle conflicts of the second kind. According to the first strategy one can be "credulous" and always trigger a maximal subset of inference rules; i.e., one always commits to as many conclusion as possible subject to the consistency requirement. But an equally acceptable alternative is to be "skeptical" or "cautious" and withhold commitment to a conclusion in the presence of conflicting rules. (Further considerations may also play a role: for instance, in the case of inheritance networks one usually wants more "specific" information to override more "generic" information.)

The difference between the two approaches is best seen in such cases as the well-known "Nixon diamond" of Fig. 1. In the presence of two conflicting defeasible inferences (and in the absence of other considerations such as specificity), the credulous approach would have us commit to one or the other conclusion, whereas the skeptical approach would have us withhold commitment.

In the literature, conceptual reasons have been offered to prefer the cautious approach, or at least for giving it equal standing to the credulous one (see for instance [14] as a general reference). However, a directly skeptical approach has only been developed in the case of acyclic inheritance networks by Horty et al. [15]. All other approaches enforce skepticism in a roundabout way, considering the *intersection* of all credulous extensions. In particular, there are no available frameworks for default logic embodying the "cautious" or "skeptical" approach in a direct manner.

This paper introduces a notion of "general extension" for default logic that tries to capture some of the intuitions behind skepticism. Whereas conflicts with strict information are handled in the usual manner (thereby retaining the fundamental *desideratum* of

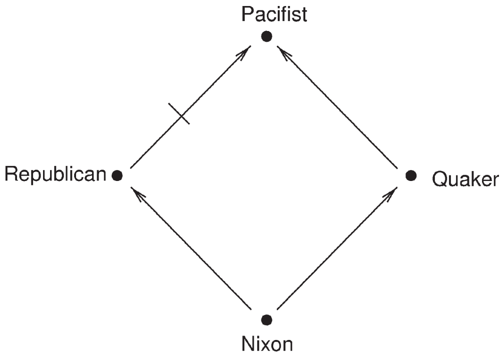

Fig. 1. The Nixon diamond.

nonmonotonicity), conflicts among defeasible rules are handled cautiously. This means, among other things, that default rules need not be explicitly triggered in order to prevent the conclusions of conflicting rules from being endorsed.

There are two major technical fall-outs of the adoption of the notion of general extension. First, according to the present approach, every default theory turns out to have a general extension. In this, default logic as based on general extensions differs from Reiter's original proposal [30] and is more in line with Łukaszewicz [18,19] and Brewka [6], both subsumed by Delgrande et al. [7] (in fact, Delgrande et al. explicitly claim their approach to provide an “amalgamation” of those of Łukaszewicz and Brewka [7, p. 198]). The basic intuition behind general extensions is however quite different from these other proposals. (The last section of the paper contains a comparison of default logic based on general extension to some of the other approaches in the literature.)

The second technical fall-out is the identification of a well-behaved relation of defeasible consequence for default theories. The abstract notion of a relation of defeasible consequence first appears with Gabbay [12]. In fact, one could argue that the whole point of having a nonmonotonic *logic* is to allow us to draw *inferences*, and this in turn requires that we have a relation $\mid\!\sim$ holding between sentences, such that $\varphi \mid\!\sim \psi$ roughly represents the fact that ψ is a defeasible or nonmonotonic consequence of φ. Gabbay [12] identifies certain principles that seem necessary for such a relation to be well-behaved, most notably the following:

(1) *Reflexivity*: $\varphi \mid\!\sim \varphi$;

(2) *Stability* or *Cautious Monotonicity*: $\dfrac{\varphi \mid\!\sim \psi,\ \varphi \mid\!\sim \chi}{\varphi \wedge \psi \mid\!\sim \chi}$;

(3) *Cut*: $\dfrac{\varphi \mid\!\sim \psi,\ \varphi \wedge \psi \mid\!\sim \chi}{\varphi \mid\!\sim \chi}$.

As observed by Makinson [20] the principle of Cautious Monotonicity is the inverse of Cut, a fact which is hidden in the case of classical logic in which Cautious Monotonicity is subsumed by a more general principle of Weakening.

As borne out by the analysis of Makinson [20], there seems to be no natural way using the notions of extension available in the literature to define a reasonably well-behaved (i.e., satisfying the above three principles) relation of defeasible consequence for default logic (skeptical or otherwise). The only plausible route appears to be the already mentioned definition of a relation $\mid\!\sim$ obtaining between a default theory and a proposition φ precisely when φ is supported by *all* extensions of the theory. In other words, φ is defeasibly inferable from a given default theory if it occurs in the intersection of all its extensions.

However, even discounting feasibility worries and other conceptual issues, there appears to be a major problem with this approach. Makinson [20, p. 60] shows that defining $\mid\!\sim$ as the intersection of all the extension (in Reiter's sense) of a default theory, one cannot satisfy Cautious Monotonicity. This is a substantial drawback.

It is then significant that using general extensions, the relation $\mid\!\sim$ defined as the intersection of the extensions, turns out to have all three properties identified by D. Gabbay. Moreover, in the important special case of semi-normal default theories, we have the existence of a unique privileged extension of the theory (the *least* extension), which makes the definition of $\mid\!\sim$ particularly simple. Admittedly, this can have, at least in some cases, some counter-intuitive consequences (see Section 4). While we argue that, once the cautious motivation of the theory is properly taken into account, such consequences are less counter-intuitive than it might appear at first, we also review alternative ways to identify a privileged extension for a default theory in such a way as to obviate some of the counter-intuitive consequences.

The paper is organized as follows. In Section 1 we illustrate the notion of general extension for the particularly simple, but already significative case of categorical (i.e., prerequisite-free) default theories. In Section 2 we take up several examples of categorical theories. In Section 3 we introduce the notion of general extension for arbitrary default theories, presenting some examples in Section 4. In Section 5 we show how to define a relation of defeasible consequence having Gabbay's three properties. Section 6 presents alternative developments including how to modify the definitions to overcome a certain degree of counter-intuitiveness. Finally, in Section 7, we draw conclusions and comparisons to other approaches. Appendix A contains proofs of selected theorems and Appendix B sketches how to give a "transfinite" version of the present approach.

1. Categorical default theories

Fix a propositional language $\mathcal{L}$, obtained from an infinite list of propositional variables $p, q, r, \ldots$ by the propositional connectives $\wedge, \vee, \neg, \rightarrow$. The language $\mathcal{L}$ is endowed with the usual (i.e., two-valued) semantics. In particular, we use $\models$ to denote the usual relation of logical consequence. If S and W are sets of propositional axioms from $\mathcal{L}$ and φ a formula from $\mathcal{L}$, then we write $S \models_W \varphi$ as shorthand for $S \cup W \models \varphi$. A set S of sentences is W-consistent if and only if $S \not\models_W p \wedge \neg p$.

A *default theory* is a pair (W, Δ), where W is a set of propositional axioms (a "world description") and Δ is a *finite* set of defaults. In turn, a *default* δ is an expression of the form

$$\frac{\zeta : \eta}{\theta},$$

where ζ, η, θ are sentences from $\mathcal{L}$. The intuitive meaning of a default δ is that *if* ζ *has been established, and* $\neg\eta$ *has not been established, then assume* θ ("*by default*"). The expressions ζ, η, θ are called the *pre-requisite*, the *justification*, and the *conclusion* of δ, respectively. If

$$\delta = \frac{\zeta : \eta}{\theta},$$

then we put $P(\delta) = \zeta$, $J(\delta) = \eta$, and $C(\delta) = \theta$. If Γ is a set of defaults, we will write $C(\Gamma)$ for $\{C(\delta)\colon\ \delta \in \Gamma\}$, and similarly for $P(\Gamma)$ and $J(\Gamma)$.

In the particular case in which a default δ has no pre-requisite (because it is of the form

$$\frac{: \eta}{\theta}$$

or, equivalently, because $P(\delta)$ is a tautology), then we say that δ is *categorical*. If all defaults in Δ are categorical, then we say that (W, Δ) is categorical. As we will see, categorical default theories form a natural and well-behaved class. Other special cases: a default δ is *normal* if $J(\delta)$ and $C(\delta)$ are logically equivalent, and it is *semi-normal* if $J(\delta)$ logically implies $C(\delta)$ (typically, because $J(\delta)$ is a conjunction one of whose conjuncts is $C(\delta)$).

Definition 1.1. Let S be a set of $\mathcal{L}$-sentences, W a world description, and δ a default. Then we say:

(1) δ is *admissible* in S (relative to W) if and only if $S \models_W P(\delta)$;
(2) δ is *conflicted* in S (relative to W) if and only if $S \models_W \neg C(\delta)$;
(3) δ is *pre-empted* in S (relative to W) if and only if $S \models_W \neg J(\delta)$.

If Γ is a set of defaults, we say that δ is admissible, conflicted, or pre-empted in Γ according as δ is admissible, pre-empted or conflicted in $C(\Gamma)$.

Reiter [30] defines the notion of an *extension* for a default theory. Intuitively, an extension for a default theory (W, Δ) is a deductively closed, consistent set of formulas containing W and the consequents of a maximal subset of Δ. Here we introduce the slightly simpler notion of a *classical extension* for a default theory. Similar simplifications have been advocated in [5,8,27].

Definition 1.2. A set Γ of defaults is a *classical extension* for (W, Δ) if and only if it satisfies $\Gamma = \bigcup_{n \geqslant 0} \Gamma_n$, where $\Gamma_0 = \emptyset$, and

$$\Gamma_{n+1} = \bigl\{\delta \in \Delta\colon\ C(\Gamma_n) \models_W P(\delta) \;\&\; C(\Gamma) \not\models_W \neg J(\delta)\bigr\};$$

(notice the occurrence of Γ in the definition of Γ_{n+1}). In other words, Γ_{n+1} is the set of defaults admissible in Γ_n that are not pre-empted in Γ.

Observe that if Γ is a classical extension for (W, Δ) then $\{\varphi\colon\ C(\Gamma) \models_W \varphi\}$ is an extension in Reiter's sense. In general, a default theory might have zero, one, or more than one extension.

We record here the following well-known facts about classical extensions.

Theorem 1.3. *Let Γ be a classical extension for (W, Δ).*

(1) *$C(\Gamma)$ is W-inconsistent if and only if W is inconsistent.*

(2) *If δ is admissible in Γ, but neither conflicted nor pre-empted in Γ, then $\delta \in \Gamma$.*

We now proceed to generalize this notion of extension, beginning with the somewhat simpler particular case of categorical theories.

Let (W, Δ) be a categorical default theory. Recall that this means that the defaults in Δ have no pre-requisite. In the general case of a default theory, extensions have to be defined in the somewhat peculiar form of Definition 1.2, since extensions have to be "grounded", where a classical extension Γ is grounded if it is a minimal set of defaults admissible but not pre-empted in Γ. Definition 1.2 is formulated in such a way as to enforce this form of minimality by letting no more defaults become admissible than "have to". When a theory is categorical groundedness is no longer a concern, and it is immediate to check that Γ is a classical extension for a categorical theory (W, Δ) if and only if it is a solution to the following fixpoint equation:

$$\Gamma = \{\delta \in \Delta\colon\ C(\Gamma) \not\models_W \neg J(\delta)\}.$$

In other words, Γ contains all and only those defaults that are not pre-empted in Γ. We now set out to generalize this notion.

Definition 1.4. A *general extension* for a categorical default theory (W, Δ) is a pair (Γ^+, Γ^-) of disjoint sets of defaults from Δ, simultaneously satisfying the following two fixpoint equations:

$$\begin{aligned}\Gamma^+ &= \{\delta\colon\ C(\Gamma^+) \not\models_W \neg C(\delta)\ \&\ C(\Delta - \Gamma^-) \not\models_W \neg J(\delta)\};\\ \Gamma^- &= \{\delta\colon\ C(\Gamma^+) \models_W \neg C(\delta) \text{ or } C(\Gamma^+) \models_W \neg J(\delta)\}.\end{aligned}$$

In other words, Γ^+ is the set of all defaults that are neither conflicted in Γ^+ nor pre-empted in $(\Delta - \Gamma^-)$, while Γ^- is the set of all defaults that are either conflicted or pre-empted in Γ^+.

Observe that the components Γ^+ and Γ^- of the extension are required to be disjoint. This corresponds to the intuition that Γ^+ is the set of the (definitely) triggered defaults and Γ^- is the set of the (definitely) excluded defaults, and that no default can be both explicitly triggered and explicitly ruled out. From a purely technical point of view, it is possible to develop the theory without such a disjointness requirement, and this would correspond to a "four-valued" intuition such as that of Belnap [4]. We do not purse this possibility here.

In the next two theorems we show that general extensions indeed generalize the notion of classical extension, and that general extensions always exist. Proofs are given in Appendix A.

Theorem 1.5. *Let Γ^+ be a classical extension for a categorical default theory (W, Δ), and put*

$$\Gamma^- = \{\delta\colon\ \delta \textit{ conflicted or pre-empted in } \Gamma^+\},$$

then (Γ^+, Γ^-) is a general extension for (W, Δ).

Theorem 1.6. *Every categorical default theory has a general extension.*

As mentioned, a full proof of the theorem can be found in the Appendix A. It is, however, interesting intuitively to characterize the process by which such an extension can be obtained. A general extension (Γ^+, Γ^-) can be constructed "from below" in stages, as the limit of the sequences Γ_n^+ and Γ_n^- (for $n \geqslant 0$). For the starting point, we put $\Gamma_0^+ = \Gamma_0^- = \emptyset$.

For the inductive step, let Γ_{n+1}^+ be a *maximal* set of defaults from Δ such that: (i) the conclusions of defaults in $\Gamma_n^+ \cup \Gamma_{n+1}^+$ form a consistent set; and (ii) no default $\delta \in \Gamma_{n+1}^+$ is pre-empted in $\Delta - \Gamma_n^-$. (Observe that such a maximal subset need not be unique, so that the process is not deterministic—this will be important for the definition of the relation $\mid\!\sim$.) Also put

$$\Gamma_{n+1}^- = \{\delta\colon\ \delta \text{ conflicted or pre-empted in } \Gamma_{n+1}^+\}.$$

The sequences are increasing and disjoint, and their limit gives a general extension.

The existence of extensions is a desirable feature that is missing in Reiter's original formulation, but that can be found for instance in all three proposals of Łukaszewicz [18,19], Brewka [6], and Delgrande et al. [7]. But as the examples discussed in the following section will make clear, the intuitions at the basis of these proposals are different from the present one, and the formal mathematical properties of the frameworks correspondingly different.

2. Examples

In this section we take up a few examples in order to compare the notion of general extension given here to other notions available in the literature. As before, "classical extension" here refers to the notion of extension in the sense of Reiter [30], as simplified in Definition 1.2.

Example 2.1. Consider first the default theory (W, Δ), where W is empty and Δ comprises the two defaults:

$$\frac{:p}{\neg q} \quad \text{and} \quad \frac{:q}{\neg p}.$$

As is well known, this theory has two classical extensions, according to which default is triggered. In addition to these, the theory has one general extension, in which no default is triggered, and none is ruled out.

Example 2.2. Consider the default theory in which W is empty, and Δ comprises only the default

$$\frac{:p}{\neg p}.$$

This theory has no classical extensions, but it has one general extension, namely, $(\emptyset, \emptyset)$, in which the default is neither triggered nor ruled out.

Example 2.3. Consider now the theory where W is empty, but Δ comprises the defaults

$$\frac{:p}{\neg p} \quad \text{and} \quad \frac{:q}{r}.$$

As before this theory has no classical extensions, because of the first default. It does have one general extension, namely

$$\left(\frac{:q}{r}, \emptyset\right).$$

The first default cannot be triggered, but there is no obstacle that prevents triggering the second.

Example 2.4. Consider a different sort of theory, in which W is empty and $\Delta = \{\delta_1, \ldots, \delta_n\}$, for some $n > 1$. Suppose also that for all k such that $1 \leqslant k < n$,

$$\delta_k = \frac{:p_k}{\neg p_{k+1}},$$

whereas

$$\delta_n = \frac{:p_n}{\neg p_1}.$$

So defaults are organized in a “loop”, in which the conclusion of each one of them pre-empts the next default, and the conclusion of the last one pre-empts the first. It is easy to check that this theory has no classical extensions if n is odd. (It has, of course one general extension even when n is odd, namely $(\emptyset, \emptyset)$.) On the other hand, if n is even, say $n = 2m$, beside the general extension $(\emptyset, \emptyset)$, there are two classical extensions, namely

$$\bigl(\{\delta_{2k+1}\colon 0 \leqslant k < m\}, \{\delta_{2k+2}\colon 0 \leqslant k < m\}\bigr),$$

and the one with the two components switched around,

$$\bigl(\{\delta_{2k+2}\colon 0 \leqslant k < m\}, \{\delta_{2k+1}\colon 0 \leqslant k < m\}\bigr).$$

Indeed, this result is consistent with similar results in [1,27].

So far we have only been concerned with existence of extensions, but this is by no means the only question that can be asked when a notion of extension is proposed. In the remaining part of this section we take up a few “benchmark” examples, mostly drawn from [7], in order to assess the behavior of the proposed notion.

