On the Semantics of Supernormal Defaults*

Stefan Brass
Institut fur Informatik, Universitat Hannover
Lange Laube 22, D-30159 Hannover, Fed. Rep. Germany
E-mail: sb@informatik.uni-hannover[.dbpj.de

Abstract

Our aim is to clarify which nonmonotonic consequence relation $\mid\!\sim_\Delta$ is given by a set Δ of "supernormal" defaults, i.e. defaults of the form $(true : \delta)/\delta$.

There are in fact a number of proposals for $\mid\!\sim_\Delta$ (e.g. the skeptical and the credulous semantics). In this paper we look at the space of all possible default semantics and try to characterize the known ones by their properties, especially the valid deduction rules.

For instance, it seems reasonable to require that any useful semantics should coincide with the original CWA if this is consistent. We might also want to allow proofs by case analysis. Then we get the skeptical semantics (assuming some other very natural deduction rules).

Our results are in fact completeness proofs for "natural deduction systems" based on different default semantics.

1 Introduction

In this paper, we consider "supernormal" [Brewka, 199l] defaults, i.e. defaults of the form *(true* : $\delta)/\delta$. For $\delta \equiv \neg p(X_1, \ldots, X_n)$ these are RKITER'S CWA-defaults which formalize the common implicit assumption of negations. But we allow any quantifier-free formula for δ, e.g.

$$flies(X) \leftarrow bird(X).$$

The intuitive semantics is that this rule should be assumed for as many *X* as possible. Although these defaults are very simple compared to the full default logic [Reiter, 1980], surprisingly many examples can be formalized with them. This restriction was suggested and investigated in [Poole, 1988; Brass and Lipeck, 1989, Brewka, 1991; Delgrande and Jackson, 1991; Dix, 1992). The advantages are that the definitions are much simpler, certain abnormalities can be avoided, and there are reasonably efficient theorem provers [Przymusinski, 1989; Ginsberg, 1989; Baker and Ginsberg, 1989; Brass, 1992a; Brass and Lipeck, 1992].

Of course, it is also possible to specify axioms φ, i.e. formulae which have to be fully satisfied (and usually there are many more axioms than defaults). In the flying ability example, the following axioms could be used:

$$bird(tweety).$$
$$penguin(polly).$$
$$bird(X) \leftarrow penguin(X).$$
$$\neg flies(X) \leftarrow penguin(X).$$

Given defaults Δ and axioms Φ, one can answer queries ψ. In the classical example the query could be

$$bird(X) \wedge flies(X)$$

("which birds can fly?"). Now a correct answer is a substitution θ for the variables of the query, e.g. $\langle X \triangleleft tweety \rangle$, such that the corresponding instance of the query follows from the axioms and the defaults, i.e.

$$\Phi \mid\!\sim_\Delta \psi\theta.$$

Of course, the consequence relation $\mid\!\sim_\Delta$ respecting the defaults Δ still has to be defined.

In the above example, it seems clear that any reasonable semantics should allow to conclude *flies(tweeiy)* and not *flies(polly).* But as soon as we have conflicts between defaults, e.g. because of disjunctive or incomplete information, there are different solutions. Probably everybody knows the credulous and the skeptical semantics, but there are other possible definitions of $\mid\!\sim_\Delta$ (see section 2). And, of course, we have to ask whether there are reasonable semantics which we do not know yet.

Therefore we try to approach this problem in a more abstract way, and look at the space of all possible default semantics. We can classify them by means of the valid deduction rules (and other properties, such as the preservation of consistency). There is a rich literature on nonmonotonic consequence relations (e.g. [Gabbay, 1985; Shoham, 1987; Makinson, 1989; Kraus *et al,* 1990; Brass, 1990; Dix, 1991; Makinson, 1992]), but our way of deriving a default semantics from given properties seems to be novel.

Of course, the soundness of certain deduction rules has been proved or disproved for the known default semantics. The aim of this paper is to show the completeness

*This work was partially supported by the CEC under the ESPRIT Working Group 607] 1S-CORE (Information Systems — Correctness and REusability), coordinated by AMILCAR SERNADAS.

of deduction systems for different default semantics. For automated theorem provers, the algorithms in the above cited literature may be better suited. But for proofs "by hand" the rules discussed here allow a much higher level with far less steps.

Additionally, our results show that if we require certain natural properties, the corresponding default semantics is uniquely determined. So the default semantics that we know are not as casual as it may seem at first.

