DEFAULTS AS GENERALIZED QUANTIFIERS

Karl Schlechta
Laboratoire d'Informatique de Marseille, URA CNRS 1787
Université de Provence, Case 32
3, Pl. Victor Hugo
F-13331 Marseille Cedex 3, France
ks@gyptis.univ-mrs.fr

December 15, 2008

Abstract

We interpret (open normal) defaults as generalized First Order Logic-quantifiers, give a semantics and a corresponding sound and complete axiom system. Nested and negated defaults are admissible and have a clear meaning. Moreover, the logic provides a notion of consistency for default theories, which is used for a theory revision approach in an order sorted language.

1 INTRODUCTION

The work described here had two aims at the outset:

First, to give sentences like "It is not true that the museum is normally closed on Mondays." meaning. Thus, we wanted to be able to handle negated defaults in a meaningful way.

Second, the default reasoning was to take place in the framework of an order sorted language, and the natural specificity criterion given by sorts should be exploited in conflict resolution.

To answer the first requirement, we have given open normal defaults (following Reiter's terminology, see [Rei80]) an interpretation as generalized quantifiers: "Normally, birds fly" is translated into "Most birds fly" or "The elements of a large or important subset of the set of birds fly", and written $\nabla x bird(x) : fly(x)$, in an extension of the language of first order logic. (We will use "large" and "important" etc. interchangeably - what really matters will be the formal definition.) This is given a clear semantics in the form of a system of such "important" subsets of the set of birds. The system will be defined to be almost a filter. We thus add additional structure - in the form of such a system - to first order models to provide a semantics for the new quantifier. In such an enlarged structure, the default "Normally, birds fly" will hold, iff there is an "important" subset of the set of birds, all of whose elements fly.

As we work in the full first order framework, we have achieved far more than we first aimed for: Not only negated defaults - in our birds example, there would then be no important subset of "birds" all of whose elements fly - have a clear meaning, but so do arbitrary boolean combinations of defaults, nested defaults, arbitrary combinations of defaults with classical quantifiers etc.

To complete the picture, we extend first order logic to a sound and complete axiomatisation for our semantics, allowing us not only to derive defaults - e.g., if "normally, $\phi(x)$", and $\phi(x)$ implies $\psi(x)$, then also "normally, $\psi(x)$" - but giving us a notion of consistency of default theories too. For instance, the theory containing the axioms "normally, $\phi(x)$", and "normally, $\neg\phi(x)$", is simply ruled out as inconsistent, as, we think, it should be. Thus, we have also obtained a criterion of the admissible - the consistent - any logic should have.

It is precisely this notion of consistency, which is exploited to answer the second problem: Inheritance and conflict resolution of defaults by specificity. Consistency allows us to detect conflicts (e.g. of two or more defaults), they are resolved by neglecting the less specific ones and proceeding sceptically when both are of equal specificity. We thus take a Theory Revision approach, using the consistency criterion of our generalized quantifier logic and exploiting specificity to select the "better" information.

Before going more into details, we outline the framework of the IBM LILOG project, into which default reasoning was to be incorporated.

Defaults in the LILOG Framework IBM Germany's Stuttgart LILOG project is a natural language (LILOG = Linguistics + Logic) project, "simulating" a travel guide. A German text describing a city (Düsseldorf) is fed into the computer, and natural language questions can then be asked about the city. Typical sentences and questions are e.g. (in translation) "The Hetjens Museum was inaugurated in 1909 and is situated in the Palais Nesselrode, which was built in the 18th century." and "Where is the Hetjens Museum?". Of course, there are hard linguistic restrictions on texts and questions, but these need not concern us here. Typical non-monotonic information in the present context is e.g.: "Museums are normally closed on Mondays", "Buildings normally have windows.", connected by "hard" (=classical) information as: "All museums are buildings.", but also exceptions in the sense of "This museum is not normally closed on Mondays.".

More precisely, LILOG's knowledge representation language is an order sorted language, a (logically equivalent, but computationally non-trivial) extension of first order predicate calculus, see [BHR90]. Sorts are e.g. "Streets", "Buildings", "Museums" etc. They are partially ordered by set-inclusion. This inclusion relation is determined by one component of the inference system, we can assume it as given and fixed. (In particular, it is monotonic - in contrast to defeasible inheritance, where the subset relation is itself defeasible.) There is hard information beyond the subset relation, represented by classical first order formulas, as "Museums have entrances.", and "soft" (=default) information. The default information is "appended" to sorts as in the examples above, to "Museums" and "Buildings", it speaks about the elements of those sorts.

Obviously, conflicts between hard and default information and within default information alone can arise. (For simplicity, the hard information alone is assumed to be conflict-free.) Basically, they have the form of the famous Tweety preclusion diagram and the Nixon diamond:

Tweety is a penguin, and all penguins are birds (Hard information). Most birds fly, most penguins don't (Soft information). Does Tweety fly? Answer: No, the more specific (by set inclusion) information should win.

Nixon is both Republican and Quaker (Hard). Most Quakers are pacifists, most Republicans not (Soft). Is Nixon a pacifist? Answer: Don't know, no specificity criterion is available.

We can thus make our task more precise:

We have a partially ordered set of information (which we treat as axioms). More specific information is considered more reliable, but comparison is not always possible. Hard (=classical) information is the most reliable one. If there are any conflicts, we have to single out a reasonable subset of the information which does not contain any conflicts any more.

The notion of conflict is made precise in the obvious way: A conflict is an inconsistency of defaults (and classical information) in our logic of generalized quantifiers. We can thus consider our problem as one of Theory Revision: Given a logic, a partial order on a (possibly) inconsistent set of axioms, single out a reasonable consistent subset of these axioms.

Remark: This is a generalization over "traditional" Theory Revision (in the sense of Gärdenfors, Makinson et al., see e.g. [AGM85], [Gär88], [GM88], [Mak85]), as our order is only partial ("epistemic entrenchment" orders are total), and Theory Revision was primarily intended for monotonic logic.

We have tried to obey the following three postulates of NMR:

Postulate 1: In case of conflict, "better" information wins over "lesser" information.

Postulate 2: In case of conflict between information of the same quality, we may

- come to no conclusion
- give one or both information (extensions)
- give the common part of both

This choice should be made sensibly and coherently.

Postulate 3: Contradictions should be limited, no EFQ (ex falso quod libet).

The reader will notice that the adopted solution (Approach 3 in Chapter 3 of the paper) offers a lot of flexibility, e.g. in the choice of the function f introduced at the beginning of Section 3.2.

Further Motivation to Consider Defaults as Generalized Quantifiers As the interpretation of defaults as generalized quantifiers is not a very standard one - to our knowledge, we were the first to suggest it, see e.g. the presentation in [Lor89] or [Lor90] - we will now further motivate it.

Let me first emphasize that we do not think that we have necessarily covered every aspect of defaults. Considering defaults as inference rules has a dynamic aspect, an "inference pressure", which our formalism does not have: Apply the default as much as you can! On the other hand, e.g the Reiter formalism [Rei80] lacks the static, minimal requirement: If you say that normally birds fly, you have to produce at least one such flying bird, and even more, a substantial subset of the set of birds must be able to fly. If we accept that these are minimal conditions justifying a sentence like "normally $\phi(x)$", then we conclude that "normally $\phi(x)$" and "normally $\neg\phi(x)$" are impossible, on the premise that two such substantial subsets (of one base set) have non-empty intersection. In our opinion, this minimal condition - $\phi(x)$ has to hold for a substantial subset to justify the sentence "normally $\phi(x)$" - leads naturally to a formalisation of "normally" via "important" subsets: "normally $\phi(x)$", or, in our notation $\nabla x\phi(x)$, holds iff there is an important subset A (of the universe, or of the set birds etc.) such that for all $x \in A$ $\phi(x)$ holds.

We have chosen a very weak formalisation of this concept, weaker than a filter: If B is the base set, then any system $\mathcal{N}(B)$ of important or large subsets of B has to satisfy

1. $B \in \mathcal{N}(B)$
2. $A \subseteq A' \subseteq B$, $A \in \mathcal{N}(B) \rightarrow A' \in \mathcal{N}(B)$
3. $A, A' \in \mathcal{N}(B) \rightarrow A \cap A' \neq \emptyset$.

We do not think that any weaker system could still capture the notion of "important" subsets. In particular, 3. may be strengthened to the filter property $A, A' \in \mathcal{N}(B) \rightarrow A \cap A' \in \mathcal{N}(B)$, but this would e.g. preclude a straightforward probablistic reading as e.g $A \in \mathcal{N}(B)$ iff A is more than half of B. Note however, that our properties capture e.g. a "prototypical" reading: if $b \in B$ is such a prototypical element of B, then $\{A \subseteq B : b \in A\}$ is a system as required. As we have seen, property 3. is strong enough to derive that "normally $\phi(x)$" and "normally $\neg\phi(x)$" together are inconsistent.

Choosing the system very weak has several advantages: First, it is easier to extend a weak system by adding additional axioms on the proof theoretical, stronger requirements on the system of important subsets on the semantic side preserving soundness and completeness than to retract some of such axioms and properties. Second, we have thus given a kind of generic interpretation, which can easily be modified to suit special domain (i.e. semantic) purposes or personal tastes. Third, we are not forced to make many "philosophical" commitments, but instead provide a tool kit to incorporate the desired postulates. Such extensions will be discussed towards the end of Chapter 2.

In our approach, we have relegated the dynamic aspects of defaults to the third Chapter, where we use - essentially - a maximal coherent subset of the information chosen according to specificity, and to the implementation, discussed in [Lor90], where as many x as possible are put into $\phi(x)$, when the default "normally $\phi(x)$" demands it. But, and this is important, we do not *start* empty-handed: Our minimal requirement, that $\phi(x)$ holds on an important subset has to be satisfied by any model of the default "normally $\phi(x)$", and then we try to improve an already "good" situation.

Let me rephrase things. Reading a default "normally $\phi(x)$" as $\forall x \phi(x)$, and then taking elements away from $\{x : \phi(x)\}$ to make the theory consistent seems to do injustice to the universal quantifier: We write down something we do not really mean. Starting on the other hand with the empty set trying to put as many elements into $\{x : \phi(x)\}$ to translate "normally $\phi(x)$" seems to violate an essential, static aspect of "normally": We should not say "normally" when we are not even shure that there is some x with $\phi(x)$. Somewhere in between $\forall x \phi(x)$ and $\exists x \phi(x)$ has to be the (static) meaning of "normally". But our generalized quantifier ∇ does just that: $\forall x \phi(x) \rightarrow \nabla x \phi(x)$ and $\nabla x \phi(x) \rightarrow \exists x \phi(x)$ are axioms in our theory, or, semantically, an important subset is a non-empty subset (and a little more). We are thus between $\forall$ and $\exists$, and, if you wish, as close to $\exists$ as reasonably possible.