Example 2.5 (*Semi-monotonicity*). Consider the default theory (W_1, Δ_1), where $W_1 = \emptyset$, and Δ_1 comprises:

$$\frac{:p \wedge \neg q}{p}.$$

Theory (W_2, Δ_2) is obtained by putting $W_2 = W_1$, and by adding to Δ_1 the default

$$\frac{:q}{q}.$$

Semi-monotonicity is a property first singled out by Reiter. According to this property the extensions of a theory increase monotonically with the size of Δ: if Γ is an extension for (W, Δ) and $\Delta \subseteq \Delta'$, then there is an extension Γ' for (W, Δ') such that $\Gamma \subseteq \Gamma'$.

By considering the above example, we see that semi-monotonicity fails for the notion of general extension of Section 3. The theory (W_1, Δ_1) has one general extension, triggering its unique default. When the second default is thrown in, however, we lose the first one, in the sense that the unique extension of the theory triggers the second default and rules out the first (the first default is now potentially pre-empted). In this, the version of default logic presented here agrees with Reiter's Default Logic; by contrast, Constrained Default Logic [7] is semi-monotonic.

Semi-monotonicity is regarded as a desirable property, because it is a form of locality: extensions for large theories can be approximated by forming extensions for smaller ones. Reiter's Default Logic is semi-monotonic in the case of normal theories, in which case we can apply the approximating process just mentioned. Although general extensions fail to give a semi-monotonic framework for default logic, they still allow us to construct extensions "from below", although in a different sense. Moreover, in the case of semi-normal theories, as we will see, the extension thus obtained is unique.

Example 2.6 (*Weak orthogonality of extensions*). Consider the default theory with $W = \emptyset$ and Δ comprising the defaults:

$$\frac{:p \wedge \neg q}{\neg q} \quad \text{and} \quad \frac{:\neg p \wedge q}{\neg p}.$$

Orthogonality is the property that any two distinct extensions are inconsistent. Delgrande et al. [7] propose a similar notion, more appropriate in the context of Constrained Default Logic, i.e., Weak Orthogonality: this is the property according to which any two sets of constraints corresponding to distinct extensions are inconsistent.

Although weak orthogonality seems an appropriate feature in the case of the family of logics inspired by "commitment to the assumptions" (such as [6,7,18,19]), it appears to be less so in the present context. Indeed, it is possible to argue that (weak) orthogonality, or even maximality of extensions is not a desirable feature given the "three-valued" intuition at the basis of the notion of general extension. In fact, orthogonality fails for general extensions, for the above theory has two maximal (mutually inconsistent) extensions, whose intersection is however again an extension.

Example 2.7 (*Commitment to assumptions*). This is the "broken arms" example of Poole [28]. Consider the theory (W, Δ) where $W = \{\neg q \vee \neg s\}$, and Δ comprises the defaults:

$$\frac{:p \wedge q}{p} \quad \text{and} \quad \frac{:r \wedge s}{r}.$$

Here the intuitive interpretation is as follows: we are told that either the left arm is broken or the right arm is broken $(\neg q \vee \neg s)$; the first default asserts that a person's left arm is usable (p) unless it is broken $(\neg q)$, while the second default asserts that a person's right arm is usable (r) unless it is broken $(\neg s)$.

Reiter's notion of extension here gives that both arms are usable, although we are told that at least one is broken. There is in fact a unique extension in Reiter's sense in which both defaults are fired, the justification of each one of them being consistent both with what is known and with the conclusions of fired defaults. The situation is essentially the same with the notion of general extension.

As Delgrande et al. [7] point out, the problem here is that the consistency of each justification is tested individually, and not—as it were—wholesale. Indeed, this provides motivation for the adoption of the second construal of prerequisites we mentioned, i.e., as working hypotheses rather than mere consistency conditions. However, if we do adopt the first (an equally viable option, as Delgrande et al. [7] admit), then this example loses much of its intuitive punch.

3. Grounded extensions

In this section we show how to generalize the results of the previous section to default theories that might not be categorical. The treatment in this section will reveal that prerequisites play an unexpectedly subtle role, in more than one way. It will also turn out that the definitions and theorems in this section are somewhat more complicated than their analogues relative to categorical default theories. Accordingly, to simplify matters somewhat, and we no loss in generality, we restrict our attention to default theories (W, Δ) for which W is propositionally consistent.

Definition 3.1. Let (W, Δ) be a default theory and $\Gamma \subseteq \Delta$. Then $\mathsf{Ad}(\Gamma)$ is the set of defaults (from Δ) that are *admissible* in Γ:

$$\mathsf{Ad}(\Gamma) = \{\delta\colon C(\Gamma) \models_W P(\delta)\}.$$

We are going to be particularly interested in sets Γ of defaults such that $\Gamma \subseteq \mathsf{Ad}(\Gamma)$, i.e., every $\delta \in \Gamma$ is admissible in Γ. Recall that classical extension for arbitrary default theories are required to be grounded. We now isolate and justify this notion of groundedness.

Definition 3.2. A set Θ of defaults is *grounded* if and only if $\Theta = \bigcup_{n \geqslant 0} \Theta_n$, where $\Theta_0 = \emptyset$, and

$$\Theta_{n+1} = \{\delta \in \Theta\colon C(\Theta_n) \models_W P(\delta)\}.$$

I.e., Θ_{n+1} is the set of all defaults from Θ that are admissible in Θ_n.

It is clear that if Γ is grounded then $\Gamma \subseteq \mathsf{Ad}(\Gamma)$. Notice that for any default theory (W, Δ) there will be several grounded sets $\Theta \subseteq \Delta$. But just because Θ is grounded does not mean that it is a set of defaults that we need to accept (that are triggered in some extension): groundedness provides only a (weak) necessary condition: classical extensions are grounded in the above sense (proof in Appendix A).

Theorem 3.3. *Let Γ be a classical extension for a default theory (W, Δ); then Γ is grounded.*

We are now ready to give the definition of general extensions for arbitrary default theories. We will identify such extensions with triples $(\Gamma^+, \Gamma^-, \Gamma^*)$ of sets of defaults, in which Γ^+ is the set of defaults explicitly triggered, Γ^- is the set of defaults explicitly conflicted in Γ^+ or pre-empted in Γ^+, and Γ^* (the set of "potentially admissible" defaults) expresses the *degree of caution* of the extension. In general, the larger Γ^*, the smaller Γ^+, so that the size of Γ^* is proportional to the *caution* of Γ^+ and hence it is inversely related to its credulousness. In the formal development, as we will see (Theorem 3.8), this is reflected in the fact that general extensions can be obtained by an iterative process which is *monotonic* in Γ^+ and Γ^-, but *anti-monotonic* in Γ^*.

Definition 3.4. A *general extension* for a default theory (W, Δ) is a triple $(\Gamma^+, \Gamma^-, \Gamma^*)$ of sets of defaults from Δ, such that:

- Γ^+ and Γ^- are disjoint;
- the following two fixpoint equations are simultaneously satisfied:

$$\Gamma^+ = \bigl\{\delta\colon\ C\bigl(\Gamma^+\bigr) \models_W P(\delta)\ \&\ C\bigl(\Gamma^+\bigr) \not\models_W \neg C(\delta)\ \& \\ C\bigl(\Gamma^* - \Gamma^-\bigr) \not\models_W \neg J(\delta)\bigr\};$$

$$\Gamma^- = \bigl\{\delta\colon\ C\bigl(\Gamma^+\bigr) \models_W \neg C(\delta) \text{ or } C\bigl(\Gamma^+\bigr) \models_W \neg J(\delta)\bigr\};$$

- $\mathsf{Ad}(\Gamma^+) \subseteq \Gamma^* \subseteq \{\delta\colon\ C(\Gamma^+) \not\models_W \neg P(\delta)\}$.

In words: Γ^+ is the set of all defaults admissible in Γ^+ but neither conflicted in Γ^+ nor pre-empted in $(\Gamma^* - \Gamma^-)$; Γ^- is the set of all defaults either conflicted or pre-empted in Γ^+; and Γ^* is *a* set of defaults containing all default admissible in Γ^+ and only defaults whose pre-requisites are consistent with $C(\Gamma^+)$ (the latter condition is intended to capture the fact that δ is "potentially" admissible in Γ^+).

From the definition it follows immediately that if $(\Gamma^+, \Gamma^-, \Gamma^*)$ is a general extension, every default δ in Γ^+ is admissible in Γ^+, i.e., $\Gamma^+ \subseteq \mathsf{Ad}(\Gamma^+)$. Recall that we restricted out attention to default theories (W, Δ) with W propositionally consistent: we can see now that nothing is lost by that assumption, for if W is propositionally inconsistent, then $(\emptyset, \Delta, \emptyset)$ is the only extension of the theory. The following theorem shows that this definition of general extension coincides with Definition 1.4 in the case of categorical theories.

Theorem 3.5. *Let (W, Δ) be a categorical default theory and (Γ^+, Γ^-) an extension in the sense of Definition* 1.4. *Then $(\Gamma^+, \Gamma^-, \Delta)$ is an extension in the sense of Definition* 3.4.

We are going to be interested in extensions that are *minimal* according to the ordering defined below.

Definition 3.6. Let $(\Gamma^+, \Gamma^-, \Gamma^*)$ and $(\Theta^+, \Theta^-, \Theta^*)$ be 3-tuples of sets of defaults. Define an ordering by putting $(\Gamma^+, \Gamma^-, \Gamma^*) \leqslant (\Theta^+, \Theta^-, \Theta^*)$ iff:

$$\Gamma^+ \subseteq \Theta^+ \quad \text{and} \quad \Gamma^- \subseteq \Theta^- \quad \text{and} \quad \Theta^* \subseteq \Gamma^*.$$

We now state the analogues of Theorems 1.5 and 1.6 (proofs in Appendix A).

Theorem 3.7. *Let* $\Gamma^+ = \bigcup_{n \geqslant 0} \Gamma_n^+$ *be a classical extension for a default theory* (W, Δ). *Put*:

$$\Gamma^- = \{\delta\colon \delta \text{ conflicted or pre-empted in } \Gamma^+\};$$
$$\Gamma^* = \{\delta\colon \delta \text{ admissible in } \Gamma^+\} = \mathsf{Ad}(\Gamma^+).$$

Then $(\Gamma^+, \Gamma^-, \Gamma^*)$ *is a general extension for* (W, Δ).

Theorem 3.8. *Let* (W, Δ) *be a default theory. Then*:
(i) *every default theory has an iteratively definable general extension*;
(ii) *such an extension is* $\leqslant$*-minimal*; *and*
(iii) *any minimal extension for* (W, Δ) *is grounded.*

Theorem 3.9. *Let* $(\Gamma^+, \Gamma^-, \Gamma^*)$ *be a* $\leqslant$*-minimal extension for a default theory* (W, Δ). *Then* $(\Gamma^+, \Gamma^-, \Gamma^*)$ *can be obtained as the limit of an inductive construction of the kind given in the proof of Theorem* 3.8.

4. Examples, continued

Example 4.1 (*Cumulativity*). Consider the theory (W, Δ), where $W = \emptyset$ and Δ comprises the two defaults:

$$\delta_1 = \frac{:p}{p} \quad \text{and} \quad \delta_2 = \frac{p \vee q : \neg p}{\neg p}.$$

Cumulativity is the property that if a default is triggered in all the extensions of a theory, adding that default's conclusion to the world description W should give a theory that has the same extensions as the original one. This is related to the property of cautious monotonicity (if a "theorem" is added back to the set of facts from which it was deduced, then the set of "theorems" should not change), but it is not equivalent to it, as we will see.

In the above example, using Reiter's notion of extension, cumulativity fails: the above theory has one extension Γ triggering the first default, therefore $C(\Gamma) \models_W p \vee q$; but the second default cannot be triggered because pre-empted. Now consider the theory $(W \cup \{p \vee q\}, \Delta)$. This modified theory gains an extension, in which the second default rather than the first is triggered.

Let us see what happens with general extensions. The theory (W, Δ) has two general extensions: a minimal one, $(\emptyset, \emptyset, \Delta)$ and a nonminimal one, $(\delta_1, \delta_2, \delta_2)$, in which δ_1 is triggered and δ_2 ruled out.

When we add $p \vee q$ back into the world description W, we obtain a new extension, $(\delta_2, \delta_1, \delta_1)$, triggering the second default and ruling out the first. Therefore, cumulativity fails also for the notion of general extension. Notice that this counterexample to cumulativity is also exactly the counterexample used by Makinson [20] to establish that defeasible consequence defined as intersection of classical extensions is not cautiously monotonic. We will see, however, that this does not preclude Cautious Monotonicity from holding when defeasible consequence is defined as intersection of general extensions, since adding the "theorem" $p \vee q$ back into the world description W gives *new* extensions, but no new *minimal ones*. In particular, although the new theory has more extensions than the old one, the *intersection* of the extensions (i.e., the set defaults triggered in *every* extension) is the same.

It is worth noting here that Brewka's Cumulative Default Logic is, of course, cumulative, whereas for Delgrande's et al. Constrained Default Logic cumulativity fails except in the categorical (prerequisite-free) case.

Example 4.2. Consider a default theory (W, Δ), where W is empty and Δ comprises the two defaults

$$\delta_1 = \frac{: p}{p} \quad \text{and} \quad \delta_2 = \frac{q : \neg p}{\neg p}.$$

This example is particularly interesting, since it highlights the intuitions behind the present approach and the extent to which it delivers counter-intuitive results. The point is that in the (unique) minimal extension of (W, Δ), neither default is triggered and neither is ruled out. This is due to the basic "skeptical" intuition behind general extensions: defaults that are *potentially* pre-empted are not triggered. In the minimal extension, δ_1 is potentially pre-empted by $\neg p$ (the conclusion of δ_2), and δ_2 is potentially pre-empted by p (the conclusion of δ_1). The way things are set up, it does not matter that there is no way δ_2 will ever be triggered, since its pre-requisite q is not entailed by the conclusions of any of the defaults in Δ.

To some extent this state of affairs is unsatisfying: since δ_2 cannot possibly be triggered, it should not be allowed to prevent the triggering of δ_1. Let us notice that there seems to be different intuitions at work here: on the one hand there is the idea that "potentially pre-empted" defaults should not be triggered, on the other hand there is the intuition that defaults that are not even potentially admissible should be ruled out. In some cases, such as the present one, these intuitions appear to be in conflict.

Notice that the theory also has a *nonminimal* extension that avoids the problem by triggering δ_1 and (consequently) ruling out δ_2. For the time being, this is where we are going to leave the matter. We will come back to this question in Section 6, and indicate alternative ways in which this problem can be solved.

Example 4.3. Consider the default theory (W, Δ) with $W = \{p\}$ and Δ comprising the following defaults:

$$\begin{aligned}
\delta_1 &= \frac{p : q}{q}, \\
\delta_2 &= \frac{q : q}{r}, \\
\delta_3 &= \frac{q : q}{\neg r}, \\
\delta_4 &= \frac{: \neg q}{s}, \\
\delta_5 &= \frac{: \neg s}{t}.
\end{aligned}$$

This example is tended to highlight features of general extensions that are not brought out by the previous examples. Let us consider the different defaults. Default δ_1 is always triggered: its pre-requisite is entailed by W, and it cannot be pre-empted since the negation of its justification q is not entailed by the conclusions of any of the defaults. Things are different for the pair of defaults δ_2 and δ_3: once δ_1 is triggered, they are both admissible, and neither is (potentially) pre-empted. However, they cannot both be triggered since they have contradictory conclusions. This state of affairs is a “diamond”, analogous to the one of Fig. 1. At the diamond, extensions split. Finally consider the two defaults δ_4 and δ_5: the latter cannot be triggered initially, because potentially pre-empted by δ_4. However, as soon as δ_1 is triggered, δ_4 is ruled out (because pre-empted), removing any obstacles to the triggering of δ_5. In conclusion, the theory (W, Δ) has two minimal extensions, one extension triggering δ_1, δ_2, and δ_5; and the other extension triggering δ_1, δ_3, and δ_5. In the first extension δ_3 and δ_4 are ruled out and in the second δ_2 and δ_4 are ruled out.

5. Defeasible consequence

First, it will be convenient to introduce an abbreviated notation for extensions: we will use boldface uppercase Greek letters to stand for triples of sets of defaults, as in $\boldsymbol{\Gamma} = (\Gamma^+, \Gamma^-, \Gamma^*)$. Similarly, given sequences of sets of defaults Γ_n^+, Γ_n^-, and Γ_n^*, we write $\boldsymbol{\Gamma}_n$ for $(\Gamma_n^+, \Gamma_n^-, \Gamma_n^*)$. If $\boldsymbol{\Gamma}_n$ is a sequence of sets of defaults ($n \geqslant 0$), we write $\lim \boldsymbol{\Gamma}_n$ for the unique $\boldsymbol{\Gamma} = (\Gamma^+, \Gamma^-, \Gamma^*)$ such that:

$$\Gamma^+ = \bigcup_{n \geqslant 0} \Gamma_n^+, \qquad \Gamma^- = \bigcup_{n \geqslant 0} \Gamma_n^-, \qquad \Gamma^* = \bigcap_{n \geqslant 0} \Gamma_n^*.$$

Finally, to simplify notation still further we write just $W + \varphi$ in place of $W \cup \{\varphi\}$.

Now we take up a property of general extensions closely related to cumulativity. Appendix A contains a proof of the following theorem.

