What does a conditional knowledge base entail? *

Daniel Lehmann† Menachem Magidor‡

Abstract

This paper presents a logical approach to nonmonotonic reasoning based on the notion of a nonmonotonic consequence relation. A conditional knowledge base, consisting of a set of conditional assertions of the type **if** ... **then** ..., represents the explicit defeasible knowledge an agent has about the way the world generally behaves. We look for a plausible definition of the set of all conditional assertions entailed by a conditional knowledge base. In a previous paper [17], S. Kraus and the authors defined and studied *preferential* consequence relations. They noticed that not all preferential relations could be considered as reasonable inference procedures. This paper studies a more restricted class of consequence relations, *rational* relations. It is argued that any reasonable nonmonotonic inference procedure should define a rational relation. It is shown that the rational relations are exactly those that may be represented by a *ranked* preferential model, or by a (non-standard) probabilistic model. The rational closure of a conditional knowledge base is defined and shown to provide an attractive answer to the question of the title. Global properties of this closure operation are proved: it is a cumulative operation. It is also computationally tractable. This paper assumes the underlying language is propositional.

1 Introduction

1.1 Background

Inference is the process of achieving explicit information that was only implicit in the agent's knowledge. Human beings are astoundingly good at infering useful and very often reliable information from knowledge that seems mostly irrelevant, sometimes erroneous and even self-contradictory. They are even better at correcting inferences they learn to be in contradiction with reality. It is

*This work was partially supported by grant 351/89 from the Basic Research Foundation, Israel Academy of Sciences and Humanities and by the Jean and Helene Alfassa fund for research in Artificial Intelligence. Its final version was prepared while the first author was visiting the Laboratoire d'Informatique Théorique et de Programmation, Univ. Paris 6

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

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

already a decade now that Artificial Intelligence has realized that the analysis of models of such inferences was a major task.

Many nonmonotonic systems have been proposed as formal models of this kind of inferences. The best known are probably: circumscription [22], the modal systems of [23] and [24], default logic [29] and negation as failure [5]. An up-to-date survey of the field of nonmonotonic reasoning may be found in [30]. Though each of these systems is interesting per se, it is not clear that any one of them really captures the whole generality of nonmonotonic reasoning. Recently (see in particular the panel discussion of [36]) a number of researchers expressed their disappointment at existing systems and suggested that no purely logical analysis could be satisfactory.

This work tries to contradict this pessimistic outlook. It takes a purely logical approach, grounded in A. Tarski's framework of consequence relations [35] and studies the very general notion of a *sensible* conclusion. It seems that this is a common ground that can be widely accepted: all reasonable inference systems draw only sensible conclusions. On the other hand, as will be shown, the notion of a sensible conclusion has a non-trivial mathematical theory and many interesting properties are shared by all ways of drawing sensible conclusions.

The reader is referred to [17] for a full description of background, motivation and the relationship of the present approach with previous work in Conditional Logic. We only wish to add here that, even though the present work will be compared explicitly only with previous work of E. Adams, some of the intuitions developed here are related with intuitions exposed already in the first works on Conditional Logic, such as [28] or [4]. The interested reader may find many relevant articles in [15] and should in particular look at [14]. The main difference between our approach and Conditional Logic is that we take the view that the truth of a conditional assertion is *necessary*, i.e., does not depend on the state of the world. For us, worlds give truth values to propositions but not to assertions, preferential models give truth values to assertions, but not to propositions. The models we propose are therefore much simpler than those previously proposed in Conditional Logic and it is doubtful whether they can shed light on the very complex questions of interest to the Conditional Logic community.

Notations and terminology conform with those of [17], but the present paper is essentially self-contained. Preliminary versions of part of the material contained in this paper appeared in [19, 18]. In [17] it was suggested that items of *default*, i.e., *defeasible* information should be represented as *conditional assertions*, i.e., pairs of formulas. For example, the information that *birds normally fly* will be represented by the conditional assertion $b \mathrel{|\!\sim} f$, where b and f are propositional variables representing *being a bird* and *flying* respectively. A set (finite or infinite) of conditional assertions is called a conditional knowledge base (knowledge base, in short) and represents the defeasible knowledge an agent may have. The fundamental question studied in this paper is the following: given a knowledge base $\mathbf{K}$, what are the conditional assertions that should be considered as entailed, i.e., logically implied, by $\mathbf{K}$? We consider that an assertion

$\alpha \mid\!\sim \beta$ should be entailed by **K** if, on the basis of the defeasible information contained in **K** and knowing that the proposition α is true, it would be sensible to conclude (defeasibly) that β is true.

The question asked in the title and detailed just above has no simple answer and has probably no unique answer good for everyone in every situation. It may well be the case that, in different situations or for different domains of knowledge, the pragmatically *right* answers to the question of the title differ. This feeling has been recently expressed in [9]. The first part of this paper defines the notion of a *rational* set of assertions and defends the thesis that any reasonable answer to the question of the title must consist of such a set of assertions.

Thesis 1 *The set of assertions entailed by any set of assertions is rational.*

The second part of the paper describes a specific construction, *rational closure*, and shows that the rational closure of a set of assertions is rational. This construction is then studied and its value as an answer to the question of the title assessed. We think that, in many situations, this is an acceptable answer, but do not claim that it provides an answer suitable to any situation. We have just argued that such an answer probably does not exist. One of the main interests of the rational closure construction is that it provides a proof of the existence of some uniform, well-behaved and elegant way of answering the question. In doing so, we develop criteria by which to judge possible answers. We shall in particular consider properties of the mapping from **K** to the set of all the assertions it entails and prove that our construction of the rational closure satisfies them. This effort and these results have to be compared with the essential absence, for the moment, of similar results about the systems of nonmonotonic reasoning mentioned above.

1.2 Plan of this paper

We survey here the main parts of this paper. The introductions to the different sections contain a detailed description. Section 2 is devoted to preferential consequence relations. This family of relations was defined and studied in [17]. The first part of this section mainly recalls definitions and results of [17], its last part presents deep new technical results on preferential entailment that will be used in the sequel, but it may be skipped on a first reading. Section 3 presents the restricted family of relations that is of interest to us: rational relations. This family was first defined, but not studied, in [17, Section 5.4]. The main result of this section is a representation theorem characterizing rational relations in terms of ranked models. Section 4 shows that entailment with respect to ranked models is exactly entailment with respect to preferential models and provides an alternative proof of E. Adams' [1] characterization of preferential entailment in terms of his probabilistic semantics. Appendix B describes a family of models based on non-standard (in the sense of A. Robinson) probability models and shows that these models provide another exact representation for rational consequence relations. This provides us with a strong justification for considering rational relations. Section 5 draws on all previous sections and is the heart of this paper. It proposes an answer to the question of the title. The notion of rational closure is first defined abstractly and global properties proved. It is then showed that finite knowledge bases have a rational closure and a model-theoretic construction is provided. An efficient algorithm is proposed for computing the rational closure of a finite knowledge base. We then discuss some examples, remark that rational closure does not provide for inheritance of generic properties to exceptional classes, and finally propose a second thesis.

2 Preferential relations and models

2.1 Introduction

The first part of this section, i.e., Sections 2.2–2.3, recalls definitions and results of [17] and provides an example (new) of a preferential relation that cannot be defined by a well-founded model. Then, in Section 2.4, the definition and some properties of preferential entailment are recalled from [17] and some new remarks included. Preferential entailment is a fundamental notion that is used throughout the paper. The last three sections are essentially independent of each other. They present an in-depth study of preferential entailment. In a first reading, they should probably be read only cursorily. The results of Section 2.5 expand on part of [18] and are used in Section 4.2. Section 2.6 presents a new technique to study preferential entailment (i.e., ranking). It is fundamental from Section 5.6 and onwards. Section 2.7 shows that preferential entailment is in the class co-NP, and hence is an co-NP-complete problem. A preliminary version of this last result appeared in [18].

2.2 Preferential relations

Our first step must be to define a language in which to express the basic propositions. In this paper Propositional Calculus is chosen. Let $\mathcal{L}$ be the set of well formed propositional formulas (thereafter formulas) over a set of propositional variables. If the set of propositional variables chosen is finite, we shall say that $\mathcal{L}$ is logically finite. The classical propositional connectives will be denoted by $\neg, \vee, \wedge, \rightarrow$ and $\leftrightarrow$. The connective $\rightarrow$ therefore denotes material implication. Small Greek letters will be used to denote formulas.

A world is an assignment of truth values to the propositional variables. The set $\mathcal{U}$ is the set of all worlds. The satisfaction of a formula by a world is defined as usual. 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 $\forall u \in \mathcal{U},\ u \models \alpha$.

If α and β are formulas then the pair $\alpha \mid\!\sim \beta$ (read "from α sensibly conclude β") is called a conditional assertion. A conditional assertion is a syntactic object to which the reader may attach any meaning he wants, but 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 represents the information I have about the true state of the world, I will jump to the conclusion that β is true. A conditional knowledge base is any set of conditional assertions. Typically it is a finite set, but need not be so. Conditional knowledge bases seem to provide a terse and versatile way of specifying defeasible information. They correspond to the explicit information an agent may have.

Certain well-behaved sets of conditional assertions will be deemed worthy of being called *consequence relations*. We shall use the notation usual for binary relations to describe consequence relations. So, if $\mid\!\sim$ is a consequence relation, $\alpha \mid\!\sim \beta$ indicates that the pair $\langle \alpha, \beta \rangle$ is in the consequence relation $\mid\!\sim$ and $\alpha \not\mid\!\sim \beta$ indicates it is not in the relation. Consequence relations correspond to the implicit information an intelligent agent may have. Consequence relations are typically infinite sets.

Certain especially interesting properties of sets of conditional assertions (i.e., binary relations on $\mathcal{L}$) will be described and discussed now. They are presented in the form of inference rules. Consequence relations are expected to satisfy those properties.

$$\frac{\models \alpha \leftrightarrow \beta \ \ , \ \ \alpha \mid\!\sim \gamma}{\beta \mid\!\sim \gamma} \qquad \textbf{(Left Logical Equivalence)} \tag{1}$$

$$\frac{\models \alpha \rightarrow \beta \ \ , \ \ \gamma \mid\!\sim \alpha}{\gamma \mid\!\sim \beta} \qquad \textbf{(Right Weakening)} \tag{2}$$

$$\alpha \mid\!\sim \alpha \qquad \textbf{(Reflexivity)} \tag{3}$$

$$\frac{\alpha \mid\!\sim \beta \ \ , \ \ \alpha \mid\!\sim \gamma}{\alpha \mid\!\sim \beta \wedge \gamma} \qquad \textbf{(And)} \tag{4}$$

$$\frac{\alpha \mid\!\sim \gamma \ \ , \ \ \beta \mid\!\sim \gamma}{\alpha \vee \beta \mid\!\sim \gamma} \qquad \textbf{(Or)} \tag{5}$$

$$\frac{\alpha \mid\!\sim \beta \ \ , \ \ \alpha \mid\!\sim \gamma}{\alpha \wedge \beta \mid\!\sim \gamma} \qquad \textbf{(Cautious Monotonicity)} \tag{6}$$

Definition 1 *A set of conditional assertions that satisfies all six properties above is called a* preferential *consequence relation.*

A more leisurely introduction with motivation may be found in [17] where a larger family of consequence relations, that of cumulative relations, was also studied. This family is closely related to the cumulative inference operations studied by D. Makinson in [21]. The attentive reader of [17] may have noticed that, there, we reserved ourselves an additional degree of freedom, that we have denied ourselves here. There, we allowed $\mathcal{U}$ to be a subset of the set of all worlds and considered the $\models$ symbol appearing in **Left Logical Equivalence** and in **Right Weakening** to be interpreted relatively to this subset. This was felt necessary to deal with *hard constraints*. In this work, we shall suppose that a hard constraint α is interpreted as the *soft* constraint, i.e., the assertion, $\neg\alpha \mid\!\sim \mathbf{false}$, which was recognized as equivalent to considering $\mathcal{U}$ to be the set of all worlds satisfying α in [17, page 174]. The second proposal there, i.e., to consider α to be part of the facts, would not be consistent with our treatment of rational closure.

For the reader's ease of mind we shall mention two important derived rules. Both **S** and **Cut** are satisfied by any preferential relation.

$$\frac{\alpha \wedge \beta \mid\!\sim \gamma}{\alpha \mid\!\sim \beta \rightarrow \gamma} \qquad (\mathbf{S}) \tag{7}$$

$$\frac{\alpha \wedge \beta \mid\!\sim \gamma \quad , \quad \alpha \mid\!\sim \beta}{\alpha \mid\!\sim \gamma} \qquad (\mathbf{Cut}) \tag{8}$$

The rule of Cut is presented here in a form that is not the most usual one. Notice, in particular, that we require the left-hand side of the second assumption to be part of the left-hand side of the first assumption. This version of Cut is close to the original form proposed by G. Gentzen. The following form, more usually used now, is *not* acceptable since it implies monotonicity.

$$\frac{\alpha \wedge \beta \mid\!\sim \gamma \quad , \quad \alpha' \mid\!\sim \beta}{\alpha \wedge \alpha' \mid\!\sim \gamma} \tag{9}$$

2.3 Preferential models and representation theorem

The following definitions are also taken from [17] and justified there. We shall define a class of models that we call *preferential* since they represent a slight variation on those proposed in [33]. The differences are nevertheless technically important.

Preferential models give a model-theoretic account of the way one performs nonmonotonic inferences. The main idea is that the agent has, in his mind, a partial ordering on possible states of the world. State s is less than state t, if, in the agent's mind, s is *preferred* to or more *natural* than t. The agent is willing to conclude β from α, if all *most natural* states that satisfy α also satisfy β.

Some technical definitions are needed. Let U be a set and $\prec$ a strict partial order on U, i.e., a binary relation that is antireflexive and transitive.

Definition 2 *Let $V \subseteq U$. We shall say that $t \in V$ is minimal in V iff there is no $s \in V$, such that $s \prec t$. We shall say that $t \in V$ is a minimum of V iff for every $s \in V, s \neq t$, we have $t \prec s$.*

Definition 3 *Let $V \subseteq U$. We shall say that V is smooth iff $\forall t \in V$, either $\exists s$ minimal in V, such that $s \prec t$ or t is itself minimal in V.*

We may now define the family of models we are interested in.

Definition 4 *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 satisfying the following* smoothness condition*: $\forall \alpha \in L$, the set of states $\widehat{\alpha} \stackrel{\text{def}}{=} \{s \mid s \in S,\ s \models\!\!\!\equiv \alpha\}$ is smooth, where $\models\!\!\!\equiv$ is defined as $s \models\!\!\!\equiv \alpha$ (read s satisfies α) iff $l(s) \models \alpha$. The model W will be said to be finite iff S is finite. It will be said to be well-founded iff $\langle S, \prec \rangle$ is well-founded, i.e., iff there is no infinite descending chain of states.*

The smoothness condition is only a technical condition. It is satisfied in any well-founded preferential model, and, in particular, in any finite model. When the language $\mathcal{L}$ is logically finite, we could have limited ourselves to finite models and forgotten the smoothness condition. Nevertheless, Lemma 1 will show that, in the general case, for the representation result of Theorem 1 to hold we could not have required preferential models to be well-founded. 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 is true for the larger family of models, where $\prec$ is just any binary relation (Definitions 2 and 3 also make sense for any binary relation $\prec$). In such a case, obviously, the smoothness condition cannot be dropped even for finite models. The completeness result holds, obviously, also for the larger family, but is less interesting.

We shall now describe the consequence relation defined by a model.

Definition 5 *Suppose a model $W = \langle S, l, \prec \rangle$ and $\alpha, \beta \in L$ are 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 $\widehat{\alpha}$, $s \models\!\!\!\equiv \beta$.*

If $\alpha \mid\!\sim_W \beta$ we shall say that the model W satisfies the conditional assertion $\alpha \mid\!\sim \beta$, or that W is a model of $\alpha \mid\!\sim \beta$.

The following theorem characterizes preferential consequence relations.

Theorem 1 (Kraus, Lehmann and Magidor) *A binary relation $\mid\!\sim$ on $\mathcal{L}$ is a preferential consequence relation iff it is the consequence relation defined by some preferential model. If the language $\mathcal{L}$ is logically finite, then every preferential consequence relation is defined by some finite preferential model.*

The next result shows we could not have restricted ourselves to well-founded models.

Lemma 1 *There is a preferential relation that is defined by no well-founded preferential model.*

Proof: Let $\mathcal{L}$ be the propositional calculus on the variables p_i, $i \in \omega$ (ω is the set of natural numbers). We shall consider the model $W \stackrel{\text{def}}{=} \langle V, l, \prec \rangle$ where V is the set $\{s_i \mid i \in \omega \cup \{\infty\}\}$, $s_i \prec s_j$ iff $i > j$ (i.e., there is an infinite descending chain of states with a bottom element) and $l(s_i)(p_j)$ is **true** iff $j \geq i$, for $i \in \omega \cup \{\infty\}$ and $j \in \omega$. The smoothness property is satisfied since the only subsets of V that do not have a minimum are infinite sets A that do not contain s_∞ and any $\alpha \in L$ that is satisfied in all states of such a set A is also satisfied in s_∞. The model W defines a preferential relation $\mid\!\sim_W$ such that $\forall i \in \omega$, $p_i \mid\!\sim_W p_{i+1}$ and $p_{i+1} \mid\!\sim_W \neg p_i$, but $p_0 \not\mid\!\sim_W$ **false**. But clearly, any preferential model defining such a relation must contain an infinite descending chain of states. ∎

We do not know of any direct characterization of those relations that may be defined by well-founded preferential models. But Lemma 3 will show that many relations may be defined by well-founded preferential models. It is clear, though, that the canonical preferential model provided by the proof of Theorem 1 is rarely well-founded. Consider, for example, the preferential closure of the empty knowledge base on a logically infinite language $\mathcal{L}$. It may be defined by some well-founded preferential model (the order $\prec$ is empty). But its canonical model is not well-founded (consider states whose second components are larger and larger disjunctions). We may only make the following obvious remark: if the underlying language $\mathcal{L}$ is logically finite, then all canonical models are well-founded.

2.4 Preferential entailment

Now that we have a proof-theoretic definition of a class of relations, a class of models and a representation theorem relating them, it is natural to put down the following definition. It will serve us as a first approximate answer to the question of the title.

Definition 6 *The assertion $\mathcal{A}$ is* preferentially entailed *by* $\mathbf{K}$ *iff it is satisfied by all preferential models of* $\mathbf{K}$*. The set of all conditional assertions that are preferentially entailed by* $\mathbf{K}$ *will be denoted by* $\mathbf{K}^p$*. The preferential consequence relation* $\mathbf{K}^p$ *is called the preferential closure of* $\mathbf{K}$*.*

In [17] it was noted that the characterization of preferential consequence relations obtained in Theorem 1 enables us to prove the following.

Theorem 2 *Let* $\mathbf{K}$ *be a set of conditional assertions, and* $\mathcal{A}$ *a conditional assertion. The following conditions are equivalent:*

1. *$\mathcal{A}$ is preferentially entailed by* $\mathbf{K}$*, i.e.,* $\mathcal{A} \in \mathbf{K}^p$
2. *$\mathcal{A}$ has a proof from* $\mathbf{K}$ *in the system* $\mathbf{P}$ *consisting of the Rules 1 to 6.*

The following compactness result follows.

Corollary 1 (compactness) **K** *preferentially entails* $\mathcal{A}$ *iff a finite subset of* **K** *does.*

The following also follows from Theorem 2.

Corollary 2 *The set* $\mathbf{K}^p$*, considered as a consequence relation, is a preferential consequence relation, therefore there is a preferential model that satisfies exactly the assertions of* $\mathbf{K}^p$*. If* **K** *is itself a preferential consequence relation then* $\mathbf{K} = \mathbf{K}^p$*. The set* $\mathbf{K}^p$ *grows monotonically with* **K**.

We see that the operation $\mathbf{K} \mapsto \mathbf{K}^p$ is a compact monotonic consequence operation in the sense of Tarski [35]. We have a particular interest in finite knowledge bases. It is therefore useful to put down the following definition.

Definition 7 *A preferential consequence relation is* finitely generated *iff it is the preferential closure of a finite knowledge base.*

Lemma 3 will show that finitely generated relations have interesting properties. In [17], it was shown that any preferential relation defines a strict ordering on formulas by: $\alpha < \beta$ iff $\alpha \vee \beta \mathrel{|\!\sim} \alpha$ and $\alpha \vee \beta \mathrel{|\!\not\sim} \beta$.

Definition 8 *A preferential relation is* well-founded *iff the strict ordering relation* $<$ *it defines is well-founded.*

The following is easy to show.

Lemma 2 *A preferential relation is well-founded iff the canonical model built in the proof of Theorem 1 is well-founded.*

We noticed, at the end of Section 2.3, that not all preferential relations that may be defined by well-founded preferential models are well-founded.

