Nonmonotonic Reasoning, Preferential Models and Cumulative Logics *

Sarit Kraus,[†] Daniel Lehmann[‡] and Menachem Magidor[§]

Abstract

Many systems that exhibit nonmonotonic behavior have been described and studied already in the literature. The general notion of nonmonotonic reasoning, though, has almost always been described only negatively, by the property it does not enjoy, i.e. monotonicity. We study here general patterns of nonmonotonic reasoning and try to isolate properties that could help us map the field of nonmonotonic reasoning by reference to positive properties. We concentrate on a number of families of nonmonotonic consequence relations, defined in the style of Gentzen [13]. Both proof-theoretic and semantic points of view are developed in parallel. The former point of view was pioneered by D. Gabbay in [10], while the latter has been advocated by Y. Shoham in [38]. Five such families are defined and characterized by representation theorems, relating the two points of view. One of the families of interest, that of preferential relations, turns out to have been studied by E. Adams in [2]. The *preferential* models proposed here are a much stronger tool than Adams' probabilistic semantics. The basic language used in this paper is that of propositional logic. The extension of our results to first order predicate calculi and the study of the computational complexity of the decision problems described in this paper will be treated in another paper.

1 Introduction

1.1 Nonmonotonic reasoning

Nonmonotonic logic is the study of those ways of inferring additional information from given information that do not satisfy the monotonicity property satisfied by all methods based on classical (mathematical) logic. In Mathematics, if a conclusion is warranted on the basis of certain premises, no additional premises will ever invalidate the conclusion.

In everyday life, however, it seems clear that we, human beings, draw sensible conclusions from what we know and that, on the face of new information, we often have to take back previous conclusions, even when the new information we gathered in no way made us want to take back our previous assumptions. For example, we may hold the assumption that most birds fly, but that penguins are birds that do not fly and, learning that Tweety is a bird, infer that it flies. Learning that Tweety is a penguin, will in no way make us change our mind about the fact that most birds fly and that penguins are birds that do not fly, or about the fact that Tweety is a bird. It should make us abandon our conclusion about its flying capabilities, though.

*This work was partially supported by the Jean and Helene Alfassa fund for research in Artificial Intelligence

†Institute for Advanced Computer Studies and Department of Computer Science, University of Maryland, College Park, MD 20742 U.S.A. This work was done while the author was at Hebrew University.

‡Department of Computer Science, Hebrew University, Jerusalem 91904 (Israel)

§Department of Mathematics, Hebrew University, Jerusalem 91904 (Israel)

It is most probable that intelligent automated systems will have to do the same kind of (nonmonotonic) inferences.

Many researchers have proposed systems that perform such nonmonotonic inferences. The best known are probably: negation as failure [5], circumscription [29], the modal system of [30], default logic [35], autoepistemic logic [31] and inheritance systems [45]. Each of those systems is worth studying by itself, but a general framework in which those many examples could be compared and classified is missing. We provide here a first attempt at such a general framework, concentrating on properties that are or should be enjoyed by at least important families of nonmonotonic reasoning systems. An up-to-date survey of the field of nonmonotonic reasoning may be found in [36].

1.2 Nonmonotonic consequence relations

The idea that the best framework to study the deduction process is that of consequence relations dates back to A. Tarski [42], [41] and [43] (see [44] for an English translation) and G. Gentzen [12] (see [13] for an English translation and related papers). For an up-to-date view on monotonic consequence relations, the reader may consult [3]. Tarski studied the consequences of arbitrary sets of formulas whereas Gentzen restricted himself to finite such sets. In the presence of compactness, the difference between the two approaches is small for monotonic consequence relations. For nonmonotonic relations, many different notions of compactness come to mind, and the relation between Tarski's infinitistic approach and Gentzen's finitistic approach is much less clear. We develop here a finitistic approach in the style of Gentzen. In [28], D. Makinson developed, in parallel with and independently from our effort, an infinitistic view of nonmonotonic consequence relations. Later efforts in this direction, by M. Freund and D. Lehmann [9], have benefited from the results presented here.

D. Gabbay [10] was probably the first to suggest to focus the study of nonmonotonic logics on their consequence relations. This is a bold step to take since some of the nonmonotonic systems mentioned above were not meant to define a consequence relation, as was soon noticed by D. Israel in [17]. D. Gabbay asked the question: what are the minimal conditions a consequence relation should satisfy to represent a bona fide nonmonotonic *logic*? He proposed three: reflexivity (see equation 1 in section 3.1), cut (see equation 4) and weak monotonicity (see equation 5). Weak monotonicity has, since, been renamed cautious monotonicity by D. Makinson [28] and we shall follow this last terminology, notwithstanding the fact that D. Makinson has now opted for the term *cumulative monotony.* D. Gabbay argued for his three conditions on proof-theoretic grounds but provided no semantics against which to check them. He also assumed a *poor* underlying language for propositions, a language without classical propositional connectives. In [28], D. Makinson proposed a semantics for Gabbay's logic and proved a completeness result, for a *poor* language. His models have a definitely syntactic flavor, whereas the models presented here seem more truly semantic and more easily suggest rules of inference.

Independently, Y. Shoham in [39] and [38] proposed a general model theory for nonmonotonic inference. He suggested models that may be described as a set of worlds equipped with a *preference* relation: the preference relation is a partial order and a world v is preferable, in the eyes of the reasoner, to some other world w if he considers v to be more *normal* than w. One would then, in the model, on the basis of a proposition α, conclude, defeasibly, that a proposition β is true if all worlds that satisfy α and are *most normal* among worlds satisfying α also satisfy β. Shoham claimed that adequate semantics could be given to known nonmonotonic systems by using such a preference relation. He assumed a *rich* underlying language for propositions, containing all classical propositional connectives. The idea that nonmonotonic deduction should be modeled by some *normality* relation between worlds is very natural and may be traced back to J. McCarthy. It appears also in relation with epistemic logic in [15]. One of the conclusions of this paper will be that none of the nonmonotonic systems defined so far in the literature, except those based on conditional logic described in [6], [7] and [34], may represent all nonmonotonic inference systems that may be defined by preferential models. The framework of preferential models, therefore, has an expressive power that cannot be captured by negation as failure, circumscription, default logic or autoepistemic logic. We do not claim that all this expressive power is needed, but will claim that the systems mentioned above lack expressive power.

The main point of this work, therefore, is to characterize the consequence relations that can be defined by models similar to Shoham's in terms of proof-theoretic properties. To this end Gabbay's conditions have to be augmented. The class of models corresponding exactly to Gabbay's conditions is also characterized. The elucidation of the relations between proofs and models that is achieved in this paper will allow for the design of decision procedures tuned to different restrictions on the language of propositions or the knowledge bases. Such decision procedures (or heuristics) could be the core of automated engines of sensible inferences. This paper will not propose any specific system of nonmonotonic reasoning. Important steps towards such a system, taken after obtaining the results reported here but before the final redaction of this paper, are reported in [19], [22] and [20].

At this point it could be useful to state the philosophy of this paper concerning the relative importance of proof-theory and semantics. We consider, in this paper, the axiomatic systems as the main object of interest (contrary to the point of view expressed in [24] for example). The different families of models described in this paper and that provide semantics to the axiomatic systems are not considered to be an ontological justification for our interest in the formal systems, but only as a technical tool to study those systems and in particular settle questions of interderivability and find efficient decision procedures. Preliminary versions of the material contained in this paper appeared in [21] and [18].

1.3 Conditional logic

In this subsection, the relation between our work and conditional logic will be briefly surveyed. Since the link, we claim, is mainly at the level of the formal systems and not at the semantic level, the reader uninterested in conditional logic may easily skip this subsection.

This work stems from a very different motivation than the vast body of work concerned with conditional logic and its semantics, (see in particular [40],[24] and [23]) which is surveyed in [32]. Two main differences must be pointed at. The first difference is that conditional logic considers a binary intensional connective that can be embedded inside other connectives and even itself, whereas we consider a binary relation symbol that is part of the meta-language. The second difference is that the semantics of the conditional implication of conditional logic is essentially different from ours. In conditional logic the formula $\alpha > \beta$ is interpreted to mean *if α were (or was) true and the situation were as close as possible, under this hypothesis, to what it really is, then β would be true.* For us $\alpha \mathrel{|\!\sim} \beta$ means that α is a good enough reason to believe β, or that β is a plausible consequence of α. The main difference is that conditional logic refers implicitly to the *actual* state of the world whereas we do not. M. Ginsberg's [14] proposal to harness conditional logic to nonmonotonic reasoning was clearly set with the former semantics in mind, and that explains our disagreements concerning the desirability of certain rules, e.g., the rule of **Rational Monotonicity** (see equation (25)).

One of the logical systems, **P**, studied in this paper turns out to be the flat (i.e. non-nested) fragment of a conditional logic studied by J. Burgess in [4] and by F. Veltman in [47]. Because of their richer language, the semantics proposed in those papers are more complex than ours: a ternary relation of accessibility between worlds is used in place of our binary preference relation. Moreover, the semantics of J. Burgess are quite different from ours in some other aspects; our semantics are closer to F. Veltman's (private communication from J. van Benthem) and to those studied by J. van Benthem in [46]. There are some connections between one of our completeness proofs and theirs, but the restricted language considered here simplifies the models and the proof a great deal. Our completeness result cannot be derived from the completeness result of [4] since the latter concerns a extended language and it is not clear that a proof in the extended language may be translated in the restricted one.

This very fragment had been considered by E. Adams in [2] (see also [1] for an earlier version and motivation). E. Adams' purpose was to propose probabilistic semantics for indicative conditionals and not the study of nonmonotonic logics. Recently J. Pearl and H. Geffner [34] have built upon E. Adams' logics, our system **P**, and his motivation in an effort to provide a system for nonmonotonic reasoning. For a gentle introduction, see chapter 10 of [33]. The semantics proposed here are not probabilistic. Probabilistic semantics that are equivalent with a restricted family of models (ranked models) will be described elsewhere. The preferential models presented in this paper provide a much sharper understanding of the system **P** than can obtained by Adams' methods.

1.4 Plan of this paper

This paper first describes the syntax proposed and compares it to more classical nonmonotonic systems. Five logical systems and families of models are then presented in turn and five soundness and completeness results are proven. The first system, **C**, corresponds to D. Gabbay's proposal. The second, stronger, system, **CL**, includes a rule of inference that seems original, and corresponds to models that seem to be more natural. None of those systems above assumes, in any essential way the existence of the classical logical connectives, if one allows a finite set of formulas to appear on the left of our symbol $\mid\!\sim$. The systems below assume the classical connectives. The third, stronger, system, **P**, is the central system of this paper. It has particularly appealing semantics. The fourth system, **CM**, is stronger than **CL** but incomparable with **P**. It provides an example of a monotonic system that is weaker than classical logic. The last one of those systems, **M**, is stronger than all previous systems and equivalent to classical propositional logic.

2 The language, comparison with other systems

2.1 Our language

The first step is to define a language in which to express the basic propositions. We shall assume that a set L of well formed formulas (thereafter formulas) is given. It is very important, from section 5 on, to assume that L is closed under the classical propositional connectives. They will be denoted by $\neg, \vee, \wedge, \rightarrow$ and $\leftrightarrow$. Negation and disjunction will be considered as the basic connectives and the other ones as defined connectives. The connective $\rightarrow$ therefore denotes material implication. Small greek letters will be used to denote formulas. Since no rule relating to the quantifiers will be discussed in this paper, the reader may as well think of L as the set of all propositional formulas on a given set of propositional variables.

With the language L, we assume semantics given by a set $\mathcal{U}$, the elements of which will be referred to as worlds, and a binary relation of *satisfaction* between worlds and formulas. The set $\mathcal{U}$ is the universe of reference, it is the set of all worlds that we shall consider possible. If L is the set of all propositional formulas on a given set of propositional variables, $\mathcal{U}$ is a subset of the set of all assignments of truth values to the propositional variables. We reserve to ourselves the right to consider universes of reference that are strict subsets of the set of all models of L. In this way we shall be able to model *strict* constraints, such as *penguins are birds*, in a simple and natural way, by restricting $\mathcal{U}$ to the set of all worlds that satisfy the material implication $penguin \rightarrow bird$. Typical universes of reference are given by the set of all propositional worlds that satisfy a given set of formulas.

We shall assume that the satisfaction relation behaves as expected as far as propositional connectives are concerned. If $u \in \mathcal{U}$ and $\alpha, \beta \in L$ we write $u \models \alpha$ if u satisfies α and assume:
1) $u \models \neg\alpha$ iff $u \not\models \alpha$.
2) $u \models \alpha \vee \beta$ iff $u \models \alpha$ or $u \models \beta$.

The notions of satisfaction of a set of formulas, validity of a formula and satisfiability of a set of formulas are defined as usual. We shall write $\models \alpha$ if α is valid, i.e. iff $\forall u \in \mathcal{U}$, $u \models \alpha$, and write $\alpha \models \beta$ for $\models \alpha \rightarrow \beta$.

We shall also make the following **assumption of compactness**[1]: a set of formulas is satisfiable if all of its finite subsets are.

Classical theorems of compactness show that if we take L to be a propositional calculus or a first order predicate calculus and $\mathcal{U}$ to be the set of all models that satisfy a given set of formulas, then the assumption of compactness described above is satisfied. Notice that the set of valid formulas, in our sense, is not, in general, closed under substitutions.

All that is done in the sequel depends on the choice of L and $\mathcal{U}$, though we shall often forget this dependence. For this work, the basic language L may well be fixed, but we shall sometimes have to consider different universes of reference. As noticed above, if Γ is a set of formulas then the subset of $\mathcal{U}$ that contains only the worlds that satisfy Γ (this set of worlds will be denoted by $\mathcal{U}_\Gamma$) is a suitable universe.

