A nonmonotonic extension of KLM Preferential Logic P

Laura Giordano[1], Valentina Gliozzi[2], Nicola Olivetti[3], and Gian Luca Pozzato[2]

[1] Dip. di Informatica - U. Piemonte O. - Alessandria - Italy - laura@mfn.unipmn.it
[2] Dip. Informatica - Univ. di Torino - Italy {gliozzi,pozzato}@di.unito.it
[3] LSIS-UMR CNRS 6168 - Marseille - France - nicola.olivetti@univ-cezanne.fr

Abstract. In this paper, we propose the logic $\mathbf{P}_{min}$, which is a nonmonotonic extension of Preferential logic $\mathbf{P}$ defined by Kraus, Lehmann and Magidor (KLM). In order to perform nonmonotonic inferences, we define a "minimal model" semantics. Given a modal interpretation of a minimal A-world as $A \wedge \Box\neg A$, the intuition is that preferred, or minimal models are those that minimize the number of worlds where $\neg\Box\neg A$ holds, that is of A-worlds which are not minimal. We also present a tableau calculus for deciding entailment in $\mathbf{P}_{min}$.

1 Introduction

In the early 90s [1] Kraus, Lehmann and Magidor (from now on KLM) proposed a formalization of nonmonotonic reasoning that was early recognized as a landmark. Their work led to a classification of nonmonotonic consequence relations, determining a hierarchy of stronger and stronger systems. The so called *KLM properties* have been widely accepted as the "conservative core" of default reasoning: they are properties that any concrete reasoning mechanism should satisfy. In KLM framework, defeasible knowledge is represented by a (finite) set of nonmonotonic conditionals or assertions of the form $A \mid\sim B$, whose reading is - depending on the context - *typically, As are Bs* or *normally, if A is true, also B is true*. By using a conditional, we can therefore express sentences as *artists are typically not rich* or *normally, if students work, they pass the exams*. The operator "$\mid\sim$" is nonmonotonic, in the sense that $A \mid\sim B$ does not imply $A \wedge C \mid\sim B$. By using the operator $\mid\sim$, one can consistently represent information that would be inconsistent, if interpreted in classical terms. For instance, a knowledge base Γ may consistently contain the conditionals: $artist \mid\sim \neg rich$, $artist \wedge successful \mid\sim rich$, expressing the fact that typically artists are not rich, except if they are successful, in which case they are rich. Observe that if $\mid\sim$ were interpreted as classical (or intuitionistic) implication, the knowledge base would be consistent only in case successful artists did not exist, which is clearly an unwanted condition.

In KLM framework, one can derive new conditional assertions from the knowledge base by means of a set of inference rules. The set of adopted rules defines some fundamental types of inference systems. The two systems that had the biggest echo in the literature are Preferential **P** and Rational **R** ([2–7], for a broader bibliography see [8]). Halpern and Friedman [9] have shown that **P** and **R** are natural and general systems: **P** (likewise **R**) is complete with respect to a wide spectrum of semantics, from ranked models, to parametrized probabilistic structures, ϵ-semantics and possibilistic structures. As an example of inference, from the knowledge base Γ above, in **P** one would derive that

$artist \mid\sim \neg successful$, i.e. typically artists are not successful (successful artists being an exception).

From a semantic point of view, to each logic corresponds a kind of models, namely a class of possible-world structures equipped with a preference relation among worlds. More precisely, for **P** we have models with a preference relation $<$ (an irreflexive and transitive relation) on worlds. For the stronger **R** the preference relation is further assumed to be *modular*. In both cases, the meaning of a conditional assertion $A \mid\sim B$ is that B holds in the *most preferred* (or *minimal* or *typical*) worlds where A holds.

The main weakness of KLM systems is that they are *monotonic*: what can be inferred from a set of conditional assertions Γ can still be inferred from any set of assertions Γ' that includes Γ. In contrast, nonmonotonic inferences allow to draw a conclusion from a knowledge base Γ *in the absence of information to the contrary*, and to possibly withdraw the conclusion on the basis of more complete information. In the example above, if we knew that typically art students are artists, and indeed non successful artists, a nonmonotonic mechanism would allow us to conclude that they are not rich, similarly to typical artists. And it would allow us to make this conclusion in a nonmonotonic way, thus leaving us the freedom to withdraw our conclusion in case we learned that art students are atypical artists and are indeed rich. In order to make this kind of inference, the core set of KLM properties must be enriched with *nonmonotonic* inference mechanisms. This is the purpose of this paper.

Here, we take Preferential logic **P** as the underlying monotonic system in KLM framework. The choice of considering **P** rather than the stronger **R** is motivated by the exigence of avoiding some counterintuitive inferences supported by **R** as highlighted in [10]. This choice is also one of the main differences between our approach and existing nonmonotonic extensions of KLM systems, such as rational closure [4] and 1-entailment [3] that are based on the system **R**.

In this paper, we present the logic $\mathbf{P}_{min}$ that results from the addition to the system **P** of a nonmonotonic inference mechanism. $\mathbf{P}_{min}$ is based on a *minimal model approach*, based on the idea of restricting one's attention to models that contain as little as possible of non typical (or non minimal) worlds. The minimal models associated with a given knowledge base Γ being independent from the minimal models associated to Γ' for $\Gamma \subseteq \Gamma'$, $\mathbf{P}_{min}$ allows for nonmonotonic inferences. In the example above, in $\mathbf{P}_{min}$, we would derive that typically, art students are not rich ($art\_students \mid\sim \neg rich$). Moreover, we would no longer make the inference, should we discover that indeed $art\_students \mid\sim rich$.
We provide a decision procedure for checking satisfiability and validity in $\mathbf{P}_{min}$. Our decision procedure has the form of a tableaux calculus, with a two-step construction: the idea is that the top level construction generates open branches that are candidates to represent models, whereas the auxiliary construction checks whether a candidate branch represents a minimal model. Our procedure can be used to determine constructively an upper bound of the complexity of $\mathbf{P}_{min}$. Namely, we obtain that checking entailment for $\mathbf{P}_{min}$ is in Π_2, thus it has the same complexity as standard nonmonotonic (skeptical) mechanisms.

2 KLM Preferential Logic P

In this section, we recall the axiomatizations and semantics of the KLM logic **P**. For a complete picture of KLM systems, see [1, 4]. The language of KLM logics consists just of conditional assertions $A \mid\sim B$. We consider here a richer language allowing boolean combinations of assertions and propositional formulas. As a consequence, we can have in our knowledge base negated conditionals, and we can handle more complex inferences than what can be done within the original KLM framework. Our language $\mathcal{L}$ is defined from a set of propositional variables ATM, the boolean connectives and the conditional operator $\mid\sim$. We use $A, B, C, ...$ to denote propositional formulas (that do not contain conditional formulas), whereas $F, G, ...$ are used to denote all formulas (including conditionals); $\Gamma, \Delta, ...$ represent sets of formulas. The formulas of $\mathcal{L}$ are defined as follows: if A is a propositional formula, $A \in \mathcal{L}$; if A and B are propositional formulas, $A \mid\sim B \in \mathcal{L}$; if F is a boolean combination of formulas of $\mathcal{L}$, $F \in \mathcal{L}$.

The axiomatization of **P** consists of all axioms and rules of propositional calculus together with the following axioms and rules. We use $\vdash_{PC}$ to denote provability in the propositional calculus, and $\vdash$ to denote provability in **P**:

- REF. $A \mid\sim A$ (reflexivity)
- LLE. If $\vdash_{PC} A \leftrightarrow B$, then $\vdash (A \mid\sim C) \rightarrow (B \mid\sim C)$ (left logical equivalence)
- RW. If $\vdash_{PC} A \rightarrow B$, then $\vdash (C \mid\sim A) \rightarrow (C \mid\sim B)$ (right weakening)
- CM. $((A \mid\sim B) \wedge (A \mid\sim C)) \rightarrow (A \wedge B \mid\sim C)$ (cautious monotonicity)
- AND. $((A \mid\sim B) \wedge (A \mid\sim C)) \rightarrow (A \mid\sim B \wedge C)$
- OR. $((A \mid\sim C) \wedge (B \mid\sim C)) \rightarrow (A \vee B \mid\sim C)$

REF states that A is always a default conclusion of A. LLE states that the syntactic form of the antecedent of a conditional formula is irrelevant. RW describes a similar property of the consequent. This allows to combine default and logical reasoning [9]. CM states that if B and C are two default conclusions of A, then adding one of the two conclusions to A will not cause the retraction of the other conclusion. AND states that it is possible to combine two default conclusions. OR states that it is allowed to reason by cases: if C is the default conclusion of two premises A and B, then it is also the default conclusion of their disjunction.

