




ELSEVIER

Contents lists available at ScienceDirect

Artificial Intelligence

www.elsevier.com/locate/artint

On rational entailment for Propositional Typicality Logic [☆]Richard Booth ^a, Giovanni Casini ^b, Thomas Meyer ^{c,*}, Ivan Varzinczak ^d^a Cardiff University, United Kingdom^b CSC, University of Luxembourg, Luxembourg^c University of Cape Town and CAIR, South Africa^d CRIL, Univ. Artois & CNRS, France

ARTICLE INFO

Article history:

Received 12 September 2018

Received in revised form 22 September 2019

Accepted 25 September 2019

Available online 27 September 2019

Keywords:

Knowledge representation and reasoning

Non-monotonic reasoning

Preferential semantics

Typicality

Rationality

ABSTRACT

Propositional Typicality Logic (PTL) is a recently proposed logic, obtained by enriching classical propositional logic with a typicality operator capturing the most typical (alias normal or conventional) situations in which a given sentence holds. The semantics of PTL is in terms of ranked models as studied in the well-known KLM approach to preferential reasoning and therefore KLM-style rational consequence relations can be embedded in PTL. In spite of the non-monotonic features introduced by the semantics adopted for the typicality operator, the obvious Tarskian definition of entailment for PTL remains monotonic and is therefore not appropriate in many contexts. Our first important result is an impossibility theorem showing that a set of proposed postulates that at first all seem appropriate for a notion of entailment with regard to typicality cannot be satisfied simultaneously. Closer inspection reveals that this result is best interpreted as an argument for advocating the development of more than one type of PTL entailment. In the spirit of this interpretation, we investigate three different (semantic) versions of entailment for PTL, each one based on the definition of rational closure as introduced by Lehmann and Magidor for KLM-style conditionals, and constructed using different notions of minimality.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Propositional Typicality Logic (PTL) [3,4] is a recently proposed logic allowing for the representation of and reasoning with an explicit notion of *typicality*. It is obtained by enriching classical propositional logic with a *typicality operator* \bullet , the intuition of which is to refer to those most typical (or normal or conventional) situations in which a given sentence holds. PTL is characterised using a preferential semantics similar to that originally proposed by Shoham [42] and extensively developed by Kraus et al. [33] and Lehmann and Magidor [35] in the propositional case, with close connections to the formalisms developed by Pearl and Goldszmidt [39,40], and by others [6,14,15,31,41,19–21] in more expressive languages.

In spite of the non-monotonic features introduced by the adoption of a preferential semantics for \bullet , the obvious definition of entailment for PTL, i.e., the one based on a Tarskian notion of logical consequence, remains monotonic. Of course, such a notion of entailment is inappropriate in non-monotonic contexts, in particular when reasoning about typicality, as is already clear from an enriched version of the classical Tweety example: If birds typically fly, and penguins are birds (and that is

[☆] The present submission is an extended and elaborated version of a paper presented at the 24th International Joint Conference on Artificial Intelligence (IJCAI 2015) [5].

* Corresponding author.

E-mail address: tmeyer@cs.uct.ac.za (T. Meyer).

all we know), we would expect to be able to conclude that typical penguins are typical birds, and therefore that typical penguins fly. Learning that penguins typically do not fly should lead us to conclude that penguins are not typical birds, and to retract the conclusions about typical penguins being typical birds, and about typical penguins flying.

In this paper, we investigate three semantic versions of entailment for PTL, constructed using three different forms of minimality. All these are based on the notion of rational closure as defined by Lehmann and Magidor [35] for KLM-style conditionals in a propositional setting. We show that they can be viewed as distinct extensions of rational closure, equivalent with respect to the conditional language originally proposed by Kraus et al., but different in the PTL framework.

We shall study the aforementioned forms of entailment in an abstract formal setting, obtained by proposing a set of postulates that, at first glance, seem appropriate for any notion of entailment with regard to typicality. Our first important result is a negative one, though. It is an impossibility result proving that the set of postulates cannot all be satisfied simultaneously. A more detailed analysis of the result shows that, instead of being viewed as negative, this result should rather be interpreted as an indication that PTL allows for different types of entailment, corresponding to different subsets of the full set of postulates we provide. In line with this argument, we define three types of entailment for PTL corresponding to distinct subsets of the postulates, referred to as *LM-entailment*, *PT-entailment*, and *PT'-entailment*, a modification of the latter. Our argument for more than one type of entailment for the same logic is in line with the proposal put forward by Lehmann in the context of entailment for conditional knowledge bases, where he proposes both *prototypical reasoning* and *presumptive reasoning* as acceptable forms of entailment [34]. We elaborate on this point in Section 8, but the gist of the argument is the acknowledgement of the existence of more than one form of entailment for the same representational formalism.

The remainder of the present paper is structured as follows. Section 2 provides the background and notation for the rest of the work. In Section 3 we discuss the complexities surrounding a notion of entailment for PTL. In Section 4 we put forward our postulates and show the impossibility result. In Section 5 we define LM-entailment while Section 6 is devoted to the definition of PT-entailment, and Section 7 to the definition of PT'-entailment. Section 8 addresses the implications of the impossibility result, making the case for three forms of PTL entailment. Section 9 discusses related work, while Section 10 concludes and discusses future work.

2. Logical preliminaries

Let \mathcal{P} be a finite set of propositional *atoms* with at least two elements.¹ We use p, q, \dots as meta-variables for atoms. Propositional sentences (and, in later sections, sentences of the richer language we shall introduce in Section 2.3 below) are denoted by α, β, \dots , and are recursively defined in the usual way: $\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \top \mid \perp$. All the other Boolean connectives ($\vee, \rightarrow, \leftrightarrow, \dots$) are defined in terms of \neg and \wedge in the standard way. With \mathcal{L} we denote the set of all propositional sentences.

We denote by \mathcal{U} the set of all propositional *valuations* $v : \mathcal{P} \rightarrow \{0, 1\}$, i.e., $\mathcal{U} := \{0, 1\}^{\mathcal{P}}$. Whenever it eases the presentation, we shall represent valuations as sets of literals (i.e., atoms or negated atoms), with each literal indicating the truth-value of the respective atom. Thus, for the logic generated from $\mathcal{P} = \{p, q\}$, the valuation in which p is true and q is false will be represented as $\{p, \neg q\}$. Satisfaction of a sentence $\alpha \in \mathcal{L}$ by $v \in \mathcal{U}$ is defined in the usual truth-functional way and is denoted by $v \models \alpha$.

2.1. KLM-style rational conditionals

In the conditional logic investigated by Kraus et al. [33], often referred to as the *KLM approach*, one is interested in (defeasible) conditionals of the form $\alpha \sim \beta$, read as “typically, if α , then β ” (or, depending on the example at hand, as “ α s are typically β s” and variants thereof). For instance, if $\mathcal{P} = \{b, f, p\}$, where b, f and p stand for, respectively, “being a bird”, “being able to fly”, and “being a penguin”, the following are examples of defeasible conditionals: $b \sim f$ (birds typically fly), $p \wedge b \sim \neg f$ (penguins that are birds typically do not fly).

Kraus et al. put forward the following list of properties that the conditional \sim ought to satisfy in order to be considered as appropriate in a non-monotonic setting (these properties have been discussed at length in the non-monotonic reasoning community and we shall not do so here):

$$\begin{array}{llll}
 \text{(Ref)} & \alpha \sim \alpha & \text{(LLE)} & \frac{\models \alpha \leftrightarrow \beta, \alpha \sim \gamma}{\beta \sim \gamma} \\
 & & \text{(And)} & \frac{\alpha \sim \beta, \alpha \sim \gamma}{\alpha \sim \beta \wedge \gamma} \\
 \text{(Or)} & \frac{\alpha \sim \gamma, \beta \sim \gamma}{\alpha \vee \beta \sim \gamma} & \text{(RW)} & \frac{\alpha \sim \beta, \models \beta \rightarrow \gamma}{\alpha \sim \gamma} \\
 & & \text{(CM)} & \frac{\alpha \sim \beta, \alpha \sim \gamma}{\alpha \wedge \beta \sim \gamma}
 \end{array}$$

A conditional satisfying such properties is called a *preferential conditional*. We can require \sim to satisfy other properties as well, one of which is rational monotonicity:

¹ This (reasonable) assumption is needed for technical reasons.

2	{b, f, p}
1	{b, ¬f, ¬p}, {b, ¬f, p}
0	{¬b, ¬f, ¬p}, {¬b, f, ¬p}, {b, f, ¬p}

Fig. 1. A ranked interpretation for $\mathcal{P} = \{b, f, p\}$.

$$(RM) \frac{\alpha \vdash \gamma, \alpha \not\vdash \neg\beta}{\alpha \wedge \beta \vdash \gamma}$$

A preferential conditional also satisfying (RM) is called a *rational conditional*.

The semantics of KLM-style rational conditionals is given by structures called *ranked interpretations* [35]:

Definition 2.1 (*Ranked interpretation*). A **ranked interpretation** \mathcal{R} is a function from \mathcal{U} to $\mathbb{N} \cup \{\infty\}$ satisfying the following *convexity* property: for every $i \in \mathbb{N}$, if $\mathcal{R}(v) = i$, then, for every j such that $0 \leq j < i$, there is a $v' \in \mathcal{U}$ for which $\mathcal{R}(v') = j$.

Observe that \mathcal{R} generates a modular order $<_{\mathcal{R}}$ on \mathcal{U} as follows: $u <_{\mathcal{R}} v$ if and only if $\mathcal{R}(u) < \mathcal{R}(v)$ (where $i < \infty$ for every $i \in \mathbb{N}$). If there is no ambiguity, we will omit the subscript and refer to the modular order as $<$.²

In a ranked interpretation \mathcal{R} the intuition is that valuations lower down in the ordering are deemed more normal (or typical) than those higher up, with those with an infinite rank (a rank of ∞) being regarded as so atypical as to be impossible.

The *possible* valuations in \mathcal{R} are defined as follows: $\mathcal{U}^{\mathcal{R}} := \{u \in \mathcal{U} \mid \mathcal{R}(u) < \infty\}$. Given $\alpha \in \mathcal{L}$, we let $\llbracket \alpha \rrbracket^{\mathcal{R}} := \{v \in \mathcal{U}^{\mathcal{R}} \mid v \Vdash \alpha\}$. Note that it may be possible that $\mathcal{R}(u) = \infty$ for every $u \in \mathcal{U}$, and therefore that $\mathcal{U}^{\mathcal{R}} = \emptyset$.

Given $\alpha, \beta \in \mathcal{L}$, we say \mathcal{R} satisfies (is a ranked model of) the conditional $\alpha \vdash \beta$ (denoted $\mathcal{R} \Vdash \alpha \vdash \beta$) if all the $<$ -minimal α -valuations also satisfy β , i.e., if $\min_{<} \llbracket \alpha \rrbracket^{\mathcal{R}} \subseteq \llbracket \beta \rrbracket^{\mathcal{R}}$. We say \mathcal{R} is a ranked model of a set of conditionals \mathcal{C} if $\mathcal{R} \Vdash \alpha \vdash \beta$ for every $\alpha \vdash \beta \in \mathcal{C}$, and that a set of conditionals \mathcal{C} is *satisfiable* only if it has a ranked model \mathcal{R} for which $\mathcal{U}^{\mathcal{R}} \neq \emptyset$. Observe that if \mathcal{C} is unsatisfiable, it has as its only ranked model the ranked interpretation \mathcal{R} for which $\mathcal{U}^{\mathcal{R}} = \emptyset$.

Sometimes it is convenient to represent a ranked interpretation \mathcal{R} as a partition $(L_0, \dots, L_{n-1}, L_{\infty})$ of \mathcal{U} where, for $i \in \mathbb{N} \cup \{\infty\}$, $L_i = \{u \in \mathcal{U} \mid \mathcal{R}(u) = i\}$ and where n is some $i \in \mathbb{N}$ for which $L_i = \emptyset$. That is, for each $i \in \{0, \dots, n-1, \infty\}$, L_i is the set of all valuations of rank i . We refer to such a ranked interpretation as an *n-rank* interpretation.

Observe that the partition above has a finite number of cells, but includes the possibility for some of the L_i s to be empty. This is necessary for three reasons. First, the cell L_{∞} (the set of all impossible valuations) may be empty. Second, it may be the case that $L_{\infty} = \mathcal{U}$. That is, it may be that all valuations are impossible. Third, as we shall see below, this representation will often be used to *compare* ranked interpretations. In cases where such ranked interpretations do not have the same number of non-empty cells, this representation allows us to represent them as having the same (finite) number of cells, say $(L_0, \dots, L_{n-1}, L_{\infty})$ and $(M_0, \dots, M_{n-1}, M_{\infty})$, where n is the smallest integer such that $L_i = M_i = \emptyset$.

Fig. 1 depicts an example of a ranked interpretation for $\mathcal{P} = \{b, f, p\}$ satisfying both $b \vdash f$ and $p \wedge b \vdash \neg f$. (In our graphical representations of the ranked interpretations we frequently omit the rank ∞ .)

For a better understanding of the reasons behind the aforementioned properties and the semantic constructions, the reader is referred to the work of Kraus et al. [33,35].

2.2. Rational closure

Given a set of conditionals \mathcal{C} , reasoning in the KLM framework amounts to the derivation of new conditionals from \mathcal{C} . Towards this end, Lehmann and Magidor [35] proposed what they refer to as *rational closure*. Here we focus on the semantic version of rational closure they present.

Their idea was to define a preference relation \trianglelefteq_{LM} over the set of possible ranked interpretations and then to base entailment on choosing only the most preferred, i.e., minimal w.r.t. \trianglelefteq_{LM} , ranked models of \mathcal{C} .