If α and β are formulas, then the pair $\alpha \mathrel{|\!\sim} \beta$ (read *if* α*, normally* β, or β *is a plausible consequence of* α) is called a conditional assertion (assertion in short). The formula α is the antecedent of the assertion, β is its consequent. The meaning we attach to such an assertion, and against which the reader should check the logical systems to be presented in the upcoming sections, is the following: if α is true, I am willing to (defeasibly) jump to the conclusion that β is true. Our choice, then, is to look at *normally* as some binary notion. It is clear that efforts to understand *normally* as some unary notion, e.g. translating *if* α*, normally* β as $\mathcal{N}(\alpha \rightarrow \beta)$ or as $\alpha \rightarrow \mathcal{N}\beta$ for some unary modal operator cannot be expressive enough. *Consequence relations* are sets of conditional assertions. Not all such sets, though, seem to be worthy of that name and our use of the term for any such set is running against a fairly well-established terminology. The term *conditional assertion* is taken from [37] (p. 417).

We hope that, by considering nonmonotonic consequence as a meta-notion, but allowing basic propositions on a rich language, we strike at the right language. It allows a new approach of questions about computational complexity (see [25] for some general decidability results), but this is left for future work.

2.2 Pragmatics

We shall now briefly sketch why we think that the study of nonmonotonic consequence relations will be a benefit to the field of automated nonmonotonic reasoning. The queries one wants to ask an automated knowledge base are formulas (of L) and query β should be interpreted as: *is* β *expected to be true*? To answer such a query the knowledge base will apply some inference procedure to the information it has. We shall now propose a description of the different types of information a knowledge base has.

The first type of information (first in the sense it is the more stable, changes less rapidly) is coded in the universe of reference $\mathcal{U}$, that describes both hard constraints (e.g. dogs are mammals) and points of definition (e.g. youngster is equivalent to not adult). Equivalently, such information will be given by a set of formulas defining $\mathcal{U}$ to be the set of all worlds that satisfy all the formulas of this set.

The second type of information consists of a set of conditional assertions describing the soft constraints (e.g. birds normally fly). This set describes what we know about the way the world generally behaves. This set of conditional assertions will be called the knowledge base, and denoted by $\mathbf{K}$.

The third type of information describes our information about the specific situation at hand (e.g. it is a bird). This information will be represented by a formula.

Our decision to consider the first type of information as a separate type is not the only possible way to

[1] The compactness assumption is needed only to treat consequence relations defined as the set of all assertions entailed by infinite sets of conditional assertions.

go. One could, equivalently, treat a formula α of the first type as the conditional assertion $\neg\,\alpha \mathrel{|\!\sim}$ **false**. One could also have decided to introduce all information of the third type as information of the first type.

Our inference procedure will work in the following way, to answer query β. In the context of the universe of reference $\mathcal{U}$, it will try to deduce (in a way that is to be discovered yet) the conditional assertion $\alpha \mathrel{|\!\sim} \beta$ from the knowledge base **K**. This is a particularly elegant way of looking at the inference process: the inference process deduces conditional assertions from sets of conditional assertions. Clearly any system of nonmonotonic reasoning may be considered in this way. So, we may look at circumscription, default logic and other systems as mechanisms to deduce conditional assertions from sets of conditional assertions. We shall now briefly investigate the expressive power of some of those systems in this light.

2.3 Expressiveness of our language

We shall now compare the expressive power of the language proposed here, i.e. conditional assertions, to that of previous approaches. Our purpose is to show that circumscription, autoepistemic logic and default logic all suffer from fundamental weaknesses, either in their expressive capabilities or in their treatment of conditional information. Let α, β and γ be formulas. We shall concentrate on the comparison of two different conditional assertions. The conditional assertion $\mathcal{A}$ is $\gamma \wedge \alpha \mathrel{|\!\sim} \beta$. The conditional assertion $\mathcal{B}$ is $\gamma \mathrel{|\!\sim} \neg\alpha \vee \beta$, i.e. $\gamma \mathrel{|\!\sim} \alpha \rightarrow \beta$. To simplify matters we shall just treat the special case when the formula γ is a tautology. In this case $\mathcal{A}$ is $\alpha \mathrel{|\!\sim} \beta$ and $\mathcal{B}$ is **true** $\mathrel{|\!\sim} \neg\alpha \vee \beta$.

The assertion $\mathcal{A}$ expresses that *if α, normally β*. Assertion $\mathcal{B}$ expresses that *Normally, if α is true then β is true*. Those assertions have very different meanings, at least when α is normally false. Assertion $\mathcal{A}$ says that in this exceptional case when α is true, one also expects β to be true. Assertion $\mathcal{B}$, on the other hand, is automatically verified if α is normally false. In any case it seems that it is perfectly possible that $\mathcal{B}$ does not say anything about cases when α is true (if these are exceptional). Take for example α to be *it is a penguin* and β to be *it flies*. If we talk about birds, it seems perfectly reasonable to accept $\mathcal{B}$ which says that *normally, either it is not a penguin or it flies*, since normally birds fly (and normally birds are not penguins, but this remark is not necessary). Nevertheless, one should be reluctant to accept $\mathcal{A}$ which says *penguins normally fly*. It seems clear to us, then, that $\mathcal{A}$ and $\mathcal{B}$ have different meanings and that $\mathcal{A}$ does not follow from $\mathcal{B}$. We agree with Y. Shoham, and this opinion will be supported in the sequel, to say that $\mathcal{B}$ should follow from $\mathcal{A}$, but we do not have to argue that case now. In the main system to be presented in this paper, **P**, the assertion $\mathcal{A}$ is strictly stronger than $\mathcal{B}$. In the weaker systems **C** and **CT**, $\mathcal{A}$ and $\mathcal{B}$ are incomparable. In **CM**, $\mathcal{B}$ is strictly stronger than $\mathcal{A}$, and this is one of the reasons we shall reject it as a system of nonmonotonic reasoning. It is only in **M**, which is equivalent to classical logic, that $\mathcal{A}$ and $\mathcal{B}$ are equivalent.

Let us consider now the expression of $\mathcal{A}$ and $\mathcal{B}$, first using circumscription. For circumscription, $\mathcal{A}$ would be expressed as: $\alpha \wedge \neg \mathit{abnormal} \rightarrow \beta$. In fact, since there would probably be a number of different abnormalities floating around, we probably should have written: $\alpha \wedge \neg \mathit{abnormal}_{543} \rightarrow \beta$, but this is not significant. On the other hand $\mathcal{B}$ would be written as: $\neg \mathit{abnormal} \rightarrow (\alpha \rightarrow \beta)$. One immediately notices that the two formulations are logically equivalent. We conclude that circumscription would need some additional mechanism to distinguish between $\mathcal{A}$ and $\mathcal{B}$. In practise, the user of circumscription would give different priorities (relative to the priorities of abnormalities of the other assertions of the knowledge base), to the two abnormalities considered here; but there is no standard procedure to determine priorities.

Let us now use autoepistemic logic. The assertion $\mathcal{A}$ would be expressed as: $\alpha \wedge \mathcal{M}(\beta) \rightarrow \beta$. On the other hand $\mathcal{B}$ would be expressed as: $\mathcal{M}(\alpha \rightarrow \beta) \rightarrow (\alpha \rightarrow \beta)$. Since the modality $\mathcal{M}$ is interpreted as $\neg\,\mathcal{K}\neg$ for some epistemic modality $\mathcal{K}$ it satisfies: $\alpha \wedge \mathcal{M}(\beta) \rightarrow \mathcal{M}(\alpha \rightarrow \beta)$. We immediately notice that, for autoepistemic logic, $\mathcal{B}$ is strictly stronger than $\mathcal{A}$. This is not what we expect.

Let us try default logic now. The natural translation of $\mathcal{A}$ in default logic would be the normal default:

$(\alpha, \beta; \beta)$, whose meaning is *if α has been concluded to be true and β is consistent with what has been concluded so far, conclude that β is true*. The assertion $\mathcal{B}$ would be expressed as: $(\mathbf{true}, \alpha \rightarrow \beta; \alpha \rightarrow \beta)$, which means that in any situation in which $\alpha \rightarrow \beta$ is consistent, one should (or could) conclude this last formula to be true. We immediately see that in all situations in which α has been already concluded to be true, both defaults act exactly in the same way, which seems very questionable. In situations in which α has been concluded to be false, the first default is inapplicable, whereas the second default may be applied but yields a trivial result (we do not get any new information from applying it). Again, both defaults are equivalent, but, in this case, this seems to fit our intuition. In situations in which neither α nor its negation have been concluded, the first default cannot be applied. For the second default, in certain situations it cannot be applied either, but in others it may be applied and yields non trivial conclusions. We conclude from this study that in some situations both defaults are equivalent, in others the second is more powerful than the first one. Again this is not what we expected. A particularly spectacular case of this problem occurs when β is a logical contradiction. The assertion $\alpha \mathrel{|\!\sim} \mathbf{false}$ has a very clear meaning: it says *if α, normally anything*. It expresses the very strong statement that we are willing to disconsider completely the possibility of α being true. To see that this may express very useful information, just think of α as *I am the Queen of England*. Most people would probably be willing to have $\alpha \mathrel{|\!\sim} \mathbf{false}$ in their personal knowledge base. As remarked above, this corresponds to restricting $\mathcal{U}$ to those worlds that do not satisfy α. Now, the translation, as a normal default, of such an assertion, which is: $(\alpha, \mathbf{false}; \mathbf{false})$, is never applicable since **false** is never consistent with anything. Therefore this default gives no information at all. Somehow, all the strength of our assertion has been lost in the translation.

We hope to have convinced the reader that one should look for formalisms in which the distinction between $\mathcal{A}$ and $\mathcal{B}$ is clear and understandable.

3 Cumulative reasoning

3.1 Cumulative consequence relations

We shall, first, study the weakest of our logical systems. It embodies what we think, at this moment, in agreement with D. Gabbay [10], are the rock-bottom properties without which a system should not be considered a logical system. This appreciation probably only reflects the fact that, so far, we do not know anything interesting about weaker systems. The order of the exposition, roughly from weaker to stronger systems, is aimed at minimizing repetitions: rules that may be derived in weaker systems may also be derived in stronger ones.

We shall name this system **C**, for *cumulative*. It is closely related to the cumulative inference studied by D. Makinson in [28], and seems to correspond exactly, to what D. Gabbay proposed in [10]. The system **C** consists of a number of inference rules and an axiom schema.

Definition 1 *A consequence relation $\mathrel{|\!\sim}$ is said to be* cumulative *iff it contains all instances of the* **Reflexivity** *axiom and is closed under the inference rules of* **Left Logical Equivalence**, **Right Weakening**, **Cut** *and* **Cautious Monotonicity** *that will be described below.*

We shall now describe and discuss the axioms and rules mentioned above and some derived rules. The purpose of the discussion is to weight the meaning of the axioms and rules when the relation $\ldots \mathrel{|\!\sim} \ldots$ is interpreted as *if ..., normally ...*.

(1) $\alpha \mathrel{|\!\sim} \alpha$ **(Reflexivity)**

Reflexivity seems to be satisfied universally by any kind of reasoning based on some notion of consequence. Relations that do not satisfy it, probably express some notion of theory change. It corresponds to the axiom ID of conditional logic.

The next two rules express the influence of the underlying logic, defined by the universe $\mathcal{U}$, on the notion of plausible consequence. Their role is similar to that of the rules of consequence of [16].

$$(2) \qquad \frac{\models \alpha \leftrightarrow \beta \;\;,\;\; \alpha \mathrel{|\!\sim} \gamma}{\beta \mathrel{|\!\sim} \gamma} \qquad \textbf{(Left Logical Equivalence)}$$

Left Logical Equivalence expresses the requirement that logically equivalent formulas have exactly the same consequences and corresponds to rule RCEA of conditional logic. The consequences of a formula should depend on its meaning, not on its form. In the presence of the other rules of **C**, it could have been weakened to: from $\alpha \wedge \beta \mathrel{|\!\sim} \gamma$ conclude $\beta \wedge \alpha \mathrel{|\!\sim} \gamma$.

The next rule, **Right Weakening** expresses the fact that one must be ready to accept as plausible consequences all that is logically implied by what one thinks are plausible consequences. In other words, plausible consequences are closed under logical consequences. It corresponds to the rule RCK of conditional logic.

$$(3) \qquad \frac{\models \alpha \rightarrow \beta \;\;,\;\; \gamma \mathrel{|\!\sim} \alpha}{\gamma \mathrel{|\!\sim} \beta} \qquad \textbf{(Right Weakening)}$$

Right Weakening obviously implies that one may replace logically equivalent formulas by one another on the right of the $\mathrel{|\!\sim}$ symbol. **Reflexivity** and **Right Weakening** already imply that $\alpha \mathrel{|\!\sim} \beta$ if $\alpha \models \beta$. All nonmonotonic systems proposed so far in the literature satisfy **Reflexivity**, **Left Logical Equivalence** and **Right Weakening**.

Our next rule is named **Cut** because of its similarity to Gentzen's *Schnitt*.

$$(4) \qquad \frac{\alpha \wedge \beta \mathrel{|\!\sim} \gamma \;\;,\;\; \alpha \mathrel{|\!\sim} \beta}{\alpha \mathrel{|\!\sim} \gamma} \qquad \textbf{(Cut)}$$

It expresses the fact that one may, in his way towards a plausible conclusion, first add an hypothesis to the facts he knows to be true and prove the plausibility of his conclusion from this enlarged set of facts and then deduce (plausibly) this added hypothesis from the facts. This is a valid way of reasoning in monotonic logic, and, as will be seen soon, its validity does not imply monotonicity, therefore it seems to us quite reasonable to accept it. Its meaning, it should be stressed, is that a plausible conclusion is as secure as the assumptions it is based on. Therefore it may be added (this is the origin of the term cumulative) into the assumptions. There is no loss of confidence along the chain of derivations. One may well be unwilling to accept such a principle and think that, on the contrary, no conclusion of a derivation is ever as secure as the assumptions. Indeed, recently, D. Gabbay [11] suggested to replace **Cut** by a weaker rule. In this paper, we shall study only systems that validate **Cut**. Our conclusion is that there are many interesting nonmonotonic systems that satisfy **Cut**. It should be mentioned that some probabilistic interpretations invalidate **Cut** (Adams' validates it), e.g. interpreting a conditional assertion $\alpha \mathrel{|\!\sim} \beta$ as meaning that the corresponding conditional probability $p(\beta \mid \alpha)$ is larger than some $q < 1$.