Lemma 3 *Any finitely generated preferential relation is defined by some well-founded preferential model.*

Proof: Let **K** be any finite set of assertions. Let $L_i, i \in \omega$ be an infinite sequence of larger and larger *logically finite* sublanguages of L such that every L_i contains all the formulas appearing in the assertions of **K** and such that L is the union of the L_i's. By Theorem 1, for each i there is a *finite* preferential model W'_i that defines the preferential closure of **K** over L_i. Let W_i be the finite preferential model (over L) obtained by extending the labeling function of W'_i to the variables of $L - L_i$ in some arbitrary way. Clearly W_i is a preferential model of **K**. Let W be the structure obtained by putting all the W_i's one alongside the other (the partial ordering $\prec$ on W never relates states belonging to W_i's with different i's). The structure W is well-founded, therefore satisfies the smoothness condition and is a preferential model. Any assertion that is preferentially entailed by **K** (over L) is satisfied by every W_i, and is therefore satisfied by W. For any assertion $\mathcal{A}$ that is not preferentially entailed by $\mathbf{K}$ one may find a language L_i large enough to include the formulas of $\mathcal{A}$. Over L_i, the assertion $\mathcal{A}$ is not preferentially entailed by $\mathbf{K}$, by Theorem 2, since a proof in the small language is a proof in the larger one. Therefore W'_i does not satisfy $\mathcal{A}$. We conclude that W_i does not satisfy $\mathcal{A}$ and that W does not satisfy $\mathcal{A}$. ∎

2.5 Some properties of preferential entailment

The following result, Theorem 3, is new. It is important for several reasons. It uses the semantic representation of Theorem 1 and a direct proof using only proof-theoretic arguments seems difficult. It will be used in Section 4.2. Its Corollary 3 should provide a starting point for the application to preferential entailment of methods based on or related to resolution. First a definition.

Definition 9 *If a formula α is such that $\alpha \not\mathrel{|\!\sim} \mathbf{false}$, we shall say that α is consistent (for the consequence relation $\mathrel{|\!\sim}$). A formula is consistent for a model iff it is consistent for the consequence relation defined by the model, or equivalently iff there is a state in the model that satisfies α.*

We shall now define a basic operation on preferential models. Suppose M is a preferential model $\langle V, l, \prec \rangle$. For $s, t \in V$ we shall write $s \preceq t$ iff $s \prec t$ or $s = t$. Let α be a formula and $u \in V$ be a minimal element of $\widehat{\alpha}$. Let $\prec^u_\alpha$ be the strict partial order obtained from $\prec$ by making u a minimum of $\widehat{\alpha}$, i.e., $s \prec^u_\alpha t$ iff $s \prec t$ or $s \preceq u$ and there exists a state $w \in \widehat{\alpha}$ such that $w \preceq t$. The following lemma describes the properties of the construction described above.

Lemma 4 *The structure $M^u_\alpha \stackrel{\text{def}}{=} \langle V, l, \prec^u_\alpha \rangle$ is a preferential model. The consequence relation defined by M^u_α extends the consequence relation defined by M. In this model u is a minimum of $\widehat{\alpha}$. Both models have the same set of consistent formulas.*

Proof: It is easy to see that $\prec^u_\alpha$ is irreflexive and transitive. It is also easy to see that, under $\prec^u_\alpha$, u is a minimum of $\widehat{\alpha}$. We want to show now that, for any $\beta \in L$, the set $\widehat{\beta}$ is smooth, under $\prec^u_\alpha$. Let $s \in \widehat{\beta}$. Since $\widehat{\beta}$ is smooth under $\prec$, there is a state t, minimal under $\prec$ in $\widehat{\beta}$ such that $t \preceq s$. If t is still minimal in $\widehat{\beta}$ under $\prec^u_\alpha$, then we are done. If not, there is some state $v \in \widehat{\beta}$ such that $v \preceq u$ and $v \prec^u_\alpha s$. Since $\widehat{\beta}$ is smooth under $\prec$, there is a state w, minimal in $\widehat{\beta}$ under $\prec$ such that $w \prec v$. Since $w \prec u$, w must be minimal in $\widehat{\beta}$ also under $\prec^u_\alpha$. But $w \prec^u_\alpha s$. We have shown that $\widehat{\beta}$ is smooth under $\prec^u_\alpha$. To see that the consequence relation defined by M^u_α extends the one defined by M, just notice that, since $\prec^u_\alpha$ extends $\prec$, all minimal elements under the former are also minimal under the latter. Lastly, since M and M^u_α have exactly the same set of worlds and the same labeling function, they define exactly the same set of consistent formulas. ∎

Theorem 3 *Let* $\mathbf{K}$ *be a knowledge base and* $\alpha \mid\!\sim \beta$ *an assertion that is* not *preferentially entailed by* $\mathbf{K}$*. The formulas that are inconsistent for the preferential closure of* $\mathbf{K} \cup \{\alpha \mid\!\sim \neg\beta\}$ *are those that are inconsistent for the preferential closure of* $\mathbf{K}$*.*

Proof: Suppose that $\alpha \mid\!\sim \beta$ is not preferentially entailed by $\mathbf{K}$. Then, let $W = \langle S, l, \prec \rangle$ be the preferential model the existence of which is guaranteed by Theorem 1 and that defines $\mathbf{K}^p$. The model W does not satisfy $\alpha \mid\!\sim \beta$. There is therefore a minimal element $s \in S$ of $\widehat{\alpha}$ that does not satisfy β. Consider now the model $W' \stackrel{\text{def}}{=} W^s_\alpha$. By Lemma 4 this is a preferential model that satisfies all the assertions satisfied by W, therefore it satisfies all the assertions of $\mathbf{K}$. Since s is the only minimal element of $\widehat{\alpha}$, it satisfies $\mathbf{K} \cup \{\alpha \mid\!\sim \neg\beta\}$. Suppose γ is inconsistent for $(\mathbf{K} \cup \{\alpha \mid\!\sim \neg\beta\})^p$. Then it must be inconsistent for W'. By Lemma 4 it is inconsistent for W, therefore inconsistent for $\mathbf{K}^p$. ∎

Corollary 3 *Let* $\mathbf{K}$ *be a conditional knowledge base and* $\alpha \mid\!\sim \beta$ *a conditional assertion. The assertion* $\alpha \mid\!\sim \beta$ *is preferentially entailed by* $\mathbf{K}$ *iff the assertion* $\alpha \mid\!\sim$ **false** *is preferentially entailed by* $\mathbf{K} \cup \{\alpha \mid\!\sim \neg\beta\}$*.*

Proof: The *only if* part follows immediately from the soundness of the **And** rule. The *if* part, follows immediately from Theorem 3. ∎

2.6 The rank of a formula

In this section, we introduce a powerful tool for studying preferential entailment. Given a knowledge base, we shall attach an ordinal, its rank, to every formula. We shall prove an important result concerning those ranks, and, in particular, show that a knowledge base $\mathbf{K}$ and its preferential closure $\mathbf{K}^p$ define the same ranks.

Definition 10 *Let* $\mathbf{K}$ *be a conditional knowledge base (i.e., a set of conditional assertions) and* α *a formula. The formula* α *is said to be* exceptional *for* $\mathbf{K}$ *iff* $\mathbf{K}$ *preferentially entails the assertion* **true** $\mid\!\sim \neg\alpha$*. The conditional assertion* $\mathcal{A} \stackrel{\text{def}}{=} \alpha \mid\!\sim \beta$ *is said to be exceptional for* $\mathbf{K}$ *iff its antecedent* α *is exceptional for* $\mathbf{K}$*.*

The set of all assertions of $\mathbf{K}$ that are exceptional for $\mathbf{K}$ will be denoted by $E(\mathbf{K})$. Notice that $E(\mathbf{K}) \subseteq \mathbf{K}$. If all assertions of $\mathbf{K}$ are exceptional for $\mathbf{K}$, i.e., if $\mathbf{K}$ is equal to $E(\mathbf{K})$, we shall say that $\mathbf{K}$ is completely exceptional. The empty knowledge base is completely exceptional. Notice that, in the definition above, $\mathbf{K}$ may be replaced by its preferential closure $\mathbf{K}^p$.

Given a conditional knowledge base $\mathbf{K}$ (not necessarily finite), we shall now define by ordinal induction an infinite non-increasing sequence of subsets of $\mathbf{K}$. Let C_0 be equal to $\mathbf{K}$. For any successor ordinal $\tau + 1$, $C_{\tau+1}$ will be $E(C_\tau)$ and for any limit ordinal τ, C_τ is the intersection of all C_ρ for $\rho < \tau$. It is clear that, after some point on, all C's are equal and completely exceptional (they may be empty, but need not be so). We shall say that a formula α has rank τ (for $\mathbf{K}$) iff τ is the least ordinal for which α is not exceptional for C_τ. A formula that is exceptional for all C_τ's is said to have no rank. Notice that such a formula is exceptional for a totally exceptional knowledge base. The following is a fundamental lemma on preferential entailment. It says that, as far as preferential entailment is concerned, non-exceptional assertions cannot help deriving exceptional assertions. The notion of rank defined above proves to be a powerful tool for studying preferential entailment.

Lemma 5 *Let τ be an ordinal. Let $\mathbf{K}$ be a conditional knowledge base and $\mathcal{A}$ a conditional assertion whose antecedent has rank larger or equal to τ (or has no rank). Then $\mathcal{A}$ is preferentially entailed by C_0 iff it is preferentially entailed by C_τ.*

Proof: The *if* part follows from the fact that C_τ is a subset of C_0. The *only if* part is proved by induction on the length of the proof of $\mathcal{A}$ from C_0. If the proof has length one, i.e., $\mathcal{A}$ is obtained by **Reflexivity** or is an assertion of C_0, then the result is obvious. If the last step of the proof is obtained by **Right Weakening** or **And**, the result follows from a trivial use of the induction hypothesis. If the last step of the proof is obtained by **Left Logical Equivalence**, the result follows from the induction hypothesis and the fact that, if α and α' are logically equivalent then α and α' have the same rank. If the last step is a use of **Or**, and $\mathcal{A}$ is of the form $\alpha \vee \beta \mid\!\sim \gamma$ then just remark that the rank of the disjunction $\alpha \vee \beta$ is the smaller of the ranks of α and β. Both α and β have therefore a rank larger or equal to τ and one concludes by the induction hypothesis. If the last step is a use of **Cautious Monotonicity**, and $\mathcal{A}$ is of the form $\alpha \wedge \beta \mid\!\sim \gamma$, where $\alpha \mid\!\sim \beta$ and $\alpha \mid\!\sim \gamma$ are preferentially entailed (with short proofs) by C_0, let σ be the rank of α. By the induction hypothesis C_σ preferentially entails $\alpha \mid\!\sim \beta$. Since α is not exceptional for C_σ, we conclude that $\alpha \wedge \beta$ is not exceptional for C_σ, and therefore has rank σ. But $\alpha \wedge \beta$ has rank larger or equal to τ. Therefore $\tau \leq \sigma$. The formula α has rank larger or equal to τ and we may apply the induction hypothesis to conclude that both $\alpha \mid\!\sim \beta$ and $\alpha \mid\!\sim \gamma$ are preferentially entailed by C_ρ. ▮

Lemma 6 *Let $\mathbf{K}$ and $\mathbf{K}'$ be knowledge bases such that $\mathbf{K} \subseteq \mathbf{K}' \subseteq \mathbf{K}^p$. For any formula, the rank it is given by $\mathbf{K}'$ is equal to the rank it is given by $\mathbf{K}$.*

Proof: Using Lemma 5, one shows by ordinal induction that $C_\tau \subseteq C'_\tau \subseteq (C_\tau)^p$. ▮

The following definition will be useful in Section 5.6.

Definition 11 *A knowledge base $\mathbf{K}$ is said to be* admissible *iff all formulas that have no rank for $\mathbf{K}$ are inconsistent for $\mathbf{K}$.*

We shall immediately show that many knowledge bases are admissible.

Lemma 7 *If the preferential closure of* $\mathbf{K}$ *is defined by some well-founded preferential model, then* $\mathbf{K}$ *is admissible. In particular, any finite knowledge base is admissible.*

Proof: We have noticed, in Lemma 6 that ranks are stable under the the replacement of a knowledge base by its preferential closure. Let P be the preferential closure of $\mathbf{K}$. Suppose P is defined by some well-founded preferential model W. Suppose α has no rank. We must show that no state of W satisfies α. As noticed above, there is an ordinal τ such that C_τ is completely exceptional and α is exceptional for C_τ. We shall show that no state of W satisfies a formula that is exceptional for C_τ. Indeed, if there were such a state, there would be such a minimal state, s, since W is well-founded. But W is a model of C_τ and no state below s satisfy any antecedent of C_τ, since C_τ is totally exceptional. Therefore the preferential model consisting of s alone is a model of C_τ. But, in a model of C_τ, no minimal state satisfy a formula that is exceptional for C_τ. A contradiction. It follows now from Lemma 3 that any finite knowledge base is admissible. ∎

2.7 Computing preferential entailment

This section is devoted to the study of the computational complexity of preferential entailment. It is not needed in the sequel. We shall explain in Section 3.5 why preferential entailment is not the *right* notion of entailment to answer the question of the title, nevertheless preferential entailment is a central concept and it is therefore worthwhile studying its computational complexity. The results here are quite encouraging: the problem is in co-NP, i.e., in the same polynomial class as the problem of deciding whether a propositional formula is valid.

Lemma 8 *Let* $\mathbf{K}$ *be a finite conditional knowledge base and* $\alpha \mathrel{|\!\sim} \beta$ *a conditional assertion that is not preferentially entailed by* $\mathbf{K}$*. There is a finite totally (i.e., linearly) ordered preferential model of* $\mathbf{K}$ *no state of which satisfies* α *except the top state. This top state satisfies* α *and does not satisfy* β*.*

Proof: Let $\mathcal{L}' \subseteq \mathcal{L}$ be a logically finite language, large enough to contain α, β and all the assertions of $\mathbf{K}$. Let us now consider $\mathcal{L}'$ to be our language of reference. Clearly, $\alpha \mathrel{|\!\sim} \beta$ is not preferentially entailed by $\mathbf{K}$, since a proof over the smaller language is a proof over the larger language. By Theorem 1, there is a finite preferential model W (over $\mathcal{L}'$) of $\mathbf{K}$ that does not satisfy $\alpha \mathrel{|\!\sim} \beta$. In W, there is therefore a state s, minimal in $\widehat{\alpha}$, that satisfies α but does not satisfy β. Consider the submodel W' obtained by deleting all states of W that are not below or equal to s. It is clearly a finite preferential model of $\mathbf{K}$, with a top state that satisfies α but not β. Let V be obtained by imposing on the states of W' any total ordering that respects the partial ordering of W'. Since there are only finitely many states in V, the smoothness condition is verified and V is a preferential model (on $\mathcal{L}'$). It is a model of $\mathbf{K}$ but not of $\mathcal{A}$. Now we may extend the labeling function of V to the propositional variables of $\mathcal{L}$ that are not in $\mathcal{L}'$ any way we want, to get the model requested. Notice that the model obtained satisfies the smoothness condition because it is finite. ∎

Theorem 4 *There is a non-deterministic algorithm that, given a finite set* $\mathbf{K}$ *of conditional assertions and a conditional assertion* $\mathcal{A}$*, checks that* $\mathcal{A}$ *is not preferentially entailed by* $\mathbf{K}$*. The running time of this algorithm is polynomial in the size of* $\mathbf{K}$ *(sum of the sizes of its elements) and* $\mathcal{A}$*.*

Proof: Let $\mathbf{K}$ be $\{\gamma_i \mathrel{|\!\sim} \delta_i\}_{i=1}^{N}$. Let $I \subseteq \{1,\ldots,N\}$ be a set of indices. We shall define: $\varphi_I \stackrel{\text{def}}{=} \bigvee_{i \in I} \gamma_i$ and $\psi_I \stackrel{\text{def}}{=} \bigwedge_{i \in I} (\gamma_i \rightarrow \delta_i)$. A sequence is a sequence of pairs (I_i, f_i) for $i = 0,\ldots,n$, where $I_i \subseteq I$ and f_i is a world. Let α and β be in $\mathcal{L}$.

Definition 12 *A sequence* (I_i, f_i)*,* $i = 0,\ldots,n$*, is a witness for* $\alpha \mathrel{|\!\sim} \beta$ *(we mean a witness that* $\alpha \mathrel{|\!\sim} \beta$ *is not preferentially entailed by* $\mathbf{K}$ *) iff*

1. $f_k \models \psi_{I_k}, \ \ \forall k = 0,\ldots,n$
2. $f_k \models \varphi_{I_k}, \ \ \forall k = 0,\ldots,n-1$
3. $I_{k+1} = I_k \bigcap \{j \mid f_k \not\models \gamma_j\}, \ \ \forall k = 0,\ldots,n-1$
4. $I_0 = \{1,\ldots,N\}$
5. $f_k \not\models \alpha, \ \ \forall k = 0,\ldots,n-1$
6. $f_n \models \alpha \wedge \neg\beta$.

We must check that: witnesses are short and a conditional assertion has a witness iff it is not preferentially entailed by $\mathbf{K}$. For the first point, just remark that, for $k = 0,\ldots,n-1$ the inclusion $I_k \supset I_{k+1}$ is strict because of items 3 and 2. The length of the sequence is therefore bounded by the number of assertions in $\mathbf{K}$. But, each pair has a short description. For the second point, suppose first there is a witness for $\alpha \mathrel{|\!\sim} \beta$. Then the ranked model W consisting of worlds $f_0,\ldots,f_n$ where $f_k \prec f_{k+1}$ for $k = 0,\ldots,n-1$ satisfies $\mathbf{K}$ but not $\alpha \mathrel{|\!\sim} \beta$. That it does not satisfy $\alpha \mathrel{|\!\sim} \beta$ is clear from items 5 and 6. Let us check that W satisfies $\gamma_i \mathrel{|\!\sim} \delta_i$. If none of the f_k's, $k = 0,\ldots,n$ satisfies γ_i then W satisfies $\gamma_i \mathrel{|\!\sim} \eta$ for any η in $\mathcal{L}$. Suppose therefore that j is the smallest k for which $f_j \models \gamma_i$. We must show that $f_j \models \delta_i$. But, by items 4 and 3 $i \in I_j$ and by item 1, $f_i \models \delta_i$.

Suppose now that $\alpha \mathrel{|\!\sim} \beta$ is not preferentially entailed by some given finite $\mathbf{K}$. By Lemma 8, there is a finite linearly ordered model W of $\mathbf{K}$, no state of which satisfies α, except the top state that is labeled by a world m that satisfies

$\alpha \wedge \neg\beta$. Let $I_0 \stackrel{\rm def}{=} \{1, \ldots, N\}$. It is easy to see that (*remark* 1): if V is any preferential model of $\mathbf{K}$, for any set $I \subseteq I_0$, V satisfies $\varphi_I \mathrel{|\!\sim} \psi_I$. Let us now consider first the set $\widehat{\alpha \vee \varphi_{I_0}}$. It cannot be empty, therefore it has a unique minimal state. Let f_0 be the label of this state. We must consider two cases. First suppose that $f_0 \models \alpha$. Then f_0 is minimal in $\widehat{\alpha}$ and therefore must be m. In such a case (I_0, m) is a witness. The only thing to check is that item 1 is satisfied. Indeed either $m \models \varphi_{I_0}$ and we conclude by remark 1 or $\widehat{\varphi_{I_0}} = \emptyset$ and m satisfies none of the γ_i's. Let us deal now with the case $f_0 \not\models \alpha$. We shall build a sequence beginning by (I_0, m). Since m does not satisfy α, it must satisfy φ_{I_0}, which takes care of item 2. Remark 1 takes care of item 1. Let us now define $I_1 = I_0 \bigcap \{j \mid f_0 \not\models \gamma_j\}$. I_1 is strictly smaller than I_0. We may now consider the set $\widehat{\alpha \vee \varphi_{I_1}}$. It is not empty and therefore has a unique minimal element and we may, in this way, go on and build a proof for $\alpha \mathrel{|\!\sim} \beta$. ∎

Since it is clear that preferential non-entailment is at least as hard as satisfiability (consider assertions with antecedent **true**), we conclude that it is an NP-complete problem, i.e., that preferential entailment is co-NP-complete. A remark of J. Dix that will be explained at the end of Section 5.8 shows that preferential entailment is reducible to the computation of rational closure and that this reduction, when applied to Horn formulas, requires only the consideration of Horn formulas. It follows that, if we restrict ourselves to *Horn* assertions, computing preferential entailment has only polynomial complexity.

3 Rationality

3.1 Introduction

In this section we explain why not all preferential relations represent reasonable nonmonotonic inference procedures. We present some additional principles of nonmonotonic reasoning and discuss them. Those principles are structurally different from the rules of preferential reasoning, since they are not of the type: deduce some assertion from some other assertions. Sections 3.2 and 3.3 present weak principles. Some results are proven concerning those principles. Deeper results on those principles, found after a first version of this paper had been circulated, appear in [11]. Our central principle is presented in Section 3.4. Those principles were first described in [17] but the technical results presented here are new. In 3.5, the value of preferential entailment as an answer to the question of the title is discussed. Our conclusion is that it is not a satisfactory answer, since it does not provide us with a rational relation. Then, in Section 3.6, a restricted family of preferential models, the family of ranked models, is presented and a representation theorem is proved. The result is central to this paper but the proof of the representation theorem may be skipped on a first reading. The representation theorem appeared in [19]. The family of ranked models is closely related to, but different from, a family studied in [7] and Section 3.7 explains the differences.

3.2 Negation Rationality

In [17, Section 5.4], it was argued that not all preferential consequence relations represented reasonable inference operations. Three *rationality* properties were discussed there, and it was argued that all three were desirable. Those properties do not lend themselves to be presented as *Horn* rules (deduce the presence of an assertion in a relation from the presence of other assertions) but have the form: deduce the absence of an assertion from the absence of other assertions. All of them are implied by **Monotonicity**. The reader may find the discussion of [17] useful. Here technical results will be described. The first property considered is the following.

$$\frac{\alpha \wedge \gamma \not\mid\!\sim \beta \ , \ \alpha \wedge \neg\gamma \not\mid\!\sim \beta}{\alpha \not\mid\!\sim \beta} \qquad \textbf{(Negation Rationality)} \qquad (10)$$

Lemma 9 *There is a preferential relation that does not satisfy* **Negation Rationality***.*

Proof: Take a preferential model containing four states: $s_i, i = 0, \ldots 3$, with $s_0 \prec s_1$ and $s_2 \prec s_3$. Let the even states be the only states satisfying q and s_0 and s_3 be the only states satisfying p. One easily verifies that the consequence relation defined by this model is such that **true** $\mid\!\sim q$, but $p \not\mid\!\sim q$ and $\neg p \not\mid\!\sim q$. ∎

No semantic characterization of relations satisfying **Negation Rationality** is known. It has been shown in [17] that the consequence relation defined by *Circumscription* does not always satisfy **Negation Rationality**.