The semantics of **P** is defined by considering possible world structures with a *preference relation* $w < w'$ among worlds, whose meaning is that w is preferred to (or more typical than) w'. $A \mid\sim B$ holds in a model $\mathcal{M}$ if B holds in all *minimal worlds* (with respect to the relation $<$) where A holds.

Definition 1 ((Preferential models). A preferential *model is a triple* $\mathcal{M} = \langle \mathcal{W}, <, V \rangle$, *where* $\mathcal{W}$ *is a non-empty set of items that we call worlds or individuals, according to the kind of knowledge we want to model,* $<$ *is an irreflexive and transitive relation on* $\mathcal{W}$, *and* V *is a function* $V : \mathcal{W} \longmapsto 2^{ATM}$, *which assigns to every world* w *the set of atoms holding in that world.*
Starting from V, *we define the evaluation function* $\models_{\mathbf{P}}$ *for all formulas.* $\models_{\mathbf{P}}$ *is defined in the standard way for boolean combinations of formulas. Furthermore, for* A *propositional, we define* $Min_<(A) = \{w \in \mathcal{W} \mid \mathcal{M}, w \models_{\mathbf{P}} A$ *and* $\forall w'$, $w' < w$ *implies* $\mathcal{M}, w' \not\models_{\mathbf{P}} A\}$. *We define:* $\mathcal{M}, w \models_{\mathbf{P}} A \mid\sim B$ *if for all* w', *if* $w' \in Min_<(A)$ *then* $\mathcal{M}, w' \models_{\mathbf{P}} B$.

The relation $<$ is assumed to satisfy the smoothness condition*: if $\mathcal{M}, w \models_{\mathbf{P}} A$, then $w \in Min_<(A)$ or $\exists w' \in Min_<(A)$ s.t. $w' < w$.*

A formula is valid in a model *$\mathcal{M}$ ($\mathcal{M} \models_{\mathbf{P}} F$) if it is satisfied by all worlds in $\mathcal{M}$, and it is simply* valid *if it is valid in every model. A formula is* satisfiable *if its negation is not valid. $\mathcal{M}$ is a model of a set of formulas Γ ($\mathcal{M} \models_{\mathbf{P}} \Gamma$) if all formulas in Γ are valid in $\mathcal{M}$.*

Observe that the above definition of preferential model extends the definition by KLM in order to cope with boolean combinations of formulas. Notice also that the truth conditions for conditional formulas are given with respect to individual possible worlds for uniformity sake. Since the truth value of a conditional only depends on global properties of $\mathcal{M}$, we have that: $\mathcal{M}, w \models_{\mathbf{P}} A \mid\!\sim B$ iff $\mathcal{M} \models_{\mathbf{P}} A \mid\!\sim B$. Moreover, by the smoothness condition, we have that $Min_<(A) = \emptyset$ implies that no world in $\mathcal{M}$ satisfy A. Thus the logic can represent genuine S5-existential and universal modalities, given respectively by $\neg(A \mid\!\sim \bot)$ and $\neg A \mid\!\sim \bot$. Finally:

Definition 2 (Entailment in P). *A formula F is entailed in* **P** *by Γ ($\Gamma \models_{\mathbf{P}} F$) if it is valid in all models of Γ.*

In this paper, we consider a special kind of preferential models, based on multi-linear frames, that is to say structures where worlds are ordered in several linear chains. As shown in [8], the restriction to this kind of special models is fully legitimate, as they validate the same formulas as general models as defined in Definition 1 above. This restriction is partially motivated by algorithmic and efficiency considerations, as our tableaux calculus is based on them.

Definition 3 (Preferential multi-linear model). *A finite preferential model $\mathcal{M} = (\mathcal{W}, <, V)$ is multi-linear if the set of worlds $\mathcal{W}$ can be partitioned into a set of components $\mathcal{W}_i$ for $i = 1, \dots, n$, that is $\mathcal{W} = \mathcal{W}_1 \cup \dots \cup \mathcal{W}_n$ and 1) the relation $<$ is a total order on each $\mathcal{W}_i$; 2) elements in two different components $\mathcal{W}_i$ and $\mathcal{W}_j$ are incomparable w.r.t. $<$.*

Theorem 1 (Theorem 2.5 in [8]). *Let Γ be a set of formulas, if Γ is satisfiable with respect to the logic* **P***, then it has a multi-linear model.*

From now on, we shall refer to this definition of preferential model **P**.

For the purpose of the calculus, we consider an extension $\mathcal{L}^{\mathbf{P}}$ of the language $\mathcal{L}$ by formulas of the form $\Box A$, where A is propositional, whose intuitive meaning is that $\Box A$ holds in a world w if A holds in all the worlds preferred to w (i.e. in all w' such that $w' < w$). This language corresponds to the language of the tableaux calculus for KLM logics introduced in [8]. We extend the notion of preferential model to provide an evaluation of boxed formulas as follows:

Definition 4 (Truth condition of modality $\Box$). *$\mathcal{M}, w \models_{\mathbf{P}} \Box A$ if, for every $w' \in \mathcal{W}$, if $w' < w$ then $\mathcal{M}, w' \models_{\mathbf{P}} A$.*

From definition of $Min_<(A)$, it follows that for any formula A, we have that $w \in Min_<(A)$ iff $\mathcal{M}, w \models_{\mathbf{P}} A \land \Box\neg A$.

3 The Logic $\mathbf{P}_{min}$

$\mathbf{P}_{min}$ is a nonmonotonic extension of $\mathbf{P}$, based on the idea of considering only $\mathbf{P}$ models that, roughly speaking, minimize the non-typical objects (or the non typical worlds). Given a set of formulas Γ, we consider a finite set $\mathcal{L}_\Box$ of formulas: these are the formulas whose non typical instances we want to minimize, when considering which inferences we can draw from Γ. We assume that the set $\mathcal{L}_\Box$ contains at least all formulas A such that the conditional $A \mid\sim B$ occurs in Γ. We have seen that, in a model $\mathcal{M}$, w is a minimal (typical) A-world, i.e. $w \in Min_<(A)$, when w is an $A-$world (i.e., $\mathcal{M}, w \models_{\mathbf{P}} A$) and it is minimal (typical), i.e., $\mathcal{M}, w \models_{\mathbf{P}} \Box\neg A$. Minimizing non typical (non minimal) worlds w.r.t. $\mathcal{L}_\Box$ therefore amounts to minimizing worlds satisfying $\neg\Box\neg A$ for $A \in \mathcal{L}_\Box$. We define the set $\mathcal{M}^{\Box^-}_{\mathcal{L}_\Box}$ of negated boxed formulas holding in a model, relative to the formulas in $\mathcal{L}_\Box$. Given a model $\mathcal{M} = \langle \mathcal{W}, <, V\rangle$, we let: $\mathcal{M}^{\Box^-}_{\mathcal{L}_\Box} = \{(w, \neg\Box\neg A) \mid \mathcal{M}, w \models_{\mathbf{P}} \neg\Box\neg A, \text{ with } w \in \mathcal{W}, A \in \mathcal{L}_\Box\}$.