The relation \trianglelefteq_{LM} can be described as follows.

Definition 2.2 (*LM-preference*). Let $\mathcal{R}_1 = (L_0, \dots, L_{n-1}, L_{\infty})$ and $\mathcal{R}_2 = (M_0, \dots, M_{n-1}, M_{\infty})$ be any pair of ranked interpretations. Then,

$$\mathcal{R}_1 \trianglelefteq_{LM} \mathcal{R}_2 \quad \text{if} \quad \text{either} \quad L_i = M_i \text{ for all } i \in \{0, \dots, n-1, \infty\}, \\ \text{or} \quad L_j \supseteq M_j \text{ for the smallest } j \geq 0 \text{ s.t. } L_j \neq M_j.$$

$\mathcal{R}_1 \triangleleft_{LM} \mathcal{R}_2$ if and only if $\mathcal{R}_1 \trianglelefteq_{LM} \mathcal{R}_2$ and not $\mathcal{R}_2 \trianglelefteq_{LM} \mathcal{R}_1$.

² Recall that, given a set X , $< \subseteq X \times X$ is modular if and only if there is total order \leq on a set Ω and a ranking function $rk : X \mapsto \Omega$ s.t. for every $x, y \in X$, $x < y$ if and only if $rk(x) < rk(y)$.

\preceq_{LM} forms a partial order over ranked interpretations, and, for every satisfiable set of conditionals \mathcal{C} , there exists a unique \preceq_{LM} -minimum element $\mathcal{R}^{rc}(\mathcal{C})$ among all the ranked models of \mathcal{C} (see Proposition A.2 in Appendix A). We will refer to this element as the *LM-minimum*.

This is not exactly the semantic representation defined by Lehmann and Magidor, but this representation can easily be derived from other work on rational closure, such as that of Booth and Paris [2] (see Appendix A).

Proposition 2.1. *Given a set of conditionals \mathcal{C} and a conditional $\alpha \sim \beta$, $\alpha \sim \beta$ is in the rational closure of \mathcal{C} iff:*

1. \mathcal{C} is unsatisfiable; or
2. $\mathcal{R}^{rc}(\mathcal{C}) \Vdash \alpha \sim \beta$.

The idea is that those ranked interpretations should be preferred in which as many valuations as possible are judged to be as plausible as the background knowledge \mathcal{C} allows. Observe also that one of the consequences of this ordering is that, all other things being equal, a ranked interpretation in which a valuation is deemed to be possible will be preferred over one in which the same valuation is seen as impossible.

Then the rational closure of \mathcal{C} is the set $\sim_{\mathcal{C}}^{rc} := \{(\alpha, \beta) \mid \mathcal{R}^{rc}(\mathcal{C}) \Vdash \alpha \sim \beta\}$. Rational closure is commonly viewed as the *basic* (although certainly not the only acceptable) form of entailment over propositional conditional knowledge bases, on which other, more venturous, forms of entailment can be constructed. It is therefore an appropriate choice on which to base our investigations into versions of entailment for PTL.

2.3. Propositional Typicality Logic

PTL [3] is a logical formalism explicitly allowing for the representation of and reasoning about a notion of *typicality*. Syntactically, it extends classical propositional logic with a *typicality operator* \bullet , the intuition of which is to capture the most typical (alias normal or conventional) situations or worlds. Here we shall briefly present the main results about PTL relevant for our purposes.

The language of PTL, denoted by \mathcal{L}^\bullet , is recursively defined by:

$$\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \top \mid \perp \mid \bullet\alpha$$

As before, p denotes an atom and all the other Boolean connectives are defined in terms of \neg and \wedge .

Let $\mathcal{P} = \{b, f, o, p\}$, where b, f and p are as before and o represents “being an ostrich”. The following are examples of \mathcal{L}^\bullet -sentences: $\bullet b$ (being a typical bird), $o \rightarrow \neg\bullet b$ (ostriches are not typical birds), $(p \vee o) \leftrightarrow (b \wedge \bullet\neg f)$ (being a penguin or an ostrich is equivalent to being a bird and being a typical non-flying creature).

Intuitively, a sentence of the form $\bullet\alpha$ is understood to refer to the typical situations in which α holds. Note that α can itself be a \bullet -sentence. The semantics of PTL is also in terms of ranked interpretations (see Definition 2.1). Satisfaction is defined inductively in the classical way, adding the following condition: $v \Vdash \bullet\alpha$ if $v \Vdash \alpha$ and there is no v' such that $v' \prec v$ and $v' \Vdash \alpha$. That is, given \mathcal{R} , $\llbracket \bullet\alpha \rrbracket^{\mathcal{R}} := \min_{\prec} \llbracket \alpha \rrbracket^{\mathcal{R}}$. In the ranked interpretation \mathcal{R} of Fig. 1, we have $\llbracket \bullet b \rrbracket^{\mathcal{R}} = \{b, f, \neg p\}$, $\llbracket \bullet p \rrbracket^{\mathcal{R}} = \{b, \neg f, p\}$ and $\llbracket \bullet(b \wedge \neg f) \rrbracket^{\mathcal{R}} = \{b, \neg f, p\}$.

We say that $\alpha \in \mathcal{L}^\bullet$ is *satisfiable* in a ranked interpretation \mathcal{R} if $\llbracket \alpha \rrbracket^{\mathcal{R}} \neq \emptyset$, otherwise α is *unsatisfiable* in \mathcal{R} . We say that \mathcal{R} is a *ranked model* of α (denoted $\mathcal{R} \Vdash \alpha$) if $\llbracket \alpha \rrbracket^{\mathcal{R}} = \mathcal{U}^{\mathcal{R}}$. Observe that when $\mathcal{U}^{\mathcal{R}} = \emptyset$, then \mathcal{R} is a model of every $\alpha \in \mathcal{L}^\bullet$.

For $X \subseteq \mathcal{L}^\bullet$ we define $Mod(X) := \{\mathcal{R} \mid \mathcal{R} \Vdash \alpha \text{ for every } \alpha \in X\}$. X is *satisfiable* iff X has at least one model \mathcal{R} for which $\mathcal{U}^{\mathcal{R}} \neq \emptyset$. Observe that if X is unsatisfiable, it has as its only ranked model the ranked interpretation \mathcal{R} for which $\mathcal{U}^{\mathcal{R}} = \emptyset$. A PTL *knowledge base* is a set of sentences $\mathcal{KB} \subseteq \mathcal{L}^\bullet$.

A useful property of the typicality operator \bullet is that it allows us to express KLM-style conditionals. That is, for every ranked interpretation \mathcal{R} and every $\alpha, \beta \in \mathcal{L}$, $\mathcal{R} \Vdash \alpha \sim \beta$ if and only if $\mathcal{R} \Vdash \bullet\alpha \rightarrow \beta$. The converse does not hold since it can be shown that there are \mathcal{L}^\bullet -sentences that cannot be expressed as a set of KLM-style \sim -statements on \mathcal{L} . To give an example (taken from Booth et al. [4]), assuming $\mathcal{P} = \{p, q\}$ then $\bullet p$ is one such sentence, since $\bullet p$ has exactly four ranked models, corresponding to the cases in which $\mathcal{U}^{\mathcal{R}}$ is respectively taken to be (1) $\{\{p, q\}, \{p, \neg q\}\}$, (2) $\{\{p, q\}\}$, (3) $\{\{p, \neg q\}\}$ and (4) \emptyset (and where, in each case the ordering $\prec_{\mathcal{R}}$ is taken to be empty). Yet there exists no set X of KLM-style \sim -statements with exactly these models.

The representation result below, extending Theorem 3.12 of Lehmann and Magidor [35] to \mathcal{L}^\bullet , shows that the formalisation of the KLM rational conditional \sim inside PTL is appropriate.

Observation 1 (Booth et al. [4], Corollary 22). Let \mathcal{R} be a ranked interpretation and let $\sim_{\mathcal{R}} := \{(\alpha, \beta) \mid \alpha, \beta \in \mathcal{L}^\bullet \text{ and } \mathcal{R} \Vdash \bullet\alpha \rightarrow \beta\}$. Then $\sim_{\mathcal{R}}$ is a rational conditional. Conversely, for every rational conditional \sim , there exists a ranked interpretation \mathcal{R} such that, for every $\alpha, \beta \in \mathcal{L}^\bullet$, $\alpha \sim \beta$ if and only if $\mathcal{R} \Vdash \bullet\alpha \rightarrow \beta$.

For more details on PTL and the aforementioned properties, the reader is referred to the work by Booth et al. [4].

2	{b, f, p}
1	{b, ¬f, ¬p}, {b, ¬f, p}
0	{¬b, ¬f, ¬p}, {¬b, f, ¬p}, {b, f, ¬p}

Fig. 2. A ranked model of $\mathcal{KB}_1 = \{p \rightarrow b, \bullet b \rightarrow f\}$ satisfying neither $\bullet p \rightarrow \bullet b$ nor $\bullet p \rightarrow f$.

3. The entailment problem for PTL

The purpose of this section is to provide a more formal motivation for the remainder of the paper. From the perspective of knowledge representation and reasoning (KR&R), a central issue is that of what it means for a PTL sentence to *follow* from a PTL knowledge base \mathcal{KB} . An obvious approach to the matter is to embrace the notion of entailment advocated by Tarski [43] and largely adopted in the logic-based KR&R community.

Definition 3.1 (*Ranked entailment and consequence*). Let \mathcal{KB} be a PTL knowledge base and $\alpha \in \mathcal{L}^*$. We say \mathcal{KB} **ranked-entails** α (noted $\mathcal{KB} \approx_0 \alpha$) if $\text{Mod}(\mathcal{KB}) \subseteq \text{Mod}(\alpha)$. Its associated **ranked consequence** operator is defined by setting $C_{n_0}(\mathcal{KB}) := \{\alpha \in \mathcal{L}^* \mid \mathcal{KB} \approx_0 \alpha\}$.

As we shall see below, this version of entailment is *not* appropriate in the context of PTL for a number of reasons. For one, consider the following definition of a conditional induced from a set of PTL sentences.

Definition 3.2 (*Induced conditional relation*). Let $\mathcal{KB} \subseteq \mathcal{L}^*$. We define $\vdash_{\mathcal{KB}} := \{(\alpha, \beta) \mid \alpha, \beta \in \mathcal{L} \text{ and } \bullet \alpha \rightarrow \beta \in \mathcal{KB}\}$.

It is worth investigating whether $\vdash_{C_{n_0}(\mathcal{KB})}$ is rational for a PTL knowledge base \mathcal{KB} , i.e., whether it satisfies all the KLM properties for rationality from Section 2.1. The following proposition, which mimics a similar result by Lehmann and Magidor in the propositional case, shows that this is not the case:

Observation 2 (*Booth et al. [4], Proposition 25*). For a PTL knowledge base \mathcal{KB} , $\vdash_{C_{n_0}(\mathcal{KB})}$ is a preferential conditional, but is not necessarily a rational conditional.

Hence, ranked consequence as defined above delivers an induced defeasible conditional that is preferential but that need not be rational. This forms an argument against ranked entailment being an appropriate notion of entailment for PTL.

One of the principles to give serious consideration when investigating PTL entailment is the *presumption of typicality* [34, p. 63]. Informally, this means that one should assume that every situation is as typical as possible. Sections 4 and 6 contain a formalisation of this principle. For now, we illustrate it with an example.

Example 3.1. Let $\mathcal{KB}_1 = \{p \rightarrow b, \bullet b \rightarrow f\}$ (penguins are birds, and typical birds fly). Given just this information about birds and penguins, it is reasonable to expect $\bullet p \rightarrow \bullet b$ (typical penguins are typical birds), and therefore $\bullet p \rightarrow f$ (typical penguins fly), to follow from \mathcal{KB}_1 . With ranked entailment, these requirements are not met, as there is a ranked model of \mathcal{KB}_1 , depicted in Fig. 2, invalidating the expected conclusions. This is so because ranked entailment, being a Tarskian relation, is not ampliative, i.e., it does not allow for venturing beyond what necessarily follows from the knowledge base. \square

Besides requiring PTL entailment to be ampliative, we also want it to be *defeasible*, that is, the conclusions derived under the presumption of typicality in an ampliative way can be retracted in case of new conflicting information. This is illustrated by the following example.

Example 3.2. Assume $\bullet p \rightarrow \bullet b$ and $\bullet p \rightarrow f$ (somehow) could follow from \mathcal{KB}_1 in Example 3.1, but then we are informed that typical penguins do not fly. That is, let $\mathcal{KB}_2 = \mathcal{KB}_1 \cup \{\bullet p \rightarrow \neg f\}$. While we want $p \rightarrow \neg \bullet b$ (penguins are not typical birds) to follow from \mathcal{KB}_2 , we do *not* want $\bullet p \rightarrow f$ to follow from \mathcal{KB}_2 , which is not possible with ranked entailment. \square

4. Towards a notion of entailment for PTL

We have seen that ranked entailment has some serious drawbacks in a non-monotonic context. Therefore, the question as to what logical consequence in PTL should mean remains mostly unanswered so far. In this section, we first specify and discuss a list of postulates formalising the requirements motivated in the last section and that, at first glance, seem reasonable for an appropriate notion of entailment in PTL. In the subsequent section, we consider specific alternatives to ranked entailment and check them against our postulates.

We start by introducing some notation. With $\approx_\gamma \subseteq \mathcal{P}(\mathcal{L}^*) \times \mathcal{L}^*$, we denote any entailment relation on the language of PTL. Given an entailment relation \approx_γ , its associated *consequence* operator is defined in the usual way by setting, for each $\mathcal{KB} \subseteq \mathcal{L}^*$, $C_{n_\gamma}(\mathcal{KB}) := \{\alpha \in \mathcal{L}^* \mid \mathcal{KB} \approx_\gamma \alpha\}$.