3.3 Disjunctive Rationality

The next property is the following.

$$\frac{\alpha \not\mid\!\sim \gamma \ , \ \beta \not\mid\!\sim \gamma}{\alpha \vee \beta \not\mid\!\sim \gamma} \qquad \textbf{(Disjunctive Rationality)} \qquad (11)$$

We may prove the following.

Lemma 10 *Any preferential relation that satisfies* **Disjunctive Rationality** *satisfies* **Negation Rationality***.*

Proof: Suppose $\alpha \wedge \gamma \not\mid\!\sim \beta$ and $\alpha \wedge \neg\gamma \not\mid\!\sim \beta$. By **Disjunctive Rationality**, we conclude that $\alpha \wedge \gamma \vee \alpha \wedge \neg\gamma \not\mid\!\sim \beta$. We conclude by **Left Logical Equivalence**. ∎

Lemma 11 *There is a preferential relation that satisfies* **Negation Rationality** *but does not satisfy* **Disjunctive Rationality***.*

Proof: Let us consider the following preferential model W. The model W has four states: a_0, a_1, b_0, b_1. The ordering is: $a_0 \prec a_1$ and $b_0 \prec b_1$. The language has three propositional variables: p, q and r. The two states a_1 and b_1 (the top states) are labeled with the same world that satisfies only p and q. State a_0 is labeled with the world that satisfies only p and r and the state b_0 with the world that satisfies only q and r. The preferential relation defined by W does not satisfy **Disjunctive Rationality** but satisfies **Weak Rationality**. For the first claim, notice that: $p \vee q \mathrel{|\!\sim}_W r$ but $p \mathrel{|\!\not\sim}_W r$ and $q \mathrel{|\!\not\sim}_W r$. For the second claim, suppose $\alpha \mathrel{|\!\sim}_W \gamma$, but $\alpha \wedge \beta \mathrel{|\!\not\sim}_W \gamma$. Then it must be the case that there is a minimal state of $\widehat{\alpha}$ that does not satisfy β and, above it, a state that is minimal in $\widehat{\alpha \wedge \beta}$. This last state must be labeled by a world that is the label of no minimal state of $\widehat{\alpha}$. Therefore, $\widehat{\alpha}$ must contain all four states of W, and $\widehat{\alpha \wedge \beta}$ must contain either the two top states alone or the two top states and one of the bottom states. In each case it is easy to see that $\alpha \wedge \neg\beta \mathrel{|\!\sim}_W \gamma$ since the minimal states of $\widehat{\alpha \wedge \neg\beta}$ are all also minimal in $\widehat{\alpha}$. ∎

No semantic characterization of relations satisfying **Disjunctive Rationality** was known at the time this paper was elaborated. M. Freund [10] has now provided a very elegant such characterization together with an alternative proof of our Theorem 5; the canonical model he builds is essentially the same as ours.

3.4 Rational Monotonicity

The last property is the following.

$$\frac{\alpha \wedge \beta \mathrel{|\!\not\sim} \gamma \quad , \quad \alpha \mathrel{|\!\not\sim} \neg\beta}{\alpha \mathrel{|\!\not\sim} \gamma} \qquad \textbf{(Rational Monotonicity)} \qquad (12)$$

This rule is similar to the thesis **CV** of conditional logic (see [25]). The reader is referred to [17, Section 5.4] for a discussion of our claim that reasonable consequence relations should satisfy **Rational Monotonicity**. Some researchers in Conditional Logic (J. Pollock in particular) have objected to **CV** as a valid thesis for (mainly subjunctive) conditionals. Echoes of this debate may be found in [12, end of Section 4.4]. The objections to **CV** that hold in the conditional logic framework do not hold for us, though their consideration is recommended to the reader. The most attractive feature of **Rational Monotonicity** is probably that it says that an agent should not have to retract any previous defeasible conclusion when learning about a new fact the negation of which was not previously derivable. In [32], K. Satoh aptly decided to call nonmonotonic reasoning that validates **Rational Monotonicity** *lazy*. The rule of **Rational Monotonicity** should be distinguished from the following rule, which is satisfied by any preferential relation.

$$\frac{\alpha \wedge \beta \mathrel{|\!\sim} \neg\gamma \quad , \quad \alpha \mathrel{|\!\not\sim} \neg\beta}{\alpha \mathrel{|\!\not\sim} \gamma} \qquad (13)$$

Definition 13 *A rational consequence relation is a preferential relation that satisfies* **Rational Monotonicity***.*

Two different representation theorems will be proved about rational relations, in Sections 3.6 and in Appendix B. The last one seems to provide evidence that reasonable inference procedures validate **Rational Monotonicity** and that all rational relations represent reasonable inference procedures.

Lemma 12 *A rational relation satisfies* **Disjunctive Rationality***.*

Proof: Suppose $\alpha \not\mathrel{|\!\sim} \gamma$ and $\beta \not\mathrel{|\!\sim} \gamma$. By **Left Logical Equivalence** we have $(\alpha \vee \beta) \wedge \alpha \not\mathrel{|\!\sim} \gamma$. If we have $\alpha \vee \beta \not\mathrel{|\!\sim} \neg\alpha$, then we could conclude by **Rational Monotonicity** that $\alpha \vee \beta \not\mathrel{|\!\sim} \gamma$. Suppose then that $\alpha \vee \beta \mathrel{|\!\sim} \neg\alpha$. If we had $\alpha \vee \beta \mathrel{|\!\sim} \gamma$, we would conclude by preferential reasoning that $\beta \not\mathrel{|\!\sim} \gamma$. ∎

Lemma 13 (David Makinson) *There is a preferential relation satisfying* **Disjunctive Rationality** *that is not rational.*

Proof: We shall build a preferential model that defines a consequence relation satisfying **Disjunctive Rationality** but not **Rational Monotonicity**. Let $\mathcal{L}$ be the propositional calculus on the three variables: p_0, p_1, p_2. Let $\mathcal{U}$ contain all propositional worlds on those variables. Let S contain three elements: s_i for $i = 0, 1, 2$ and $l(s_i)$ satisfy only p_i. The partial order $\prec$ is such that $s_1 \prec s_2$ and no other pair satisfies the relation. This defines a preferential model W. First we shall show that the consequence relation defined by W does not satisfy **Rational Monotonicity**. Indeed, we have both $p_0 \vee p_1 \vee p_2 \mathrel{|\!\sim}_W \neg p_2$ and $p_0 \vee p_1 \vee p_2 \not\mathrel{|\!\sim}_W p_1$. Nevertheless, we also have $\neg p_1 \wedge (p_0 \vee p_1 \vee p_2) \not\mathrel{|\!\sim}_W \neg p_2$. Let us show now that any preferential model that does not satisfy **Disjunctive Rationality** must have at least 4 states. Suppose $\alpha \vee \beta \mathrel{|\!\sim} \gamma$, but $\alpha \not\mathrel{|\!\sim} \gamma$ and $\beta \not\mathrel{|\!\sim} \gamma$. The last two assumptions imply the existence of states a and b, minimal in $\widehat{\alpha}$ and $\widehat{\beta}$ respectively and that do not satisfy γ, and therefore are not minimal in $\widehat{\alpha \vee \beta}$. Those states are different, since any state minimal in both $\widehat{\alpha}$ and $\widehat{\beta}$ would be minimal in $\widehat{\alpha \vee \beta}$. By the smoothness condition there must be a state a' minimal in $\widehat{\alpha \vee \beta}$ and such that $a' \prec a$. Clearly a' satisfies γ and does not satisfy α (since a is minimal in $\widehat{\alpha}$) but satisfies β. Similarly there must be a state b' minimal in $\widehat{\alpha \vee \beta}$ and such that $b' \prec b$ and b' satisfies γ and does not satisfy β, but satisfies α. It is left to show that all four states are different. We have already noticed that $a \neq b$. The states a' and b' satisfy γ and are therefore different from a and b. But b' satisfies α and a' does not and therefore $a' \neq b'$. ∎

3.5 Discussion of Preferential Entailment

We may now assess preferential entailment as a possible answer to the question of the title. Corollary 2 explains why the notion of preferential entailment cannot be the one we are looking for: the relation $\mathbf{K}^p$ can be any preferential relation and is not in general rational. For typical $\mathbf{K}$'s, $\mathbf{K}^p$ fails to satisfy a large number of instances of **Rational Monotonicity** and is therefore highly unsuitable. One particularly annoying instance of this is the following. Suppose a conditional knowledge base $\mathbf{K}$ contains one single assertion $p \mid\sim q$ where p and q are different propositional variables. Let r be a propositional variable, different from p and q. We intuitively expect the assertion $p \wedge r \mid\sim q$ to follow from $\mathbf{K}$. The rationale for that has been discussed extensively in the literature and boils down to this: since we have no information whatsoever about the influence of r on objects satisfying p it is sensible to assume that it has no influence and that there are normal p-objects that satisfy r. The normal $p \wedge r$-objects are therefore normal p-objects and have all the properties enjoyed by normal p-objects. Nevertheless it is easy to check that $p \wedge r \mid\sim q$ is not in $\mathbf{K}^p$. The problem lies, at least in part, with the fact that $\mathbf{K}^p$ is not rational, since any rational relation containing $p \mid\sim q$, must contain $p \wedge r \mid\sim q$ unless it contains $p \mid\sim \neg r$.

In conclusion, it seems that the set of conditional assertions entailed by $\mathbf{K}$ should be larger and more *monotonic* than the set $\mathbf{K}^p$. It should also be rational. This question will be brought up again in Section 5 and a solution will be proposed.

3.6 Ranked models and a representation theorem for rational relations

In this section a family of preferential models will be defined and it will be shown that the relations defined by models of this family are exactly the rational relations.

Lemma 14 *If $\prec$ is a partial order on a set V, the following conditions are equivalent.*

1. *for any $x, y, z \in V$ if $x \not\prec y$, $y \not\prec x$ and $z \prec x$, then $z \prec y$*
2. *for any $x, y, z \in V$ if $x \prec y$, then, either $z \prec y$ or $x \prec z$*
3. *for any $x, y, z \in V$ if $x \not\prec y$ and $y \not\prec z$, then $x \not\prec z$*
4. *there is a totally ordered set Ω (the strict order on Ω will be denoted by $<$) and a function $r : V \mapsto \Omega$ (the ranking function) such that $s \prec t$ iff $r(s) < r(t)$.*

The proof is simple and will not be given. A partial order satisfying any of the conditions of Lemma 14 will be called *modular* (this terminology is proposed in [12] as an extension of the notion of modular lattice of [13]).

Definition 14 *A ranked model W is a preferential model $\langle V, l, \prec \rangle$ for which the strict partial order $\prec$ is modular.*

Those models are called ranked since the effect of function r of property 4 of Lemma 14 is to rank the states: a state of smaller rank being more normal than a state of higher rank. We shall always suppose that a ranked model W comes equipped with a totally ordered set Ω and a ranking function r. Notice that we still require W to satisfy the smoothness condition. It is easy to see that for any subset T of V and any $t \in T$, t is minimal in T iff $r(t)$ is the minimum of the set $r(T)$. It follows that all minimal elements of T have the same image by r. The smoothness condition is then equivalent to the following: for any formula $\alpha \in L$, if $\widehat{\alpha}$ is not empty, the set $r(\widehat{\alpha})$ has a minimum. The smoothness condition is always verified if Ω is a well-ordered set. The reader may check that the preferential model W defined in the proof of Lemma 1 is ranked (it is even totally ordered). It follows that there are rational relations that are defined by no well-founded ranked model. The following is a soundness result.

Lemma 15 *If W is a ranked model, the consequence relation $\mathrel{|\!\sim}_W$ it defines is rational.*

Proof: It is enough to show that $\mathrel{|\!\sim}_W$ satisfies **Rational Monotonicity**. For this, the smoothness condition is not needed; it is needed, though, for the soundness of **Cautious Monotonicity**. Suppose W is a ranked model. We shall use the notations of Definition 14. Suppose also that $\alpha \mathrel{|\!\sim}_W \gamma$ and $\alpha \mathrel{|\!\not\sim}_W \neg\beta$. From this last assumption we conclude that there is a minimal element of $\widehat{\alpha}$ that satisfies β. Let $t \in V$ be such a state. Let $s \in V$ be a minimal element of $\widehat{\alpha \wedge \beta}$. Since $t \in \widehat{\alpha \wedge \beta}$, $t \not\prec s$ and $r(s) \leq r(t)$. But this implies that s is minimal in $\widehat{\alpha}$: any state u such that $u \prec s$ satisfies $r(u) < r(s)$ and therefore $r(u) < r(t)$ and $u \prec t$. Since $\alpha \mathrel{|\!\sim}_W \gamma$, $s \models\!\!\!\equiv \gamma$. ∎

We shall show now that the converse of Lemma 15 holds. We shall first mention four derived rules of preferential logic. In fact the first three of these rules are even valid in cumulative logic (see [17, Section 3]). Their proof (either proof-theoretic or model-theoretic) is straightforward and is omitted.

Lemma 16 *The following rules are derived rules of preferential logic:*

$$\frac{\alpha \mathrel{|\!\sim} \mathbf{false}}{\alpha \wedge \beta \mathrel{|\!\sim} \mathbf{false}} \tag{14}$$

$$\frac{\alpha \vee \beta \mathrel{|\!\sim} \neg\beta}{\alpha \mathrel{|\!\sim} \neg\beta} \tag{15}$$

$$\frac{\alpha \vee \beta \vee \gamma \mathrel{|\!\sim} \neg\alpha \wedge \neg\beta}{\beta \vee \gamma \mathrel{|\!\sim} \neg\beta} \tag{16}$$

$$\frac{\alpha \vee \beta \mathrel{|\!\sim} \neg\alpha}{\alpha \vee \beta \vee \gamma \mathrel{|\!\sim} \neg\alpha} \tag{17}$$

We shall now derive a property of rational relations.

Lemma 17 *If $\mid\!\sim$ is a rational relation, then the following rule is valid:*

$$\frac{\alpha \vee \gamma \mid\!\sim \neg\alpha \quad , \quad \beta \vee \gamma \not\mid\!\sim \neg\beta}{\alpha \vee \beta \mid\!\sim \neg\alpha} \tag{18}$$

Proof: From the first hypothesis, by Rule 17 one deduces $\alpha \vee \beta \vee \gamma \mid\!\sim \neg\alpha$. From the second hypothesis, by Rule 16 one deduces $\alpha \vee \beta \vee \gamma \not\mid\!\sim \neg(\alpha \vee \beta)$. If one applies now **Rational Monotonicity**, one gets the desired conclusion. ▮

For the completeness result, we proceed in the style of L. Henkin. Completeness proofs in this style have been used in conditional logics since [34]. Since a number of technical lemmas are needed, we have relegated them to Appendix A and state here the characterization theorem.

Theorem 5 *A binary relation $\mid\!\sim$ on $\mathcal{L}$ is a rational consequence relation iff it is the consequence relation defined by some ranked model. If the language $\mathcal{L}$ is logically finite, then every rational consequence relation is defined by some finite ranked model.*

Proof: The *if* part is Lemma 15. For the *only if* part, let $\mid\!\sim$ be a consequence relation satisfying the rules of **R**. The relation $\mid\!\sim$ defines a structure W as described in Appendix A. By Lemmas 41 and 36, W is a ranked model. We claim that, for any $\alpha, \beta \in L$, $\alpha \mid\!\sim \beta$ iff $\alpha \mid\!\sim_W \beta$. Suppose first $\alpha \mid\!\sim \beta$. By Lemma 42, if $< m, \beta >$ is minimal in $\widehat{\alpha}$, then m is normal for α. We conclude that $m \models \beta$ and $< m, \beta > \models\!\!\!\equiv \beta$. Therefore $\alpha \mid\!\sim_W \beta$. Suppose now that $\alpha \mid\!\sim_W \beta$. Let m be a normal world for α. By Corollary 7, the pair $< m, \alpha >$ is minimal in $\widehat{\alpha}$ and therefore $m \models \beta$. All normal worlds for α therefore satisfy β and Lemma 8 of [17] implies that $\alpha \mid\!\sim \beta$. For the last sentence of the theorem, notice that, if the language $\mathcal{L}$ is logically finite, the model W is finite. ▮

As we remarked just prior to Lemma 15, the theorem would not hold had we required models to be well-founded.

3.7 Comparison with Delgrande's work

The system of proof-rules and the models presented above may be compared with the results of J. Delgrande in [7, 8]. The general thrust is very similar but differences are worth noticing. A first difference is in the language used. Delgrande's language differs from this paper's in three respects: his work is specifically tailored to first-order predicate calculus, whereas this work deals with propositional calculus; he allows negation and disjunction of conditional assertions, which are not allowed in this paper; he allows nesting of conditional operators in the language, though his completeness result is formulated only for unnested formulas. Therefore Delgrande's central completeness result in [7], only shows that any proposition in which there is no nesting of conditional operators (let us call those propositions *flat*) that is valid has a proof from the axioms and rules of his system. But this proof may use propositions that are not flat. The completeness results reported here show that valid assertions have proofs that contain only flat assertions.

A second difference is that Delgrande's logical system is different from ours: Delgrande's logic **N** does not contain **Cautious Monotonicity**. Our class of ranked models is more restricted than his class of models: our models are required to obey the smoothness condition and Delgrande's are not. One may also notice that our logic enjoys the finite model property, but Delgrande's does not. This difference between our two logical systems may sound insignificant when one remarks that many instances of the rule of **Cautious Monotonicity** may be derived from **Rational Monotonicity**, and are therefore valid in Delgrande's system **N**. What we mean is that if $\alpha \mid\!\sim \beta$ and $\alpha \mid\!\sim \gamma$ then, if $\alpha \not\mid\!\sim \neg\beta$ one may conclude $\alpha \wedge \beta \mid\!\sim \gamma$ by **Rational Monotonicity** rather than by **Cautious Monotonicity**. But if $\alpha \mid\!\sim \neg\beta$, and therefore $\alpha \mid\!\sim$ **false** one cannot conclude. The Rule 14 is sound in preferential logic but not in Delgrande's logic. A proof will soon be given. We want to remark here that Rule (14) is very natural, since the meaning of $\alpha \mid\!\sim$ **false** is that if α is true than anything may be true. It therefore means that it is absolutely unthinkable that α be true. In such a case we would expect $\alpha \wedge \beta$ to be also absolutely unthinkable.

Let us show now that Rule (14) is not valid for Delgrande's structures. Consider the following structure. Let the set V consists of one infinite descending chain: $\prec$ is a total ordering. Suppose now that the top element of V is the only state that satisfies the propositional variable p. In this structure $\widehat{\mathbf{true}}$ is V and has no minimal element, therefore **true**$\mid\!\sim$ **false**. But $\widehat{p}$ consists only of the top element and has a minimal point and therefore $p \not\mid\!\sim$ **false**. We have shown that the Rule 14 is not valid for Delgrande's structures. This example also shows that Delgrande's logic does not posses the finite model property.

A third difference is that his definition of the set of conditional assertions entailed by a conditional knowledge base is different from the one presented here, at least at first sight.

4 Ranked entailment

4.1 Introduction

After having defined the family of relations and the family of models we are interested in, we proceed to study the notion of entailment provided by those models. Our main result is presented in 4.2. It is negative, in the sense that this entailment is equivalent to preferential entailment. A preliminary version of this result may be found in [19]. The collapsing of the two notions of entailment, as opposed to the two different classes of relations represented, sheds a new light on the results of [1]. Section 4.3 describes the probabilistic semantics given to preferential entailment by E. Adams in [1] and shows how the result of

Section 4.2 provides an alternative proof for Adams' results. The results of this section were contained in [19].

4.2 Ranked entailment is preferential entailment

In the discussion of Section 3.5, we expressed the wish that the set of assertions entailed by a conditional knowledge base $\mathbf{K}$ be rational and larger than $\mathbf{K}^p$. A natural candidate would be the set of all assertions that are satisfied in all ranked models that satisfy the assertions of $\mathbf{K}$. This is an intersection of rational relations. This proposal fails in a spectacular way. Problems with this proposal have been noted in [8, Section 4]. It is also easy to see that the intersection of rational relations may fail to be rational. Theorem 6 shows this failure to be total.