Definition 5 (Preferred and minimal models). *Given a model $\mathcal{M} = \langle \mathcal{W}_\mathcal{M}, <_\mathcal{M}, V_\mathcal{M}\rangle$ of a given set of formulas Γ and a model $\mathcal{N} = \langle \mathcal{W}_\mathcal{N}, <_\mathcal{N}, V_\mathcal{N}\rangle$ of Γ, we say that $\mathcal{M}$ is preferred to $\mathcal{N}$ with respect to $\mathcal{L}_\Box$, and we write $\mathcal{M} <_{\mathcal{L}_\Box} \mathcal{N}$, if: $\mathcal{W}_\mathcal{M} = \mathcal{W}_\mathcal{N}$, and $\mathcal{M}^{\Box^-}_{\mathcal{L}_\Box} \subset \mathcal{N}^{\Box^-}_{\mathcal{L}_\Box}$.*
A model $\mathcal{M}$ is a minimal model *for Γ (with respect to $\mathcal{L}_\Box$) if it is a model of Γ and there is no a model $\mathcal{M}'$ of Γ such that $\mathcal{M}' <_{\mathcal{L}_\Box} \mathcal{M}$.*

We can now define entailment in $\mathbf{P}_{min}$ as follows:

Definition 6 (Minimal Entailment in $\mathbf{P}_{min}$). *A formula F is minimally entailed in $\mathbf{P}_{min}$ by a set of formulas Γ with respect to $\mathcal{L}_\Box$ if it is valid in all models satisfying Γ that are minimal with respect to $\mathcal{L}_\Box$. We write $\Gamma \models^{\mathcal{L}_\Box}_{min} F$.*

It can be easily verified that this notion of entailment is nonmonotonic, as opposed to entailment in $\mathbf{P}$ that was monotonic. Indeed, in $\mathbf{P}_{min}$ the formulas minimally entailed by Γ might no longer be entailed by Γ': $\Gamma \subset \Gamma'$. Some significant properties of $\mathbf{P}_{min}$ are expressed in the next propositions. In the following, $\models_{PC}$ is entailment in classical logic.

Proposition 1. *Take a set of formulas Γ and a formula F. (i) If Γ has a model, then Γ has a minimal model with respect to any $\mathcal{L}_\Box$. (ii) Let us replace all formulas of the form $A \mid\sim B$ in Γ with $A \rightarrow B$, and call Γ' the resulting set of formulas. Similarly let F' be obtained from F by replacing conditional $\mid\sim$ by classical implication $\rightarrow$. If $\Gamma \models^{\mathcal{L}_\Box}_{min} F$ then $\Gamma' \models_{PC} F'$.*

Proof. (i) is follows immediately from the finite model property of the logic P: given a finite model $\mathcal{M}$ of Γ the set of models $\mathcal{M}' <_{\mathcal{L}_\Box} \mathcal{M}$ is finite thus there exists a minimal model. Concerning (ii) let v be a classical model of Γ' (a propositional evaluation), define the trivial multi-linear model $\mathcal{M} = \langle \mathcal{W}_\mathcal{M}, <_\mathcal{M}, V_\mathcal{M}\rangle$, where $\mathcal{W}_\mathcal{M} = \{x_v\}$, $<_\mathcal{M} = \emptyset$, $V_\mathcal{M}(x_v) = \{P | v(P) = 1\}$. Observe that for every propositional formula G, $Min_<(G) = \{x_v\}$ if $v(G) = 1$ and $Min_<(G) = \emptyset$ otherwise. For every $A \mid\sim B \in \Gamma$, we know that $v(A \rightarrow B) = 1$; thus trivially for all $w' \in Min_<(A)$, $\mathcal{M}, w' \models_{\mathbf{P}} B$, that is $\mathcal{M}, x_v \models_{\mathbf{P}} A \mid\sim B$. Moreover, $\mathcal{M}^{\Box^-}_{\mathcal{L}_\Box} = \emptyset$.

Thus $\mathcal{M}$ is a P-minimal model of Γ. Now we just apply the hypothesis and we obtain that $\mathcal{M}, x_v \models_{\mathbf{P}} F$, whence $v(F) = 1$. ■

The proposition (ii) states a kind of coherence with respect to classical logic: the inferences we make according to $\mathbf{P}_{min}$ are always a subset of inferences that we would make classically by interpreting $\mid\!\sim$ as classical implication. Obviously, the converse of (ii) is generally false. However, when Γ is just a set of positive conditionals, and F is a positive conditional, the converse of (ii) also holds.

Proposition 2. *Let Γ be a set of positive conditionals and let us denote by Γ' its classical counterpart as in the previous proposition. (i) if Γ' is inconsistent then also Γ is inconsistent. (ii) if $\Gamma' \models_{PC} A \rightarrow B$ then $\Gamma \models^{\mathcal{L}_\Box}_{min} A \mid\!\sim B$.*

Proof. (i) We proceed by contrapositive. Suppose Γ has a $\mathbf{P}$ model $\mathcal{M} = \langle \mathcal{W}_\mathcal{M}, <_\mathcal{M}, V_\mathcal{M} \rangle$. Take any world $x \in \mathcal{W}$ that is minimal wrt. $<$ (thus it has no predecessors); for every formula C we have $\mathcal{M}, x \models_{\mathbf{P}} \Box\neg C$. Thus for every formula C, $x \in Min_<(C)$ iff $\mathcal{M}, x \models_{\mathbf{P}} C$. Since $\mathcal{M} \models_{\mathbf{P}} \Gamma$, for every $C \mid\!\sim D \in \Gamma$, for every $w \in \mathcal{W}$, if $w \in Min_<(C)$ then $\mathcal{M}, w \models_{\mathbf{P}} D$. Therefore, if $\mathcal{M}, x \models_{\mathbf{P}} C$, since $x \in Min_<(C)$, we obtain $\mathcal{M}, x \models_{\mathbf{P}} D$, that is $\mathcal{M}, x \models_{\mathbf{P}} C \rightarrow D$. Thus x (as a propositional evaluation) satisfies every implication in Γ', and Γ' is consistent.

(ii) Suppose $\Gamma' \models_{PC} A \rightarrow B$. By (i) we can assume that Γ' is consistent, otherwise by (i) also Γ is inconsistent, and the result trivially follows. Therefore let v_0 be a propositional evaluation that satisfies Γ'. Let now $\mathcal{M} = \langle \mathcal{W}_\mathcal{M}, <_\mathcal{M}, V_\mathcal{M} \rangle$ be any $\mathbf{P}_{min}$-minimal model of Γ. Define $\mathcal{M}' = \langle \mathcal{W}_\mathcal{M}, <', V' \rangle$, where $<'$ is the empty relation and for all $w \in \mathcal{W}_\mathcal{M}$, $V'(w) = \{P \mid v_0(P) = 1\}$. We have that for all $w \in \mathcal{W}_\mathcal{M}$ and for every $C \rightarrow D \in \Gamma'$, $\mathcal{M}', w \models_{\mathbf{P}} C \rightarrow D$. Thus every implication of Γ' is valid in $\mathcal{M}'$. But, for every formula C $\mathcal{M}', w \models_{\mathbf{P}} \Box\neg C$, so that $w \in Min^{\mathcal{M}'}_<(C)$ iff $\mathcal{M}', w \models_{\mathbf{P}} C$. We can conclude that $\mathcal{M}' \models_{\mathbf{P}} \Gamma$. Furthermore, by minimality of $\mathcal{M}$ we cannot have $\mathcal{M}' <_{\mathcal{L}_\Box} \mathcal{M}$. Thus for every formula C, we have also $\mathcal{M}, w \models_{\mathbf{P}} \Box\neg C$. By reasoning as above, we infer that $w \in Min^{\mathcal{M}}_<(C)$ iff $\mathcal{M}, w \models_{\mathbf{P}} C$. Since $\mathcal{M} \models_{\mathbf{P}} \Gamma$, this entails that $\mathcal{M} \models_{\mathbf{P}} \Gamma'$. From this, we conclude that $\mathcal{M} \models_{\mathbf{P}} A \rightarrow B$, so that finally $\mathcal{M} \models_{\mathbf{P}} A \mid\!\sim B$. ■

The above properties entail that $\mathbf{P}_{min}$ behaves essentially as classical logic (and in particular $\mid\!\sim$ collapses to classical implication) for knowledge bases that are just a set of *positive conditionals*. As a consequence, $\mathbf{P}_{min}$ behaves monotonically for this restricted kind of knowledge bases, with respect to the *addition of positive conditionals*. Of course this is a very special case.