Theorem 5.1. *Let (W, Δ) be a default theory, and suppose $\boldsymbol{\Gamma}$ is a 3-tuple of subsets of Δ such that $C(\Gamma^+) \models_W \varphi$. Then $\boldsymbol{\Gamma}$ is a general extension for (W, Δ), if and only if it is a general extension for $(W + \varphi, \Delta)$.*

The above result shows that any extension $(\Gamma^+, \Gamma^-, \Gamma^*)$ remains such when we adjoin one of its consequences to the world description. This by itself does not tell us anything about extensions other than $(\Gamma^+, \Gamma^-, \Gamma^*)$. In particular, it does not establish cumulativity. Recall that cumulativity is the property that if φ is supported in every extension of a default theory (W, Δ), then the theory $(W + \varphi, \Delta)$ has the same extensions as (W, Δ). Cumulativity fails for general extensions just as for classical extensions (see Example 4.1). But this, as we will see, still allows us to define a cautiously monotonic relation of defeasible consequence based on general extensions. It is perhaps worth recalling that Example 4.1 is precisely Makinson's counterexample to cautious monotony for defeasible consequence based on classical extensions.

We also record here that Cumulativity fails for Constrained Default Logic, which is, on the other hand, both Semi-monotonic and Orthogonal. A comparison of General Extensions and Constrained Default Logic is summarized in Fig. 2.

For the purposes of this section, in the light of the characterization of minimal extensions given in Theorem 3.9, we refer to a sequence $\boldsymbol{\Gamma}_n$ as a *construction sequence* for a theory (W, Δ), if the following conditions are met:

(1) $\Gamma_0^+ = \Gamma_0^- = \emptyset$;
(2) Γ_0^* is the set of all defaults from Δ whose pre-requisite is W-consistent;
(3) Γ_{n+1}^+ is a maximal set of defaults such that:
  (a) $C(\Gamma_n^+ \cup \Gamma_{n+1}^+)$ is consistent;
  (b) every $\delta \in \Gamma_{n+1}^+$ is admissible in Γ_n^+;
  (c) no default $\delta \in \Gamma_{n+1}^+$ is pre-empted in $\Gamma_n^* - \Gamma_n^-$;
(4) $\Gamma_{n+1}^- = \{\delta\colon \delta$ pre-empted or conflicted in $\Gamma_{n+1}^+\}$;
(5) $\Gamma_{n+1}^* = \{\delta\colon C(\Gamma_{n+1}^+) \not\models_W \neg P(\delta)\}$.

Definition 5.2. Let (W, Δ) be a default theory. Then for any sentence φ, we say that φ is a *defeasible consequence* of (W, Δ), written $(W, \Delta) \mid\!\sim \varphi$, if and only if $C(\Gamma^+) \models_W \varphi$, whenever $\boldsymbol{\Gamma}$ is a $\leqslant$-minimal extension of (W, Δ).

So we define $\mid\!\sim$ skeptically by taking the intersection of all (minimal) extensions of the theory, i.e., all extensions that can be obtained by an inductive process of the same kind of the one given in the proof of Theorem 3.8. The following then is an immediate consequence of the definitions and Theorem 3.8.

Theorem 5.3. *If $(W, \Delta) \mid\!\sim \varphi$, then for every construction sequence $\boldsymbol{\Gamma}_0, \boldsymbol{\Gamma}_1, \ldots$ there is n such that $C(\Gamma_n^+) \models_W \varphi$.*

	Cumulative	Semi-monotonic	Orthogonal	Cautiously monotonic
Gen Ext	No	No	No	Yes
ConDL	No	Yes	Yes	?

Fig. 2. Comparison of General Extensions and Constrained Default Logic.

We need to verify that $\mathrel{|\!\sim}$ is indeed a well-behaved notion of defeasible consequence, in that it satisfies the conditions of Reflexivity, Stability, and Cut. However, such properties were formulated for a relation $\mathrel{|\!\sim}$ that holds between sentences, whereas the relation defined above holds between a default theory and a sentence. Accordingly, we need to reformulate such properties in order to take this fact into account.

(1) *Reflexivity*: if $\varphi \in W$ then $(W, \Delta) \mathrel{|\!\sim} \varphi$;

(2) *Stability* or *Cautious Monotonicity*: $\dfrac{(W, \Delta) \mathrel{|\!\sim} \varphi,\ (W, \Delta) \mathrel{|\!\sim} \psi}{(W + \varphi, \Delta) \mathrel{|\!\sim} \psi}$;

(3) *Cut*: $\dfrac{(W, \Delta) \mathrel{|\!\sim} \varphi,\ (W + \varphi, \Delta) \mathrel{|\!\sim} \psi}{(W, \Delta) \mathrel{|\!\sim} \psi}$.

Theorem 5.4. *Suppose* $(W, \Delta) \mathrel{|\!\sim} \varphi$ *and let* $\boldsymbol{\Gamma} = \lim \boldsymbol{\Gamma}_n$ *be a minimal extension for* $(W + \varphi, \Delta)$*. Then there is a minimal extension* $\boldsymbol{\Theta} = \lim \boldsymbol{\Theta}_n$ *for* (W, Δ) *such that* $\boldsymbol{\Gamma} \leqslant \boldsymbol{\Theta}$.

This theorem shows that although minimal extensions for (W, Δ) start out more slowly than extensions for $(W + \varphi, \Delta)$, they eventually “catch up”. Of course, the proof of Theorem 5.4 depends crucially on the minimality of $\boldsymbol{\Gamma}$.

Theorem 5.5. *Suppose* $(W, \Delta) \mathrel{|\!\sim} \varphi$ *and let* $\boldsymbol{\Gamma} = \lim \boldsymbol{\Gamma}_n$ *be a minimal extension for* $(W + \varphi, \Delta)$*. Then there is a minimal extension* $\boldsymbol{\Theta} = \lim \boldsymbol{\Theta}_n$ *for* (W, Δ) *such that* $\boldsymbol{\Theta} \leqslant \boldsymbol{\Gamma}$.

Theorem 5.5 is the converse of the previous Theorem 5.4 and, in a sense, the crucial result in the paper. Theorem 5.5 is what is needed to establish Cautious Monotonicity for $\mathrel{|\!\sim}$ and, characteristically, this is where the mathematics itself becomes interesting. The reader is referred to the Appendix A for the proof: here we only note that in a somewhat unexpected twist, the proof of Theorem 5.5 makes use of Theorem 5.4.

Notice that if $(W, \Delta) \mathrel{|\!\sim} \varphi$ then (by the two theorems above) every minimal extension $\boldsymbol{\Gamma}$ for $(W + \varphi)$ is between two minimal extensions $\boldsymbol{\Theta}$ and $\boldsymbol{\Pi}$ for (W, Δ), i.e., $\boldsymbol{\Theta} \leqslant \boldsymbol{\Gamma} \leqslant \boldsymbol{\Pi}$. By minimality we have $\boldsymbol{\Theta} = \boldsymbol{\Gamma} = \boldsymbol{\Pi}$, so $\boldsymbol{\Gamma}$ is an extension for (W, Δ). This shows that every minimal extension for $(W + \varphi, \Delta)$ is a minimal extension for (W, Δ).

The converse also holds: let $\boldsymbol{\Gamma}$ be a minimal extension for (W, Δ). Then, by Theorem 5.1, $\boldsymbol{\Gamma}$ is an extension for $(W + \varphi, \Delta)$: we need to see that it is minimal. If not, then there is an extension $\boldsymbol{\Theta}$ for $(W + \varphi, \Delta)$ that is properly below $\boldsymbol{\Gamma}$: $\boldsymbol{\Theta} < \boldsymbol{\Gamma}$. By Theorem 5.5 there is an extension $\boldsymbol{\Pi}$ for (W, Δ) such that $\boldsymbol{\Pi} \leqslant \boldsymbol{\Theta}$. By transitivity, we get $\boldsymbol{\Pi} < \boldsymbol{\Gamma}$, contradicting the hypothesis that $\boldsymbol{\Gamma}$ is a minimal extension for (W, Δ). This shows that every minimal extension for (W, Δ) is also a minimal extension for $(W + \varphi, \Delta)$. Thus we have established the following:

Theorem 5.6. *Suppose* $(W, \Delta) \mathrel{|\!\sim} \varphi$*. Then* (W, Δ) *and* $(W + \varphi, \Delta)$ *have exactly the same* minimal *extensions.*

It is now clear why failure of cumulativity does not preclude the relation of defeasible consequence (as based on general extensions) from being cautiously monotonic: by adding the "theorem" φ back into the world description W we do get new extensions, but such new extensions are not minimal. This fact is exploited in the proof of Theorem 5.7 below.

Theorem 5.7. *The relation* $\mid\!\sim$ *satisfies the properties of Cut, Reflexivity, and Stability.*

6. Alternative developments

In this section we consider alternative developments of the theory of the foregoing pages. First, we take up the case of semi-normal default theories. Such theories are shown to have a unique minimal extension, so that defeasible consequence appears to have a particularly simple definition for semi-normal default theories. Second, we identify certain particular, "optimal", extensions of semi-normal default theories and argue that in some cases they might provide more intuitive results, especially in the light of Example 4.2. It is not clear, however, that optimal extension can still give rise to a cautiously monotonic notion of defeasible consequence. Finally, we identify a variant of the notion of general extension that, although slightly more complicated than the one of Section 3, seems also to take care of the problematic case of Example 4.2.

6.1. Semi-normal theories

Recall that a default is semi-normal if its justification implies the conclusion, e.g., because it contains the conclusion as a conjunct. A theory all of whose defaults are semi-normal is also called semi-normal. Semi-normal default theories form a natural and well-behaved class. They were first introduced by Reiter and Criscuolo [31] to handle conflicts among defaults and to block certain unwanted instances of transitivity of implication. In this sense, they seem sufficient for most purposes in knowledge representation. However, with Reiter's notion of extension, and in contrast to the class of *normal* default theories, semi-normal default theories might fail to have an extension. It then becomes interesting to notice that using the notion of general extension not only are these theories guaranteed to have an extension (as they are with many other notions of extension), but they are guaranteed to have a uniquely minimal one.

Theorem 6.1. *Every semi-normal default theory has a unique minimal extension.*

The reader is referred to Appendix A for a complete proof, but it is worth noting here the basic conceptual fact that the proof exploits. We have seen that the proof of the fact that every default theory has a general extension proceeds by a genuine inductive process. However, such a process is nondeterministic. In particular, in the categorical case, we begin by putting $\Gamma_0^+ = \Gamma_0^- = \emptyset$, and for the inductive step we choose as Γ_{n+1}^+ a maximal set of defaults all having the following properties: (i) the conclusions of defaults in $\Gamma_n^+ \cup \Gamma_{n+1}^+$ form a consistent set; and (ii) no default in Γ_{n+1}^+ is pre-empted in $\Delta - \Gamma_n^-$ (the other set Γ_{n+1}^- is determined by the choice of Γ_{n+1}^+).

But now we are dealing with a semi-normal default theory. In particular, no *semi-normal* default can be conflicted (relative to any set of defaults) without being already pre-empted (relative to that set of defaults). It follows that the above nondeterministic definition can be replaced by a deterministic one, letting Γ^+_{n+1} be the set of *all* defaults that are not pre-empted in $\Delta - \Gamma^-_n$. By the previous remark, this ensures consistency of the set of conclusions as well.

This appears to vindicate the view, put forward by Horty et al. [15], that a cautious theory of defeasible inheritance should be *directly skeptical*, by defining the set of skeptically acceptable conclusions in a direct way, and not via detour through the intersection of all credulous extensions. Such a view has been challenged by Makinson and Schlechta [21], but now it appears that at least in the case of semi-normal default theories such a directly skeptical approach is indeed viable. In fact, in the case of a default theory (W, Δ) having a *least* extension $(\Gamma^+, \Gamma^-, \Gamma^*)$, we have that for any φ,

$$(W, \Delta) \mathrel{|\!\sim} \varphi \iff C\bigl(\Gamma^+\bigr) \models_W \varphi.$$

Then the set of the skeptically acceptable conclusions warranted by (W, Δ) can be obtained by simply giving the (deterministic) inductive process of the proof of Theorem 3.8, as amended in Theorem 6.1.

Moreover, it is possible to piggy-back a relation of defeasible consequence for arbitrary default theories on the relation $\mathrel{|\!\sim}$ defined above, restricted to semi-normal ones. In other words, for each set Δ of defaults let $\mathsf{SN}(\Delta)$ the result of replacing each default

$$\frac{\alpha : \beta}{\gamma} \quad \text{by} \quad \frac{\alpha : \beta \wedge \gamma}{\gamma}.$$

Then one could define for an arbitrary default theory (W, Δ), $(W, \Delta) \mathrel{|\!\sim}^{\mathsf{SN}} \varphi$ iff

$$\bigl(W, \mathsf{SN}(\Delta)\bigr) \mathrel{|\!\sim} \varphi.$$

Indeed, the switch from a default theory to its semi-normalized version is a natural and well-motivated move (see Reiter and Criscuolo [31] or Delgrande et al.—in fact, the latter argue that it is reasonable to replace arbitrary defaults by categorical, semi-normal ones [7, p. 193]).

6.2. *Optimal extensions*

We noticed, in discussing Example 4.2, that some of the problems with the minimal extension of the theory discussed there could be obviated by stepping up to a nonminimal extension. The problem, of course, is that in general there will be several nonminimal extensions to choose from, and there seems to be no principled way to single one out as privileged.

It turns out that this problem has a solution in the case of semi-normal default theories, for which we can identify, following Manna and Shamir [22], certain extensions as the "optimal" fixpoints of a certain monotonic operator. We begin by introducing such an operator, along with the necessary auxiliary notions. In this section, given a default theory (W, Δ) we refer to arbitrary triples $\boldsymbol{\Gamma} = (\Gamma^+, \Gamma^-, \Gamma^*)$ of subsets of Δ as *pseudo-extensions*. If a pseudo-extension $\boldsymbol{\Gamma}$ satisfies the additional condition that Γ^+ and Γ^- are disjoint, and $\Gamma^+ \subseteq \Gamma^*$, then it is called a *potential extension*. Moreover, as we did also elsewhere in the paper, we assume that the background world-description W is consistent—otherwise, there is only one extension $(\emptyset, \Delta, \emptyset)$.

Definition 6.2. Given a semi-normal default theory (W, Δ), define an operator $\boldsymbol{\tau}$ over the class of pseudo-extensions for (W, Δ) by putting $\boldsymbol{\tau}(\boldsymbol{\Gamma}) = \boldsymbol{\Theta}$, where:

$$\begin{aligned}
\Theta^+ = \{\delta\colon\ &\delta \text{ is admissible in } \Gamma^+ \text{ and}\\
&\delta \text{ is not pre-empted in } \Gamma^* - \Gamma^- \text{ relative to } W\};\\
\Theta^- = \{\delta\colon\ &\delta \text{ pre-empted or conflicted in } \Theta^+ \text{ relative to } W\};\\
\Theta^* = \{\delta\colon\ &P(\delta) \text{ is consistent with } C(\Theta^+) \text{ relative to } W\}.
\end{aligned}$$

It is immediate to verify that $\boldsymbol{\tau}$ is a monotone operator over the class of the pseudo-extensions, i.e., if $\boldsymbol{\Gamma} \leqslant \boldsymbol{\Theta}$ then $\boldsymbol{\tau}(\boldsymbol{\Gamma}) \leqslant \boldsymbol{\tau}(\boldsymbol{\Theta})$. We are going to be interested in the fixed-points of $\boldsymbol{\tau}$, i.e., pseudo-extensions $\boldsymbol{\Theta}$ such that $\boldsymbol{\tau}(\boldsymbol{\Theta}) = \boldsymbol{\Theta}$. Of course, these fixed-points need not be extensions in the usual sense: $\boldsymbol{\tau}$ will have many fixed points $\boldsymbol{\Theta}$ in which Θ^+ and Θ^- will not be disjoint, or in which Θ^+ is not a subset of $\Theta^* - \Theta^-$.

This is where Manna and Shamir [22] come to the rescue. Following their terminology, say that a collection $\mathcal{C}$ of pseudo-extensions is $\leqslant$*-consistent* if and only if any two pseudo-extensions in $\mathcal{C}$ have an upper bound in $\mathcal{C}$: i.e., if $\boldsymbol{\Gamma}$ and $\boldsymbol{\Theta}$ are in $\mathcal{C}$ then there is a pseudo-extension $\boldsymbol{\Pi}$ such that both $\boldsymbol{\Gamma} \leqslant \boldsymbol{\Pi}$ and $\boldsymbol{\Theta} \leqslant \boldsymbol{\Pi}$.

Now let $\mathcal{PE}$ be the set of all *potential* extensions for (W, Δ). It is then possible to verify that:

(C1) Any $\leqslant$-consistent subset of $\mathcal{PE}$ has a least upper bound in $\mathcal{PE}$.
(C2) Any nonempty subset of $\mathcal{PE}$ has a greatest lower bound in $\mathcal{PE}$.

Definition 6.3. Let $\boldsymbol{\Gamma}$ be a fixed point of $\boldsymbol{\tau}$. Then $\boldsymbol{\Gamma}$ is called *intrinsic* if for any other fixed point $\boldsymbol{\Theta}$ of $\boldsymbol{\tau}$, the set $\{\boldsymbol{\Gamma}, \boldsymbol{\Theta}\}$ is $\leqslant$-consistent. The largest intrinsic fixed point of $\boldsymbol{\tau}$ is called *optimal*. (See [22, p. 419].)