Lemma 18 *Let $\mathcal{E}$ be any preferential relation. There exists a rational extension of $\mathcal{E}$ for which a formula is inconsistent only if it is inconsistent for $\mathcal{E}$.*

Proof: Let us choose some enumeration of triples of formulas α, β and γ in which every triple appears an unbounded number of times. Let K_0 be equal to $\mathcal{E}$. At every step i we define K_{i+1} in the following way. Let α, β and γ be the triple enumerated at step i. Unless K_i contains the pair $\alpha \mid\sim \gamma$ but contains neither $\alpha \wedge \beta \mid\sim \gamma$ nor $\alpha \mid\sim \neg\beta$, we shall take K_{i+1} to be equal to K_i. If K_i satisfies the condition above, we shall take K_{i+1} to be the preferential closure of $K_i \cup \{\alpha \mid\sim \beta\}$. Notice that, by **Cautious Monotonicity**, $\alpha \wedge \beta \mid\sim \gamma$ will enter K_{i+1}. It is clear that the K_i's provide an increasing sequence of preferential extensions of $\mathcal{E}$. Let K_∞ be the union of all the K_i's. Clearly K_∞ is a preferential extension of $\mathcal{E}$. By construction, and since we took care of removing all counter-examples to the rule of **Rational Monotonicity**, K_∞ is a rational consequence relation. We claim that a formula α is inconsistent for K_∞ (i.e., $\alpha \mid\sim$ **false** is in K_∞) only if it is inconsistent already for $\mathcal{E}$. Indeed, if α is inconsistent for K_∞ it must be inconsistent for some K_i, but Theorem 3 shows that, by construction, all K_i's have the same inconsistent formulas. ∎

Theorem 6 *If the assertion $\mathcal{A}$ is satisfied by all* ranked *models that satisfy all the assertions of* $\mathbf{K}$*, then it is satisfied by all* preferential *such models.*

Proof: Let $\mathcal{A} \stackrel{\text{def}}{=} \delta \mid\sim \varepsilon$ be as in the hypotheses. Let $\mathcal{E}$ be the rational extension of the preferential closure of $\mathbf{K} \cup \{\delta \mid\sim \neg\varepsilon\}$, whose existence is asserted by Lemma 18. The assertion $\delta \mid\sim \varepsilon$ is in $\mathcal{E}$ since it is in any rational relation that extends $\mathbf{K}$, by Theorem 5. Since $\delta \mid\sim \neg\varepsilon$ is obviously in $\mathcal{E}$, we conclude that δ is inconsistent for $\mathcal{E}$ and therefore inconsistent for the preferential closure of $\mathbf{K} \cup \{\delta \mid\sim \neg\varepsilon\}$. By Corollary 3, $\mathcal{A}$ is preferentially entailed by $\mathbf{K}$. ∎

4.3 Comparison with Adams' probabilistic entailment

In this section we shall show that ranked models are closely related to Adams' probabilistic entailment described in [1]. Theorem 6, then, provides an alternative proof of Adams' axiomatic characterization of probabilistic entailment. There are some technical differences between Adams' framework and ours since Adams insists on allowing formulas as conditional assertions: for him the formula α is a synonym for **true** $\mid\sim \alpha$. We also insist on studying infinite knowledge bases whenever possible, where Adams restricts himself to finite knowledge bases.

A probability assignment for the language $\mathcal{L}$ is a probability measure on $\mathcal{L}$ yielded by some probability measure given on $\mathcal{U}$. E. Adams proposed the following definitions.

Definition 15 (Adams) *A probability assignment p for the language $\mathcal{L}$ is said to be* proper *for a conditional assertion $\alpha \mid\sim \beta$ iff $p(\alpha) > 0$. It is proper for a set of conditional assertions iff it is proper for each element. If p is proper for $\mathcal{A} \stackrel{\text{def}}{=} \alpha \mid\sim \beta$, we shall use $p(\mathcal{A})$ to denote the conditional probability $p(\beta \mid \alpha)$.*

Definition 16 (Adams) *Let $\mathbf{K}$ be a set of conditional assertions. We shall say that $\mathbf{K}$ is* probabilistically consistent *if and only if for any real number $\epsilon > 0$ there exists a probability assignment p for $\mathcal{L}$ that is proper for $\mathbf{K}$ and such that, for all $\mathcal{A}$ in $\mathbf{K}$, one has $p(\mathcal{A}) \geq 1 - \epsilon$.*

Definition 17 *Let $\mathbf{K}$ be a set of conditional assertions and $\mathcal{A}$ a conditional assertion. We shall say that $\mathbf{K}$* probabilistically entails *$\mathcal{A}$ iff for all $\epsilon > 0$ there exists $\delta > 0$ such that for all probability assignments p for $\mathcal{L}$ which are proper for $\mathbf{K}$ and $\mathcal{A}$, if $p(\mathcal{B}) \geq 1 - \delta$ for all $\mathcal{B}$ in $\mathbf{K}$, then $p(\mathcal{A}) \geq 1 - \epsilon$.*

In [1], Adams studies extensively the relations between the two notions of probabilistic consistency and probabilistic entailment, at least for finite sets of conditional assertions. Here we shall only show the fundamental relation that exists between Adams' notions and ours. First, we shall make three easy remarks. The first one concerns only probabilistic notions and was claimed by Adams for finite knowledge bases but is true in general.

Lemma 19 *A set $\mathbf{K}$ of conditional assertions is probabilistically inconsistent iff it probabilistically entails any conditional assertion.*

Our second remark provides a first link between probabilistic notions and the notions introduced in this paper. It is essentially the soundness part of Adams' soundness and completeness result (see beginning of proof of 4.2 at page 62 of [1]). This is the easy direction.

Lemma 20 *Any conditional assertion preferentially entailed by $\mathbf{K}$ is probabilistically entailed by $\mathbf{K}$.*

Our third remark is the following.

Lemma 21 *If the conditional assertion $\alpha \mid\!\sim \beta$ is in* **K** *and $\alpha \mid\!\sim$* **false** *is preferentially entailed by* **K** *then* **K** *is probabilistically inconsistent.*

Proof: Under the assumptions of the lemma, Lemma 20 shows that $\alpha \mid\!\sim$ **false** is probabilistically entailed by **K**. But for any probability assignment p that is proper for $\alpha \mid\!\sim \beta$, $p(\alpha \mid\!\sim \mathbf{false})$ is defined and equal to 0. Since **K** probabilistically entails $\alpha \mid\!\sim$ **false** we conclude that there is an $\epsilon > 0$ such that no probability assignment that is proper for **K** and $\alpha \mid\!\sim \beta$ gives probabilities larger than $1 - \epsilon$ to all assertions of **K**. Since any probability assignment that is proper for **K** is also proper for $\alpha \mid\!\sim \beta$ the conclusion is proved. ∎

We shall now prove the converse of Lemma 20 in the case **K** is finite and probabilistically consistent. The basic remark is the following. Suppose W is a finite (i.e., the set S of states is finite) ranked model. Let $\epsilon > 0$ be some real number. We shall describe a probability measure p_ϵ on S. The first principle that will be used in defining p_ϵ is that all states of the same rank will have equal probabilities. The second principle is that the weight of the set of all states of rank n, w_n will be such that $\frac{w_{n+1}}{w_n} = \epsilon$. The intuitive meaning of this choice (since ϵ will approach zero) is that normal states are more probable than exceptional states. There is clearly exactly one probability measure satisfying both principles above, for any given finite ranked model. The probability measure p_ϵ, defined on states, yields a probability measure on formulas. It is clear that a formula α has probability zero under p_ϵ iff α is inconsistent in W, i.e., $\alpha \mid\!\sim_W$ **false**. Suppose α is consistent. Let us consider the conditional probability of β given α, which is well defined. If $\alpha \mid\!\sim_W \beta$ then this conditional probability is larger than $1 - \epsilon - \epsilon^2 - \epsilon^3 - \ldots$ and therefore approaches one when ϵ approaches zero. On the other hand, if $\alpha \not\mid\!\sim_W \beta$, then this conditional probability cannot exceed $1 - \frac{1}{m}$ where m is the number of states at the rank which is minimal for α. It is therefore bounded away from 1 when ϵ approaches 0.

Lemma 22 *Let a finite probabilistically consistent knowledge base* **K** *be given and suppose $\mathcal{A} \stackrel{\mathrm{def}}{=} \alpha \mid\!\sim \beta$ is not preferentially entailed by* **K**. *Then,* **K** *does not probabilistically entail $\mathcal{A}$.*

Proof: Let **K** and $\mathcal{A}$ be as described in the lemma. Let $\mathcal{L}'$ be some logically finite sublanguage of $\mathcal{L}$ that contains α, β and all propositions appearing in **K**. Relative to $\mathcal{L}'$, the hypotheses of the lemma are still true. By Lemmas 18 and Theorem 3, there is a rational relation $\mathcal{E}$ that contains **K** and $\alpha \mid\!\sim \neg\beta$ and for which a formula is inconsistent only if it is inconsistent for $\mathbf{K}^p$. By Theorem 5, there is a finite ranked model, W that satisfies $\alpha \mid\!\sim \neg\beta$ and all assertions of **K** but whose inconsistent formulas are exactly those of $\mathbf{K}^p$. Since α is not inconsistent for $\mathbf{K}^p$, the model W does not satisfy $\alpha \mid\!\sim \beta$. Let W' be the model obtained by extending the labeling function of W to the full language $\mathcal{L}$ in an arbitrary way. We shall now apply the construction of p_ϵ described above on the model W'.

Using the model W' and a sequence of ϵ's approaching zero, we define a sequence of probability measures p_ϵ. Let us show that all assignments p_ϵ are proper for $\mathbf{K}$ and $\mathcal{A}$. If $\gamma \in L'$, the assignment p_ϵ gives zero probability to γ iff γ is inconsistent in W, i.e., inconsistent for $\mathbf{K}^p$. But $\mathbf{K}$ is probabilistically consistent and, by Lemma 21, p_ϵ is proper for $\mathbf{K}$. Since W' does not satisfy $\mathcal{A}$, its antecedent cannot be inconsistent in W' and p_ϵ is proper for $\mathcal{A}$ too. When ϵ approaches zero, the conditional probabilities corresponding to each assertion of $\mathbf{K}$ approach 1 and the conditional probability corresponding to A is bounded away from 1. ∎

That the result cannot be extended to infinite sets of conditional assertions follows from Adams' remark that his notion of probabilistic consistency does not enjoy the compactness property and from Corollary 1. Adams' example [1, pages 51–52] is closely related to the construction of Lemma 1. The results of Adams presented in this section have been interpreted, in particular by [26], to mean that probabilistic semantics validate preferential reasoning. We certainly agree. But the results that will be presented now show, in our opinion, that probabilistic semantics support the claim that inference procedures should not only be preferential but also rational. Indeed we show, in Appendix B, that some very natural probabilistic models always define rational relations and that, when the language $\mathcal{L}$ is countable, all rational relations may be defined by such models. Those models are non-standard probability spaces, in the sense of A. Robinson. Since no use of those models will be made in the paper, their treatment has been relegated to an appendix.

5 The rational closure of a conditional knowledge base

5.1 Introduction

So far, we have argued for Thesis 1 and gathered much knowledge about rational relations, showing in particular that there is no obvious way to define a notion of closure satisfying Thesis 1. In this section we shall show that there is a natural notion of closure (called rational closure) that satisfies Thesis 1. We shall study it and prove that it possesses many very elegant mathematical properties. We shall, then, evaluate the value of rational closure as an answer to the question of the title. In conclusion, we shall propose Thesis 2, that claims that any satisfactory answer is a superset of rational closure. In other terms we think that any reasonable system should endorse any assertion contained in the rational closure, but it may also endorse some additional assertions. At present, we do not know of any natural construction satisfying Thesis 1 other than rational closure.

A first possible answer is rejected in 5.2. This result appeared in [18]. The remainder of the paper describes rational closure. In 5.3 a partial ordering between rational relations is defined, which captures the notion of a relation being preferable to (i.e., smaller, less adventurous, more reasonable than) another one. The rational closure of a knowledge base is then defined in 5.4 as the rational extension of a knowledge base that is preferable in the ordering defined in Section 5.3 to all other rational extensions. Not every knowledge base has a rational closure, but in Section 5.6 it will be shown that any admissible (see Definition 11) knowledge base has a rational closure. By Lemma 7, then, any finite knowledge base has a rational closure. We claim that the rational closure of a knowledge base, when it exists, provides a reasonable answer to the question of the title. Global properties of the operation of rational closure are described in 5.5. These results, concerning the global behavior of a nonmonotonic inference operation, are the first of their kind. In 5.6 an algorithmic construction of the rational closure of an admissible knowledge base is described. This algorithmic description essentially replaces and improves upon the proof-theoretic description of [18]. A corrected and generalized model-theoretic construction, first described in [18] is proposed in 5.7. Section 5.8 presents an algorithm to compute the rational closure of a finite knowledge base and discusses complexity issues. Section 5.9 discusses the appeal of rational closure and provides some examples. Section 6 concludes by considering topics for further research. In [27] J. Pearl proposes his own version of the rational closure construction that had been described in [18].

5.2 Perfect extensions

All that has been done so far does not allow us to give a satisfactory answer to the question of the title. Let $\mathbf{K}$ be a set of conditional assertions. We would like to define a consequence relation $\overline{\mathbf{K}}$, the rational closure of $\mathbf{K}$, that contains all the conditional assertions that we intuitively expect to follow from $\mathbf{K}$. At this point the reader should be convinced that $\overline{\mathbf{K}}$ should be a rational consequence relation that extends $\mathbf{K}$. Any such relation obviously also extends $\mathbf{K}^p$. It seems that we would also like this rational extension of $\mathbf{K}$ to be as small as possible. Unfortunately Theorem 6 shows that the intersection of all rational extensions of $\mathbf{K}$ is exactly $\mathbf{K}^p$ and therefore not in general rational and highly unsuitable as shown in Section 3.5. There is obviously a maximal such extension: the full consequence relation, (i.e., $\alpha \mid\!\sim \beta$ for all α, β in $\mathcal{L}$) but this is certainly not the one we are looking for. Can we find out a number of properties that we would like $\overline{\mathbf{K}}$ to possess, in order to, at least, narrow the field of possibilities? We shall look both for *local* properties of $\overline{\mathbf{K}}$ with respect to $\mathbf{K}$ and for *global* properties of the mapping $\mathbf{K} \mapsto \overline{\mathbf{K}}$. The sequel will present a proposal for the definition of $\overline{\mathbf{K}}$ and proofs that it enjoys both local and global (in particular a strong form of cumulativity) properties.

If $\mathbf{K}^p$ happens to be rational, then we probably have no reason to look further and should take $\overline{\mathbf{K}}$ to be equal to $\mathbf{K}^p$. If $\mathbf{K}^p$ is not rational, then there is an assertion $\alpha \mathrel{|\!\sim} \beta$ in $\mathbf{K}^p$, and a formula γ such that neither $\alpha \wedge \gamma \mathrel{|\!\sim} \beta$ nor $\alpha \mathrel{|\!\sim} \neg\gamma$ are in $\mathbf{K}^p$. It seems that the right thing to do, in most such cases, is to introduce $\alpha \wedge \gamma \mathrel{|\!\sim} \beta$ in $\overline{\mathbf{K}}$. One may try to require that any assertion in $\overline{\mathbf{K}} - \mathbf{K}^p$ be of the form $\alpha \wedge \gamma \mathrel{|\!\sim} \beta$ where $\alpha \mathrel{|\!\sim} \beta$ is in $\mathbf{K}^p$, i.e., that any assertion in $\overline{\mathbf{K}} - \mathbf{K}^p$ have *support* in $\mathbf{K}^p$. It will be shown that this may well be impossible. Let us encapsulate this idea in definitions.

Definition 18 *An assertion $\alpha \mathrel{|\!\sim} \beta$ is said to be supported by (or in) $\mathbf{K}^p$ iff there is a formula γ such that $\alpha \models \gamma$ and $\gamma \mathrel{|\!\sim} \beta$ is in $\mathbf{K}^p$.*

Definition 19 *A rational extension $\mathbf{K}'$ of $\mathbf{K}$ is called perfect iff every assertion of $\mathbf{K}'$ is supported by $\mathbf{K}^p$.*

We may present the following disappointing result.

Lemma 23 *There is a finite conditional knowledge base that has no rational perfect extension.*

Proof: Let $\mathcal{L}$ be the set of all propositional formulas built out of the set of four propositional variables: $\{a, b, c, d\}$. Let W be the preferential model with three states: $\{s, t, u\}$, in which $s \prec t$ (and this is the only pair in the relation $\prec$) and s satisfies only a, t satisfies only b and u satisfies only c and d. Let $\mathbf{K}$ be the set of all conditional assertions satisfied in W. We claim that $\mathbf{K}$ has no rational perfect extension. Notice, first, that W satisfies $a \vee b \mathrel{|\!\sim} \neg b$. This assertion is therefore in $\mathbf{K}$. Any ranked model satisfying $a \vee b \mathrel{|\!\sim} \neg b$ must satisfy at least one of the following two assertions: $a \vee c \mathrel{|\!\sim} \neg c$ or $b \vee c \mathrel{|\!\sim} \neg b$. Any rational extension of $\mathbf{K}$ must therefore contain one of the two assertions above. But $a \vee c \mathrel{|\!\sim} \neg c$ has clearly no support in $\mathbf{K}^p$ and therefore any perfect rational extension of $\mathbf{K}$ must contain: $b \vee c \mathrel{|\!\sim} \neg b$. But W satisfies $c \mathrel{|\!\sim} d$ and any ranked model satisfying both $b \vee c \mathrel{|\!\sim} \neg b$ and $c \mathrel{|\!\sim} d$ must also satisfy $b \vee c \mathrel{|\!\sim} d$. Any perfect rational extension of $\mathbf{K}$ must therefore contain this last formula but it clearly lacks support in $\mathbf{K}^p$. We conclude that $\mathbf{K}$ has no perfect rational extension. ∎

It is therefore reasonable to look for less than perfect extensions. Let us first examine perfection concerning two special kinds of formulas. The following is easily proved.

Lemma 24 *An assertion of the form $\alpha \mathrel{|\!\sim}$ **false** is supported by $\mathbf{K}^p$ iff it is in $\mathbf{K}^p$. An assertion of the form **true** $\mathrel{|\!\sim} \alpha$ is supported by $\mathbf{K}^p$ iff it is in $\mathbf{K}^p$.*

We shall propose a construction of $\overline{\mathbf{K}}$ such that $\overline{\mathbf{K}}$ does not contain any formula of the form $\alpha \mathrel{|\!\sim}$ **false** or of the form **true** $\mathrel{|\!\sim} \alpha$ that is not in $\mathbf{K}^p$.

5.3 Ordering rational relations

In this section we shall define a strict partial ordering between rational relations. This ordering captures the notion of a relation being preferable to, i.e., less adventurous than another one. An intuitive explanation will be given immediately after the definition. For the rest of this section we shall write $\alpha < \beta$ *for (or in)* K to mean *the assertion* $\alpha \vee \beta \mid\!\sim \neg\beta$ *is in* K. We shall write $\alpha \leq \beta$ for (or in) K when it is *not* the case that $\beta{<}\alpha$ in K.

Definition 20 *Let K_0 and K_1 be two rational consequence relations. We shall say that K_0 is preferable to K_1 and write $K_0 \prec K_1$ iff:*

1. *there exists an assertion $\alpha \mid\!\sim \beta$ in $K_1 - K_0$ such that for all γ such that $\gamma < \alpha$ for K_0, and for all δ such that $\gamma \mid\!\sim \delta$ is in K_0, we also have $\gamma \mid\!\sim \delta$ in K_1, and*
2. *for any γ, δ if $\gamma \mid\!\sim \delta$ is in $K_0 - K_1$ there is an assertion $\rho \mid\!\sim \eta$ in $K_1 - K_0$ such that $\rho < \gamma$ for K_1.*

The intuitive explanation behind Definition 20 is the following. Suppose two agents, who agree on a common knowledge base, are discussing the respective merits of two rational relations K_0 and K_1. A typical attack would be: *your relation contains an assertion, $\alpha \mid\!\sim \beta$, that mine does not contain* (and therefore contains unsupported assertions). A possible defense against such an attack could be: *yes, but your relation contains an assertion $\gamma \mid\!\sim \delta$ that mine does not, and you yourself think that γ refers to a situation that is more usual than the one refered to by α.* Such a defense must be accepted as valid. Definition 20 exactly says that the proponent of K_0 has an attack that the proponent of K_1 cannot defend against (this is part 1) but that he (i.e., the proponent of K_0) may find a defense against any attack from the proponent of K_1 (this is part 2 of the definition).

Lemma 25 *The relation $\prec$ between rational consequence relations is irreflexive and transitive.*

Proof: Irreflexivity follows immediately from Condition 1. For transitivity, let us suppose that $K_0 \prec K_1$, with $\alpha \mid\!\sim \beta$ the witness promised by Condition 1 and that $K_1 \prec K_2$ with $\gamma \mid\!\sim \delta$ as a witness. Our first step will be to show that there exists an assertion $\varepsilon \mid\!\sim \zeta$ in $K_2 - K_0$ such that $\varepsilon \leq \alpha$ in K_0 and $\varepsilon \leq \gamma$ in K_1. We shall have to consider many different cases.

1. Suppose $\gamma < \alpha$ in K_2.
   (a) If $\gamma < \alpha$ is not in K_0, then $\gamma < \alpha$ is a suitable $\varepsilon \mid\!\sim \zeta$.
   (b) If $\gamma < \alpha$ is in K_0, then $\gamma \mid\!\sim \delta$ is a suitable $\varepsilon \mid\!\sim \zeta$, since if it were in K_0 it would be in K_1.

2. Suppose therefore that $\gamma < \alpha$ is not in K_2, i.e., for K_2, $\alpha \leq \gamma$.

   (a) If $\gamma < \alpha$ is in K_1, then it is in $K_1 - K_2$ and there is an assertion $\xi \mathrel{|\!\sim} \eta$ in $K_2 - K_1$ such that $\xi < \gamma \vee \alpha$ in K_2.

      i. If $\xi < \gamma \vee \alpha$ is not in K_0, then it is a suitable $\varepsilon \mathrel{|\!\sim} \zeta$.

      ii. If $\xi < \gamma \vee \alpha$ is in K_0, then $\xi < \alpha$ in K_0 and we have both that $\xi < \gamma \vee \alpha$ in K_1 and that $\xi \mathrel{|\!\sim} \eta$ cannot be in K_0, otherwise it would be in K_1. We conclude that $\xi \mathrel{|\!\sim} \eta$ is a suitable $\varepsilon \mathrel{|\!\sim} \zeta$.