It is easy to see that circumscription satisfies **Cut**, at least when all models that have to be considered are finite. In [28], D. Makinson shows that default logic satisfies **Cut** too. The following example should help convince the reader to endorse **Cut**. Suppose we tell you *we expect it will be raining tonight* and *if it rains tonight, normally Fireball should win the race tomorrow*. Wouldn't you conclude that we think that *normally, Fireball should win the race tomorrow*?

Our last rule, named **Cautious Monotonicity**, is taken from D. Gabbay [10]. It corresponds to axiom A3 of Burgess' system $\mathcal{S}$ in [4]. The same property is named *triangulation* in [34].

$$(5) \quad \frac{\alpha \mathrel{|\!\sim} \beta \ , \ \alpha \mathrel{|\!\sim} \gamma}{\alpha \wedge \beta \mathrel{|\!\sim} \gamma} \qquad \textbf{(Cautious Monotonicity)}$$

Cautious Monotonicity expresses the fact that learning a new fact, the truth of which could have been plausibly concluded should not invalidate previous conclusions. It is a central property of all the systems considered here. The origin of the term *cautious monotonicity* will be explained in section 3.3. The probabilistic semantics that invalidates **Cut** also invalidates **Cautious Monotonicity**. In [28], D. Makinson showed that default logic, even when defaults are normal, does not always satisfy **Cautious Monotonicity**. Circumscription, though, satisfies it, at least when all models considered are finite. What are our reasons to accept **Cautious Monotonicity**? On the general level, D. Gabbay's argumentation seems convincing: if α is reason enough to believe β and also to believe γ, then α and β should also be enough to make us believe γ, since α was enough anyway and, on this basis, β was expected. From a pragmatic point of view **Cautious Monotonicity** is very important since we typically learn new facts and we would like to minimize the updating we have to make to our beliefs. **Cautious Monotonicity** and **Cut** together tell us, as will be made clear in lemma 1, that if the new facts learned were expected to be true, nothing changes in our beliefs. This will help minimizing the updating. From a semantic point of view, we want to argue the case for **Cautious Monotonicity** on the following example. Suppose we tell you *we expect it will be raining tonight* and *normally, Fireball should win the race tomorrow*. Wouldn't you conclude that we think that *even if it rains tonight, normally Fireball should win the race tomorrow*?

Lemma 1 *The rules of* **Cut** *and* **Cautious Monotonicity** *may be expressed together by the following principle: if* $\alpha \mathrel{|\!\sim} \beta$ *then the plausible consequences of* α *and of* $\alpha \wedge \beta$ *coincide.*

Let us now consider some rules that may be derived in **C**.

3.2 Derived rules of C

The first rule corresponds to CSO of conditional logic and expresses the fact that two propositions that are plausible consequences of each other, have exactly the same plausible consequences.

$$(6) \quad \frac{\alpha \mathrel{|\!\sim} \beta \ , \ \beta \mathrel{|\!\sim} \alpha \ , \ \alpha \mathrel{|\!\sim} \gamma}{\beta \mathrel{|\!\sim} \gamma} \qquad \textbf{(Equivalence)}$$

The second rule corresponds to CC of conditional logic and expresses the fact that the conjunction of two plausible consequences is a plausible consequence.

$$(7) \quad \frac{\alpha \mathrel{|\!\sim} \beta \ , \ \alpha \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \beta \wedge \gamma} \qquad \textbf{(And)}$$

The third rule amounts to *modus ponens* in the consequent.

$$(8) \quad \frac{\alpha \mathrel{|\!\sim} \beta \rightarrow \gamma \ , \ \alpha \mathrel{|\!\sim} \beta}{\alpha \mathrel{|\!\sim} \gamma} \qquad \textbf{(MPC)}$$

The fourth rule is perhaps less expected and brought up here to show that **C** is not as weak as one could think. It will be put to use in section 5.3.

$$(9) \qquad \frac{\alpha \vee \beta \mathrel{|\!\sim} \alpha \quad , \quad \alpha \mathrel{|\!\sim} \gamma}{\alpha \vee \beta \mathrel{|\!\sim} \gamma}$$

Lemma 2 **Equivalence**, **And**, **MPC** *and (9) are derived rules of the system* **C**.

Proof: For **Equivalence**, use first **Cautious Monotonicity** to show that $\alpha \wedge \beta \mathrel{|\!\sim} \gamma$, then **Left Logical Equivalence** to get $\beta \wedge \alpha \mathrel{|\!\sim} \gamma$ and then conclude by **Cut**.

For **And**, first use **Cautious Monotonicity** to show $\alpha \wedge \beta \mathrel{|\!\sim} \gamma$. Then, since $\alpha \wedge \beta \wedge \gamma \models \beta \wedge \gamma$, we have: $\alpha \wedge \beta \wedge \gamma \mathrel{|\!\sim} \beta \wedge \gamma$. Using **Cut** we conclude that $\alpha \wedge \beta \mathrel{|\!\sim} \beta \wedge \gamma$ and the desired conclusion is obtained by one more use of **Cut**.

For **MPC**, use **And** and **Right Weakening**.

For (9), remark that, since $\alpha \models \alpha \vee \beta$, we have $\alpha \mathrel{|\!\sim} \alpha \vee \beta$. This, with the hypotheses, enables us to apply **Equivalence** and conclude. ∎

3.3 Monotonicity

We shall now justify the term **Cautious Monotonicity** and introduce four new rules. They cannot be derived in **C**. The first rule is **Monotonicity**, or **Left Strengthening**.

$$(10) \qquad \frac{\models \alpha \rightarrow \beta \quad , \quad \beta \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \gamma} \qquad \textbf{(Monotonicity)}$$

It is clear that both **Left Logical Equivalence** and **Cautious Monotonicity** are special cases of **Monotonicity**. This explains the name **Cautious Monotonicity**.

The next rule corresponds to the *easy* half of the deduction theorem.

$$(11) \qquad \frac{\alpha \mathrel{|\!\sim} \beta \rightarrow \gamma}{\alpha \wedge \beta \mathrel{|\!\sim} \gamma} \qquad \textbf{(EHD)}$$

The next two rules have been considered by many.

$$(12) \qquad \frac{\alpha \mathrel{|\!\sim} \beta \quad , \quad \beta \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \gamma} \qquad \textbf{(Transitivity)}$$

$$(13) \qquad \frac{\alpha \mathrel{|\!\sim} \beta}{\neg\beta \mathrel{|\!\sim} \neg\alpha} \qquad \textbf{(Contraposition)}$$

It is easy to find apparent counter-examples to each one of the last four rules in the folklore of nonmonotonic reasoning. The next lemmas will explain why. Let us notice that, nevertheless, adding the first three of those rules to the system **C** leaves us with a system, **CM**, that is strictly weaker than classical monotonic logic, as will be seen in section 6. The next lemmas will describe some of the relations between the rules above.

Lemma 3 *In the presence of the rules of* **C**, *the rules of* **Monotonicity**, **EHD**, *and* **Transitivity** *are all equivalent.*

Proof: We shall not mention the uses of **Reflexivity**, **Left Logical Equivalence** and **Right Weakening**. **Monotonicity** implies **EHD**, using **And**. **EHD** implies **Monotonicity**. **Monotonicity** implies **Transitivity**, using **Cut**. **Transitivity** implies **Monotonicity**. ▮

Lemma 4 *In the presence of* **Left Logical Equivalence** *and* **Right Weakening**, **Contraposition** *implies* **Monotonicity**.

Proof: Use **Contraposition**, then **Right Weakening** and **Contraposition** again. ▮

The results of section 6 show that **Monotonicity** does not imply **Contraposition** even in the presence of the rules of **C**. Since **Monotonicity** seems counter-intuitive in nonmonotonic systems, the two lemmas above show we should not accept **EHD**, **Transitivity** or **Contraposition**.

3.4 Cumulative models

We shall now develop a semantic account of cumulative reasoning, i.e. reasoning using the rules of the system **C**. We shall define a family of models (without any reference to the rules of **C**) and show how each model defines a consequence relation. We shall then show that each model of the family defines a cumulative consequence relation (this is a soundness result) and that every cumulative consequence relation is defined by some model of the family (this is a completeness result, or a representation theorem).

Let us, first, describe the models informally. A model essentially consists of a set of states (they represent possible states of affairs, including perhaps the state of mind or knowledge of the reasoner) and a binary relation on those states. The relation represents the preferences the reasoner may have between different states: he could for example prefer the states in which he is rich to the ones in which he is poor, and prefer the states in which he knows he is rich to those in which he is rich but does not know about it. More realistically, one could prefer states in which Tweety is a bird and flies to those in which Tweety is a bird but does not fly. The reasoner, described by a model, accepts a conditional assertion $\alpha \mathrel{|\!\sim} \beta$ iff all those states that are *most preferred* among all states satisfying α, satisfy β. The reader should notice we have not yet said what is a state and what formulas are satisfied by a state.

We shall not define further the notion of a state, but suppose that every state is, in a model, labeled with a set of worlds (intuitively the set of all worlds the reasoner thinks are possible in this state). Modal logicians will identify our labels as S5 models. Considering a binary relation on states labeled by sets of worlds, instead of considering a binary relation on sets of worlds, gives us an additional degree of freedom in building models: the same set of worlds may appear at many states (that are not equivalent from the point of view of the binary relation). This additional freedom is vital for the representation theorem to hold, and was missing from Shoham's account [38].

Some technical definitions are needed first.

Definition 2 *Let $\prec$ be a binary relation on a set U. We shall say that $\prec$ is* asymmetric *iff $\forall s, t \in U$ such that $s \prec t$, we have $t \not\prec s$.*

Definition 3 *Let $V \subseteq U$ and $\prec$ a binary relation on U. We shall say that $t \in V$ is* minimal *in V iff $\forall s \in V$, $s \not\prec t$. We shall say that $t \in V$ is a* minimum *of V iff $\forall s \in V$ such that $s \neq t$, $t \prec s$.*

The reader has noticed that, though the last definitions sound familiar in the case the relation $\prec$ is a strict partial order, we intend to use them for arbitrary relations.

Definition 4 *Let $P \subseteq U$ and $\prec$ a binary relation on U. We shall say that P is* smooth *iff $\forall t \in P$, either $\exists s$ minimal in P, such that $s \prec t$ or t is itself minimal in P.*

We shall use the following lemma, the proof of which is obvious.

Lemma 5 *Let U be a set and $\prec$ an asymmetric binary relation on U. If U has a minimum it is unique, it is a minimal element of U and U is smooth.*

Definition 5 *A* cumulative *model is a triple $\langle S, l, \prec \rangle$, where S is a set, the elements of which are called states, $l : S \mapsto 2^{\mathcal{U}}$ is a function that labels every state with a non-empty set of worlds and $\prec$ is a binary relation on S, satisfying the* **smoothness condition** *that will be defined below in definition 7.*

The relation $\prec$ represents the reasoner's preference among states. The fact that $s \prec t$ means that, in the agent's mind, s is *preferred* to or more *natural* than t. As will be formally defined below, the agent is willing to conclude β from α, if all *most natural* states which satisfy α also satisfy β.

Definition 6 *Let $\langle S, l, \prec \rangle$ be as above. If α is a formula, we shall say that $s \in S$ satisfies α and write $s \models\!\equiv \alpha$ iff for every world $m \in l(s)$, $m \models \alpha$. The set: $\{s \mid s \in S,\ s \models\!\equiv \alpha\}$ of all states that satisfy α will be denoted by $\widehat{\alpha}$.*

Definition 7 (smoothness condition) *A triple $\langle S, l, \prec \rangle$ is said to satisfy the smoothness condition iff, $\forall \alpha \in L$, the set $\widehat{\alpha}$ is smooth.*

The smoothness condition is necessary to ensure the validity of **Cautious Monotonicity**. It is akin to the *limit assumption* of Stalnaker [40] and Lewis [24], but it is defined in a more general context. Smoothness is the property called, contrary to mathematical usage, *well-foundedness* in [8] and in [26].

We shall now describe how a cumulative model defines a consequence relation.

Definition 8 *Suppose a cumulative model $W = \langle S, l, \prec \rangle$ is given. The consequence relation defined by W will be denoted by $\mid\!\sim_W$ and is defined by: $\alpha \mid\!\sim_W \beta$ iff for any s minimal in $\widehat{\alpha}$, $s \models\!\equiv \beta$.*

Definition 9 *A triple $\langle S, l, \prec \rangle$ is said to be a* strong *cumulative model iff*

1. *the relation $\prec$ is asymmetric*
2. *for each formula α, the set $\widehat{\alpha}$ has a minimum.*

It is clear that strong cumulative models are cumulative models, i.e. satisfy the smoothness condition. The definition of cumulative models and the consequence relations they define seems quite natural, i.e. a preference relation on epistemic states, but one should not forget that the preference relation $\prec$ is not required to be a partial order and that in triples (even when the set of states is finite) in which the relation $\prec$ is not a partial order, the smoothness condition is, in general, not an easy thing to check.

3.5 Characterization of cumulative relations

In this section we shall characterize the relation between cumulative consequence relations and cumulative models. The first lemma is obvious.

Lemma 6 *Let $W = \langle S, l, \prec \rangle$ be a cumulative model. For $\alpha, \beta \in L$, $\widehat{\alpha \wedge \beta} = \widehat{\alpha} \cap \widehat{\beta}$.*

Lemma 7 (Soundness) *For any cumulative model W, the consequence relation $\mid\!\sim_W$ it defines is a cumulative relation, i.e. all the rules of the system* **C** *are satisfied by the relations defined by cumulative models.*

Proof: The proof is easy and we shall only treat **Cut** and **Cautious Monotonicity**. The smoothness condition is needed only for dealing with **Cautious Monotonicity**.