Theorem 6.4. *Let (W, Δ) be a default theory, and $\mathcal{PE}$ the collection of all potential extensions for (W, Δ). Then $\mathcal{PE}$ contains an optimal fixed point of $\boldsymbol{\tau}$.*

This is just an application of Theorem 3 of Manna and Shamir [22, p. 417] using conditions (C1) and (C2).

Theorem 6.5. *Let (W, Δ) be a semi-normal theory, and suppose $\boldsymbol{\Gamma}$ is a potential extension that is also a fixed point of $\boldsymbol{\tau}$. Then $\boldsymbol{\Gamma}$ is a general extension for (W, Δ).*

It follows from the theorem that every semi-normal default theory has a unique "optimal" extension. It is easy to check that the nonminimal extension of the theory of Example 4.2 is also the unique optimal extension. It appears then that optimal extensions, at least in the case of semi-normal default theories provide a viable alternative to minimal extensions, an alternative, moreover, that seems capable of avoiding any residual degree of counter-intuitiveness of minimal extensions.

It is immediate to use optimal extensions to define a relation of defeasible consequence: just say that $(W, \Delta) \mid\!\sim \varphi$ if and only if $C(\Gamma^+) \models_W \varphi$, where $\boldsymbol{\Gamma}$ is the optimal extension of (W, Δ). However, we do not know at this stage if optimal extensions can support a relation of defeasible consequence that also satisfies Cautious Monotonicity. The crucial role played by minimality in establishing Cautious Monotonicity when $\mid\!\sim$ is defined using minimal extensions makes it at least not obvious that this property carries over to optimal extensions.

6.3. "New twist" extensions

The problematic case given in Example 4.2 suggests a variant of our notion of extension. Recall that the problem with Example 4.2 appeared to be that we allowed defaults in Γ^* even though they were not even potentially admissible, in that their pre-requisite was not a logical consequence of conclusions of defaults in Δ. Now, as we remarked, this is in keeping with the intuition that potentially pre-empted defaults should not be triggered, independently of whether the pre-empting defaults are themselves potentially admissible.

However, the further intuition that defaults that are not even potentially admissible should be ruled out can be incorporated in a rather straightforward way into our definition of extension. This has the cost of complicating the proofs somewhat, but appears at the same time to take care of whatever residual degree of counter-intuitiveness the notion of extension might have. Below we consider general extensions with this "new twist", appropriately referred to as "NT-extensions".

Definition 6.6. An *NT-extension* for a default theory (W, Δ) is a triple $(\Gamma^+, \Gamma^-, \Gamma^*)$ of sets of defaults from Δ, such that:
- Γ^+ and Γ^- are disjoint;
- the following two fixpoint equations are simultaneously satisfied:

$$\Gamma^+ = \big\{\delta\colon\ \delta \text{ admissible in } \Gamma^+ \,\&\, \delta \text{ not conflicted in } \Gamma^+ \,\& \\ \delta \text{ not pre-empted in } \Gamma^* - \Gamma^-\big\};$$

$$\Gamma^- = \big\{\delta\colon\ \delta \text{ conflicted or pre-empted in } \Gamma^+\big\} \cup \\ \big\{\delta\colon\ \neg P(\delta) \text{ is } W\text{-consistent with } C\big(\Gamma^* - \Gamma^-\big)\big\};$$

- $\mathsf{Ad}(\Gamma^+) \subseteq \Gamma^* \subseteq \{\delta\colon\ P(\delta) \text{ is } W\text{-consistent with } C(\Gamma^+)\}$.

Thus, Γ^+ is the set of all defaults admissible in Γ^+ but neither conflicted in Γ^+ nor pre-empted in $(\Gamma^* - \Gamma^-)$. On the other hand, Γ^- is the set of all defaults either (i) conflicted or pre-empted in Γ^+ or (ii) whose pre-requisite is not implied by the potentially triggered defaults. Finally, Γ^* is *a* set of defaults containing all default admissible in Γ^+ and whose pre-requisite is consistent with $C(\Gamma^+)$.

It is now upon us to show that the theory of general extension carries over to NT-extensions. However, there is no obstacle to doing so, except perhaps the cost of the added complication in proofs resulting from the "new twist". As evidence of this, we give the extension-existence theorem for NT-extensions, whose proof can be found in the Appendix A.

Theorem 6.7. *Let (W, Δ) be a default theory. Then there is an iteratively definable NT-extension $\boldsymbol{\Theta}$ for (W, Δ).*

To see how this works, let us take up again Example 4.2. Recall that we had a default theory (W, Δ) with W empty and Δ comprising the defaults:

$$\delta_1 = \frac{:p}{p} \quad \text{and} \quad \delta_2 = \frac{q : \neg p}{\neg p}.$$

Since $C(\Delta) \not\models_W q$, the default δ_2 is not even potentially admissible in Δ and therefore not even potentially admissible in any $\Theta^* \subseteq \Delta$. It follows that if $\boldsymbol{\Theta}$ is an NT-extension for the theory, then $\delta_2 \in \Theta^-$. Thus, δ_1 is no longer pre-empted in $\Theta^* - \Theta^-$, and therefore $\delta_1 \in \Theta^+$, as one would expect. Indeed, it is easy to see that $(\delta_1, \delta_2, \delta_1)$ is the only NT-extension of the theory. However, as we know, it is not the only general extension. Therefore it appears that the net effect of switching to NT-extensions is to eliminate some of the minimal extensions. In this sense, this solution to the problem presented by Example 4.2 is analogous to the one in which one steps up to a nonminimal (e.g., optimal) general extension.

7. Conclusions and comparisons

As we have seen, the present approach has two main technical fallouts, namely the fact that extensions are always guaranteed to exist and the fact that general extensions allow for the definition of a well-behaved relation of defeasible consequence.

7.1. Existence of extensions

The former problem, that of the existence of extensions, has long been considered one of the basic problems of default logic. This problem has been analyzed from the point of view of computational complexity in [9,10,16,32]. In these works, broad classes of default theories have been singled out, for which extensions (in Reiter's sense) can be proved to exist. One such class comprises the *ordered default theories* of [9,32]. An even more general approach is advocated in [27], where it is shown that *even default theories* always have extensions. Indeed, there are reasons to believe that, within the framework of Reiter's approach, the result of [27] is optimal and cannot be improved upon.

Several proposals have been put forward for somewhat different notions of extensions for default theories to solve this problem. As mentioned extensions always exist according to the notions of extensions proposed by Łukaszewicz [18,19], Brewka [6], and Delgrande et al. [7]. However, these proposals are characterized by an essentially different interpretation of rôle of the justification of a default. According to the original intuition of Reiter [30] (which is adhered to in this paper) the justification of a default is to be interpreted as a mere *consistency condition* on the "triggering" of the default.

This allows, for instance, for the simultaneous triggering of defaults having mutually inconsistent justifications (provided, of course, that the conclusions are not in turn also inconsistent). This sometimes leads to counter-intuitive results as in the "broken-arms" example of Poole [28] (discussed in Section 2).

There is a second interpretation, championed for instance by Delgrande et al. [7], that rather views the justification of a default as a "working hypothesis", whose truth (as opposed to its mere consistency) is to be assumed until and unless there is information to the contrary. This intuition quite naturally leads to the so called "commitment to the justifications", one of whose consequences, for instance, is that defaults with mutually inconsistent justifications cannot be simultaneously fired, in spite of the fact that their conclusions might be consistent. In other words, the intuition behind this interpretation is that we cannot entertain inconsistent working hypotheses, and that therefore the truth of these working hypotheses should be assumed not only in each individual firing of a default, but also across firings of different defaults.

It is important to observe that extensions (in any of the senses available in the literature) for default theories can rarely be constructed (when they exist) by means of a *cumulative process*, of the sort in which defaults are successively assessed for some kind of property that can guarantee their belonging to the extension being constructed. On the contrary, in most cases, we first have to "guess" a set Θ of defaults and then check that it does indeed satisfy the equation defining extensions.

This seems to be connected, at an intuitive level, with the fact that the notions of extension employed by Reiter and many others are intrinsically "two-valued", in the sense that such extensions contain the consequences of a maximal set of defaults whose justifications are consistent with the extension itself. This means, among other things, that the triggering of a default can only be prevented if its justification or conclusion is explicitly refuted. Let us informally refer to an approach to default logic as "bold" if it shares this feature that any admissible default not explicitly pre-empted or conflicted is triggered. The formal counterpart to this idea is that *one* set of defaults is used to accomplish a *two-fold* task, namely the specification of which defaults are triggered and which defaults are pre-empted.

On the other hand, the notion of extension employed in this paper is essentially "three-valued", similarly to the one put forward in [1] for defeasible inheritance networks, and equally inspired by an analogy to Kripke's [17] approach to the theory of truth. When using general extensions, a default having a sentence φ as its justification (and whose conclusion is otherwise consistent) is prevented from being triggered if and only if $\neg\varphi$ is not explicitly rejected. With Reiter's notion of extension, a default having a sentence φ as its justification can be prevented from being triggered if and only if $\neg\varphi$ is explicitly asserted. As a further difference, contrary to Reiter's extensions, extensions of the sort proposed here can in many cases be obtained as the limit of a genuine inductive construction. By this we mean that general extensions can be constructed "from below" in stages. With the usual notion of extension such a construction from below is possible only when the theory is *semi-monotonic*, a desirable feature that is not always easy to enforce (see [7] for a discussion of the desirability of semi-monotonicity and the ways it can be achieved).

The intuitions at the basis of the notion of general extension here proposed are quite different also from those underlying the approach of Łukaszewicz, Brewka, or Delgrande et al., and this in spite of the superficial resemblances. If Reiter's notion of extension is to be generalized in such a way as to allow every default theory to have an extension, our proposal generalizes in a different direction from the one of mentioned above, a direction giving rise, as we have already seen, to quite different mathematical properties (some of these properties are summarized in Fig. 2).

This is perhaps the right place to point out that an approach somewhat similar to ours has been developed for logic programming. A "three-valued" or "well-founded" semantics for logic programs has been developed by van Gelder et al. in [13], and has found what is perhaps the most general formulation in the notion of "stable model" of Fitting [11]. And of course, similarities between logic programs and default theories have long been known (see Przymusiński [29], who applies such three-valued models to defeasible representational formalisms, although not explicitly to default logic).

However, the approach of this paper differs from the above, first of all, in being "three valued" only in inspiration. In particular, the underlying logic is thoroughly classical. The "three valued" intuition is rather cashed out by identifying extensions for default logic not with sets of sentences or defaults, but with *pairs* (or, later, *triples*) of such sets. This allows for a more concrete representation of extensions, whose properties can then be more easily investigated. Moreover, the present approach gives rise, when applied to arbitrary default theories (as opposed to "categorical" or "prerequisite-free" ones), to both *monotonic* and *anti-monotonic* processes that in general do not seem to arise in the well-founded approaches to logic programming, except in the general approach of Fitting [11].

7.2. Defeasible consequence—again

We now come to the second technical fallout of the notion of general extension, namely the definition of a well-behaved relation of defeasible consequence.

Using the notions of extension for Default Logic available in the literature, there are several ways in which we can define, given a default theory having multiple incomparable extensions, a relation $\mid\!\sim$:

1. We can decide to be *credulous* and say that $(W, \Delta) \mid\!\sim \varphi$ precisely when φ follows from (the set of conclusions of defaults in) *some* extension of (W, Δ).
2. We can arbitrarily pick an extension Γ among the many possible, and decide that $(W, \Delta) \mid\!\sim \varphi$ just in case φ follows from (conclusions of defaults in) Γ.
3. We can be *skeptical* and say that $(W, \Delta) \mid\!\sim \varphi$ just in case φ follows from (conclusions of defaults in) *all* extensions of (W, Δ).

All three alternatives have drawbacks. Alternative (2) is not acceptable unless we have a principled way to make such a choice of Γ. Alternative (1) can lead us sometimes to endorse contradictory statements. Alternative (3) is the one that best resonates with certain intuitions about defeasible reasoning, e.g., the fact that defeasible reasoners should be cautious in drawing their inferences (see Horty et al. [15] for a general argument in favor of skepticism in defeasible reasoning), but it seems to be going about it the wrong way. It does not appear to be an appropriate "implementation" of skepticism to generate *all* possible extensions of a theory and then take the intersection. The contrast here is with a "directly skeptical" approach that would generate the set of conclusions that are skeptically acceptable *without* going through the detour of first generating all credulous extensions. Moreover, if feasibility of computation is an issue at all, the "intersection-of-extensions" approach is by far the least resource-oriented. To clinch matters, as shown by Makinson [20], the intersection of extensions approach using Reiter's notion, simply does not give rise to a cautiously monotonic relation of defeasible consequence for default logic.

The cautiously three-valued character of general extensions is responsible for the fact that at least in some important cases there is a natural definition of a notion of logical consequence for default theories. In particular, as we have seen, in the case of semi-normal default theories, there is a privileged (unique minimal) general extension which can be used to define a notion of defeasible consequence. (It is also worth remarking that at least in the case of semi-normal default theories the least extension can be obtained by means of a genuine inductive construction—in this being similar to, e.g., "prerequisite-free constrained default logic" (PfConDL) of [7], which also allows for a similar construction, but using quite different underlying intuitions.)

In contrast, there is no privileged extension in any of the senses of Reiter [30], Łukaszewicz [18,19], Brewka [6] or Delgrande et al. [7]. In fact, extensions are always taken to be maximal, meaning that any two distinct extensions are $\subseteq$-incomparable, if not outright mutually inconsistent.

Acknowledgement

I thank Mel Fitting for illuminating conversations on the semantics of logic programs, as well as Dagfinn Føllesdal and Horacio Arlo Costa for helpful comments. I am indebted to two anonymous referees for the journal *Artificial Intelligence* for providing useful feedback and criticisms, and especially for pointing out the crucial Example 4.2.

Appendix A. Proofs of selected theorems

Theorem 1.5. *Let Γ^+ be a classical extension for a categorical default theory (W, Δ), and put*

$$\Gamma^- = \{\delta\colon\ \delta \text{ conflicted or pre-empted in } \Gamma^+\},$$

then (Γ^+, Γ^-) is a general extension for (W, Δ).

Proof. From the hypothesis that Γ^+ is a classical extension we know that it satisfies

$$\Gamma^+ = \{\delta \in \Delta\colon\ C(\Gamma^+) \not\models_W \neg J(\delta)\}.$$

First we note that Γ^+ and Γ^- are disjoint: if δ belongs to both Γ^+ and Γ^- then (from the latter) we have that either δ is conflicted in Γ^+ or δ is pre-empted in Γ^+. The second alternative is impossible, since δ is in Γ^+ (and hence *not* pre-empted in Γ^+); therefore it can only be that δ is conflicted in Γ^+, i.e., $C(\Gamma^+) \models_W \neg C(\delta)$. Since δ itself is in Γ^+, we have that $C(\Gamma^+)$ is W-inconsistent and any default is pre-empted in an inconsistent set of sentences. It would follow that δ is pre-empted in Γ^+ after all, which we have already ruled out. We conclude that there cannot be any such δ.

So we need to show that the pair (Γ^+, Γ^-) satisfies the pair of fixpoint equations of Definition 1.4. In turn, we observe that the second of such equations, namely

$$\Gamma^- = \bigl\{\delta\colon\ C\bigl(\Gamma^+\bigr) \models_W \neg C(\delta) \text{ or } C\bigl(\Gamma^+\bigr) \models_W \neg J(\delta)\bigr\},$$

is satisfied by definition of Γ^-. So we need to establish that

$$\Gamma^+ = \bigl\{\delta\colon\ C\bigl(\Gamma^+\bigr) \not\models_W \neg C(\delta) \,\&\, C\bigl(\Delta - \Gamma^-\bigr) \not\models_W \neg J(\delta)\bigr\}.$$

So what we need to show is that δ is not pre-empted in Γ^+ if and only if it is neither conflicted in Γ^+ nor pre-empted in $\Delta - \Gamma^-$:

$$C\bigl(\Gamma^+\bigr) \not\models_W \neg J(\delta) \iff C\bigl(\Gamma^+\bigr) \not\models_W \neg C(\delta) \,\&\, C\bigl(\Delta - \Gamma^-\bigr) \not\models_W \neg J(\delta).$$

In turn, this claim breaks down into the following, which need to be separately established:

(1) $C(\Gamma^+) \models_W \neg J(\delta)$ only if $C(\Gamma^+) \models_W \neg C(\delta)$ or $C(\Delta - \Gamma^-) \models_W \neg J(\delta)$;
(2) $C(\Gamma^+) \models_W \neg C(\delta)$ only if $C(\Gamma^+) \models_W \neg J(\delta)$;
(3) $C(\Delta - \Gamma^-) \models_W \neg J(\delta)$ only if $C(\Gamma^+) \models_W \neg J(\delta)$.

Part (1) establishes the converse implication, while parts (2) and (3) together suffice for the direct implication.

For part (1), assume $C(\Gamma^+) \models_W \neg J(\delta)$; then there are defaults $\delta_1, \ldots, \delta_n \in \Gamma^+$ such that $C(\delta_1), \ldots, C(\delta_n) \models_W \neg J(\delta)$; since Γ^+ is a classical extension, $\delta_1, \ldots, \delta_n$ are neither conflicted nor pre-empted in Γ^+, whence by definition of Γ^- we have $\delta_1, \ldots, \delta_n \notin \Gamma^-$, which is to say $C(\Delta - \Gamma^-) \models_W \neg J(\delta)$. Part (1) follows.