2 Default Semantics

In this section, we define the notion of a default semantics and give some concrete examples.

We assume a signature Σ as given, which declares the predicates and constants of the application area. As a useful and common simplification, we exclude function symbols and require that Σ is finite. Instead of clauses, we allow arbitrary quantifier-free formulae, because they may be built by the deduction rules (also conjunctive defaults are sometimes useful, and $\delta_1 \wedge \delta_2$ is not equivalent to the two single defaults).

Definition 2.1 (Default): *A default is a quantifier-free Σ-formula δ.*

Definition 2.2 (Consequence Relation): *A consequence relation is any relation $\mathrel{|\!\!\sim}$ between sets Φ of quantifier-free Σ-formulae (the axioms) and variable-free Σ-formulae ψ (the query after replacing the variables by the answer-constants).*

We do not consider answers containing variables. This is again a simplification, but it is very natural, at least if the axioms, the defaults and the query are range-restricted (see [Brass, 1992b]).

Definition 2.3 (Default Semantics): *A default semantics is a mapping which assigns to every set Δ of defaults a consequence relation $\mathrel{|\!\!\sim}_\Delta$.*

Most default semantics refer to the single instances of the defaults, not to the default rules themselves:

Definition 2.4 (Default Instances): *We write Δ^* for the set of ground instances of formulae from Δ.*

One can view Δ as a mere shorthand for Δ^*, since in the following, we will only refer to Δ^* (and call the $\delta \in \Delta^*$ "defaults", too).

The simplest default semantics is the CWA [Reiter, 1978], which assumes a default instance δ (e.g., a negative ground literal) iff it is consistent with the axioms Φ, i.e. its negation does not follow from Φ:

Definition 2.5 (CWA): $\Phi \mathrel{|\!\!\sim}_{cwa(\Delta)} \psi$ *iff*

$$\Phi \cup \{\delta \in \Delta^* \mid \Phi \nvdash \neg\delta\} \vdash \psi.$$

It is well known that this definition may lead to inconsistencies if there are conflicting defaults. The simplest example is $\Phi := \{p \vee q\}$ and $\Delta := \{\neg p, \neg q\}$.

Therefore, we have to check the consistency of a set of defaults at a time:

Definition 2.6 (Extension): *$E \subseteq \Delta^*$ is a Δ-extension of Φ iff*

- *$\Phi \cup E$ is consistent, and*
- *$\Phi \cup E \cup \{\delta\}$ is inconsistent for every $\delta \in \Delta^* - E$.*

So an extension is simply a $\subseteq$-maximal set of default instances which is consistent with Φ. Note that our extensions are not closed under logical consequences. But if one takes $Th(\Phi \cup E)$, one gets exactly the original extensions from [Reiter, 1980], as proven in [Dix, 1992].

Of course, there may be more than one extension, but we need a single $\mathrel{|\!\!\sim}_\Delta$. There are different solutions to this problem, e.g. the careful, skeptical, and credulous semantics, and the CWA-semantics also fits nicely in this scheme (see lemma 2.8):

Definition 2.7 (Different Default Semantics):

- $\Phi \mathrel{|\!\!\sim}_{caref(\Delta)} \psi \ :\Longleftrightarrow\ \psi \in Th(\Phi \cup \bigcap E)$
- $\Phi \mathrel{|\!\!\sim}_{skept(\Delta)} \psi \ :\Longleftrightarrow\ \psi \in \bigcap Th(\Phi \cup E)$
- $\Phi \mathrel{|\!\!\sim}_{cred(\Delta)} \psi \ :\Longleftrightarrow\ \psi \in \bigcup Th(\Phi \cup E)$
- $\Phi \mathrel{|\!\!\sim}_{cwa'(\Delta)} \psi \ :\Longleftrightarrow\ \psi \in Th(\Phi \cup \bigcup E)$

(where $\bigcup$ and $\bigcap$ are taken over all extensions E).

Note that the careful semantics is not the same as the skeptical one: For instance, in the above example, we have the two extensions $\{\neg p\}$ and $\{\neg q\}$. Then the skeptical semantics will allow us to conclude $\neg p \vee \neg q$. The careful semantics may only use defaults contained in the intersection of the extensions (which is empty in this case). Given a set of negation defaults, the skeptical semantics turns disjunctions into "exclusive or", while the careful semantics interprets them as "inclusive or" (corresponding to the GCWA [Minker, 1982])[1].

