

Normality in Update Semantics

Reinhard Blutner
Max Planck Gesellschaft, Berlin

1. Introduction: The relationship between meaning and inference

If we know the meanings of the sentences of a (natural) language \mathcal{L} , do we know then (in principle) what the valid inferences in the language are? In asking this question we view humans not only as speakers but also as reasoners. The question was asked first in symbolic logic and analytic philosophy, but it is essential also for natural language semantics and computational linguistics. The relationship between meaning and inference is an important one not only because we must be able to reason with the semantic representations we obtain from utterances but also because inference is intimately bound up with the interpretation process itself.

Perhaps the question asked at the beginning must be answered differently depending on which part of the language vocabulary is concerned. With regard to the *logical* vocabulary of natural language the classical answer is summarized in the following thesis:

Logicity: The meanings of the logical units of a (natural) language \mathcal{L} determine what the valid logical inferences (entailments) in that language are.

Obviously, this thesis can be derived from two traditional assumptions: (A) identifying meanings with truth-conditions and (B) defining the notion of entailment as validity of the conclusion in every possible case where the premises hold. From knowing the meanings of the premises of an argument, we can determine (in principle) all the possible cases where the premises are true and then check whether in all these cases the conclusion is true as well. If this check is positive (negative) the argument is (in)valid.

The traditional view leads to a notion of entailment which obeys the principle of monotonicity: If a set of premises Γ entails a formula Φ and Γ' extends Γ , i.e. $\Gamma \subseteq \Gamma'$, then Γ' entails Φ . Now consider the case of *common sense entailment* where this principle clearly is wrong as the following example illustrates:

Ravens are (normally) black
Albin is a raven

(Presumably) Albin is black

Ravens are (normally) black
Albin is a raven

Albin is not black
(Presumably) Albin is black invalid

In order to describe/explain such inference patterns based on the meaning of the involved expressions at least one of the traditional assumptions (A) and (B) must be given up. One possibility is to retain (A) but to modify (B) by approaching the

notion of entailment using a restricted set of *preferred* models. This approach has been called preference semantics. Another approach is update semantics which already abandons the assumption of identifying meanings with truth-conditions. Instead, the update approach identifies meanings with condition of information change and it also uses a more dynamically coloured notion of entailment.

The logical particles of natural language that I will investigate in this study are certain adverbs of quantification as *normally*, *typically*, and *commonly*. Capturing the subtle differences in the meanings of these adverbs are not envisaged here. Instead, I will concentrate myself upon the supposed common semantic core of these adverbs of generic quantification.

The aims of the present paper are threefold: First, I want to demonstrate that the thesis of logicity can be satisfied in connection with the phenomenon of nonmonotonicity if the paradigm of dynamic semantics is used; it cannot be satisfied in case of preference semantics. Second, I will demonstrate that the approach of preference semantics leads to some further problems: It cannot describe the coexistence of highly plausible patterns of inference like GRADED NORMALITY and LOGICAL CLOSURE IN THE CONSEQUENT; it cannot give an explanatory adequate treatment of the pattern of inference called PENGUIN PRINCIPLE. Finally, I want to develop a version of update semantics (called *update semantics for prototypes*) that obeys the doctrine of logicity and deals with the mentioned problems in a fairly explanatory way.

In Section 2 and 3 I introduce the general ideas of preference semantics and dynamic semantics, respectively. Section 4 summarizes some of the crucial patterns of inference that are involved in reasoning with generic sentences such as ϕ s are (normally) ψ s. In Section 5 several models of nonmonotonic reasoning are compared and three challenging problems of these models are isolated. Section 6 develops the update semantics for prototypes and draws some conclusions.

2. Preference semantics

SHOHAM (1988) has provided a unifying semantical framework for the definition of nonmonotonic logics. Starting from any standard logic with a model-theoretic semantics, a *preference logic* is obtained by adding a strict partial preference order on interpretations to the logic. The preferred models are those maximum with respect to the given ordering. We can now define two notions of entailment.

Definition 1: Let \mathcal{L} be a language with a model-theoretic semantics and \sqsubset a strict partial order on the interpretations of \mathcal{L} (i.e. an irreflexive, asymmetric, and transitive relation). Furthermore, let Γ be a set of formulae, Φ a formula of \mathcal{L} .

- (a) Γ **classically entails** Φ (written $\Gamma \models \Phi$) iff the models of Γ are a subset of the models of Φ ;
- (b) Γ **preferentially entails** Φ (written $\Gamma \succ_{\sqsubset} \Phi$) iff the preferred models of Γ are a subset of the models of Φ .

The most obvious illustration for preference semantics is MCCARTHY's (1980, 1986) (variable) predicate circumscription. Here \mathcal{L} is a first order language with a distinguished predicate letter P whose extension is made minimal while the elements of a certain predicate set v (not containing P) are allowed to vary during the minimization. Depending on P and v the partial order on the interpretations of \mathcal{L} is defined as follows: $M \sqsubset_{P,v} M'$ iff (i) M and M' have the same domain of individuals, (ii) all predicate symbols of \mathcal{L} other than P and v have the same extensions in M and M' , (iii) the extension of P in M' is contained in its extension in M . It is evident that the preferred models of a set of formulae Γ are those where the extension of the predicate P is minimized.

To give an example let us translate generic sentences such as *ravens are (normally) black* as $\forall x(\text{RAVEN}(x) \wedge \neg \text{AB}(x) \rightarrow \text{BLACK}(x))$. Intuitively: Ravens who are not abnormal are black. AB then is the distinguished predicate letter whose extension should be minimized while the extension of BLACK is allowed to vary during the minimization. This makes as few ravens as possible abnormal (with regard to the feature BLACK). Now let us consider the following set of premises:

$$\Gamma_1 = \{\forall x(\text{RAVEN}(x) \wedge \neg \text{AB}(x) \rightarrow \text{BLACK}(x)), \text{RAVEN}(\text{ALBIN})\}$$

Using Definition 1 it is clear that Γ_1 preferentially entails $\text{BLACK}(\text{ALBIN})$ (of course, this conclusion is not classically entailed by Γ_1). Now let us extend Γ_1 as in the introductory example of Section 1:

$$\Gamma_2 = \{\forall x(\text{RAVEN}(x) \wedge \neg \text{AB}(x) \rightarrow \text{BLACK}(x)), \text{RAVEN}(\text{ALBIN}), \neg \text{BLACK}(\text{ALBIN})\}$$

Obviously, the extension Γ_2 does not preferentially entail $\text{BLACK}(\text{ALBIN})$. These observations demonstrate the nonmonotonicity of the notion of preferential entailment and they verify an important pattern of defeasible reasoning with generic sentences called DEFEASIBLE MODUS PONENS.