   (b) Suppose therefore that $\gamma < \alpha$ is not in K_1, i.e., for K_1, like for K_2, $\alpha \leq \gamma$.

      i. If $\alpha \mathrel{|\!\sim} \beta$ is in K_2, then it is a suitable $\varepsilon \mathrel{|\!\sim} \zeta$.

      ii. If $\alpha \mathrel{|\!\sim} \beta$ is not in K_2, then it is in $K_1 - K_2$ and there is a $\xi \mathrel{|\!\sim} \eta$ in $K_2 - K_1$ such that $\xi < \alpha$ in K_2.

         A. If $\xi < \alpha$ is not in K_0, then it is a suitable $\varepsilon \mathrel{|\!\sim} \zeta$ since, in K_1, $\xi \vee \alpha \leq \alpha \leq \gamma$.

         B. If $\xi < \alpha$ is in K_0 then $\xi \mathrel{|\!\sim} \eta$ is a suitable $\varepsilon \mathrel{|\!\sim} \zeta$ since $\xi \mathrel{|\!\sim} \eta$ cannot be in K_0, otherwise it would be in K_1.

We have now proved the existence of an assertion $\varepsilon \mathrel{|\!\sim} \zeta$ with the desired properties. Let us proceed to the proof that $K_0 \prec K_2$. For property 1, we claim that $\varepsilon \mathrel{|\!\sim} \zeta$ provides a suitable witness. It is indeed in $K_2 - K_0$ by construction. Suppose now that $\xi < \varepsilon$ in K_0. Then $\xi < \alpha$ in K_0 and therefore $\xi < \varepsilon$ in K_1. Therefore $\xi < \gamma$ in K_1. If $\xi \mathrel{|\!\sim} \eta$ is in K_0, then it must be in K_1 since $\xi < \alpha$ in K_0 and also in K_2 since $\xi < \gamma$ in K_1. This concludes the verification of Condition 1.

For Condition 2, suppose that $\varphi \mathrel{|\!\sim} \theta$ is in $K_0 - K_2$. We have to find a $\xi \mathrel{|\!\sim} \eta$ in $K_2 - K_0$ such that $\xi < \varphi$ in K_2. We shall consider a number of different cases.

1. If $\varepsilon < \varphi$ in K_2, then $\varepsilon \mathrel{|\!\sim} \zeta$ is a suitable $\xi \mathrel{|\!\sim} \eta$.

2. Suppose then that $\varepsilon < \varphi$ is not in K_2, i.e., $\varphi \leq \varepsilon$ for K_2.

   (a) Suppose, first that $\varepsilon < \varphi$ is in K_1, therefore in $K_1 - K_2$. There is then an assertion $\rho \mathrel{|\!\sim} \tau$ in $K_2 - K_1$ such that $\rho < \varepsilon \vee \varphi$ in K_2.

      i. If $\rho < \varepsilon \vee \varphi$ is in K_0, then $\rho < \alpha$ in K_0 and we conclude that $\rho \mathrel{|\!\sim} \tau$ is not in K_0, otherwise it would be in K_1. We conclude that $\rho \mathrel{|\!\sim} \tau$ is a suitable $\xi \mathrel{|\!\sim} \eta$.

      ii. If $\rho < \varepsilon \vee \varphi$ is not in K_0, then it is a suitable $\xi \mathrel{|\!\sim} \eta$.

   (b) Suppose, then, that $\varepsilon < \varphi$ is not in K_1, i.e., for K_1, $\varphi \leq \varepsilon$.

      i. If $\varepsilon < \varphi$ is in K_0, then it is in $K_0 - K_1$. Therefore there is an assertion $\rho \mathrel{|\!\sim} \tau$ in $K_1 - K_0$ such that $\rho < \varepsilon \vee \varphi$ in K_1. But then $\rho < \gamma$ in K_1 and we conclude that $\rho \mathrel{|\!\sim} \tau$ is in K_2 and that $\rho < \varphi$ in K_2. The assertion $\rho \mathrel{|\!\sim} \tau$ is a suitable $\xi \mathrel{|\!\sim} \eta$.

ii. Suppose, then that, on the contrary, $\varepsilon < \varphi$ is not in K_0, i.e., $\varphi \leq \varepsilon$ in K_0, as in K_1 and K_2.

A. Suppose first that $\varphi \mathrel{|\!\sim} \theta$ is in K_1, therefore in $K_1 - K_2$. Then, there is an assertion $\rho \mathrel{|\!\sim} \tau$ in $K_2 - K_1$, such that $\rho < \varphi$ in K_2. There are two cases. If $\rho < \varphi$ is in K_0, then $\rho < \varepsilon$ in K_0, and $\rho \mathrel{|\!\sim} \tau$ is not in K_0, otherwise it would be in K_1, since $\rho < \varepsilon \leq \alpha$ in K_0. The assertion $\rho \mathrel{|\!\sim} \tau$ is a suitable $\xi \mathrel{|\!\sim} \eta$. If, on the other hand $\rho < \varphi$ is not in K_0, then it is in $K_2 - K_0$, and it is a suitable $\xi \mathrel{|\!\sim} \eta$.

B. Suppose now that $\varphi \mathrel{|\!\sim} \theta$ is not in K_1, therefore in $K_0 - K_1$. There is an assertion $\rho \mathrel{|\!\sim} \tau$ in $K_1 - K_0$, such that $\rho < \varphi$ in K_1. But $\rho < \varphi \leq \varepsilon \leq \gamma$ in K_1 and since $\rho \mathrel{|\!\sim} \tau$ is in K_1 it must be in K_2. Also, since $\rho < \varphi$ is in K_1, it must be in K_2. We see that $\rho \mathrel{|\!\sim} \tau$ is a suitable $\xi \mathrel{|\!\sim} \eta$.

▮

5.4 Definition of rational closure

We may now define the rational closure of a knowledge base.

Definition 21 *Let* $\mathbf{K}$ *be an arbitrary knowledge base. If there is a rational extension* $\overline{\mathbf{K}}$ *of* $\mathbf{K}$ *that is preferable to all other rational extensions of* $\mathbf{K}$*, then* $\overline{\mathbf{K}}$ *will be called the rational closure of* $\mathbf{K}$*.*

Notice first that the rational closure of a knowledge base is unique, if it exists, since *preference* is a partial ordering. Notice then that there are knowledge bases that do not have a rational closure. Example 1 will show this. In Section 5.6 we shall show that admissible knowledge bases, including all finite knowledge bases, have a rational closure.

Example 1 Let $\mathcal{L}$ be the propositional calculus built upon the variables p_n where n is an arbitrary *integer* (i.e., positive or negative). Let N be the knowledge base that contains all assertions of the form $p_n \mathrel{|\!\sim} p_{n+2}$ and of the form $p_n \mathrel{|\!\sim} \neg p_{n-2}$ for all integers n. We shall show that N has no rational closure.

We shall first prove a lemma about invariance of the operation of rational closure under renaming of the proposional variables. This lemma is of independent interest.

Definition 22 *1. A renaming of the propositional calculus* $\mathcal{L}$ *is a bijection of the propositional variables.*

2. Let f *be a renaming of* $\mathcal{L}$*. The formula obtained from* α *by substituting* $f(p)$ *for the propositional variable* p *will be denoted by* $f(\alpha)$*.*

3. *Let f be as above and $\alpha \mathrel{|\!\sim} \beta$ a conditional assertion. The assertion $f(\alpha) \mathrel{|\!\sim} f(\beta)$ will be denoted by $f(\alpha \mathrel{|\!\sim} \beta)$.*

4. *Let f be as above and K a consequence relation. The relation $f(K)$ will be defined by $f(K) = \{f(\alpha \mathrel{|\!\sim} \beta) \mid \alpha \mathrel{|\!\sim} \beta \in K\}$.*

Lemma 26 (Invariance under renaming) *Let f be a renaming of $\mathcal{L}$.*

1. *Let K_0 and K_1 be rational consequence relations. Then $K_0 \prec K_1$ iff $f(K_0) \prec f(K_1)$.*

2. *Let K be a consequence relation and $\overline{K}$ its rational closure ,then $f(\overline{K})$ is the rational closure of $f(K)$*

3. *Let K be a consequence relation which is invariant under f, namely $f(K) = K$,then its rational closure (if it exists) is invariant under f.*

Proof: The proof is immediate from the definitions, noting that f is also a bijection of the set of all consequence relations. ∎

Lemma 27 *The knowledge base N defined above has no rational closure.*

Proof: We shall reason by contradiction. Suppose R is the rational closure of N. From Lemma 30 in the sequel (the proof of which does not depend on the present lemma), we know that there is no assertion of the form $\alpha \mathrel{|\!\sim}$ **false** in R that is not in N^p. Using a construction very similar to the one used in the proof of Lemma 1, one may build, for any integer n, a preferential model of N, containing a top state that satisfies p_n. Therefore, for any n, p_n is consistent for R.

Remember that $a < b$ is the assertion $a \vee b \mathrel{|\!\sim} \neg b$. We shall write $a < b$ to mean that $a \vee b \mathrel{|\!\sim} \neg b$ is in R. If follows from results of Section 3.6 that, on formulas that are consistent for R, the relation $<$ is a strict modular ordering. Notice, also, that, for any n, the assertion $p_{n+2} < p_n$ belongs to N^p, since both $p_n \mathrel{|\!\sim} p_{n+2}$ and $p_{n+2} \mathrel{|\!\sim} \neg p_n$ are in N^p. Therefore $p_{n+2} < p_n$. There are, in R, two infinite (in both directions) chains (for $<$), one containing the variables of odd index, the other one containing those of even index. Since $<$ is modular, we may consider only four cases:

1. For every even n and odd k, $p_n > p_k$. Let f be the renaming of $\mathcal{L}$ given by $f(m) = m + 1$. Clearly $f(N) = N$. Hence, by Lemma 26 we must have $f(R) = R$. But this last statement implies $p_n > p_k$ for even n and odd k, and therefore implies that all p_n's are inconsistent for R. A contradiction.

2. For every even n and odd k, $p_n < p_k$. The argument is exactly as in case 1, systematically interchanging 'odd' and 'even'.

3. There is an odd k ,and there are even m and n such that $p_n < p_k < p_m$. In this case, define a renaming f by $f(l) = l$ for odd l and $f(l) = l + m - n$ for even l. The contradiction is as above by noting that f transforms $p_n < p_k$ into $p_m < p_k$.
4. None of the above is true. In such a case one may see that there must exist an even m ,and odd i and j such that $p_i < p_m < p_j$. The argument is exactly as in case 3, systematically interchanging 'odd' and 'even'.

▮

5.5 Global properties of the operation of rational closure

First, we show that rational closure possesses a *loop* property analogous to the property discussed in [17, Section 4]. This is a powerful property that one is happy to have.

Lemma 28 (Loop property) *Let K_i for $i = 0, \ldots, n-1$ be knowledge bases such that, for any i, $K_{i+1} \subseteq \overline{K_i}$, where addition is understood modulo n. Then for any i, j, one has $\overline{K_i} = \overline{K_j}$.*

Proof: Let $K \preceq K'$ mean that either $K \prec K'$ or $K = K'$. Since $\overline{K_i}$ is a rational extension of K_{i+1}, we have $\overline{K_{i+1}} \preceq \overline{K_i}$, for all i's (modulo n). We conclude that the rational closures of all the $\mathbf{K}_i$'s are equal. ▮

The following property of *reciprocity* is the special case $n = 2$.

Corollary 4 *If $X \subseteq \overline{Y}$ and $Y \subseteq \overline{X}$, then $\overline{X} = \overline{Y}$.*

The following property of *cumulativity* is equivalent to reciprocity in the presence of inclusion (i.e., $K \subseteq \overline{K}$).

Corollary 5 *If $X \subseteq Y \subseteq \overline{X}$ then $\overline{X} = \overline{Y}$.*

The meaning of Corollary 5 is that one may add to a knowledge base anything that is in its rational closure without changing this closure. We may now show that, in two different respects, rational closure is close to being perfect.

Lemma 29 *The consequence relation $\overline{\mathbf{K}}$, if it exists, contains an assertion of the form* **true** $\mid\sim \alpha$ *only if this assertion is in $\mathbf{K}^p$.*

Proof: Suppose an assertion of the form above is in $\overline{\mathbf{K}}$. We shall show that it must be in any rational extension of $\mathbf{K}$ and will conclude by Theorem 6. Suppose K' is a rational extension of $\mathbf{K}$ and **true** $\mid\sim \alpha$ is in $\overline{\mathbf{K}} - K'$. Since $\overline{\mathbf{K}} \prec K'$, we know there is an assertion $\gamma \mid\sim \delta$ in $K' - \overline{\mathbf{K}}$ such that $\gamma <$ **true** in K'. But this means **true** $\mid\sim$ **false** is in K', and contradicts the fact that **true** $\mid\sim \alpha$ is not in K'. ▮

Lemma 30 *The consequence relation $\overline{\mathbf{K}}$, if it exists, contains an assertion of the form $\alpha \mid\sim$* **false** *only if this assertion is in $\mathbf{K}^p$.*

Proof: Let V and T be preferential models defining the relations $\overline{\mathbf{K}}$ and $\mathbf{K}^p$ respectively. Such models exist by Theorem 1. Let U be the model obtained by putting T on top of V, i.e., every state of V is less than every state of T. One easily sees that U satisfies the smoothness property and is therefore a preferential model. It defines a preferential relation S. An assertion of the form $\alpha \mid\sim \mathbf{false}$ is in S only if it is in $\mathbf{K}^p$, since T is a submodel of U. If α is not inconsistent in $\overline{\mathbf{K}}$ then for any β, $\alpha \mid\sim \beta$ is in S iff it is in $\overline{\mathbf{K}}$. By Lemma 18, there is a rational extension R of S with the same set of inconsistent formulas. If one looks at the construction described in the proof of this lemma, one sees that it will add to S only assertions with antecedent inconsistent in $\overline{\mathbf{K}}$. Therefore, if α is not inconsistent in $\overline{\mathbf{K}}$, for any β, $\alpha \mid\sim \beta$ is in R iff it is in $\overline{\mathbf{K}}$. Now, R is a rational extension of $\mathbf{K}$. If R is equal to $\overline{\mathbf{K}}$, we are through. Suppose not. Then we have $\overline{\mathbf{K}} \prec R$. Suppose now that $\alpha \mid\sim \mathbf{false}$ is in $\overline{\mathbf{K}} - R$. There must be an assertion $\gamma \mid\sim \delta$ in $R - \overline{\mathbf{K}}$ such that $\gamma < \alpha$ in R. But $\gamma \mid\sim \delta$ in $R - \overline{\mathbf{K}}$ implies that $\gamma \mid\sim \mathbf{false}$ is in $\overline{\mathbf{K}}$ and $\gamma \mid\sim \delta$ is in $\overline{\mathbf{K}}$. A contradiction. ∎

5.6 Admissible knowledge bases and their rational closure

In this section, we show that an admissible (see Definition 11) knowledge base has a rational closure and that this rational closure may be defined in terms of the ranks of the formulas, as defined in Section 2.6. This provides a useful and elegant characterization of the rational closure of an admissible knowledge base.

Theorem 7 *Let* $\mathbf{K}$ *be an admissible conditional knowledge base. The rational closure* $\overline{\mathbf{K}}$ *of* $\mathbf{K}$ *exists and is the set* S *of all assertions* $\alpha \mid\sim \beta$ *such that either*

1. *the rank of* α *is strictly less than the rank of* $\alpha \wedge \neg\beta$ *(this includes the case* α *has a rank and* $\alpha \wedge \neg\beta$ *has none), or*
2. α *has no rank (In this case* $\alpha \wedge \neg\beta$ *has no rank either).*

Proof: Suppose indeed that every formula consistent with $\mathbf{K}^p$ has a rank. We have many things to check. First let us prove that S contains $\mathbf{K}$. If $\alpha \mid\sim \beta$ is in $\mathbf{K}$ and α has rank τ, then C_τ contains $\alpha \mid\sim \beta$ and entails $\mathbf{true} \mid\sim \alpha \rightarrow \beta$. Therefore $\alpha \wedge \neg\beta$ is exceptional for C_τ, and has rank strictly larger than τ.

We should now check that S is rational. For **Left Logical Equivalence**, **Right Weakening** and **Reflexivity** the proof is easy. For **Cautious Monotonicity**, notice that if $\alpha \mid\sim \beta$ is in S, then α and $\alpha \wedge \beta$ have the same rank. For **And** and **Or**, notice that the rank of a disjunction is the smaller of the ranks of its components. For **Rational Monotonicity**, notice that if $\alpha \mid\sim \neg\beta$ is not in S, then α and $\alpha \wedge \beta$ have the same rank.

We must now check that if R is a rational extension of $\mathbf{K}$ that is different from S then $S \prec R$. Let R be such an extension. We shall first show that S and R must agree on all assertions whose antecedents have no rank (the notion of rank is always defined by reference to $\mathbf{K}$). Indeed, by construction, any such assertion is in S, and it is preferentially entailed by $\mathbf{K}$ since $\mathbf{K}$ is admissible. It is therefore in R. We conclude that S and R must differ for some assertion that has rank. Let τ be the smallest rank at which S and R differ, i.e., the smallest rank of an α such that there is a β such that $\alpha \mid\!\sim \beta \in (S - R) \cup (R - S)$. We have two cases to consider, either there is a formula α of rank greater or equal to τ such that, for all formulas β of rank greater or equal to τ, $\alpha \leq \beta$ in R, or there is no such formula.

Suppose there is such an α. Our first claim is that, for any β of rank greater than τ, the assertion $\alpha < \beta$ is in R. Consider indeed a ranked model W that defines R. Let W'' be the supermodel obtained from W by adding to W, at each level l a state labeled with world w for all worlds w that label a state of rank less than l in W. It is clear that W'' is ranked and defines the same relation as W, and, in W'', every label that appears at some level l also appears at all greater levels. Let W' be the submodel of W'' that contains all those states of level (rank in W'') greater or equal to the minimal level l at which some state satisfies α. It clearly satisfies the smoothness property (for this we needed to go through the construction of W''). Since a formula is satisfied in W'' at some level less than l iff it is of rank less than τ, no antecedent of an assertion of C_τ is satisfied at any level less than l. But W'' is a model of $\mathbf{K}$ and therefore W' is a model of C_τ. But C_τ preferentially entails $\mathbf{true} \mid\!\sim \neg\beta$. The model W' therefore satisfies $\mathbf{true} \mid\!\sim \neg\beta$ but not $\mathbf{true} \mid\!\sim \neg\alpha$. It therefore also satisfies $\alpha < \beta$. But the antecedent of this last assertion has rank greater or equal to τ, and therefore no state of W'' that is not in W' satisfies it. Therefore $\alpha < \beta$ is satisfied by W'' and is an element of R. satisfies Our second claim is that there is an assertion $\gamma \mid\!\sim \delta$ in $R - S$, such that γ is of rank τ and $\gamma \leq \alpha$ in R. We consider two cases.

1. There is an assertion $\xi \mid\!\sim \eta$ in $S - R$ with ξ of rank τ. Then $\xi \wedge \neg\eta$ has rank greater than τ, and by our first claim, $\alpha < \xi \wedge \neg\eta$ is in R. But $\xi \mid\!\sim \eta$ is not in R, and therefore we must have $\alpha < \xi$ for R. This last assertion is not in S since both α and ξ have rank τ. The assertion $\alpha < \xi$ is a suitable $\gamma \mid\!\sim \delta$.

2. There is an assertion $\xi \mid\!\sim \eta$ in $R - S$ with ξ of rank τ. If $\xi \leq \alpha$ in R, then $\xi \mid\!\sim \eta$ is a suitable $\gamma \mid\!\sim \delta$. Suppose, then, that $\alpha < \xi$ in R. Since ξ has the same rank as α, $\alpha < \xi$ is in $R - S$ and a suitable $\gamma \mid\!\sim \delta$.

We may now conclude that $S \prec R$. The assertion $\gamma \mid\!\sim \delta$ fulfills the requirements of Condition 1 of Definition 20, since γ has rank τ. For Condition 2, suppose $\xi \mid\!\sim \eta$ is in $S - R$, then ξ must be of rank greater or equal to τ and $\xi \wedge \neg\eta$ is of rank greater than τ. By our first remark we conclude that $\alpha < \xi \wedge \neg\eta$ for R. It is a matter of elementary properties of rational relations to check that if $\alpha < \xi \wedge \neg\eta$ is in R, but $\xi \mid\!\sim \eta$ is not, then $\alpha < \xi$ for R. Since $\gamma \leq \alpha$ in R, we conclude that $\gamma < \xi$ for R.

Suppose now that there is no such α. Take any formula γ of rank τ. There is a formula δ of rank greater or equal to τ such that $\delta < \gamma$ for R. But this assertion is then in $R - S$. It satisfies Condition 1 of Definition 20, since its antecedent has rank τ. Suppose now $\varphi \mid\sim \theta$ is in $S - R$. Then φ is of rank at least τ. If it is of rank τ, there is a formula π of rank at least τ such that $\pi < \varphi$ is in R, but not in S and this provides the witness requested by Condition 2. If it is of rank greater than τ, then the assertion $\delta < \gamma$ defined just above will do. ∎

5.7 A model-theoretic description of rational closure

We shall describe here a model-theoretic construction that transforms a preferential model W into a ranked model W' by letting all states of W sink as low as they can respecting the order of W, i.e., ranks the states of W by their height in W. We shall show that, under certain conditions, the model W' defines the rational closure of the relation defined by W. This construction is clearly interesting only when the model W is well-founded. We know that, in this case, the relation defined by W indeed possesses a rational closure (Theorem 7 and Lemma 7). It would have been pleasant to be able to prove the validity of such a construction on an arbitrary well-founded preferential model. Unfortunately we are not able to show this in general, but need to suppose, in addition, that the preferential relation defined by W is well-founded (see Definition 8). This is quite a severe restriction, since we have seen at the end of Section 2.3 that finitely-generated relations on arbitrary languages $\mathcal{L}$ are not always well-founded. When the language $\mathcal{L}$ is logically finite, we know all preferential relations are well-founded. Given a well-founded preferential relation, the construction may be applied to any of its well-founded models.