Let me further point out here that the interpretation of defaults as the classical universal quantifier precludes in general a Theory Revision approach: *Any* default theory thus read will usually be inconsistent, by the mere presence of exceptions. Our formalism allows us to detect finer default inconsistencies of the form "normally $\phi(x)$" and "normally $\neg\phi(x)$". $\forall x \phi(x)$ and $\neg\phi(a)$ already would be classically inconsistent, but, of course - as you shall see in the formal part - $\nabla x \phi(x)$ and $\phi(a)$ are very well compatible, $\nabla x \phi(x)$ and $\nabla x \neg\phi(x)$, however, are not. Consequently, a classical interpretation "sees" too many inconsistencies, precisely because it is too rigorous, "normally" is not necessarily "all", our weaker reading, "normally" as something like a qualitative "most", penetrates through the outer layer of classical or pseudo-conflicts (pseudo in our context) to those which may exist on the level of normality itself.

In summary, we might say that we have tried to cover the minimal properties of the static part of "normally", giving an interpretation in an extended First Order Logic model, where $\mathcal{M} \models$ "*normally* $\phi(x)$" iff there is a large subset A of $\mathcal{M}'s$ universe such that for all $a \in A$ $\mathcal{M} \models \phi(a)$. This is a "local" interpretation in one model, not a "global" interpretation as in preferential structures - a further expression of the static aspect we are trying to cover: In *this* extended First Order Logic model, the minimal requirements justifying to say "normally $\phi(x)$" hold.

Technical Introduction We now give a very brief overview of the technical development. A detailed introduction to the central technical part Section 2.1 is to be found below.

We augment proof theory and semantics of classical First Order Logic to deal with the (open normal, in the sense of [Rei80]) default "normally, $\phi(x)$ holds". We say that $\phi(x)$ holds normally (in a model) iff there is a "large" subset of the universe of that model in which $\phi(x)$ "really" holds. In other words, "normally" is read as a generalized quantifier. In that sense, our semantics is local, as we work in one model only, and give a direct interpretation of defaults, in contrast to the various types of preferential models (see e.g. [Bou90], [BS85], [Del87], [Del88], [KLM90], [Lif85], [Lif86], [LM92], [McC80], [McC86], [Sch92], [Sho87]) which are global, working with several classical models. We first describe such systems of large subsets (Definition 2.1), and consider as basic semantic structures pairs $\mathcal{M} =< M, \mathcal{N}(M) >$, where M is a model of classical First Order Logic, and $\mathcal{N}(M)$ is such a system of large subsets of (the universe of) M. Validity of $\nabla x \phi(x)$ (= $\phi(x)$ holds normally) in $\mathcal{M}$ is defined in Definition 2.3. Suitable axioms for the quantifier ∇ are given in Definition 2.4, soundness and completeness of the axioms with respect to the semantics are shown in Lemma 2.4 and Theorem 5 there. The system permits full nestedness and boolean combinations of defaults. In particular, a notion of consistency of defaults results. The rest of Section 2.1 is devoted to extensions of the basic idea: (Open normal) defaults with prerequisites are seen as restricted generalized quantifiers, defaults of various strengths of "normality" are introduced, soundness and completeness shown.

In conclusion, we examine various strengthenings of the axioms, which correspond to widespread use of defaults and other systems of defeasible reasoning discussed in the literature.

In Chapter 3, we use the notion of consistency of defaults given in 2.1 to choose a reasonable consistent subset from a possibly inconsistent set of default information. We exploit the reliability relation given by the (partial) order of specificity in the sorted language of the LILOG project. The basic idea is to consider minimal inconsistent subsets, and to eliminate at least one suitably chosen (according to specificity, and determined by the function f there) element from each. Several approaches are discussed, one of which has essentially been implemented.

2 DEFAULTS AS GENERALIZED QUANTIFIERS

Overview of this Chapter In this chapter, we first take up the informal discussion of the introduction in somewhat more formal terms. Moreover, we also present several problems with the common use of defaults. We then proceed to the formal part of rigorous definitions, theorems and proofs.

The common use of defaults seems to presuppose strong assumptions on the structure of the universe - and seems thus to have common aspects with learning, and even philosophy of science. E.g. applying *both* defaults "normally $\phi(x)$" and "normally $\psi(x)$" seems to presuppose that not only "normally $\phi(x)$" and "normally $\psi(x)$", but also "normally $\phi(x) \wedge \psi(x)$". Even if our base system is too weak to permit such reasoning, we can easily adapt it by introducing a corresponding axiom schema, and strengthening the system of "important" subsets to a filter. Problems of this kind are discussed in the introductory part of this Chapter.

Section 2.1 contains the central definitions and results of this paper. We introduce $\mathcal{N}-systems$, which formalize the notion of a "large" or "important" subset, introduce the generalized quantifier ∇ - read e.g. "for almost all". $\mathcal{N}-systems$ then provide us with a semantics: A $\nabla-structure$ (or -model) is a classical first order structure with an $\mathcal{N}-system$ over its universe, and a formula $\nabla x\phi(x)$ is defined to hold in an $\nabla-structure$ iff there is a large subset A of the universe, i.e. some A in the $\mathcal{N}-system$ over the universe, such that for all elements $a \in A$ $\phi(a)$ holds. An axiomatisation is given, soundness and completeness for our semantics is shown. (For simplicity, we first treat the case of normal open defaults without prerequisites - corresponding to $\nabla x \phi(x)$ in our notation - the extensions to those with prerequisites is straightforward: $\nabla x \phi(x) : \psi(x)$ is interpreted as a generalized quantifier restricted to $\{x : \phi(x)\}$.

The weakness - choosing $\mathcal{N}-systems$ weaker than filters, and making no but trivial connections between the $\mathcal{N}-systems$ over different subsets of the universe for the relativized $\nabla - quantifier$ - is deliberate: We make as little "philosophical" commitments as possible, and permit easy and straightforward strengthenings into many different directions. In particular, the axiom systems P and R of [KLM90] and [LM92] can be incorporated easily into our system. Several such extensions are discussed in Section 2.2.

Problems in the use of defaults Before we proceed to a formal discussion of reasoning about normality, we further elaborate on the general picture and problems in defeasible reasoning.

We work on a background of naive "Normal Case Semantics". Thus, not all default theories will be meaningful to us, e.g. $\{\frac{:\phi(x)}{\phi(x)}, \frac{:\neg\phi(x)}{\neg\phi(x)}\}$ will make no sense. Normally $\phi(x)$, and normally $\neg\phi(x)$, at the same time can't be, as e.g. probabilistic reasoning with normally corresponding to a probability $> 1/2$ shows. We shall consider here only defaults of the form $\frac{\psi(x):\phi(x)}{\phi(x)}$ or $\frac{:\phi(x)}{\phi(x)}$, with at least one free variable. The latter will be abbreviated by $\nabla x \phi(x)$, and the former by $\nabla x \psi(x) : \phi(x)$. Defaults of the form $\frac{:\phi}{\phi}$, where ϕ has no free variables, will not be treated. Intuitively, we would read them as implicitely quantifying over all models under consideration.

We have treated $\nabla x \phi(x)$ as some kind of quantifier which lies between $\exists$ and $\forall$. It seems - at least to us - unreasonable to say that normally $\phi(x)$ holds, when there is no x such that $\phi(x)$. On the other hand, if $\forall x \phi(x)$ holds, it seems reasonable to say that ϕ normally holds. So, we have $\forall x \phi(x) \rightarrow \nabla x \phi(x) \rightarrow \exists x \phi(x)$. As said above, it seems unreasonable to say that normally $\phi(x)$ holds, and at the same time, that normally $\neg\phi(x)$ holds. We thus have the axiom $\nabla x \phi(x) \rightarrow \neg \nabla x \neg \phi(x)$, i.e. $\{\nabla x \phi(x), \nabla x \neg \phi(x)\}$ will be inconsistent.

So far, we have only treated normal defaults without prerequisites, i.e. of the type $\frac{:\phi(x)}{\phi(x)}$. How do we interpret a normal default with prerequisite, i.e. of the type $\frac{\psi(x):\phi(x)}{\phi(x)}$? We shall read it as a bounded quantifier, write $\nabla x \in \{y : \psi(y)\} : \phi(x)$, abbreviated $\nabla x\psi(x) : \phi(x)$, and understand it as saying that all x which satisfy ψ, will normally satisfy ϕ too. Of course, this is very different from $\nabla x(\psi(x) \rightarrow \phi(x))$. The former will be true iff the set X of things which satisfy ψ and do not satisfy ϕ, is small *in the set of things which satisfy* ψ, whereas the latter will be true iff X is small *in the universe.* Again, we shall have rules like $\forall x(\psi(x) \rightarrow \phi(x)) \rightarrow \nabla x\psi(x) : \phi(x)$ $\rightarrow [\exists x\psi(x) \rightarrow \exists x(\psi(x) \wedge \phi(x))]$.

We shall not give other types of defaults an (intuitive) semantics here, but mention that a seminormal default might be seen as having a computational meaning: To assume the full strength of the supposedly true situation might be too costly in case of error and revision.

We turn to the discussion of some problems in *applying* defaults.

A Problem of Language and Homogenousness Consider the world of birds, and the default rule that birds normally fly, $\nabla x fly(x)$. Suppose we know that penguins, emus and ostriches don't fly, albatrosses do, and we have the axiom $\forall x[largebird(x) \leftrightarrow$ penguin(x) $\vee$ emu(x) $\vee$ ostrich(x) $\vee$ albatros(x)]. Suppose we are told of some particular bird Tweety largebird(Tweety). Will we conclude that Tweety flies? Formally - i.e. in the sense of Reiter's Default Theory, see [Rei80] - yes, as it is consistent, since Tweety may be an albatros. Intuitively, no, because the non-flyers are a majority among the largebirds. So, normally, birds fly, but it is not the case that large birds normally fly, maybe even large birds normally don't fly. Contrast this with the colour of birds: We (as ornithological laymen) tend to believe that the colour of birds' feathers has nothing to do with their flying capabilities. So, if we know that Tweety is white-feathered, we will still intuitively believe that it flies. Suppose I tell you now that Tweety is white-feathered, lays spreckeled eggs of light blue and white colour, lives in Africa, has a long red beak, black feet, and a blue tail. What do you suspect, will Tweety fly? You might feel uncertain now, as the description is so narrow as to fit maybe only one species - which does not fly. The story looks like a trap (if you don't believe so, would you take a bet?).