Of course, the above list is not complete. For instance, it is common to select one of the extensions by some external means (e.g., "implementation defined", or by priorities between the defaults, etc.).

We now prove the above characterization of the CWA:

Lemma 2.8: $\Phi \mathrel{|\!\!\sim}_{cwa(\Delta)} \psi \iff \Phi \mathrel{|\!\!\sim}_{cwa'(\Delta)} \psi$.

Proof: The CWA assumes those $\delta \in \Delta^*$ s.t. $\Phi \nvdash \neg\delta$, i.e. $\Phi \cup \{\delta\}$ is consistent. But every consistent set of defaults can be extended to an extension [Brass and Lipeck, 1989]. The other direction is even simpler. □

3 Nonmonotonic Consequence Relations

In this section, we will consider deduction rules for nonmonotonic consequence relations and comment on their validity for the default semantics of the last section. The

[1] JÜRGEN DIX pointed out that $\Phi := \{p \vee q,\ r \leftarrow p \wedge q\}$, $\Delta := \{\neg p, \neg q, \neg r\}$ shows that *sometimes* even the careful semantics turns "$\vee$" into exclusive or ($\neg r$ is assumed and allows to conclude $\neg p \vee \neg q$). This is another indication that there may be many more useful default semantics than we currently know. (WGCWA protects disjunctions, but violates EQ (see later) and cannot be directly generalized to arbitrary Δ.)

rules REF, RW, AND, EQ, OR are taken from [Kraus *et al.*, 1990].

First, we should expect that $\mid\!\sim_\Delta$ allows at least to conclude the (instances of) preconditions, so that no information is lost:

Definition 3.1 (REF): $\mid\!\sim_\Delta$ *is reflexive iff*

$$\psi \in \Phi^* \implies \Phi \mid\!\sim_\Delta \psi.$$

In this and all of the following definitions we of course require that the implication holds for all sets of quantifier-free formulae Φ and all variable-free formulae ψ (φ ranges over all quantifier-free formulae, and δ over the elements of Δ^*).

All of the default semantics from section 2 have the property REF and it is hard to think of any reason why it should be violated (at least if Φ is consistent).

Next, we should require that $\mid\!\sim_\Delta$ behaves "logical" in the following sense:

Definition 3.2 (RW): $\mid\!\sim_\Delta$ *allows right weakening iff*

$$\Phi \mid\!\sim_\Delta \psi,\ \{\psi\} \vdash \psi' \implies \Phi \mid\!\sim_\Delta \psi'.$$

This rule is sound for all of the above default semantics.

A stronger requirement would be to allow arbitrary logical consequences of what we have already derived:

$$\begin{array}{ll} & \Phi \mid\!\sim_\Delta \psi_1,\ \ldots,\ \Phi \mid\!\sim_\Delta \psi_n,\ \{\psi_1, \ldots, \psi_n\} \vdash \psi \\ \implies & \Phi \mid\!\sim_\Delta \psi. \end{array}$$

This follows from RW together with the following rule:

Definition 3.3 (AND): $\mid\!\sim_\Delta$ *is closed under conjunction iff*

$$\Phi \mid\!\sim_\Delta \psi_1,\ \Phi \mid\!\sim_\Delta \psi_2 \implies \Phi \mid\!\sim_\Delta \psi_1 \wedge \psi_2.$$

All of the proposed semantics except the credulous one are closed under conjunction. Given $\Phi := \{p \vee q\}$, $\Delta := \{\neg p, \neg q\}$, we can credulously conclude $\neg p$ and q (from $E_1 := \{\neg p\}$) and $\neg q$ and p (from $E_2 := \{\neg q\}$); but we cannot conclude $p \wedge \neg p$, i.e. *false*.

The following is a weaker property, which is satisfied also by the credulous semantics:

Definition 3.4 (CLC): $\mid\!\sim_\Delta$ *allows conjunction with logical consequences iff*

$$\Phi \mid\!\sim_\Delta \psi_1,\ \Phi \vdash \psi_2 \implies \Phi \mid\!\sim_\Delta \psi_1 \wedge \psi_2.$$

The credulous semantics satisfies also the following restriction of AND to non-conflicting defaults:

Definition 3.5 (DA): $\mid\!\sim_\Delta$ *is closed under conjunction for non-conflicting defaults iff*

$$\begin{array}{ll} & \Phi \mid\!\sim_\Delta \delta_1,\ \ldots,\ \Phi \mid\!\sim_\Delta \delta_n,\ \Phi \nvdash \neg\delta_1 \vee \cdots \vee \neg\delta_n \\ \implies & \Phi \mid\!\sim_\Delta \delta_1 \wedge \cdots \wedge \delta_n. \end{array}$$

Of course, a default semantics should only look at the logical contents of Φ, not at the way the axioms are written down:

Definition 3.6 (EQ): $\mid\!\sim_\Delta$ *is compatible with left logical equivalence iff*

$$\Phi_1 \vdash \Phi_2,\ \Phi_2 \vdash \Phi_1,\ \Phi_1 \mid\!\sim_\Delta \psi \implies \Phi_2 \mid\!\sim_\Delta \psi.$$

The following property allows a case analysis for disjunctive preconditions:

Definition 3.7 (OR): $\mid\!\sim_\Delta$ *supports disjunctions in the preconditions iff*

$$\begin{array}{ll} & \Phi \cup \{\varphi_1\} \mid\!\sim_\Delta \psi,\ \Phi \cup \{\varphi_2\} \mid\!\sim_\Delta \psi \\ \implies & \Phi \cup \{\varphi_1 \vee \varphi_2\} \mid\!\sim_\Delta \psi. \end{array}$$

The careful semantics does not have this property, as can be seen in the standard example with $\Delta := \{\neg p, \neg q\}$: Here $\{p\} \mid\!\sim_{caref(\Delta)} \neg p \vee \neg q$ and $\{q\} \mid\!\sim_{caref(\Delta)} \neg p \vee \neg q$ both hold, but $\{p \vee q\} \mid\!\sim_{caref(\Delta)} \neg p \vee \neg q$ does not.

A weaker property, which is satisfied by the careful semantics, is the following:

Definition 3.8 (DOR): $\mid\!\sim_\Delta$ *supports disjunctions in the preconditions for deriving defaults iff*

$$\Phi \cup \{\varphi_1\} \mid\!\sim_\Delta \delta,\ \Phi \cup \{\varphi_2\} \mid\!\sim_\Delta \delta \implies \Phi \cup \{\varphi_1 \vee \varphi_2\} \mid\!\sim_\Delta \delta.$$

Of course, there are many more such rules, e.g. cumulation. But we do not need them in the results of section 4.

What is missing are properties allowing to conclude certain defaults. Up to now, the logical implication $\vdash$ would be perfect — it satisfies all the rules. But, of course, it is not a reasonable default semantics.

Naturally, we should only require that a default is assumed if this is really obvious. The results would be void if we use the default semantics itself in its own characterization.

We believe that the original CWA is the only sensible semantics if it is consistent. Obviously, we cannot assume more defaults, because this would immediately ruin the consistency (even if we would assume no other default). But there is also no reason why we should not adopt a default contained in the consistent CWA. And intuitively, a default should be assumed if there is no evidence to the contrary.

Definition 3.9 (CWA): $\mid\!\sim_\Delta$ *is weakly compatible with the CWA iff*

$$\Phi \mid\!\sim_{cwa(\Delta)} \delta,\ \Phi \not\mid\!\sim_{cwa(\Delta)} \mathit{false} \implies \Phi \mid\!\sim_\Delta \delta.$$

This property is satisfied by all the default semantics of section 2, which is an additional argument for its adequacy.

If we want to use it in deductions, we need a global consistency proof for the CWA, which might be quite difficult. The following property avoids this and looks only at the default we want to conclude and at possibly conflicting defaults:

Definition 3.10 (UCD): $\mid\!\sim_\Delta$ *satisfies uncritical defaults iff* $\Phi \mid\!\sim_\Delta \delta$ *for every* $\delta \in \Delta^*$ *and consistent* Φ *such that there are no* $\delta_1, \ldots, \delta_n \in \Delta^*$, $n \geq 0$ *with*

- $\Phi \cup \{\delta_1, \ldots, \delta_n\}$ *is consistent, and*
- $\Phi \cup \{\delta_1, \ldots, \delta_n\} \vdash \neg\delta$.

This property is satisfied by all the known default semantics. And it is very natural, because one typically explains the applicability of a default by an argument like "well, there is no way to conclude $\neg\delta$, so we can assume δ". The property UCD formalizes this (especially the "no way" has to take into account other defaults).

But the property UCD is strictly stronger than the CWA property, so it is automatically more questionable.