Following the tradition in the non-monotonic reasoning literature, the obvious starting point is to consider some of the basic properties of classical consequence operators.

P1 For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$, $\mathcal{KB} \subseteq Cn_\gamma(\mathcal{KB})$ (Inclusion)

P2 For every $\mathcal{KB}, \mathcal{KB}' \subseteq \mathcal{L}^\bullet$,
if $\mathcal{KB} \subseteq \mathcal{KB}' \subseteq Cn_\gamma(\mathcal{KB})$, then $Cn_\gamma(\mathcal{KB}') = Cn_\gamma(\mathcal{KB})$ (Cumulativity)

Note that Cumulativity and Inclusion imply Idempotence. Idempotence, formalised as

For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$, $Cn_\gamma(\mathcal{KB}) = Cn_\gamma(Cn_\gamma(\mathcal{KB}))$ (Idempotence)

can be derived from Cumulativity by setting $\mathcal{KB}' = Cn_\gamma(\mathcal{KB})$, and letting Inclusion impose the satisfaction of the antecedent. Idempotence indicates that a consequence operator behaves as a ‘once-off’ operation, that is, as a closure operator. There is agreement in the literature that both Inclusion and Cumulativity are desirable properties to have [36, p. 43].

Ranked entailment, as defined in Section 3, satisfies Properties **P1** and **P2**. Nevertheless, $Cn_0(\cdot)$, the associated consequence relation of ranked entailment, also satisfies the classical property of Monotonicity: If $\mathcal{KB} \subseteq \mathcal{KB}'$, then $Cn_0(\mathcal{KB}) \subseteq Cn_0(\mathcal{KB}')$. As seen in Example 3.2, this is a property that we do not want $Cn_\gamma(\cdot)$ to satisfy (certainly not in general).

So, we require $Cn_\gamma(\cdot)$ to be a *non-monotonic* consequence operator. This amounts to requiring $Cn_\gamma(\cdot)$ to satisfy the following two postulates:

P3 For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$, $Cn_0(\mathcal{KB}) \subseteq Cn_\gamma(\mathcal{KB})$ (Ampliativeness)

P4 For some $\mathcal{KB}, \mathcal{KB}' \subseteq \mathcal{L}^\bullet$, $\mathcal{KB} \subseteq \mathcal{KB}'$ but $Cn_\gamma(\mathcal{KB}) \not\subseteq Cn_\gamma(\mathcal{KB}')$ (Defeasibility)

Ampliativeness, a property generalising supra-classicality [37] (where the basic underlying entailment relation is classical), says that $Cn_\gamma(\cdot)$ should be at least as venturesome as its underlying ranked entailment. Defeasibility specifies that $Cn_\gamma(\cdot)$ should be flexible enough to disallow previously derived conclusions in the light of new (possibly conflicting) information. In Example 3.1, assuming $\bullet p \rightarrow f \in Cn_\gamma(\mathcal{KB}_1)$ is the case, then $\bullet p \rightarrow f$ should no longer be concluded if $\bullet p \rightarrow \neg f$ is added to \mathcal{KB}_1 . Note that adding Defeasibility to Ampliativeness actually implies a *strict* version of Ampliativeness which says $Cn_\gamma(\cdot)$ should in some cases be *more* venturesome than its underlying ranked entailment. (Since, if $Cn_\gamma(\mathcal{KB}) = Cn_0(\mathcal{KB})$ for all \mathcal{KB} , then $Cn_\gamma(\cdot)$ is just ranked entailment, which is monotonic.)

P1, **P2** and **P3** together imply that the closure operation $Cn_\gamma(\cdot)$ gives as output a theory that is closed under $Cn_0(\cdot)$.

Lemma 4.1. *If $Cn_\gamma(\cdot)$ satisfies **P1**, **P2** and **P3**, then, for every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$,*

$$Cn_\gamma(\mathcal{KB}) = Cn_0(Cn_\gamma(\mathcal{KB}))$$

Proof. $Cn_0(\cdot)$ is a Tarskian consequence relation (see Definition 3.1), and, as such, it satisfies *Inclusion*. That is, for every set of formulas \mathcal{S} , $\mathcal{S} \subseteq Cn_0(\mathcal{S})$. To see it, it is sufficient to check that, according to Definition 3.1, for every $\alpha \in \mathcal{S}$, $\mathcal{S} \approx_0 \alpha$. Hence, since $Cn_0(\cdot)$ satisfies Inclusion, $Cn_\gamma(\mathcal{KB}) \subseteq Cn_0(Cn_\gamma(\mathcal{KB}))$.

By P3 we have $Cn_0(Cn_\gamma(\mathcal{KB})) \subseteq Cn_\gamma(Cn_\gamma(\mathcal{KB}))$, that, by Idempotence (consequence of **P1** and **P2**), implies $Cn_0(Cn_\gamma(\mathcal{KB})) \subseteq Cn_\gamma(\mathcal{KB})$. \square

Similarly to KLM in the propositional case, we would ideally like the defeasible conditional induced by $Cn_\gamma(\mathcal{KB})$ (see Definition 3.2) to satisfy all the rationality properties:

P5 For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$, $\vdash_{Cn_\gamma(\mathcal{KB})}$ is a rational conditional relation on \mathcal{L} (Conditional Rationality)

As observed above, **P5** requires the defeasible conditional induced by $Cn_\gamma(\mathcal{KB})$ to be rational—that is, to satisfy all the rationality properties. But from Theorem 3.12 of Lehmann and Magidor [35] it follows that every rational defeasible conditional can be obtained from a single ranked interpretation. So, from this it follows that requiring the defeasible conditional induced by $Cn_\gamma(\mathcal{KB})$ to be rational amounts to requiring that the defeasible conditional be generated by a single ranked interpretation. That is, by courtesy of this result, **P5** can also be rephrased as follows:

P5' For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$, there is a ranked interpretation \mathcal{R} s.t., for every $\alpha, \beta \in \mathcal{L}$, $\alpha \vdash_{Cn_\gamma(\mathcal{KB})} \beta$ if and only if $\mathcal{R} \Vdash \bullet \alpha \rightarrow \beta$.
(\vdash Single Model)

The next postulate we consider, which is easily shown to be a strengthening of **P5**, simply applies this same requirement, not just to defeasible statements, but to all statements expressible in PTL:

P6 For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$, there is a ranked interpretation \mathcal{R} s.t., for all $\alpha \in \mathcal{L}^\bullet$, $\alpha \in Cn_\gamma(\mathcal{KB})$ if and only if $\mathcal{R} \Vdash \alpha$ (Single Model)

An important special case of a PTL knowledge base is when the individual elements of \mathcal{KB} correspond to KLM-style conditionals.

Definition 4.1 ((Propositional) conditional knowledge base). A PTL knowledge base \mathcal{KB} will be called a (propositional) **conditional knowledge base** if each element of \mathcal{KB} is of the form $\bullet\alpha \rightarrow \beta$, for $\alpha, \beta \in \mathcal{L}$.

The next postulate says that if \mathcal{KB} is a propositional conditional knowledge base, then the result should coincide with Lehmann and Magidor's definition of rational closure:

P7 For every conditional knowledge base \mathcal{KB} , $\vdash_{Cn_7(\mathcal{KB})} = \vdash_{\mathcal{KB}}^{rc}$ (Respects Rational Closure)

P7 implies **P4**, since rational closure is a non-monotonic closure operation.

The following property was shown by Lehmann and Magidor to be satisfied by the rational closure for *conditional* knowledge bases.

P8 For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$ and $\alpha \in \mathcal{L}$, $\alpha \in Cn_7(\mathcal{KB})$ if and only if $\alpha \in Cn_0(\mathcal{KB})$ (Strict Entailment)

P8 states that $Cn_7(\cdot)$ should coincide with ranked entailment for those sentences not involving typicality. The motivation for Strict Entailment is twofold. First, it is a proposal for ranked entailment to be the lower bound for entailment w.r.t. classical sentences (those not involving typicality), a proposal that is not controversial. But secondly, it also requires entailment of classical sentences to correspond to exactly those sanctioned by ranked entailment. This can be viewed as adhering to the principle of *minimal change*. Being Tarskian, ranked entailment is monotonic, and the argument is therefore that, while non-monotonicity may be applicable for sentences involving typicality, it should not be applicable to classical statements.

We are also interested in a couple of progressively weaker versions of Strict Entailment (the reasons for that will become clear later on). The first restricts it to hold only when \mathcal{KB} is a conditional knowledge base.

P9 For every conditional knowledge base \mathcal{KB} and $\alpha \in \mathcal{L}$, $\alpha \in Cn_7(\mathcal{KB})$ if and only if $\alpha \in Cn_0(\mathcal{KB})$ (Conditional Strict Entailment)

Note that **P7** also implies **P9**. To see this, first it is easy to check that every propositional formula α is equivalent to the PTL formula $\bullet\neg\alpha \rightarrow \perp$.

Proposition 4.1. For every formula $\alpha \in \mathcal{L}$ and every ranked interpretation \mathcal{R} , $\mathcal{R} \Vdash \alpha$ iff $\mathcal{R} \Vdash \bullet\neg\alpha \rightarrow \perp$.

Proof. $\mathcal{R} \Vdash \alpha$ implies $\llbracket \alpha \rrbracket^{\mathcal{R}} = U^{\mathcal{R}}$, that is equivalent to $\mathcal{R} \Vdash \neg\alpha \rightarrow \perp$, that, in turn, implies $\mathcal{R} \Vdash \bullet\neg\alpha \rightarrow \perp$. In the opposite direction, $\mathcal{R} \Vdash \bullet\neg\alpha \rightarrow \perp$ means that for every $u \in U^{\mathcal{R}}$, $u \Vdash \bullet\neg\alpha$. $u \Vdash \bullet\neg\alpha$ for every $u \in U^{\mathcal{R}}$ implies that for every $u \in U^{\mathcal{R}}$, $u \Vdash \neg\alpha$: if we had a valuation v satisfying $\neg\alpha$ in some cell L_i , with $i < \infty$, we would either have that $v \Vdash \bullet\neg\alpha$, or there would be a valuation v' in some L_j , $j < i$, such that $v' \Vdash \bullet\neg\alpha$. Consequently, $u \Vdash \alpha$ for every $u \in U^{\mathcal{R}}$, that is, $\mathcal{R} \Vdash \alpha$. \square

P7 implies that, for every $\alpha \in \mathcal{L}$ and every conditional knowledge base \mathcal{KB} , $(\alpha, \perp) \in \vdash_{Cn_7(\mathcal{KB})}$ iff $(\alpha, \perp) \in \vdash_{\mathcal{KB}}^{rc}$. A well-known result by Lehmann and Magidor [35, Lemma 5.16] states that for every $\alpha \in \mathcal{L}$ and every conditional knowledge base \mathcal{KB} , $\alpha \vdash \perp$ is in the rational closure of \mathcal{KB} iff $\alpha \vdash \perp$ is a ranked consequence of \mathcal{KB} , that is, $(\alpha, \perp) \in \vdash_{\mathcal{KB}}^{rc}$ iff $(\alpha, \perp) \in \vdash_{Cn_0(\mathcal{KB})}$. Hence we have that for every $\alpha \in \mathcal{L}$ and every conditional knowledge base \mathcal{KB} , $(\alpha, \perp) \in \vdash_{Cn_7(\mathcal{KB})}$ iff $(\alpha, \perp) \in \vdash_{Cn_0(\mathcal{KB})}$, that, together with Proposition 4.1, implies **P9**.

In turn, **P9** implies that entailment for PTL coincides with classical propositional entailment in the case of propositional knowledge bases, as formalised by the next property.

P9' For every $\mathcal{KB} \subseteq \mathcal{L}$ and $\alpha \in \mathcal{L}$, $\alpha \in Cn_7(\mathcal{KB})$ if and only if \mathcal{KB} entails α in classical propositional logic. (Classical Entailment)

Since for every $\mathcal{KB} \cup \{\alpha\} \subseteq \mathcal{L}$, \mathcal{KB} entails α in classical propositional logic if and only if $\alpha \in Cn_0(\mathcal{KB})$, and any $\alpha \in \mathcal{L}$ is equivalent to $\bullet\neg\alpha \rightarrow \perp$, **P9'** is indeed a weakening of **P9** (provided that **P8** also holds).

Finally, we consider another property shown by Lehmann and Magidor to be satisfied by the rational closure for conditional knowledge bases.

P10 For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$ and $\alpha \in \mathcal{L}$, $\bullet\top \rightarrow \alpha \in Cn_7(\mathcal{KB})$ if and only if $\bullet\top \rightarrow \alpha \in Cn_0(\mathcal{KB})$ (Typical Entailment)

The motivation for **P10** is similar to that for **P8** above in that we want to constrain what should hold in the most typical situations. That is, given a knowledge base, the property speaks to which formulas of the form $\bullet\top \rightarrow \alpha$ should follow. Ranked entailment clearly provides a lower bound for such a kind of statement, but **P10** also proposes to consider ranked

entailment as the upper bound, thereby requiring that the set of statements $\bullet\top \rightarrow \alpha$ entailed by a knowledge base should correspond exactly to those sanctioned by ranked entailment. The argument for this is that ranked entailment is monotonic and, applying the principle of minimal change, it is only when dealing with atypical situations that ranked entailment is not always sufficient.

Although these postulates all seem reasonable on their own, it turns out that they cannot all be satisfied simultaneously. In fact, this impossibility result already holds for a strict subset of the postulates.