Another, more formal, and more drastic example. Suppose we know fly(a), and someone tells us Tweety=a $\vee \neg fly(Tweety)$. It is consistent to assume Tweety=a, thus that it flies, so we assume it and conclude that

Tweety flies (and maybe that Tweety=a!).

Analysis of the problem Formal reasoning with the default $\nabla x fly(x)$ presupposes that logically consistent definable subsets are "probabilistically consistent" too. I.e., if things normally fly in the universe of birds, so do the elements of all subsets which are definable by some predicate, and which might consistently contain some flyer. If the language is empty, then there is no problem, but the richer the language (see the African bird above), the more there might be counterintuitive results, unless we somehow stop the applicability of the default $\nabla x fly(x)$. In other words, the conventional formal use of the default $\nabla x fly(x)$ presupposes that every definable subset $\{x : \phi(x)\}$ of the universe which is consistent with the flyer subset {x:fly(x)} behaves just as the universe itself. We may also say that ϕ is an irrelevant property for $\nabla x fly(x)$.

Reasoning with Defaults: Iterability Suppose we have the default theory $\nabla x\phi(x)$, $\nabla x\psi(x)$, and conclude $\phi(a) \wedge \psi(a)$ by default. Thus, we conclude not only that normally $\phi(x)$ and, normally, $\psi(x)$, but, in addition, normally $\phi(x)$ and $\psi(x)$. In other words, the set $\{x : \phi(x) \wedge \psi(x)\}$ has to be large in the universe, i.e. the intersection of $\{x : \phi(x)\}$ and $\{x : \psi(x)\}$ is large. We call this iterability. It is easy to think of many cases where each of the sets $\{x : \phi_i(x)\}$ is large in the universe, $\bigcap\{\{x : \phi_i(x)\} : i \in I\}$ is not empty (thus it will escape logical inconsistency), but small in the universe, thus iterability does not hold.

Our base axiom system permits e.g. neither $\nabla x\phi(x) \wedge \nabla x\psi(x) \rightarrow \nabla x\phi \wedge \psi(x)$ nor $\nabla x\phi(x) \rightarrow \nabla x\sigma(x) : \phi(x)$, but, of course, they can be added if so desired.

The discussion on iterability and homogenousness in the case of defaults with prerequisites has now to be carried out inside $\{x : \psi(x)\}$. The modifications are straightforward (for $\nabla x\psi(x) : \phi(x)$ and $\nabla x\psi'(x) : \phi'(x)$, we have to consider $\{x : \psi \wedge \psi'(x)\}$). We see again the advantage of a system which is poor by the number of its axioms and thus avoids excessive commitments, but whose base structure - i.e. whose language, in our case extended First Order Logic language - is expressive enough to permit multiple and strong extensions in a natural and straightforward way.

2.1 Semantics and Proof Theory

Overview of this Section The basic definition of the paper is that of an $\mathcal{N}-system$ (over a set M), it is given at the beginning of this section. We recollect that an $\mathcal{N}-system$ over M is supposed to capture the notion of a "large" or "important" subset, or one containing the "important" or "typical" elements of a base set M. Let me emphasize again that we have made the properties of an $\mathcal{N}-system$ deliberately extremely weak, weaker than a filter. Thus, our approach is generic on the sense that, if need be, we can strengthen the notion of an $\mathcal{N}-system$, add the corresponding axioms to the proof theory and still have a sound and complete system. Such modifications would be harder to achieve if, instead, we had to take away some undesirable strong properties. The axioms for an $\mathcal{N}-system$ of large subsets over M are:

1. M is large
2. supersets of large subsets are large
3. the intersection of two large subsets is non-empty.

The first property could be replaced by making the system itself non-empty.

The third property is the most interesting one. A filter would require the intersection of two large subsets to be large itself, we content ourselves with non-empty intersection. This permits e.g. a simple probabilistic interpretation by "more than half". It is property 3. which permits - in the intended interpretation - to deduce that "normally $\phi(x)$" and "normally $\neg\phi(x)$" together can't be: there is no element which satisfies both $\phi(x)$ and $\neg\phi(x)$.

The further development proceeds in two stages. We first treat normal open defaults without prerequisites, and then extend the discussion to normal open defaults with prerequisites (and $\mathcal{N}-families$, see below). All essential techniques and ideas are present in the case without prerequisite, and the reader so inclined may just leaf through the later subsections.

We introduce the new (unbounded) quantifier ∇ into the language (and its dual $\Diamond$ for proof theoretical purposes only). The crucial definition linking language and semantics is Definition 3, where we define an $\mathcal{N}-model$ and validity of a $\nabla-formula$ in an additional inductive step: An $\mathcal{N}-model$ is a classical first order structure, with an $\mathcal{N}-system$ of large subsets over its universe M. $\nabla x\phi(x)$ is defined to hold in the $\mathcal{N}-model$, iff there is a large subset $A \subseteq M$ such that for all $a \in A$ $\phi(a)$ "really" holds. As ϕ was not necessarily a classical formula, we can treat nested $\nabla' s$. Moreover, the new quantifier is fully embedded into the classical setting, so we can form boolean combinations, quantify classically over defaults in the sense of e.g. $\exists x \nabla y \phi(x, y)$ etc., and give these formulas a precise meaning, preserving the constructive spirit of First Order Logic. This seems to me one advantage over "global" Kripke style semantics where we have to "look elsewhere" for the interpretation of normality. Here, we can say "look, you see its holds *in this universe* on a large subset, so it normally holds here".

A corresponding axiomatisation follows in Definition 4: The first axiom says that implication preserves normality: If $\phi(x)$ normally holds, and $\phi(x)$ implies $\psi(x)$, then $\psi(x)$ normally holds. This corresponds to the second property of $\mathcal{N} - systems$. The second axiom says that normally ϕ and normally $\neg\phi$ can't be (we have discussed this above), and the third that the $\nabla - quantifier$ lies between the two classical ones, corresponding to the fact that M is a large subset of itself. (The second half of the axiom could be derived by $\neg\exists x \phi(x) \rightarrow \forall x \neg\phi(x) \rightarrow \nabla x \neg\phi(x) \rightarrow \neg\nabla x \phi(x)$ - we prefer to state it explicitely.) The last two axioms are auxiliary.

Lemma 2 states some basic and trivial consequences of the axioms. We then give a normal form for $\nabla - formulas$ to facilitate the completeness proof.

The central result of the paper is given in Lemma 4: A consistent $\nabla - theory$ has a model. The idea is to consider first order consequences of pairs of $\nabla - formulas$: Assume T to be deductively closed under our axiomatisation. Let $\nabla x \psi(x)$, $\nabla y \psi'(y) \in$ T with, for simplicity, ψ, ψ' classical formulas. By Lemma 2, a) $\nabla x \psi(x) \wedge \nabla y \psi'(y) \rightarrow \exists x(\psi \wedge \psi')(x)$. We now take a classical structure M satisfying all those first order consequences, say $M \models \psi \wedge \psi'(c_{\psi \wedge \psi'})$. For fixed $\nabla x \psi(x) \in T$, let $X_{\nabla x \psi(x)} := \{c_{\psi \wedge \psi'} : \nabla x \psi'(x) \in T\}$. $X_{\nabla x \psi(x)}$ will be one of the large subsets of the system to be constructed. It remains to show that the intersections of those sets are non-empty, and that the defined structure really is a model of T. But this is not difficult.

The soundness and completeness theorem is a direct consequence of this Lemma.

The extension to normal open defaults with prerequisites is straightforward: $\nabla x \phi(x) : \psi(x)$ ("if $\phi(x)$, then normally $\psi(x)$") is interpreted as a generalized quantifier relativized to $\{x : \phi(x)\}$, i.e. we consider an $\mathcal{N} - system$ not over the whole universe, but only over the subset where $\phi(x)$ holds.

Again, we keep our system deliberately very weak, in the sense that we do not demand any connections between the $\mathcal{N}-systems$ over the different $\{x : \phi(x)\}$ and $\{x : \phi'(x)\}$ - besides the trivial ones when e.g. $\{x : \phi(x)\} = \{x : \phi'(x)\}$. We even do not postulate $A \in \mathcal{N}(B) \wedge A \subseteq B' \subseteq B \rightarrow A \in \mathcal{N}(B')$, which a purely quantitative reading would justify: A large subset of B is a fortiori large in B', when $A \subseteq B' \subseteq B$. Again, we want to leave open all possibly intended developments and strengthenings.

The next extension concerns different degrees of normality. This is again straightforward and the reader is referred directly to this subsection.

Semantics

Definition 2.1 *Call $\mathcal{N}(M) \subseteq \mathcal{P}(\mathcal{M})$ (= the powerset of M) a $\mathcal{N}-system$ over M iff*

1. $M \in \mathcal{N}(M)$
2. $A \in \mathcal{N}(M)$, $A \subseteq B \subseteq M \rightarrow B \in \mathcal{N}(M)$
3. $A, B \in \mathcal{N}(M) \rightarrow A \cap B \neq \emptyset$ *if* $M \neq \emptyset$ *(thus,* $\emptyset \notin \mathcal{N}(M)$, *if* $M \neq \emptyset$*)*
   (Note that this is weaker than the corresponding axiom for filters.)

To facilitate proofs and enable normal forms, we introduce a complementary quantifier, $\Diamond$, too, with the meaning $\Diamond x\phi(x) :\leftrightarrow \neg\nabla x\neg\phi(x)$. The intuitive reading of $\Diamond x\phi(x)$ is thus roughly: "for at least a few x, $\phi(x)$ holds".

Remark: Our semantics covers the two extremes: - fix one element a of the universe U, then $\{A \subseteq U\colon\ a \in A\}$ will be a $\mathcal{N}-system$ - let some probability measure be given on U, then $\{A \subseteq U\colon\ p(A) > 0.5\}$ will be a $\mathcal{N}-system$. (Note, however, that the former can also be expressed by a suitable point measure on U.) We can thus cover both the "prototypical" and the "average" case.

Definition 2.2 *We augment the language of first order logic by the new quantifiers: If ϕ and ψ are formulas, then so are $\nabla x\phi(x)$, $\Diamond x\phi(x)$, $\nabla x\phi(x) : \psi(x)$, $\Diamond x\phi(x) : \psi(x)$ for any variable x. We call any formula of $\mathcal{L}$, possibly containing ∇ or $\Diamond$ a $\nabla-\mathcal{L}-formula$.*

Definition 2.3 *Let $\mathcal{L}$ be a first order language, and M be a $\mathcal{L}-structure$. Let $\mathcal{N}(M)$ be a $\mathcal{N}-$ system over M. Define $< M, \mathcal{N}(M) > \models \phi$ for any $\nabla - \mathcal{L} - formula$ inductively as usual, with two additional induction steps:*

- $< M, \mathcal{N}(M) > \models \nabla x\phi(x)$ *iff there is* $A \in \mathcal{N}(M)$ *such that* $\forall a \in A$ $(< M, \mathcal{N}(M) > \models \phi[a])$
- $< M, \mathcal{N}(M) > \models \Diamond x\phi(x)$ *iff* $\{a \in M\colon < M, \mathcal{N}(M) > \models \neg\phi[a]\} \notin \mathcal{N}(M)$.