For **Cut**, suppose all minimal elements of $\widehat{\alpha}$ satisfy β and all minimal elements of $\widehat{\alpha \wedge \beta}$ satisfy γ. Any minimal element of $\widehat{\alpha}$ satisfies β and therefore satisfies $\alpha \wedge \beta$. Since it is minimal in $\widehat{\alpha}$ and $\widehat{\alpha \wedge \beta} \subseteq \widehat{\alpha}$, it is also minimal in $\widehat{\alpha \wedge \beta}$.

For **Cautious Monotonicity**, the smoothness condition is needed. Suppose that $\alpha \mid\!\sim_W \beta$ and $\alpha \mid\!\sim_W \gamma$. We have to prove that $\alpha \wedge \beta \mid\!\sim_W \gamma$, i.e., that, for any s minimal in $\widehat{\alpha \wedge \beta}$, $s \models\!\!\!\equiv \gamma$. Such an s is in $\widehat{\alpha}$. We shall prove that it is minimal in $\widehat{\alpha}$. By the smoothness condition, if it were not minimal, there would be an s' minimal in $\widehat{\alpha}$ such that $s' \prec s$. But $\alpha \mid\!\sim_W \beta$ therefore $s' \models\!\!\!\equiv \beta$ and then $s' \in \widehat{\alpha} \cap \widehat{\beta}$. By lemma 6 we conclude that s' is in $\widehat{\alpha \wedge \beta}$, in contradiction with the minimality of s in this set. Therefore s is minimal in $\widehat{\alpha}$ and, since $\alpha \mid\!\sim_W \gamma$, one concludes: $s \models\!\!\!\equiv \gamma$. ∎

We now intend to show that, given any cumulative relation $\mid\!\sim$, one may build a cumulative model W that defines a consequence relation $\mid\!\sim_W$ that is exactly $\mid\!\sim$. Suppose, therefore, that $\mid\!\sim$ satisfies the rules of **C**. All definitions will be relative to this relation.

Definition 10 *The world $m \in \mathcal{U}$ is said to be a* normal *world for α iff $\forall \beta \in L$ such that $\alpha \mid\!\sim \beta$, $m \models \beta$.*

So, a world is normal for a formula if it satisfies all of its plausible consequences. Obviously, if the consequence relation we start from satisfies **Reflexivity**, a normal world for α satisfies α.

Lemma 8 *Suppose a consequence relation $\mid\!\sim$ satisfies* **Reflexivity**, **Right Weakening** *and* **And**, *and let $\alpha, \beta \in L$. All normal worlds for α satisfy β iff $\alpha \mid\!\sim \beta$.*

Proof: The *if* part follows from definition 10. Let us show the *only if* part. Suppose $\alpha \not\mid\!\sim \beta$, we shall build a normal world for α that does not satisfy β. Let $\Gamma_0 \stackrel{\text{def}}{=} \{\neg\beta\} \cup \{\delta \mid \alpha \mid\!\sim \delta\}$. It is enough to show that Γ_0 is satisfiable. Suppose not, then, by the compactness assumption, there exists a finite subset of Γ_0 that is not satisfiable and therefore a finite set $D \subseteq \{\delta \mid \alpha \mid\!\sim \delta\}$ such that $\models \bigwedge_{\delta \in D} \delta \rightarrow \beta$. Now, $\models \alpha \rightarrow \left(\bigwedge_{\delta \in D} \delta \rightarrow \beta\right)$ and, by **Reflexivity** and **Right Weakening** $\alpha \mid\!\sim \left(\bigwedge_{\delta \in D} \delta \rightarrow \beta\right)$. But, using **And** one gets $\alpha \mid\!\sim \bigwedge_{\delta \in D} \delta$. Then, using **MPC** (the proof of lemma 2 shows that only **And** and **Right Weakening** are needed to derive **MPC**), one concludes $\alpha \mid\!\sim \beta$, a contradiction. ∎

Definition 11 *We shall say that α is* equivalent *to β and write $\alpha \sim \beta$ iff $\alpha \mid\!\sim \beta$ and $\beta \mid\!\sim \alpha$.*

Lemma 9 *$\alpha \sim \beta$ iff $\forall \gamma \quad \alpha \mid\!\sim \gamma \Leftrightarrow \beta \mid\!\sim \gamma$. The relation $\sim$ is therefore an equivalence relation.*

Proof: The *if* part follows from **Reflexivity** and the *only if* part from **Equivalence**. ∎

The equivalence class of a formula α, under $\sim$, will be denoted by $\bar{\alpha}$.

Definition 12 *$\bar{\alpha} \leq \bar{\beta}$ iff $\exists \alpha' \in \bar{\alpha}$ such that $\beta \mathrel{|\!\sim} \alpha'$.*

It is clear that the definition of $\leq$ makes sense, i.e. does not depend on the choice of the representatives α and β. The relation $\leq$ is reflexive but is not in general transitive.

Lemma 10 *The relation $\leq$ is antisymmetric.*

Proof: Suppose $\bar{\alpha} \leq \bar{\beta}$ and $\bar{\beta} \leq \bar{\alpha}$. There are formulas $\alpha', \alpha'' \in \bar{\alpha}$ and $\beta', \beta'' \in \bar{\beta}$ such that: $\beta' \mathrel{|\!\sim} \alpha''$ and $\alpha' \mathrel{|\!\sim} \beta''$. By lemma 9, $\beta'' \mathrel{|\!\sim} \alpha''$ and $\alpha'' \mathrel{|\!\sim} \beta''$. Therefore $\alpha'' \sim \beta''$, and $\bar{\alpha} = \bar{\beta}$. ∎

The cumulative model W will be defined the following way: $W \stackrel{\mathrm{def}}{=} \langle S, l, \prec \rangle$, where $S \stackrel{\mathrm{def}}{=} L/\sim$ is the set of all equivalence classes of formulas under the relation $\sim$, $l(\bar{\alpha}) \stackrel{\mathrm{def}}{=} \{m \mid m \text{ is a normal world for } \alpha\}$ and $\bar{\alpha} \prec \bar{\beta}$ iff $\bar{\alpha} \leq \bar{\beta}$ and $\bar{\alpha} \neq \bar{\beta}$ (the relation $\leq$ has been defined in definition 12). One easily checks the definition of l does not depend on the choice of the representative α and that $\prec$ is asymmetric.

Lemma 11 *For any $\alpha \in L$ the state $\bar{\alpha}$ is a minimum of $\widehat{\alpha}$.*

Proof: Indeed suppose $s \neq \bar{\alpha}$ and $s \in \widehat{\alpha}$. This last assumption implies, by the definition of $\widehat{\alpha}$, that every world of $l(s)$ satisfies α. Let $s = \bar{\beta}$. By the definition of l, every normal world for β satisfies α. By lemma 8, $\beta \mathrel{|\!\sim} \alpha$, and therefore $\bar{\alpha} \leq s$. Since $s \neq \bar{\alpha}$ we conclude $\bar{\alpha} \prec s$. ∎

It follows from lemma 11 and the fact that $\prec$ is asymmetric that the model W defined above is a strong cumulative model. We may now prove what we wanted to achieve.

Lemma 12 *$\alpha \mathrel{|\!\sim} \beta$ iff $\alpha \mathrel{|\!\sim}_W \beta$.*

Proof: Lemmas 11 and 5 imply that the only minimal state of $\widehat{\alpha}$ is $\bar{\alpha}$, therefore $\alpha \mathrel{|\!\sim}_W \beta$ iff all normal worlds for α satisfy β and lemma 8 implies the conclusion. ∎

Theorem 1 (Representation theorem for cumulative relations) *A consequence relation is a cumulative consequence relation iff it is defined by some cumulative model.*

Proof: The *if* part is lemma 7. The *only if* part follows from the construction of W and lemma 11 (that shows W is a cumulative model) and lemma 12. ∎

One may remark that the representation result proved is a bit stronger than what is claimed in theorem 1: any cumulative consequence relation is the consequence relation defined by a strong cumulative model. It is now easy to study the notion of entailment yielded by cumulative models.

Corollary 1 *Let **K** be a set of conditional assertions, and $\alpha, \beta \in L$, the following conditions are equivalent. In case they hold we shall say that **K** cumulatively entails $\alpha \mathrel{|\!\sim} \beta$.*

1. *for all cumulative models V such that $\mathrel{|\!\sim}_V$ contains **K**, $\alpha \mathrel{|\!\sim}_V \beta$*
2. *$\alpha \mathrel{|\!\sim} \beta$ has a proof from **K** in the system **C**.*

Proof: From lemma 7 one sees that 2) implies 1). For the other direction, suppose 2) is not true. The smallest consequence relation closed under the rules of **C** that contains **K** is a cumulative consequence relation that does not contain $\alpha \mid\sim \beta$. By theorem 1, there is a cumulative model that defines it. This model shows property 1) does not hold. ▮

Corollary 2 *Let* **K** *be a set of conditional assertions. There is a cumulative model that satisfies exactly those assertions that are cumulatively entailed by* **K**.

The following compactness result follows.

Corollary 3 (compactness) **K** *entails* $\alpha \mid\sim \beta$ *iff a finite subset of* **K** *does.*

Proof: Proofs are always finite and therefore use only a finite number of assumptions from **K**. ▮

To conclude our study of cumulative reasoning, let us say that the system **C** provides an interesting general setting in which to study nonmonotonic reasoning. Weaker systems are probably very different from systems that are at least as strong as **C**. The system **C** is probably too weak to be the backbone of realistic inference systems and cumulative models are quite cumbersome to manipulate. The next section will propose nicer models and an additional rule of inference.

4 Cumulative reasoning with Loop

4.1 Cumulative ordered models

The original motivation for the study of the system **CL**, that will be proposed in this section, stems from semantic considerations. Later on, a number of results, including the result that will be described in section 5.6, which says that, if one restricts oneself to Horn assertions, then the system **CL** is as strong as **P**, seemed to point out that **CL** is worth studying.

Looking back on the cumulative models of definition 5, one may wonder why we did not require the binary relation $\prec$ to be a strict partial order. We could have required it to be asymmetric without jeopardizing the representation theorem, but the construction of section 3.5 builds a model in which $\prec$ is not always transitive. Nevertheless, preferences could probably be assumed to be transitive and, most important, transitivity of $\prec$ eases enormously the task of checking the smoothness condition: if $\prec$ is a partial order (strict), then all finite models (models in which the set of states is finite) satisfy the smoothness condition, and even all well-founded models (in which there is no infinite descending $\prec$-chain) do. Could we have required $\prec$ to be a partial order? In other terms, are there rules that are not valid for cumulative models in general but are valid for all cumulative models the preference relation of which is a strict partial order? We shall now give a positive answer to this last question and exactly characterize this sub-family of cumulative models.

Definition 13 *A cumulative* ordered *model is a cumulative model in which the relation* $\prec$ *is a strict partial order.*

4.2 The system CL

The following rule, named **Loop** after its form, will be shown to be the exact counterpart of transitivity of the preference relation in the models. It says that, if propositions may be arranged in a loop, in a way each one is a plausible consequence of the previous one, then each one of them is a plausible consequence of any one of them, i.e. they are all equivalent in the sense of **Equivalence**.

Definition 14 *The system* **CL** *consists of all the rules of* **C** *and the following:*

$$(14)\quad \frac{\alpha_0 \mathrel{|\!\sim} \alpha_1 \;,\; \alpha_1 \mathrel{|\!\sim} \alpha_2 \;,\ldots,\; \alpha_{k-1} \mathrel{|\!\sim} \alpha_k \;,\; \alpha_k \mathrel{|\!\sim} \alpha_0}{\alpha_0 \mathrel{|\!\sim} \alpha_k} \qquad (\textbf{Loop})$$

A consequence relation that satisfies all rules of **CL** *is said to be* loop-cumulative*.*

Lemma 13 *The following is a derived rule of* **CL***, for any* $i, j = 0, \ldots k$*.*

$$(15)\quad \frac{\alpha_0 \mathrel{|\!\sim} \alpha_1 \;,\; \alpha_1 \mathrel{|\!\sim} \alpha_2 \;,\ldots,\; \alpha_{k-1} \mathrel{|\!\sim} \alpha_k \;,\; \alpha_k \mathrel{|\!\sim} \alpha_0}{\alpha_i \mathrel{|\!\sim} \alpha_j}$$

Proof: It is clear that, because of the invariance of the assumptions under cyclic permutation, the conclusion of **Loop**, could as well have been $\alpha_{i+1} \mathrel{|\!\sim} \alpha_i$, for any $i = 0, \ldots, k$ (addition is understood modulo $k+1$). From **Equivalence** one can then conclude $\alpha_i \mathrel{|\!\sim} \alpha_j$, for any $i, j = 0, \ldots, k$. ∎

It seems the rule **Loop** has never been considered in the literature. We feel it is an acceptable principle of nonmonotonic reasoning. It is particularly interesting that **Loop** does not mention any of the propositional connectives.