Lemma 3.11: *If $\vdash\!\sim_\Delta$ has the property UCD, then it has also the property CWA.*

Proof: The original CWA is consistent if and only if there are no defaults $\delta_1, \ldots, \delta_m \in \Delta^*$, $m \geq 2$, such that $\Phi \vdash \neg\delta_1 \vee \cdots \vee \neg\delta_m$ and $\Phi \not\vdash \neg\delta_i$ for $i = 1, \ldots, m$ (this is a simple generalization of a result from [Shepherdson, 1984]). Then obviously only the case $n = 0$ in definition 3.10 is interesting and it collapses to the definition of the CWA. □

The property UCD in fact requires that at least all the defaults contained in the intersection of the extensions can be concluded by the default semantics:

Lemma 3.12: *If $\delta \in \Delta^*$ can be concluded by the property UCD, then it is contained in every extension.*

Proof: Suppose it would not be contained in an extension E. Then choose $\{\delta_1, \ldots, \delta_n\} = E$ (E is finite because of our constraints on Σ, but this lemma holds also in the general case, one only has to apply the compactness theorem). By definition $\Phi \cup E$ is consistent and $\Phi \cup E \cup \{\delta\}$ is inconsistent, i.e. $\Phi \cup E \vdash \neg\delta$. □

Therefore characterizations of the careful semantics using UCD are not very interesting (we will give one using CWA and DOR).

Finally, the following rule formalizes the idea that defaults are assumed in the absence of information to the contrary:

Definition 3.13 (IMD): *$\vdash\!\sim_\Delta$ is inverse monotonic for defaults iff*

$$\Phi \cup \{\varphi\} \vdash\!\sim_\Delta \delta \implies \Phi \vdash\!\sim_\Delta \delta.$$

This is a very strong property, since it ignores conflicts with additionally assumable defaults. The original CWA and the credulous semantics have this property, the skeptical and the careful semantics violate it (consider $\Phi := \{p \vee q\}$, $\varphi := p$, and $\Delta := \{\neg p, \neg q\}$).

Table 1 summarizes which deduction rules are valid for which default semantics. We see that the interesting properties for classifying default semantics are AND, OR and IMD. All other deduction rules hold for all of the four semantics (but they are needed in the completeness proofs of the next section).

	caref	skept	cred	cwa
REF	•	•	•	•
RW	•	•	•	•
AND	•	•	—	•
CLC	•	•	•	•
DA	•	•	•	•
EQ	•	•	•	•
OR	—	•	•	•
DOR	•	•	•	•
CWA	•	•	•	•
UCD	•	•	•	•
IMD	—	—	•	•

Table 1: Validity of deduction rules

4 Characterizations

This section contains our main results which characterize the default semantics by their properties.

These characterizations are of the form "the semantics X is the weakest one satisfying the properties Y". We call a default semantics weaker than another one iff all the conclusions valid with respect to the former are also valid with respect to the latter:

Definition 4.1 (Weaker Than): *$\vdash\!\sim_\Delta$ is (non-strictly) weaker than $\vdash\!\sim'_\Delta$ iff*

$$\Phi \vdash\!\sim_\Delta \psi \implies \Phi \vdash\!\sim'_\Delta \psi.$$

For instance, the semantics studied here form a chain in the sequence $\vdash\!\sim_{caref(\Delta)}$, $\vdash\!\sim_{skept(\Delta)}$, $\vdash\!\sim_{cred(\Delta)}$, $\vdash\!\sim_{cwa(\Delta)}$ (from weakest to strongest).

The weakest default semantics satisfying certain deduction rules is that one which allows to conclude exactly the derivable formulae. Therefore a characterization of this form simply states the soundness and completeness of the deduction rules.

Theorem 4.2: *$\vdash\!\sim_{skept(\Delta)}$ is the weakest consequence relation satisfying REF, RW, AND, EQ, OR, and CWA.*

Proof: The soundness of these rules is trivial or already proven in the literature. So we only have to prove the completeness, i.e. that there is no weaker semantics.

Let $\vdash\!\sim_\Delta$ be any other consequence relation with these properties, and let $\Phi \vdash\!\sim_{skept(\Delta)} \psi$. We have to show that $\Phi \vdash\!\sim_\Delta \psi$. The proof is by induction on the number of extensions of Φ (there are only finitely many because of our restrictions on Σ).