Lemma 2.1 $< M, \mathcal{N}(M) > \models \Diamond x\phi(x)$ *iff*
$\forall A \in \mathcal{N}(M)\exists a \in A(< M, \mathcal{N}(M) > \models \phi[a])$. $\square$

Proof Theory

Definition 2.4 *Let any axiomatization of predicate calculus be given. Augment this with the axiom schemata*

1. $\nabla x\phi(x) \wedge \forall x(\phi(x) \rightarrow \psi(x)) \rightarrow \nabla x\psi(x)$
2. $\nabla x\phi(x) \rightarrow \neg\nabla x\neg\phi(x)$
3. $\forall x\phi(x) \rightarrow \nabla x\phi(x) \rightarrow \exists x\phi(x)$
4. $\Diamond x\phi(x) :\leftrightarrow \neg\nabla x\neg\phi(x)$
5. $\nabla x\phi(x) \leftrightarrow \nabla y\phi(y)$ *if x does not occurr free in $\phi(y)$ and y does not occurr free in $\phi(x)$*

(for all ϕ, ψ). We also denote the corresponding notion of derivability by $\vdash_\nabla$.

Lemma 2.2 *The following formulae are derivable:*

1. $\nabla x\phi(x) \wedge \nabla x\psi(x) \rightarrow \exists x(\phi \wedge \psi)(x)$

2. $\nabla x\phi(x) \land \neg\nabla x\psi(x) \to \exists x(\phi \land \neg\psi)(x)$

3. $\neg\nabla x\neg\phi(x) \to \exists x\phi(x)$

4. $\Diamond x\phi(x) \to \exists x\phi(x)$

5. $\nabla x\phi(x) \land \Diamond x\psi(x) \to \exists x(\phi \land \psi)(x)$

6. $\forall x(\phi(x) \leftrightarrow \psi(x)) \to (\nabla x\phi(x) \leftrightarrow \nabla x\psi(x)) \land (\Diamond x\phi(x) \leftrightarrow \Diamond x\psi(x))$

7. $\forall x\phi(x) \to \Diamond x\phi(x)$

It is usually *not* derivable: $\Diamond x\phi(x) \land \Diamond x\psi(x) \to \exists x(\phi \land \psi)(x)$. (To see this, use Theorem 5 below and argue semantically.) □

Soundness and Completeness To prepare the proof of completeness, we introduce $\nabla - normal$ forms ($\nabla - NF$).

Definition 2.5 *ϕ is in $\nabla - normal$ form ($\nabla - NF$) iff*

1. *ϕ contains only $\neg$, $\land$, $\lor$ as propositional operators*

2. *only atomic First Order Logic formulas are in the scope of $\neg$.*

Lemma 2.3 *For every ϕ there is ϕ' in $\nabla - NF$ such that $\vdash_\nabla \phi \leftrightarrow \phi'$.*

Proof By induction on the depth of $\nabla + \Diamond$ - nesting.
Case 1: depth=0: This is a classical result, take e.g. disjunctive prenex normal form (PNF).
Case 2: Let the depth of ϕ be n+1, and the result be proven up to depth n. We take an $\vdash_\nabla -equivalent$ ϕ" in e.g. disjunctive PNF, treating the subformulas within the outmost ∇ and $\Diamond$ quantifiers like classical atomic formulas, so ϕ" is of the form $Q_1 \ldots Q_n[\phi_1 \lor \ldots \lor \phi_m]$, the Q_i classical quantifiers, the ϕ_i of the form $\phi_{i1} \land \ldots \land \phi_{ik_i}$, where the ϕ_{ij} are either classical (negated) atomic formulas, or of the form $\nabla x\psi(x,\overline{y})$, $\neg\nabla x\psi(x,\overline{y})$, $\Diamond x\psi(x,\overline{y})$, or $\neg\Diamond x\psi(x,\overline{y})$. The negation can be passed through by $\vdash_\nabla \neg\nabla x\psi(x,\overline{y}) \leftrightarrow \Diamond x\neg\psi(x,\overline{y})$ and

$\vdash_\nabla \neg\Diamond x\psi(x,\overline{y}) \leftrightarrow \nabla x\neg\psi(x,\overline{y})$. By induction hypothesis, the $(\neg)\psi(x,\overline{y})$ can be transformed into an $\vdash_\nabla -equivalent$ $\psi'(x,\overline{y})$ in $\nabla - NF$. Axioms 1 and 5 give the result. □

Lemma 2.4 *Let T be a* $\nabla - \mathcal{L} - theory$*. Then T is consistent under the axioms of Definition 4 iff T has a model as defined in Definition 3.*

Proof The consistency of T when it has a model is trivial.

Let T be a $\vdash_\nabla -consistent$ $\nabla - \mathcal{L} - theory$. We have to show that it has a model. Throughout the proof, let " $\vdash_\nabla -consistent$" be abbreviated by "consistent". We give a constructive proof, to make the reader comfortable with the new logic. By the above, assume wlog. that all $\phi \in T$ are in $\nabla - NF$.

We first construct a consistent $T' \supseteq T$.

We add $c_\alpha : \alpha < \kappa$ new constants to $\mathcal{L}$, where κ is the size of $\mathcal{L}$, and inductively construct $T' = \bigcup\{T_\gamma : \gamma < \beta\}$ (T_γ ascending, β large enough) with $T_0 := T$, by adding new formulas to T, preserving consistency. (For simplicity, we omit the exact enumeration process - it does not matter anyway.) Let $\phi \in T_\gamma$, depending on the topmost operator, we add 0, 1, or several new formulas. It should be noted that all added formulas are in $\nabla - NF$ too.

Case 1: $\phi = \neg\psi$: We do nothing, by $\nabla - NF$, ψ is a classical atomic formula
Case 2: $\phi = \psi \wedge \psi'$: We add ψ, ψ', obviously preserving consistency.
Case 3: $\phi = \psi \vee \psi'$: Both $T_\gamma + \psi$ and $T_\gamma + \psi'$ can't be inconsistent, as $\phi \in T_\gamma$, so add one (or both) which preserves consistency.
Case 4: $\phi = \forall x\psi(x)$: Add all $\psi(c_\alpha)$, $\alpha < \kappa$
Case 5: $\phi = \exists x\psi(x)$: Add some $\psi(c_\alpha)$ which preserves consistency
Case 6: $\phi = \nabla x\psi(x)$: Add $\exists x\psi(x)$, and for each $\nabla y\psi'(y) \in T_\gamma$ $\exists x(\psi \wedge \psi')(x)$ and for each $\Diamond y\psi'(y) \in T_\gamma$ $\exists x(\psi \wedge \psi')(x)$. (after suitable renaming, preserving consistency by Lemma 2)
Case 7: $\phi = \Diamond x\psi(x)$: Add $\exists x\psi(x)$, and for each $\nabla y\psi'(y) \in T_\gamma$ $\exists x(\psi \wedge \psi')(x)$ (after suitable renaming, preserving consistency by Lemma 2)

In case 6 and 7, we mark all new $\exists x\psi(x)$ / $\exists x(\psi \wedge \psi')(x)$ as children of $\phi = \nabla x\psi(x)$ / $\phi = \nabla x\psi(x)$ and $\phi = \nabla x\psi'(x)$ etc.

Let $T' := \bigcup\{T_\gamma : \gamma < \beta\}$ β large enough, and $T" \subseteq T'$ be the set of First Order Logic-formulas of T'. By First Order Logic completeness, T' has a model M with universe U, where each $u \in U$ is denoted by some c_α.

Next, we define $\mathcal{N}(U)$.

Case 1: T' contains no $\nabla x\psi(x)$: Set $\mathcal{N}(U) := \{U\}$.

Case 2: Otherwise. Let $\nabla x\psi(x)$ be in T', and its children be $\exists x\psi(x)$, $\exists x(\psi \wedge \psi_i)(x)$, $i \in I$ (with $\nabla y\psi_i(y)$ / $\Diamond y\psi_i(y) \in$ T'), so there are $\psi(c_\alpha)$, $(\psi \wedge \psi_i)(c_{\alpha_i})$ $\in$ T'. Let $X_{\nabla x\psi(x)} := \{c_\alpha\} \cup \{c_{\alpha_i} : i \in I\}$ (we identify the c_α with their interpretation), and set $\mathcal{N}(U) := \{V \subseteq U$: $X_{\nabla x\psi(x)} \subseteq V$ for some $\nabla x\psi(x) \in T'\}$ Obviously, for $\nabla x\psi(x)$, $\nabla x\psi'(x) \in$ T', $X_{\nabla x\psi(x)} \cap X_{\nabla x\psi'(x)} \neq \emptyset$, as they have the common child $\exists x(\psi \wedge \psi')(x)$, so $\mathcal{N}(U)$ is a $\mathcal{N} - system$.

It remains to show that T holds in $\mathcal{M} :=< M, \mathcal{N}(U) >$. We show by induction on the complexity of ϕ that all $\phi \in T'$ hold in $\mathcal{M}$.

The atomic case is trivial, so are the cases $\neg$, $\wedge$, $\vee$, $\forall$, $\exists$. Consider now $\nabla x\psi(x)$. Note that for each $c_\alpha \in X_{\nabla x\psi(x)}$, $\mathcal{M} \models \psi(c_\alpha)$ by induction hypothesis. But $X_{\nabla x\psi(x)} \in \mathcal{N}(U)$, so $\mathcal{M} \models \nabla x\psi(x)$. Finally, consider $\Diamond x\psi(x)$.

Case 1: $\mathcal{N}(U) = \{U\}$. $\Diamond x\psi(x)$ has the child $\exists x\psi(x)$, so $\mathcal{M} \models \psi(c_\alpha)$ for some c_α, so $\mathcal{M} \models \Diamond x\psi(x)$ by Lemma 1.