A raven may be abnormal with respect to the hue of his feathers but normal in some other aspect. Thus we need different abnormality predicates AB_i for different aspects (usually one for each generic sentence). However, the unrestricted minimization of these predicates leads to several problems. For instance, from *ravens are (normally) black*, represented

$$\forall x(\text{RAVEN}(x) \wedge \neg \text{AB}_1(x) \rightarrow \text{BLACK}(x))$$

and *ravens can (normally) fly*, represented

$$\forall x(\text{RAVEN}(x) \wedge \neg \text{AB}_2(x) \rightarrow \text{CAN-FLY}(x))$$

it does not follow with our preference semantics that ravens are (normally) black and can fly, represented

$$\forall x(\text{RAVEN}(x) \wedge \neg \text{AB}_2(x) \rightarrow \text{BLACK}(x) \wedge \text{CAN-FLY}(x)).$$

I find this prediction counterintuitive. As a second example consider the following set of premises:

ravens are (normally) black
albino ravens are (normally) not black
all albino ravens are ravens
Albin is an albino raven

which is represented as

$$\begin{aligned} \forall x(\text{RAVEN}(x) \wedge \neg \text{AB}_1(x) &\rightarrow \text{BLACK}(x)) \\ \forall x(\text{ALBINO-RAVEN}(x) \wedge \neg \text{AB}_2(x) &\rightarrow \neg \text{BLACK}(x)) \\ \forall x(\text{ALBINO-RAVEN}(x) &\rightarrow \text{RAVEN}(x)) \\ \text{ALBINO-RAVEN}(\text{ALBIN}) \end{aligned}$$

From these premises it does not follow with our preference semantics that *Albin is not black* (represented $\neg \text{BLACK}(\text{ALBIN})$). Thus, the present form of preferential entailment does not capture the pattern of defeasible reasoning which sometimes has been called *elsewhere condition*, *principle of specificity*, or *penguin principle*.

The crucial point with respect to these examples is that there must be additional constraints that restrict the extensions of the abnormality predicates during the minimization process. In the latter case, for example, MCCARTHY (1986) has introduced the notion of *prioritized circumscription*. Prioritized circumscription excludes "unwanted" minimal models by realizing priorities when multiple predicates are to be minimized. In the example under discussion the abnormality predicate related to the more specific category (i.e. AB_2) has priority over the abnormality predicate related to the less specific one (i.e. AB_1). Thus, in the selected models Albin comes out as a normal Albino raven and as an abnormal raven. In order to choose the priorities appropriately, however, it needs an external mechanism that operates outside of the compositional module of semantic interpretation. In other words: The meanings of generic sentences are units that contain free parameters. These parameters can only be fixed by an external (context-dependent) mechanism. This suggests that the thesis of *logicality* cannot be satisfied in the case of circumscription. In Section 5 I will consider other models of nonmonotonic reasoning that can be mimicked by preference semantics and where the same conclusion can be drawn.

For comparing different systems of monotonic and/or nonmonotonic reasoning it is useful to focus the study of the corresponding logical systems on general properties of their consequence relations. In the thirties TARSKI (1930, 1935) stated quite generally some minimal requirements which a deductive consequence relation \models must fulfill if it is truly to be a logical notion. GABBAY (1985) was probably the first who has stated and investigated some minimal requirements which a nonmonotonic consequence relation \triangleright should

satisfy to represent a bona fide nonmonotonic logic. Definition 2 combines both conceptions in a general framework for nonmonotonic formalisms called *nonmonotonic calculus* (cf. MORREAU 1992).

Definition 2: Let \mathcal{L} be a language and let \models and \triangleright be relations on subsets of the formulas of \mathcal{L} . The triple $(\mathcal{L}, \models, \triangleright)$ is called a **nonmonotonic calculus** just in case

- (1) \models (the monotonic core) is Tarskian, i.e. it satisfies the following principles (here Γ and Γ' range over sets of formulas and Φ over isolated formulas):
 - REFLEXIVITY: $\Gamma \models \Gamma$
 - CUT: if $\Gamma \models \Gamma'$ and $\Gamma' \cup \Gamma'' \models \Phi$, then $\Gamma \models \Phi$
 - MONOTONICITY: if $\Gamma \models \Phi$, then $\Gamma \cup \Gamma'' \models \Phi$
- (2) \triangleright (the nonmonotonic periphery) is a cumulative consequence relation, i.e. it satisfies the following principles:
 - REFLEXIVITY $\Gamma \triangleright \Gamma$
 - CUT: if $\Gamma \triangleright \Gamma'$ and $\Gamma' \cup \Gamma'' \triangleright \Phi$, then $\Gamma \triangleright \Phi$
 - WEAK MONOTONICITY: if $\Gamma \triangleright \Phi$ and $\Gamma \triangleright \Gamma'$, then $\Gamma \cup \Gamma'' \triangleright \Phi$
- (3) The periphery \triangleright extends the core \models , i.e. if $\Gamma \models \Gamma'$, then $\Gamma \triangleright \Gamma'$

Intuitively, the monotonic core inference notion \models extracts from a set of premises Γ those conclusions which are non-defeasible with regard to any extension of Γ . The nonmonotonic inference notion \triangleright adds to the conclusions drawn by means of \models a periphery of defeasible or soft conclusions (sometimes called *plausible inferences* drawn from Γ). With regard to the nonmonotonic periphery GABBAY (1985) and others have suggested that a nonmonotonic consequent relation may reasonably be required to retain at least the part of monotonicity which the principle of WEAK MONOTONICITY (also called CAUTIOUS MONOTONICITY) would salvage.

GABBAY (1985) argued for his three conditions REFLEXIVITY, CUT, and WEAK MONOTONICITY on proof-theoretic grounds but provided no semantics against which to check them. The following theorem shows that SHOHAM's (1988) general model theory provides the warranted semantical foundation for these conditions. Moreover, it states that classical entailment and preferential entailment are related in the sense of Definition 2 and constitute a nonmonotonic calculus.

Theorem 1: Let \mathcal{L} be a language with a model-theoretic semantics and \sqsubseteq a strict partial order on the interpretations of \mathcal{L} . Furthermore, let \models be the relation of classical entailment and $\triangleright_{\sqsubseteq}$ the relation of preferential entailment (see Definition 1). Then $(\mathcal{L}, \models, \triangleright_{\sqsubseteq})$ constitutes a nonmonotonic calculus (presumed it holds some kind of *limit assumption* (STALNAKER 1968; LEWIS 1973) or *smoothness condition* (KRAUS, LEHMANN & MAGIDOR 1990) that ensures that each satisfiable set of formulae has a preferred model).

The less trivial parts of the proof are given, for example, in KRAUS, LEHMANN

& MAGIDOR (1990). The same paper contains a proof of the fact that preferential entailment satisfies the following condition (when propositional logic is the basic language):

OR: if $\Phi_1 \supset \Psi$ and $\Phi_2 \supset \Psi$, then $\Phi_1 \vee \Phi_2 \supset \Psi$.