For all these reasons, the logic $\mathbf{P}_{min}$ turns out to be quite strong. If we consider the knowledge base of the Introduction, namely: $\Gamma = \{$*artist* $\mid\!\sim \neg$*rich*, *artist* $\wedge$*successful* $\mid\!\sim$ *rich*$\}$, $\mathbf{P}_{min}$ would entail that $(artist \wedge successful) \mid\!\sim \bot$: typically, successful artists do not exist. This inference is meaningful if we consider that $\mathbf{P}_{min}$ minimizes the existence of non-typical elements of a model, and successful artists are non typical artists. As a difference with respect to classical logic, though, this is a nonmonotonic inference, since we can consistently add to the knowledge base the information that indeed successful artists exist. Observe that this fact would be expressed by a *negative conditional*.

In any case, it is true that in the meaning of a sentence like "typically, successful artists are rich", it is somewhat entailed that successful artists exist. For this reason, from now on we restrict our attentions to knowledge bases that explicitly include the information that there are typical As for all As such that the sentence "typically As are Bs" ($A \mid\sim B$) is in the knowledge base. This leads us to propose the following working definition:

Definition 7. *A well-behaved knowledge base Γ has the form $\{A_i \mid\sim B_i, \neg(A_i \mid\sim \bot) \mid i = 1, \ldots, n\}$.*

Therefore, the above knowledge base is transformed into the following one: $\Gamma' = \{artist \mid\sim \neg rich, artist \wedge successful \mid\sim rich, \neg(artist \mid\sim \bot), \neg(artist \wedge successful \mid\sim \bot)\}$. In the following, we we will consider well-behaved knowledge bases, although in principle one is free to use a knowledge base as Γ above (and to make the kind of strong inferences we mentioned). Obviously, a knowledge base may contain any formula of $\mathcal{L}$. In particular, it may contain negative conditionals, such as $\neg(artist \mid\sim french)$, saying that it is not true that typically artists are French. In the next section, we shall see that if we have these formulas, the choice of **P** rather than **R** as starting system makes a significant difference.

Example 1. Consider the example of art students mentioned in the Introduction. If we add to the previous knowledge base the information that typically, art students are non successful artists (and they do exist), we obtain the following knowledge base: $\Gamma'' = \Gamma' \cup \{art\_student \mid\sim artist \wedge \neg successful, \neg(art\_student \mid\sim \bot)\}$. It can be seen that $art\_student \mid\sim \neg rich$ is minimally entailed by Γ''.

3.1 Relations between $\mathbf{P}_{min}$ and 1-entailment and rational closure.

This is not the first nonmonotonic extension of KLM systems with a nonmonotonic inference mechanism. Lehmann and Magidor [4] propose the well-known rational closure that is a nonmonotonic mechanism built over the KLM system **R**. Pearl in [3] proposes 1-entailment that is equivalent to rational closure. The first difference between our work and both 1-entailment and rational closure is that we start from KLM system **P** (and add a nonmonotonic mechanism to it), whereas they start from the stronger system **R**.

P versus R as the starting KLM system. Our choice of starting from **P** rather than from **R** is motivated by the fact that **R** is controversial, and enforces counterintuitive inferences, as outlined in [10]. To give some evidence, consider the example above of typical artists not being rich. If we knew that there is a typical American who is an artist but who is rich, and is therefore not a typical artist (i.e. if we added $\neg(american \mid\sim (\neg artist \vee \neg rich))$ to the knowledge base), if we adopt system **R**, we are forced to conclude that $artist \mid\sim \neg american$, saying that there are no typical artists who are American. The problem here is that in **R** we conclude something about all typical artists from what we know about a single (non typical) artist. What is most important is that in system **R**, we are forced to make this conclusion in a monotonic way: the addition of

$\neg(artist \mathrel{|\!\sim} \neg american)$, saying that indeed there are typical artists who are American, would lead to an inconsistent knowledge base.

Expressive power and strength of $\mathbf{P}_{min}$, compared to 1-entailment and rational closure. Further differences between our approach and both rational closure and 1-entailment lie both in expressivity, that is the kind of assertions that can be handled, and in the inferences' strength. First of all, it is not clear how 1-entailment can deal with negated conditionals, whence how it can represent assertions of the type "*not* all typical Americans have a given feature". In contrast, we have seen that negated conditionals are part of the language of $\mathbf{P}_{min}$. Secondly, the inferences that can be done in $\mathbf{P}_{min}$ and in 1-entailment are quite different. For instance, take the classical example of penguins, also considered by Pearl in [3]. In $\mathbf{P}_{min}$, from $penguin \mathrel{|\!\sim} bird, bird \mathrel{|\!\sim} fly, penguin \mathrel{|\!\sim} \neg fly, penguin \mathrel{|\!\sim} arctic$ (typically, penguins are birds, birds fly but penguins do not fly and penguins live in the arctic), we would nonmonotonically derive that $(penguin \wedge \neg arctic) \mathrel{|\!\sim} \bot$: in the absence of information to the contrary, there are no penguins that do not live in the arctic. The inference holds both if we start from the given knowledge base, and if we consider its well-behaved extension. On the contrary, the inference does not hold in 1-entailment, nor in rational closure: in neither of the two one can derive a conditional of the form $A \mathrel{|\!\sim} \bot$ from a consistent knowledge base. On the other hand, in 1-entailment (and rational closure) from the above knowledge base one derives that $bird \wedge white \mathrel{|\!\sim} fly$, whereas this is not derivable in $\mathbf{P}_{min}$ in case we consider the well-behaved extension of the knowledge base (whereas it holds if we consider the knowledge base as it is). Therefore, $\mathbf{P}_{min}$ is significantly different from 1-entailment and rational closure as it can handle a richer language, and support different (sometime stronger) inferences. Indeed, we can prove that:

Proposition 3. *Let Γ be a set of positive conditionals.(i) $\mathbf{P}_{min}$ is strictly stronger than 1-entailment (for inferring positive conditionals); (ii) let Γ' be the well-behaved extension of Γ. $\mathbf{P}_{min}$ and 1-entailment inferences from Γ' are incomparable (for inferring positive conditionals).*

Notice that these differences would remain even if we chose $\mathbf{R}$ as the underlying monotonic system: the obtained logic would still be different from 1-entailment and rational closure.

4 A Tableaux Calculus for $\mathbf{P}_{min}$

In this section we present a tableau calculus for deciding whether a formula F is minimally entailed by a given set of formulas Γ. We introduce a labelled tableau calculus called $\mathcal{TAB}^{\mathbf{P}}_{min}$, which allows to reason about minimal models. $\mathcal{TAB}^{\mathbf{P}}_{min}$ performs a two-phase computation in order to check whether $\Gamma \models^{\mathcal{L}_T}_{min} F$. In particular, the procedure tries to build an open branch representing a minimal model satisfying $\Gamma \cup \{\neg F\}$. In the first phase, a tableau calculus, called $\mathcal{TAB}^{\mathbf{P}}_{PH1}$, simply verifies whether $\Gamma \cup \{\neg F\}$ is satisfiable in a $\mathbf{P}$ model, building candidate models. In the second phase another tableau calculus, called $\mathcal{TAB}^{\mathbf{P}}_{PH2}$, checks whether the candidate models found in the first phase are *minimal* models of Γ, i.e. for each open branch of the first phase, $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ tries to build a "smaller" model of Γ, i.e. a model whose individuals satisfy less formulas $\neg\Box\neg A$ than the corresponding candidate model. Therefore, F is minimally entailed by Γ if there is no open branch in the tableau built by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ (therefore, there are no Preferential models satisfying Γ and $\neg F$) or for each open branch built by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$, there is an open branch built by $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ (therefore, the model corresponding to the branch of $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ is not minimal, and there is a more preferred one corresponding to the open branch by $\mathcal{TAB}^{\mathbf{P}}_{PH2}$). The whole procedure $\mathcal{TAB}^{\mathbf{P}}_{min}$ is described at the end of this section (Definition 10).