For part (2), suppose $C(\Gamma^+) \models_W \neg C(\delta)$. We can assume that $C(\Gamma^+)$ is consistent, because otherwise the conclusion $C(\Gamma^+) \models_W \neg J(\delta)$ follows immediately. Then, if $C(\Gamma^+)$ is consistent, we have $\delta \notin \Gamma^+$ (since $C(\Gamma^+) \models_W \neg C(\delta)$), and since Γ^+ is a classical extension, also $C(\Gamma^+) \models_W \neg J(\delta)$, which is the desired conclusion.

Finally, for part (3), suppose δ is pre-empted in $C(\Delta - \Gamma^-)$, i.e., $C(\Delta - \Gamma^-) \models_W \neg J(\delta)$. Then there are $\delta_1, \ldots, \delta_n \notin \Gamma^-$ such that $C(\delta_1), \ldots, C(\delta_n) \models_W \neg J(\delta)$. By the definition of Γ^-, we know that $\delta_1, \ldots, \delta_n$ are neither conflicted nor pre-empted in Γ^+, and since Γ^+ is a classical extension, we have $\delta_1, \ldots, \delta_n \in \Gamma^+$ by Theorem 1.3. The desired conclusion $C(\Gamma^+) \models_W \neg J(\delta)$ follows. □

Theorem 1.6. *Every categorical default theory has a general extension.*

Proof. Let (W, Δ) be a categorical default theory; we can assume that W is consistent, because otherwise $(\emptyset, \Delta)$ is an extension. We define a general extension (Γ^+, Γ^-), where Γ^+ is the union of a sequence $\Gamma_0^+, \Gamma_1^+, \ldots$ and Γ^- is the union of the sequence $\Gamma_0^-, \Gamma_1^-, \ldots$; in turn, the sets Γ_n^+ and Γ_n^- are inductively defined using an iterative process similar to the one in [14] or [1].

For the base case put $\Gamma_0^+ = \Gamma_0^- = \emptyset$, which are obviously disjoint. For the inductive step, let Γ_{n+1}^+ be a *maximal* set of defaults such that:

(1) $C(\Gamma_n^+ \cup \Gamma_{n+1}^+)$ is consistent;
(2) no default $\delta \in \Gamma_{n+1}^+$ is pre-empted in $\Delta - \Gamma_n^-$.

(Observe that such a maximal set of defaults always exists, although it might not be unique.) For Γ_{n+1}^-, put

$$\Gamma_{n+1}^- = \bigl\{\delta\colon\ \delta \text{ pre-empted or conflicted in } \Gamma_{n+1}^+\bigr\}.$$

As in the proof of Theorem 1.5, it follows that Γ_{n+1}^{+} and Γ_{n+1}^{-} are disjoint.

Next, we show that $\Gamma_n^{+} \subseteq \Gamma_{n+1}^{+}$ and $\Gamma_n^{-} \subseteq \Gamma_{n+1}^{-}$, by induction on n. The base case for $n = 0$ is trivial. For the inductive step, assume $\Gamma_n^{+} \subseteq \Gamma_{n+1}^{+}$ and $\Gamma_n^{-} \subseteq \Gamma_{n+1}^{-}$, in order to show $\Gamma_{n+1}^{+} \subseteq \Gamma_{n+2}^{+}$ and $\Gamma_{n+1}^{-} \subseteq \Gamma_{n+2}^{-}$.

Ad $\Gamma_{n+1}^{+} \subseteq \Gamma_{n+2}^{+}$. Let $\delta \in \Gamma_{n+1}^{+}$. First we show that δ is not pre-empted in $\Delta - \Gamma_{n+1}^{-}$: for if δ were so pre-empted, then $C(\Delta - \Gamma_{n+1}^{-}) \models_W \neg J(\delta)$, and since $(\Delta - \Gamma_{n+1}^{-}) \subseteq (\Delta - \Gamma_n^{-})$ by the inductive hypothesis, also $C(\Delta - \Gamma_n^{-}) \models_W \neg J(\delta)$, which is impossible if $\delta \in \Gamma_{n+1}^{+}$. Moreover, by the consistency of $C(\Gamma_{n+1}^{+} \cup \Gamma_{n+2}^{+})$, we have that δ cannot be inconsistent with $C(\Gamma_{n+1}^{+} \cup \Gamma_{n+2}^{+})$. By the maximality of Γ_{n+2}^{+}, it follows $\delta \in \Gamma_{n+2}^{+}$.

Ad $\Gamma_{n+1}^{-} \subseteq \Gamma_{n+2}^{-}$. Suppose $\delta \in \Gamma_{n+1}^{-}$. Then δ is either conflicted or pre-empted in Γ_{n+1}^{+}, and since $\Gamma_{n+1}^{+} \subseteq \Gamma_{n+2}^{+}$, we have that δ is also conflicted or, respectively, pre-empted in Γ_{n+2}^{+}, so that $\delta \in \Gamma_{n+2}^{-}$.

This shows that the sequences Γ_n^{+} and Γ_n^{-} (for $n \geqslant 0$) are increasing. Put:

$$\Gamma^{+} = \bigcup_{n \geqslant 0} \Gamma_n^{+};$$

$$\Gamma^{-} = \bigcup_{n \geqslant 0} \Gamma_n^{-}.$$

We verify that (Γ^{+}, Γ^{-}) is an extension for (W, Δ). First we observe that Γ^{+} and Γ^{-} are disjoint (any common member δ would have had to be put in at some stage n, but this is impossible). Next, we check that the following equation is satisfied:

$$\Gamma^{+} = \bigl\{\delta\colon C\bigl(\Gamma^{+}\bigr) \not\models_W \neg C(\delta) \,\&\, C\bigl(\Delta - \Gamma^{-}\bigr) \not\models_W \neg J(\delta)\bigr\}.$$

In one direction, let $\delta \in \Gamma^{+}$ and choose k such that $\delta \in \Gamma_k^{+}$. First we show that δ cannot be conflicted in Γ^{+}, for otherwise it would have been conflicted in some Γ_n^{+}, and in particular also in $\Gamma_{\max(n,k)}^{+}$ which would then have been inconsistent, against construction. Moreover, δ cannot be pre-empted in $\Delta - \Gamma^{-}$: if it were so pre-empted, then for some $\delta_1, \dots, \delta_m \notin \Gamma^{-}$ we would have $C(\delta_1), \dots, C(\delta_m) \models_W \neg J(\delta)$. Since the sequence Γ_n^{-} is increasing, the sequence $(\Delta - \Gamma_n^{-})$ is *decreasing*, so that δ would have been pre-empted also in some (indeed: every) $\Delta - \Gamma_n^{-}$.

For the other direction, we need to show that if δ is neither conflicted in Γ^{+} nor pre-empted in $\Delta - \Gamma^{-}$ then $\delta \in \Gamma^{+}$. First observe that if δ is not conflicted in Γ^{+}, then it cannot be conflicted in any Γ_n^{+}. So, if $\delta \notin \Gamma^{+}$ it must be because δ is pre-empted in *every* $\Delta - \Gamma_n^{-}$. So for every $n \geqslant 0$, there are $\delta_1, \dots, \delta_m \notin \Gamma_n^{-}$ such that $C(\delta_1), \dots, C(\delta_m) \models_W \neg J(\delta)$. *Now for the first time we use the finiteness hypothesis for* Δ: since Δ is finite, there are finitely many m-tuple of defaults. Since the sequence $(\Delta - \Gamma_n^{-})$ is decreasing, some m-tuple of defaults pre-empting δ must be in every $(\Delta - \Gamma_n^{-})$, and hence also in $(\Delta - \Gamma^{-})$, against hypothesis. So, there must be some k such that every m-tuple $\delta_1, \dots, \delta_m$ of defaults pre-empting δ has a member in Γ_k^{-}. Then δ is not pre-empted in $(\Delta - \Gamma_k^{-})$, whence $\delta \in \Gamma_{k+1}^{+} \subseteq \Gamma^{+}$, as required.

For the second equation, we need to verify

$$\Gamma^{-} = \bigl\{\delta\colon C\bigl(\Gamma^{+}\bigr) \models_W \neg C(\delta) \text{ or } C\bigl(\Gamma^{+}\bigr) \models_W \neg J(\delta)\bigr\}.$$

In one direction, suppose $\delta \in \Gamma^-$; then for some n, $\delta \in \Gamma^-_{n+1}$, so that δ is pre-empted or conflicted in Γ^+_{n+1}. This implies that δ is also pre-empted or conflicted in Γ^+. For the other direction, suppose δ is pre-empted in Γ^+. So there are defaults $\delta_1, \ldots, \delta_m \in \Gamma^+$ such that $C(\delta_1), \ldots, C(\delta_m) \models_W \neg J(\delta)$. Find n large enough such that $\delta_1, \ldots, \delta_m \in \Gamma^+_n$. Then $\delta \in \Gamma^-_n \subseteq \Gamma^-$. The case for δ conflicted is similar. □

Theorem 3.3. *Let Γ be a classical extension for a default theory (W, Δ); then Γ is grounded.*

Proof. Since Γ is classical, $\Gamma = \bigcup_{n\geqslant 0} \Gamma_n$, where $\Gamma_0 = \emptyset$ and

$$\Gamma_{n+1} = \{\delta\colon\ \delta \text{ admissible in } \Gamma_n \text{ and not pre-empted in } \Gamma\}.$$

In order to show that Γ is grounded, put $\Theta_0 = \emptyset$ and $\Theta_{n+1} = \{\delta \in \Gamma\colon\ \delta \text{ admissible in } \Theta_n\}$.

We need to show $\Gamma = \bigcup_{n\geqslant 0} \Theta_n$, i.e.,

$$\bigcup_{n\geqslant 0} \Gamma_n = \bigcup_{n\geqslant 0} \Theta_n.$$

First we observe that the inclusion $\Theta_n \subseteq \Gamma$ holds for every n, by definition of Θ_n. It follows that $\bigcup_{n\geqslant 0} \Theta_n \subseteq \Gamma$.

Second, we show that $\Gamma \subseteq \bigcup_{n\geqslant 0} \Theta_n$. In turn, it suffices to show that $\Gamma_p \subseteq \bigcup_{n\geqslant 0} \Theta_n$ for every p. We proceed by induction on p. For the base case, $\Gamma_0 = \emptyset \subseteq \bigcup_{n\geqslant 0} \Theta_n$. For the inductive step assume $\Gamma_p \subseteq \bigcup_{n\geqslant 0} \Theta_n$. Suppose $\delta \in \Gamma_{p+1}$: in particular, δ is admissible in Γ_p. So there are defaults $\delta_1, \ldots, \delta_m \in \Gamma_p$ such that $C(\delta_1), \ldots, C(\delta_m) \models_W P(\delta)$. By the inductive hypothesis $\Gamma_p \subseteq \bigcup_{n\geqslant 0} \Theta_n$, so there must be $q \geqslant 0$ such that $\delta_1, \ldots, \delta_m \in \Theta_q$. Then $\delta \in \Theta_{q+1} \subseteq \bigcup_{n\geqslant 0} \Theta_n$, as required.

If we now put $\Theta = \bigcup_{n\geqslant 0} \Theta_n$, we have $\Gamma = \Theta$. Then we have $\Theta_0 = \emptyset$ and $\Theta_{n+1} = \{\delta \in \Theta\colon\ \delta \text{ admissible in } \Theta_n\}$. So Θ is grounded and hence so is Γ. □

Theorem 3.7. *Let $\Gamma^+ = \bigcup_{n\geqslant 0} \Gamma^+_n$ be a classical extension for a default theory (W, Δ). Put*:

$$\Gamma^- = \bigl\{\delta\colon\ \delta \text{ conflicted or pre-empted in } \Gamma^+\bigr\};$$
$$\Gamma^* = \bigl\{\delta\colon\ \delta \text{ admissible in } \Gamma^+\bigr\} = \mathsf{Ad}\bigl(\Gamma^+\bigr).$$

Then $(\Gamma^+, \Gamma^-, \Gamma^*)$ is a general extension for (W, Δ).

Proof. We modify the proof of Theorem 1.5 as needed. From the hypothesis that Γ^+ is a classical extension it follows that if $\delta \in \Gamma^-$ then $\delta \notin \Gamma^+$, i.e., Γ^+ and Γ^- are disjoint. Moreover, we know that $\Gamma^+ = \bigcup_{n\geqslant 0} \Gamma^+_n$, where $\Gamma^+_0 = \emptyset$, and

$$\Gamma^+_{n+1} = \bigl\{\delta \in \Delta\colon\ C\bigl(\Gamma^+_n\bigr) \models_W P(\delta) \;\&\; C\bigl(\Gamma^+\bigr) \not\models_W \neg J(\delta)\bigr\}.$$

So we need to show that the triple $(\Gamma^+, \Gamma^-, \Gamma^*)$ satisfies the conditions of Definition 3.4. In turn, we observe that the condition

$$\Gamma^- = \bigl\{\delta\colon\ C\bigl(\Gamma^+\bigr) \models_W \neg C(\delta) \text{ or } C\bigl(\Gamma^+\bigr) \models_W \neg J(\delta)\bigr\},$$

is satisfied by definition of Γ^-, whereas the condition

$$\mathsf{Ad}\big(\Gamma^+\big) \subseteq \Gamma^* \subseteq \big\{\delta\colon\ C\big(\Gamma^+\big) \not\models_W \neg P(\delta)\big\},$$

derives from the hypothesis that W is consistent and hence $C(\Gamma^+)$ is W-consistent. So we need to establish that

$$\Gamma^+ = \big\{\delta\colon\ C\big(\Gamma^+\big) \models_W P(\delta)\ \&\ C\big(\Gamma^+\big) \not\models_W \neg C(\delta)\ \&\ C\big(\Gamma^* - \Gamma^-\big) \not\models_W \neg J(\delta)\big\}.$$

In other words, what we need to show is that $\delta \in \Gamma^+$ if and only if it is admissible in Γ^+, but neither conflicted in Γ^+ nor pre-empted in $\Gamma^* - \Gamma^-$. In turn, this claim breaks down into the following four parts:

(1) $\delta \notin \Gamma^+$ only if $C(\Gamma^+) \not\models_W P(\delta)$ or $C(\Gamma^+) \models_W \neg C(\delta)$ or $C(\Gamma^* - \Gamma^-) \models_W \neg J(\delta)$;
(2) $C(\Gamma^+) \models_W \neg C(\delta)$ only if $\delta \notin \Gamma^+$;
(3) $C(\Gamma^* - \Gamma^-) \models_W \neg J(\delta)$ only if $\delta \notin \Gamma^+$;
(4) $C(\Gamma^+) \not\models_W P(\delta)$ only if $\delta \notin \Gamma^+$.

For part (1), assume $\delta \notin \Gamma^+$. Since Γ^+ is classical, by Theorem 1.3, either (i) δ is not admissible in Γ^+, or (ii) it is conflicted in Γ^+, or (iii) it is pre-empted in Γ^+. The conclusion follows immediately in cases (i) and (ii). For case (iii) we need to show that if δ is pre-empted in Γ^+, then it is pre-empted in $C(\Gamma^* - \Gamma^-)$. So suppose δ is pre-empted in Γ^+: then there are defaults $\delta_1, \ldots, \delta_n \in \Gamma^+$ such that $C(\delta_1), \ldots, C(\delta_n) \models_W \neg J(\delta)$. Since Γ^+ is a classical extension, $\delta_1, \ldots, \delta_n$ are neither conflicted nor pre-empted in Γ^+, whence by definition of Γ^- we have $\delta_1, \ldots, \delta_n \notin \Gamma^-$. Moreover, again because Γ^+ is classical, we have that $\delta_1, \ldots, \delta_n$ are all admissible in Γ^+, and hence $\delta_1, \ldots, \delta_n \in \Gamma^*$. It follows that $C(\Gamma^* - \Gamma^-) \models_W \neg J(\delta)$, whence part (1) follows.

For part (2), we have immediately that since Γ^+ is classical, the hypothesis that $\delta \in \Gamma^+$ implies that δ is not conflicted in Γ^+.

For part (3), suppose δ is pre-empted in $C(\Gamma^* - \Gamma^-)$, i.e.,

$$C\big(\Gamma^* - \Gamma^-\big) \models_W \neg J(\delta).$$

Then there are $\delta_1, \ldots, \delta_n \notin \Gamma^-$ such that $C(\delta_1), \ldots, C(\delta_n) \models_W \neg J(\delta)$. In particular, $\delta_1, \ldots, \delta_n$ belong to Γ^* and hence are admissible in Γ^+; moreover by the definition of Γ^-, we know that $\delta_1, \ldots, \delta_n$ are neither conflicted nor pre-empted in Γ^+. So, since Γ^+ is a classical extension, we have $\delta_1, \ldots, \delta_n \in \Gamma^+$ by Theorem 1.3. The desired conclusion $C(\Gamma^+) \models_W \neg J(\delta)$ follows.

Finally, for part (4), if $\delta \in \Gamma^+$ then δ is admissible in Γ^+, since Γ^+ is a classical extension. □

Theorem 3.8. *Let (W, Δ) be a default theory. Then*:

(i) *every default theory has an iteratively definable general extension*;
(ii) *such an extension is $\leqslant$-minimal*; *and*
(iii) *any minimal extension for (W, Δ) is grounded.*

Proof. Part (i) is similar to the proof of Theorem 1.6. Let (W, Δ) be a default theory. We define a general extension $(\Gamma^+, \Gamma^-, \Gamma^*)$, where $\Gamma^+ = \bigcup_{n \geqslant 0} \Gamma_n^+$ and $\Gamma^- = \bigcup_{n \geqslant 0} \Gamma_n^-$, whereas $\Gamma^* = \bigcap_{n \geqslant 0} \Gamma_n^*$.