Furthermore, there is a representation theorem that states that each cumulative consequence relation that satisfies OR can be represented by preferential entailment with regard to some preference order \sqsubset on interpretations to the logic. This result suggests (at least in the propositional case) that each nonmonotonic calculus that satisfies OR can be represented by preference semantics. In spite of this result it may be useful to consider alternative semantic frameworks that provide a semantic foundation for nonmonotonicity and yield the the same conditions as those formulated in Definition 2 for alternative reasons.

3. Dynamic semantics

Dynamic semantics is another general framework which can provide a semantic foundation of several aspects of nonmonotonicity. In contrast to preference semantics this approach abandons the traditional assumption of identifying meanings with truth-conditions. In dynamic semantics it is not the truth-conditional content, but the information change potential of a sentence that is regarded as constituting its meaning. The information change potential of a sentence can be considered an operation on the domain of the so-called *information states*. Depending on the definition of an information state different examples of dynamic semantics have been developed such as dynamic predicate logic (GROENENDIJK & STOKHOF 1991) and update semantics (VELTMAN 1994). These systems can be considered formalizing different aspects of the dynamics of discourse. In this section I will provide a general formulation which is oriented towards VELTMAN's (1994) update semantics.

Let σ be an information state and Φ a formula of a basic language \mathcal{L} with meaning $[\Phi]$ (an operation on information states). Let us write $\sigma[\Phi]$ for the information state that results by applying $[\Phi]$ to σ . Intuitively, $\sigma[\Phi]$ denotes the information state that results when σ is updated with Φ . Truth (or acceptance) now becomes a secondary notion which can be derived from the primary notion of information change. A formula Φ will be called true (accepted) in σ if the information conveyed by Φ is already subsumed by σ : if σ is updated with Φ , the resulting information state turns out to be σ again.

Definition 3: Let σ be an information state and Φ a formula in \mathcal{L} . Then Φ is called **true (accepted)** in σ , $\sigma \Vdash \Phi$, iff $\sigma[\Phi] = \sigma$.

Now let us borrow from VELTMAN (1994) the following two notions of entailment which I will call for obvious reasons static and dynamic entailment.

Definition 4: let be $\Gamma = (\Psi_1, \dots, \Psi_n)$ a sequence of formulas in \mathcal{L} and let Φ be an isolated formula of \mathcal{L} .

(a) (Ψ_1, \dots, Ψ_n) **statically entails** Φ (written $\Gamma \models_{\text{stat}} \Phi$) iff from $\sigma \Vdash \Psi_i$ for every $1 \leq i \leq n$ it follows that $\sigma \Vdash \Phi$.

(b) (Ψ_1, \dots, Ψ_n) **dynamically entails** Φ (written $\Gamma \supseteq_{\text{dyn}} \Phi$) iff $[\Psi_1] \dots [\Psi_n] \vdash \Phi$. (Here $\mathbf{0}$ denotes the so-called minimal information state)

A sequence of premises statically entails a conclusion iff one cannot accept all the premises without having to accept the conclusion as well. A sequence of premises dynamically entails a conclusion iff updating the minimal information state $\mathbf{0}$ with the premises in the given order yields an information state that affirms the conclusion. (Without debating the lattice-theoretic structure of information states, I will assume that there is such a thing as the minimal information state). From the definition of static entailment it immediately follows that it is Tarskian, i.e. it satisfies REFLEXIVITY, CUT, and MONOTONICITY. Note that the order of the premises in (Ψ_1, \dots, Ψ_n) is not essential for its static entailments.

For discussing the notion of dynamic entailment we have to consider the following principles:

IDEMPOTENCE: For any information state σ and formula Φ of \mathcal{L} : $\sigma[\Phi][\Phi] = \sigma[\Phi]$ (i.e. $\sigma[\Phi] \Vdash \Phi$).

COMMUTATIVITY: For any information state σ and formulae Φ_1 and Φ_2 of \mathcal{L} : $\sigma[\Phi_1][\Phi_2] = \sigma[\Phi_2][\Phi_1]$.

STABILITY: For any information state σ and formulae Φ, Ψ of \mathcal{L} : if $\sigma \Vdash \Phi$, then $\sigma[\Psi] \Vdash \Phi$.

In the present context the principle of IDEMPOTENCE can be taken for granted. Intuitively, updating an information state with Φ has the consequence that Φ becomes accepted after the information change. Assuming IDEMPOTENCE yields \supseteq_{dyn} as a reflexive relation. Furthermore, we have the following order-sensitive variants of CUT and WEAK MONOTONICITY:

CUT⁺: If $(\Psi_1, \dots, \Psi_n) \supseteq_{\text{dyn}} \Lambda$ and $(\Psi_1, \dots, \Psi_n, \Lambda) \supseteq_{\text{dyn}} \Phi$, then $(\Psi_1, \dots, \Psi_n) \supseteq_{\text{dyn}} \Phi$.

WEAK MONOTONICITY⁺: If $(\Psi_1, \dots, \Psi_n) \supseteq_{\text{dyn}} \Phi$ and $(\Psi_1, \dots, \Psi_n) \supseteq_{\text{dyn}} \Lambda$, then $(\Psi_1, \dots, \Psi_n, \Lambda) \supseteq_{\text{dyn}} \Phi$.

The order of the premises is essential. Violations of COMMUTATIVITY, for instance, make the following variant of CUT invalid in the general case:

If $(\Psi_1, \dots, \Psi_n) \supseteq_{\text{dyn}} \Lambda$ and $(\Lambda, \Psi_1, \dots, \Psi_n) \supseteq_{\text{dyn}} \Phi$, then $(\Psi_1, \dots, \Psi_n) \supseteq_{\text{dyn}} \Phi$.

However, if we assume COMMUTATIVITY to be a general principle, then the order of the premises does no matter and CUT and WEAK MONOTONICITY hold in the sense of Definition 2.

Now let us presuppose COMMUTATIVITY and IDEMPOTENCE as valid principles of information change. Then it is simple to show that \supseteq_{dyn} extends \models_{stat} . The discussion so far can be summarized in the following theorem.

Theorem 2: Let be \models_{stat} and \supseteq_{dyn} the relations of static and dynamic entailment as given in Definition 4 and let be COMMUTATIVITY and IDEMPOTENCE valid principles of information change. Then $(\mathcal{L}, \models_{\text{stat}}, \supseteq_{\text{dyn}})$ constitutes a nonmonotonic calculus.

If STABILITY is satisfied as an additional principle, then \supseteq_{dyn} becomes monotonic. The two notions of entailment, \models_{stat} and \supseteq_{dyn} , completely agree if moreover the following condition of *expressibility* is satisfied: every information state σ can be expressed as $\sigma = 0[\Gamma]$, for some set Γ of formulas of \mathcal{L} .