Lemma 14 **Loop** *is valid in all cumulative ordered models.*

Proof: Let $W = \langle S, l, \prec \rangle$ be a cumulative ordered model such that $\alpha_i \mathrel{|\!\sim}_W \alpha_{i+1}$ for $i = 0, \ldots, k$ (addition is understood modulo $k+1$) and let $s_0 \in S$ be a minimal state in $\widehat{\alpha_0}$. We shall show that $s_0 \models\!\!\!\equiv \alpha_k$. Since $\alpha_0 \mathrel{|\!\sim}_W \alpha_1$, the state s_0 must be in $\widehat{\alpha_1}$. By the smoothness condition, if s_0 is not minimal in $\widehat{\alpha_1}$ then there is a state s_1 minimal in $\widehat{\alpha_1}$ such that $s_1 \prec s_0$. Similarly, for every $i = 0, \ldots, k$ there is a state s_i minimal in $\widehat{\alpha_i}$ such that $s_i = s_{i-1}$ or $s_i \prec s_{i-1}$. Since $\prec$ is transitive, $s_k = s_0$ or $s_k \prec s_0$. But s_k is minimal in $\widehat{\alpha_k}$ and $\alpha_k \mathrel{|\!\sim}_W \alpha_0$, we conclude that $s_k \in \widehat{\alpha_0}$. But s_0 is minimal in $\widehat{\alpha_0}$, we conclude that $s_k = s_0$ and $s_0 \models\!\!\!\equiv \alpha_k$. ∎

Lemma 15 *The rule* **Loop** *is not valid in cumulative models.*

Proof: Let L be the propositional calculus on the propositional variables p_0, p_1, p_2 and $\mathcal{U}$ be the set of all propositional models on those variables. We shall build a cumulative model $V = \langle S, l, \prec \rangle$ such that $p_i \mathrel{|\!\sim}_V p_{i+1}$ for all $i = 0, \ldots, 2$ (addition is modulo 3) but $p_0 \not\mathrel{|\!\sim}_V p_2$. The set S has four states: s_i, for $i = -1, \ldots, 2$. For every $i = 0, \ldots, 2$ we have $s_{-1} \prec s_i$ and $s_{i+1} \prec s_i$. Notice that $\prec$ is not transitive. Let us now describe l. For $i = 0, \ldots, 2$, $l(s_i)$ is the set of all worlds satisfying p_i and p_{i+1}, and $l(s_{-1})$ is the set of all worlds satisfying at least two out of the three variables. First we want to show that V satisfies the smoothness condition. Clearly all subsets of S that contain s_{-1} are smooth since s_{-1} is a minimum in S. A set that contains at most two elements is always smooth. We conclude that the only subset of S that is not smooth is $A \stackrel{\text{def}}{=} \{s_0, s_1, s_2\}$. We must show that there is no formula α such that $A = \widehat{\alpha}$. Let α be any formula and let $i = 0, \ldots, 2$. If $s_i \in \widehat{\alpha}$ all worlds of $l(s_i)$ must satisfy α and by definition of l, $p_i \wedge p_{i+1} \models \alpha$. We conclude that if $A \subseteq \widehat{\alpha}$ then any world that satisfies at least two of the variables satisfies α. We conclude that $\widehat{\alpha}$ must therefore also include s_{-1}.

To see that $p_i \mathrel{|\!\sim}_V p_{i+1}$, notice that $\widehat{p_i} = \{s_{i-1}, s_i\}$ and that, since $s_i \prec s_{i-1}$, the only minimal state in $\widehat{p_i}$ is s_i that satisfies p_{i+1}. The only thing left to check is that $p_0 \not\mathrel{|\!\sim}_V p_2$. But we just noticed that the only minimal state of $\widehat{p_0}$ is s_0 and clearly $s_0 \not\models\!\!\!\equiv p_2$. ∎

4.3 Characterization of loop-cumulative consequence relations

We now want to show that, given any loop-cumulative relation $\mid\!\sim$ one may build a cumulative ordered model V such that $\mid\!\sim_V$ is equal to $\mid\!\sim$. Suppose $\mid\!\sim$ is such a relation and $W = \langle S, l, \prec \rangle$ is the cumulative model built out of $\mid\!\sim$ in section 3.5. Let $\prec^+$ be the transitive closure of $\prec$. First we shall show that, since $\mid\!\sim$ satisfies **Loop**, the relation $\prec^+$ is a strict partial order.

Lemma 16 *The relation $\prec^+$ is irreflexive and therefore a strict partial order.*

Proof: Suppose $\bar{\alpha}_0 \prec^+ \bar{\alpha}_0$. Since $\prec$ is asymmetric, it is irreflexive and t here must be some $n > 0$ such that for $i = 0, \ldots, n$, $\bar{\alpha}_i \prec \bar{\alpha}_{i+1}$ (addition is modulo n). From the definitions of $\prec$ and $\leq$, we see that, for $i = 0, \ldots, n$, there are formulas α'_i such that $\alpha_i \sim \alpha'_i$ and $\alpha_{i+1} \mid\!\sim \alpha'_i$. From lemma 9, we conclude that $\alpha'_{i+1} \mid\!\sim \alpha'_i$ for $i = 0, \ldots, n$. By **Loop** we see that $\alpha'_i \mid\!\sim \alpha'_{i+1}$ and therefore $\alpha'_{i+1} \sim \alpha'_i$ and $\alpha_i \sim \alpha_{i+1}$. But this contradicts the asymmetry of $\prec$. We have shown that $\prec$ is irreflexive. Since it is transitive by construction it is a strict partial order. ∎

Let us now define $V \stackrel{\text{def}}{=} \langle S, l, \prec^+ \rangle$ where S, l and $\prec$ are as in the definition of W.

Lemma 17 *In V, for any α, the state $\bar{\alpha}$ is a minimum of $\widehat{\alpha}$. Therefore V is a strong cumulative ordered model.*

Proof: Lemma 11 says $\bar{\alpha}$ is a minimum of $\widehat{\alpha}$ with respect to $\prec$. It is therefore a minimum with respect to any weaker relation and in particular $\prec^+$. Lemma 16 implies that $\prec^+$ is asymmetric and, by lemma 5, V satisfies the smoothness condition. ∎

Lemma 18 *$\alpha \mid\!\sim \beta$ iff $\alpha \mid\!\sim_V \beta$.*

Proof: Lemma 17 implies that the only minimal state of $\widehat{\alpha}$ is $\bar{\alpha}$, therefore $\alpha \mid\!\sim_V \beta$ iff all normal worlds for α satisfy β, and lemma 8 implies the conclusion. ∎

We may now summarize.

Theorem 2 (Representation theorem for loop-cumulative relations) *A consequence relation is a loop-cumulative relation iff it is defined by some cumulative ordered model.*

As in the cumulative case one may study the notion of entailment yielded by cumulative ordered models and obtain results that parallel corollaries 1, 2 and 3.

5 Preferential reasoning

5.1 The system P

We shall now consider a system that seems to occupy a central position in the hierarchy of nonmonotonic systems. It is strictly stronger than **CL**, but assumes the existence of disjunction in the language of formulas. We call this system **P**, for *preferential*, because its semantics, described in section 5.2, are a variation on those proposed by Y. Shoham in [38]. The differences (the distinction we make and he does not between states and worlds) are nevertheless technically important, as noticed above just before definition 2, and as will be shown at the end of section 5.2. This very system has been considered by E. Adams [2] and proposed by J. Pearl and H. Geffner [34] to serve as the *conservative core* of a nonmonotonic reasoning system. It is the flat fragment of the system $\mathcal{S}$ studied by J. Burgess in [4].

Definition 15 *The system* **P** *consists of all the rules of* **C** *and the following:*

$$(16)\quad \frac{\alpha \mathrel{|\!\sim} \gamma \ , \ \beta \mathrel{|\!\sim} \gamma}{\alpha \vee \beta \mathrel{|\!\sim} \gamma} \qquad \textbf{(Or)}$$

A consequence relation that satisfies all rules of **P** *is said to be* preferential*.*

The rule **Or** corresponds to the axiom CA of conditional logic. It says that any formula that is, separately, a plausible consequence of two different formulas, should also be a plausible consequence of their disjunction. It is a valid principle of monotonic classical reasoning and does not imply monotonicity, therefore we tend to accept it. Further consideration also seems to support **Or**: if we think that *if John attends the party, normally, the evening will be great* and also that *if Cathy attends the party, normally, the evening will be great* and hear that at least one of Cathy or John will attend the party, shouldn't we be tempted to join in? There is, though, an *epistemic* reading of $\alpha \mathrel{|\!\sim} \beta$ that invalidates the **Or** rule. If we interpret $\alpha \mathrel{|\!\sim} \beta$ as meaning: *if all I know about the world is α then it is sensible for me to suppose that β is true*, we must reject the **Or** rule. Indeed, one may imagine a situation in which α expresses a fact that can very well be true or false but the truth value of which is normally not known to me. If I knew α to be true, that would be quite an abnormal situation in which I may be willing to accept γ. If I knew α to be false, similarly, it would be an exceptional situation in which I may accept γ, but the knowledge that $\alpha \vee \neg\alpha$ is true is essentially void and certainly does not allow me to conclude that anything exceptional is happening. Notice that, in this reading, the left hand side of the symbol $\mathrel{|\!\sim}$ involves a hidden epistemic operator (the right hand side may also do so, but need not). We shall therefore defend the **Or** rule by saying that such a hidden operator should be made explicit and the example just above only invalidates the inference: from $\mathcal{K}\alpha \mathrel{|\!\sim} \gamma$ and $\mathcal{K}\beta \mathrel{|\!\sim} \gamma$ infer $\mathcal{K}(\alpha \vee \beta) \mathrel{|\!\sim} \gamma$. But nobody would defend such an inference anyway.

The interplay between **Or** and the rules of **C** makes **P** a powerful system. For example, **Loop** is a derived rule of **P**. Since this result will be obvious once we have characterized preferential relations semantically, we shall leave a proof-theoretic derivation of **Loop** in **P** for the reader to find.

We shall now put together a number of remarks revolving around the rule **Or**. Our first remark is that we may derive from **Or** a rule that is similar to the *hard* half of the deduction theorem. This rule was suggested in [39]. It is a very useful rule and expresses the fact that deductions performed under strong assumptions may be useful even if the assumptions are not known facts.

Lemma 19 *In the presence of* **Reflexivity**, **Right Weakening** *and* **Left Logical Equivalence***, the rule of* **Or** *implies the following:*

$$(17)\quad \frac{\alpha \wedge \beta \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \beta \rightarrow \gamma} \qquad \textbf{(S)}$$

S *is therefore a derived rule of* **P***.*

Proof: Suppose $\alpha \wedge \beta \mathrel{|\!\sim} \gamma$. We have $\alpha \wedge \beta \mathrel{|\!\sim} \beta \rightarrow \gamma$, by **Right Weakening**. But one has $\alpha \wedge \neg\beta \mathrel{|\!\sim} \beta \rightarrow \gamma$. One concludes by **Or** and **Left Logical Equivalence**. ∎

Our second remark is that, in the presence of **S**, the rule of **Cut** is implied by **And**. Therefore **Reflexivity**, **Left Logical Equivalence**, **Right Weakening**, **And**, **Or** and **Cautious Monotonicity** are an elegant equivalent axiomatization of the system **P**.

Lemma 20 *In the presence of* **Right Weakening**, **S** *and* **And** *imply* **Cut**.

Proof: Use **S**, **And** and **Right Weakening**. ∎

D. Makinson [27] suggested the following rule. It expresses the principle of proof by cases.

$$(18)\quad \frac{\alpha \wedge \neg\beta \mathrel{|\!\sim} \gamma \ \ , \ \ \alpha \wedge \beta \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \gamma} \qquad (\mathbf{D})$$

Lemma 21 *In the presence of* **Reflexivity**, **Right Weakening** *and* **Left Logical Equivalence**,

1. **Or** *implies* **D** *and*
2. **D** *implies* **Or** *in the presence of* **And**.

Therefore **D** *is a derived rule of* **P**.

The proof is left to the reader.

The next lemma gathers some more derived rules of the system **P**. They will be used in the proof of the representation theorem of section 5.3. The importance of these rules is mainly technical. The reader should notice that **P** is a powerful system, in which one may build quite sophisticated proofs.

Lemma 22 *The following are derived rules of* **P**:

$$(19)\quad \frac{\alpha \mathrel{|\!\sim} \gamma \ \ , \ \ \beta \mathrel{|\!\sim} \delta}{\alpha \vee \beta \mathrel{|\!\sim} \gamma \vee \delta}$$

$$(20)\quad \frac{\alpha \vee \gamma \mathrel{|\!\sim} \gamma \ \ , \ \ \alpha \mathrel{|\!\sim} \beta}{\gamma \mathrel{|\!\sim} \alpha \rightarrow \beta}$$

$$(21)\quad \frac{\alpha \vee \beta \mathrel{|\!\sim} \alpha \ \ , \ \ \beta \vee \gamma \mathrel{|\!\sim} \beta}{\alpha \vee \gamma \mathrel{|\!\sim} \alpha}$$

$$(22)\quad \frac{\alpha \vee \beta \mathrel{|\!\sim} \alpha \ \ , \ \ \beta \vee \gamma \mathrel{|\!\sim} \beta}{\alpha \mathrel{|\!\sim} \gamma \rightarrow \beta}$$

Proof: The uses of **Left Logical Equivalence** will not always be mentioned any more. For (19), use first **Right Weakening** on each of the two hypotheses and then **Or**. This seems to be a very intuitive rule that is often useful.

For (20), from the second hypothesis, using **Left Logical Equivalence** we have $(\alpha \vee \gamma) \wedge \alpha \mathrel{|\!\sim} \beta$. By **S** we conclude $\alpha \vee \gamma \mathrel{|\!\sim} \alpha \rightarrow \beta$. But, using the first hypothesis and **Cautious Monotonicity** one may now conclude.

For (21), from both hypotheses and using (19) one concludes $\alpha \vee \beta \vee \gamma \mathrel{|\!\sim} \alpha \vee \beta$. Now, using our first hypothesis and (9) we see $(\alpha \vee \beta) \vee \gamma \mathrel{|\!\sim} \alpha$. Leaving this result for a moment, notice that from the first hypothesis and $\gamma \mathrel{|\!\sim} \gamma$, using (19) we obtain $\alpha \vee \beta \vee \gamma \mathrel{|\!\sim} \alpha \vee \gamma$. Now, coming back to the result we left hanging, using **Cautious Monotonicity**, we may conclude.