Theorem 4.1. *There is no PTL consequence operator $Cn_7(\cdot)$ satisfying all of P1, P2, P3, P5, P8 and P10.*

Proof. Regarding P5, requiring $\vdash_{Cn_7(\cdot)}$ to satisfy (RM) is equivalent to requiring that, for every knowledge base \mathcal{KB} and whatever formulas α, β, γ , if $\bullet\alpha \rightarrow \gamma \in Cn_7(\cdot)$ and $\bullet\alpha \rightarrow \beta \notin Cn_7(\cdot)$, then we have $\bullet(\alpha \wedge \neg\beta) \rightarrow \gamma \in Cn_7(\cdot)$.

Assume $Cn_7(\cdot)$ satisfies the given properties, and let $\mathcal{KB} = \{\bullet\top \rightarrow p, \bullet\neg p \rightarrow \bullet q\}$. By Strict Entailment (P8), $p \notin Cn_7(\mathcal{KB})$ (because of e.g. the 2-rank model $(\{p, \neg q\}, \{\neg p, q\})$ of \mathcal{KB}). By Typical Entailment (P10), $\bullet\top \rightarrow \neg q \notin Cn_7(\mathcal{KB})$ (because of e.g. the 1-rank model $(\{p, q\}, \{p, \neg q\})$ of \mathcal{KB}). By Inclusion (P1) $\bullet\top \rightarrow p \in Cn_7(\mathcal{KB})$, and then by (RM) we must conclude that $\bullet(\top \wedge q) \rightarrow p \in Cn_7(\mathcal{KB})$, that is, $(\top \wedge q, p) \in \vdash_{Cn_7(\mathcal{KB})}$; since $\vdash_{Cn_7(\cdot)}$ must satisfy LLE, the latter implies $(q, p) \in \vdash_{Cn_7(\mathcal{KB})}$, that is, $\bullet q \rightarrow p \in Cn_7(\mathcal{KB})$.

Since by Inclusion (P1) $\bullet\neg p \rightarrow \bullet q \in Cn_7(\mathcal{KB})$, we have $\{\bullet q \rightarrow p, \bullet\neg p \rightarrow \bullet q\} \subset Cn_7(\mathcal{KB})$. Since $\bullet\neg p \rightarrow p \in Cn_0(\{\bullet q \rightarrow p, \bullet\neg p \rightarrow \bullet q\})$ and $Cn_0(\cdot)$ is monotonic, we have $\bullet\neg p \rightarrow p \in Cn_0(Cn_7(\mathcal{KB}))$. Then, by Lemma 4.1, that assumes P1, P2 and P3, we have that $\bullet\neg p \rightarrow p \in Cn_7(\mathcal{KB})$.

We have that $p \in Cn_0(\{\bullet\neg p \rightarrow p\})$ holds: let $\mathcal{R} \Vdash \bullet\neg p \rightarrow p$, and let v be a world in \mathcal{R} s.t. $v \Vdash \neg p$. v cannot satisfy $\bullet\neg p$, since we would have that $v \Vdash \neg p \wedge p$; but $v \Vdash \neg p$ and $v \not\Vdash \bullet\neg p$ implies that in \mathcal{R} there is a world w , such that $w < v$ and $w \Vdash \bullet\neg p$, that, again, implies $w \Vdash \neg p \wedge p$.

From $p \in Cn_0(\{\bullet\neg p \rightarrow p\})$, $\bullet\neg p \rightarrow p \in Cn_7(\mathcal{KB})$, and the monotonicity of ranked entailment, we must conclude also $p \in Cn_0(Cn_7(\mathcal{KB}))$, that is, by Lemma 4.1, $p \in Cn_7(\mathcal{KB})$, against P8. \square

While, at first glance, this seems to be a negative result, our contention is that it should be interpreted as an indication that a logic as expressive as PTL admits more than one form of entailment. We elaborate directly on this point in Section 8, and indirectly in Sections 5 and 6, where we define and discuss two instances of entailment for PTL.

5. LM-entailment

We now come to our first construction of an entailment relation in PTL. We first observe that there is nothing to stop us from using the preference relation \leq_{LM} (see Section 2.2) to compare ranked interpretations of any PTL knowledge base \mathcal{KB} . The question then is, does there always exist a *unique* LM-minimum element of the ranked models of \mathcal{KB} , as there does in the restricted conditional case? And if so, how can we construct it? We now answer these questions.

We assume as input a PTL knowledge base \mathcal{KB} , where each sentence $\alpha \in \mathcal{KB}$ is in *normal form*:

Definition 5.1 (Normal form). $\alpha \in \mathcal{L}^\bullet$ is in **normal form** if it is of the form $\bigwedge_{i \leq t} \bullet\theta_i \rightarrow (\phi \vee \bigvee_{i \leq s} \bullet\psi_i)$, where $t, s \geq 0$ and the θ_i, ϕ and ψ_i are all purely propositional sentences.

Theorem 5.1. *The normal form is complete for \mathcal{L}^\bullet , i.e., for every sentence $\alpha \in \mathcal{L}^\bullet$ there is a (finite) set of sentences $X \subseteq \mathcal{L}^\bullet$, each one in normal form, such that $Mod(\alpha) = Mod(\bigwedge X)$.*

Proof. From the results by Booth et al. [3, Section 4], it follows that we need only consider sentences with non-nested instances of the typicality operator. So we let α be such a sentence. We let the set of *typicality atoms* be the propositional atoms occurring in \mathcal{L}^\bullet together with every sentence of the form $\bullet\beta$ where β is a propositional sentence (we refer to the latter as *pure typicality atoms*). And we define the set of *typicality literals* in the obvious way: the set of typicality atoms and their negations. The set of *pure typicality literals* consists of the pure typicality atoms and their negations.

Now we define *typicality conjunctive normal form* as a conjunctive normal form defined on typicality atoms. It follows immediately that α can be rewritten as a sentence, say α' , in typicality conjunctive normal form. Let X' be the set of conjuncts occurring in α' . We show below how to rewrite each conjunct in X' into a sentence in normal form. The resulting set X of sentences in normal form is the set referred to above.

By construction, each sentence $\gamma \in X'$ is a disjunction of typicality literals. We separate them into three disjoint sets, the set of propositional literals, the set of positive pure typicality literals (with cardinality of, say t , where $t \geq 0$) and the set of negative pure typicality literals (with cardinality of, say s , where $s \geq 0$). Let ϕ be the disjunction of propositional literals, denote the s positive pure typicality literals by ψ_1, \dots, ψ_s , and the t negative pure typicality literals by $\theta_1, \dots, \theta_t$. It follows immediately that γ can be rewritten as the sentence $\bigwedge_{i \leq t} \theta_i \rightarrow (\phi \vee \bigvee_{i \leq s} \psi_i)$. \square

For any ranked interpretation \mathcal{R} , and $S \subseteq \mathcal{U}^{\mathcal{R}}$, let \mathcal{R}_S^∞ be the ranked interpretation such that $\mathcal{R}_S^\infty(v) = \mathcal{R}(v)$ for every $v \in S$, and $\mathcal{R}_S^\infty(v) = \infty$ for every $v \in \mathcal{U} \setminus S$. That is, \mathcal{R}_S^∞ is the ranked interpretation obtained from \mathcal{R} by turning all

\mathcal{R}_0	0	{ $\neg f, \neg p, \neg r$ }, { $\neg f, \neg p, r$ }, { $\neg f, p, \neg r$ }, { $\neg f, p, r$ }, { $f, \neg p, \neg r$ }, { $f, \neg p, r$ }, { $f, p, \neg r$ }, { f, p, r }
-----------------	---	--

\mathcal{R}_1	1	{ $\neg f, \neg p, r$ }, { $\neg f, p, r$ }, { $f, \neg p, \neg r$ }, { $f, p, \neg r$ }, { f, p, r }
	0	{ $\neg f, \neg p, \neg r$ }, { $f, \neg p, r$ }

$\mathcal{R}_{\mathcal{KB}}^*$	∞	{ $\neg f, \neg p, r$ }, { $\neg f, p, r$ }, { $f, \neg p, \neg r$ }, { $f, p, \neg r$ }, { f, p, r }
	0	{ $\neg f, \neg p, \neg r$ }, { $f, \neg p, r$ }

$\mathcal{R}_{\mathcal{KB}}^*$ with the valuations of rank ∞ omitted:

0	{ $\neg f, \neg p, \neg r$ }, { $f, \neg p, r$ }
---	--

Fig. 3. The ranked interpretations generated in Example 5.1.

valuations not in S into impossible valuations. Similarly, let \mathcal{R}_S^1 be the ranked interpretation such that $\mathcal{R}_S^1(v) = \mathcal{R}(v)$ for every $v \in S$, and $\mathcal{R}_S^1(v) = \mathcal{R}(v) + 1$ for every $v \in \mathcal{U} \setminus S$. That is, \mathcal{R}_S^1 is the ranked interpretation obtained from \mathcal{R} by increasing the rank of all valuations not in S by 1.

Given a PTL knowledge base \mathcal{KB} we now define a ranked interpretation $\mathcal{R}_{\mathcal{KB}}^*$, obtained from \mathcal{KB} , as follows:

- Step 1** Set $\mathcal{R}_0(v) := 0$ for all $v \in \mathcal{U}$, $S_0 := \emptyset$, and $i := 1$;
- Step 2** $S_1 := \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_0}$ (separate the valuations which satisfy \mathcal{KB} w.r.t. the current ranked interpretation \mathcal{R}_0 from those that do not);
- Step 3** If $S_i = S_{i-1}$, then $\mathcal{R}_{\mathcal{KB}}^* := (\mathcal{R}_i)_{S_i}^\infty$, and stop (if there is no change in the new S_i then set the rank of those valuations that do not satisfy \mathcal{KB} w.r.t. \mathcal{R}_i to ∞ , let $\mathcal{R}_{\mathcal{KB}}^*$ be the interpretation that remains, and stop);
- Step 4** Otherwise $\mathcal{R}_i := (\mathcal{R}_{i-1})_{S_i}^1$ (otherwise create a new ranked interpretation \mathcal{R}_i by increasing the rank of every valuation not in S_i by 1);
- Step 5** $S_{i+1} := \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i}$ and $i := i + 1$ (separate the valuations which satisfy \mathcal{KB} w.r.t. the current ranked interpretation \mathcal{R}_i from those that do not, and increment i);
- Step 6** Go to Step 3.

Algorithm 1 below gives a compact description of the steps above. Note that if the input to the algorithm, \mathcal{KB} , is unsatisfiable, the ranked interpretation $\mathcal{R}_{\mathcal{KB}}^*$ that it returns is such that $\mathcal{U}^{\mathcal{R}_{\mathcal{KB}}^*} = \emptyset$.

Algorithm 1: LM-minimal.

Input: \mathcal{KB}

Output: $\mathcal{R}_{\mathcal{KB}}^*$

- 1 $\mathcal{P}_{\mathcal{KB}} := \{p \mid p \text{ is a propositional letter occurring in } \mathcal{KB}\}$;
 - 2 Let \mathcal{U} be the universe of valuations for the vocabulary $\mathcal{P}_{\mathcal{KB}}$;
 - 3 $\mathcal{R}_0(v) := 0$ for every $v \in \mathcal{U}$;
 - 4 $S_0 := \emptyset$;
 - 5 $S_1 := \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_0}$;
 - 6 $i := 1$;
 - 7 **while** $S_i \neq S_{i-1}$ **do**
 - 8 $\mathcal{R}_i := (\mathcal{R}_{i-1})_{S_i}^1$;
 - 9 $S_{i+1} := \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i}$;
 - 10 $i := i + 1$;
 - 11 $\mathcal{R}_{\mathcal{KB}}^* := (\mathcal{R}_{i-1})_{S_i}^\infty$;
 - 12 **return** $\mathcal{R}_{\mathcal{KB}}^*$
-

Example 5.1. Let us assume, for the sake of the example, that we are only talking about birds. Let $\mathcal{KB} := \{\bullet \top \rightarrow (\neg p \wedge \neg r), \bullet p \rightarrow \bullet \neg f, \bullet r \rightarrow \bullet f, p \rightarrow \neg r\}$ (the most typical things are neither penguins nor robins, typical penguins are typical non-flying birds, and typical robins are typical flying birds, penguins are not robins). The procedure initialises with all valuations being assigned the rank of 0. The only valuations that satisfy all three sentences w.r.t. \mathcal{R}_0 are those satisfying both $\neg p$ and $\neg r$. Thus $S_1 := \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_0} = \{\{\neg f, \neg p, \neg r\}, \{f, \neg p, \neg r\}\}$ and so we obtain \mathcal{R}_1 by changing the rank of all valuations not in S_1 to 1. Note that $\llbracket \bullet \neg f \rrbracket^{\mathcal{R}_1} = \{\{\neg f, \neg p, \neg r\}\}$ and $\llbracket \bullet f \rrbracket^{\mathcal{R}_1} = \{\{f, \neg p, \neg r\}\}$, so we can see that none of the valuations in $\mathcal{U} \setminus S_1$ is able to satisfy either $\bullet p \rightarrow \bullet \neg f$ or $\bullet r \rightarrow \bullet f$ w.r.t. \mathcal{R}_1 . As a consequence, $S_2 := \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_1} = S_1$ and so the procedure terminates here with $\mathcal{R}_{\mathcal{KB}}^*$ as the ranked interpretation in which all valuations in S_1 ($\{\neg f, \neg p, \neg r\}$ and $\{f, \neg p, \neg r\}$) have rank 0 and all other valuations have rank ∞ . See Fig. 3 for the ranked interpretations generated by this example. \square

We now proceed to show that: (i) the algorithm always terminates if \mathcal{KB} is finite; (ii) the ranked model $\mathcal{R}_{\mathcal{KB}}^*$ it returns is a ranked model of \mathcal{KB} , and (iii) for any other ranked model \mathcal{R} of \mathcal{KB} , we have $\mathcal{R}_{\mathcal{KB}}^* \sqsubseteq_{\text{LM}} \mathcal{R}$. We know the following about (i) and (ii):

Lemma 5.1. *The following hold for each $i \geq 0$:*

1. $S_i \subseteq S_{i+1}$, i.e., $[S_0 \subseteq S_1$ and, for all $i \geq 0$, $\llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i} \subseteq \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_{i+1}}$];
2. For all $v_1, v_2 \in \mathcal{U}$, if $\mathcal{R}_i(v_1) < \mathcal{R}_i(v_2)$, then $v_1 \in \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i}$;
3. \mathcal{R}_i is a ranked interpretation.