The calculus $\mathcal{TAB}^{\mathbf{P}}_{min}$ makes use of labels to represent worlds. We consider a language $\mathcal{L}^{\mathbf{P}}$ and a denumerable alphabet of labels $\mathcal{A}$, whose elements are denoted by x, y, z, ...

4.1 Phase 1: the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH1}$

A tableau of $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ is a tree whose nodes are sets of *labelled* formulas of the form $x : A$ or $x : A \mathrel{|\!\sim} B^L$. L is a list of labels used in order to ensure the termination of the calculus. A branch is a sequence of nodes $\Gamma_1, \Gamma_2, \ldots, \Gamma_n$, where each node Γ_i is obtained from its immediate predecessor Γ_{i-1} by applying a rule of $\mathcal{TAB}^{\mathbf{P}}_{PH1}$, having Γ_{i-1} as the premise and Γ_i as one of its conclusions. A branch is closed if one of its nodes is an instance of (**AX**), otherwise it is open. A tableau is closed if all its branches are closed.

Definition 8 (Truth conditions of labelled formulas of $\mathcal{TAB}^{\mathbf{P}}_{PH1}$). *Given a model $\mathcal{M} = \langle \mathcal{W}, <, V \rangle$ and a label alphabet $\mathcal{A}$, we consider a mapping $I : \mathcal{A} \mapsto \mathcal{W}$. Given a formula α of the calculus $\mathcal{TAB}^{\mathbf{P}}_{min}$, we define $\mathcal{M} \models_{\mathbf{P}_I} x : F$ iff $\mathcal{M}, I(x) \models_{\mathbf{P}} F$. We say that a set of labelled formulas Γ is satisfiable if, for all $\alpha \in \Gamma$, we have that $\mathcal{M} \models_{\mathbf{P}_I} \alpha$, for some model $\mathcal{M}$ and some mapping I.*

In order to verify that a set of formulas Γ is unsatisfiable, we label all the formulas in Γ with a new label x, and verify that the resulting set of labelled formulas has a closed tableau. To this purpose, the rules of the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ are applied until either a contradiction is generated (**AX**) or a model satisfying Γ can be obtained from the resulting open branch. For each conditional formulas $A \mathrel{|\!\sim} B \in \Gamma$, we consider $x : A \mathrel{|\!\sim} B^{\emptyset}$.

The rules of $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ are shown in Figure 1. We define $\Gamma^{M}_{u \to x} = \{x : \neg A, x : \Box\neg A \mid u : \Box\neg A \in \Gamma\}$. Rules $(\mathrel{|\!\sim}^-)$ and $(\Box^-)$ are called *dynamic* since they introduce a new variable in their conclusions. The other rules are called *static.* We do not need any extra rule for the positive occurrences of the $\Box$ operator, since these are taken into account by the computation of $\Gamma^{M}_{u \to x}$. The (cut) rule ensures that, given any formula $A \in \mathcal{L}_{\Box}$, an open branch built by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ contains either $x : \Box\neg A$ or $x : \neg\Box\neg A$ for each label x: this is needed in order to allow $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ to check the minimality of the model corresponding to the open branch, as we will discuss later.

The calculus $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ adopts the following standard strategy: the application of the $(\Box^-)$ rule is postponed to the application of all propositional rules and to the test of whether Γ is an instance of (**AX**) or not. The calculus so obtained is sound and complete with respect to the semantics in Definition 3.

$$(\mathbf{AX})\ \Gamma, x : P, x : \neg P \quad \text{with } P \in ATM \qquad (\wedge^+)\ \frac{\Gamma, x : F \wedge G}{\Gamma, x : F, x : G} \qquad (\wedge^-)\ \frac{\Gamma, x : \neg(F \wedge G)}{\Gamma, x : \neg F \qquad \Gamma, x : \neg G} \qquad (\neg)\ \frac{\Gamma, x : \neg\neg F}{\Gamma, x : F}$$

$$(cut)\ \frac{\Gamma}{\Gamma, x : \Box\neg A \qquad \Gamma, x : \neg\Box\neg A}\ \ x \text{ occurs in } \Gamma,\ A \in \mathcal{L}_\Box$$

$$(\mid\!\sim^+)\ \frac{\Gamma, u : A \mid\!\sim B^L}{\Gamma, u : A \mid\!\sim B^{L,x}, x : \neg A \qquad \Gamma, u : A \mid\!\sim B^{L,x}, x : \neg\Box\neg A \qquad \Gamma, u : A \mid\!\sim B^{L,x}, x : B}\ \ x \text{ occurs in } \Gamma,\ x \notin L$$

$$(\mid\!\sim^-)\ \frac{\Gamma, u : \neg(A \mid\!\sim B)}{\Gamma, x : A, x : \Box\neg A, x : \neg B \qquad \Gamma, y_1 : A, y_1 : \Box\neg A, y_1 : \neg B \qquad \Gamma, y_2 : A, y_2 : \Box\neg A, y_2 : \neg B \quad \cdots \quad \Gamma, y_n : A, y_n : \Box\neg A, y_n : \neg B}\ \ \begin{array}{r} x \text{ new} \\ \text{for each } y_i \text{ occurring in } \Gamma \end{array}$$

$$(\Box^-)\ \frac{\Gamma, u : \neg\Box\neg A}{\Gamma, \Gamma^M_{u\to x}, x : A, x : \Box\neg A \qquad \Gamma, \Gamma^M_{u\to y_1}, y_1 : A, y_1 : \Box\neg A \qquad \Gamma, \Gamma^M_{u\to y_2}, y_2 : A, y_2 : \Box\neg A \quad \cdots \quad \Gamma, \Gamma^M_{u\to y_n}, y_n : A, y_n : \Box\neg A}\ \ \begin{array}{r} x \text{ new} \\ \text{for each } y_i \text{ occurring in } \Gamma\ , u \neq y_i \end{array}$$

Fig. 1. The calculus $\mathcal{TAB}^{\mathbf{P}}_{PH1}$. To save space, we omit the standard rules for $\vee$ and $\rightarrow$.

Theorem 2 (Soundness and Completeness of $\mathcal{TAB}^{\mathbf{P}}_{PH1}$). *Given a set of formulas Γ, it is unsatisfiable iff it has a closed tableau in $\mathcal{TAB}^{\mathbf{P}}_{PH1}$.*

We can show that the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ always terminates, i.e. every tableau built by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ is finite. Similarly to the calculus for **R** introduced in [8], it is easy to observe that it is useless to apply the rule $(\mid\!\sim^+)$ on the same conditional formula more than once in the same world, i.e. by using the same label x. We prevent redundant applications of $(\mid\!\sim^+)$ by keeping track of labels in which a conditional $u : A \mid\!\sim B$ has already been applied in the current branch. To this purpose, we add to each positive conditional the above mentioned list of *used* labels L; we then restrict the application of $(\mid\!\sim^+)$ only to labels not occurring in the corresponding list L. Moreover, the rule $(\mid\!\sim^-)$ can be applied only once for each negated conditional Γ. By virtue of the properties of $\Box$, no other additional machinery is required to ensure termination. The generation of infinite branches due to the interplay between rules $(\mid\!\sim^+)$ and $(\Box^-)$ cannot occur. Indeed, each application of $(\Box^-)$ to a formula $x : \neg\Box\neg A$ (introduced by $(\mid\!\sim^+)$) adds the formula $y : \Box\neg A$ to the conclusion, so that $(\mid\!\sim^+)$ can no longer consistently introduce $y : \neg\Box\neg A$. This is due to the properties of $\Box$ (no infinite descending chains of $<$ are allowed). Finally, the (cut) rule does not affect termination, since it is applied only to the finitely many formulas belonging to $\mathcal{L}_\Box$.