The Theorems 1 and 2 state that both preference semantics and dynamic semantics are qualified for constituting a nonmonotonic calculus in the sense of Definition 2. One difference between preference semantics and dynamic semantics is that the former framework needs no further restriction to reproduce the conditions of Definition 2 whereas the latter needs two further conditions that are not constitutive for this framework: IDEMPOTENCE and COMMUTATIVITY. Though it is plausible to assume IDEMPOTENCE in update semantics it is not so for COMMUTATIVITY. VELTMAN (1994), for instance, has discusses some clear violations of COMMUTATIVITY. Non-commutativity is found in connection with modal qualifications like *presumably*, *probably*, *must*, *may* or *might*. As an illustration let us consider the following two sequences (cf. VELTMAN 1994: 2):

- *Somebody is knocking at the door... Presumably, it's John... It's Mary.*
- *Somebody is knocking at the door... It's Mary ... Presumably, it's John.*

These two sentences consist of the same sentences, only the order differs. Nevertheless the first sequence makes sense, whereas the second does not. Explanation: if I hear someone knocking at the door, I may have the expectation that it is, say, John. Still, in that situation it is perfectly possible for me to find out that it is Mary who is knocking, not John. On the other hand, once I have found out that Mary is knocking at the door, it is excluded that it is John, and then it is quite absurd to say that I have the expectation that it is John.

The capacity of dynamic semantics to comprehend the phenomenon of non-commutativity in natural language semantics is a crucial advantage of the dynamic paradigm. In Section 6 I will use this capacity in order to develop an improved update semantics for adverbs of generic quantification.

4. Inference patterns for prototypes

Intuitively, there are two kinds of reasoning with sentences involving generic quantification. The first corresponds to a relation between the premises and conclusion of an argument which holds whenever truth is invariably passed from the former to the latter. Inferring *Ravens are (normally) black and can fly* from

Ravens are (normally) black and *Ravens can (normally) fly* presumably is a good candidate of this kind of reasoning. The second kind of reasoning with generic sentences enables us to jump to *reasonable* conclusions whenever we don't have unlimited resources for gathering all the facts which bear on our decisions. This non truth-preserving mode of inferring is called *plausible reasoning*. It is exemplified by inference patterns such as DEFEASIBLE MODUS PONENS and PENGUIN PRINCIPLE mentioned above.

In Section 2 the distinction between a monotonic core and a nonmonotonic periphery has been introduced in order to reflect the two modes of reasoning on a very abstract level. Furthermore, we have seen how the paradigms of preference semantics and dynamic semantics, respectively, are able to explicate the two kinds of entailment that correspond to that distinction. In this Section I will give an (incomplete) list of the patterns of monotonic (non-defeasible) and nonmonotonic (defeasible) inference involved in reasoning with generic quantification. For obvious reasons, I will call these schemes *inference patterns for prototypes*.

Let us assume a binary generalized quantifier $>$ which is the supposed common integrator of adverbs of generic quantification as *normally*, *typically*, and *commonly*. Intuitively, a formula $\phi > \psi$ (ϕ and ψ are one-place predicate expressions) is intended to mean that ϕ s are normally ψ s. Among the patterns of nondefeasible inference that the logical analysis of the adverbs of generic quantification under consideration should explain are the following two:

MI 1: DISJUNCTIVE ANTECEDENTS or DUDLEY DOORITE

$\phi_1 > \psi, \phi_2 > \psi \models (\phi_1 \vee \phi_2) > \psi$;

e.g. *Mathematicians are (normally) good musicians, physicists are (normally) good musicians* \models *mathematicians and physicists are (normally) good musicians*.

MI 2: LOGICAL CLOSURE IN THE CONSEQUENT

If $\psi_1(x), \dots, \psi_n(x) \models \gamma(x)$ then $\phi > \psi_1, \dots, \phi > \psi_n \models \phi > \gamma$;

e.g. *Ravens are (normally) black, ravens can (normally) fly* \models *Ravens are (normally) black and can fly* (CONJUNCTION IN THE CONSEQUENT);

Ravens are (normally) black \models *ravens are (normally) black or can fly* (WEAKENING IN THE CONSEQUENT).

In order to express the patterns of defeasible inference in a sufficient general way we need the notion of (logically) independent predicate expressions:

Definition 5 (from ASHER & MORREAU 1991: 26): A system $\{\phi_1, \dots, \phi_n\}$ of one-place predicate expressions of a basic language \mathcal{L} is called **logically independent** just in case for all individual constants c , each Boolean combination containing at most one instance of each of $\phi_1(c), \dots, \phi_n(c)$ is satisfiable.

Now the nonmonotonic component of any reasonable logical analysis of generic sentences should explain at least the following patterns of inference:

NMI 1: DEFEASIBLE MODUS PONENS

For any logically independent system $\{\phi, \psi\}$ of predicate expressions of \mathcal{L} and any individual constant c of \mathcal{L} it holds:

$\phi > \psi$, $\phi(c) \triangleright \psi(c)$, but not $\phi > \psi$, $\phi(c)$, $\neg \psi(c) \triangleright \psi(c)$;

e.g. *Ravens are (normally) black, Albin is a raven* \triangleright (presumably) *Albin is black*, but not: *Ravens are (normally) black, Albin is a raven, Albin is not black* \triangleright (presumably) *Albin is black*.

NMI 2: THE NIXON DIAMOND

For any logically independent system $\{\phi_1, \phi_2, \psi\}$ of predicate expressions of \mathcal{L} and any individual constant c of \mathcal{L} it holds neither

$\phi_1 > \psi$, $\phi_2 > \neg \psi$, $\phi_1(c)$, $\phi_2(c) \triangleright \psi(c)$ nor $\phi_1 > \psi$, $\phi_2 > \neg \psi$, $\phi_1(c)$, $\phi_2(c) \triangleright \neg \psi(c)$;

e.g. There is irresolvable conflict in the following: *Quakers are (normally) pacifists, Republicans are (normally) nonpacifists, Nixon is Quaker and Republican*.

NMI 3: TAXONOMIC PENGUIN PRINCIPLE

For any logically independent systems $\{\phi_1, \psi\}$ and $\{\phi_2, \psi\}$ of predicate expressions of \mathcal{L} and any individual constant c of \mathcal{L} it holds: If ϕ_2 is more specific than ϕ_1 (i.e. $\phi_2(x) \models \phi_1(x)$), then

$\phi_1 > \psi$, $\phi_2 > \neg \psi$, $\phi_2(c) \triangleright \neg \psi(c)$, but not $\phi_2 > \phi_1$, $\phi_1 > \psi$, $\phi_2 > \neg \psi$, $\phi_2(c) \triangleright \psi(c)$;

e.g. suppose as logical truth that *all penguins are birds*, then *birds can (normally) fly, penguins (normally) can't fly, Tweety is a penguin* \triangleright (presumably) *Tweety can't fly*, but not \triangleright (presumably) *Tweety can fly*.