Proof. See Appendix B.1. \square

From Item 1 in Lemma 5.1 above, we know the algorithm terminates if \mathcal{KB} is finite, since it generates a sequence of ranked interpretations (by Item 3) in which the set of valuations satisfying \mathcal{KB} increases monotonically from one ranked interpretation to the next. Since each of these is finite, and since there is a finite number of valuations, the stopping criterion in Line 7 of the algorithm is guaranteed to occur eventually.

To show that the algorithm returns a ranked model of \mathcal{KB} it suffices to show the following.

Lemma 5.2. *For every $i > 0$, $(\mathcal{R}_i)_{S_i}^\infty$ is a ranked model of \mathcal{KB} .*

Proof. See Appendix B.2. \square

So, at each stage of the algorithm, the current ranked interpretation, when those valuations not satisfying \mathcal{KB} are excluded, forms a ranked model of \mathcal{KB} . Since the output $\mathcal{R}_{\mathcal{KB}}^*$ takes precisely this form we have the following result.

Proposition 5.1. $\mathcal{R}_{\mathcal{KB}}^* \models \bigwedge \mathcal{KB}$.

Proof. Follows from Lemma 5.2 and the construction of $\mathcal{R}_{\mathcal{KB}}^*$. \square

Next we want to show that for any other ranked model \mathcal{R} of \mathcal{KB} , we have $\mathcal{R}_{\mathcal{KB}}^* \trianglelefteq_{\text{LM}} \mathcal{R}$.

Lemma 5.3. *Let $\mathcal{R}_{\mathcal{KB}}^* := (L_0, \dots, L_{n-1}, L_\infty)$ and let $\mathcal{R} := (M_0, \dots, M_{n-1}, M_\infty)$ be any other ranked model of \mathcal{KB} . Let $i \in \{0, \dots, n-1\}$. If $L_j = M_j$ for all $j < i$, then $M_i \subseteq L_i$.*

Proof. See Appendix B.3. \square

From this lemma we can state:

Proposition 5.2. *Consider any \mathcal{KB} and let \mathcal{R} be a ranked model of \mathcal{KB} . Then $\mathcal{R}_{\mathcal{KB}}^* \trianglelefteq_{\text{LM}} \mathcal{R}$.*

That Algorithm 1 runs in time that is (singly) exponential in the size of the input knowledge base \mathcal{KB} whenever \mathcal{KB} is finite is not hard to see. Let $|\mathcal{KB}| = k$ and $|\mathcal{P}_{\mathcal{KB}}| = j$. The procedure starts by computing the universe \mathcal{U} of all valuations for the vocabulary $\mathcal{P}_{\mathcal{KB}}$, and therefore we have $|\mathcal{U}| = 2^j$. Next, in the first round of the loop, each sentence in \mathcal{KB} has to be checked against all of the exponentially many valuations in \mathcal{U} , which amounts to $k \times 2^j$ verifications. In the worst-case scenario, only one valuation is kept at level 0, with all the others moved up to level 1. In the next round, each sentence in \mathcal{KB} has to be checked against the $2^j - 1$ valuations at level 1, but also against the only valuation at level 0, because the truth of \bullet -sentences in a model also depends on those valuations that are lower down in the model. This amounts to $k \times 2^j$ verifications, which in the worst case will again result in a single valuation kept at level 1 with all the $2^j - 2$ ones moved up to level 2, and a number of $k \times 2^j$ checks to be performed in the next round. By repeating this argument one can see that, in the worst case, the algorithm will have built a ranked interpretation consisting of 2^j layers, each one containing a single valuation, i.e., a linear ordering on the 2^j valuations. This process will have involved 2^j runs, each run requiring $k \times 2^j$ valuation checks to create a new layer. It remains to know the cost of checking whether a sentence is satisfied by a valuation in a ranked model. In the first run of the loop, namely when there is a single layer, since the preference relation at this stage is empty, each of such verifications amounts to a propositional verification, which is a polynomial-time task. From the second run of the loop onward, i.e., when truth depends on the lower layers, we have that all valuations at the lower layers have to be inspected, which in the worst case amounts to $m \times 2^j$ checks to be performed, with m the number of sub-formulas of the one being checked. Putting the results together, we have that in the worst case there are a maximum of 2^j runs of the main loop, each with $k \times 2^j$ checks, and each of such valuation checks taking at most $m' \times 2^j$ operations, with m' the number of sub-formulas in \mathcal{KB} , i.e., $m' = 2^\ell$, for some ℓ . Hence the algorithm runs in $2^j \times (k \times 2^j) \times (2^\ell \times 2^j) = k \times 2^{3j+\ell}$, and is therefore in EXPTIME.

We are now in a position to define our first form of entailment for PTL.

Definition 5.2 (*LM-entailment*). Let $\mathcal{KB} \subseteq \mathcal{L}^\bullet$ and $\alpha \in \mathcal{L}^\bullet$. We say \mathcal{KB} **LM-entails** α , denoted $\mathcal{KB} \approx_{\text{LM}} \alpha$, if $\mathcal{R}_{\mathcal{KB}}^* \Vdash \alpha$. Its corresponding consequence operator is defined as $Cn_{\text{LM}}(\mathcal{KB}) := \{\alpha \in \mathcal{L}^\bullet \mid \mathcal{R}_{\mathcal{KB}}^* \Vdash \alpha\}$.

The next result outlines which properties from the previous section are satisfied by $Cn_{\text{LM}}(\cdot)$.

Theorem 5.2. $Cn_{\text{LM}}(\cdot)$ satisfies **P1–P7**, **P9**, and **P10**, but not **P8**.

Proof. For **P1**, Proposition 5.1 guarantees that $\mathcal{R}_{\mathcal{KB}}^*$ is a model of \mathcal{KB} . About **P2**, by Proposition 5.2, $\mathcal{R}_{\mathcal{KB}}^*$ is the LM-minimum model of \mathcal{KB} . If $\mathcal{KB} \subseteq \mathcal{KB}' \subseteq Cn_{\text{LM}}(\mathcal{KB})$, then $\text{Mod}(\mathcal{KB}') \subseteq \text{Mod}(\mathcal{KB})$ and $\mathcal{R}_{\mathcal{KB}}^* \in \text{Mod}(\mathcal{KB}')$; consequently $\mathcal{R}_{\mathcal{KB}}^*$ must also be the LM-minimum model of \mathcal{KB}' . For **P3**, note that $\mathcal{R}_{\mathcal{KB}}^*$ is a ranked model of \mathcal{KB} (Lemma 5.1, Item 3, plus Proposition 5.1), and so if $\alpha \in Cn_0(\mathcal{KB})$, then $\alpha \in \mathcal{R}_{\mathcal{KB}}^*$. **P4** is an immediate consequence of the satisfaction of **P7**.³ **P5** is an immediate consequence of the satisfaction of **P6**. The latter holds by definition of $Cn_{\text{LM}}(\mathcal{KB})$. For **P7**, see Section 2.2. **P9** is an immediate consequence of the satisfaction of **P7**.

Now consider **P10**. From right to left, it is an immediate consequence of **P3**. From left to right, assume there is a formula $\bullet \top \rightarrow \alpha$ that is in $Cn_{\text{LM}}(\mathcal{KB})$, but not in $Cn_0(\mathcal{KB})$. It means that there is a ranked model \mathcal{R} of \mathcal{KB} that has in its lower layer a propositional valuation v s.t. $v \Vdash \neg \alpha$; but, given that the model $\mathcal{R}_{\mathcal{KB}}^*$ defining $Cn_{\text{LM}}(\mathcal{KB})$ is the LM-minimum model of \mathcal{KB} , then also the lower layer of $\mathcal{R}_{\mathcal{KB}}^*$ must contain the valuation v , against the hypothesis.

Failure of **P8** can be seen in Example 5.1. There we have $\neg p \in Cn_{\text{LM}}(\mathcal{KB})$ (there is no penguin) because $\neg p$ holds in both valuations occurring in $\mathcal{R}_{\mathcal{KB}}^*$. Thus LM-entailment forces us to infer $\neg p$ from \mathcal{KB} . But $\neg p \notin Cn_0(\mathcal{KB})$, because there does exist a ranked model \mathcal{R} of \mathcal{KB} for which $\llbracket p \rrbracket^{\mathcal{R}} \neq \emptyset$, for instance the model \mathcal{R}_2 appearing in Example 6.1 below. \square

In summary then, LM-entailment satisfies all our postulates, except for Strict Entailment (**P8**). Lest this be seen as a negative result, bear in mind that LM-entailment satisfies Conditional Strict Entailment (**P9**), the weakened version of Strict Entailment, and therefore also Classical Entailment.

In the next section we turn to a form of entailment satisfying Strict Entailment, but at the price of having to forego Conditional Rationality, and therefore the Single Model postulate as well.

6. PT-entailment

In this section we consider another option for entailment based on a version of minimality, and derived from the characterisation of rational closure by Giordano et al. [30,32]. The general idea is to respect the principle of *presumption of typicality* (see Section 3). We shall refer to this form of entailment as *Presumption of Typicality* entailment, shortened to *PT-entailment*. Such a principle indicates the way in which the property (RM) should be satisfied. If we have $\alpha \sim \gamma$ in our knowledge base \mathcal{KB} , then, in order to satisfy (RM), we have to add either $\alpha \sim \neg \beta$ or $\alpha \wedge \beta \sim \gamma$. The presumption of typicality requires that, whenever possible, we prefer the latter (that corresponds to a constrained application of monotonicity) over the former. Semantically, given the ranked models of a knowledge base \mathcal{KB} , this corresponds to considering only those models in which every valuation is taken as typical as possible, that is, it is ‘pushed downward’ in the model as much as possible, modulo the satisfaction of \mathcal{KB} .

In order to identify the interpretations that are necessary for the definition of a notion of entailment, we introduce a preference relation $\trianglelefteq_{\text{PT}}$ on the set of ranked interpretations that follows directly from the presumption of typicality.

Definition 6.1 (*Relation $\trianglelefteq_{\text{PT}}$*). For two ranked interpretations \mathcal{R}_1 and \mathcal{R}_2 , $\mathcal{R}_1 \trianglelefteq_{\text{PT}} \mathcal{R}_2$ if and only if for every $w \in \mathcal{U}$, $\mathcal{R}_1(w) \subseteq \mathcal{R}_2(w)$. $\mathcal{R}_1 \prec_{\text{PT}} \mathcal{R}_2$ if and only if $\mathcal{R}_1 \trianglelefteq_{\text{PT}} \mathcal{R}_2$ and not $\mathcal{R}_2 \trianglelefteq_{\text{PT}} \mathcal{R}_1$.

It is easy to check that $\trianglelefteq_{\text{PT}}$ is a pre-order. Consistent with the principle of presumption of typicality, as a guideline in the choice of the relevant interpretations, the relation $\trianglelefteq_{\text{PT}}$ can be used to identify the relevant interpretations for the definition of a notion of entailment: we choose the models of \mathcal{KB} in which the valuations are presumed to be as typical as possible, that is, the relevant models are those that are in $\min_{\trianglelefteq_{\text{PT}}} \text{Mod}(\mathcal{KB})$. Then, \mathcal{KB} entails α if and only if α holds in all the (preferred) models in $\min_{\trianglelefteq_{\text{PT}}} \text{Mod}(\mathcal{KB})$. We will sometimes refer to the models in $\min_{\trianglelefteq_{\text{PT}}} \text{Mod}(\mathcal{KB})$ as the *PT-minimal* models of \mathcal{KB} . Note that if \mathcal{KB} is unsatisfiable, it has exactly one PT-minimal model, namely the ranked interpretation \mathcal{R} for which $\mathcal{U}^{\mathcal{R}} = \emptyset$.

If we consider knowledge bases composed only of classical non-monotonic conditionals $\alpha \sim \beta$, Giordano et al. have proved that for every satisfiable knowledge base there is a unique PT-minimal model [32, Theorem 1], and that such a PT-minimal model characterises the rational closure of the knowledge base [32, Theorem 2]. Given such results, it is quite immediate to prove that, given a satisfiable conditional knowledge base, its PT-minimal model corresponds to the LM-minimal model.

³ For this conclusion we need the requirement (specified in Section 2) that \mathcal{P} contains at least two elements.

Proposition 6.1. Let \mathcal{KB} be a satisfiable conditional knowledge base. A ranked interpretation \mathcal{R} is \mathcal{KB} 's PT-minimal model iff it is \mathcal{KB} 's LM-minimal model.

Proof. If \mathcal{KB} is a satisfiable conditional knowledge base, then it has a unique LM-minimal model \mathcal{R} (see Proposition A.2) and a unique PT-minimal model \mathcal{R}' [32, Theorem 1]. \mathcal{R} and \mathcal{R}' are equivalent, in the sense that they satisfy exactly the same conditionals, since they both characterise the rational closure of \mathcal{KB} (see Proposition 2.1 here for LM-minimality and the theorem by Giordano and others for PT-minimality [32, Theorem 2]).