Theorem 3 (Termination of $\mathcal{TAB}^{\mathbf{P}}_{PH1}$). *Let Γ be a set of labelled formulas, then any tableau generated by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ for Γ is finite.*

Let us now conclude this section by refining the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ in order to obtain a systematic procedure that allows the satisfiability problem of a set of formulas Γ to be decided in nondeterministic polynomial time. Let n be the size of the starting set Γ of which we want to verify the satisfiability. The number of applications of the rules is proportional to the number of labels introduced in the tableau. In turn, this is $O(2^n)$ due to the interplay between the rules $(\mid\!\sim^+)$ and $(\Box^-)$. Hence, the complexity of $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ is exponential in n.

Similarly to what done in [8], in order to obtain an **NP** procedure we take advantage of the following facts:

1. Negated conditionals do not interact among themselves, thus they can be handled separately and eliminated always as a first step;
2. We can replace the $(\Box^-)$ by a stronger rule $(\Box^-)_s$ which allows to build *mutlilinear* models, as defined in Definition 3.

Regarding (2), we can adopt the following strengthened version of $(\Box^-)$. We use $\Gamma^{-i\Box}_{u\rightarrow x}$ to denote $\{x : \neg\Box\neg A_j \vee A_j \mid u : \neg\Box\neg A_j \in \Gamma \wedge j \neq i\}$.

$$\frac{\Gamma, u : \neg\Box\neg A_1, u : \neg\Box\neg A_2, \ldots, u : \neg\Box\neg A_m}{\Gamma, x : A_k, x : \Box\neg A_k, \Gamma^M_{u\rightarrow x}, \Gamma^{-k\Box}_{u\rightarrow x} \qquad \ldots \qquad \Gamma, y_i : A_k, y_i : \Box\neg A_k, \Gamma^M_{u\rightarrow y_i}, \Gamma^{-k\Box}_{u\rightarrow y_k}} (\Box^-_s)$$

where x is a new label, for each y_i occuring in Γ and for all $k = 1, 2, \ldots, m$. Rule $(\Box^-_s)$ contains: - m branches, one for each $u : \neg\Box\neg A_k$ in Γ, where a new minimal world x is created for A_k (i.e. $x : A_k$ and $x : \Box\neg A_k$ are added), and for all other $u : \neg\Box\neg A_j$, either $x : A_j$ holds in that world or the formula $x : \neg\Box\neg A_j$ is recorded; - other $m \times l$ branches, where l is the number of labels occurring in Γ, one for each label y_i and for each $u : \neg\Box\neg A_k$ in Γ; in these branches, a given y_i is chosen as a minimal world for A_k, that is to say $y_i : A_k$ and $y_i : \Box\neg A_k$ are added, and for all other $u : \neg\Box\neg A_j$, either $y_i : A_j$ holds in that world or the formula $y_i : \neg\Box\neg A_j$ is recorded.
The rule $(\Box^-_s)$ is sound with respect to multilinear models. The advantage of this rule over the original $(\Box^-)$ rule is that all the negated box formulas labelled by u are treated in one step, introducing only a new label x in (some of) the conclusions, at the price of building an exponential number of branches. Instead, the original rule $(\Box^-)$, introduces one new world *for each* $u : \neg\Box\neg A_k$ and, applied m times on all of them, m new worlds are introduced in each branch.

In the following, we describe a rule application's strategy that allows us to decide the satisfiability of a set of formulas Γ in non-deterministic polynomial time. As a first step, the strategy applies the boolean rules as long as possible. In case of branching rules, this operation nondeterministically selects (guesses) one of the conclusions of the rules. For each negated conditional, the strategy applies the rule $(\mid\!\sim^-)$ to it, generating a set Γ' that does not contain any negated conditional. On this node, the strategy applies the static rules as far as possible, then it iterates the following steps until either the current node contains an axiom or it does not contain negated boxed formulas $u : \neg\Box\neg A$: 1. apply the $(\Box^-_s)$ rule by guessing a branch; 2. apply the static rules as far as possible to the obtained node. If the final node does not contain an axiom, then we conclude that Γ' is satisfiable.

The overall complexity of the strategy can be estimated as follows. Consider that $n = \mid \Gamma \mid$. There are at most $O(n)$ negated conditionals, therefore the initial application of $(\mid\!\sim^-)$ introduces $O(n)$ labels in the obtained node Γ'. The number of different negated box formulas is at most $O(n)$, too. The strategy builds a tableau branch for Γ' by alternating applications of $(\Box^-_s)$ and of static rules (to saturate the obtained nodes). In case of branching rules, this saturation non-deterministically selects (guesses) one of the conclusions of the rules. Consider $\{u : \neg\Box\neg A_1, u : \neg\Box\neg A_2, \ldots, u : \neg\Box\neg A_m\} \subseteq \Gamma'$, and consider a branch generated by the application of the $(\Box^-_s)$ rule. In the worst case, a new label x_1 is

introduced. Suppose also that the branch under consideration is the one containing $x_1 : A_1$ and $x_1 : \Box\neg A_1$. The $(\Box_s^-)$ rule can then be applied to formulas $x : \neg\Box\neg A_k$, introducing also a further new label x_2. However, by the presence of $x_1 : \Box\neg A_1$, the rule $(\Box_s^-)$ can no longer consistently introduce $x_2 : \neg\Box\neg A_1$, since $x_2 : \Box\neg A_1 \in \Gamma^M_{x_1 \rightarrow x_2}$. Therefore, if $(\Box_s^-)$ can be applied (at most) n times in x_1 (one for each different negated boxed formula), then the rule can be applied (at most) $n-1$ times in x_2, and so on. Therefore, at most $O(n)$ new labels are introduced by $(\Box_s^-)$ in each branch.
As the number of different subformulas of Γ is at most $O(n)$, in all steps involving the application of static rules, there are at most $O(n)$ applications of these rules. Therefore, the length of the tableau branch built by the strategy is $O(n^2)$. Finally, we observe that all the nodes of the tableau contain a number of formulas which is polynomial in n, therefore to test that a node contains an axiom has at most complexity polynomial in n. The above strategy allows to prove that:

Theorem 4 (Complexity of Phase 1). *The problem of deciding the satisfiability of a set of formulas Γ is in* NP.

Notice that the above strategy is able to build branches of polynomial length thanks to the presence of the rule (*cut*). Indeed, the key point is that, when the rule $(\Box_s^-)$ building multilinear models is applied to a given label u, *all negated boxed formulas* $u : \neg\Box\neg A_k$ belong to current set of formulas. It could be the case that, after an application of $(\Box_s^-)$ by using u, the same label u is used in one of the conclusions of another application of $(\Box_s^-)$, say to some x_i. Therefore, the application of static rules could introduce $u : \neg\Box\neg A$, and a further application of $(\Box_s^-)$ could be needed. However, since (*cut*) is a static rule, and since $A \in \mathcal{L}_\Box$ because A is the antecedent of (at least) the conditional formula $A \mid\!\sim B$ generating $\neg\Box\neg A$, either $u : \neg\Box\neg A$ or $u : \Box\neg A$ have already been introduced in the branch *before* the second application of $(\Box_s^-)$, which is a dynamic rule.

4.2 Phase 2: the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH2}$

Let us now describe the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ which, for each open branch $\mathbf{B}$ built by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$, verifies if it is a minimal model of the initial set of formulas Γ.

Definition 9. *Given an open branch* $\mathbf{B}$ *of a tableau built from* $\mathcal{TAB}^{\mathbf{P}}_{PH1}$*, we define: (i)* $\mathcal{D}(\mathbf{B})$ *as the set of labels occurring on* $\mathbf{B}$*; (ii)* $\mathbf{B}^{\Box^-} = \{x : \neg\Box\neg A \mid x : \neg\Box\neg A$ *occurs in* $\mathbf{B}\}$.

A tableau of $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ is a tree whose nodes are pairs of the form $\langle \Gamma \mid K \rangle$, where Γ is a set of labelled formulas of $\mathcal{L}^{\mathbf{P}}$, whereas Γ contains formulas of the form $x : \neg\Box\neg A$, with $A \in \mathcal{L}_\Box$.