For the base case put $\Gamma_0^+ = \emptyset$, $\Gamma_0^- = \{\delta\colon \delta \text{ conflicted or pre-empted in } \Gamma_0^+\}$, and $\Gamma_0^* = \{\delta\colon \Gamma_0^+ \not\models_W \neg P(\delta)\}$.

For the inductive step, let Γ_{n+1}^+ be a *maximal* set of defaults such that:

(1) $C(\Gamma_n^+ \cup \Gamma_{n+1}^+)$ is consistent;
(2) every $\delta \in \Gamma_{n+1}^+$ is admissible in Γ_n^+;
(3) no default $\delta \in \Gamma_{n+1}^+$ is pre-empted in $\Gamma_n^* - \Gamma_n^-$.

(Observe again that such a maximal set of defaults need not be unique.) Moreover, put:

$$\Gamma_{n+1}^- = \{\delta\colon \delta \text{ pre-empted or conflicted in } \Gamma_{n+1}^+\};$$
$$\Gamma_{n+1}^* = \{\delta\colon C(\Gamma_{n+1}^+) \not\models_W \neg P(\delta)\}.$$

Next, we show that the sequences Γ_n^+ and Γ_n^- are *increasing*, whereas the sequence Γ_n^* is *decreasing*. We proceed by induction on n and show that:

(1) $\Gamma_n^+ \subseteq \Gamma_{n+1}^+$;
(2) $\Gamma_n^- \subseteq \Gamma_{n+1}^-$;
(3) $\Gamma_{n+1}^* \subseteq \Gamma_n^*$.

Base case for $n = 0$: trivially we have $\Gamma_0^+ \subseteq \Gamma_1^+$. Moreover, if $\delta \in \Gamma_0^-$, then δ is conflicted or pre-empted in Γ_0^+: then it is also conflicted or pre-empted in Γ_1^+, so that $\delta \in \Gamma_1^-$. This shows $\Gamma_0^- \subseteq \Gamma_1^-$. Finally, if $C(\Gamma_0^+) \models_W \neg P(\delta)$, then also $C(\Gamma_1^+) \models_W \neg P(\delta)$, so that $\Gamma_1^* \subseteq \Gamma_0^*$.

For the inductive step, assume $\Gamma_n^+ \subseteq \Gamma_{n+1}^+$ and $\Gamma_n^- \subseteq \Gamma_{n+1}^-$, as well as $\Gamma_{n+1}^* \subseteq \Gamma_n^*$, in order to show the analogue properties for $n+1$ and $n+2$.

Ad $\Gamma_{n+1}^+ \subseteq \Gamma_{n+2}^+$. Let $\delta \in \Gamma_{n+1}^+$. First we show that δ is not pre-empted in $\Gamma_{n+1}^* - \Gamma_{n+1}^-$: for if δ were so pre-empted, then $C(\Gamma_{n+1}^* - \Gamma_{n+1}^-) \models_W \neg J(\delta)$, and since by the inductive hypothesis $(\Delta - \Gamma_{n+1}^-) \subseteq (\Delta - \Gamma_n^-)$, as well as $\Gamma_{n+1}^* \subseteq \Gamma_n^*$ also $C(\Gamma_n^* - \Gamma_n^-) \models_W \neg J(\delta)$, which is impossible if $\delta \in \Gamma_{n+1}^+$. Moreover, by the W-consistency of $C(\Gamma_{n+1}^+ \cup \Gamma_{n+2}^+)$, we have that δ cannot be W-inconsistent with $C(\Gamma_{n+1}^+ \cup \Gamma_{n+2}^+)$. By the maximality of Γ_{n+2}^+, it follows $\delta \in \Gamma_{n+2}^+$.

Ad $\Gamma_{n+1}^- \subseteq \Gamma_{n+2}^-$. Suppose $\delta \in \Gamma_{n+1}^-$. Then δ is either conflicted or pre-empted in Γ_{n+1}^+, and since $\Gamma_{n+1}^+ \subseteq \Gamma_{n+2}^+$, we have that δ is also conflicted or, respectively, pre-empted in Γ_{n+2}^+, so that $\delta \in \Gamma_{n+2}^-$.

Finally, ad $\Gamma_{n+2}^* \subseteq \Gamma_{n+1}^*$. This follows from the fact that $\Gamma_{n+1}^+ \subseteq \Gamma_{n+2}^+$, so that if $C(\Gamma_{n+1}^+) \models_W \neg P(\delta)$ then also $C(\Gamma_{n+2}^+) \models_W \neg P(\delta)$.

This shows that the sequences Γ_n^+ and Γ_n^- are increasing whereas Γ_n^* is decreasing (for $n \geqslant 0$). It is not hard to see that the sequences Γ_n^+ and Γ_n^- are disjoint. Now put:

$$\Gamma^+ = \bigcup_{n \geqslant 0} \Gamma_n^+;$$
$$\Gamma^- = \bigcup_{n \geqslant 0} \Gamma_n^-;$$
$$\Gamma^* = \bigcap_{n \geqslant 0} \Gamma_n^*.$$

Observe that $C(\Gamma^+)$ must be consistent, given the assumption that W is consistent. We verify that $(\Gamma^+, \Gamma^-, \Gamma^*)$ is an extension for (W, Δ). First we remark that the sets Γ^+ and Γ^- are disjoint, since they are the limits of disjoint sequences. Then we check that the following equation is satisfied:

$$\Gamma^+ = \bigl\{\delta\colon C\bigl(\Gamma^+\bigr) \models_W P(\delta) \mathbin{\&} C\bigl(\Gamma^+\bigr) \not\models_W \neg C(\delta) \mathbin{\&} C\bigl(\Gamma^* - \Gamma^-\bigr) \not\models_W \neg J(\delta)\bigr\}.$$

In one direction, let $\delta \in \Gamma^+$ and choose k such that $\delta \in \Gamma_k^+$. First we observe that δ cannot be conflicted in Γ^+, for otherwise it would have been conflicted in some Γ_n^+, and in particular also in $\Gamma_{\max(n,k)}^+$ which would then have been inconsistent, against construction. Moreover, δ must be admissible in Γ_k^+ and hence in Γ^+. Finally, δ cannot be pre-empted in $\Gamma^* - \Gamma^-$: if it were so pre-empted, then for some $\delta_1, \ldots, \delta_m \in (\Gamma^* - \Gamma^-)$ we would have $C(\delta_1), \ldots, C(\delta_m) \models_W \neg J(\delta)$. Since the sequence Γ_n^- is increasing, the sequence $(\Delta - \Gamma_n^-)$ is *decreasing*, so that δ would have been pre-empted also in every $(\Gamma_n^* - \Gamma_n^-)$, including $(\Gamma_k^* - \Gamma_k^-)$.

For the other direction, we need to show that if δ is admissible in Γ^+ but neither conflicted in Γ^+ nor pre-empted in $\Gamma^* - \Gamma^-$ then $\delta \in \Gamma^+$. First observe that if δ is not conflicted in Γ^+, then it cannot be conflicted in any Γ_n^+. Moreover, if δ is admissible in Γ^+ it must be admissible in Γ_k^+ for every k greater than or equal to some p. So, if $\delta \notin \Gamma^+$ it must be because δ is pre-empted in *every* $\Gamma_{n+p}^* - \Gamma_{n+p}^-$. So for every $n \geqslant 0$, there are $\delta_1, \ldots, \delta_m \in (\Gamma_{n+p}^* - \Gamma_{n+p}^-)$ such that $C(\delta_1), \ldots, C(\delta_m) \models_W \neg J(\delta)$. *Now we use the finiteness hypothesis for* Δ: since Δ is finite, and the sequence $(\Gamma_n^* - \Gamma_n^-)$ is decreasing, there is some m-tuple $\delta_1, \ldots, \delta_m$ pre-empting δ and some $q \geqslant 0$ such that $\delta_1, \ldots, \delta_m \in (\Gamma_p^* - \Gamma_p^-)$ for all $p \geqslant q$. Then δ is pre-empted in $\Gamma^* - \Gamma^-$, against assumption. So $\delta \in \Gamma_{k+1}^+ \subseteq \Gamma^+$, as required.

For the second equation, we need to verify

$$\Gamma^- = \bigl\{\delta\colon C\bigl(\Gamma^+\bigr) \models_W \neg C(\delta) \text{ or } C\bigl(\Gamma^+\bigr) \models_W \neg J(\delta)\bigr\}.$$

In one direction, suppose $\delta \in \Gamma^-$; then for some n, $\delta \in \Gamma_{n+1}^-$, so that δ is pre-empted or conflicted in Γ_{n+1}^+. This implies that δ is also pre-empted or conflicted in Γ^+. For the other direction, suppose δ is pre-empted in Γ^+. So there are defaults $\delta_1, \ldots, \delta_m \in \Gamma^+$ such that $C(\delta_1), \ldots, C(\delta_m) \models_W \neg J(\delta)$. Find n large enough such that $\delta_1, \ldots, \delta_m \in \Gamma_n^+$. Then $\delta \in \Gamma_n^- \subseteq \Gamma^-$. The case for δ conflicted is similar.

Finally, the condition

$$\mathsf{Ad}\bigl(\Gamma^+\bigr) \subseteq \Gamma^* \subseteq \bigl\{\delta\colon C\bigl(\Gamma^+\bigr) \not\models_W \neg P(\delta)\bigr\},$$

is satisfied by construction. For the first inclusion we observe that if δ is admissible in Γ^+ then $C(\Gamma^+) \not\models_W \neg P(\delta)$ given that $C(\Gamma^+)$ is consistent. For the second inclusion, if $C(\Gamma^+) \models_W \neg P(\delta)$, then for some n $C(\Gamma_n^+) \not\models_W \neg P(\delta)$, whence $\delta \notin \Gamma_n^*$; this implies $\delta \notin \bigcap_{n \geqslant 0} \Gamma_n^* = \Gamma^*$.

We now take up part (ii): we need to show that the extension obtained in the construction given above is $\leqslant$-minimal. We do this by assuming that for some extension $(\Theta^+, \Theta^-, \Theta^*)$, we have

$$\bigl(\Theta^+, \Theta^-, \Theta^*\bigr) \leqslant \bigl(\Gamma^+, \Gamma^-, \Gamma^*\bigr),$$

and showing that $(\Theta^+, \Theta^-, \Theta^*) = (\Gamma^+, \Gamma^-, \Gamma^*)$. In turn, by induction on n, it suffices to show:

(1) $\Gamma_n^+ \subseteq \Theta^+$;
(2) $\Gamma_n^- \subseteq \Theta^-$;
(3) $\Theta^* \subseteq \Gamma_n^*$.

In the case for $n = 0$, parts (1) and (2) are immediate. To show $\Theta^* \subseteq \Gamma_0^*$ suppose that $\delta \in \Theta^*$. Now, Γ_0^* is the set of defaults whose pre-requisite is W-consistent (with Γ_0^+). If $P(\delta)$ were W-inconsistent, then $C(\Theta^+) \models_W \neg P(\delta)$, which implies $\delta \notin \Theta^*$.

For the inductive step, assume parts (1)–(3) for n, in order to establish them for $n+1$. From (2) and (3) for n we have $(\Theta^* - \Theta^-) \subseteq (\Gamma_n^* - \Gamma_n^-)$. Now suppose $\delta \in \Gamma_{n+1}^+$; then: (i) δ is admissible in Γ_n^+ and hence also in Θ^+. (ii) δ is not pre-empted in $(\Gamma_n^* - \Gamma_n^-)$ and hence not pre-empted in $(\Theta_n^* - \Theta_n^-)$, either. It remains to show that δ is not conflicted in Θ^+; if it were so conflicted, since by hypothesis $\Theta^+ \subseteq \Gamma^+$, δ would also be conflicted in Γ_m^+, for some $m \geqslant n+1$. Since $\delta \in \Gamma_{n+1}^+ \subseteq \Gamma_m^+$, we would have $C(\Gamma_m^+)$ W-inconsistent, which is impossible in our construction. This concludes the inductive step for (1).

The inductive step for (2) follows immediately from the inductive hypothesis. For (3), we need to show $\Theta^* \subseteq \Gamma_{n+1}^*$. If $\delta \notin \Gamma_{n+1}^*$, then $C(\Gamma_{n+1}^+) \models_W \neg P(\delta)$, and since $\Gamma_{n+1}^+ \subseteq \Theta^+$, also $C(\Theta^+) \models_W \neg P(\delta)$, so that $\delta \notin \Theta^*$.

Finally, for part (iii) of the theorem, we show that Γ^+ is grounded. Put $\Theta_0 = \emptyset$, and $\Theta_{n+1} = \{\delta \in \Gamma^+\colon\ \delta \text{ admissible in } \Theta_n\}$. If we let $\Theta = \bigcup_{n \geqslant 0} \Theta_n$, then as before it suffice to show $\Theta = \Gamma^+$. As in Theorem 3.3, $\Theta \subseteq \Gamma^+$ is immediate. For the converse inclusion, we can show by induction on n that $\Gamma_n^+ \subseteq \Theta$, whence $\bigcup_{n \geqslant 0} \Gamma_n^+ \subseteq \Theta$. □

Theorem 3.9. *Let $(\Gamma^+, \Gamma^-, \Gamma^*)$ be a $\leqslant$-minimal extension for a default theory (W, Δ). Then $(\Gamma^+, \Gamma^-, \Gamma^*)$ can be obtained as the limit of an inductive construction of the kind given in the proof of Theorem* 3.8.

Proof. Suppose we have defined sets of defaults Θ_n^+, Θ_n^-, and Θ_n^* (the first two subsets of Γ^+, Γ^-, respectively, and the third a superset of Γ^*), in such a way as to conform to the construction of Theorem 3.8. Then, if we put:

$$\Theta^+ = \bigcup_{n \geqslant 0} \Theta_n^+, \qquad \Theta^- = \bigcup_{n \geqslant 0} \Theta_n^*, \qquad \Theta^* = \bigcap_{n \geqslant 0} \Theta_n^*,$$

we have that $(\Theta^+, \Theta^-, \Theta^*) \leqslant (\Gamma^+, \Gamma^-, \Gamma^*)$, and moreover $(\Theta^+, \Theta^-, \Theta^*)$ is an extension. By minimality of $(\Gamma^+, \Gamma^-, \Gamma^*)$, we have

$$\left(\Theta^+, \Theta^-, \Theta^*\right) = \left(\Gamma^+, \Gamma^-, \Gamma^*\right),$$

and the conclusion follows.

So we need to show how to define Θ_n^+, Θ_n^-, and Θ_n^*. We define $\Theta_0^+ = \Theta_0^- = \emptyset$, and $\Theta_0^* =$ the set of defaults whose prerequisite is consistent. For the inductive step, we pick as Θ_{n+1}^+ the set of defaults, *all drawn from Γ^+*, that are admissible at the previous stage:

$$\Theta_{n+1}^+ = \left\{\delta \in \Gamma^+\colon\ C\left(\Theta_n^+\right) \models_W P(\delta)\right\}.$$

In order to show that this conforms to the construction, we need to establish that Θ_{n+1}^+ is a maximal set of defaults having the following properties: (i) $C(\Theta_n^+ \cup \Theta_{n+1}^+)$ is

W-consistent; (ii) every $\delta \in \Theta^+_{n+1}$ is admissible in Θ^+_n; and (iii) no default in Θ^+_{n+1} is pre-empted in $(\Theta^*_n - \Theta^-_n)$.

Clearly Θ^+_{n+1} has properties (i)–(iii); we need to show that it is maximal with those properties. Suppose for contradiction that some $\delta \notin \Theta^+_{n+1}$ could be adjoined to Θ^+_{n+1} preserving (i)–(iii). Then, in particular, δ would be admissible in Θ^+_n, and since $\delta \notin \Theta^+_{n+1}$, it must be that $\delta \notin \Gamma^+$.

On the other hand, since $\Theta^+_n \subseteq \Gamma^+$, we have that δ is admissible in Γ^+; and since $(\Gamma^* - \Gamma^-) \subseteq (\Theta^*_n - \Theta^-_n)$, we have that δ cannot be pre-empted in $(\Gamma^* - \Gamma^-)$ (not being pre-empted in the larger set). But since we assumed $(\Gamma^+, \Gamma^-, \Gamma^*)$ to be an extension, we would have $\delta \in \Gamma^+$, contradicting the conclusion reached at the end of the previous paragraph.

We conclude that Θ^+_{n+1} is a maximal subset of Δ satisfying (i)–(iii). Now, by defining Θ^-_{n+1} and Θ^*_{n+1} in the obvious way, we obtain that the sequence thus obtained conforms to the construction of Theorem 3.8. The limit $(\Theta^+, \Theta^-, \Theta^*)$ is therefore an extension and by construction,

$$\left(\Theta^+, \Theta^-, \Theta^*\right) \leqslant \left(\Gamma^+, \Gamma^-, \Gamma^*\right).$$

Minimality of $(\Gamma^+, \Gamma^-, \Gamma^*)$ gives the desired conclusion. □

Theorem 5.1. *Let* (W, Δ) *be a default theory, and suppose* $(\Gamma^+, \Gamma^-, \Gamma^*)$ *is a* 3*-tuple of subsets of* Δ *such that* $C(\Gamma^+) \models_W \varphi$. *Then* $(\Gamma^+, \Gamma^-, \Gamma^*)$ *is a general extension for* (W, Δ), *if and only if it is a general extension for* $(W \cup \{\varphi\}, \Delta)$.