In order to show that they are exactly the same model, we just need to prove that whenever two ranked interpretations \mathcal{R} and \mathcal{R}' satisfy exactly the same set of conditionals, then they are the same interpretation. Let $\mathcal{R} = (L_0, \dots, L_{n-1}, L_\infty)$ and $\mathcal{R}' = (M_0, \dots, M_{n-1}, M_\infty)$.

First of all, we prove $U^{\mathcal{R}} = U^{\mathcal{R}'}$: let $v \in U^{\mathcal{R}}$ and $v \notin U^{\mathcal{R}'}$, and let \bar{v} be the characteristic formula of the valuation v ; we would have $\mathcal{R}' \models \bar{v} \perp$ and $\mathcal{R} \not\models \bar{v} \perp$, against the hypothesis that \mathcal{R} and \mathcal{R}' satisfy the same set of conditionals. $U^{\mathcal{R}} = U^{\mathcal{R}'}$ immediately implies that $L_\infty = M_\infty$.

We conclude the proof by induction on the rank of the cells below ∞ . Given a cell $L_i = \{v_1, \dots, v_n\}$, let $\bar{L}_i := (\bar{v}_1 \vee \dots \vee \bar{v}_n)$.

Assume $L_0 \neq M_0$, that is, w.l.o.g., there is a v s.t. $v \in L_0$ and $v \notin M_0$. That implies $\mathcal{R}' \models \bar{v} \perp$, while $\mathcal{R} \not\models \bar{v} \perp$, against the hypothesis.

Given a number $j \leq (i-1)$, let $L_k = M_k$ for every k s.t. $0 \leq k < j$, but $L_j \neq M_j$, that is, w.l.o.g., there is a v s.t. $v \in L_j$ and $v \notin M_j$. That implies $\mathcal{R}' \models \neg(\bigvee_{0 \leq k < j} \bar{L}_k) \perp$, while $\mathcal{R} \not\models \neg(\bigvee_{0 \leq k < j} \bar{L}_k) \perp$, against the hypothesis. Since all their cells must contain the same valuations, \mathcal{R} and \mathcal{R}' are the same model. \square

Despite Proposition 6.1, given the extra expressive power of PTL, we obtain the surprising result that the two semantic constructions are not equivalent anymore. Moreover, in the present context, this notion of minimality can give back a number of minimal models, as the following example shows.

Example 6.1. Consider the knowledge base \mathcal{KB} from Example 5.1. Then, one can see that $\min_{\leq_{PT}} \text{Mod}(\mathcal{KB}) = \{\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3\}$, where:

\mathcal{R}_1 :	0	{ $\neg f, \neg p, \neg r$ }, { $f, \neg p, \neg r$ },
-------------------	---	--

\mathcal{R}_2 :	2	{ $f, p, \neg r$ }
	1	{ $\neg f, \neg p, \neg r$ }, { $\neg f, p, \neg r$ }
	0	{ $f, \neg p, \neg r$ }

\mathcal{R}_3 :	2	{ $\neg f, \neg p, r$ }
	1	{ $f, \neg p, r$ }, { $f, \neg p, \neg r$ }
	0	{ $\neg f, \neg p, \neg r$ }

In Example 6.1, note that \mathcal{R}_1 is the LM-minimum of \mathcal{KB} . In fact, it is easy to check from the characterisation of rational closure in Section 2.2 and Definition 6.1 that the LM-minimum of \mathcal{KB} is always in $\min_{\leq_{PT}} \text{Mod}(\mathcal{KB})$.

Proposition 6.2. For every knowledge base \mathcal{KB} , the LM-minimum of \mathcal{KB} is in $\min_{\leq_{PT}} \text{Mod}(\mathcal{KB})$.

Proof. Consider the definition of the preference relation for LM-minimality.

$$\mathcal{R}_1 \leq_{LM} \mathcal{R}_2 \text{ if and only if } \begin{array}{l} \text{either } L_i = M_i \text{ for all } i \in \{0, \dots, n-1, \infty\}, \\ \text{or } L_j \supseteq M_j \text{ for the smallest } j \geq 0 \text{ s.t. } L_j \neq M_j, \end{array}$$

where $\mathcal{R}_1 = (L_0, \dots, L_{n-1}, L_\infty)$ and $\mathcal{R}_2 = (M_0, \dots, M_{n-1}, M_\infty)$. The result follows from the fact that if $\mathcal{R}_1 \prec_{PT} \mathcal{R}_2$ then $\mathcal{R}_1 \prec_{LM} \mathcal{R}_2$. To see that this holds, assume $\mathcal{R}_1 \prec_{PT} \mathcal{R}_2$. Then $\mathcal{R}_1(w) \leq \mathcal{R}_2(w)$ for all $w \in \mathcal{U}$, with $\mathcal{R}_1(w') < \mathcal{R}_2(w')$ for at least one $w' \in \mathcal{U}$. From the latter, we know we cannot have $L_i = M_i$ for all i , so let $j \geq 0$ be minimal such that $L_j \neq M_j$. To show the conclusion $\mathcal{R}_1 \prec_{LM} \mathcal{R}_2$ we must show $L_j \supseteq M_j$, so let $u \in M_j$. Then $\mathcal{R}_2(u) = j$. Since $\mathcal{R}_1 \prec_{PT} \mathcal{R}_2$ we know $\mathcal{R}_1(u) \leq j$. But if $\mathcal{R}_1(u) = k < j$ then $u \in L_k = M_k$ (by minimality of j), contradicting $u \in M_j$. Hence $\mathcal{R}_1(u) = j$, i.e., $u \in L_j$ as required.

Knowing that $\mathcal{R}_1 \prec_{PT} \mathcal{R}_2$ implies $\mathcal{R}_1 \prec_{LM} \mathcal{R}_2$, it is easy to conclude our proof. Let \mathcal{R} be the LM-minimum of \mathcal{KB} , but not an element of $\min_{\leq_{PT}} \text{Mod}(\mathcal{KB})$. That is, there is an $\mathcal{R}^* \in \text{Mod}(\mathcal{KB})$ s.t. $\mathcal{R}^* \prec_{PT} \mathcal{R}$, that implies $\mathcal{R}^* \prec_{LM} \mathcal{R}$, thus contradicting the LM-minimality of \mathcal{R} . \square

We are now ready for the definition of our second type of entailment.

Definition 6.2 (PT-entailment). Let $\mathcal{KB} \subseteq \mathcal{L}^\bullet$ and $\alpha \in \mathcal{L}^\bullet$. We say \mathcal{KB} **PT-entails** α , denoted $\mathcal{KB} \approx_{\text{PT}} \alpha$, if and only if $\min_{\leq_{\text{PT}}}(\text{Mod}(\mathcal{KB})) \subseteq \text{Mod}(\alpha)$.

Its corresponding consequence operator $Cn_{\text{PT}}(\cdot)$ is inferentially weaker than $Cn_{\text{LM}}(\cdot)$, since it is defined on a possibly larger set of models.

Proposition 6.3. $Cn_{\text{PT}}(\cdot)$ satisfies **P1–P4** and **P7–P10**.

Proof. P1. $Cn_{\text{PT}}(\mathcal{KB})$ is defined using only models of \mathcal{KB} .

P2. Since $\mathcal{KB} \subseteq \mathcal{KB}' \subseteq Cn_{\text{PT}}(\mathcal{KB})$, we have $\min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}) \subseteq \text{Mod}(Cn_{\text{PT}}(\mathcal{KB})) \subseteq \text{Mod}(\mathcal{KB}') \subseteq \text{Mod}(\mathcal{KB})$. It is sufficient to prove that $\min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}') = \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB})$.

Let \mathcal{R} be a model of \mathcal{KB} and \mathcal{KB}' s.t. $\mathcal{R} \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}')$ and $\mathcal{R} \notin \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB})$. That is, there must be a model \mathcal{R}' of \mathcal{KB} s.t. $\mathcal{R}' \prec_{\text{PT}} \mathcal{R}$ and $\mathcal{R}' \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB})$. However, since $\min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}') \subseteq \text{Mod}(\mathcal{KB}')$, \mathcal{R}' is also a model of \mathcal{KB}' that is PT-preferred to \mathcal{R} , that is, it cannot be the case that $\mathcal{R} \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}')$. Inversely, let \mathcal{R} be a model of \mathcal{KB} and \mathcal{KB}' s.t. $\mathcal{R} \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB})$ and $\mathcal{R} \notin \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}')$. That is, there must be a model \mathcal{R}' of \mathcal{KB}' s.t. $\mathcal{R}' \prec_{\text{PT}} \mathcal{R}$ and $\mathcal{R}' \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}')$. However, since $\mathcal{KB} \subseteq \mathcal{KB}'$, \mathcal{R}' is also a model of \mathcal{KB} that is PT-preferred to \mathcal{R} , that is, it cannot be the case that $\mathcal{R} \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB})$.

Hence, for every $\mathcal{KB}, \mathcal{KB}'$ s.t. $\mathcal{KB} \subseteq \mathcal{KB}' \subseteq Cn_{\text{PT}}(\mathcal{KB})$, it must be $\min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}) = \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}')$, that implies **P2**.

P3. Every model in $\min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB})$ is by definition a ranked model of \mathcal{KB} . Hence if $\alpha \in Cn_0(\mathcal{KB})$, i.e., α is true in all ranked models of \mathcal{KB} , then it is true in all ranked models in $\min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB})$, i.e., $\alpha \in Cn_{\text{PT}}(\mathcal{KB})$.

P4. It is an immediate consequence of the satisfaction of **P7**.⁴

P7. See the analogous result by Giordano et al. [32, Section 2.3.2]; in particular Theorem 2, that implies that in case of a conditional KB the use of PT-minimality leads to a single minimal model, characterising Rational Closure.

P8. Let α be a propositional formula s.t. $\alpha \notin Cn_0(\mathcal{KB})$: then there is a ranked model \mathcal{R} of \mathcal{KB} s.t. $\mathcal{R}(v) < \infty$ for some v s.t. $v \Vdash \neg\alpha$. Either \mathcal{R} is a PT-minimal model of \mathcal{KB} itself, or there is a PT-minimal model \mathcal{R}' of \mathcal{KB} s.t. $\mathcal{R}' \triangleleft_{\text{PT}} \mathcal{R}$; that is, it must be the case that $\mathcal{R}'(v) < \infty$ for some model $\mathcal{R}' \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB})$, that in turn implies that $\alpha \notin Cn_{\text{PT}}(\mathcal{KB})$.

P9. It is an immediate consequence of the satisfaction of **P7**, as explained in Section 4, immediately after introducing **P9**.

P10. It is a direct consequence of Proposition 6.2 and the satisfaction of **P10** for LM-entailment. \square

Unfortunately, *Conditional Rationality* (**P5**) is not valid and therefore, neither is the Single Model postulate (**P6**).

Theorem 6.1. There is some \mathcal{KB} such that the conditional induced by $Cn_{\text{PT}}(\mathcal{KB})$ is not a rational conditional.

To see this, consider Example 6.1: we have $\bullet \rightarrow \neg p \rightarrow \neg r \in Cn_{\text{PT}}(\mathcal{KB})$ (typical non-penguins are not robins). This is because we have $\min_{\prec_{\mathcal{R}_i}} \llbracket \neg p \rrbracket^{\mathcal{R}_i} \subseteq \llbracket \neg r \rrbracket^{\mathcal{R}_i}$ for each $i = 1, 2, 3$. However both $\bullet \rightarrow \neg p \rightarrow \neg f \notin Cn_{\text{PT}}(\mathcal{KB})$ and $\bullet \rightarrow (\neg p \wedge f) \rightarrow \neg r \notin Cn_{\text{PT}}(\mathcal{KB})$. The former holds because, e.g., $\min_{\prec_{\mathcal{R}_1}} \llbracket \neg p \rrbracket^{\mathcal{R}_1} \not\subseteq \llbracket \neg f \rrbracket^{\mathcal{R}_1}$, the latter because $\min_{\prec_{\mathcal{R}_3}} \llbracket \neg p \wedge f \rrbracket^{\mathcal{R}_3} \not\subseteq \llbracket \neg r \rrbracket^{\mathcal{R}_3}$. This means the rational monotonicity property (RM) is not satisfied.

On the other hand, observe that $\neg p \notin Cn_{\text{PT}}(\mathcal{KB})$. Recall from the proof of Theorem 5.2 that we used the fact that $\neg p \in Cn_{\text{LM}}(\mathcal{KB})$ to show that LM-entailment does not satisfy Strict Entailment (**P8**).

7. PT-entailment

As we have shown above, relying on LM-minimality results in the loss of property **P8** (Strict Entailment), while using PT-minimality results in the loss of the uniqueness of the minimal model (**P6**) and the rationality of our conditional reasoning (**P5**). To summarise, on the one hand LM-minimality, failing to satisfy **P8**, can potentially enforce classical propositional information that is not a necessary consequence of the knowledge base. On the other hand, PT-minimality can be inferentially too weak. In this section we consider a third possibility, aimed at strengthening the inferential power while still preserving the satisfaction of **P8**. This third proposal is based on using the same approach as in PT-minimality, but among the PT-minimal models we consider only the “biggest” ones, that is, the ones with the *maximal* sets of possible valuations (w.r.t. \subseteq). This should allow us to augment the inferential power (we define the entailment relation using fewer models), while still preserving **P8** (we consider all the biggest models, that is, the models that assume as little propositional knowledge as possible). We now analyse this option.

We let $\min_{\supseteq_{\text{PT}}} \text{Mod}(\mathcal{KB}) := \{\mathcal{R} \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}) \mid \text{there is no } \mathcal{R}' \in \min_{\leq_{\text{PT}}} \text{Mod}(\mathcal{KB}) \text{ s.t. } \mathcal{U}^{\mathcal{R}'} \supset \mathcal{U}^{\mathcal{R}}\}$.