The basic idea of $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ is as follows. Given an open branch $\mathbf{B}$ built by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ and corresponding to a model $\mathcal{M}^{\mathbf{B}}$ of $\Gamma \cup \{\neg F\}$, $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ checks whether $\mathcal{M}^{\mathbf{B}}$ is a minimal model of Γ by trying to build a model of Γ which is preferred to $\mathcal{M}^{\mathbf{B}}$. Obviously, since the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ and the strategy on its application build *multilinear* models, the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ also adopts the

$$(\mathbf{AX})\ \langle \Gamma, x : P, x : \neg P \mid K \rangle \quad P \in ATM$$

$$(\mathbf{AX}_{\Box^-})\ \langle \Gamma, u : \neg\Box\neg P \mid K \rangle \quad u : \neg\Box\neg A \notin K$$

$$(\mathbf{AX}_{\emptyset})\ \langle \Gamma \mid \emptyset \rangle$$

$$(\neg)\ \frac{\langle \Gamma, x : \neg\neg F \mid K \rangle}{\langle \Gamma, x : F \mid K \rangle}$$

$$(\wedge^+)\ \frac{\langle \Gamma, x : F \wedge G \mid K \rangle}{\langle \Gamma, x : F, x : G \mid K \rangle}$$

$$(\wedge^-)\ \frac{\langle \Gamma, x : \neg(F \wedge G) \mid K \rangle}{\langle \Gamma, x : \neg F \mid K \rangle \qquad \langle \Gamma, x : \neg G \mid K \rangle}$$

$$(cut)\ \frac{\langle \Gamma \mid K \rangle}{\langle \Gamma, x : \Box\neg A \mid K \rangle \qquad \langle \Gamma, x : \neg\Box\neg A \mid K \rangle} \quad \text{if } x : \neg\Box\neg A \notin \Gamma \text{ and } x : \Box\neg A \notin \Gamma,\ x \in \mathcal{D}(\mathbf{B}), A \in \mathcal{L}_{\Box}$$

$$(\mid\!\sim^+)\ \frac{\langle \Gamma, u : A \mid\!\sim B^{L} \mid K \rangle}{\langle \Gamma, u : A \mid\!\sim B^{L,x}, x : \neg A \mid K \rangle \qquad \langle \Gamma, u : A \mid\!\sim B^{L,x}, x : \neg\Box\neg A \mid K \rangle \qquad \langle \Gamma, u : A \mid\!\sim B^{L,x}, x : B \mid K \rangle} \quad x \in \mathcal{D}(\mathbf{B}) \text{ and } x \notin L$$

$$(\mid\!\sim^-)\ \frac{\langle \Gamma, u : \neg(A \mid\!\sim B) \mid K \rangle}{\langle \Gamma, y_1 : A, y_1 : \Box\neg A, y_1 : \neg B \mid K \rangle \qquad \langle \Gamma, y_2 : A, y_2 : \Box\neg A, y_2 : \neg B \mid K \rangle \quad \cdots \quad \langle \Gamma, y_n : A, y_n : \Box\neg A, y_n : \neg B \mid K \rangle} \quad \text{for all } y_i \in \mathcal{D}(\mathbf{B})$$

$$(\Box_s^-)\ \frac{\langle \Gamma, u : \neg\Box\neg A_1, u : \neg\Box\neg A_2, \ldots, u : \neg\Box\neg A_m \mid K, u : \neg\Box\neg A_1, u : \neg\Box\neg A_2, \ldots, u : \neg\Box\neg A_m \rangle}{\langle \Gamma, \Gamma^M_{u \to y_i}, y_i : A_1, y_i : \Box\neg A_1, \Gamma^{-1^{\Box}}_{u \to y_i} \mid K \rangle \qquad \langle \Gamma, \Gamma^M_{u \to y_i}, y_i : A_2, y_i : \Box\neg A_2, \Gamma^{-2^{\Box}}_{u \to y_i} \mid K \rangle \cdots \langle \Gamma, \Gamma^M_{u \to y_i}, y_i : A_m, y_i : \Box\neg A_m, \Gamma^{-m^{\Box}}_{u \to y_i} \mid K \rangle} \quad \text{for all } y_i \in \mathcal{D}(\mathbf{B}), u \neq y_i$$

Fig. 2. The calculus $\mathcal{TAB}^{\mathbf{P}}_{PH2}$.

same strategy on the order of application of the rules and the refined rule $(\Box_s^-)$ in order to build only such multilinear models.

Checking (un)satisfiability of $\langle \Gamma \mid \mathbf{B}^{\Box^-} \rangle$ allows to verify whether the candidate model $\mathcal{M}^{\mathbf{B}}$ is minimal. More in detail, $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ tries to build an open branch containing all the objects appearing on $\mathbf{B}$, i.e. those in $\mathcal{D}(\mathbf{B})$. To this aim, the dynamic rules use labels in $\mathcal{D}(\mathbf{B})$ instead of introducing new ones in their conclusions. The additional set Γ of a tableau node, initialized with $\mathbf{B}^{\Box^-}$, is used in order to ensure that any branch **B'** built by $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ is preferred to **B**, that is **B'** only contains negated boxed formulas occurring in **B** and there exists at least one $x : \neg\Box\neg A$ that occurs in **B** and *does not occur* in **B'**. The rules of $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ are shown in Figure 2.

More in detail, the rule $(\mid\!\sim^-)$ is applied to a set of formulas containing a formula $u : \neg(A \mid\!\sim B)$; it introduces $y : A, y : \Box\neg A$ and $y : \neg B$, where $y \in \mathcal{D}(\mathbf{B})$, instead of y being a new label. The choice of the label y introduces a branching in the tableau construction. The rule $(\mid\!\sim^+)$ is applied in the same way as in $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ to *all the labels of* $\mathcal{D}(\mathbf{B})$ (and not only to those appearing in the branch). The rule $(\Box_s^-)$ is applied to a node $\langle \Gamma, u : \neg\Box\neg A_1, u : \neg\Box\neg A_2, \ldots, u : \neg\Box\neg A_m \mid K \rangle$, when $\{u : \neg\Box\neg A_1, u : \neg\Box\neg A_2, \ldots, u : \neg\Box\neg A_m\} \subseteq K$, i.e. when all the formulas $u : \neg\Box\neg A_1, u : \neg\Box\neg A_2, \ldots, u : \neg\Box\neg A_m$ also belong to the open branch **B**. In this case, the rule introduces a branch on the choice of the individual $y_i \in \mathcal{D}(\mathbf{B})$ which is preferred to u and is such that A_k and $\Box\neg A_k$ hold in y_i, recording, for all the other formulas $\neg\Box\neg A_j$, that either $y_i : A_j$ or $y_i : \neg\Box\neg A_j$ holds. Notice that, as a difference with the rule $(\Box_s^-)$ reformulated for $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ in Phase 1, here the branches generating a new label x for each negated boxed formula $u : \neg\Box\neg A_k$ are no longer introduced.

In case a tableau node has the form $\langle \Gamma, u : \neg\Box\neg A \mid K \rangle$, and $u : \neg\Box\neg A \notin K$, then $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ detects an inconsistency, called $(\mathbf{AX})_{\Box^-}$: this corresponds to the situation in which $u : \neg\Box\neg A$ does not belong to **B**, while $\Gamma, u : \neg\Box\neg A$ is satisfiable in a model $\mathcal{M}$ only if $\mathcal{M}$ contains $u : \neg\Box\neg A$, and hence only if $\mathcal{M}$ is *not* preferred to the model represented by $\mathbf{B}$.

The calculus $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ also contains the closing condition $(\mathbf{AX})_{\emptyset}$. Since each application of $(\Box^-_s)$ removes the principal formula $u : \neg\Box\neg A$ from the set Γ, when Γ is empty all the negated boxed formulas occurring in $\mathbf{B}$ also belong to the current branch. In this case, the model built by $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ satisfies the same set of negated boxed formulas (for all labels) as $\mathbf{B}$ and, thus, it is not preferred to the one represented by $\mathbf{B}$.