Proof. We do the case for a categorical default theory, the general case being similar. Since $C(\Gamma^+) \models_W \varphi$, any default δ is conflicted in Γ^+ relative to W if and only if it is conflicted in Γ^+ relative to $W \cup \{\varphi\}$. Thus, in order to establish the theorem it suffices to establish that δ is pre-empted in $\Delta - \Gamma^-$ relative to W if and only if δ is so pre-empted relative to $W \cup \{\varphi\}$.

One direction is immediate: if δ is pre-empted in $\Delta - \Gamma^-$ relative to W then it is still so pre-empted relative to $W \cup \{\varphi\}$, by monotonicity of classical logic.

For the converse, assume δ is pre-empted in $\Delta - \Gamma^-$ relative to $W \cup \{\varphi\}$. Then:

$$C\left(\Delta - \Gamma^-\right) \models_{W \cup \{\varphi\}} \neg J(\delta); \tag{A.1}$$

But by hypothesis, $C(\Gamma^+) \models_W \varphi$, and by disjointness of Γ^+ and Γ^-, also $\Gamma^+ \subseteq (\Delta - \Gamma^-)$. By monotonicity of classical logic, $C(\Delta - \Gamma^-) \models_W \varphi$, whence by (A.1) and Cut (for classical logic), also $C(\Delta - \Gamma^-) \models_W \neg J(\delta)$, as desired. □

Theorem 5.4. *Suppose* $(W, \Delta) \mid\!\sim \varphi$ *and let* $\boldsymbol{\Gamma} = \lim \boldsymbol{\Gamma}_n$ *be a minimal extension for* $(W + \varphi, \Delta)$. *Then there is a minimal extension* $\boldsymbol{\Theta} = \lim \boldsymbol{\Theta}_n$ *for* (W, Δ) *such that* $\boldsymbol{\Theta} \leqslant \boldsymbol{\Gamma}$.

Proof. Since $(W, \Delta) \mid\!\sim \varphi$ and every default theory has a minimal extension, let $\boldsymbol{\Pi} = \lim \boldsymbol{\Pi}_n$ be a minimal extension for (W, Δ) such that $C(\Gamma^+) \models_W \varphi$. Let $k > 0$ be an integer such that already $C(\Gamma^+_k) \models_W \varphi$.

We are going to define a construction sequence $\boldsymbol{\Theta}_n$. For $m \leqslant k$ we put $\boldsymbol{\Theta}_m = \boldsymbol{\Pi}_m$. For $k+n$ (where $n > 0$) we put:

$$\begin{aligned}
\Theta^+_{k+n} \;=\;& \text{a maximal subset of } \Delta \text{ extending } \Gamma^+_n\text{, such that:}\\
& \text{(A)}\quad C\big(\Theta^+_{k+n-1} \cup \Theta^+_{k+n}\big) \text{ is } W\text{-consistent;}\\
& \text{(B)}\quad \text{every } \delta \in \Theta^+_{k+n} \text{ is admissible in } \Theta^+_{k+n-1}\text{;}\\
& \text{(C)}\quad \text{no } \delta \in \Theta^+_{k+n} \text{ is pre-empted in } \Theta^*_{k+n-1} - \Theta^-_{k+n-1}\text{;}\\
\Theta^-_{k+n} \;=\;& \big\{\delta\colon\ \delta \text{ pre-empted or conflicted in } \Theta^+_{k+n} \text{ relative to } W\big\};\\
\Theta^*_{k+n} \;=\;& \big\{\delta\colon\ P(\delta) \text{ is consistent with } C(\Theta^+_{k+n}) \text{ relative to } W\big\}.
\end{aligned}$$

Since any maximal subset of Δ extending Γ^+_n and having properties (A), (B), and (C) is also a maximal subset of Δ having properties (A), (B), and (C), we obtain immediately that $\boldsymbol{\Theta}_n$ is a construction sequence and that $\boldsymbol{\Theta} = \lim \Theta_n$ is a minimal extension for (W, Δ).

So we need to show $\boldsymbol{\Gamma} \leqslant \boldsymbol{\Theta}$; in turn, it suffices to show that $\boldsymbol{\Gamma}_n \leqslant \boldsymbol{\Theta}_{k+n}$. This we do by induction on n.

Case $n = 0$. Since we have $\Gamma^+_0 = \Gamma^-_0 = \emptyset$, all we need to show is that $\Theta^*_{k+0} \subseteq \Gamma^*_0 = \{\delta\colon\ P(\delta) \text{ is } W+\varphi\text{-consistent}\}$. Now,

$$\Theta^*_{k+0} = \{\delta\colon\ P(\delta) \text{ is } W+\varphi\text{-consistent}\}.$$

Since by construction $C(\Theta^+_{k+0}) \models_W \varphi$, we have that if $P(\delta)$ is $W+\varphi$-inconsistent then it is also W-inconsistent with $C(\Theta^+_{k+0})$. So if $\delta \notin \Gamma^*_0$ then also $\delta \notin \Theta^*_{k+0}$, as desired. This shows $\boldsymbol{\Gamma}_0 \leqslant \boldsymbol{\Theta}_{k+0}$.

Case $n+1$. Assume $\boldsymbol{\Gamma}_n \leqslant \boldsymbol{\Theta}_{k+n}$, in order to show $\boldsymbol{\Gamma}_{n+1} \leqslant \boldsymbol{\Theta}_{k+n+1}$. By construction, Θ^+_{k+n+1} extends Γ^+_{n+1}. In order to prove $\Gamma^-_{n+1} \subseteq \Theta^-_{n+1}$, recall that by definition:

$$\begin{aligned}
\Theta^-_{k+n+1} &= \big\{\delta\colon\ \delta \text{ pre-empted or conflicted in } \Theta^+_{k+n+1} \text{ relative to } W\big\};\\
\Gamma^-_{n+1} &= \big\{\delta\colon\ \delta \text{ pre-empted or conflicted in } \Gamma^+_{n+1} \text{ relative to } W+\varphi\big\}.
\end{aligned}$$

Now, if δ is pre-empted or conflicted in Γ^+_{n+1} relative to $W+\varphi$, since $C(\Theta^+_{k+n+1}) \models_W \varphi$ and $\Gamma^+_{n+1} \subseteq \Theta^+_{k+n+1}$, δ is also pre-empted or (respectively) conflicted in Θ^+_{k+n+1} relative to W. So $\Gamma^-_{n+1} \subseteq \Theta^-_{n+1}$.

Finally, ad $\Theta^*_{k+n+1} \subseteq \Gamma^*_{n+1}$. We have:

$$\begin{aligned}
\Theta^*_{k+n+1} &= \big\{\delta\colon\ P(\delta) \text{ is } W\text{-consistent with } \Theta^+_{k+n+1}\big\};\\
\Gamma^*_{n+1} &= \big\{\delta\colon\ P(\delta) \text{ is } W+\varphi\text{-consistent with } \Gamma^+_{n+1}\big\}.
\end{aligned}$$

Suppose $\delta \notin \Gamma^*_{n+1}$; then $C(\Gamma^+_{n+1}) \models_{W+\varphi} \neg P(\delta)$. Since $\Gamma^+_{n+1} \subseteq \Theta^+_{k+n+1}$ and $C(\Theta^+_{k+n+1}) \models_W \varphi$, also $C(\Theta^+_{k+n+1}) \models_W \neg P(\delta)$, whence $\delta \notin \Theta^*_{k+n+1}$ as desired.

This concludes the induction, showing $\boldsymbol{\Gamma} \leqslant \boldsymbol{\Theta}$. □

Theorem 5.5. *Suppose* $(W, \Delta) \mathrel{|\!\sim} \varphi$ *and let* $\boldsymbol{\Gamma} = \lim \boldsymbol{\Gamma}_n$ *be a minimal extension for* $(W+\varphi, \Delta)$. *Then there is a minimal extension* $\boldsymbol{\Theta} = \lim \boldsymbol{\Theta}_n$ *for* (W, Δ) *such that* $\boldsymbol{\Theta} \leqslant \boldsymbol{\Gamma}$.

Proof. Since $\boldsymbol{\Gamma}$ is a minimal extension for $(W+\varphi, \Delta)$ and $(W, \Delta) \mid\!\sim \varphi$, we know from the previous Theorem 5.4, that there is a minimal extension $\boldsymbol{\Pi}$ for (W, Δ) above $\boldsymbol{\Gamma}$ (in the $\leqslant$ ordering). Pick such an extension, which will remain fixed for the rest of the proof. We will show that there is an extension for (W, Δ) below $\boldsymbol{\Gamma}$: of course it will follow that there is a *minimal* extension for (W, Δ) below $\boldsymbol{\Gamma}$.

Such an extension $\boldsymbol{\Theta}$ that is below $\boldsymbol{\Gamma}$ (in the ordering $\leqslant$) will be defined as the limit of an increasing sequence, i.e., $\boldsymbol{\Theta} = \lim \boldsymbol{\Theta}_n$. For $n = 0$ we put $\Theta_0^+ = \Theta_0^- = \emptyset$, and $\Theta_0^* = \{\delta\colon\ P(\delta) \text{ is } W\text{-consistent}\}$.

For the inductive step:

$$\begin{aligned}\Theta_{n+1}^+ = {} & \text{a maximal subset of } \Pi^+ \text{ such that:}\\ & \text{(A)}\quad C(\Theta_n^+ \cup \Theta_n^+) \text{ is } W\text{-consistent};\\ & \text{(B)}\quad \text{every } \delta \text{ in } \Theta_{n+1}^+ \text{ is admissible in } \Theta_n^+ \text{ relative to } W;\\ & \text{(C)}\quad \text{no } \delta \in \Theta_{n+1}^+ \text{ is pre-empted in } \Theta_n^* - \Theta_n^-, \text{ relative to } W.\end{aligned}$$

Having done this, we define Θ_{n+1}^- and Θ_{n+1}^* as usual as the set of defaults pre-empted or conflicted in Θ_{n+1}^+ relative to W, and (respectively) as the set of defaults whose prerequisite is consistent with Θ_{n+1}^+ relative to W.

So we can put $\boldsymbol{\Theta} = \lim \boldsymbol{\Theta}_n$. We need to establish the following:

(1) the sequence $\boldsymbol{\Theta}_n$ is increasing in the ordering $\leqslant$; and
(2) $\boldsymbol{\Theta}$ is an extension for (W, Δ);
(3) $\boldsymbol{\Theta} \leqslant \boldsymbol{\Gamma}$.

The proof for the first two items is precisely similar to the corresponding items in the proof of Theorem 3.8, and we skip it. We concentrate on the last, crucial, item.

To show $\boldsymbol{\Theta} \leqslant \boldsymbol{\Gamma}$ suffices to show $\boldsymbol{\Theta}_n \leqslant \boldsymbol{\Gamma}$ by induction on n. The case for $n = 0$ is easy. We have $\Theta_0^+ = \emptyset \subseteq \Gamma^+$, and similarly for Θ_0^-. Now $\{\delta\colon\ P(\delta) \text{ is } W\text{-consistent}\} = \Theta_0^*$: if $P(\delta)$ is W-inconsistent then it is also $W + \varphi$ inconsistent with $C(\Gamma^+)$, so if $\delta \notin \Theta_0^*$ then $\delta \notin \Gamma^*$, i.e., $\Gamma^* \subseteq \Theta_0^*$.

Now the inductive step for $n+1$. We assume that $\delta \in \Theta_{n+1}^+$ in order to show that $\delta \in \Gamma^+$; in turn, this is established by proving that δ is admissible in Γ^+ relative to $W + \varphi$, not conflicted in Γ^+ relative to $W + \varphi$ and not pre-empted in $\Gamma^* - \Gamma^-$ relative to $W + \varphi$.

(1) δ is admissible in Θ_n^+ relative to W and hence (since $\Theta_n^+ \subseteq \Gamma^+$) also admissible in Γ^+ relative to $W + \varphi$.

(2) By inductive hypothesis, we have $(\Gamma^* - \Gamma^-) \subseteq (\Theta_n^* - \Theta_n^-)$. Suppose for contradiction that δ is pre-empted in $(\Gamma^* - \Gamma^-)$ relative to $W + \varphi$. Then

$$\bigl(\Gamma^* - \Gamma^-\bigr) \models_{W+\varphi} \neg J(\delta).$$

We show $(\Gamma^* - \Gamma^-) \models_W \varphi$, whence

$$\bigl(\Gamma^* - \Gamma^-\bigr) \models_W \neg J(\delta),$$

which in turn gives (by monotonicity of $\models_W$) $(\Theta_n^* - \Theta_n^-) \models_W \neg J(\delta)$, which is impossible.

To prove $(\Gamma^* - \Gamma^-) \models_W \varphi$: we have that $\boldsymbol{\Pi}$ is a minimal extension for (W, Δ): since $(W, \Delta) \mid\!\sim \varphi$, we have $C(\Pi^+) \models_W \varphi$. Since $\boldsymbol{\Pi}$ is an extension, all defaults in

Π^+ are admissible in Π^+ and hence they belong to Π^*; since Π^+ and Π^- are disjoint, we have $\Pi^+ \subseteq (\Pi^* - \Pi^-)$. It follows that $(\Pi^* - \Pi^-) \models_W \varphi$, and since $\boldsymbol{\Gamma} \leqslant \boldsymbol{\Pi}$ also $(\Gamma^* - \Gamma^-) \models_W \varphi$, as desired.

(3) Finally, we show that δ is not conflicted in Γ^+ relative to $W + \varphi$. If δ were so conflicted, then $C(\Gamma^+) \models_{W+\varphi} \neg C(\delta)$, and since $\Gamma^+ \subseteq \Pi^+$, also $C(\Pi^+) \models_{W+\varphi} \neg C(\delta)$.

But now, as before, $C(\Pi^+) \models_W \varphi$, so $C(\Pi^+) \models_W \neg C(\delta)$. In other words, δ is conflicted in Π^+ relative to W, which is impossible given that $\delta \in \Theta^+_{n+1} \subseteq \Pi^+$ and $\boldsymbol{\Pi}$ is an extension for (W, Δ).

As mentioned, this gives $\Theta^+_{n+1} \subseteq \Gamma^+$. From this, we can obtain (using the fact that $S \models_W \psi$ implies $S \models_{W+\varphi} \psi$) that $\Theta^-_{n+1} \subseteq \Gamma^-$ and $\Gamma^* \subseteq \Theta^*_{n+1}$. In turn, this gives $\boldsymbol{\Theta} \leqslant \boldsymbol{\Gamma}$. □

Theorem 5.7. *The relation $\mid\!\sim$ satisfies the properties of Cut, Reflexivity, and Stability.*

Proof. We take up the different properties in turn. As we will see, Theorem 5.6 will play a crucial rôle.

(1) *Reflexivity*: we need to show that if $\varphi \in W$ then $(W, \Delta) \mid\!\sim \varphi$. This follows immediately: if $(\Gamma^+, \Gamma^-, \Gamma^*)$ is any extension (minimal or otherwise) of the theory, and $\varphi \in W$, then in particular $C(\Gamma^+) \models_W \varphi$, so that $(W, \Delta) \mid\!\sim \varphi$.

(2) *Stability* or *Cautious Monotonicity*: we need to show

$$\frac{(W, \Delta) \mid\!\sim \varphi, \quad (W, \Delta) \mid\!\sim \psi}{(W + \varphi, \Delta) \mid\!\sim \psi}.$$

Since $(W, \Delta) \mid\!\sim \varphi$ and $(W, \Delta) \mid\!\sim \psi$, then for every $\leqslant$-minimal extension $\boldsymbol{\Gamma} = (\Gamma^+, \Gamma^-, \Gamma^*)$ of (W, Δ) we have $C(\Gamma^+) \models_W \varphi \wedge \psi$. By Theorem 5.6 each such extension is also a $\leqslant$-minimal extension of $(W + \varphi, \Delta)$, and conversely. By monotonicity of classical logic, $C(\Gamma^+) \models_{W+\varphi} \psi$, whence $(W + \varphi, \Delta) \mid\!\sim \psi$, as desired.