The corresponding entailment relation $\approx_{\text{PT}'}$ can be defined as follows.

Definition 7.1 (PT'-entailment). Let $\mathcal{KB} \subseteq \mathcal{L}^\bullet$ and $\alpha \in \mathcal{L}^\bullet$. We say \mathcal{KB} **PT'-entails** α , denoted $\mathcal{KB} \approx_{\text{PT}'} \alpha$, if and only if $\min_{\supseteq_{\text{PT}}} \text{Mod}(\mathcal{KB}) \subseteq \text{Mod}(\alpha)$.

⁴ As in Theorem 5.2, for this conclusion we need the requirement (specified in Section 2) that \mathcal{P} contains at least two elements.

For example, in Example 6.1 we would consider only \mathcal{R}_2 and \mathcal{R}_3 .

Note that if \mathcal{KB} is unsatisfiable then $\min_{\leq_{PT}}^{\supset} Mod(\mathcal{KB})$ is a singleton set containing the ranked interpretation \mathcal{R} for which $\mathcal{U}^{\mathcal{R}} = \emptyset$. Also, recall from Section 6 that for every satisfiable conditional knowledge base \mathcal{KB} there is a single PT-minimal model [32, Theorem 1], that characterises the rational closure of \mathcal{KB} [32, Theorem 2]. Such a single PT-minimal model is by definition also the only PT'-minimal model of \mathcal{KB} , and consequently, in case of conditional knowledge bases, PT'-entailment also characterises the rational closure.

Our first result regarding PT'-entailment is that it is inferentially stronger than PT-entailment.

Proposition 7.1. (i) For every $\mathcal{KB} \subseteq \mathcal{L}^\bullet$ and every $\alpha \in \mathcal{L}^\bullet$, if $\mathcal{KB} \models_{PT} \alpha$ then $\mathcal{KB} \models_{PT'} \alpha$. (ii) There exists some $\mathcal{KB}' \subseteq \mathcal{L}^\bullet$ and $\beta \in \mathcal{L}^\bullet$ such that $\mathcal{KB}' \models_{PT'} \beta$ and $\mathcal{KB}' \not\models_{PT} \beta$.

Proof. (i). Note that, since $\min_{\leq_{PT}}^{\supset} Mod(\mathcal{KB}) \subseteq \min_{\leq_{PT}} Mod(\mathcal{KB})$ for every \mathcal{KB} , $\models_{PT} \subseteq \models_{PT'}$. (ii). Observe from Example 7.1, here below, that $\mathcal{KB}' \models_{PT'} \bullet \top \rightarrow \neg f$ but $\mathcal{KB}' \not\models_{PT} \bullet \top \rightarrow \neg f$. \square

Example 7.1. Consider the knowledge base $\mathcal{KB}' := \{\bullet \top \rightarrow (\neg p \wedge \neg r), \bullet p \rightarrow \neg f, \bullet r \rightarrow \bullet f, p \rightarrow \neg r\}$, which is a modified version of the knowledge \mathcal{KB} from Example 5.1. The only difference is that now we state that typical penguins are non-flying birds, not that they are typical non-flying birds.

Then, one can check that $\min_{\leq_{PT}}^{\supset} Mod(\mathcal{KB}') = \{\mathcal{R}_1, \mathcal{R}_2\}$, where:

\mathcal{R}_1 :	2	{f, p, $\neg r$ }
	1	{ $\neg f$, p, $\neg r$ }
	0	{ $\neg f$, $\neg p$, $\neg r$ }, {f, $\neg p$, $\neg r$ },

\mathcal{R}_2 :	2	{ $\neg f$, $\neg p$, r}, {f, p, $\neg r$ }
	1	{f, $\neg p$, r}, {f, $\neg p$, $\neg r$ }, { $\neg f$, p, $\neg r$ }
	0	{ $\neg f$, $\neg p$, $\neg r$ }

while $\min_{\leq_{PT}}^{\supset} Mod(\mathcal{KB}') = \{\mathcal{R}_2\}$, since $\mathcal{U}^{\mathcal{R}_2} \supset \mathcal{U}^{\mathcal{R}_1}$.

Unfortunately, while PT'-entailment is an improvement over PT-entailment in terms of inferential strength, it is weaker than PT-entailment when it comes to the satisfaction of the list of desirable properties. That is, it satisfies, and fails to satisfy, the same properties as PT-entailment, except for Typical Entailment (P10), which PT-entailment satisfies, but PT'-entailment does not.

Proposition 7.2. $Cn_{PT'}(\cdot)$ satisfies P1–P4 and P7–P9, but does not satisfy P5, P6, and P10.

Proof. Regarding P1, P2, P3, P4, and P9 the proof for $Cn_{PT'}(\cdot)$ is the same as for $Cn_{PT}(\cdot)$ (Proposition 6.3 above).

Regarding the failure of P5, consider Example 6.1. In this example, while $\min_{\leq_{PT}} Mod(\mathcal{KB}) = \{\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3\}$, we have that $\min_{\leq_{PT}}^{\supset} Mod(\mathcal{KB}) = \{\mathcal{R}_2, \mathcal{R}_3\}$. We can use the same case used in the proof of Theorem 6.1: we have $\mathcal{KB} \models_{PT'} \bullet (\neg p) \rightarrow \neg r$, but neither $\mathcal{KB} \models_{PT'} \bullet (\neg p) \rightarrow \neg f$, nor $\mathcal{KB} \models_{PT'} \bullet (\neg p \wedge f) \rightarrow \neg r$ hold.

The failure of P5 immediately implies the failure of P6.

P7. As pointed out in Proposition 6.3, in case we are dealing with a conditional KB, deciding PT-minimality over a satisfiable conditional KB gives back a single minimal model, characterising Rational Closure. It follows immediately that such a model is also the only PT'-minimal one.

P8. Again, it follows from the satisfaction of P8 for PT-entailment (see Proposition 6.3). Let \mathcal{KB} be a knowledge base and α be a propositional formula. If there is a PT-minimal model \mathcal{R} s.t. $\mathcal{R}(v) \leq \infty$ for some $v \not\models \alpha$, then, by definition, there must be also in $\min_{\leq_{PT}}^{\supset} Mod(\mathcal{KB})$ a model \mathcal{R}' of \mathcal{KB} s.t. $\mathcal{R}'(v) \leq \infty$.

For the failure of P10, we consider Example 7.1 and the case used in the proof of Proposition 7.1: $\mathcal{KB}' \models_{PT'} \bullet \top \rightarrow \neg f$ but, since $\mathcal{KB}' \not\models_{PT} \bullet \top \rightarrow \neg f$ and \models_{PT} satisfies Ampliativeness (P3), $\bullet \top \rightarrow \neg f$ is not in $Cn_0(\mathcal{KB}')$. \square

8. Making sense of the impossibility result

Theorem 4.1 in Section 4 shows that there is no PTL consequence operator satisfying all of our postulates—more specifically, none satisfying P1, P2, P3, P5, P8, and P10. This raises the important question of which of these postulates ought to be foregone in the search for an appropriate form of PTL entailment. In trying to find an answer to this question, it is useful to consider the three forms of entailment we proposed in the previous sections. The answer seems to be that it makes sense to consider (at least) two forms of entailment for PTL, represented here by LM-entailment and PT-entailment. PT'-entailment is not viewed as a viable option, given that it satisfies fewer properties than PT-entailment. In essence then,

it comes down to a choice between having a form of entailment that satisfies Strict Entailment (PT-entailment), and one that satisfies the Single Model postulate and Conditional Rationality, i.e., is based on a rational conditional (LM-entailment).

The advantage of LM-entailment is that it satisfies all postulates except for Strict Entailment, which includes not only Single Model and Conditional Rationality, but also Conditional Strict Entailment and Classical Entailment, the weakened versions of Strict Entailment. On the other hand, the argument for PT-entailment is that the Single Model property is too restrictive in the context of full PTL, and ought to be dropped. That is, in a logic as expressive as PTL in which there are not any restrictions on the typicality operator, any form of entailment based on minimality, and adhering to the presumption of typicality, as outlined in Section 6, is likely to violate the Single Model property.

The point of view that different forms of entailment can be appropriate in enriched versions of propositional logic, particularly enriched versions dealing with aspects of typicality, is not surprising, nor new. Lehmann [34] makes the case for two forms of entailment for the conditional logic discussed in Section 2.1 on which PTL is based. He draws a distinction between *prototypical reasoning*, corresponding to rational closure as discussed in Section 2.2, and *presumptive reasoning*.

The intuition underlying prototypical reasoning is that conclusions to be drawn are constrained by the typicality of the objects under consideration. To make matters more concrete, suppose we know that birds typically fly, that birds typically have wings, that robins are birds, that penguins are birds, and that penguins typically don't fly. Robins can be regarded as typical birds and therefore inherit the properties of typical birds, such as having wings. Penguins, on the other hand, should be regarded as atypical birds since they typically cannot fly, and therefore do not inherit the properties of a typical bird, such as having wings. This is to be contrasted with presumptive reasoning, a more permissive form of reasoning for which the intuition is to draw conclusions unless we have specific information to the contrary. In our example above presumptive reasoning would allow us to conclude that penguins typically have wings (since we do not have explicit information contradicting that conclusion), thereby distinguishing it from prototypical reasoning.

Our argument here is *not* that the relationship between PT-entailment and LM-entailment is analogous to the relationship between prototypical reasoning and presumptive reasoning, although it is true that LM-entailment can be viewed as a refinement of PT-entailment (yielding more conclusions), just as presumptive reasoning is a refinement of prototypical reasoning. Rather, the important point is that differences in context will determine which form of entailment is appropriate. It is our contention that the same principle applies to the differences between LM-entailment and PT-entailment. Below we discuss the technical differences between the two forms of entailment and then provide an example to illustrate the principle.

As we have seen above, the difference between these two forms of entailment comes down to a choice between Strict Entailment on the one hand, Conditional Rationality (and Single Model) on the other hand. Employing LM-entailment ensures that we remain rational (i.e., satisfying all the KLM properties), but at the cost of going beyond Tarskian monotonicity for typicality-free sentences. Conversely, making use of PT-entailment allows us to remain Tarskian for typicality-free sentences, but forces us to forego rationality, and in particular, the rational monotonicity property RM. Intuitively then, LM-entailment is the more permissive form of entailment here. Not only do we remain rational, unlike PT-entailment, but we do so at the cost of allowing the entailment of more typicality-free sentences than permitted by PT-entailment. The example below is indicative of the factors to take into account when deciding, in a specific context, which of LM-entailment or PT-entailment is the more appropriate form of reasoning.

Example 8.1. Consider again the knowledge base $\mathcal{KB} := \{\bullet T \rightarrow (\neg p \wedge \neg r), \bullet p \rightarrow \bullet \neg f, \bullet r \rightarrow \bullet f, p \rightarrow \neg r\}$ introduced in Example 5.1. From Examples 5.1 and 6.1 it is not hard to verify that both LM-entailment and PT-entailment sanction the conclusion that typical non-robins are not penguins ($\mathcal{KB} \models_{LM} \bullet (\neg r) \rightarrow \neg p$ and $\mathcal{KB} \models_{PT} \bullet (\neg r) \rightarrow \neg p$), and do *not* allow for the entailment that typical non-robins can fly ($\mathcal{KB} \not\models_{LM} \bullet (\neg r) \rightarrow f$ and $\mathcal{KB} \not\models_{PT} \bullet (\neg r) \rightarrow f$). This leaves us with a choice. On the one hand it is reasonable to conclude from this that typical *non-flying* non-robins are not penguins (that is, $\bullet (\neg r \wedge \neg f) \rightarrow \neg p$). In fact, rational monotonicity requires of us to be able to draw this conclusion. But in order to do so, we need to be able to conclude that there are no penguins, which violates Strict Entailment. This is the route followed by LM-entailment. The other option would be to insist that we do not have enough information to conclude that there are no penguins, but in the process of doing so, also forego the conclusion that typical non-flying non-robins are not penguins. That is, we insist on Strict Entailment at the expense of rational monotonicity. This is the path followed by PT-entailment. \square

9. Related work

To the best of our knowledge, the first attempt to formalise an explicit notion of typicality in defeasible reasoning was that by Delgrande [24]. Given the strong links between our constructions and the KLM approach, most of the remarks in the comparison made by Lehmann and Magidor [35, Section 3.7] are applicable in comparing Delgrande's approach to ours and therefore we shall not repeat them here.

Crocco and Lamarre [23] as well as Boutilier [6] have explored the links between conditionals and notions of normality similar to the one we investigate here. In particular, Boutilier defines a family of conditional logics of normality in which a statement of the form "if α , then normally β " is formalised via a binary modality \Rightarrow as a conditional $\alpha \Rightarrow \beta$. Here we achieve the same with a unary operator.

Roughly speaking, Boutilier's semantic intuition is the same as that of KLM (and therefore the same as ours). The main difference is that Boutilier defines a conditional connective \Rightarrow in the language, whereas Kraus et al. define \sim at a meta-

level to the language. In this respect, Boutilier's approach is more general in that it allows for nested conditionals. If these are omitted, i.e., if one works in the 'flat' conditional logic in which \Rightarrow is the main connective and no nesting is allowed, then one gets the same results for preferential entailment with both systems. So Boutilier achieves with modalities (he works in a bi-modal language) what Kraus and colleagues achieve with a (meta-level) preference order.