Case 2: $\mathcal{N}(U) \neq \{U\}$. Let $V \in \mathcal{N}(U)$, so there is some $X_{\nabla x\psi'(x)} \subseteq V$, $\nabla x\psi'(x) \in T'$. $\Diamond x\psi(x)$ and $\nabla x\psi'(x)$ have the common child $\exists x(\psi \wedge \psi')(x)$, so there is some $c_\alpha \in X_{\nabla x\psi'(x)}$ with $\mathcal{M} \models (\psi \wedge \psi')(c_\alpha)$ by induction hypothesis. As this holds for all such V, $\mathcal{M} \models \Diamond x\psi(x)$ by Lemma 1 again. □

Theorem 2.5 *The axioms given in Definition 4 are sound and complete for the semantics of Definition 3, "they capture the $\mathcal{N} - semantics$ of ∇".*

Proof Let $T \not\models \phi$. Then there is a model M, such that $M \models T \wedge \neg\phi$. Thus, $Con(T \wedge \neg\phi)$, so $T \not\vdash \phi$. The other direction is analogous. □

Extension to Normal Defaults with Prerequisites

Definition 2.6 *Call $\mathcal{N}^+(M) =< \mathcal{N}(N) : N \subseteq M >$ a $\mathcal{N}^+ - system$ over M iff for each $N \subseteq M$ $\mathcal{N}(N)$ is a $\mathcal{N} - system$ over N. (It suffices to consider the definable subsets of M.)*

Definition 2.7 *Extend the logic of first order predicate calculus by adding the axiom schemata*

1. (a) $\nabla x\phi(x) \leftrightarrow \nabla x(x = x) : \phi(x)$

   (b) $\forall x(\sigma(x) \leftrightarrow \tau(x)) \wedge \nabla x\sigma(x) : \phi(x) \rightarrow \nabla x\tau(x) : \phi(x)$

2. $\nabla x\phi(x) : \psi(x) \wedge \forall x(\phi(x) \wedge \psi(x) \rightarrow \vartheta(x)) \rightarrow \nabla x\phi(x) : \vartheta(x)$

3. $\exists x\phi(x) \wedge \nabla x\phi(x) : \psi(x) \rightarrow \neg\nabla x\phi(x) : \neg\psi(x)$

4. $\forall x(\phi(x) \rightarrow \psi(x)) \rightarrow \nabla x\phi(x) : \psi(x) \rightarrow [\exists x\phi(x) \rightarrow \exists x(\phi(x) \wedge \psi(x))]$

5. $\Diamond x\phi(x) : \psi(x) \leftrightarrow \neg\nabla x\phi(x) : \neg\psi(x)$

6. $\nabla x\phi(x) : \psi(x) \leftrightarrow \nabla y\phi(y) : \psi(y)$ *(under the usual caveat for substitution.)*

(for all ϕ, ψ, ϑ, σ, τ*).*

Lemma 2.6 *The following are derivable:*

1. *the axioms of Definition 4, and the formulae of Lemma 2 (via the above 2.7.1.a and the corresponding relativized versions).*

2. *the relativized versions of Lemma 2, where the existential statements have to be weakened by an existential assumption as in Axiom 2.7.4.*

□

Definition 2.8 *Let* $\mathcal{L}$ *be a first order language, and* M *a* $\mathcal{L} -$ *structure. Let* $\mathcal{N}^+(M)$ *be a* $\mathcal{N}^+ -$ *system over M. Define* $< M, \mathcal{N}^+(M) > \models \phi$ *for any formula inductively as usual, with the additional induction steps:*

1. $< M, \mathcal{N}^+(M) > \models \nabla x\phi(x)$ *iff there is* $A \in \mathcal{N}(M)$ *such that* $\forall a \in A$ $(< M, \mathcal{N}^+(M) > \models \phi[a])$.

2. $< M, \mathcal{N}^+(M) > \models \Diamond x\phi(x)$ *iff* $\{a \in M : < M, \mathcal{N}^+(M) > \models \neg\phi[a]\} \notin \mathcal{N}(M)$.

3. $< M, \mathcal{N}^+(M) > \models \nabla x \phi(x) : \psi(x)$ *iff there is* $A \in \mathcal{N}(\{x :< M, \mathcal{N}^+(M) > \models \phi(x)\})$ *such that* $\forall a \in A$ $(< M, \mathcal{N}^+(M) > \models \psi[a])$

4. $< M, \mathcal{N}^+(M) > \models \Diamond x \phi(x) : \psi(x)$ *iff* $\{a \in M$: $< M, \mathcal{N}^+(M) > \models \phi[a] \wedge \neg \psi[a]\} \notin \mathcal{N}(\{x :< M, \mathcal{N}^+(M) > \models \phi(x)\})$.

Theorem 2.7 *The axioms of Definition 7 capture* $\mathcal{N}^+ - semantics$ *of* ∇.

Proof Fix any $\phi(x)$. The proof of Lemma 4 shows how to construct a model for all $\nabla x \phi(x) : \psi_i(x)$. Axiom 7.1 shows that equivalent ϕ will give the same construction of normal subsets of {x: $M \models \phi(x)$}. The additional assumption $\exists x \phi(x)$ in 7.4 was not needed in Definition 4, because the domain of a classical model is always non-empty. □

Extension to $\mathcal{N} - families$ So far, we can describe normal cases and the classical situation. We would like to generalize now to $\mathcal{N} - families$, sequences of $\mathcal{N} - systems$ of increasing strength, i.e. "approaching" hard=classical information. For a motivation, the reader is referred to a semantics for defeasible inheritance, discussed in [Sch90b].

Definition 2.9 1. *Let* γ *be any ordinal. Call* $< \mathcal{N}_i(c) : i < \gamma >$, $\mathcal{N}_i(c) \subseteq \mathcal{P}(c)$ *a* $\mathcal{N} - family$ *over c iff*

   (a) $c \in \mathcal{N}_i(c)$ *for all* i

   (b) $a \in \mathcal{N}_i(c)$, $a \subseteq b \subseteq c \rightarrow b \in \mathcal{N}_i(c)$ *for all* $i < \gamma$

   (c) $< \mathcal{N}_i(c) : i < \gamma >$ *is decreasing, i.e.* $a \in \mathcal{N}_i(c), j < i \rightarrow a \in \mathcal{N}_j(c)$

   (d) $a \in \mathcal{N}_i(c)$, $b \in \mathcal{N}_j(c) \rightarrow a \cap b \neq \emptyset$ *for all i,j, if* $c \neq \emptyset$.

2. *Call* $\mathcal{N}^+(M) =< \mathcal{N}(N) : N \subseteq M >$ *a* $\mathcal{N}^+ - family$ *over M iff for each* $N \subseteq M$ $\mathcal{N}^+(N)$ *is a* $\mathcal{N} - family$ *over N.*

Remark 2.8 *We have now formalized* γ *degrees of normality. Condition d. says, that all* $\mathcal{N}_i(c)$ $(\mathcal{N}_i(c))$ *are* $\mathcal{N}$ $(\mathcal{N}^+) - systems$ *over c.*

Definition 2.10 *We introduce $\gamma * 2$ quantifiers ∇^i, $\diamondsuit^i$, $i < \gamma$ into first order predicate calculus, and the following axioms: Let any axiomatization of predicate calculus be given. Augment this with the axiom schemata*

1. *For $\mathcal{N} - families$*
   (a) $\nabla^i x\phi(x) \wedge \forall x(\phi(x) \rightarrow \psi(x)) \rightarrow \nabla^i x\psi(x)$
   (b) $\nabla^i x\phi(x) \rightarrow \neg\nabla^i x\neg\phi(x)$
   (c) $\forall x\phi(x) \rightarrow \nabla^i x\phi(x) \rightarrow \exists x\phi(x)$
   (d) $\nabla^i x\phi(x) \rightarrow \nabla^j x\phi(x)$ *for* $j < i$
   (e) $\diamondsuit^i x\phi(x) \leftrightarrow \neg\nabla^i x\neg\phi(x)$ *(for all* ϕ, ψ, $i, j < \gamma$*).*
   (f) $\nabla^i x\phi(x) \leftrightarrow \nabla^i y\phi(y)$ *for safe substitution*
2. *For $\mathcal{N}^+ - families$:*
   (a) i. $\nabla^i x\phi(x) \leftrightarrow \nabla^i x(x = x) : \phi(x)$
   ii. $\forall x(\sigma(x) \leftrightarrow \tau(x)) \wedge \nabla^i x\sigma(x) : \phi(x) \rightarrow \nabla^i x\tau(x) : \phi(x)$
   (b) $\nabla^i x\phi(x) : \psi(x) \wedge \forall x(\phi(x) \wedge \psi(x) \rightarrow \vartheta(x)) \rightarrow \nabla^i x\phi(x) : \vartheta(x)$
   (c) $\exists x\phi(x) \wedge \nabla^i x\phi(x) : \psi(x) \rightarrow \neg\nabla^i x\phi(x) : \neg\psi(x)$
   (d) $\forall x(\phi(x) \rightarrow \psi(x)) \rightarrow \nabla^i x\phi(x) : \psi(x) \rightarrow [\exists x\phi(x) \rightarrow \exists x(\phi(x) \wedge \psi(x))]$
   (e) $\nabla^i x\phi(x) : \psi(x) \rightarrow \nabla^j x\phi(x) : \psi(x)$ *for* $j < i$
   (f) $\diamondsuit^i x\phi(x) : \psi(x) \leftrightarrow \neg\nabla^i x\phi(x) : \neg\psi(x)$
   (g) $\nabla^i x\phi(x) : \psi(x) \leftrightarrow \nabla^i y\phi(y) : \psi(y)$ *for safe substitution*

(for all ϕ, ψ, ϑ, σ, τ, $i, j < \gamma$*).*

Definition 2.11 *Let $\mathcal{L}$ be a first order language, and M a $\mathcal{L}-structure$. Let $\mathcal{N}(M)$ $(\mathcal{N}^+(M))$ be a $\mathcal{N}$ $(\mathcal{N}^+(M))-family$ over M. The definitions of $< M, \mathcal{N}(M) > \models \phi$ etc. are straightforward, like $< M, \mathcal{N}(M) > \models \nabla^i x\phi(x)$ iff there is $A \in \mathcal{N}_i(M)$ such that $\forall a \in A$ $(< M, \mathcal{N}(M) > \models \phi[a])$ etc.*

Theorem 2.9

1. *The axioms of Definition 2.10,1) capture $\mathcal{N}-family-semantics$ of ∇^i.*
2. *The axioms of Definition 2.10,2) capture $\mathcal{N}^+-family-semantics$ of ∇^i,*

Proof Analogous to the proofs of Lemma 4 and Theorem 7. □

Remark 2.10 *To improve performance of a system, we might have several layers of information. The very general (and simple) information might be formulated with $\nabla^0-quantifiers$, the more specific one with ∇^1- quantifiers etc., and the most specific (and hard) information classically, and all levels kept apart in the axiomatisation. Thus, when pressed for a fast answer, the system might "jump to a conclusion" by using only the $\nabla^0-information$, without checking consistency (i.e. without using ∇^i- $(i > 0)$ and hard information).*

2.2 Strengthening the Axioms

Overview of this Section In this section, we discuss various strengthenings of the axioms. They consist in adding properties to the single $\mathcal{N}-systems$ (corresponding to part of "iterability" - see below), in particular making $\mathcal{N}-systems$ filters, or coherence properties between the $\mathcal{N}-systems$ over different subsets of the universe for the relativized case. They are motivated partly by the common *use* of defaults, partly by discussion elsewhere in the literature. In particular, we show how to interpret the axiom systems P and R of [KLM90] and [LM92] in our first order setting, several of those axioms trivially hold in our system, "And" corresponds again to the filter property, and others translate into more subtle interdependencies among the $\mathcal{N} - systems$ of different subsets.