(3) *Cut*: we need to show

$$\frac{(W, \Delta) \mid\!\sim \varphi, \quad (W + \varphi, \Delta) \mid\!\sim \psi}{(W, \Delta) \mid\!\sim \psi}.$$

Let $\boldsymbol{\Gamma} = (\Gamma^+, \Gamma^-, \Gamma^*)$ be $\leqslant$-minimal extension of (W, Δ): we need to show $C(\Gamma^+) \models_W \psi$. By Theorem 5.1 and the first premise of Cut $\boldsymbol{\Gamma}$ is also a $\leqslant$-minimal extension of $(W + \varphi, \Delta)$. By the second premise of Cut, $C(\Gamma^+) \models_{W+\varphi} \psi$. But also $C(\Gamma^+) \models_W \varphi$: using Cut for classical logic, we have $C(\Gamma^+) \models_W \psi$, as desired. □

Theorem 6.1. *Every semi-normal default theory has a unique minimal extension.*

Proof. We do the case for (W, Δ) categorical, the general case being similar. We define, as in the proof of Theorem 1.6 a sequence of pairs of sets (Γ^+_n, Γ^-_n) of defaults, and we begin by putting $\Gamma^+_0 = \Gamma^-_0 = \emptyset$. For the inductive step we put:

$$\Gamma^+_{n+1} = \{\delta\colon \delta \text{ not pre-empted in } \Delta - \Gamma^-_n\};$$

$$\Gamma^-_{n+1} = \{\delta\colon \delta \text{ pre-empted or conflicted in } \Gamma^+_{n+1}\}.$$

We assume that the sequence (Γ_n^+, Γ_n^-) coincides with the one in the proof of Theorem 1.6 up to stage n, and show that this must be the case also at stage $n+1$. Observe that the inductive hypothesis yields, in particular, that $C(\Gamma_n^+)$ is W-consistent.

Now, Γ_{n+1}^+ has the property of being a maximal set of defaults not pre-empted in $\Delta - \Gamma_n^-$ (being the set of *all* such defaults). So if we can show that it also has the further property that $C(\Gamma_n^+ \cup \Gamma_{n+1}^+)$ is W-consistent, it will follow that Γ_{n+1}^+ is a maximal set of defaults having the *two* mentioned properties, and we will have recovered the construction given in the proof of Theorem 1.6.

So we show that $C(\Gamma_n^+ \cup \Gamma_{n+1}^+)$ is W-consistent. Suppose for contradiction that this fails. We know that $C(\Gamma_n^+)$ is W-consistent by itself, so that if $C(\Gamma_n^+ \cup \Gamma_{n+1}^+)$ is W-inconsistent, it must be that $\Gamma_{n+1}^+ \neq \emptyset$. We will contradict this last fact, showing that $\Gamma_{n+1}^+ = \emptyset$.

From the hypothesis that $C(\Gamma_n^+ \cup \Gamma_{n+1}^+)$ is W-inconsistent, it follows that there are defaults $\delta_1, \ldots, \delta_k \in \Gamma_{n+1}^+$ such that

$$C\bigl(\Gamma_n^+\bigr) \models_W \neg\bigl(C(\delta_1) \wedge \cdots \wedge C(\delta_k)\bigr). \tag{2}$$

Now we show that for each default δ_i among $\delta_1, \ldots, \delta_k$, we have $\delta_i \in (\Delta - \Gamma_n^-)$. Reasoning by reductio, suppose $\delta_i \in \Gamma_n^-$. Then δ_i is either conflicted or pre-empted in Γ_n^+. But δ_i is semi-normal, so that it cannot be conflicted without being already pre-empted. So this implies that δ_i is pre-empted in Γ_n^+, and since $\Gamma_n^+ \subseteq (\Delta - \Gamma_n^-)$, it follows that δ_i is pre-empted in $(\Delta - \Gamma_n^-)$. But by definition, this is equivalent to $\delta_i \notin \Gamma_{n+1}^+$, against assumption. We conclude that $\delta_1, \ldots, \delta_k \in (\Delta - \Gamma_n^-)$.

From (2) we have that

$$C(\delta_1), \ldots, C(\delta_n) \models_W \neg C\bigl(\Gamma_n^+\bigr),$$

where $\neg C(\Gamma_n^+)$ is to be construed as the negation of the conjunction of conclusions of defaults in Γ_n^+. Since $\delta_1, \ldots, \delta_k \in (\Delta - \Gamma_n^-)$, also

$$C\bigl(\Delta - \Gamma_n^-\bigr) \models_W \neg C\bigl(\Gamma_n^+\bigr).$$

But Γ_n^+ and Γ_n^- are disjoint, so that $\Gamma_n^+ \subseteq (\Delta - \Gamma_n^-)$. We conclude that $C(\Delta - \Gamma_n^-)$ must be W-inconsistent, so that any default is pre-empted in $(\Delta - \Gamma_n^-)$, whence $\Gamma_{n+1}^+ = \emptyset$. This is the contradiction we sought.

We conclude that the construction given here coincides with the one given in the proof of Theorem 1.6 in the case of semi-normal theories (or Theorem 3.8 in the noncategorical case), and therefore the limit of the sequence of pairs of sets of defaults yields a general extension. The process is now deterministic, so this extension is unique. We already know from Theorem 3.8 that any extension obtained in this way is minimal. Theorem 3.9 gives uniqueness. □

Theorem 6.5. *Let* (W, Δ) *be a semi-normal theory, and suppose* $\boldsymbol{\Gamma}$ *is a potential extension that is also a fixed point of* $\boldsymbol{\tau}$*. Then* $\boldsymbol{\Gamma}$ *is a general extension for* (W, Δ).

Proof. Let $\boldsymbol{\Theta}$ be a fixed point of $\boldsymbol{\tau}$. First we observe that since W is consistent, Θ^+ must also be W-consistent: if not, then any δ is pre-empted in Θ^+ relative to W, so $\Theta^- = \Delta$, and by disjointness $\Theta^+ = \emptyset$. Therefore, if Θ^+ is W-inconsistent it must be that W is already inconsistent.

Given that $\boldsymbol{\Theta}$ is a fixed point, one easily verifies the equations defining extensions. Similarly, it's immediate to see that if $P(\delta)$ is W-inconsistent then $\delta \notin \Theta^*$. So all is left to verify is that Θ^* contains all δ admissible in Θ^+: but if $C(\Theta^+) \models_W P(\delta)$ then by W-consistency of Θ^+ also $C(\Theta^+) \not\models_W \neg P(\delta)$, i.e., $P(\delta)$ is W-consistent with Θ^+, whence $\delta \in \Theta^*$ as desired. □

Theorem 6.7. *Let* (W, Δ) *be a default theory. Then there is an iteratively definable NT-extension* $\boldsymbol{\Theta}$ *for* (W, Δ).

Proof. We construct an NT-extension for (W, Δ) iteratively by putting: $\Theta_0^+ = \Theta_0^- = \emptyset$, and $\Theta_0^* = \{\delta\colon\ P(\delta) \text{ is } W\text{-consistent}\}$.

For the inductive step, as in the proof of Theorem 3.8, we put:

$$\begin{aligned} \Theta_{n+1}^+ = {} & \text{a maximal set of defaults such that:} \\ & \text{(A)} \quad C(\Theta_n^+ \cup \Theta_n^+) \text{ is } W\text{-consistent;} \\ & \text{(B)} \quad \text{every } \delta \text{ in } \Theta_{n+1}^+ \text{ is admissible in } \Theta_n^+ \text{ relative to } W; \\ & \text{(C)} \quad \text{no } \delta \in \Theta_{n+1}^+ \text{ is pre-empted in } \Theta_n^* - \Theta_n^-, \text{ relative to } W. \end{aligned}$$

Now for the new twist: we put

$$\Theta_{n+1}^- = \bigl\{\delta\colon\ \delta \text{ conflicted or pre-empted in } \Theta_{n+1}^+\bigr\} \cup \bigl\{\delta\colon\ C\bigl(\Theta^* - \Theta^-\bigr) \not\models_W P(\delta)\bigr\}.$$

On the other hand, the definition of Θ^* is the usual one: Θ^* is the set of defaults δ whose pre-requisite $P(\delta)$ is W-consistent with Θ^+.

The first thing to prove is that the sequence we obtain is increasing, i.e., that $\boldsymbol{\Theta}_n \leqslant \boldsymbol{\Theta}_{n+1}$ for each n. This can be shown by induction on n, just like in the proof of Theorem 3.8, except at the inductive step, where one shows $\Theta_{n+1}^- \subseteq \Theta_{n+2}^-$. This we do in some detail. So suppose $\delta \in \Theta_{n+1}^-$; to show $\delta \in \Theta_{n+2}^-$, we distinguish two cases.

(a) Suppose δ is pre-empted or conflicted in Θ_n^+; then using the inductive hypothesis we obtain that δ is pre-empted or conflicted in Θ_{n+1}^+ and hence $\delta \in \Theta_{n+2}^-$.

(b) The other case is $C(\Theta_n^* - \Theta_n^-) \not\models_W P(\delta)$. By the inductive hypothesis, $\Theta_{n+1}^* \subseteq \Theta_n^*$ and $\Theta_n^- \subseteq \Theta_{n+1}^-$. It follows that $(\Theta_{n+1}^* - \Theta_{n+1}^-)$ is a subset of $(\Theta_n^* - \Theta_n^-)$, so that $C(\Theta_{n+1}^* - \Theta_{n+1}^-) \not\models_W P(\delta)$. So again we have $\delta \in \Theta_{n+2}^-$.

Finally, we need to show that $\boldsymbol{\Theta} = \lim \boldsymbol{\Theta}_n$ is an NT-extension. The only added complication over the proof of Theorem 3.8 is when we check that

$$\text{if } C(\Theta^* - \Theta^-) \not\models_W P(\delta) \text{ then } \delta \in \Theta^-.$$

We can assume that δ is not conflicted or pre-empted in Θ^+ and hence not in any Θ_n^+, for if it is then we immediately get $\delta \in \Theta^-$. So suppose that $\delta \notin \Theta^-$. Then for each n we have $C(\Theta_n^* - \Theta_n^-) \models_W P(\delta)$. But the sequence $C(\Theta_n^* - \Theta_n^-)$ is *decreasing* in n, so eventually there must be a tuple of defaults $\delta_1, \ldots, \delta_k$ and some n such that

$$\delta_1, \ldots, \delta_k \in \bigcap_{m \geqslant n} \bigl(\Theta_m^* - \Theta_m^-\bigr),$$

and moreover

$$C(\delta_1), \ldots, C(\delta_k) \models_W P(\delta).$$

Then $\delta_1, \ldots, \delta_k \in (\Theta^* - \Theta^-)$ whence $C(\Theta^* - \Theta^-) \models_W P(\delta)$, as required. □

Appendix B. Infinitely many defaults

Here we indicate briefly how to extend the present approach to default theories (W, Δ) comprising infinitely many defaults. For simplicity, we consider only categorical default theories, which form a natural and well-behaved class. As observed in Section 1 (and in particular in the proof of Theorem 1.6), the hypothesis that Δ is finite is used only once, in order to obtain a certain combinatorial fact. The general case where Δ is infinite can be handled as follows.

Definition B.1. Let (W, Δ) be a categorical default theory, and $\delta, \gamma \in \Delta$. Then δ is *below* γ, denoted $\delta \prec \gamma$ iff and only if there are defaults $\delta_1, \ldots, \delta_n$ such that:

$$C(\delta_1), \ldots, C(\delta_n), C(\delta) \models_W \neg J(\gamma),$$

but for no *proper* subset $\Gamma_0 \subset \{\delta_1, \ldots, \delta_n, \delta\}$ we have $C(\Gamma_0) \models_W \neg J(\gamma)$.

In other words, $\delta \prec \gamma$ iff δ is part of a *minimal* tuple of defaults pre-empting γ.

Definition B.2. Let (W, Δ) be a categorical default theory. For each $\delta \in \Delta$, put:

$$\mathsf{rk}(\delta) = \sup\bigl\{\mathsf{rk}(\gamma) + 1\colon\ \gamma \in \Delta \ \&\ \gamma \prec \delta\bigr\},$$

if such a sup exists, and $\mathsf{rk}(\delta) = \infty$ otherwise (in which case we say that $\mathsf{rk}(\delta)$ is undefined).

Similarly, put $\mathsf{rk}(\Delta) = \sup\{\mathsf{rk}(\delta)\colon\ \mathsf{rk}(\delta) \text{ is defined}\}$.

It is then possible to carry out the construction of Theorem 1.6 transfinitely through the ordinal $\mathsf{rk}(\Delta)$ appropriately taking unions at limit stages. It is not difficult to see that this yields a general extension. The crucial step is that if $\delta_1, \ldots, \delta_n$ is a (minimal) n-tuple pre-empting δ then $\mathsf{rk}(\delta_i) < \mathsf{rk}(\delta)$ for every i, so we can apply an appropriate inductive hypothesis to conclude that such an n-tuple must have a member in Γ^-.

As an added bonus, we have the following: for any default theory (W, Δ), let the *well-founded part* of Δ be $\Theta = \{\delta\colon\ \mathsf{rk}(\delta) \text{ is defined}\}$. Then it can be verified that any general extension for (W, Δ) is a *classical* extension for (W, Θ).

References

[1] G.A. Antonelli, Defeasible inheritance on cyclic networks, Artificial Intelligence 92 (1997) 1–23.

[2] G.A. Antonelli, Defeasible reasoning as a cognitive model, in: The Parikh Project. Seven Papers in Honour of Rohit, Uppsala Prints and Preprints in Philosophy, No. 18, Department of Philosophy, Uppsala University, 1996, pp. 1–13.

[3] G.A. Antonelli, Virtuous circles, in: A. Gupta, A. Chapuis (Eds.), Special issue of the Journal of the Indian Council of Philosophical Research on "Circularity, Definition, and Truth", to appear.

[4] N.D. Belnap, A useful four-valued logic, in: J.M. Dunn, G. Epstein (Eds.), Modern Uses of Multiple-Valued Logic, Reidel, Dordrecht, 1977.

[5] S. Brass, On the semantics of supernormal defaults, in: Proc. IJCAI-93, Chambery, France, 1993, Vol. I, Morgan Kaufmann, San Mateo, CA, 1993, pp. 578–583.

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

[7] J.P. Delgrande, T. Schaub, W.K. Jackson, Alternative approaches to default logic, Artificial Intelligence 70 (1994) 167–237.

[8] J. Dix, Default theories of poole-type and a method for constructing cumulative versions of default logic, in: B. Neumann (Ed.), Proc. 10th European Conference on Artificial Intelligence (ECAI-92), Wiley, New York, 1992, pp. 289–293.

[9] D. Etherington, Formalizing nonmonotonic reasoning systems, Artificial Intelligence 31 (1987) 41–85.

[10] Y. Dimopoulos, V. Magirou, A graph theoretic approach to default logic, Inf. Comput., to appear.

[11] M. Fitting, The family of stable models, J. Logic Programming 17 (1993) 197–225.

[12] D.M. Gabbay, Theoretical foundations for nonmonotonic reasoning in expert systems, in: K. Apt (Ed.), Logics and Models of Concurrent Systems, Springer, Berlin, 1985, pp. 439–459.

[13] A. van Gelder, K.A. Ross, J.S. Schlipf, The well-founded semantics for general logic programs, J. ACM (1990).

[14] J.F. Horty, Some direct theories of nonmonotonic inheritance, in: D.M. Gabbay, C.J. Hogger, J.A. Robinson (Eds.), Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3, Nonmonotonic Reasoning and Uncertain Reasoning, Oxford Univ. Press, Oxford, 1994, pp. 111–187.

[15] J.F. Horty, R.H. Thomason, D.S. Touretzky, A skeptical theory of inheritance in nonmonotonic semantic networks, Artificial Intelligence 42 (1990) 311–348.

[16] H. Kautz, B. Selman, Hard problems for simple default logic, Artificial Intelligence 49 (1991) 243–279.

[17] S. Kripke, Outline of a theory of truth, J. Philos. 72 (1975) 690–716.

[18] W. Łukaszewicz, Considerations on default logic, in: Proc. Nonmonotonic Reasoning Workshop, New Paltz, NY, 1984, pp. 165–193.

[19] W. Łukaszewicz, Considerations on default logic, Computational Intelligence 4 (1988) 1–16.

[20] D. Makinson, General patterns in nonmonotonic reasoning, in: D.M. Gabbay, C.J. Hogger, J.A. Robinson (Eds.), Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3, Nonmonotonic Reasoning and Uncertain Reasoning, Oxford Univ. Press, Oxford, 1994.

[21] D. Makinson, K. Schlechta, Floating conclusions and zombie paths: Two deep difficulties in the "directly skeptical" approach to defeasible inheritance nets, Artificial Intelligence 48 (1991) 199–209.

[22] Z. Manna, A. Shamir, The theoretical aspects of the optimal fixedpoint, SIAM J. Comput. 5 (3) (1986) 414–426.

[23] J. McCarthy, Epistemological problems of artificial intelligence, in: Proc. IJCAI-77, Cambridge, MA, MIT Press, 1977.

[24] J. McCarthy, Circumscription—A form of nonmonotonic reasoning, Artificial Intelligence 13 (1980) 27–39, 171–172.

[25] D. McDermott, J. Doyle, Nonmonotonic logic I, Artificial Intelligence 13 (1980) 41–72.

[26] R.C. Moore, Semantical considerations on nonmonotonic logic, Artificial Intelligence 25 (1985) 75–94.

[27] C.H. Papadimitriou, M. Sideri, Default theories that always have extensions, Artificial Intelligence 69 (1994) 347–357.

[28] D.L. Poole, A logical framework for default reasoning, Artificial Intelligence 36 (1988) 27–48.

[29] T. Przymusiński, Three-valued nonmonotonic formalisms and semantics of logic programs, Artificial Intelligence 49 (1991) 309–343.

[30] R. Reiter, A logic for default reasoning, Artificial Intelligence 13 (1980) 81–132.

[31] R. Reiter, G. Criscuolo, On interacting defaults, in: Proc. IJCAI-81, Vancouver, BC, 1981, pp. 270–276.

[32] B. Selman, Tractable default reasoning, Ph.D. Thesis, University of Toronto, Toronto, Ont., 1990.