NMI 4: WEAK PENGUIN PRINCIPLE

For any logically independent system $\{\phi_1, \phi_2, \psi\}$ of predicate expressions of \mathcal{L} and any individual constant c of \mathcal{L} it holds

$\phi_2 > \phi_1$, $\phi_1 > \psi$, $\phi_2 > \neg \psi$, $\phi_2(c) \triangleright \neg \psi(c)$, but not

$\phi_2 > \phi_1$, $\phi_1 > \psi$, $\phi_2 > \neg \psi$, $\phi_2(c) \triangleright \psi(c)$;

e.g. *Students are (normally) adults, adults are (normally) employed, students are (normally) not employed, Sam is a student* \triangleright (presumably) *Sam is not employed*, but not \triangleright (presumably) *Sam is employed*.

NMI 5: GRADED NORMALITY

For any logically independent system $\{\phi_1, \phi_2, \psi\}$ of predicate expressions of \mathcal{L} and any individual constant c of \mathcal{L} it holds

$\phi > \psi_1$, $\phi > \psi_2$, $\phi(c)$, $\neg \psi_1(c) \triangleright \psi_2(c)$, but not

$\phi > \psi_1$, $\phi > \psi_2$, $\phi(c)$, $\neg \psi_1(c) \triangleright \neg \psi_2(c)$

e.g. *Ravens are (normally) black, ravens can (normally) fly, Albin is a raven, Albin is not black* \triangleright (presumably) *Albin can fly*, and does not \triangleright (presumably) *Albin cannot fly*.

NMI 6: SEPARATION EFFECT

For any logically independent systems $\{\phi, \psi_1\}$ and $\{\phi, \psi_2\}$ of predicate expressions of \mathcal{L} and any individual constant c of \mathcal{L} it holds: whenever $\{\psi_1(c), \neg\psi_2(c)\}$ is logically consistent, then

$\phi > (\psi_1 \vee \psi_2), \phi(c), \neg\psi_1(c) \models \psi_2(c)$, but not $\phi > \psi_1, \phi(c), \neg\psi_1(c) \models \psi_2(c)$.
e.g. *Tigers (normally) have four or five legs, Shere Khan is a tiger. Shere Khan does not have four legs \models (presumably) Shere Khan has five legs*, but not: *Tigers (normally) have four legs, Shere Khan is a tiger, Shere Khan does not have four legs \models (presumably) Shere Khan has five legs*.

All of these patterns of inference are familiar from the literature. An possible exception is the last pattern of defeasible reasoning which is due to VELTMAN (1991). VELTMAN (1994) has discussed this pattern in connection with the intuitive validity of LOGICAL CLOSURE IN THE CONSEQUENT (MI 2). In the subsequent discussion I comment on some of the intuitive and theoretic problems that are associated with these patterns of inference.

5. Some models of nonmonotonic reasoning and three challenging problems

There are several models of nonmonotonic reasoning that provide, in some more or less direct sense, a definition of the generic quantifier $>$ and capture at least some of the inference patterns for prototypes listed above. In this section I will assume some familiarity with these models (for introductory material see, for example, BREWKA 1991; MORREAU 1992).

The models considered here are the following:

- (a) *Circumscription*: (MCCARTHY 1980; 1986)
 $\phi > \psi =_{\text{def}} \forall x(\phi(x) \wedge \neg \text{Ab}_i(x) \rightarrow \psi(x))$
- (b) *Default Logic*: (REITER 1980)
 $\phi > \psi =_{\text{def}} \phi(x) : \mathbf{M} \psi(x) / \psi(x)$
- (c) *Autoepistemic Logic*: (MOORE 1985)
 $\phi > \psi =_{\text{def}} \Box \phi(x) \wedge \neg \Box \neg \psi(x) \rightarrow \Box \psi(x)$
- (d) *Update Semantics* (VELTMAN 1994)
 $\phi > \psi : \Phi \rightarrow \Psi$
- (e) *Common Sense Entailment* (ASHER & MORREAU 1991)
 $\phi > \psi =_{\text{def}} \phi(x) >_x \psi(x)$

Among these models the first three can be seen as to be compatible with the unifying framework of preference semantics. In the case of *circumscription*, the compatibility with preference semantics is perfectly clear. With respect to *default logic* and *autoepistemic logic* SHOHAM (1988) has shown that these systems can be described - at least approximately - by a preference ordering on the interpretations of a corresponding language \mathcal{L} (see also BREWKA 1991). VELTMAN's (1994) *update semantics* is the first (and perhaps the best hitherto existing) representative of dynamic semantics that deals with inference patterns for prototypes. *Common Sense Entailment* (ASHER & MORREAU 1991) rests upon

elements of the dynamic approach. Strictly speaking, however, it is neither a representative of dynamic semantics nor a representative of preference semantics, but it should be seen as a theory *sui generis* involving a three-step design for dynamic information models: (i) a set of information states containing a particular information state representing ignorance or minimal knowledge, (ii) an update function, (iii) assuming individuals to be normal representatives of kinds - modelled by defining the notion of (iteratized) normalization.

Figure 1 compares the mentioned models of nonmonotonic reasoning with respect to the inference patterns for prototypes listed in Section 4.

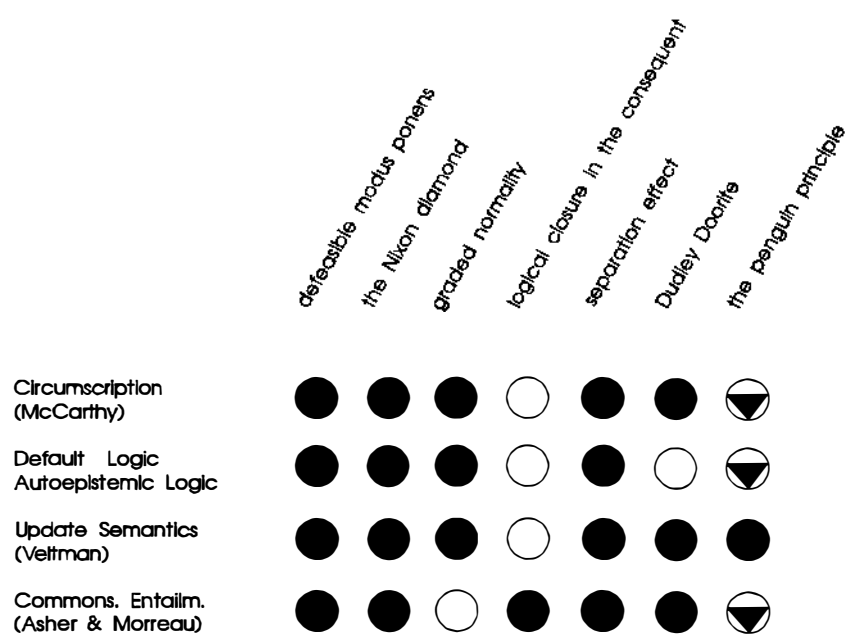


Figure 1: Some models for representing nonmonotonic reasoning patterns (● : the corresponding inference pattern has been correctly described; ○ : the inference pattern cannot be described correctly; ◐ : in order to describe the inference pattern correctly some external mechanism is needed)

Each of the models captures the classical patterns of nonmonotonic reasoning DEFEASIBLE MODUS PONENS and the NIXON DIAMOND. Noteworthy, the same holds for the pattern of inference which I have called SEPARATION EFFECT.