Let P be a well-founded preferential relation and $W = \langle S, l, \prec \rangle$ any well-founded preferential model that defines P. We shall define, for any ordinal τ, two sets of states: U_τ and V_τ. Those sets satisfy, for any τ, $U_\tau \subseteq V_\tau \subseteq U_{\tau+1}$. The set U_τ contains, in addition to the elements of previous V's, the states that are minimal among those states not previously added. The set V_τ contains, in addition to the states of U_τ, all states that satisfy only formulas already satisfied by states previously considered.

$$U_\tau \stackrel{\text{def}}{=} \bigcup_{\rho<\tau} V_\rho$$

$$\cup \{s \in S \mid \forall t \in S \text{ such that } t \prec s, \text{there is a } \rho < \tau \text{ such that } t \in V_\rho\} \quad (19)$$

$$V_\tau \stackrel{\text{def}}{=} \{s \in S \mid \forall \alpha \in \mathcal{L} \text{ such that } s \models \alpha, \exists t \in U_\tau \text{ such that } t \models \alpha\} \quad (20)$$

Since the model W is well-founded, every state $s \in S$ is in some V_τ. Let the height of a state $s \in S$ (in W) be the least ordinal τ for which $s \in V_\tau$. We shall now show that there is a close relationship between the rank of a formula α in P (see definition following Definition 10) and the height in W of the states that satisfy α. For any ordinal τ, we shall denote by W_τ the substructure of W consisting of all states of height larger or equal to τ. Notice that, since W is well-founded, W_τ is a preferential model. Notice also that all elements of $U_\tau - \bigcup_{\rho<\tau} V_\rho$ are minimal elements of W_τ.

Lemma 31 *Let τ be an ordinal. Let α be a formula of rank at least τ and β be any formula.*

1. *No state of height less than τ satisfies α.*
2. *The model W_τ satisfies $\alpha \mid\!\sim \beta$ iff $\alpha \mid\!\sim \beta$ is preferentially entailed by C_τ.*

In particular, if α has no rank, no state in S satisfies α.

Proof: It proceeds by simultaneous ordinal induction on τ. Suppose both claims have been proved for all ordinals $\rho < \tau$. Let us prove our first claim. Since α has rank at least τ, for any ρ, $\rho < \tau$, C_ρ preferentially entails **true** $\mid\!\sim \neg\alpha$. By the induction hypothesis (item 2), W_ρ satisfies **true** $\mid\!\sim \neg\alpha$. Therefore no state of $U_\rho - \bigcup_{\sigma<\rho} V_\sigma$ satisfies α. If there were a state s of height $\rho < \tau$ satisfying α, there would be a state t of $U_\rho - \bigcup_{\sigma<\rho} V_\sigma$ satisfying α. We conclude that no state of height less than τ satisfies α.

For the second claim, by Lemma 5, $\alpha \mid\!\sim \beta$ is preferentially entailed by C_0 (i.e., in P, i.e., satisfied by W) iff it is preferentially entailed by C_τ. By the first claim, $\alpha \mid\!\sim \beta$ is satisfied by W iff it is satisfied by W_τ. ∎

Lemma 32 *A formula α has rank τ in P iff there is a state $s \in S$ of height τ that satisfies α and there is no such state of height less than τ.*

Proof: We shall prove the *only if* part. The *if* part is then obvious. First, remark that if $\mid\!\sim$ is a preferential relation that contains the assertion $\alpha \vee \beta \mid\!\sim \neg\alpha$, then it contains the assertion **true** $\mid\!\sim \neg\alpha$. This is easily shown by preferential reasoning. Suppose now that α has rank τ. Lemma 31 shows that no state of height less than τ satisfies α. We must show that there is a state of height τ satisfying α. Let β be any formula of rank larger or equal to τ that is minimal with respect to $<$ among those formulas. There is such a formula since the set is not empty (α is there) and $<$ is well-founded. Since α is not exceptional for C_τ, the assertion **true** $\mid\!\sim \neg\alpha$ is not preferentially entailed by C_τ and therefore the assertion $\alpha \vee \beta \mid\!\sim \neg\alpha$ is not preferentially entailed by C_τ. But $\alpha \vee \beta$ has rank τ and, by Lemma 31, W_τ does not satisfy $\alpha \vee \beta \mid\!\sim \neg\alpha$. There is, therefore, in W_τ a state s satisfying α such that no state t in W_τ, $t \prec s$, satisfies β. We shall show that s is minimal in W_τ and has therefore height τ. Suppose s is not minimal in W_τ. There would be a state t minimal in W_τ such that $t \prec s$. This state t has height τ and, by construction, it satisfies some formula β' that is not satisfied at any smaller height. By Lemma 31, β' has rank larger or equal to τ, and the formula $\beta \vee \beta'$ has rank larger or equal to τ. Since $\beta \vee \beta' \leq \beta$, the minimality of β implies that $\beta \leq \beta \vee \beta'$. In other terms, $\beta \vee \beta' \mathrel{|\!\sim} \beta$. But the state t, in W, satisfies β' and is minimal among states satisfying $\beta \vee \beta'$. Therefore t satisfies β. A contradiction. ∎

Lemma 32 shows that, given a well-founded preferential relation (resp. a finite knowledge base), and a well-founded preferential model W for it (resp. for its preferential closure), one may build a ranked model for its rational closure by ranking the states of W by their depth.

5.8 Computing rational closure

We shall now provide an algorithm for deciding whether an assertion is in the rational closure of a finite knowledge base. The notation $E(C)$ has been defined following Definition 10. Lemma 7 and Theorem 7 show that, given a finite knowledge base $\mathbf{K}$ and an assertion $\alpha \mathrel{|\!\sim} \beta$ the following algorithm is adequate.

$C = \mathbf{K}$;

while α is exceptional for C and $E(C) \neq C$, $C := E(C)$;

if $\alpha \wedge \neg\beta$ is exceptional for C then answer yes else answer no.

The only thing left for us to implement is checking whether a formula is exceptional for a given finite knowledge base. The next lemma shows this is easily done.

Definition 23 *Let $\mathcal{A}$ be the conditional assertion $\alpha \mathrel{|\!\sim} \beta$. The material counterpart of $\mathcal{A}$, denoted by $\tilde{\mathcal{A}}$, is the formula $\alpha \rightarrow \beta$, where $\rightarrow$, as usual, denotes material implication. If $\mathbf{K}$ is a set of assertions, its material counterpart $\tilde{\mathbf{K}}$ is the set of material counterparts of $\mathbf{K}$.*

Lemma 33 *Let $\mathbf{K}$ be a conditional knowledge base and α a formula. Then $\tilde{\mathbf{K}} \models \alpha$ iff $\mathbf{K}$ preferentially entails $\mathbf{true} \mathrel{|\!\sim} \alpha$.*

Proof: The *if* part follows from the fact that any world satisfying $\tilde{\mathbf{K}}$ and not α provides a one state preferential model satisfying $\mathbf{K}$ and not satisfying $\mathbf{true} \mathrel{|\!\sim} \alpha$. For the *only if* part suppose $\tilde{\mathbf{K}} \models \alpha$. By compactness, there is a finite subset of $\tilde{\mathbf{K}}$ that entails α. By rules **S**, **And** and **Right Weakening** we conclude that $\mathbf{K}$ preferentially entails $\mathbf{true} \mathrel{|\!\sim} \alpha$. ∎

Corollary 6 *Let $\mathbf{K}$ be a conditional knowledge base and α a formula. The formula α is exceptional for $\mathbf{K}$ iff $\tilde{\mathbf{K}} \models \neg\alpha$.*

We see that, if $\mathbf{K}$ contains n assertions, in the previous algorithm, we may go over the while loop at most $O(n)$ times. Each time we shall have to consider at most $n+1$ formulas and decide whether they are exceptional or not. The whole algorithm needs at most $O(n^2)$ such decisions. In the most general case all such decisions are instances of the satisfiability problem for propositional calculus, therefore solvable in non-deterministic polynomial time (in the size of the knowledge base $\mathbf{K}$ times the size of the formulas involved). Therefore, even in the most general case, the problem is not much more complex than the satisfiability problem for propositional calculus. These results may be improved if we restrict ourselves to assertions of a restricted type. For example, if the assertions of $\mathbf{K}$ are of the Horn type (we mean their material counterpart is a Horn formula), and the assertion $\alpha \mathrel{|\!\sim} \beta$ is of the same type, then, since each decision may be taken in polynomial deterministic time, the whole algorithm runs in deterministic polynomial time. The complexity discussion above is mainly of theoretical interest. The important practical question is: given a fixed large knowledge base, what information, of reasonable size, should be precomputed to allow efficient answers to queries of the type: is an $\mathcal{A}$ in the rational closure? The pre-computation of the different C_n sub-bases would already reduce the exponent of n in the complexity of the algorithm by one.

J. Dix noticed that the algorithm just presented for computing the rational closure of a finite knowledge base may be used to compute the preferential closure of such a knowledge base, since, by Corollary 3 and Lemmas 7 and 30, the assertion $\alpha \mathrel{|\!\sim} \beta$ is in $\mathbf{K}^p$ iff the assertion $\alpha \mathrel{|\!\sim} \mathbf{false}$ is in the rational closure of the knowledge base $\mathbf{K} \cup \{\alpha \mathrel{|\!\sim} \neg\beta\}$.

5.9 A discussion of rational closure

We have so far shown that rational closure provides a mathematically elegant and effective answer to the question of the title that satisfies Thesis 1. It is now time to evaluate whether it provides an answer that matches our intuitions. We shall first present two now classical knowledge bases, describe their rational closure and examine whether they fit our intuitions. Then, we shall discuss the way rational closure treats inheritance of generic properties to abnormal individuals. Finally, we shall try to address the question of whether our formalism is suitable to describe domain knowledge.

Example 2 (Nixon diamond) Let our knowledge base consist of the following two assertions.

1. $\mathit{republican} \mathrel{|\!\sim} \neg \mathit{pacifist}$
2. $\mathit{quaker} \mathrel{|\!\sim} \mathit{pacifist}$

It is easy to see that none of the assertions of the base is exceptional, but that the formula $\mathit{republican} \wedge \mathit{quaker}$ is exceptional. From this we deduce that neither the assertion $\mathit{republican} \wedge \mathit{quaker} \mathrel{|\!\sim} \mathit{pacifist}$ nor the assertion $\mathit{republican} \wedge \mathit{quaker} \mathrel{|\!\sim} \neg \mathit{pacifist}$ is in the rational closure. This seems the intuitively correct decision in the presence of contradictory information. In fact, if we know somebody to be both a Quaker and a Republican, we (i.e., rational closure) shall draw about him only conclusions that are logically implied by our information. Rational closure endorses $\mathit{worker} \wedge \mathit{republican} \mathrel{|\!\sim} \neg \mathit{pacifist}$, meaning that, since we have no information on the pacifism of workers, we shall assume that Republican workers behave as Republicans in this respect. We (i.e., rational closure) also endorse $\mathit{pacifist} \mathrel{|\!\sim} \neg \mathit{republican}$, meaning we are ready to use contraposition in many circumstances. We do not have $\neg \mathit{pacifist} \mathrel{|\!\sim} \mathit{republican}$, though, and quite rightly, since Republicans may well be a small minority among non-pacifists. We have $\mathbf{true} \mathrel{|\!\sim} \neg\, (\mathit{republican} \wedge \mathit{quaker})$, meaning we think being both a Republican and a Quaker is exceptional. We endorse $\mathit{republican} \mathrel{|\!\sim} \neg \mathit{quaker}$ and $\mathit{quaker} \mathrel{|\!\sim} \neg \mathit{republican}$, that are also intuitively correct conclusions. If we add to our knowledge base the fact that rich people are typically Republicans, we shall deduce that rich people are typically not pacifists, meaning we endorse a restricted form of transitivity. We shall also deduce that Quakers are typically not rich, which is perhaps more debatable. We shall not conclude anything about the pacifism of rich Quakers though, since rich Quakers are exceptional. We shall not conclude anything either concerning rich Quakers that are not Republicans, which is more debatable. If we want to conclude that rich non-Republican Quakers are pacifists, we should add this assertion explicitly to the knowledge base. The addition will not interfere with previously discussed assertions.

Example 3 (Penguin triangle) Let our knowledge base consist of the following three assertions.

1. $\mathit{penguin} \mathrel{|\!\sim} \mathit{bird}$
2. $\mathit{penguin} \mathrel{|\!\sim} \neg \mathit{fly}$
3. $\mathit{bird} \mathrel{|\!\sim} \mathit{fly}$

The first two assertions are exceptional, the last one is not. It follows that we (i.e., rational closure) endorse the following assertions: $\mathit{fly} \mathrel{|\!\sim} \neg \mathit{penguin}$ (a case of contraposition), $\neg \mathit{fly} \mathrel{|\!\sim} \neg \mathit{bird}$ (another case of contraposition), $\neg \mathit{fly} \mathrel{|\!\sim} \neg \mathit{penguin}$ (penguins are exceptional, even among non-flying objects), $\mathit{bird} \mathrel{|\!\sim} \neg \mathit{penguin}$ (penguins are exceptional birds), $\neg \mathit{bird} \mathrel{|\!\sim} \neg \mathit{penguin}$ (penguins are exceptional also among non-birds), $\mathit{bird} \wedge \mathit{penguin} \mathrel{|\!\sim} \neg \mathit{fly}$ (this is an intuitively correct preemption: we prefer specific information to non-specific information), $\mathit{penguin} \wedge \mathit{black} \mathrel{|\!\sim} \neg \mathit{fly}$ (black penguins don't fly either, since they are normal penguins), $\mathit{bird} \wedge \mathit{green} \mathrel{|\!\sim} \mathit{fly}$ (green birds are normal birds).

The following assertions are not endorsed: $\mathit{bird} \wedge \neg \mathit{fly} \mathrel{|\!\sim} \mathit{penguin}$ (there could be non-flying birds other than penguins), $\mathit{bird} \wedge \neg \mathit{fly} \mathrel{|\!\sim} \neg \mathit{penguin}$ (seems intuitively clear), $\mathit{penguin} \mathrel{|\!\sim} \mathit{fly}$ (obviously).

A more general reflexion suggests the following. Theorem 7 shows that, in the rational closure, no information about normal cases may be relevant to abnormal cases. It is a very intriguing question whether human beings obey this rule of reasoning or not. A specific example has been discussed by J. Pearl in a personal communication. It probably goes back to A. Baker. Suppose we know that most Swedes are blond and tall. If we are going to meet Jon, whom we know to be short and to come from Sweden, should we necessarily expect him to be fair? The answer endorsed by rational closure is *not necessarily*, since short Swedes are exceptional and we have no specific information about such cases. We do not know how people generally handle this and, even if we knew, it is not clear that AI systems should react in exactly the same way: people are, after all, notoriously bad with statistical information. The answer to the question how should people behave in this case, if they were smart and had all the relevant information, depends on the sociobiology of the Swedish population and is not relevant either. There is very solid ground, though, to claim that, in the framework described here, in which a knowledge base contains only positive conditional assertions, the only sensible way to handle this problem is not to expect anything about the color of Jon's hair. The reason is that, if we ever find out that most short Swedes are blond (or dark, for that matter) it will be easy enough to add this information to our knowledge base. On the contrary, had we chosen to infer that Jon is expected to be blond, and had we found out that half of the short Swedes only are fair, we would not have been able to correct our knowledge base to remove the unwanted inference: adding the fact that most short Swedes are not blond being obviously incorrect.

Since, by looking at a number of examples, we have gathered some experience on the behavior of rational closure, we would like to propose the following strengthening of Thesis 1.

Thesis 2 *The set of assertions entailed by any set of assertions* $\mathbf{K}$ *is a rational superset of* $\overline{\mathbf{K}}$.

Thesis 2 means that a reasonable system should endorse any assertion contained in the rational closure, but it may also endorse some additional assertions, as long as it defines a rational relation. The search for natural constructions satisfying Thesis 2, but providing more *inheritance* than rational closure is open.

The main question that has not been addressed yet is whether conditional knowledge bases are suitable to describe domain knowledge. Undoubtedly much work still has to be done before we may answer this question satisfactorily. We shall only try to express here why we think the answer may well be positive. Representing common sense knowledge is far from trivial in any one of the existing formalisms, such as Circumscription or Default Logic. Indeed to represent any substantive piece of common sense knowledge in one of those formalisms, one needs to be an expert at the mechanics of the formalism used, and they differ greatly from one formalism to the next. Deciding on the different abnormality predicates in Circumscription and the relations between them, or working out the default rules in Default Logic so as to ensure the correct precedence of defaults needs the hand of an expert. In the formalism proposed here, conditional knowledge bases, the treatment is much simpler since abnormality predicates do not appear explicitely and the default information is described in a much poorer language than Default Logic. We rely on the general algorithm for computing rational closure (or some other algorithm that will be found suitable) to deal in a mechanical, uniform and tractable manner with the interactions between different pieces of default information. The fact that our language of assertions is much poorer than other formalisms seems to us to be a great asset.

Nevertheless, it is probable that the size of useful conditional knowledge bases will be very large. Indeed, in our approach, adding new assertions to the knowledge base may solve almost any problem. Two main topics for further research may then be delineated. The first one is to find practical ways to avoid having to look at the whole knowledge base before answering any query. The set of assertions constituting a knowledge base will have to be structured (off-line, once and for all) in such a way that irrelevant assertions do not have to be looked at. The second one is to find lucid and compact descriptions of large conditional knowledge bases. This will involve looking seriously into the question: where does the conditional knowledge come from? Different answers may be appropriate in different domains: it may well be that conditional knowledge is derived from causal knowledge in ways that are different from those in which it is derived from conventions of speech or statistical information.

6 Conclusion

We have presented a *mathematically tractable* framework for nonmonotonic reasoning that can be proved to possess many pragmatically attractive features. Its computational complexity compares favorably with that of most well-established systems. In many cases the intuitively correct answer is obtained. In others, the answer given and the way it was obtained provide an interesting point of view on the knowledge base. Much more practical experience is needed before one may assess the pragmatic value of the approach. The task of extending the results presented here to first-order languages is not an easy one. First steps towards this goal are described in [20].

7 Acknowledgements

David Makinson suggested importing the thesis **CV** of conditional logic into the study of nonmonotonic consequence relations, i.e., suggested to consider what is called here rational relations. He conjectured that the corresponding family of models was that of ranked models. He was also instrumental in stressing the importance of studying global properties of nonmonotonic inference operations.

Discussions with the following people helped us to disprove hasty conjectures, putting this work in perspective, and improve the presentation of this paper: Johan van Benthem, Michael Freund, Haim Gaifman, Hector Geffner, Matthew Ginsberg, David Israel, Sarit Kraus, John McCarthy and Judea Pearl. Karl Schlechta's suggestions and Jürgen Dix's remarks on a previous draft have been very useful. Finally, this paper has been fortunate to receive attention and care of a rare quality from two anonymous referees. We want to thank them.

References

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

[2] Peter Cheeseman. In defense of *an inquiry into computer understanding. Computational Intelligence*, 4(1):129–142, February 1988.

[3] Peter Cheeseman. An inquiry into computer understanding. *Computational Intelligence*, 4(1):58–66, February 1988.

[4] R. Chisholm. The contrary-to-fact conditional. *Mind*, 55:289–307, 1946. reprinted in Readings in Philosophical Analysis, edited by H. Feigl and W. Sellars, Appleton-Century-Crofts, New York, 1949, pp. 482–497.

[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] N. J. Cutland. Non standard measure theory and its applications. *Bulletin of the London Mathematical Society*, pages 529–589, 1983.

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

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

[9] Jon Doyle and Michael P. Wellman. Impediments to universal preference-based default theories. *Artificial Intelligence*, 49(1–3):97–128, May 1991.

[10] Michael Freund. A semantic characterization of disjunctive relations. In Philippe Jorrand and J. Kelemen, editors, *Proceedings of FAIR'91, Lecture Notes in Artificial Intelligence Vol. 535*, pages 72–83, Smolenice, Czechoslovakia, September 1991. Springer Verlag.

[11] Michael Freund, Daniel Lehmann, and Paul Morris. Rationality, transitivity and contraposition. *Artificial Intelligence*, 1991? in print.

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

[13] George Grätzer. *Lattice Theory*. W. H. Freeman, San Francisco, 1971.

[14] William L. Harper. A sketch of some recent developments in the theory of conditionals. In William L. Harper, Robert Stalnaker, and Glenn Pearce, editors, *Ifs: Conditionals, Belief, Decision, Chance and Time*, volume 15 of *The University of Western Ontario Series in Philosophy of Science*, chapter Introduction, pages 3–38. D. Reidel, Dordrecht, Boston, London, 1981.

[15] William L. Harper, Robert Stalnaker, and Glenn Pearce, editors. *Ifs: Conditionals, Belief, Decision, Chance and Time*, volume 15 of *The University of Western Ontario Series in Philosophy of Science*. D. Reidel, Dordrecht, Boston, London, 1981.