If there is only one extension E, then the CWA is consistent, so we get all $\delta \in E$ by the property CWA. Since $\Phi \cup E \vdash \psi$, by HERBRAND's theorem $\Phi^* \cup E \vdash \psi$, and it is easy to derive ψ with REF, RW and AND.

Now assume that there are at least two extensions and let E be one of them. Let $\delta_1, \ldots, \delta_n \in \Delta^*$ be those defaults which are not contained in E, but in some other extension. Then we consider $\Phi_i := \Phi \cup \{\delta_i\}$ and $\Phi' := \Phi \cup \{\neg\delta_1 \wedge \cdots \wedge \neg\delta_n\}$.

Obviously, an extension of one of the Φ_i is also an extension of Φ. So it cannot have more extensions and it follows that $\Phi_i \mid\!\sim_{skept(\Delta)} \psi$. On the other hand, Φ_i has at least one fewer extension, namely E. Therefore we can apply the inductive hypothesis and conclude $\Phi_i \mid\!\sim_\Delta \psi$.

The only remaining case is Φ'. It has the extension E and only that (all defaults not contained in E are explicitly excluded). So we get $\Phi' \mid\!\sim_\Delta \psi$ (as in the base step).

Now we apply iteratively the property OR to "disjoin" $\Phi_1, \ldots, \Phi_n, \Phi'$. This results in the formula

$$\delta_1 \vee \cdots \vee \delta_n \vee (\neg\delta_1 \wedge \cdots \wedge \neg\delta_n)$$

which is a tautology, and therefore can be removed by the property EQ. So we finitely get $\Phi \mid\!\sim_\Delta \psi$. □

We can use also the property UCD instead of CWA:

Corollary 4.3: $\mid\!\sim_{skept(\Delta)}$ *is the weakest consequence relation satisfying REF, RW, AND, EQ, OR, and UCD.*

Proof: This follows directly from theorem 4.2 because on the one hand, UCD implies CWA (lemma 3.11), and on the other hand, $\mid\!\sim_{skept(\Delta)}$ has this property (lemma 3.12). □

A characterization of the careful semantics uses the weaker property DOR instead of OR:

Theorem 4.4: $\mid\!\sim_{caref(\Delta)}$ *is the weakest consequence relation satisfying REF, RW, AND, EQ, DOR, and CWA.*

Proof: The proof is very similar to that of theorem 4.2. But instead of directly trying to conclude ψ, we derive the defaults $\delta_1, \ldots, \delta_n \in \Delta^*$ which are contained in the intersection of the extensions and were used to derive ψ. For them, we can use DOR instead of OR. The rest directly follows from REF, RW, and AND. □

As explained above, a characterization of the careful semantics through the property UCD is nearly trivial:

Corollary 4.5: $\mid\!\sim_{caref(\Delta)}$ *is the weakest consequence relation satisfying REF, RW, AND, and UCD.*

The credulous and the CWA semantics allow the use of the rule IMD:

Theorem 4.6: $\mid\!\sim_{cred(\Delta)}$ *is the weakest consequence relation satisfying REF, RW, CLC, DA, CWA, and IMD. This result holds also with CWA replaced by UCD.*

Proof: Let $\Phi \mid\!\sim_{cred(\Delta)} \psi$. So there is an extension E of Φ with $\Phi \cup E \vdash \psi$, i.e. there are $\delta_1, \ldots, \delta_n \in E$ and $\varphi_1, \ldots, \varphi_m \in \Phi^*$ with $\varphi_1 \wedge \cdots \wedge \varphi_m \wedge \delta_1 \wedge \cdots \wedge \delta_n \vdash \psi$. We have to show that $\Phi \mid\!\sim_\Delta \psi$ for any consequence relation $\mid\!\sim_\Delta$ with the given properties.

Let $\delta'_1, \ldots, \delta'_k$ be those defaults not contained in E, but in some other extension. Let $\Phi' := \Phi \cup \{\neg\delta'_1, \ldots, \neg\delta'_k\}$. As above, we find that Φ' has only one extension, namely E. So the CWA-property or the UCD-property allow to conclude $\Phi' \mid\!\sim_\Delta \delta_i$ for $i = 1, \ldots, n$.