For (22), from the second hypothesis one has $(\beta \vee \gamma) \wedge (\alpha \vee \beta \vee \gamma) \mathrel{|\!\sim} \beta$. By **S**: $\alpha \vee \beta \vee \gamma \mathrel{|\!\sim} (\beta \vee \gamma) \rightarrow \beta$. By **Right Weakening**, one may then obtain $\alpha \vee \beta \vee \gamma \mathrel{|\!\sim} \gamma \rightarrow \beta$ But from the two hypotheses, using (19), one obtains: $\alpha \vee \beta \vee \gamma \mathrel{|\!\sim} \alpha \vee \beta$. Using **Cautious Monotonicity** on those last two results, we obtain: $\alpha \vee \beta \mathrel{|\!\sim} \gamma \rightarrow \beta$. Using the first hypothesis and **Cautious Monotonicity** one concludes. ∎

5.2 Preferential Models

We may now describe our version of preferential models. Preferential models are cumulative ordered models in which states are labeled by single worlds (and not sets of worlds). The reasoner has, then, essentially, a preference over worlds (except that the same world may label different states). We may now define the family of models we are interested in.

Definition 16 *A* preferential *model W is a triple $\langle S, l, \prec \rangle$ where S is a set, the elements of which will be called states, $l : S \mapsto \mathcal{U}$ assigns a world to each state and $\prec$ is a strict partial order on S (i.e. an irreflexive, transitive relation), satisfying the* **smoothness condition** *of definition 7.*

Notice that, for a preferential model, $s \models\!\!\!\mid \alpha$ iff $l(s) \models \alpha$. The smoothness condition, here, as explained in section 4.1, is only a technical condition needed to deal with infinite sets of formulas, it is always satisfied in any preferential model in which S is finite, or in which $\prec$ is well-founded (i.e. no infinite descending chains). The requirement that the relation $\prec$ be a strict partial order has been introduced only because such models are nicer and the smoothness condition is easier to check on those models, but the soundness result of lemma 24 is true for the larger family of models, where $\prec$ is just any binary relation. In such a case, obviously, the smoothness condition cannot be dropped even for finite models. The completeness result of theorem 3 holds obviously, too, for the larger family, but is less interesting. Preferential models, since they are cumulative models, define consequence relations as in definition 8.

Y. Shoham, in [38], proposed a more restricted notion of preferential models. He required the set of states S to be a subset of the universe $\mathcal{U}$ and the labeling function l to be the identity. He also required the relation $\prec$ to be a well-order. Any one of those two requirements would make the representation theorem incorrect. The second point is treated in [22]. For the first point, we leave it as an exercise to the reader to show that the following model has no equivalent model in which no label appears twice. Let L be the propositional calculus on two variables p and q. Let S have four states: $s_0 \prec s_2$ and $s_1 \prec s_3$. Let s_0 satisfy p and $\neg q$, s_1 satisfy $\neg p$ and $\neg q$ and s_2 and s_3 both satisfy p and q.

5.3 Characterization of preferential consequence relations

Our first lemma is obvious. It does not hold in cumulative models and should be contradistincted with lemma 6.

Lemma 23 *Let $W = \langle S, l, \prec \rangle$ be a preferential model. For any $\alpha, \beta \in L$, $\widehat{\alpha \vee \beta} = \widehat{\alpha} \cup \widehat{\beta}$.*

Lemma 24 (Soundness) *For any preferential model W, the consequence relation $\mathrel|\joinrel\sim_W$ it defines is a preferential relation, i.e. all the rules of the system* **P** *are satisfied by the relations defined by preferential models.*

Proof: Indeed, as we remarked above, the fact that $\prec$ is a partial order is not used at all. Since a preferential model is a cumulative model, in light of lemma 7, we only need to check the validity of **Or**. Suppose a preferential model $W = \langle S, l, \prec \rangle$ and $\alpha, \beta, \gamma \in L$ are given. Suppose that $\alpha \mathrel|\joinrel\sim_W \gamma$ and $\beta \mathrel|\joinrel\sim_W \gamma$. Any state minimal in $\widehat{\alpha \vee \beta}$ is, by lemma 23, minimal in the set $\widehat{\alpha} \cup \widehat{\beta}$, and therefore minimal in any of the subsets it belongs to. ▮

We shall now begin the proof of the representation theorem. Let us, first, define a relation among formulas, that will turn out to be a pre-ordering whenever the relation $\mathrel|\joinrel\sim$ satisfies **P**.

Definition 17 *We say that α is not less* ordinary *than β and write $\alpha \leq \beta$ iff $\alpha \vee \beta \mid\!\sim \alpha$.*

Indeed, if we would conclude that α is true on the basis that either α or β is true, this means that the former is not more out of the ordinary than the latter. Notice that, if $\mid\!\sim$ satisfies **Reflexivity** and **Left Logical Equivalence**, then for any $\alpha, \beta \in L$, $\alpha \vee \beta \leq \alpha$.

Lemma 25 *If the relation $\mid\!\sim$ is preferential, the relation $\leq$ is reflexive and transitive.*

Proof: Reflexivity follows from **Left Logical Equivalence** and **Reflexivity**. Transitivity follows from (21) of lemma 22. ∎

From now on, and until theorem 3, we shall suppose that the relation $\mid\!\sim$ is preferential.

Lemma 26 *If $\alpha \leq \beta$ and m is a normal world for α that satisfies β, then m is a normal world for β.*

Proof: Suppose $\beta \mid\!\sim \delta$. By (20) of lemma 22, we have $\alpha \mid\!\sim \beta \rightarrow \delta$. If m is normal for α it must satisfy $\beta \rightarrow \delta$, and since it satisfies β, it must satisfy δ. ∎

Lemma 27 *If $\alpha \leq \beta \leq \gamma$ and m is a normal world for α that satisfies γ then it is a normal world for β.*

Proof: By lemma 26, it is enough to show that m satisfies β. By (22) of lemma 22 we have $\alpha \mid\!\sim \gamma \rightarrow \beta$, but m is a normal world for α that satisfies γ, therefore it must satisfy β. ∎

We may now describe the preferential model we need for the representation result. Remember that we start from any preferential relation $\mid\!\sim$. We then consider the following model: $W \stackrel{\text{def}}{=} \langle S, l, \prec \rangle$ where

1. $S \stackrel{\text{def}}{=} \{< m, \alpha > \mid m \text{ is a normal world for } \alpha\}$,
2. $l(< m, \alpha >) = m$ and
3. $< m, \alpha > \prec < n, \beta >$ iff $\alpha \leq \beta$ and $m \not\models \beta$.

The first thing we want to show is that W is a preferential model, i.e. that $\prec$ is a strict partial order and that W satisfies the smoothness condition. We shall then show that the relation $\mid\!\sim_W$ is exactly $\mid\!\sim$.

Lemma 28 *The relation $\prec$ is a strict partial order, i.e. it is irreflexive and transitive.*

Proof: The relation $\prec$ is irreflexive since $< m, \alpha > \prec < m, \alpha >$ would imply $m \not\models \alpha$, but m is a normal world for α, and since $\alpha \mid\!\sim \alpha$ by **Reflexivity**, it satisfies α. It is left to show that $\prec$ is transitive. Suppose $< m_0, \alpha_0 > \prec < m_1, \alpha_1 >$ and $< m_1, \alpha_1 > \prec < m_2, \alpha_2 >$. By the definition of $\prec$ we have $\alpha_0 \leq \alpha_1$ and $\alpha_1 \leq \alpha_2$. From this we may conclude two things. First, by lemma 25 we conclude $\alpha_0 \leq \alpha_2$. Secondly, since m_0 is a normal world for α_0 that does not satisfy α_1, we may conclude by lemma 27 that it does not satisfy α_2. ∎

We are now going to characterize all minimal states in sets of the form $\widehat{\alpha}$.

Lemma 29 *In the model W, $< m, \beta >$ is minimal in $\widehat{\alpha}$ iff $m \models \alpha$ and $\beta \leq \alpha$.*

Proof: For the *if* part, suppose $m \models \alpha$ and $\beta \leq \alpha$. Clearly $m \in \widehat{\alpha}$. Suppose now that $< n, \gamma > \prec < m, \beta >$ and $n \models \alpha$. We would have $\gamma \leq \beta \leq \alpha$, n normal for γ, and $n \not\models \beta$ and $m \models \alpha$. This stands in contradiction with lemma 27.

For the *only if* part, suppose $< m, \beta >$ is minimal in $\widehat{\alpha}$. Clearly $m \models \alpha$. Suppose n is a normal world for $\alpha \vee \beta$ that does not satisfy β (it is not claimed that such a normal world exists). Since $\alpha \vee \beta \leq \alpha$, we must have $< n, \alpha \vee \beta > \prec < m, \beta >$. But n is a normal world for $\alpha \vee \beta$ that does not satisfy β and therefore must satisfy α. This stands in contradiction with the minimality of $< m, \beta >$ in $\widehat{\alpha}$. We conclude that every normal world for $\alpha \vee \beta$ satisfies β. By lemma 8, $\alpha \vee \beta \mathrel{|\!\sim} \beta$. ∎

We shall now prove that W satisfies the smoothness condition.

Lemma 30 *For any $\alpha \in L$, $\widehat{\alpha}$ is smooth.*

Proof: Suppose $< m, \beta > \in \widehat{\alpha}$, i.e., $m \models \alpha$. If $\beta \leq \alpha$ then, by lemma 29 $< m, \beta >$ is minimal in $\widehat{\alpha}$. On the other hand, if $\alpha \vee \beta \mathrel{|\!\not\sim} \beta$ then by lemma 8 there is a normal world n for $\alpha \vee \beta$ such that $n \not\models \beta$. But $\alpha \vee \beta \leq \beta$ and therefore $< n, \alpha \vee \beta > \prec < m, \beta >$. But, $n \models \alpha \vee \beta$ and $n \not\models \beta$ therefore $n \models \alpha$. Since $\alpha \vee \beta \leq \alpha$, Lemma 29 enables us to conclude that $< n, \alpha \vee \beta >$ is minimal in $\widehat{\alpha}$. ∎

We have shown that W is a preferential model. We shall now show that $\mathrel{|\!\sim}_W$ is exactly the relation $\mathrel{|\!\sim}$ we started from.

Lemma 31 *If $\alpha \mathrel{|\!\sim} \beta$, then $\alpha \mathrel{|\!\sim}_W \beta$.*

Proof: We must show that all minimal states of $\widehat{\alpha}$ satisfy β. Suppose $< m, \gamma >$ is minimal in $\widehat{\alpha}$. Then m is a normal world for γ that satisfies α. By lemma 29, $\gamma \leq \alpha$ and therefore, by lemma 26, m is a normal world for α. ∎

Lemma 32 *If $\alpha \mathrel{|\!\sim}_W \beta$, then $\alpha \mathrel{|\!\sim} \beta$.*

Proof: It follows from the definition of the relation $\prec$ (lemma 29 could also be used, but is not really necessary here) that, given any normal world m for α, $< m, \alpha >$ is minimal in $\widehat{\alpha}$. If $\alpha \mathrel{|\!\sim}_W \beta$, β is satisfied by all normal worlds for α, and we may conclude by lemma 8. ∎

We may now state the main result of this section.

Theorem 3 (Representation theorem for preferential relations) *A consequence relation is a preferential consequence relation iff it is defined by some preferential model.*

Proof: The *if* part is Lemma 24. For the *only if* part, let $\mathrel{|\!\sim}$ be any consequence relation satisfying the rules above and let W be defined as above. Lemmas 28 and 30 show that W is a preferential model. Lemmas 31 and 32 show that it defines an consequence relation that is exactly $\mathrel{|\!\sim}$. ∎

As in the cumulative and cumulative ordered cases we may study the notion of preferential entailment and obtain results similar to Corollaries 1, 2 and 3.

5.4 Some rules that cannot be derived in P

Is **P** a reasonable system for nonmonotonic reasoning? We think a good reasoning system should validate all the rules of **P**. Notice that all the rules we have considered so far are of the form: from the presence of certain assertions in the consequence relation, deduce the presence of some other assertion. After careful consideration of many other rules of this form, we may say we have good reason to think that there are no rules of this type that should be added. Certain principles of reasoning that seem appealing, though, fail to be validated by certain preferential consequence relations. This means, in our sense, that many agents that reason in a way that is fully consistent with all the rules of **P**, nevertheless behave irrationally. We shall show that circumscription does not, in general, satisfy even the weakest of the principles we shall present. The reader will notice that the form of these principles is different from that of all the rules previously discussed: from the *absence* of certain assertions in the relation, we deduce the *absence* of some other assertion.

(23) $\dfrac{\alpha \wedge \gamma \not\mid\!\sim \beta \ , \ \alpha \wedge \neg\gamma \not\mid\!\sim \beta}{\alpha \not\mid\!\sim \beta}$ (**Negation Rationality**)

(24) $\dfrac{\alpha \not\mid\!\sim \gamma \ , \ \beta \not\mid\!\sim \gamma}{\alpha \vee \beta \not\mid\!\sim \gamma}$ (**Disjunctive Rationality**)

(25) $\dfrac{\alpha \wedge \beta \not\mid\!\sim \gamma \ , \ \alpha \not\mid\!\sim \neg\beta}{\alpha \not\mid\!\sim \gamma}$ (**Rational Monotonicity**)

Each one of those rules is implied by **Monotonicity** and therefore expresses some kind of restricted monotonicity. Any rational reasoner should, in our opinion, support them, and we shall, now, explain and justify them. The rule of **Negation Rationality** says that inferences are not made solely on the basis of ignorance. If we accept that β is a plausible consequence of α, we must either accept that it is a plausible consequence of $\alpha \wedge \gamma$ or accept that it is a plausible consequence of $\alpha \wedge \neg\gamma$. Indeed, suppose we hold that *normally, the party should be great*, but that we do not hold that *even if Peter comes to the party, it will be great*, i.e. we seriously doubt the party could stand Peter's presence. It seems we could not possibly hold that we also seriously doubt that the party could stand Peter's absence. If we do not expect the party to be great if Peter is there and do not expect it to be great if Peter is not there, how could we expect it to be great? After all, either Peter is going to be there or he is not. It is, though, easy to find examples of preferential models that define consequence relations that do not satisfy **Negation Rationality**.