[16] H. J. Keisler. *Foundations of Infinitesimal Calculus*. Prindle, Weber&Schmidt Inc., Boston, 1976.

[17] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1–2):167–207, July 1990.

[18] 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.

[19] 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.

[20] Daniel Lehmann and Menachem Magidor. Preferential logics: the predicate calculus case. In *Proceedings of the Third Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 57–72, Monterey, California, March 1990. Morgan Kaufmann.

[21] 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.

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

[23] Drew McDermott and Jon Doyle. Non-monotonic logic I. *Artificial Intelligence*, 25:41–72, 1980.

[24] Robert C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94, 1985.

[25] 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.

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

[27] Judea Pearl. System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proceedings of the Third Conference on Theoretical Aspects of Reasoning About Knowledge*, pages 121–135, Monterey, California, March 1990. Morgan Kaufmann.

[28] Frank Plumpton Ramsey. General propositions and causality (1925). In D. H. Mellor, editor, *Foundations: Essays in Philosophy, Logic, Mathematics and Economics*, pages 237–257. Routledge and K. Paul, London, 1978.

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

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

[31] Abraham Robinson. *Non-standard Analysis.* North-Holland, Amsterdam, 1966.

[32] Ken Satoh. A probabilistic interpretation for lazy nonmonotonic reasoning. Technical report, Institute for New Generation Computer Technology, 1-4-28 Mita, Minato-ku, Tokyo 108, Japan, December 1989.

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

[34] Robert C. Stalnaker and Richmond H. Thomason. A semantic analysis of conditional logic. *Theoria*, 36:23–42, 1970.

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

[36] Moshe Y. Vardi, editor. *Proceedings of the Second Conference on Theoretical Aspects of Reasoning About Knowledge*, Monterey, California, March 1988. Morgan Kaufmann.

A Lemmas needed to prove Theorem 5

Let us suppose that some rational consequence relation $\mid\!\sim$ is given. The notion of a consistent formula has been presented in Definition 9. Let S denote the set of all consistent formulas. Let us now recall Definition 10 of [17].

Definition 24 *The world $m \in \mathcal{U}$ is a normal world for α iff $\forall \beta \in L$ such that $\alpha \mid\!\sim \beta$, $m \models \beta$.*

The following is an easy corollary of Lemma 8 of [17].

Lemma 34 *A formula is consistent iff there is a normal world for it.*

We shall now define a pre-order relation on the set S.

Definition 25 *Where $\alpha, \beta \in S$, we shall say that α is not more exceptional than β and write $\alpha \mathcal{R} \beta$ iff $\alpha \vee \beta \not\mid\!\sim \neg\alpha$.*

Lemma 35 *The relation $\mathcal{R}$ is transitive.*

Proof: Straightforward from Lemma 17. The fact that the relation $\mathcal{R}$ was restricted to the set S is not used here and $\mathcal{R}$ would have been transitive also on the whole language $\mathcal{L}$. ▮

Lemma 36 *Let $\alpha, \beta \in S$. Either $\alpha \mathcal{R} \beta$ or $\beta \mathcal{R} \alpha$ (or both). In particular $\mathcal{R}$ is reflexive.*

Proof: The proof proceeds by contradiction. Suppose we have $\alpha \not\mathcal{R} \beta$ and $\beta \not\mathcal{R} \alpha$. Then we have $\alpha \vee \beta \mid\!\sim \neg\alpha$ and $\alpha \vee \beta \mid\!\sim \neg\beta$. By **And** and **Reflexivity** we have $\alpha \vee \beta \mid\!\sim \neg\alpha \wedge \neg\beta \wedge (\alpha \vee \beta)$, and therefore $\alpha \vee \beta \mid\!\sim$ **false** and, by Rule 14, $(\alpha \vee \beta) \wedge \beta \mid\!\sim$ **false**. Therefore $\beta \mid\!\sim$ **false**, contradicting $\beta \in S$. The fact that $\mathcal{R}$ was restricted to S is crucial here. ▮

The following will be useful in the sequel.

Lemma 37 *If $\alpha \mathcal{R} \beta$, any normal world for α that satisfies β is normal for β.*

Proof: Suppose $\alpha \mathcal{R} \beta$, m is normal for α and satisfies β. Let γ be such that $\beta \mid\!\sim \gamma$. We must show that $m \models \gamma$. Since m is normal for α and satisfies β, it is enough to show that $\alpha \mid\!\sim \beta \rightarrow \gamma$. But, $\beta \mid\!\sim \gamma$ implies, by **Left Logical Equivalence**, $(\alpha \vee \beta) \wedge \beta \mid\!\sim \gamma$. By the rule **S** of [17], one then obtains $\alpha \vee \beta \mid\!\sim \beta \rightarrow \gamma$. But, by definition of $\mathcal{R}$, $\alpha \vee \beta \not\mid\!\sim \neg\alpha$ and, by **Rational Monotonicity** one deduces $(\alpha \vee \beta) \wedge \alpha \mid\!\sim \beta \rightarrow \gamma$. ▮

Definition 26 *Let $\alpha, \beta \in S$. We shall say that α is as exceptional as β and write $\alpha \sim \beta$ iff $\alpha \mathcal{R} \beta$ and $\beta \mathcal{R} \alpha$.*

Since $\mathcal{R}$ is reflexive and transitive, the relation $\sim$ is an equivalence relation. The equivalence class of a formula α will be denoted by $\overline{\alpha}$ and E will denote the set of equivalence classes of formulas of S under $\sim$. We shall write $\overline{\alpha} \leq \overline{\beta}$ iff $\alpha \mathcal{R} \beta$ and we shall write $\overline{\alpha} < \overline{\beta}$ iff $\overline{\alpha} \leq \overline{\beta}$ and $\alpha \not\sim \beta$. This notation should cause no confusion with a similar notation used with a different meaning, in the context of preferential relations, in [17] and in Section 2.4. By Lemmas 35 and 36, the relation $<$ is a strict total order on the set E.

Lemma 38 *Let α, β be consistent formulas. If $\overline{\beta} < \overline{\alpha}$ then $\beta \mathrel{|\!\sim} \neg\alpha$.*

Proof: The assumption implies that $\alpha \not\mathcal{R} \beta$, i.e., $\alpha \vee \beta \mathrel{|\!\sim} \neg\alpha$. Rule (15) implies the conclusion. ∎

Lemma 39 *Let α, β be consistent formulas. If there is a normal world for α that satisfies β, then $\overline{\beta} \leq \overline{\alpha}$.*

Proof: If there is a normal world for α that satisfies β, then we conclude by Lemma 38 that $\alpha \mathrel{|\!\not\sim} \neg\beta$. ∎

Let W be the ranked model $\langle V, l, \prec \rangle$, where $V \subseteq \mathcal{U} \times S$ is the set of all pairs $< m, \alpha >$ such that m is a normal world for α, $l(< m, \alpha >)$ is defined to be m and $\prec$ is defined as $< m, \alpha > \prec < n, \beta >$ iff $\overline{\alpha} < \overline{\beta}$. To show that W is a ranked model, we must prove that it satisfies the smoothness condition.

Lemma 40 *In W, the state $< m, \alpha >$ is minimal in $\widehat{\beta}$ iff $m \models \beta$ and $\overline{\beta} = \overline{\alpha}$*

Proof: First notice that $< m, \alpha > \in \widehat{\beta}$ iff $m \models \beta$. For the *only if* part, suppose that $< m, \alpha >$ is minimal in $\widehat{\beta}$. The world m is normal for α and satisfies β. By Lemma 39 we conclude that $\overline{\beta} \leq \overline{\alpha}$. But, since β is consistent, by Lemma 34 there is a normal world n for β. The pair $< n, \beta >$ is an element of V that satisfies β and, by the minimality of $< m, \alpha >$ in $\widehat{\beta}$, $< n, \beta > \not\prec < m, \alpha >$, i.e., $\overline{\beta} \not< \overline{\alpha}$, i.e., $\overline{\alpha} \leq \overline{\beta}$. We conclude $\overline{\beta} = \overline{\alpha}$. For the *if* part, suppose that m is a normal world for α that satisfies β and that $\overline{\beta} = \overline{\alpha}$. If n is normal for γ and $< n, \gamma > \prec < m, \alpha >$ then $\overline{\gamma} < \overline{\alpha}$ and therefore $\overline{\gamma} < \overline{\beta}$. By Lemma 38 $\gamma \mathrel{|\!\sim} \neg\beta$ and n, which is normal for γ, cannot satisfy β. The state $< m, \alpha >$ is then minimal in $\widehat{\beta}$. ∎

The following is an immediate corollary of Lemma 40.

Corollary 7 *If m is a normal world for α the pair $< m, \alpha >$ is a state of V and is minimal in $\widehat{\alpha}$.*

Proof: Suppose m is normal for α. First, since there is a normal world for α, $\alpha \in S$ and the pair $< m, \alpha >$ is in V. Since m is normal for α it satisfies α. ∎

We may now prove that the model W satisfies the smoothness property and defines the consequence relation $\mathrel{|\!\sim}$.

Lemma 41 *Let α be a consistent formula. The set $\widehat{\alpha} \subseteq V$ is smooth.*

Proof: Suppose $< m, \beta > \in \widehat{\alpha}$. Then, m is a normal world for β that satisfies α and, by Lemma 39, $\overline{\alpha} \leq \overline{\beta}$. If $\overline{\alpha} = \overline{\beta}$, then, by Lemma 40, $< m, \beta >$ is minimal in $\widehat{\alpha}$. Otherwise, $\overline{\alpha} < \overline{\beta}$. In this case, let n be any world normal for α (there is such a world since α is consistent). The pair $< n, \alpha >$ is minimal in $\widehat{\alpha}$ by Lemma 40 and $< n, \alpha > \prec < m, \beta >$. ∎

Lemma 42 *If $< m, \alpha >$ is minimal in $\widehat{\beta}$, then m is normal for β.*

Proof: Suppose $< m, \alpha >$ is minimal in $\widehat{\beta}$. By Lemma 40 $\overline{\alpha} = \overline{\beta}$. Therefore $\alpha \mathcal{R} \beta$. But m is normal for α and satisfies β, and Lemma 37 implies that m is normal for β. ∎

B Non-standard probabilistic semantics

B.1 Introduction

We shall describe now, in Definition 33 another family of probabilistic models, they provide much more direct semantics for nonmonotonic reasoning than Adams', at the price of using the language of non-standard (in the sense of A. Robinson) probability theory. The purpose of this section is to provide additional evidence in support of Thesis 1. We shall show that rational relations are exactly those that may be defined by non-standard *probabilistic* models. In other terms, if, given a probability distribution, we decide to accept the assertion $\alpha \mathrel{|}\!\sim \beta$ iff the conditional probability of β given α is very close to one, then the consequence relation we define is rational. On the other hand, any rational relation may be defined, in such a way, by some probability distribution. The results presented in the appendix are not used in the body of the paper. A different representation theorem for rational relations, also based on Theorem 5, in terms of one-parameter families of standard probabilistic models has been proved recently by K. Satoh [32]. Results relating the semantics of conditionals and non-Archimedean probabilities seem to have been obtained by R. Giles around 1980.

There is a school of thought in Artificial Intelligence, represented in particular by [3, 2], that denies the validity of the logical approach to modeling common-sense reasoning. The alternative suggested is the Bayesian probabilistic approach. Namely, the only way in which we should make sensible inferences from our knowledge α is by estimating the conditional probability of the required conclusion β given our knowledge α, and then adopting β if we are satisfied that this conditional probability is close enough to 1. We believe that this approach may run into considerable practical difficulties, the choice being between keeping an explicit data base of these many conditional probabilities or estimating them from a small sample. The chief source of difficulty here is that knowing the probability of α and β tells you very little about the probability of their intersection.

But we shall not argue the matter in detail here. The main purpose of this section is to show that rational knowledge bases may be considered to come from such a probabilistic model, if we let the cut-off point of how close the conditional probability of β given α has to be before we are ready to adopt β as a sensible consequence of α, approach 1 as a limit. Namely, β is a sensible consequence of α, iff the conditional probability is *infinitesimally* close to 1. In order to have an interesting theory, there must be probabilities that are not standard real numbers, but belong to a richer system of numbers, containing some infinitesimally small numbers.

We shall show that this approach allows one to keep a probabilistic intuition while thinking about common-sense reasoning, namely think about $\alpha \mid\!\sim \beta$ as meaning that the conditional probability of β given α is large, and still defines a well-behaved consequence relation that is not necessarily monotonic. Note that if one considers a standard probabilistic model and accepts $\alpha \mid\!\sim \beta$ as satisfied by the model iff the conditional probability $Pr(\alpha/\beta) > 1 - \epsilon$, for some choice of a positive ϵ one obtains a consequence relation that is not well-behaved. For instance, one may have $\alpha \mid\!\sim \beta$ and $\alpha \mid\!\sim \gamma$ satisfied by the model, while $\alpha \mid\!\sim \beta \wedge \gamma$ is not satisfied. If, on the other one hand, one chooses ϵ to be 0, one obtains a well-behaved consequence relation, but this relation is always monotonic, and the entailment defined is classical entailment (read $\mid\!\sim$ as material implication). J. McCarthy told us he suggested considering non-standard probabilistic models long ago, but, as far as we know, this suggestion has not been systematically pursued.

The structure of this section is as follows: first we shall, briefly, survey the basic notions of non-standard analysis. We shall also introduce non-standard probability spaces. Then we shall introduce non-standard probabilistic models for non-monotonic reasoning, define the consequence relation given by such a model and prove that any consequence relation given by a non-standard probability model is rational. Lastly we shall show that the axioms are complete for this interpretation, i.e., any rational consequence relation can be represented as the consequence relation given by some non-standard probability model, at least in the case the language $\mathcal{L}$ is countable. If $\mathcal{L}$ is not countable, an easy counter-example shows the result does not hold, but we shall not elaborate in this paper.

B.2 Non-Standard Analysis

Non-standard analysis was invented by Abraham Robinson in order to give a rigorous development of analysis in which limiting processes are replaced by behaviour at the infinitesimally small, e.g., the derivative becomes a quotient of the change in the function divided by the change in the argument, when the argument is infinitesimally increased. In this section we shall give a very brief introduction to the basic ideas. The reader interested in a full treatment can consult A. Robinson's [31] or Keisler's [16] books on the topic. More advanced topics related to non-standard probability theory are surveyed in [6].

The basic idea of non-standard analysis is to extend the real numbers to a larger ordered field while preserving many of the basic properties of the reals. Therefore, we consider a structure of the form

$$\mathcal{R}^* = \langle R^*, +^*, \times^*, <^*, 0, 1\rangle$$

such that $\mathcal{R}^*$ is an elementary extension of the standard real numbers, namely $\mathbf{R} \subset R^*$, the operations and the order relation of $\mathcal{R}^*$ extend those of $\mathbf{R}$ and for every first order formula Φ

$$\mathcal{R}^* \models \Phi(x_1, \ldots, x_n) \text{ iff } \mathbf{R} \models \Phi(x_1, \ldots, x_n)$$

for $x_1, \ldots, x_n \in \mathbf{R}$. Since we would like to consider not only properties of the real numbers, but real valued functions, functions from real valued functions into reals, and so on, we shall consider a richer structure: the superstructure of the real numbers.

Definition 27 *The superstructure of the set X is $V_\infty(X) = \bigcup_{n=0}^{\infty} V_n$ where V_n are defined by induction:*

- $V_0 = X$
- $V_{n+1} = \mathcal{P}(V_n) \cup V_n$ *where $\mathcal{P}(Y)$ is the power set of Y.*

Note that the superstructure of X contains all the relations on X, all n-valued functions from X into X, etc. In a non-standard model of the real numbers we would like to have a non-standard counterpart to any standard member of the superstructure of the real numbers. Note that the set theoretical relation $\in$ makes sense in the superstructure of X. Recall that a formula of the first order language having only $\in$ as a non logical constant is called bounded if is constructed by the usual connectives and *bounded* quantifiers, namely $(\forall x \in y)$ and $(\exists x \in y)$ meaning respectively: $\forall x$ *if* $x \in y$ *then* ... and $\exists x\ x \in y \wedge \ldots$.

Definition 28 *A non-standard model of analysis is an ordered field $\mathcal{R}^*$ that is a* proper *extension of the ordered field of the reals, together with a map * from the superstructure of $\mathbf{R}$ into the superstructure of $\mathcal{R}^*$, such that for every* bounded *formula $\Phi(x_1 \ldots x_n)$:*

$$V_\infty(\mathbf{R}) \models \Phi(a_1 \ldots a_n) \textit{ iff } V_\infty(\mathcal{R}^*) \models \Phi(a_1^* \ldots a_n^*) \ (\textbf{Leibniz Principle})$$

and such that for $x \in \mathbf{R}$ $x^* = x$ (we assume that * transforms the standard operations of $\mathbf{R}$ into those of $\mathcal{R}^*$).

The Leibniz principle guarantees that the non-standard counterpart of any standard notion (namely its *) preserves many of the properties of the standard object. In particular it is an object of the same kind: for example if A is a set of functions from **R** to **R**, then A^* is a set of functions from $\mathcal{R}^*$ into $\mathcal{R}^*$. As another example consider the absolute value as a function from **R** to **R**. In **R** it has the property

$$(\forall x \in \mathbf{R})(\mid x \mid \geq 0 \wedge (\mid x \mid = 0 \leftrightarrow x = 0))$$

Then by the Leibniz principle

$$(\forall x \in \mathcal{R}^*)\,(\mid x \mid^* \geq^* 0 \wedge (\mid x \mid^* = 0 \leftrightarrow x = 0))$$

In fact since the * versions of the standard arithmetic operations and relations (like $\leq, \geq, >, <$) are so similar to the standard ones (they extend them) we shall simplify the notation by dropping the *, letting the context determine whether we mean the standard operation on **R**, or its extension to $\mathcal{R}^*$. The next theorem shows that this is not a formal game:

Theorem 8 (Robinson) *There exists a non-standard model for analysis.*

The proof is an application of the compactness theorem. The extension of **R**, $\mathcal{R}^*$, is not unique but nothing in the following arguments depends on the particular choice of the non-standard extension of **R**. So fix one such extension $\mathcal{R}^*$.

Definition 29

1. *$x \in \mathcal{R}^*$, $x \neq 0$, is called* finite *if $\mid x \mid < y$ for some $y \in \mathbf{R}$, or, equivalently, if $\mid x \mid < n$ for some natural number n.*
2. *$x \in \mathcal{R}^*$ is called* infinitesimal *if for all ϵ in $\mathbf{R}$, $\epsilon > 0$, $\mid x \mid < \epsilon$. Following our definition 0 is infinitesimal.*
3. *$x \in V_\infty(\mathcal{R}^*)$ is called* internal *if $x \in y^*$ for some $y \in V_\infty(\mathbf{R})$. The set of internal objects is denoted by V_∞^*.*
4. *$x \in V_\infty(\mathcal{R}^*)$ is standard if $x = y^*$ for some $y \in V_\infty(\mathbf{R})$.*

It follows easily, from the fact that $\mathcal{R}^*$ is a proper extension of **R**, that there are infinitesimal, as well as infinite, members of $\mathcal{R}^*$. In fact x is infinitesimal iff $1/x$ is infinite. If **N** is the set of natural numbers, one can show that **N** is a proper subset of $\mathbf{N}^*$ and every member of $\mathbf{N}^* - \mathbf{N}$ is called a non-standard natural number.

Lemma 43

1. *The sum, product and difference of two infinitesimals is infinitesimal.*
2. *The product of an infinitesimal and a finite member of $\mathcal{R}^*$ is infinitesimal.*

Theorem 9 (Robinson's Overspill Principle) *Let $\langle A_n \mid n \in \mathbf{N}\rangle$ be a sequence of members of $V_k(\mathbf{R})$ for some $k \in \mathbf{N}$. Assume also that, for all $n \in \mathbf{N}$, $A_n \neq \emptyset$ and $A_{n+1} \subseteq A_n$. Then $\bigcap_{n \in \mathbf{N}} A_n^*$ is not empty.*

Sketch of proof : Note that a sequence of elements of $V_k(\mathbf{R})$ can be considered to be a function from $\mathbf{N}$ into $V_k(\mathbf{R})$, and therefore it is a member of $V_\infty(\mathbf{R})$. Hence $\langle A_n \mid n \in \mathbf{N}\rangle^*$ makes sense and it is a function from $\mathbf{N}^*$ into $V_k(\mathbf{R})^*$. Its value at $h \in \mathbf{N}^*$ will be denoted by $(A)^*_h$. Note that $(A)^*_n = A^*_n$ for $n \in \mathbf{N}$. Let $h \in \mathbf{N}^* - \mathbf{N}$. One can easily check, using the Leibniz principle, that $(A)^*_h$ is not empty and that for $n \in \mathbf{N}$ $(A)^*_h \subseteq A^*_n$, hence $\bigcap_{n\in\mathbf{N}} A^*_n$ is not empty. ∎

We can now define the notion of non-standard probability space, which is like a standard (finitely additive) probability space, except that the values of the probability function are in $\mathcal{R}^*$.