The latter pattern has been introduced by VELTMAN (1994) in order to offer an argument against WEAKENING OF THE CONSEQUENT (a special case of the general pattern LOGICAL CLOSURE IN THE CONSEQUENT, MI 2) in theories of defeasible reasoning involving generics. Consider the two generic sentences

Tigers normally have four legs.
Tigers normally have four or five legs.

According to the general principle **MI 2** the first sentence makes a statement which is logically stronger than that made by the second sentence. Despite of this fact the second sentence seems to suggest conclusions that are not sanctioned by the first sentence. To give an example, VELTMAN (1994) assumes an tiger called *Shere Khan* who does not have four legs. Then the second but not the first sentence suggests to assume that *Shere Khan* has five legs (cf. the pattern called SEPARATION EFFECT, **NMI 6**). VELTMAN (1994) concludes from these observations, which are quite right, I think, that the principle **MI 2** must be given up (especially the special case of this principle called WEAKENING OF THE CONSEQUENT). However, as argued by MORREAU (1992), this conclusion is not convincing. It is the nonmonotonic notion of consequence which is involved in the *Shere Khan*-example and not the monotonic one. The possibility that a conclusion like *Shere Khan has five legs* can be withdrawn on moving from a set of premises to a logically stronger set (by substituting *tigers normally have four legs* for *tigers normally have four or five legs*) is exactly the sort of thing that a defeasible consequence notion makes room for. Consequently, I think VELTMAN's example should not be taken as an argument against LOGICAL CLOSURE IN THE CONSEQUENT, but as a challenge to develop a theory that simultaneously satisfies LOGICAL CLOSURE IN THE CONSEQUENT and the SEPARATION EFFECT. As can be seen from Figure 1, ASHER & MORREAU's (1991) *Commonsense Entailment* is such a theory. In Section 6 another one is developed.

The second challenge has to do with the coexistence of GRADED NORMALITY and LOGICAL CLOSURE IN THE CONSEQUENT. Figure 1 tells us that none of the known models satisfies both patterns simultaneously. The question arises of whether this is an accidental situation or whether this happens for systematic reasons. MORREAU (1992: 146) has proved a theorem that shows the second possibility is right: given a nonmonotonic calculus in the sense of Definition 2 and given that \supseteq (the nonmonotonic periphery) is closed under *modus ponens*, then **MI 2** (LOGICAL CLOSURE IN THE CONSEQUENT) and **NMI 5** (GRADED NORMALITY) cannot be satisfied simultaneously. Each of the given models of defeasible reasoning is conform to the principles underlying the premises of MORREAU's theorem. Consequently, in such cases the incompatibility of **MI 2** and **NMI 5** is a matter of logical necessity. Furthermore, this result can be generalized with respect to *all* models that are formulated within the paradigm of preference semantics (as a consequence of Theorem 1).

With regard to dynamic semantics the situation looks different. Only if we have COMMUTATIVITY as a valid principle of information change, then the premises of MORREAU's theorem are satisfied, i.e. we are concerned with a nonmonotonic calculus in the strict sense of Definition 2 (cf. Theorem 2). VELTMAN's (1994) theory is an example for a dynamic approach where COMMUTATIVITY has been assumed to be a valid principle of information change when it comes to evaluate adverbs of generic quantification. I take it as a

challenge to develop a version of *update semantics* where COMMUTATIVITY is violated in some cases, and, as a consequence, MI 2 and NMI 5 can be satisfied simultaneously. In the next section I will present such a theory.

The third challenge concerns the question of explanatory adequacy. Take as an example the pattern of defeasible inference NMI 4 (WEAK PENGUIN PRINCIPLE). None of the known models deals with this pattern of inference in a truly explanatory way. In *Prioritized Circumscription*, for example, one has to fix a "minimization strategy" saying which predicates are to be minimized and in which order (cf. Section 2). The intuitive idea that more specific information takes precedence over less specific then must be explicitly built into this strategy. Certainly, in this way an "extraneous ghost" can steer the mechanism toward the intuitively desired conclusion (expressed by the *penguin principle*), but this is not a true explanation in any interesting scientific sense. It appears to me that the same evaluation is valid in case of the other models that can be described (at least approximately) by the preferential approach, especially (prioritized) default logic and (hierarchical) autoepistemic logic (see BREWKA 1991). Taken together, these observations provide circumstantial evidence that the thesis of logicity cannot be satisfied in case of preference semantics if this approach is to describe the inference patterns for prototypes.

However, even "ghost-free" models such as ASHER & MORREAU's (1991) *Commonsense Entailment* and VELTMAN's (1994) *Update Semantics* fail to give a true explanation of the PENGUIN PRINCIPLE in the general case. At best, these models can be seen as deriving this principle from a related condition on another level of description (e.g. from a condition on model frames). How to explain the PENGUIN PRINCIPLE from independently motivated conditions of information update? I consider this a fundamental challenge and an opportunity that dynamic semantics should make use of.

6. An update semantic for prototypes

The final part of the paper develops a version of dynamic semantics that deals with the three challenges mentioned before. To simplify matters, I use VELTMAN'S (1994) language \mathcal{L}_3^{At} as basic language. This language is an extension of ordinary propositional logic \mathcal{L}_0^{At} with At as its nonlogical vocabulary (the atomic sentences), and has in its logical vocabulary two additional unary operators *necessarily* and *presumably*, one additional binary operator " \rightarrow " ($\Phi \rightarrow \Psi$ is to be read as Φ *normally implies* Ψ), and a punctuation sign denoted by ";" (for sequence).

In accordance with STALNAKER (1968), VELTMAN (1994), and others an epistemic approach to semantics is adopted. The fundamental notion of the analysis is that of a *knowledge state*. First of all, a knowledge state is a subset s of a given domain W of possible worlds and may be thought of as representing the *factual* knowledge of an epistemic subject. In addition to this "external" component, there is an "internal" component in the semantics which in a sense mirrors a person's expectations and presumptions. This component uses several layers of possible hypotheses $\Delta_1/\dots/\Delta_n$ to represent different degrees of

reliability. (This idea corresponds to RESCHER's (1964) ordered subtheories and BREWKA's (1991) level default theories, but it deviates in significant respects from VELTMAN's (1994) *expectation frames*). Each part Δ_i of the "internal" component contains propositions (set of possible worlds) that encode the default rules (at level i) that the agent is acquainted with. Rules at level i are of higher reliability than rules at level j if $i < j$. Furthermore, a proposition w is assumed which is determined by the content of the strict rules that the agent is acquainted with. These rules do not allow for exceptions; consequently, it is assumed that $s \subseteq w$. In some sense the proposition w can be seen to encode the rules with the highest reliability.