Theorem 5 (Soundness and completeness of $\mathcal{TAB}^{\mathbf{P}}_{PH2}$). *Given a set of labelled formulas Γ and a formula F, an open branch $\mathbf{B}$ built by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ for $\Gamma \cup \{\neg F\}$ is satisfiable by an injective mapping in a minimal model of Γ iff the tableau in $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ for $\langle \Gamma \mid \mathbf{B}^{\Box^-} \rangle$ is closed.*

$\mathcal{TAB}^{\mathbf{P}}_{PH2}$ always terminates. Indeed, only a finite number of labels can occur on the branch (only those in $\mathcal{D}(\mathbf{B})$ which is finite). Moreover, the side conditions on the application of the rules copying their principal formulas in their conclusion(s) prevent the uncontrolled application of the same rules.

The overall procedure $\mathcal{TAB}^{\mathbf{P}}_{min}$ is defined as follows:

Definition 10. *Given a set of formulas Γ and a formula F, the calculus $\mathcal{TAB}^{\mathbf{P}}_{min}$ checks whether $\Gamma \models^{\mathcal{L}_\Box}_{min} F$ by means of the following procedure:* (phase 1) *the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ is applied to $\Gamma \cup \{\neg F\}$; if, for each branch $\mathbf{B}$ built by $\mathcal{TAB}^{\mathbf{P}}_{PH1}$, either (i) $\mathbf{B}$ is closed or (ii)* (phase 2) *the tableau built by the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ for $\langle \Gamma \mid \mathbf{B}^{\Box^-} \rangle$ is open, then $\Gamma \models^{\mathcal{L}_\Box}_{min} F$, otherwise $\Gamma \not\models^{\mathcal{L}_\Box}_{min} F$.*

$\mathcal{TAB}^{\mathbf{P}}_{min}$ is a sound and complete decision procedure for verifying if a formula F can be minimally entailed from a set of formulas Γ. It can be shown that:

Theorem 6 (Complexity of Phase 2). *The problem of verifying that a branch $\mathbf{B}$ represents a minimal model for Γ in $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ is in* NP *in the size of $\mathbf{B}$.*

By Theorems 4 and 6, we can prove that:

Theorem 7 (Complexity of $\mathcal{TAB}^{\mathbf{P}}_{min}$). *The problem of deciding whether $\Gamma \models^{\mathcal{L}_\Box}_{min} F$ is in Π_2.*

Example 2. As an example, let $\Gamma = \{artist \mid\!\sim \neg rich, \neg(artist \mid\!\sim \bot)\}$. We show that $\Gamma \models^{\mathcal{L}_\Box}_{min} artist \wedge blond \mid\!\sim \neg rich$ by means of the calculus $\mathcal{TAB}^{\mathbf{P}}_{min}$. To save space, we write A for *artist*, R for *rich* and B for *blond*. We consider $\mathcal{L}_\Box = \{A, A \wedge B\}$. The tableau $\mathcal{TAB}^{\mathbf{P}}_{PH1}$ starts by two applications of $(\mid\!\sim^-)$ on the initial set $\Gamma \cup \{\neg(A \wedge B \mid\!\sim \neg R)\}$, obtaining the node $\{x : A, x : \Box\neg A, x : \neg\bot, y : A \wedge B, y : \Box\neg(A \wedge B), y : \neg\neg R\}$. We apply $(\mid\!\sim^+)$, $(\neg)$, $(\wedge^+)$, and (cut). Disregarding the nodes that are instances of $(\mathbf{AX})$, the only left branch contains $\Gamma = \{x : A, x : \Box\neg A, x : \neg\bot, x : \neg R, x : \Box\neg(A \wedge B), y : A, y : B, y : \Box\neg(A \wedge B), y : R, y : \neg\Box\neg A\}$. Now we can apply the rule $(\Box^-_s)$, obtaining two conclusions. The first one adds to Γ the formulas $z : A, z : \Box\neg A, z : \neg(A \wedge B), z : \Box\neg(A \wedge B)$ where z is a new label, the other one intoduces the same formulas labelled by x itself. We then apply the rules ($\mid\!\sim^+$) and ($\wedge^-$). Disregarding branches containing axioms, the only two open branches, say $\mathbf{B_1}$ and $\mathbf{B_2}$, contain $\Gamma_1 = \Gamma \cup \{z : A, z : \Box\neg A, z : \neg(A \wedge B), z : \Box\neg(A \wedge B), z : \neg B, z : \neg R\}$ and $\Gamma_2 = \Gamma \cup \{x : \neg(A \wedge B), x : \Box\neg(A \wedge B), x : \neg B\}$, respectively, and their nodes cannot be further expanded. We now apply the calculus $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ to both $\mathbf{B_1}$ and $\mathbf{B_2}$. Let us start with the latter one. The tableau starts with $\langle A \mid\!\sim \neg R, \neg(A \mid\!\sim \bot) \mid y : \neg\Box\neg A\rangle$. Applications of ($\mid\!\sim^-$), ($\mid\!\sim^+$) and ($cut$) lead to an open branch containing $x : A, x : \Box\neg A, x : \neg R, x : \Box\neg(A \wedge B), y : \neg A, y : \Box\neg A, y : \Box\neg(A \wedge B), z : \neg A, z : \Box\neg A, z : \Box\neg(A \wedge B)$. Similarly for $\mathbf{B_1}$. Since the tableaux in $\mathcal{TAB}^{\mathbf{P}}_{PH2}$ for $\mathbf{B_1}$ and $\mathbf{B_2}$ are open, these two branches are closed. Thus the whole procedure $\mathcal{TAB}^{\mathbf{P}}_{min}$ verifies that $\Gamma \models^{\mathcal{L}_\Box}_{min} A \wedge B \mid\!\sim \neg R$.

5 Conclusion

In this paper, we have proposed the logic $\mathbf{P}_{min}$, which is a nonmonotonic extension of the KLM system $\mathbf{P}$. Our choice of starting from $\mathbf{P}$ rather than from the stronger $\mathbf{R}$ is that $\mathbf{P}$ avoids the unwanted inferences that $\mathbf{R}$ entails, and that are outlined in [10]. The system $\mathbf{P}_{min}$ turns out to be both more expressive and stronger, when compared to both 1-entailment and rational closure. We have also provided a two-phase tableau calculus for checking entailment in $\mathbf{P}_{min}$. Our procedure can be used to determine constructively an upper bound of the complexity of $\mathbf{P}_{min}$, namely that checking entailment is in Π_2.

Acknowledgements. The work has been supported by Project "MIUR PRIN08 LoDeN: Logiche Descrittive Nonmonotone: Complessità e implementazioni".

References

1. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence **44**(1-2) (1990) 167–207
2. Benferhat, S., Dubois, D., Prade, H.: Nonmonotonic reasoning, conditional objects and possibility theory. Artificial Intelligence **92**(1-2) (1997) 259–276
3. Pearl, J.: System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In: Proceedings of TARK. (1990) 121–135
4. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial Intelligence **55**(1) (1992) 1–60
5. Arieli, O., Avron, A.: General patterns for nonmonotonic reasoning: From basic entailments to plausible relations. Logic Journal of the IGPL **8**(2) (2000) 119–148
6. Dubois, D., Fargier, H., Perny, P.: Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach. Artificial Intelligence **148**(1-2) (2003) 219–260
7. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Reasoning About Typicality in Preferential Description Logics. In: *JELIA08. LNAI 5293, pp. 192-205.* (2008)
8. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Analytic Tableaux Calculi for KLM Logics of Nonmonotonic Reasoning. ACM ToCL **10**(3) (2009)
9. Friedman, N., Halpern, J.Y.: Plausibility measures and default reasoning. Journal of the ACM **48**(4) (2001) 648–685
10. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.L.: Preferential vs Rational Description Logics: which one for Reasoning About Typicality? ECAI 2010.