The ease with which we can incorporate these strengthenings is, of course, due to the fact that our system was chosen so weak at the outset, while making the language very expressive.

The Details In the introduction to this chapter, we have discussed problems of homogenousness and iterability. What do they look like in Normal Case language? But, this is simple, iterability will be discussed presenty, and homogenousness with respect to $\sigma(x)$ amounts to $\nabla x\phi(x) \rightarrow \nabla x\sigma(x) : \phi(x)$. We can add these as single axioms, or as axiom schemata for all formulae σ, and all defaults $\nabla\phi$, $\nabla\psi$, when justified. So, we are allowed to apply the conjunction of two defaults only if we have deduced $\nabla x(\phi \wedge \psi)(x)$. And, if we know something which is more than a mere tautology about a, let's say $T \vdash \sigma(a)$, then we are allowed to apply the default $\nabla\phi$ only if we can deduce $T \vdash \nabla x\sigma(x) : \phi(x)$.

We can now give a formal definition of iterability within our extended language. If a default theory is closed under iterability, and the defaults are of a simple form, then the construction of a model is especially easy, this is the content of the subsequent lemma.

Definition 2.12 *We call the following axiom schema Iterability for defaults without prerequisites:* $\nabla x\phi(x) \wedge \nabla x\psi(x) \rightarrow \nabla x(\phi \wedge \psi)(x)$

Lemma 2.11 *Let D be a finite set of defaults of the form* $\nabla x\phi_i(x)$, $i < n$, *where* ϕ_i *is a classical formula, and let T' be a set of classical formulae, and suppose that D is closed under iterability. Then* $D \cup T'$ *is consistent iff* $T' \cup \{\exists x(\phi_1 \wedge \ldots \wedge \phi_{n-1})(x)\}$ *is consistent.*

Proof Trivial. □

In the presence of defaults with prerequisites, iterability takes a slightly extended form. In addition, "normal use" of defaults sanctions still another axiom schema, which we call chaining. Again, we have a simplified model construction for simple defaults closed under iterability.

Definition 2.13 *We call the following axiom schemata Iterability for defaults with prerequisites:* $\nabla x\phi(x) \wedge \nabla x\sigma(x) : \psi(x) \to \nabla x\sigma(x) : (\phi \wedge \psi)(x)$ $\nabla x\sigma(x) : \psi(x) \wedge \nabla x\sigma'(x) : \psi'(x) \to \nabla x(\sigma \wedge \sigma')(x) : (\psi \wedge \psi')(x)$

Definition 2.14 *We call the following axiom schemata Chaining:* $\nabla x\phi(x) \wedge \nabla x\phi(x) : \psi(x) \to \nabla x\psi(x)$ $\nabla x\sigma(x) : \psi(x) \wedge \nabla x\psi(x) : \psi'(x) \to \nabla x\sigma(x) : \psi'(x)$

Lemma 2.12 *Let D be a finite set of defaults of the form* $\nabla x\phi_i(x)$, $i < n$, *let P be a finite set of defaults of the form* $\nabla x\sigma_i(x) : \psi_i(x)$, $i < m$, *where all* ϕ_i, σ_i, ψ_i *are classical formulae, and let T' be a set of classical formulae, and suppose that D and P are closed under iterability. Then* $D \cup P \cup T'$ *is consistent iff the following set of (classical) formulae is consistent: C:= T'* $\cup$ $\{\exists x(\bigwedge\{\phi_i : i < n\})(x)\} \cup \{ \exists x(\bigwedge\{\sigma_i : i \in p\})(x) \to \exists x(\bigwedge\{\sigma_i : i \in p\} \wedge \bigwedge\{\psi_i : i \in p\} \wedge \bigwedge\{\phi_i : i < n\})(x)$*:* $p \subseteq m$ $\}$

Proof " $\to$ " is trivial. " $\leftarrow$ ": Let M be a (first order) model of C. For $p \subseteq m$, let p_σ:= $\{ x \in M$: $M \models (\bigwedge\{\sigma_i : i \in p\})(x) \}$, p_ψ:= $\{ x \in M$: $M \models (\bigwedge\{\sigma_i : i \in p\} \wedge \bigwedge\{\psi_i : i \in p\} \wedge \bigwedge\{\phi_i : i < n\})(x) \}$ As M is a model of C, $p_\sigma \neq \emptyset \to p_\psi \neq \emptyset$. s_i:= $\{ x \in M$: $M \models \sigma_i(x) \}$ Let R:= $\{ p \subseteq m$: p is maximal such that $p_\sigma \neq \emptyset \}$ Note that $s_i = \emptyset$ iff there is no $p \in R$ such that $i \in p$. Define now $\mathcal{N}(M)$:= $\{ a \subseteq M$: {x: $M \models (\bigwedge\{\phi_i : i < n\})(x) \} \subseteq$ a },

$$\mathcal{N}(s_i) := \begin{cases} \{s_i\} & \text{iff } s_i = \emptyset \\ \{a \subseteq s_i : \bigcup\{p_\psi : i \in p, p \in R\} \subseteq a\} & \text{otherwise} \end{cases}$$

$\mathcal{N}(X)$:= {X} for all other subsets $X \subseteq M$. First, we show that this is indeed a $\mathcal{N}^+ - system$. As $M \models C$, this is trivial for $\mathcal{N}(M)$. Suppose now $s_i \neq \emptyset$. Thus, there is $p \in R$, $i \in p$. Consequently, $p_\psi \neq \emptyset$, $p_\psi \subseteq s_i$, and $\mathcal{N}(s_i)$ is a $\mathcal{N} - system$ for s_i. Next, we show that $< M, \mathcal{N}^+ > \models T' \cup D \cup P$. Let $\nabla x\phi_i(x) \in D$. As {x: $M \models (\bigwedge\{\phi_j : j < n\})(x)\} \subseteq$ {x: $M \models \phi_i(x)\}$, $< M, \mathcal{N}^+ > \models \nabla x\phi_i(x)$. Let $\nabla x\sigma_i(x) : \psi_i(x) \in P$. If $s_i = \emptyset$, $< M, \mathcal{N} > \models \nabla x\sigma_i(x) : \psi_i(x)$. Suppose now $s_i \neq \emptyset$, and let $i \in p$, $p \in R$. Then $p_\psi \subseteq \{x \in s_i$: $M \models \psi_i(x)\}$, thus $\{x \in s_i$: $M \models \psi_i(x)\} \in \mathcal{N}(s_i)$ and $< M, \mathcal{N}^+ > \models \nabla x\sigma_i(x) : \psi_i(x)$. $\Box$

Definition 2.15 *Call $\mathcal{N}^+(M)$ a smooth $\mathcal{N}^+ -$ system over M, iff it is a $\mathcal{N}^+ -$ system over M, and for each $N, N' \subseteq M$, $N \subseteq N'$, $A \in \mathcal{N}(N')$, $A \subseteq N$ implies $A \in \mathcal{N}(N)$. This will be captured by adding the axiom schema $\nabla x \phi(x) : \psi(x) \wedge \forall x(\sigma(x) \rightarrow \phi(x)) \wedge \forall x(\phi(x) \wedge \psi(x) \rightarrow \sigma(x)) \rightarrow \nabla x \sigma(x) : \psi(x)$.*

An Alternative Semantics for a Predicate Logic Version of P and R

We conclude by giving predicate logic versions of the systems P and R (see [KLM90] and [LM92]) an alternative semantics by translating $\phi(x) \mid\sim \psi(x)$ into $\nabla \phi(x) : \psi(x)$.

For the convenience of the reader, we repeat these axiom systems in their propositional form:

Definition 2.16 *(Strictly speaking - i.e. in object language -, the axioms are rules, and a system X of $\alpha \mid\sim \beta'$s is said to satisfy P (or R), iff it is closed under the corresponding rules.)*

1. *Right Weakening: $\models \alpha \rightarrow \beta$ and $\gamma \mid\sim \alpha$ entails $\gamma \mid\sim \beta$*
2. *Reflexivity: $\alpha \mid\sim \alpha$*
3. *And: $\alpha \mid\sim \beta$ and $\alpha \mid\sim \gamma$ entails $\alpha \mid\sim \beta \wedge \gamma$*
4. *Or: $\alpha \mid\sim \gamma$ and $\beta \mid\sim \gamma$ entails $\alpha \vee \beta \mid\sim \gamma$*
5. *Left Logical Equivalence: $\models \alpha \leftrightarrow \beta$ and $\beta \mid\sim \gamma$ entails $\alpha \mid\sim \gamma$*
6. *Cautious Monotony: $\alpha \mid\sim \beta$ and $\alpha \mid\sim \gamma$ entails $\alpha \wedge \beta \mid\sim \gamma$*
7. *Rational Monotony: $\alpha \mid\sim \beta$ entails $\alpha \mid\sim \neg\gamma$ or $\alpha \wedge \gamma \mid\sim \beta$*

The system P consists of 1-6, R of 1-7

We see that Right Weakening, Reflexivity, and Left Logical Equivalence are already built into our system.

The others have to be introduced by additional axiom schemata, like $\nabla x \alpha(x) : \beta(x) \rightarrow (\nabla x \alpha(x) : \neg\gamma(x) \vee \nabla x(\alpha \wedge \gamma)(x) : \beta(x))$ for Rational Monotony. The translation is obvious, so we discuss the semantical counterparts.

"And" says that $\mathcal{N} - systems$ are closed under finite intersections, i.e. $\omega - complete$ filters. (This is part of the above iterability.)