Definition 6: Let be $W = \mathcal{P}(At)$ the set of possible worlds. Then $\sigma = \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle$ is a **knowledge state** (on the basis of a layered system of expectations) iff

- (a) $s \subseteq w \subseteq W$
- (b) Δ_i ($1 \leq i \leq n$) are sets of propositions

The knowledge state $\mathbf{0} = \langle W, W, [\emptyset] \rangle$ is called the minimal state; the knowledge state $\mathbf{1} = \langle \emptyset, \emptyset, [\emptyset] \rangle$ is called the absurd state.

Furthermore, the notions of *resulting expectation sets* (corresponding to RESCHER's (1964) notion of *preferred maximally consistent subsets*) and *weak acceptability* are defined as follows:

Definition 7: A knowledge state $\sigma = \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle$ has a **resulting expectation set** Δ' (of propositions) iff $\Delta' = \Delta_1' \cup \dots \cup \Delta_n'$ for some $\Delta_i' \subseteq \Delta_i$ ($1 \leq i \leq n$) and for all k ($1 \leq k \leq n$) $\Delta_1' \cup \dots \cup \Delta_k'$ is a maximal subset of $\Delta_1 \cup \dots \cup \Delta_k$ that is consistent with the proposition s .

Definition 8: A proposition δ is **weakly accepted** in σ iff $\delta \supseteq s \cap [\cap \Delta']$ for each resulting expectation set Δ' of σ .

Now we are ready to explain our update clauses for the formulae of the basic language \mathcal{L}_3^{At} . The update clauses for ordinary propositional formulae Φ , for *necessarily*(Φ), *presumably*(Φ), and $(\Phi; \Psi)$ are simply translations of VELTMAN's (1994) update clauses into the present framework. Let Φ be a formula of \mathcal{L}_0^{At} and let $\llbracket \Phi \rrbracket$ designate the proposition expressed by Φ . Notice that $\llbracket \Phi \rrbracket$ can be calculated using the usual recursive conditions: $\llbracket p \rrbracket = \{z \in W : p \in z\}$ for atoms p , $\llbracket \neg \phi \rrbracket = W \setminus \llbracket \phi \rrbracket$, $\llbracket \phi \wedge \psi \rrbracket = \llbracket \phi \rrbracket \cap \llbracket \psi \rrbracket$, $\llbracket \phi \vee \psi \rrbracket = \llbracket \phi \rrbracket \cup \llbracket \psi \rrbracket$ for any formulae Φ and Ψ von \mathcal{L}_0^{At} . Now let us assume the following: $\sigma[\Phi]$ changes the s -component of the knowledge state into $s \cap \llbracket \Phi \rrbracket$ (leaving everything else unchanged); $\sigma[necessarily](\Phi)$ changes the s -component into $s \cap \llbracket \Phi \rrbracket$ and the w -component into $w \cap \llbracket \Phi \rrbracket$; $\sigma[presumably](\Phi)$ describes a *test* and checks whether $\llbracket \Phi \rrbracket$ is weakly accepted in σ or not; the punctuation sign ";" designates functional composition, i.e. $\sigma[x_1; x_2] = \sigma[x_1][x_2]$.

The crucial deviation from VELTMAN (1994) concerns the update clause for $\Phi \rightarrow \Psi$. For its formulation we need an operator G on knowledge states that de-

letes the factual knowledge from σ , i.e. $G(\langle s, w, [\Delta_1/\dots/\Delta_n] \rangle) = \langle w, w, [\Delta_1/\dots/\Delta_n] \rangle$. It is intuitively plausible to consider $\Phi \rightarrow \Psi$ accepted in σ iff the following condition is satisfied: the proposition $\llbracket \Psi \rrbracket$ is weakly accepted in $G(\sigma)[\Phi]$. I tend to consider this condition a variant on the Ramsey test in line with the rephrasing of $\Phi \rightarrow \Psi$ as asserting that Φ in the normal course of events tends to imply Ψ . For convenience, let us call this condition the Ramsey condition.

According to the present dynamic approach, the update clause for $\Phi \rightarrow \Psi$ describes a kind of learning with *minimal* effort. Being more precise, it is assumed that a knowledge state σ does not change by updating with $\Phi \rightarrow \Psi$ when the corresponding Ramsey condition already is satisfied with respect to σ . In the other case it is assumed that a new (default) rule with content $\llbracket \Phi \rightarrow \Psi \rrbracket$ is added to the hypotheses actually at the *smallest level of reliability* where successful learning is possible (i.e. the Ramsey condition is satisfied with respect to the resulting knowledge state).

It can be shown that successful learning is always possible when $w \cap \llbracket \Phi \rightarrow \Psi \rrbracket \neq \emptyset$. (In the case that $w \cap \llbracket \Phi \rightarrow \Psi \rrbracket = \emptyset$ let us identify $\sigma[\Phi \rightarrow \Psi]$ with the *absurd* knowledge state 1). In case of normality expressions with "conflicting information" (consider, for instance, the pair $BIRD \rightarrow CAN-FLY$, $PENGUIN \rightarrow \neg CAN-FLY$) it is this idea of *adding information at the smallest level of reliability where successful updating is possible* that leads to the effect that more specific information takes precedence over less specific one. Consequently, the pattern of defeasible reasoning called PENGUIN PRINCIPLE can be explained without further *ad hoc*-stipulations, specifically, without an external "ghost" that fixes the relevant parameters.

In order to give an exact formulation of the update clause for " \rightarrow " it is useful to consider the following notion of the *i-refinement of a layered system of expectations*:

DEFINITION 9: Let $[\Delta_1/\dots/\Delta_n]$ be a layered system of expectations, δ a proposition, and $0 \leq i \leq n$. Then

$$[\Delta_1/\dots/\Delta_n] \cup_i \delta = \begin{cases} \{\{\delta\}/\Delta_1/\dots/\Delta_n\} & \text{if } i=0 \\ [\Delta_1/\dots/\Delta_i \cup \{\delta\}/\dots/\Delta_n] & \text{else} \end{cases}$$

is called the **i-refinement of a layered system of expectations** with δ .

Adding information at the *smallest* level of reliability now simply means performing the *maximal* i-refinement with the corresponding proposition.