We shall even show, now, that circumscriptive reasoning does not always obeys **Negation Rationality**. Suppose our language has two unary predicate symbols *special* and *beautiful*, and one individual constant a. We know that, *normally an object is not special*, i.e. we circumscribe by minimizing the extension of *special*, keeping *beautiful* constant. Take α to be **true** and β to be $\neg special(a)$. Indeed, without any information, we shall suppose that a is not special. But take γ to be $beautiful(a) \leftrightarrow special(a)$. If we had the information that a is beautiful if and only if it is special, we could not conclude that a is not special anymore, since it could well be beautiful, i.e. there are two *minimal* models that must be considered: the first one with a neither beautiful nor special and the second one with a beautiful and special. On the other hand, had we had the informatiom that either a is beautiful or it is special but not both, we could not have concluded that it is not special either, since it could well not be beautiful. It seems that circumscription may lead to unexpected conclusions. The example presented here is a simplification, due to M. Ginsberg, of an example due to the second author. If we try to understand where circumscription differs from intuitive reasoning, we probably will have to say that, even with the knowledge that a is special if and only if it is beautiful, we would have kept the expectation that it is not special, and therefore gained the expectation that it is not beautiful. Similarly, with the knowledge that a is either special or beautiful but not both, we would have kept the expectation that it is not special and therefore formed the expectation that it is beautiful.

The rule of **Disjunctive Rationality** says that inferences made from a disjunction of propositions must be supported by at least one of the component propositions. Again, this seems like a reasonable requirement. If we do not hold that *if Peter comes to the party, it will be great* and do not hold that *if Cathy comes to the party, it will be great*, how could we hold that *if at least one of Peter or Cathy comes, the party will be great*?

In this example, the reader may prefer to read *even if* instead of *if*, but the conclusion stands anyway. It is easy to see that **Disjunctive Rationality** implies **Negation Rationality**. The second author recently showed that **Disjunctive Rationality** is strictly stronger than **Negation Rationality**.

The rule of **Rational Monotonicity** is similar to the axiom CV of conditional logic. It expresses the fact that only additional information the negation of which was expected should force us to withdraw plausible conclusions previously drawn. It is an important tool in minimizing the updating we have to do when learning new information. Suppose we hold that *normally, the party will be great* but do not hold that *even if Peter comes, the party will be great*, i.e. we think Peter's presence could well spoil the party, shouldn't we hold that *normally, Peter will not come to the party*? One easily shows that, in the presence of the rules of **C**, **Rational Monotonicity** implies **Disjunctive Rationality**. D. Makinson proved that **Rational Monotonicity** is strictly stronger than **Disjunctive Rationality** and conjectured a model-theoretic characterization of preferential relations that satisfy **Rational Monotonicity**. The second author proved the corresponding representation result in the case the language L is finite. The third author lifted the restriction on L. These results will appear in a separate paper.

5.5 Examples: diamonds and triangles

We shall now show what preferential reasoning may provide in the setting of two toy situations that have become classics in the literature. First the so-called *Nixon diamond*. Suppose our knowledge base **K** contains the four assertions that follow. The reader may read *teen-ager* for t, *poor* for p, *student* for s and *employed* for e.

1. $t \mid\!\sim p$
2. $t \mid\!\sim s$
3. $p \mid\!\sim e$
4. $s \mid\!\sim \neg e$

It is easy to see, by describing suitable preferential models, that no assertion that would look like some kind of contradiction is preferentially entailed by **K**. In particular neither $t \mid\!\sim e$, nor $t \mid\!\sim \neg e$ is preferentially entailed by **K**. We cannot conclude, from the information given above, that teen-agers are normally employed, neither can we conclude that they generally are not employed. This seems much preferable than the consideration of multiple extensions. This weakness of the system **P** seems to be exactly what we want. Nevertheless, preferential reasoning allows for some quite subtle conclusions. For example the following assertions are preferentially entailed by **K**: **true** $\mid\!\sim \neg t$ (*normally, people are not teen-agers*), **true** $\mid\!\sim \neg(p \wedge s)$ (*normally, people are not poor students*). The following assertions are not preferentially entailed: $s \mid\!\sim \neg p$ (*students, normally are not poor*), or $p \mid\!\sim \neg s$ (*poor persons are normally not students*), and we feel indeed that there is not enough information in **K** to justify them. An example of an assertion that is not preferentially entailed by **K** but we think should follow from **K** is: $a \wedge p \mid\!\sim e$, since a is not mentioned in **K**. The reader may consult [20] for a possible solution.

A second classical example is the *penguin triangle*. Suppose our knowledge base **K** contains the three assertions that follow. The reader may read *penguin* for p, *flies* for f, and *bird* for b.

1. $p \mid\!\sim b$
2. $p \mid\!\sim \neg f$
3. $b \mid\!\sim f$

It is easy to see, by describing suitable preferential models, that no assertion that would lead to some kind of contradiction is preferentially entailed by **K**. In particular $p \mathrel{|\!\sim} f$ is not preferentially entailed by **K**. On the other hand, the following assertions are preferentially entailed by **K** and we leave it to the reader to show that they are satisfied by all preferential models that satisfy **K**:

1. $p \wedge b \mathrel{|\!\sim} \neg f$
2. $f \mathrel{|\!\sim} \neg p$
3. $b \mathrel{|\!\sim} \neg p$
4. $b \vee p \mathrel{|\!\sim} f$
5. $b \vee p \mathrel{|\!\sim} \neg p$

The reader should remark that no *multiple extension* problem arises here and that preferential reasoning correctly chooses the most specific information and in effect pre-empts the application of a less specific default.

5.6 Horn assertions

In this section we shall show that, if we consider only assertions of a restricted type (i.e. Horn assertions), then the system **P** is no stronger than **CL**. For this result we shall need the full strength of theorem 2. To keep notations simple, let us suppose L is a propositional language.

Definition 18 *An assertion $\alpha \mathrel{|\!\sim} \beta$ will be called a Horn assertion iff the antecedent α is a conjunction of zero or more propositional variables and the consequent β is either a single propositional variable or the formula* **false***.*

The crucial remark is the following.

Lemma 33 *If W is a cumulative ordered model, there is a preferential model V such that $\mathrel{|\!\sim}_W$ and $\mathrel{|\!\sim}_V$ coincide as far as Horn assertions are concerned.*

Proof: Let W be the model $\langle\, S, l, \prec\rangle$. We shall define V to be the model $\langle\, S, l', \prec\rangle$, where l' is defined in the following way. For any $s \in S$ and for any propositional variable p, $l'(s) \models p$ iff for every $u \in l(s)$, $u \models p$, in other words iff $s \models\!\!\equiv p$ in W. It is clear that, if α is a conjunction of propositional variables then the sets $\widehat{\alpha}$ in W and V coincide. Therefore, if W satisfies the smoothness condition, so does V and $\mathrel{|\!\sim}_W$ and $\mathrel{|\!\sim}_V$ agree on Horn formulas. ▮

Theorem 4 *Let* **K** *be a knowledge base containing only Horn assertions, and $\mathcal{A}$ a Horn assertion. If the assertion $\mathcal{A}$ may be derived from* **K** *in the system* **P***, then it may be derived from* **K** *in the system* **CL***.*

Proof: Suppose $\mathcal{A}$ cannot be derived in **CL**. By the representation theorem 2, there is a cumulative ordered model W that satisfies all the assertions of **K**, but does not satisfy $\mathcal{A}$. By lemma 33, there is a preferential model V that satisfies **K**, but does not satisfy $\mathcal{A}$. We conclude, by the soundness part of theorem 3, that $\mathcal{A}$ cannot be derived in **P**. ▮

6 Cumulative monotonic reasoning

6.1 The system CM

In section 3.3, three rules were shown equivalent in the presence of the rules of **C**. We shall now study the system obtained by adding those rules (or one of them) to the system **C**. One obtains a system that is strictly stronger than **CL**, but incomparable with **P**. It is corresponds to some natural family of models.

Definition 19 *The system* **CM** *contains all the rules of* **C** *and the rule of* **Monotonicity***, defined in (10). A consequence relation that satisfies all the rules of* **CM** *is said to be* cumulative monotonic.

In fact, **Left Logical Equivalence** and **Cautious Monotonicity** are now redundant, since they follow from **Monotonicity**. From lemma 3, one sees that **EHD** and **Transitivity** are derived rules of **CM**. It is obvious that **Loop** is also a derived rule of **CM** (by **Transitivity**). It is not difficult to find preferential models that do not satisfy **Monotonicity** and we conclude that **CM** is strictly stronger than **CL** and not weaker than **P**.

6.2 Simple cumulative models

Definition 20 *A cumulative model will be called a* simple *cumulative model iff the binary relation $\prec$ on its states is empty.*

A simple cumulative model is a cumulative ordered model. The smoothness condition is always satisfied in such a model. It is very easy to see that the consequence relation defined by any simple cumulative model satisfies **Monotonicity**. It is not difficult to find simple cumulative models that do not satisfy certain instances of the **Or** rule. We conclude that **P** and **CM** are incomparable. It is also easy to find such models that do not satisfy certain instances of **Contraposition**.

6.3 Characterization of monotonic cumulative consequence relations

Theorem 5 (Representation theorem for cumulative monotonic relations) *A consequence relation is cumulative monotonic iff it is defined by some simple cumulative model.*

Proof: It has been noticed above that the *if* part is trivial. For the *only if* part, suppose $\mid\!\sim$ is a consequence relation that satisfies the rules of **CM**. Let $W \stackrel{\text{def}}{=} \langle A, l, \emptyset \rangle$, where $A \subseteq L$ is the set of all formulas α such that $\alpha \not\mid\!\sim$ **false** and $l \stackrel{\text{def}}{=} \{m \mid m$ is a normal world for $\alpha\}$. Lemma 8 implies that all labels are non-empty. By lemma 8, for any formula α, $\widehat{\alpha} = \{\beta \mid \beta \mid\!\sim \alpha\}$. Since all states of $\widehat{\alpha}$ are minimal in $\widehat{\alpha}$, we see that $\alpha \mid\!\sim_W \beta$ iff for all γ such that $\gamma \mid\!\sim \alpha$ and all normal worlds m for γ, $m \models \beta$. By lemma 8 this last condition is equivalent to $\gamma \mid\!\sim \beta$ and we have: $\alpha \mid\!\sim_W \beta$ iff for any γ, $\gamma \mid\!\sim \alpha \Rightarrow \gamma \mid\!\sim \beta$. Suppose $\alpha \mid\!\sim \beta$, take any γ such that $\gamma \mid\!\sim \alpha$, we have by **Transitivity**, a derived rule of **CM**, $\gamma \mid\!\sim \beta$. Therefore $\alpha \mid\!\sim_W \beta$. Suppose now that $\alpha \mid\!\sim_W \beta$, then, by taking $\gamma = \alpha$ one sees that $\alpha \mid\!\sim \beta$. ▮

As in the cumulative, cumulative ordered and preferential cases, one may study the notion of entailment yielded by simple cumulative models and obtain results similar to Corollaries 1, 2 and 3.

7 Monotonic reasoning

7.1 The system M

The results presented in this section are probably folklore. They are presented here for completeness' sake.

Definition 21 *The system* **M** *consists of all the rules of* **C** *and the rule of* **Contraposition**. *A consequence relation that satisfies all the rules of* **M** *is said to be* monotonic.

Lemma 4 and the results to come will show that the system **M** is strictly stronger than **P** and **CM**.

Lemma 34 *The rule* **Or** *is a derived rule of* **M**.

Proof: Use **Contraposition** twice, then **And** and finally **Contraposition**. ▮

Lemma 35 *A consequence relation is monotonic iff it satisfies* **Reflexivity**, **Right Weakening**, **Monotonicity**, **And** *and* **Or**.

Proof: The *only if* part follows from lemmas 2, 4 and 34. For the *if* part, notice, first, that **Left Logical Equivalence** and **Cautious Monotonicity** are special cases of **Monotonicity**. The remark preceding lemma 20 shows that all rules of **P** may be derived from the rules above. We must now show that **Contraposition** may be derived from the rules of **P** and **Monotonicity**. Suppose $\alpha \mid\sim \beta$. By **S**, one has **true** $\mid\sim \alpha \rightarrow \beta$. By **Right Weakening**, we conclude **true** $\mid\sim \neg\beta \rightarrow \neg\alpha$. By **Monotonicity**, we have $\neg\beta \mid\sim \neg\beta \rightarrow \neg\alpha$. We conclude by **Reflexivity** and **MPC**. ▮

7.2 Simple preferential models

The account of monotonic reasoning that we propose is essentially the following. The agent has in mind a set of possible worlds V: this is the set of worlds the agent thinks are possible in practice. This set V is a subset of the set $\mathcal{U}$ of all logically possible worlds. The agent is willing to conclude β from α if all worlds of V that satisfy α also satisfy β.

Definition 22 *A* simple *preferential model is a preferential model in which the binary relation $\prec$ is empty.*

A simple preferential model is a simple cumulative model in which the labeling function l labels each state with a single world. Since repeated labels are obviously useless we could, as well, have considered a model to be a subset of $\mathcal{U}$.

7.3 Characterization of monotonic consequence relations

Theorem 6 (Representation theorem for monotonic relations) *A consequence relation is monotonic iff it is defined by some simple preferential model.*

Proof: The proof of the *if* part is trivial. For the *only if* part we shall build a simple preferential model for any given monotonic consequence relation $\mid\!\sim$. Let $V \stackrel{\text{def}}{=} \{m \in \mathcal{U} \mid \forall\alpha,\beta \in L, \text{ if } \alpha \mid\!\sim \beta \text{ then } m \models \alpha \rightarrow \beta\}$ and let $W \stackrel{\text{def}}{=} \langle V, l \rangle$ where l is the identity function. So, $m \equiv \alpha$ iff $m \models \alpha$.