"Or", Cautious and Rational Monotony concern the relation of the $\mathcal{N} - systems$ of the subsets:

Or: $A \in \mathcal{N}(X)$, $B \in \mathcal{N}(Y) \rightarrow A \cup B \in \mathcal{N}(X \cup Y)$

Cautious Monotony: $A, B \in \mathcal{N}(X) \rightarrow A \cap B \in \mathcal{N}(A)$

Rational Monotony: $A \in \mathcal{N}(X)$, $B \subseteq X \rightarrow B \in \mathcal{N}(X)$ or $A - B \in \mathcal{N}(X - B)$

3 SCEPTICAL REVISION OF PARTIALLY ORDERED DEFAULTS

Overview of this Chapter Recollect from Chapter 1 that the LILOG setting consisted of an order sorted language, which determines the strength of default information through the specificity of the sort to which the default is applied. This partial order on the default information will now be exploited to select, in case of conflict, a reasonable consistent subset of the information. Recall also that "conflict" is made precise as inconsistency in our generalized quantifier logic introduced in Chapter 2.

We present and discuss four alternative approaches to Theory Revision, and decide for the third, which seems to present the intuitively best results. More precisely, this solution is a generic one, as it still leaves much liberty of choice for the "parameters" f and $\triangleleft$ (see below).

As usual in Theory Revison, we try to retain as much information as possible - under some restrictions. In particular, when two formulas of equal strength are in conflict, we will be fair and exclude both. On the other hand, if ϕ and ψ are in conflict, and so are ψ and σ, we will sometimes like to preserve ϕ and σ, eliminating just ψ, and thus preserve a maximum of information. This is basically the problem we investigate here: Given such ϕ, ψ, σ and the partial order on formulas, what are the intuitions that might guide us in the selection of a subset $\{\phi, \psi, \sigma\}$, and how can we formalize our choice? In our proposed solution, not only the relation between $\{\phi, \psi\}$ and $\{\psi, \sigma\}$ - e.g., is ϕ stronger than σ? - will enter the picture, but also the certainty (to be called definiteness of choice below) with which we would eliminate e.g. ψ from $\{\phi, \psi\}$.

3.1 Introduction

We will treat here a special case of theory revision based on axiom systems. Consider a situation of a (possibly inconsistent) partially ordered set of defaults. Our task will be to choose a consistent subset of those defaults, - where consistency is determined by the above discussed logic - using the given partial order, but being fair otherwise. Consequently, we will proceed sceptically, i.e. not necessarily choose a maximal consistent subset, as two contradictory defaults of the same or incomparable quality should both be excluded - as fairness dictates. The reader may as well assume that all formulae considered are classical ones, and the notion of consistency is the usual one. As a matter of fact, logic will play only a marginal role, being restricted to the notion of consistency.

First, we will introduce some basic definitions, and subsequently discuss several approaches to making a good and fair choice.

The reader will find many ideas from non-monotonic inheritance theories applied to and generalized in the following.

3.2 Basic Definitions and Approaches

Let Σ be a set of sorts S, partially ordered by $\leq$, and Δ a set of defaults, where each $\delta \in \Delta$ belongs to some sort s, written $s \models \delta$, thus Δ inherits the order $\leq$ from the sorts. To avoid inessential complications, we assume Σ and Δ to be finite. "$s \models \delta$" is supposed to read something like "normally, all elements of the sort s, have the property $\delta(x)$". (In addition, at sorts s we might have classical information, but we shall always assume that the total classical information is consistent, and also that the defeasible information written directly to some s is always consistent too (even when taking the classical information into account). This will simplify matters, but is not essential.) Thus, for some $x \in s$, $s < t < u$, $t \models \phi$, $u \models \psi$, it is natural to consider $\phi(x)$ to be the stronger information than $\psi(x)$, because it is the more specific information: In case of conflict, ϕ should win over ψ. (Thus, quality decreases with increase by $\leq$!) If, however, t and u are incomparable, fairness dictates that a conflict between ϕ and ψ should result in disbelief of both ϕ and ψ. We thus consider $f := \mathcal{P}(\cdot) - \{\emptyset\} \rightarrow \mathcal{P}(\cdot)$ with $f(\alpha) \subset \alpha$, where f is supposed to choose the "best" elements of α - if there are none, $f(\alpha)$ will be empty. Let further $\overline{f}(\alpha) := \alpha - f(\alpha)$.

Any such f gives rise naturally to a notion of quality of the choice: Suppose $card(\alpha) > 1$. If $card(\overline{f}(\alpha)) = 1$, then $f(\alpha)$ is a very definite choice, if $f(\alpha) = \emptyset$, i.e. $card(\overline{f}(\alpha)) = card(\alpha)$, $f(\alpha)$ is very indefinite. We can thus define the definiteness of the choice $f(\alpha)$ by

$$d(f(\alpha)) := \begin{cases} 1 & \text{iff } card(\alpha) = 1 \\ \frac{card(f(\alpha))}{card(\alpha)-1} & \text{otherwise} \end{cases}$$

(Thus, in the first case, $d(f(\alpha)) = 1$, in the second one $d(f(\alpha)) = 0$.) Speaking in terms of defeasible inheritance, $d(f(\alpha)) = 1$ corresponds to preclusion, $d(f(\alpha)) = 0$ to contradiction. We shall not use the notion of definiteness until the third approach, when we need a relation $\triangleleft$ on minimal inconsistent subsets.

Examples: Let $g(\alpha)$ be the greatest element of α - i.e. for all $x \in \alpha$, $x \neq g(\alpha)$, $x < g(\alpha)$ - if it exists.

$$f_1(\alpha) := \begin{cases} \alpha - \{g(\alpha)\} & \text{iff } g(\alpha) \text{ is defined} \\ \emptyset & \text{otherwise} \end{cases}$$

$$f_2(\alpha) := \begin{cases} \{x \in \alpha : x \text{ is no non-trivial} \\ \text{maximum, i.e.} \exists y \in \alpha (x < y) \\ \text{or } \neg \exists y \in \alpha (y < x)\} & \text{iff there are } x, y \in \alpha \\ & \text{such that } x < y \\ \emptyset & \text{otherwise} \end{cases}$$

$f_3(\alpha) := \{x \in \alpha$: x is a non-maximal element, i.e. $\exists y \in \alpha (x < y)\}$

By $f(\alpha) \subset \alpha$, $f(\alpha)$ will be consistent, if α is minimal inconsistent.

This suggests the probably simplest

Approach 1 Iterate some fixed f, starting on Δ, until consistency is reached. By finiteness of Δ and antitony of f, this will always work. This approach has some effects which may not always be desirable: Let $s < t < u$, $s \models \phi$, $t \models \neg\phi$, $u \models \psi$, where $\vdash \psi \rightarrow \phi$. ψ being the weakest information, it should always be eliminated first by f. On the other hand, one might argue that $\{\phi, \psi\}$ is a good choice, as ϕ should be accepted by being the best information, thus $\neg\phi$ should be out, leaving the way open to ψ.

Approach 2 We shall leave momentarily the above introduced function f, and even the order $\leq$ on sorts and defaults, and discuss a quite different way. We now consider arguments, which we identify with subsets of Δ, choosing the best ones, and hoping that the union of the best arguments is consistent. This meets with some problems. An order $\prec$ on arguments should respect the following properties: Let α, β etc be subsets of Δ. (Again, strength will decrease with increasing $\prec$.)
1. $\prec$ should be transitive, i.e. $\alpha \prec \beta \prec \gamma \rightarrow \alpha \prec \gamma$
2. $\alpha \prec \beta \subseteq \gamma \rightarrow \alpha \prec \gamma$. Reason: simply adding some information to an argument should not make it stronger. Adding the truth 2+2=4 to an argument should not give it more power. It is rather that the weakest part should determine its force.
3. $\forall i \in I(\alpha_i \prec \beta) \rightarrow \bigcup\{\alpha_i i \in I\} \prec \beta$: just putting arguments together should not violate a common upper bound. This seems to be well in accord with many natural definitions based purely on an order $\leq$ as given above - though not on all. (Consider e.g. $\alpha \prec \beta :\leftrightarrow \exists x \in \beta \forall y \in \alpha(y < x)$.)
4. The inverse seems to be very doubtful: $\beta \prec \bigcup\{\alpha_i : i \in I\} \rightarrow \exists i \in I(\beta \prec \alpha_i)$. This is already violated by $\alpha \prec \beta :\leftrightarrow \forall x \in \alpha \exists y \in \beta(x < y)$.

Yet rule 4 seems to suggest itself in our present approach, when trying to prove consistency: Let α, β, γ be pairwise consistent, but $\neg Con(\alpha \cup \beta \cup \gamma)$. If $\alpha \prec \beta$ and $\beta \prec \gamma$, then by 3., $\alpha \cup \gamma \prec \gamma$, and γ will be omitted in the inductive procedure, resulting in a consistent choice. If, however, $\alpha \perp \gamma$ (incom- parable), and $\beta \perp \gamma$, then maybe $\gamma \prec \alpha \cup \beta$ (unless 4. holds), so $\alpha \cup \beta$ will not be chosen, but maybe each of α, β, γ, resulting in an inconsistent theory.

Approach 3 We now work more closely again with the introduced functions f. First, a combinatorial result.

Definition 3.1 *Let D be a finite set, $M' \subseteq \mathcal{P}(\mathcal{D})$, $\lhd$ an acyclic binary relation on M', for $A \in M'$ let $\emptyset \neq X_A \subseteq A$ be defined. Define inductively for $i \in \omega$:*

M'_i:= $\{$ $A \in M'$: $A \lhd -$ *minimal in* $M' - \bigcup\{M'_j : j < i\}$,
M_i:= $\{$ $A \in M'_i$: $\forall j < i \forall B \in M_j (B \lhd A \rightarrow A \cap X_B = \emptyset)$ $\}$,
X_i:= $\bigcup\{X_A : A \in M_i\}$ *and*
$M := \bigcup\{M_i : i \in \omega\}$, $X := \bigcup\{X_i : i \in \omega\}$, *D':=D-X.*

Let, in the sequel, the situation of the definition be given.

Lemma 3.1 *For all* $A \in M'$, $A \cap X \neq \emptyset$.