Officially stated, the definition of the update semantic for the basic language \mathcal{L}_3^{At} is as follows:

DEFINITION 10: Let $\sigma = \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle$ be a knowledge state. For all ordinary propositional formulae Φ and Ψ and all sequences $(\chi_1; \chi_2)$ of $\mathcal{L}_3^{\text{At}}$ it holds:

$$\begin{aligned}
 (a) \quad \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle[\Phi] &= \begin{cases} 1 & \text{if } s \cap \llbracket \Phi \rrbracket = \emptyset \\ \langle s \cap \llbracket \Phi \rrbracket, w, [\Delta_1/\dots/\Delta_n] \rangle & \text{else} \end{cases} \\
 (b) \quad \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle[\text{nec } \Phi] &= \begin{cases} 1 & \text{if } s \cap \llbracket \Phi \rrbracket = \emptyset \\ \langle s \cap \llbracket \Phi \rrbracket, w \cap \llbracket \Phi \rrbracket, [\Delta_1/\dots/\Delta_n] \rangle & \text{else} \end{cases} \\
 (c) \quad \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle[\text{pres } \Phi] &= \begin{cases} \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle & \text{if } \llbracket \Phi \rrbracket \text{ is weakly} \\ & \text{accepted in } \sigma \\ 1 & \text{else} \end{cases} \\
 (d) \quad \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle[\Phi \rightarrow \Psi] &= \begin{cases} 1 & \text{if } w \cap \llbracket \Phi \rightarrow \Psi \rrbracket = \emptyset \\ \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle & \text{if } \llbracket \Psi \rrbracket \text{ is weakly} \\ & \text{accepted in } \mathbf{G}(\sigma)[\llbracket \Phi \rrbracket] \\ \langle s, w, [\Delta_1/\dots/\Delta_n] \cup_{i^*} \{ \llbracket \Phi \rightarrow \Psi \rrbracket \} \rangle & \text{else} \end{cases}
 \end{aligned}$$

In the last line i^* denotes the biggest natural number between 0 and n which satisfies the condition that $\llbracket \Psi \rrbracket$ is weakly accepted in the knowledge state $\langle w, w, [\Delta_1/\dots/\Delta_n] \cup_{i^*} \{ \llbracket \Phi \rightarrow \Psi \rrbracket \} \rangle[\Phi]$.

$$(e) \quad \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle[\chi_1; \chi_2] = \langle s, w, [\Delta_1/\dots/\Delta_n] \rangle[\chi_1][\chi_2]$$

It can be proved that each formula Φ of our basic language $\mathcal{L}_3^{\text{At}}$ satisfies $\sigma[\Phi] \Vdash \Phi$ (for any knowledge state σ). Consequently, the notion of dynamic entailment \supseteq_{dyn} complies with the principle of REFLEXIVITY: $\Phi \supseteq_{\text{dyn}} \Phi$. Furthermore, dynamic entailment complies with the principles of CUT⁺ and WEAK MONOTONICITY⁺ as introduced in Section 3. A particular examination of the proposed update model shows that all of the inference patterns for prototypes introduced in Section 4 hold with respect to \models_{stat} and \supseteq_{dyn} , respectively. In the case of the patterns of defeasible inference the order of the premises is important and should be as given in NMI 1 - NMI 6.

A disadvantage of the notion \supseteq_{dyn} is that it is highly non-commutative, i.e. the order of premises is essential in most cases. This disadvantage can be overcome with a notion of dynamic entailment that involves repeated updating with the premises, roughly: $\Gamma * \supseteq_{\text{dyn}} \Phi$ iff $\emptyset[\Gamma] \dots [\Gamma] \Vdash \Phi$. With respect to the notion $* \supseteq_{\text{dyn}}$ it can be shown that NMI 1 - NMI 6 hold independent on the order of the premises. Furthermore, it can be shown that \models_{stat} and $* \supseteq_{\text{dyn}}$ form a nonmonotonic calculus in the sense of Definition 2 (with CUT⁺ and WEAK MONOTONICITY⁺ instead of CUT and WEAK MONOTONICITY, respectively, in the definition of the nonmonotonic periphery).

Interestingly enough, a residue of non-commutativity remains. This residue is inevitable in order to secure the coexistence of GRADED NORMALITY and LOGICAL CLOSURE IN THE CONSEQUENT. But the order of the premises is significant only in cases where we have logical dependencies as in

Tigers normally have four legs; tigers normally have four or five legs.

References

- Asher, N., & Morreau, M. (1991), 'Commonsense entailment: A modal theory of nonmonotonic reasoning', in J. van Eijck (Ed.), *Logics in AI - Proceedings of the European workshop JELIA 90*, Springer-Verlag, Berlin.
- Brewka, G. (1991), *Nonmonotonic reasoning: logical foundations of commonsense*, Cambridge University Press, Cambridge.
- Gabbay, D. M. (1985), 'Theoretical foundations for non-monotonic reasoning in expert systems', in K. R. Apt (ed.), *Logic and models of concurrent systems*, Springer-Verlag, Berlin.
- Groenendijk, J., & Stokhof, M. (1991), 'Dynamic predicate logic', *Linguistics and Philosophy*, **14**, 39-100.
- Horty, F. J., Thomason, R. H., & Touretzky, D. S. (1987), 'A clash of intuitions: The current state of nonmonotonic multiple inheritance systems', in *Proceedings of the 10th IJCAI*, Milan, Italy.
- Konolige, K. G. (1987), 'On the relation between default and autoepistemic logic', Technical Note 407, SRI International, Menlo Park, Cal.
- Kraus, S., Lehmann, D., & Magidor, M. (1990), 'Nonmonotonic reasoning, preferential
- McCarthy, J. (1980), 'Circumscription - a form of non-monotonic reasoning', *Artificial Intelligence*, **13**, 27-39.
- (1986), 'Applications of circumscription to formalizing common-sense knowledge', *Artificial Intelligence*, **28**, 89-116.
- Moore, R. (1985), 'Semantical considerations on non-monotonic logic', *Artificial Intelligence*, **25**, 75-94.
- Morreau, M. (1992), *Conditionals in philosophy and artificial intelligence*, Ph.D. thesis, Universität Stuttgart.
- Reiter, R. (1980), A logic for default reasoning, *Artificial Intelligence*, **13**, 81-132.
- Rescher, N. (1964), *Hypothetical reasoning*, Reidel, Dordrecht.
- Shoham, Y. (1988), *Reasoning about change*, MIT Press, Cambridge, Mass.
- Stalnaker, R. C. (1968), 'A theory of conditionals', *American Philosophical Quarterly Monograph Series*, **2**, 98-112.
- Tarski, A. (1930), 'Fundamentale Begriffe der Methodologie der deduktiven Wissenschaften I', *Monatshefte für Mathematik und Physik*, **37**, 361-404.
- (1935), 'Grundzüge des Systemkalküls I', *Fund. Math.*, **25**, 503-526.
- Veltman, F. (1994), 'Defaults in update semantics', Accepted by *Journal of Philosophical Logic*. (ms. 1994, University of Amsterdam).