Definition 30 *An* $\mathcal{R}^*$-probability space *is a triple* $\langle X, \mathcal{F}, Pr\rangle$ *where* X *is a non-empty set ,* $\mathcal{F}$ *is a Boolean subalgebra of* $\mathcal{P}(X)$*, (namely* $X \in \mathcal{F}$*,* $\emptyset \in \mathcal{F}$*, and* $\mathcal{F}$ *is closed under finite unions, intersections and differences) and* Pr *is a function from* $\mathcal{F}$ *into* $\mathcal{R}^*$ *such that*

1. $Pr(A) \geq 0$ *for* $A \in \mathcal{F}$.
2. $Pr(X) = 1$
3. $Pr(A \cup B) = Pr(A) + Pr(B)$ *for* $A, B \in \mathcal{F}$*,* A *and* B *disjoint*

Note that many of the notions that are usually associated with probability spaces are immediately generalized to $\mathcal{R}^*$-probability space, like independence of 'events' (namely sets in $\mathcal{F}$) and conditional probability: if $Pr(A) \neq 0$ then the conditional probability of B given A, is

$$Pr(B \mid A) = \frac{Pr(A \cap B)}{Pr(A)}.$$

See [6] for sophisticated applications of non-standard probability spaces. A useful way of getting $\mathcal{R}^*$-probability spaces is by using hyperfinite sets, sets which are considered by $\mathcal{R}^*$ to be finite.

Definition 31 *An internal object* $A \in V^*_\infty$ *is called* hyperfinite *iff there exists a function* $f \in V^*_\infty$ *and* $h \in \mathbf{N}^*$ *such that* f *is a 1-1 mapping of* h *onto* A*. Note that we follow the usual set theoretical convention by which a natural number is identified with all smaller natural numbers. Of course here we apply this convention also to non-standard natural numbers.*

By applying the Leibniz principle we can show that if A is hyperfinite and B is an internal subset of A, then B is hyperfinite. Given an $\mathcal{R}^*$-valued function f which is internal, and A an hyperfinite subset of the domain of f, we can naturally define the 'sum' of the values of f on A, $\sum^*_{x\in A} f(x)$. $\sum^*$ is defined by taking the * of the standard operation of taking the sum of a finite set of real numbers. $\sum^*$ shares many of the properties of its standard counterpart, for example

$$\sum\nolimits^*{}_{x\in A\cup B} f(x) = \sum\nolimits^*{}_{x\in A} f(x) + \sum\nolimits^*{}_{x\in B} f(x)$$

for A, B hyperfinite and disjoint. The next definition generalizes the notion of a finite probability space.

Definition 32 (Hyperfinite Probability Space) *Let $A \in V^*_\infty$ be an hyperfinite set, let f be an internal $\mathcal{R}^*$-valued function on A, which is not constantly zero and such that for $x \in A$ $f(x) \geq 0$. Then the $\mathcal{R}^*$-*probability space generated by A and f *(denoted by $PR^*(A, f)$) is $\langle A, \mathcal{F}, Pr \rangle$ where $\mathcal{F}$ is the collection of all internal subsets of A, and* Pr *is given by*

$$Pr(B) = \frac{\sum^*{}_{x\in B} f(x)}{\sum^*{}_{x\in A} f(x)}$$

One can verify that under the conditions of Definition 32, $PR^*(A, f)$ is a $\mathcal{R}^*$-probability space.

B.3 Non-standard Probabilistic Models and Their Consequence Relations

An $\mathcal{R}^*$ probabilistic model is an $\mathcal{R}^*$-probability measure on some subset $\mathcal{M}$ of $\mathcal{U}$. Of course, we assume that for every formula of our language, α, the set $\hat{\alpha}$ is measurable, namely it is in $\mathcal{F}$. The probability measure induces a non-standard probability assignment to the formulas of the language by $Pr(\alpha) = Pr(\hat{\alpha})$. The $\mathcal{R}^*$ probabilistic model $\mathcal{M}$ is said to be *neat* if for every formula, α, if $Pr(\alpha) = 0$ then α is satisfied in no world of $\mathcal{M}$.

Definition 33 *1. Let $\mathcal{M}$ be an $\mathcal{R}^*$ probabilistic model. The conditional assertion $\alpha \mathrel{|\!\sim} \beta$ is valid in $\mathcal{M}$, $\mathcal{M} \models \alpha \mathrel{|\!\sim} \beta$, if either $Pr(\alpha) = 0$ or the conditional probability of β given α is infinitesimally close to 1, i.e., $1 - Pr(\beta \mid \alpha)$ is infinitesimal. Note that this is equivalent to saying that $Pr(\alpha) = 0$ or $Pr(\neg\beta \mid \alpha)$ is infinitesimal.*

2. The consequence relation defined by $\mathcal{M}$ is:

$$K(\mathcal{M}) = \{\alpha \mathrel{|\!\sim} \beta \mid \mathcal{M} \models \alpha \mathrel{|\!\sim} \beta\}$$

Theorem 10 (Soundness for Non-standard Probabilistic Models) *For every $\mathcal{R}^*$ probabilistic model $\mathcal{M}$, $K(\mathcal{M})$ is a rational consequence relation.*

Proof: Left Logical Equivalence, **Right Weakening**, and **Reflexivity** are immediate. **And** follows from:

$$Pr(\neg(\beta \wedge \gamma) \mid \alpha) = Pr((\neg\beta \vee \neg\gamma) \mid \alpha) \leq Pr(\neg\beta \mid \alpha) + Pr(\neg\gamma \mid \alpha)$$

and from the fact that the sum of two infinitesimals is infinitesimal. **Or** is proved by the following manipulation:

$$Pr(\neg\gamma \mid \alpha \vee \beta) \quad = \quad \frac{Pr(\neg\gamma \wedge (\alpha \vee \beta))}{Pr(\alpha \vee \beta)} \leq$$

$$\begin{aligned}\frac{Pr(\neg\gamma\wedge\alpha)}{Pr(\alpha\vee\beta)} &+ \frac{Pr(\neg\gamma\wedge\beta)}{Pr(\alpha\vee\beta)} \leq \\ \frac{Pr(\neg\gamma\wedge\alpha)}{Pr(\alpha)} &+ \frac{Pr(\neg\gamma\wedge\beta)}{Pr(\beta)} = Pr(\neg\gamma \mid \alpha) + Pr(\neg\gamma \mid \beta)\end{aligned}$$

and again using the fact that the sum of two infinitesimals is infinitesimal. We assumed above that $Pr(\alpha) > 0$ and $Pr(\beta) > 0$. If this fails then the argument is easier. We shall prove **Rational Monotonicity** by contradiction, so we assume that $\alpha \mid\!\sim \neg\beta$ is not in K($\mathcal{M}$), and that $\alpha \mid\!\sim \gamma$ is in K($\mathcal{M}$). We shall prove that $\alpha\wedge\beta \mid\!\sim \gamma$ is in K($\mathcal{M}$). We can assume that $Pr(\alpha\wedge\beta) > 0$ (hence $Pr(\alpha) > 0$) otherwise the argument is trivial.

$$\begin{aligned}Pr(\neg\gamma \mid \alpha\wedge\beta) &= \frac{Pr(\neg\gamma\wedge\alpha\wedge\beta)}{Pr(\alpha\wedge\beta)} = \\ \frac{Pr(\neg\gamma\wedge\alpha\wedge\beta)}{Pr(\alpha)} &\Big/ \frac{Pr(\alpha\wedge\beta)}{Pr(\alpha)} \leq \\ \frac{Pr(\neg\gamma\wedge\alpha)}{Pr(\alpha)} &\Big/ \frac{Pr(\alpha\wedge\beta)}{Pr(\alpha)} = \\ Pr(\neg\gamma \mid \alpha) &\times \frac{1}{Pr(\beta \mid \alpha)}\end{aligned} \tag{21}$$

Since $\alpha \mid\!\sim \neg\beta$ is not in $K(\mathcal{M})$, we get that $Pr(\beta \mid \alpha)$ is not infinitesimal, hence $\frac{1}{Pr(\beta|\alpha)}$ is finite. By Lemma 43 $Pr(\neg\gamma\wedge\alpha) \times \frac{1}{Pr(\beta|\alpha)}$ is infinitesimal. Hence by Equation 21, $\alpha\wedge\beta \mid\!\sim \gamma$ is in $K(\mathcal{M})$. **Cautious Monotonicity** now follows easily. Suppose $\alpha \mid\!\sim \gamma$ and $\alpha \mid\!\sim \beta$ are both in $K(\mathcal{M})$. If $\alpha \mid\!\sim \neg\beta$ is not in $K(\mathcal{M})$, we conclude by **Rational Monotonicity**. If $\alpha \mid\!\sim \neg\beta$ is in $K(\mathcal{M})$, we must have $Pr(\alpha) = 0$, since $Pr(\beta \mid \alpha)$ and $Pr(\neg\beta \mid \alpha)$ cannot be both infinitesimally close to 1. Therefore $Pr(\alpha\wedge\beta) = 0$ and we conclude that $\alpha\wedge\beta \mid\!\sim \gamma \in K(\mathcal{M})$. ∎

B.4 Completeness for the Non-Standard Probabilistic Interpretation

Theorem 11 *Suppose the language $\mathcal{L}$ is countable (this assumption cannot be dispensed with) and P is a rational consequence relation on $\mathcal{L}$. Let $\mathcal{R}^*$ be any non-standard model of analysis, then there exists an $\mathcal{R}^*$-probabilistic neat model $\mathcal{M}$ such that $K(\mathcal{M}) = K$.*

Proof: Let $W = \langle S, l, \prec\rangle$, with ranking function r, be a countable (i.e., S is countable) ranked model that defines the consequence relation P. The model built in the proof of Theorem 5 shows that such models exist. If S is finite, or even if each level in W is finite and W is well-founded, one may simply use the construction described just before Lemma 22, with some arbitrary infinitesimal ϵ. In case the model W is infinitely *broad*, i.e., has some level containing an infinite number of states, then the construction has to be slightly more sophisticated, but the real difficulty appears when W is not well-founded, and we have already remarked that there are rational relations that have no well-founded ranked model. Following the proof of Lemma 22 we would like to assign a (non-standard) probability distribution to the states of the model in such a way that the relative probabilty of a level to that of a lower level is infinitesimal, but, for every formula which is satified at a given level, we would like to keep its relative weight within the level non infinitesimal. To each formula we shall assign a positive real number r such that, if the formula is satisfied at level l, its relative probability within this level should be at least r. In order that these requirements not be contradictory, the sum of the r's so assigned should be at most 1. Quite arbitrarily, we pick for the i-th formula $r = 1/2^{i+1}$. Now we have to show that we can find a probability assignment satisfying these requirements. We shall define a set B_n of all probability assignments that are *good up to rank* n. An assignment which is *good for every* n will satisfy our requirements. So, we would like to intersect the B_n's. The overspill principle will tell us that this intersection is not empty.

Since S is countable we may assume that $S = \mathbf{N}$. Since every countable linear ordering may be order embedded into the real numbers, we may assume without loss of generality that the ranking function, r, is into $\mathbf{R}$. Since $\prec$ is a partial ordering of $\mathbf{N}$, $\prec^*$ is a partial ordering of $\mathbf{N}^*$ which is ranked by the ranking function r^* mapping $\mathbf{N}^*$ into $\mathcal{R}^*$.

For each formula α, let $A_\alpha = (\hat{\alpha})^*$. Note that A_α is a subset of $\mathbf{N}^*$ (but must not be a subset of $\mathbf{N}$). We can now associate a world, $\mathcal{U}_h$, with each $h \in \mathbf{N}^*$, defined by $\mathcal{U}_h \models p$ iff $h \in A_p$. It is easily checked that, for standard h (i.e., $h \in \mathbf{N}$), one has $\mathcal{U}_h = l(h)$ and that, for arbitrary h, $\mathcal{U}_h \models \alpha$ iff $h \in A_\alpha$. Our idea now is to find an h in $\mathbf{N}^*$ and an internal function f, from h into $\mathcal{R}^*$ such that, if we consider the probability distribution given by the hyperfinite probability space $PR^*(h, f)$ on the set of worlds $\{\mathcal{U}_k \mid k \in \mathbf{N}^*, k < h\}$, we shall get a probabilistic model whose consequence relation is exactly P (recall that we are identifying a member of $\mathbf{N}^*$ with the set of smaller members of $\mathbf{N}^*$). Fix an enumeration $\langle \alpha_n \mid n \in \mathbf{N} \rangle$ of all the formulas of our language. For $i \in \mathbf{N}$ let x_i be the real number such that the ranking function f maps all the states minimal in $\hat{\alpha_i}$ to it.

We are now going to define a sequence of sets of possible approximations to the object we are looking for, namely the appropriate $h \in \mathbf{N}^*$ and the appropriate f. For $n \in \mathbf{N}$, let B_n be the set of all triples of the form (k, ϵ, f) that have the following properties:

1. $k \geq n$,
2. $\epsilon \in \mathbf{R}$, $\epsilon > 0$, $\epsilon \leq 1/n$,
3. f is a function from $\mathbf{N}$ into $\mathbf{R}$ such that for any $s \in \mathbf{N}$, $f(s) > 0$,

4. for any $x, y \in \mathbf{R}$ such that $x < y$, if x and y are in the range of the ranking function r on k, then

$$\frac{\sum_{m<k, r(m)=y} f(m)}{\sum_{m<k, r(m)=x} f(m)} \leq \epsilon,$$

5. for α_i, $i < k$, if $C = \hat{\alpha_i} \cap \{j \mid r(j) = x_i\} \cap \{0, \ldots, k-1\} \neq \emptyset$, then

$$\frac{\sum_{m \in C, m<k} f(m)}{\sum_{m<k, r(m)=x_i} f(m)} \geq \frac{1}{2^{i+1}}$$

It easily follows from the definition of the sequence of sets $\langle B_n \mid n \in \mathbf{N} \rangle$ that $B_{n+1} \subseteq B_n$ for $n \in \mathbf{N}$. One may also verify from item 4 that, if $(k, \epsilon, f) \in B_n$ and if x is in the range of r on k, then:

$$\frac{\sum_{m<k, r(m)>x} f(m)}{\sum_{m<k, r(m)=x} f(m)} \leq \sum_{i=1}^{\infty} \epsilon^i = \frac{\epsilon}{1-\epsilon} \tag{22}$$

Lemma 44 *For any $n \in \mathbf{N}$, $B_n \neq \emptyset$.*

Proof: The proof is essentially similar to the remarks preceding the proof of Lemma 22 in Section 4.3. Let, indeed, W_n be the finite ranked model defined by $\langle \{0, \ldots, n-1\}, \prec, l \rangle$. We can easily arrange a probability assignment for it such that the ratio of the probability of each rank and and each smaller rank will be at most $1/n$. Within the rank we have to satisfy item 5 in the definition of B_n but we can easily arrange for $i < n$, that if α_i has a non empty intersection with this rank, then its relative probability within this rank is at least $\frac{1}{2^{i+1}}$. This may be arranged because $\sum_{i \in \mathbf{N}} \frac{1}{2^{i+1}} = 1$. If we extend this probability assignment to any function from $\mathbf{N}$ into $\mathbf{R}$, we see that $(n, \frac{1}{n}, f) \in B_n$. ∎

Once we have Lemma 44 we can use Robinson's overspill principle (Theorem 9) to show that $\cap_{n \in \mathbf{N}} B^*_n$ is not empty. So let $(\tilde{h}, \tilde{\varepsilon}, \tilde{f})$ be a member of B^*_n for every $n \in \mathbf{N}$. One can easily verify that $\tilde{h}$ is in $\mathbf{N}^*$ and that it is a non-standard natural number: indeed for every $n \in \mathbf{N}$, $\tilde{h} > n$ since $(\tilde{h}, \tilde{\varepsilon}, \tilde{f}) \in B^*_n$. Similarly $\tilde{\varepsilon}$ is a positive member of $\mathcal{R}^*$ such that for every standard natural number n we have $\tilde{\varepsilon} \leq \frac{1}{n}$, hence $\tilde{\varepsilon}$ is a positive infinitesimal. Also $\tilde{f}$ is a function from $\mathbf{N}^*$ into the positive members of $\mathcal{R}^*$, satisfying the appropriate transfer of items 5 and 4 into the context of $\mathcal{R}^*$. In particular, Equation 22 carries over and we have $x = r^*(m)$ for some $m \in \mathbf{N}^*, m < \tilde{h}$:

$$\frac{{\sum^*}_{m<\tilde{h}, r^*(m)>x} \tilde{f}(m)}{{\sum^*}_{m<\tilde{h}, r^*(m)=x} \tilde{f}(m)} \leq \frac{\tilde{\varepsilon}}{1-\tilde{\varepsilon}} \tag{23}$$

We conclude therefore that the left-hand side of Equation 23 is infinitesimal. We claim that the $\mathcal{R}^*$ probabilistic model $\mathcal{M}$ whose collection of states is $\tilde{h}$, i.e.,

$\{m \mid m \in \mathbf{N}^*, m < \tilde{h}\}$, the world associated with m is $\mathcal{U}_m$, and the probability measure is given by the hyperfinite probability space $PR^*(\tilde{h}, \tilde{f})$ is the model we are looking for. Since any f satisfying the requirements can be multiplied by any positive member of $\mathcal{R}^*$ and still satisfies the requirements, we may assume without loss of generality that

$$\sum\nolimits^*_{m \in \mathbf{N}^*, m < \tilde{h}} \tilde{f}(m) = 1$$

Note that $\mathcal{M}$ is a neat model since, if we have both $m < \tilde{h}$ and $\mathcal{U}_m \models \alpha$, we must also have $Pr(\alpha) \geq \tilde{f}(m) > 0$.

Claim 1 $K(\mathcal{M}) = K$

Proof: First note that $\alpha \mathrel{|\!\sim} \mathbf{false} \in K(\mathcal{M})$ iff $\alpha \mathrel{|\!\sim} \mathbf{false} \in K$. If $\alpha \mathrel{|\!\sim} \mathbf{false} \in K$, then $A_\alpha = \emptyset$, hence $\{m \mid m \in \mathbf{N}^*, \mathcal{U} \models \alpha\} = \emptyset$. Therefore $\alpha \mathrel{|\!\sim} \mathbf{false} \in K(\mathcal{M})$. For the other direction, if $\alpha \mathrel{|\!\sim} \mathbf{false} \notin K$, then, for some $m \in \mathbf{N}$, $U_m \models \alpha$. But m, being a standard natural number, is less than $\tilde{h}$, hence some state in $\mathcal{M}$ satisfies α. By the neatness of $\mathcal{M}$, $\alpha \mathrel{|\!\sim} \mathbf{false} \notin K(\mathcal{M})$. By the previous remark, we can now assume that $\alpha \mathrel{|\!\sim} \mathbf{false} \notin K$, hence $\hat{\alpha} \neq \emptyset$. Let $m \in \mathbf{N}$ be minimal in $\hat{\alpha}$ and let $x = r(m)$. Let $i \in \mathbf{N}$ be such that $\alpha = \alpha_{i-1}$. In particular we have:

$$(\forall y \in \mathbf{R})(y < x \Rightarrow r^{-1}(y) \cap \hat{\alpha}) = \emptyset.$$

Using the Leibniz principle we get:

$$\{m \mid m \in \mathbf{N}^*, m < \tilde{h}, \mathcal{U}_m \models \alpha\} \subseteq \{m \mid m \in \mathbf{N}^*, r^*(m) \geq x\}.$$

Let us define now

$$\eta = \sum\nolimits^*_{m \in \tilde{h}, r^*(m) > x} \tilde{f}(m)$$

and

$$\rho = \sum\nolimits^*_{m \in \tilde{h}, r^*(m) = x} \tilde{f}(m).$$

For every formula γ define:

$$\lambda(\gamma) = \sum\nolimits^*_{m \in \tilde{h}, r^*(m) = x, \mathcal{U}_m \models \gamma} \tilde{f}(m)$$

Of course one always has $Pr(\gamma) \geq \lambda(\gamma)$. Note that by Equation 23, η/ρ is infinitesimal. Also by item 5 of the definition of the sequence $\langle B_n \mid n \in \mathbf{N}\rangle$ if $A_\gamma \cap \{m \mid m \in \tilde{h}, r^*(m) = x\} \neq \emptyset$ and if $\gamma = \alpha_{j-1}, j \in \mathbf{N}$ then

$$\lambda(\gamma) \geq \rho \times \frac{1}{2^j}.$$

Now, assume $\alpha \mathrel{|\!\sim} \beta \in K$. Hence for every $m \in \tilde{h}$, if $\mathcal{U} \models \neg\beta \wedge \alpha$, we must have $r^*(m) > x$. Therefore $Pr(\neg\beta \wedge \alpha) \leq \eta$. Therefore:

$$Pr(\neg\beta \mid \alpha) = \frac{Pr(\neg\beta \wedge \alpha)}{Pr(\alpha)} \leq \frac{\eta}{\rho \times \frac{1}{2^i}} = 2^i \times \frac{\eta}{\rho}.$$

Therefore $Pr(\neg\beta \mid \alpha)$ is infinitesimal and by definition $\alpha \mid\!\sim \beta \in K(\mathcal{M})$. If $\alpha \mid\!\sim \beta \notin K$, then some $m \in \mathbf{N}$, $r^*(m) = x$ satisfies $\mathcal{U}_m \models \neg\beta \wedge \alpha$. But this m satisfies $m \in \tilde{h}$, so it is in our model. Let $j \in \mathbf{N}$ be such that $\neg\beta \wedge \alpha = \alpha_{j-1}$. Since we clearly have: $Pr(\alpha) \leq \rho + \eta$, we also have:

$$Pr(\neg\beta \mid \alpha) = \frac{Pr(\neg\beta \wedge \alpha)}{Pr(\alpha)} \geq \frac{\frac{1}{2^j} \times \rho}{\rho + \eta} \geq \frac{1}{2^{j+1}}$$

since obviously $\rho \geq \eta$. So $Pr(\neg\beta \mid \alpha)$ is *not* infinitesimal and $\alpha \mid\!\sim \beta \notin K(\mathcal{M})$. (end of proof of Claim 1) ∎

We have already noticed that $\mathcal{M}$ is a neat model. Claim 1 shows that it has the desired property. (end of proof of Theorem 11) ∎