Proof Let $A \in M'_i$, $i < \omega$. If $A \in M_i$, $X_A \subseteq X_i \subseteq X$, so $A \cap X \neq \emptyset$. If $A \in M'_i - M_i$, then there is $j < i$, $B \in M_j$, $B \lhd A$ such that $A \cap X_B \neq \emptyset$. By $X_B \subseteq X_j \subseteq X$, $A \cap X \neq \emptyset$. □

Remark 3.2 *1) It suffices to choose* $\lhd$ *well-founded in the above definition, D and M' can then be infinite, too. We do not need the more general result, however.*
2) Evidently, $\lhd$ *may be chosen as the empty relation.*
3) Let $A \in M'$. *Then* $A \in M \leftrightarrow \forall B \lhd A (B \in M \rightarrow A \cap X_B = \emptyset)$.
Proof by induction on i: Let $A \in M'_i$ *be minimal such that the result fails.* " $\rightarrow$ ": *If* $A \in M_i$ *and* $\exists B \lhd A (B \in M,\ A \cap X_B \neq \emptyset)$, *then* $B \in M_j$ *for some* $j < i$, *and* $A \notin M_i$ *by construction.* " $\leftarrow$ ": *Let* $A \in M'_i - M_i$ *such that* $\forall B \lhd A (B \in M \rightarrow A \cap X_B = \emptyset)$. *Thus, for all* $j < i$, $B \in M_j$, $B \lhd A$ $A \cap X_B = \emptyset$, *and* $A \in M_i$ *by construction again.* □

Corollary 3.3 *Let D be a finite set of formulae,* $M' \subseteq \mathcal{P}(D)$ *the set of minimal inconsistent sets of formulae from D. If X is defined as above, D' will be consistent.*

Proof *Suppose not, so D' contains some* $A \in M'$, *so* $A \cap X = \emptyset$, *contradiction.* □

Thus, letting f be as above, defining $X_A := A - f(A)$, any acyclic $\lhd$ on minimal inconsistent subsets of Δ will give a consistent subset $\Delta' \subseteq \Delta$. Choosing $\lhd$ non-empty will lead to some inconsistencies "being invisible", since some elements responsible are eliminated already. The effect can be seen easily when looking back at the example discussed in the first approach: Letting $\{\phi, \neg\phi\} \lhd \{\neg\phi, \psi\}$, $\neg\phi$ might be eliminated (by suitable f and X) as $\{\phi, \neg\phi\}$ is $\lhd -minimal$, leaving $\{\neg\phi, \psi\}$ out of consideration. The discussion of $\lhd$ will be resumed in a moment.

Next, we look at the compatibility of this approach with an order on arguments (subsets of Δ).

It is natural to define for α, $\beta \subseteq \Delta$: $\alpha \prec \beta :\leftrightarrow \alpha \subseteq f(\alpha \cup \beta)$ - leaving all logic aside for the moment. We have discussed the postulates 1.-3. for orders on arguments above, and examine now which of the above introduced f_i satisfy some or all of these postulates. f_1 will violate 2., as the addition of any element, such that the largest element does not exist any more, will show. Consider f_2, and let $\alpha := \{x\}$, $\beta := \{y, z\}$, ordered by $y < z$ (as defaults) only. Then $f(\alpha \cup \beta) = \{x, y\}$, thus $\alpha \prec \beta$ (as arguments), which seems very unnatural. Looking at $f = f_3$, which preserves all non-maximal elements, we see that the order on arguments generated by f satisfies indeed 1.-3. This is immediate by $\alpha \prec \beta \leftrightarrow \alpha \subseteq f_3(\alpha \cup \beta) \leftrightarrow \forall x \in \alpha \exists y \in \alpha \cup \beta (x < y) \leftrightarrow \forall x \in \alpha \exists y \in \beta (x < y)$, the last equivalence by transitivity of the partial order on Σ.

We have now for $\delta \in \Delta$ ($\lhd = \emptyset$ again, M' being the set of minimal inconsistent subsets of Δ, thus by $\lhd = \emptyset$ M=M'): $\delta \in \Delta' \leftrightarrow \forall A \in M'(\delta \notin X_A)$ $\leftrightarrow \forall A \in M'(\delta \in A \rightarrow \delta \in f(A)) \leftrightarrow \forall A \in M'(\delta \in A \rightarrow \{\delta\} \prec A - \{\delta\}) \leftrightarrow \forall \beta \subseteq \Delta(\neg Con(\delta, \beta) \rightarrow \{\delta\} \prec \beta)$. In the last equivalence, we use property 2.: " $\leftarrow$ " is trivial. " $\rightarrow$ ": Suppose there is β, $\neg Con(\delta, \beta)$, $\neg\{\delta\} \prec \beta$. Take $\beta' \subseteq \beta$ such that $\{\delta\} \cup \beta'$ is minimal inconsistent, so by prerequisite, $\{\delta\} \prec \beta'$, and by property 2. $\{\delta\} \prec \beta$.

Moreover, for $\alpha \subseteq \Delta'$, we have $\neg Con(\alpha \cup \beta) \rightarrow \alpha \prec \beta$: Let $\beta' \subseteq \beta$ such that $\alpha \cup \beta'$ is minimal inconsistent. As $\alpha \subseteq \Delta'$, $X_{\alpha \cup \beta'} \cap \alpha = \emptyset$, thus $\alpha \subseteq f(\alpha \cup \beta')$ and $\alpha \prec \beta'$, by 2. $\alpha \prec \beta$.

Let us resume the discussion of $\lhd$. What are the meaning and properties of the $\lhd -relation$? Let $A \lhd B$. Thus, the elimination of the inconsistency A by eliminating the elements of X_A is estimated so strong that the inconsistency B need not be considered any more. In other words, the argument $B - X_B$ against X_B looses its importance. Now, we have by f a relation $\prec$ on arguments. Moreover, we have a notion of definiteness, resulting from f as well, as introduced at the beginning of the Section. A natural definition (suggested by S.Lorenz) of $\lhd$ will be e.g.: $A \lhd B :\leftrightarrow$ 1. $\overline{f}(A) \subseteq B$, 2. $A - B \prec B$, 3. d(f(A))=1, $d(f(B)) \neq 1$. By 3., any such $\lhd$ will be acyclic, by the same reason, we can't have $A \lhd B \lhd C$. To conclude $A \lhd B \subseteq C \rightarrow A \lhd C$ for C such that $d(f(C)) \neq 1$, we need $A - C \prec C$, which will not always hold. (It will hold for $\prec$ defined by f_3.)

Approach 4 As a last alternative, we shall discuss a more "procedural approach", which tries to directly find a suitable subset $\Delta' \subseteq \Delta$. Again we suppose some subset function f to be given. Let $\phi_1 < \phi_2 < \phi_3$, $\phi_1 < \psi < \rho < \phi_4$ be defaults ordered by their sorts. The ψ, ρ are (logically) unimportant here and might be tautologies. Suppose further $\{\phi_1, \phi_2, \phi_4\}$ and $\{\phi_2, \phi_3\}$ are the minimal inconsistent subsets. Suppose further $f(\{\phi_1, \phi_2, \phi_4\}) = \{\phi_1\}$, $f(\{\phi_2, \phi_3\}) = \{\phi_2\}$, which seems a natural choice. Proceeding now by the rank of the ϕ_i in the $< -order$, at rank 2 ϕ_3 will be eliminated, as the elimination of ϕ_2 and ϕ_4 will be discovered only at rank 3. In other words, we have here a result similar to the one discussed in the first approach. (It corresponds again to choosing the empty order $\lhd$ on minimal inconsistent subsets.)

4 ACKNOWLEDGEMENTS

This work was carried out while the author was at IBM Germany, Stuttgart. I would like to thank here IBM Germany, which, in the LILOG project framework, has not only provided interesting problems, but also a generous setting which enabled fundamental research originating from questions of applications. Sven Lorenz has essentially implemented Approach 3 within the LILOG project, see [Lor90], and thus has taken much burden off the author's shoulders. Finally, I would like to thank two anonymous referees who incited me to put more prose between poetry, and thus make this text more readable - at least I hope so.

References

[AGM85] C.Alchourron, P.Gärdenfors, D.Makinson, "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions", Journal of Symbolic Logic 50 (1985), p.510-530

[BHR90] K.H.Bläsius, U.Hedtstück, C.-R.Rollinger (eds.), "Sorts and Types in Artificial Intelligence", Springer, Berlin, 1990

[Bou90] C.Boutilier, "Viewing Conditional Logics of Normality as Extensions of the Modal System S4", Toronto University, KRR-TR-90-4, June 1990

[BS85] G.Bossu, P.Siegel, "Saturation, Nonmonotonic Reasoning and the Closed-World Assumption", Artificial Intelligence 25 (1985) 13-63

[Del87] J.P.Delgrande, "A First-Order Conditional Logic for Prototypical Properties", Artificial Intelligence 33 (1987), p.105-130

[Del88] J.P.Delgrande, "An Approach to Default Reasoning Based on a First-Order Conditional Logic: Revised Report", Artificial Intelligence 36 (1988), p.63-90

[Gar88] P.Gärdenfors, "Knowledge in Flux", MIT Press, 1988

[GM88] P.Gärdenfors, D.Makinson, "Revisions of Knowledge Systems Using Epistemic Entrenchment". In: M.Y.Vardi (ed.), "Theoretical Aspects of Reasoning about Knowledge", Morgan Kaufmann, Los Altos, Cal., 1988, p.83-95

[KLM90] S.Kraus, D.Lehmann, M.Magidor, "Nonmonotonic Reasoning, Preferential Models and Cumulative Logics", Artificial Intelligence 44 (1990), p.167-207

[Lif85] V.Lifschitz, "Computing Circumscription", Proceedings IJCAI 1985, p.121-127

[Lif86] V.Lifschitz, "On the Satisfiability of Circumscription", Artificial Intelligence 28 (1986), p.17-27

[LM92] D.Lehmann, M.Magidor, "What does a conditional knowledge base entail?" Artificial Intelligence 55 (1992), p.1-60

[Lor89] S. Lorenz, "Skeptical Reasoning with Order-Sorted Defaults", Proceedings of the Workshop on Nonmonotonic Reasoning, Dec. 1989, Bonn, G.Brewka, H.Freitag eds., GMD Tech.Rep. 443, GMD, POB 1240, D-53757 St.Augustin, Germany

[Lor90] S. Lorenz, "Nichtmonotones Schliessen mit ordnungssortierten Defaults", IWBS-Report 100, IBM Germany, D-7000 Stuttgart 80, POB 800880, Germany, 1990

[Mak85] D.Makinson, "How to give it up: A Survey of Some Formal Aspects of the Logic of Theory Change"; Synthese 62 (1985), p.347-363

[McC80] J.McCarthy, "Circumscription - a form of non-monotonic reasoning", Artificial Intelligence 13 (1980), p.27-39

[McC86] J.McCarthy, "Applications of Circumscription to Formalizing Common-Sense Knowledge", Artificial Intelligence 28 (1986), p.89-116

[Rei80] R.Reiter, "A Logic for Default Reasoning", Artificial Intelligence 13 (1980), p.81-132

[Sch90b] K.Schlechta, "Semantics for Defeasible Inheritance", in: L.G.Aiello (ed.), "Proceedings ECAI 90", London, 1990, p.594-597

[Sch92] K.Schlechta, "Some Results on Classical Preferential Models", Journal of Logic and Computation, Vol.2, No.6, 1992

[Sho87] Y.Shoham, "A semantical approach to nonmonotonic logics". In Proceed. IJCAI 87, p.388-392