But then we get $\Phi \mid\!\sim_\Delta \delta_i$ for $i = 1, \ldots, n$ by the rule IMD. Since the δ_i are contained in the same extension, there is no conflict between them and rule DA results in $\Phi \mid\!\sim_\Delta \delta_1 \wedge \cdots \wedge \delta_n$. Now we use CLC to derive $\Phi \mid\!\sim_\Delta \varphi_1 \wedge \cdots \wedge \varphi_k \wedge \delta_1 \wedge \cdots \wedge \delta_n$ and finitely RW to get $\Phi \mid\!\sim_\Delta \psi$. □

If we require the stronger property AND instead of CLC and DA, we get the CWA:

Theorem 4.7: $\mid\!\sim_{cwa(\Delta)}$ *is the weakest consequence relation satisfying REF, RW, AND, CWA, and IMD. This result holds also with CWA replaced by UCD.*

Proof: Let $\Phi \mid\!\sim_{cwa(\Delta)} \psi$, so there are $\delta_1, \ldots, \delta_n \in \Delta^*$ and $\varphi_1, \ldots, \varphi_m \in \Phi^*$ with $\Phi \not\vdash \neg\delta_i$ $(i = 1, \ldots, n)$ and $\varphi_1 \wedge \cdots \wedge \varphi_m \wedge \delta_1 \wedge \cdots \wedge \delta_n \vdash \psi$. Because of lemma 2.8, there are extensions E_i containing δ_i $(i = 1, \ldots, n)$.

We again have to show $\Phi \mid\!\sim_\Delta \psi$ for any consequence relation satisfying the properties of the theorem. As in the proof of theorem 4.6, we get $\Phi \mid\!\sim_\Delta \delta_i$ $(i = 1, \ldots, n)$ by means of CWA (or UCD) and IMD. Property REF allows to conclude $\Phi \mid\!\sim_\Delta \varphi_i$ $(i = 1, \ldots, m)$, so AND results in $\Phi \mid\!\sim_\Delta \varphi_1 \wedge \cdots \wedge \varphi_m \wedge \delta_1 \wedge \cdots \wedge \delta_n$, and RW does the rest. □

So the properties REF, RW, AND, CWA, and IMD are only satisfied by $\mid\!\sim_{cwa(\Delta)}$ or a stronger semantics. But $\mid\!\sim_{cwa(\Delta)}$ already destroys the consistency, so in effect we cannot require all these properties together.

5 Conclusions

In this paper we considered different semantics for supernormal defaults and their properties as nonmonotonic consequence relations. Our main results are completeness proofs for certain deduction systems. On the one hand, this should help to better understand and motivate these semantics, and on the other hand, the deduction systems have some practical use for proofs "by hand".

Of course, we are still at the beginning of a general theory of default semantics. Many interesting extensions remain to be explored.

For instance, defaults with priorities are practically very important. If we use a prioritized CWA as basis, the characterizations of the skeptical and the careful semantics can be generalized [Brass, 1992b]. But the extension of the CWA to partially ordered defaults is not as clear as it might seem at first: We noted in [Brass, 1992a] that the constructive semantics proposed in [Brewka, 199l] is in fact different from any "preferential model" approach. So there are at least two incompatible notions of "extension" in this setting, and we should study properties formalizing the prioritization.

Of course, one should also try to generalize our results to full default logic.

Another direction for future work is to find previously unknown semantics by requiring certain properties. For instance, the rationality or rational monotonicity

$$\Phi \mid\!\sim_\Delta \psi,\ \Phi \not\mid\!\sim_\Delta \neg\varphi \implies \Phi \cup \{\varphi\} \mid\!\sim_\Delta \psi$$

is not satisfied by the skeptical and the careful semantics. So what happens if we require this rule in addition?

Note that here reference is made to $\not\vdash_\Delta$, so we cannot simply take the set of derivable formulae as the standard semantics. Other properties require that something should not be derivable (e.g., false), or give more indefinite information about the derivable formulae (e.g., the expansion property [Brass, 1990]). With such properties it might be possible to show that there is only a unique default semantics satisfying them - or none at all.

Acknowledgments

I would like to thank UDO LIPECK for suggesting the development of a system for natural deduction with defaults. He supervised my doctoral thesis which contains similar results to the ones presented here (among other things).

I would also like to thank JURGEN Dix, whose comments were most helpful.

References

[Baker and Ginsberg, 1989] Andrew B. Baker and Matthew L. Ginsberg. A theorem prover for prioritized circumscription. In Proc. 11th International Joint Conf. on Artificial Intelligence (IJCAI), pages 4C3-467, 1989.