We shall prove that $\alpha \mid\!\sim \beta$ iff $\alpha \mid\!\sim_W \beta$. If $\alpha \mid\!\sim \beta$ then by the construction of V, $\alpha \mid\!\sim_W \beta$. Suppose now that $\alpha \not\mid\!\sim \beta$, we shall show that there is a world $m \in V$ that does not satisfy $\alpha \rightarrow \beta$. Let $\Gamma_0 \stackrel{\text{def}}{=} \{\neg\beta\} \cup \{\delta \mid \alpha \mid\!\sim \delta\}$. Since $\alpha \not\mid\!\sim \beta$, Γ_0 is satisfiable (the full proof is given in a more general case in lemma 8). Let m be a world that satisfies Γ_0. We shall prove that $\forall \varphi, \psi \in L$ if $\varphi \mid\!\sim \psi$ then $m \models \varphi \rightarrow \psi$. If $\varphi \mid\!\sim \psi$ then **true** $\mid\!\sim \varphi \rightarrow \psi$ by **S** and $\alpha \mid\!\sim \varphi \rightarrow \psi$ by **Monotonicity**. Therefore, $\varphi \rightarrow \psi \in \Gamma_0$ by the definition of Γ_0, and $m \models \varphi \rightarrow \psi$. We conclude that $m \in V$ and clearly $m \equiv \alpha$ but $m \not\equiv \beta$. ▮

It will now be shown that all the constructions and results described above relativize without problems to a given set of conditional assertions.

Corollary 4 *Let* $\mathbf{K}$ *be a set of conditional assertions, and* $\alpha, \beta \in L$. *Let* $\Delta \stackrel{\text{def}}{=} \{\gamma \rightarrow \delta \mid \gamma \mid\!\sim \delta \in \mathbf{K}\}$ *and let* W *be the monotonic model* $\langle \mathcal{U}_\Delta, l \rangle$, *where* l *is the identity function. The notation* $\mathcal{U}_\Delta$ *has been defined in section 2.1. The following conditions are equivalent. If they are satisfied we shall say that* $\mathbf{K}$ *monotonically entails* $\alpha \mid\!\sim \beta$.

1. *for all monotonic models* V *such that* $\mid\!\sim_V$ *contains* $\mathbf{K}$, $\alpha \mid\!\sim_V \beta$
2. $\alpha \mid\!\sim_W \beta$
3. $\alpha \mid\!\sim \beta$ *has a proof from* $\mathbf{K}$ *in the system* $\mathbf{M}$.
4. $\alpha \rightarrow \beta$ *follows logically (with respect to* $\mathcal{U}$*) from the formulas of* Δ.

Proof: We shall first show the equivalence of 1 and 2. The relation defined in 1 is the intersection of all those monotonic consequence relations that contain $\mathbf{K}$. If V is any monotonic model such that $\mid\!\sim_V$ contains $\mathbf{K}$ then the labels of its states must be in $\mathcal{U}_\Delta$ (as defined in 2) and therefore $\mid\!\sim_V$ contains $\mid\!\sim_W$. But $\mid\!\sim_W$ contains $\mathbf{K}$ and is one of the relations considered in 1. To see the equivalence of 1 and 3, notice that the relation defined in 3 is the intersection of all those monotonic relations that contain $\mathbf{K}$. Theorem 6 implies that 1 and 3 define the same relation. The equivalence between 2 and 4 is immediate. ▮

From the equivalence of conditions 1 and 3 one easily proves the following compactness result:

Corollary 5 (compactness) $\mathbf{K}$ *monotonically entails* $\alpha \mid\!\sim \beta$ *iff a finite subset of* $\mathbf{K}$ *does.*

8 Summary, future work and conclusion

Five families of models and consequence relations have been defined and their relations will be summarized here. Each family has been characterized by a logical system and no two of those systems are equivalent. The family of cumulative models contains all other families and is characterized by the logical system **C** that consists of **Logical Left Equivalence**, **Right Weakening**, **Reflexivity**, **Cut** and **Cautious Monotonicity**. The next largest family is that of cumulative ordered models. It contains all three families not yet mentioned here. It is characterized by the logical system **CT** that contains, in addition to the rules of **C**, the rule of **Loop**. The families of simple cumulative models and of preferential models are two incomparable subfamilies of the family of cumulative ordered models. Simple cumulative models are characterized by the logical system **CM** that contains, in addition to the rules of **C**, the rule of **Monotonicity** (or equivalently, **Transitivity**). The family of preferential models, probably the most important one, is characterized by the logical system **P** that contains, in addition to the rules of **C** the rule of **Or**. The family of monotonic models is the smallest one of them all. It is contained in all other four. It is characterized by the logical system **M** that contains, in addition to the rules of **C**, both rules **Monotonicity** and **Or**.

Of those families of consequence relations, which is the best suited to represent the inferences of a nonmonotonic reasoner in the presence of a fixed knowledge base? Monotonic and and cumulative monotonic reasoning are too powerful, i.e. simple cumulative and simple preferential models are too restrictive to represent the wealth of nonmonotonic inference procedures we would like to consider. We feel that all bona fide logical systems should implement reasoning patterns that fall inside the framework of cumulative reasoning, but probably not all cumulative models represent useful nonmonotonic systems. The same may probably said about cumulative ordered models. Preferential reasoning seems to be closest to what we are looking for.

Nevertheless, many preferential reasoners lack properties that seem desirable, for example **Rational Monotonicity**. A major problem that is not solved in this paper is to describe reasonable inference procedures that would guarantee that the set of assertions that may be deduced from any conditional knowledge base satisfies the property of **Rational Monotonicity**. The second author proposed a solution to this problem in [20]. Another major problem, not solved here, is to extend the results presented here to predicate calculus and answer the question: how should quantifiers be treated, or what is the meaning of the conditional assertion $\mathit{bird}(x) \mid\sim \mathit{fly}(x)$? The second and third authors have a solution, still unpublished, to this problem too.

We hope the results presented above will convince the reader that the field of artificial nonmonotonic reasoning may benefit from the study of nonmonotonic consequence relations.

9 Acknowledgments

Discussions with and remarks from David Makinson, Johan van Benthem, David Israel, Benjamin Grosof and an anonymous referee helped us put this work in perspective and improve its presentation. They are gratefully acknowledged.

References

[1] Ernest W. Adams. Probability and the logic of conditional. In J. Hintikka and P. Suppes, editors, *Aspects of Inductive Logic*. North Holland, Amsterdam, 1966.

[2] Ernest W. Adams. *The Logic of Conditionals*. D. Reidel, Dordrecht, 1975.

[3] Arnon Avron. Simple consequence relations. LFCS Report Series 87-30, Dept. of Computer Science, Univ. of Edinburgh, June 1987.

[4] John P. Burgess. Quick completeness proofs for some logics of conditionals. *Notre Dame Journal of Formal Logic*, 22:76–84, 1981.

[5] Keith L. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logics and Data Bases*, pages 293–322. Plenum Press, 1978.

[6] James P. Delgrande. A first-order logic for prototypical properties. *Artificial Intelligence*, 33:105–130, 1987.

[7] James P. Delgrande. An approach to default reasoning based on a first-order conditional logic: Revised report. *Artificial Intelligence*, 36:63–90, August 1988.

[8] David W. Etherington, Robert E. Mercer, and Raymond Reiter. On the adequacy of predicate circumscription for closed-world reasoning. *Computational Intelligence*, 1:11–15, 1985.

[9] Michael Freund, Daniel Lehmann, and David Makinson. Canonical extensions to the infinite case of finitary nonmonotonic inference operations. In *Workshop on Nomonotonic Reasoning*, pages 133–138, Sankt Augustin, FRG, December 1989. Arbeitspapiere der GMD no. 443.

[10] Dov M. Gabbay. Theoretical foundations for non-monotonic reasoning in expert systems. In Krzysztof R. Apt, editor, *Proc. of the NATO Advanced Study Institute on Logics and Models of Concurrent Systems*, pages 439–457, La Colle-sur-Loup, France, October 1985. Springer-Verlag.

[11] Dov M. Gabbay. personal communication, March 1989.

[12] Gerhard Gentzen. Über die existenz unabhängiger axiomensysteme zu unendlichen satzsystemen. *Mathematische Annalen*, 107:329–350, 1932.

[13] Gerhard Gentzen. *The Collected Papers of Gerhard Gentzen, edited by M. E. Szabo.* North Holland, Amsterdam, 1969.

[14] Matthew L. Ginsberg. Counterfactuals. *Artificial Intelligence*, 30:35–79, 1986.

[15] Joseph Y. Halpern and Yoram Moses. Towards a theory of knowledge and ignorance: preliminary report. In *Proc. Workshop on Non-Monotonic Reasoning*, pages 125–143. AAAI, New Paltz, 1984.

[16] Christopher Anthony R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12:576–580, 1969.

[17] David J. Israel. What's wrong with nonmonotonic logic? In *Proceedings of the AAAI National Conference*, pages 99–101, 1980.

[18] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Preferential models and cumulative logics. Technical Report TR 88-15, Leibniz Center for Computer Science, Dept. of Computer Science, Hebrew University, Jerusalem, November 1988.

[19] Daniel Lehmann. Preferential models and cumulative logics. In Ehud Shapiro, editor, *Fifth Israeli Symposium on Artificial Intelligence, Vision and Pattern Recognition*, pages 365–381, Tel Aviv, Israel, December 1988. Information Processing Association of Israel.

[20] Daniel Lehmann. What does a conditional knowledge base entail? In Ron Brachman and Hector Levesque, editors, *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, Canada, May 1989. Morgan Kaufmann.

[21] Daniel Lehmann and Sarit Kraus. Non monotonic logics: Models and proofs. In *European Workshop on Logical Methods in Artificial Intelligence*, pages 58–64, Roscoff (Finistère) France, June 1988.

[22] Daniel Lehmann and Menachem Magidor. Rational logics and their models: a study in cumulative logic. Technical Report TR 88-16, Leibniz Center for Computer Science, Dept. of Computer Science, Hebrew University, Jerusalem, November 1988.

[23] David K. Lewis. Completeness and decidability of three logics of counterfactual conditionals. *Theoria*, 37:74–85, 1971.

[24] David K. Lewis. *Counterfactuals.* Harvard University Press, 1973.

[25] David K. Lewis. Intensional logics without iterative axioms. *Journal of Philosophical Logic*, 3:457–466, 1974.

[26] Vladimir Lifschitz. On the satisfiability of circumscription. *Artificial Intelligence*, 28:17–27, 1986. Research Note.

[27] David Makinson. personal communication, April 1988.

[28] David Makinson. General theory of cumulative inference. In M. Reinfrank, J. de Kleer, M. L. Ginsberg, and E. Sandewall, editors, *Proceedings of the Second International Workshop on Non-Monotonic Reasoning*, pages 1–18, Grassau, Germany, June 1988. Springer Verlag. Volume 346, Lecture Notes in Artificial Intelligence.

[29] John McCarthy. Circumscription, a form of non monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

[30] Drew McDermott and Jon Doyle. Non-monotonic logic I. *Artificial Intelligence*, 13:41–72, 1980.

[31] Robert C. Moore. Possible-world semantics for autoepistemic logic. In *Proceedings of AAAI Workshop on Nomonotonic Reasoning*, pages 396–401, New Paltz, 1984.

[32] Donald Nute. Conditional logic. In Dov M. Gabbay and Franz Guenthner, editors, *Handbook of Philosophical Logic*, chapter Chapter II.8, pages 387–439. D. Reidel, Dordrecht, 1984.

[33] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, P.O. Box 50490, Palo Alto, CA 94303, 1988.

[34] Judea Pearl and Hector Geffner. Probabilistic semantics for a subset of default reasoning. TR CSD-8700XX, R-93-III, Computer Science Dept., UCLA, March 1988.

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

[36] Raymond Reiter. *Nonmonotonic Reasoning*, volume 2 of *Annual Reviews in Computer Science*, pages 147–186. Annual Reviews Inc., 1987.

[37] Dana S. Scott. Completeness and axiomatizability. In L. Henkin & al., editor, *Proceedings of the Tarski Symposium, Proc. of Symposia in Pure Mathematics, Vol. 25*, pages 411–435, Providence, R.I., June 1971. Association for Symbolic Logic, American Mathematical Society.

[38] Yoav Shoham. A semantical approach to nonmonotonic logics. In *Proc. Logics in Computer Science*, pages 275–279, Ithaca, N.Y., 1987.

[39] Yoav Shoham. *Reasoning about Change*. The MIT Press, 1988.

[40] Robert C. Stalnaker. *A Theory of Conditionals*, volume 2 of *American Philosophical Quarterly Monograph Series (Nicholas Rescher, ed.*, pages 98–112. Blackwell, Oxford, 1968.

[41] Alfred Tarski. Fundamentale begriffe der methodologie der deduktiven wissenschaften.i. *Monatshefte für Mathematik und Physik*, 37:361–404, 1930.

[42] Alfred Tarski. Über einige fundamentale begriffe der metamatematik. *Comptes Rendus des séances de la Société des Sciences et des Lettres de Varsovie*, 23:22–29, 1930.

[43] Alfred Tarski. Grunzüge des systemenkalkül, erster teil. *Fundamenta Mathematicæ*, 25:503–526, 1935.

[44] Alfred Tarski. *Logic, Semantics, Metamathematics. Papers from 1923–1938*. Clarendon Press, Oxford, 1956.

[45] David S. Touretzky. *The Mathematics of Inheritance Systems*. Research Notes in Artificial Intelligence. Pitman, London — Morgan Kaufmann, Los Altos, 1986.

[46] Johan van Benthem. Foundations of conditional logic. *Journal of Philosophical Logic*, 13:303–349, 1984.

[47] Frank Veltman. *Logics for Conditionals*. PhD thesis, Filosofisch Instituut, Universiteit van Amsterdam, 1986.