[Brass and Lipeck, 1989] Stefan Br ass and Udo W. Lipeck. Specifying closed world assumptions for logic databases. In Janos Demetrovics and Bernhard Thaiheim, editors, 2nd Symposium on Mathematical Fundamentals of Database Systems (MFDBS'89), number 364 in LNCS, pages G8-84. Springer-Verlag, 1989.

Brass an d Lipeck, 1992] Stefan B rass and Udo W. Lipeck. Generalized bottom-up query evaluation. In Advances in Database Technology — EDBT'92, 3rd International Conference, number 580 in LNCS, pages 88-103. Springer-Verlag, 1992.

[Brass, 1990] Stefan Brass. Beginnings of a theory of general database completions. In Serge Abiteboul and Paris C. Kanellakis, editors, Third International Conference on Database Theory (ICDT'90), number 470 in LNCS, pages 349 363. Springer-Verlag, 1990.

[Brass, 1992a] Stefan Brass. Deduction with supernormal defaults. In Gerd Brewka, Klaus P. Jantke, and Peter H. Schmitt, editors, Nonmonotonic and Inductive Logics, 2nd International Workshop (NIL '91), number 659 in LNAI, pages 153 174. Springer-Verlag, 1992.

[Brass, 1992b] Stefan Brass. Defaults in Deduktiven Datenbanken. Doctoral thesis, Universitat Hannover, 1992. In German. English version in preparation.

[Brewka, 199l] Gerhard Brewka. Nonmonotonic Reasoning: Logical Foundations of Commonsense. Cambridge University Press, Cambridge, 1991.

[Delgrande and Jackson, 1991] James P. Delgrande and W. Ken Jackson. Default logic revisited. In James Allen, Richard Fikes, and Erik Sandewall, editors, Principles of Knowledge Representation and Reasoning, Proc. of the Second Int. Conf (KR'91), pages 118-127. Morgan Kaufmann, 1991.

[Dix, 1991] Jurgen Dix. Classifying semantics of logic programs. In Proc. of the 1st Int. Conf. on Logic Programming and Non-Monotonic Reasoning. MIT Press, 1991.

[Dix, 1992] Jurgen Dix. Default theories of Poole-type and a method for constructing cumulative versions of default logic. In Bernd Neumann, editor, Proc. of the 10th European Conf on Artificial Intelligence (ECAI 92), pages 289 293. John Wiley & Sons, 1992.

[Gabbay, 1985] Dov M. Gabbay. Theoretical foundations for non-monotonic reasoning in expert systems. In Krzysztof R. Apt, editor, Logics and Models of Concurrent Systems, pages 439-457. Springer-Verlag, 1985.

[Ginsberg, 1989] Matthew L. Ginsberg. A circumscriptive theorem prover. Artificial Intelligence, 39:209 230, 1989.

[Kraus et al., 1990] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence, 44:167-207, 1990.

[Makinson, 1989] David Makinson. General theory of cumulative inference. In Non-Monotonic Reasoning (2nd International Workshop), number 346 in LNAI, pages 1-18. Springer-Verlag, 1989.

[Makinson, 1992] David Makinson. General patterns in nonmonotonic reasoning. In Dov M. Gabbay, editor, Handbook of Logic in Artificial Intelligence and Logic Programming, volume 2. Oxford University Press, 1992.

[Minker, 1982] Jack Minker. On indefinite databases and the closed world assumption. In D. W. Loveland, editor, 6th Conference on Automated Deduction, number 138 in LNCS, pages 292-308. Springer-Verlag, 1982.

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

[Przymusinski, 1989] Teodor C. Przymusinski. An algorithm to compute circumscription. Artificial Intelligence, 38:49-73, 1989.

[Reiter, 1978] Raymond Reiter. On closed world data bases. In Herve Gallaire and Jack Minker, editors, Logic and Data Bases, pages 55 76, New York, 1978. Plenum.

[Reiter, 1980] Raymond Reiter. A logic for default reasoning. Artificial Intelligence, 13:81 132, 1980.

[Shepherdson, 1984] John C. Shepherdson. Negation as failure: A comparison of Clark's completed data base and Reiter's closed world assumption. The Journal of Logic Programming, 1:51 79, 1984.

[Shoham, 1987] Yoav Shoham. Nonmonotonic logics: Meaning and utility. In Proc. 10th International Joint Conference on Artificial Intelligence (1JCA1), pages 388 393, Milan, 1987.

Our papers are available from the anonymous ftp-server (130.75.26.1).