It turns out that in Boutilier's approach one cannot always capture the notion of "most typical α 's". In Boutilier's modal logic, such a set (of *most* normal α -worlds) need not exist in general. This is because Boutilier drops the smoothness condition [6, p. 103] and therefore at any point in a ranked model one can have infinitely descending chains of increasingly more normal α -worlds. If one imposes smoothness in Boutilier's approach, which can be done by e.g. requiring the ordering determined by Boutilier's \square also to be *Noetherian*, one could then define his conditional \Rightarrow more elegantly as follows:

$$\alpha \Rightarrow \beta := \bullet\alpha \rightarrow \beta \quad (1)$$

where, in Boutilier's notation, $\bullet\alpha$ would be given by

$$\bullet\alpha := \alpha \wedge \square\neg\alpha \quad (2)$$

(Of course negated conditionals of the form $\alpha \not\Rightarrow \beta$ can then be expressed as $\neg(\bullet\alpha \rightarrow \beta)$.) In adopting smoothness and defining conditionals in this way, one would expect both approaches to become equivalent modulo the underlying language – ours is propositional, whereas Boutilier's is modal. However, our statement $\bullet\alpha \rightarrow \beta$ differs from Boutilier's $\alpha \Rightarrow \beta$ in a significant way. In Boutilier's approach, a statement of the form $\alpha \Rightarrow \beta$ is true at some world (in a ranked model) if and only if it is true at *all* worlds in that ranked model [6, p. 114]. On the other hand, it is not hard to find a ranked model in which $\bullet\alpha \rightarrow \beta$ holds at a world without being true in the whole model. This establishes Boutilier's conditional as a 'global' statement, while ours has the (more general) 'local flavour'. We can easily simulate Boutilier's notion of *acceptance* [6, p. 115] by stating $\top \rightarrow (\bullet\alpha \rightarrow \beta)$.

It is also worth mentioning that our interpretation of the conditional \Rightarrow in (1) above and Boutilier's differ in another subtle way, which also relates to whether one adopts smoothness or not. In (1), $\alpha \Rightarrow \beta$ is defined as "the normal α 's are β 's", whereas, strictly speaking, Boutilier's definition of $\alpha \Rightarrow \beta$ reads as "there is a point from which $\alpha \rightarrow \beta$ is not violated". Such a 'frontier' for normality, implicitly referred to in Boutilier's definition of $\alpha \Rightarrow \beta$, is not as crisp as ours in the sense that the point where one draws the normality line might be too 'far away' (in the ordering) from the more and more normal α -worlds. One can definitely make a case for dropping the smoothness condition, but requiring it is a small price to pay given the much simpler account of typicality one obtains.

When it comes to entailment from a defeasible knowledge base, all approaches discussed above adopt a Tarskian-style notion of consequence and therefore do not go beyond ranked entailment. The move towards a non-monotonic notion of entailment and an investigation of its different facets in the context of PTL was precisely our motivation in the present work.

Giordano et al. [29] proposed the system P_{min} which is based on a language that is as expressive as the one we propose in this paper. However, they end up using a constrained form of such a language that goes only slightly beyond the expressivity of the language of KLM-style conditionals (their *well-behaved knowledge bases*). More importantly, their approach differs from ours since they build P_{min} on a semantic approach that relies on preferential models and a notion of minimality that is more akin to circumscription [38].

In a description logic setting, Giordano et al. [25] also study notions of typicality. Semantically, they do so by placing an (absolute) ordering on *objects* in first-order domains in order to define versions of defeasible subsumption relations in the description logic \mathcal{ALC} . The authors moreover extend the language of \mathcal{ALC} with an explicit typicality operator $\mathbf{T}(\cdot)$ of which the intended meaning is to single out instances of a concept that are deemed as 'typical'. That is, given an \mathcal{ALC} concept C , $\mathbf{T}(C)$ denotes the most typical individuals having the property of being C in a particular DL interpretation. It is worth pointing out, though, that most of the analysis in the work of Giordano et al. is dedicated to a constrained use of the typicality operator $\mathbf{T}(\cdot)$, that is allowed to occur only in the left-hand side of GCIs and not in the scope of other concept constructors. Not having such a syntactic constraint is a feature of our approach that we have put forward in the present work. Still in the framework of Description Logics, also Bonatti et al. [1] introduce a typicality operator $N(\cdot)$, with a meaning that mirrors the operator $\mathbf{T}(\cdot)$; also the use of the N operator is generally constrained, and their semantic framework is differs from the present one, not being preferential.

Giordano et al.'s approach has been extended in a series of papers [26,27,31,32], in particular also to deal with the computation of non-monotonic entailment from defeasible knowledge bases. In the latter case, the authors define a hyper-tableau calculus to compute the rational closure of a defeasible ontology via a minimal-model construction [30,32] that, as mentioned before, is closely related to our notion of PT-entailment. Nevertheless, that remains the only notion of non-monotonic entailment investigated by the authors. We conjecture the more expressive DL setting has the potential to give rise to a much broader spectrum of consequence relations when enriched with typicality operators, in particular when these apply not only to concepts but also to roles [44]. Nevertheless, that remains the only notion of non-monotonic entailment investigated by the authors. We conjecture the more expressive DL setting has the potential to give rise to a much broader spectrum of consequence relations when enriched with typicality operators, in particular when these apply not only to concepts but also to roles [44].

Finally, Britz and Varzinczak [11,12] investigate another, complementary aspect of defeasibility to the one here presented by introducing (non-standard) modal operators allowing us to talk about relative normality in accessible worlds. With their defeasible versions of modalities, namely \boxdot and \boxless , formalising respectively the notions of *defeasible necessity* and *distinct possibility*, it becomes possible to make statements of the form “ α holds in all of the normal (typical) accessible worlds”, thereby capturing defeasibility of what is ‘expected’ in target worlds. (Note that this is different from stating something like $\Box \bullet \alpha$, which says that all accessible worlds are typical α -worlds.) Such preferential versions of modalities allow for the definition of a family of modal logics in which defeasible modes of inference such as defeasible actions, knowledge and obligations can be expressed. These can be integrated either with existing \sim -based modal logics [14,16] or with a modal extension of our typicality operator in striving towards a comprehensive theory of defeasible reasoning in more expressive languages.

10. Conclusion

The focus of this paper is an investigation into the entailment problem for the logic PTL. We approached the problem from two angles: an abstract formal perspective, in which a set of appropriate postulates was presented and discussed, and a constructive perspective, in which three specific entailment relations were defined and studied. The primary conclusion to be drawn from this investigation is that a logic as expressive as PTL supports more than one form of entailment. This conclusion is supported from the abstract perspective via an impossibility result, as well as through the constructive approach via the definition of two of the three distinct types of PTL entailment: LM-entailment and PT-entailment. While both forms of entailment are generalisations of *rational closure*, only one, LM-entailment, retains all the rationality properties associated with rational closure, formalised as the Conditional Rationality postulate (P5). However, it does not satisfy Strict Entailment (P8), a postulate which requires an entailment relation to remain Tarskian for conclusions not involving typicality, although it satisfies weakened versions of Strict Entailment (P9 and P9’). On the other hand, the other form of entailment we studied, PT-entailment, satisfies P8, but not Conditional Rationality (P5).

The framework of Booth et al. [3,4] is, to the best of our knowledge, the first attempt to introduce a full-fledged typicality operator into propositional logic. In terms of other related work, the closest we are aware of is the restricted form of typicality for description logics by Giordano et al. [28]. However, a consequence of their restricted use of typicality is that a propositional version of their logic would correspond to a KLM-style conditional logic in which rational closure behaves well, and which is much less expressive than PTL.

Britz et al. [13] and Giordano et al. [28] have investigated the connection between the KLM approach and Gödel-Löb modal logic, which is closely related to PTL. Exploiting this connection should deliver an axiomatisation of an inference relation corresponding to ranked entailment, but it does not seem useful for modelling entailment relations based on minimisation as LM- and PT-entailment.

For future work, an obvious open question is whether our conjecture, that the subsets of postulates satisfied by LM-entailment and PT-entailment respectively provide appropriate abstract formalisations of two distinct forms of PTL entailment, can be formalised through representation theorems. From a computational perspective, it is worth investigating whether, as is the case for rational closure for conditional logics, the computation of (the different forms of) PTL entailment can be reduced to a series of classical entailment checks.

Our results in the propositional setting pave the way for an investigation of appropriate forms of entailment in other, more expressive, preferential approaches, such as preferential description logics [15,31,17,8,9,22,18] and modal logics of defeasibility [14,7,10,12]. The move to logics with more structure is of a challenging nature, and a simple rephrasing of our approach to these logics may not deliver the expected results. We are currently investigating these issues.

Declaration of competing interest

The authors confirm that there are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

Acknowledgements

The work of Giovanni Casini and of Thomas Meyer has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 690974 (MIREL project). The work of Thomas Meyer has been supported in part by the National Research Foundation of South Africa (grant No. UID 98019).

Appendix A. Proofs for Section 2

We give here a proof of Proposition 2.1. In order to do that, we need to introduce some extra notions and prove some extra propositions. First of all, analogously to the definitions for PTL introduced in Section 2.3, we say that a set of conditionals \mathcal{C} is *satisfiable* iff there is a ranked interpretation \mathcal{R} for which $\mathcal{U}^{\mathcal{R}} \neq \emptyset$ satisfying all the conditionals in it, and let $Mod(\mathcal{C})$ be the set of all the ranked models of \mathcal{C} .

We are going to use a notion of *merging* ranked interpretations.

Definition A.1 (Ranked Union). Given a set of ranked interpretations $\mathfrak{R} = \{\mathcal{R}_1, \dots, \mathcal{R}_n\}$, its **ranked union** $\mathcal{R}^{\mathfrak{R}}$ is defined as follows:

- for every $v, v' \in \mathcal{U}$, $v <_{\mathcal{R}^{\mathfrak{R}}} v'$ iff $\min\{\mathcal{R}_i(v) \mid \mathcal{R}_i \in \mathfrak{R}\} < \min\{\mathcal{R}_j(v') \mid \mathcal{R}_j \in \mathfrak{R}\}$.
- $\mathcal{R}^{\mathfrak{R}} := (L_0^{\mathcal{R}^{\mathfrak{R}}}, \dots, L_{n-1}^{\mathcal{R}^{\mathfrak{R}}}, L_{\infty}^{\mathcal{R}^{\mathfrak{R}}})$ is defined as
 - $L_{\infty}^{\mathcal{R}^{\mathfrak{R}}} := \bigcap \{L_{\infty}^{\mathcal{R}_i} \mid \mathcal{R}_i \in \mathfrak{R}\}$.
 - $L_0^{\mathcal{R}^{\mathfrak{R}}} := \min_{<_{\mathcal{R}^{\mathfrak{R}}}} (\mathcal{U} \setminus L_{\infty}^{\mathcal{R}^{\mathfrak{R}}})$; $L_1^{\mathcal{R}^{\mathfrak{R}}} := \min_{<_{\mathcal{R}^{\mathfrak{R}}}} (\mathcal{U} \setminus L_0^{\mathcal{R}^{\mathfrak{R}}} \cup L_{\infty}^{\mathcal{R}^{\mathfrak{R}}})$; and so on until $L_n^{\mathcal{R}^{\mathfrak{R}}} = \emptyset$.

Proposition A.1. Let \mathcal{C} be a satisfiable set of conditionals, and let $\mathfrak{R} := \{\mathcal{R}_1, \dots, \mathcal{R}_n\}$ be a set of models of \mathcal{C} . Then their ranked union $\mathcal{R}^{\mathfrak{R}}$ is a model of \mathcal{C} , and $\mathcal{R}^{\mathfrak{R}} \triangleleft_{\text{LM}} \mathcal{R}_i$ for every $\mathcal{R}_i \in (\mathfrak{R} \setminus \mathcal{R}^{\mathfrak{R}})$.

Proof. We first prove that $\mathcal{R}^{\mathfrak{R}} := (L_0^{\mathcal{R}^{\mathfrak{R}}}, \dots, L_{n-1}^{\mathcal{R}^{\mathfrak{R}}}, L_{\infty}^{\mathcal{R}^{\mathfrak{R}}})$ is a model of \mathcal{C} . For every $v \in L_0^{\mathcal{R}^{\mathfrak{R}}}$, it must be the case, by Definition A.1, that $v \in L_0^{\mathcal{R}_i}$ for some $\mathcal{R}_i \in \mathfrak{R}$; since such \mathcal{R}_i is a model of \mathcal{C} and $v \in L_0^{\mathcal{R}_i}$, $v \Vdash \neg\alpha \vee \beta$ for every $\alpha \sim \beta \in \mathcal{C}$.

Now, let $v \in L_i^{\mathcal{R}^{\mathfrak{R}}}$, with $0 < i < n$, s.t. $v \Vdash \alpha \wedge \neg\beta$ for some $\alpha \sim \beta \in \mathcal{C}$ (if there is no such v , $\mathcal{R}^{\mathfrak{R}}$ is necessarily a model of \mathcal{C}). Being every $\mathcal{R}_j \in \mathfrak{R}$ a model of \mathcal{C} , it must be the case, again by Definition A.1, that $h_{\mathcal{R}_j}(v) \geq i$ and $h_{\mathcal{R}_j}(v') < h_{\mathcal{R}_j}(v)$ for some v' satisfying $\alpha \wedge \beta$. Hence, it must be that $\min\{\mathcal{R}_i(v') \mid \mathcal{R}_i \in \mathfrak{R}\} < \min\{\mathcal{R}_j(v) \mid \mathcal{R}_j \in \mathfrak{R}\}$ for some v' satisfying $\alpha \wedge \beta$, that is, $v' <_{\mathcal{R}^{\mathfrak{R}}} v$, that implies $v' \in L_j^{\mathcal{R}^{\mathfrak{R}}}$, with $j < i$.

To summarise, for every $\alpha \sim \beta \in \mathcal{C}$, if there is a valuation $v \in \mathcal{R}^{\mathfrak{R}}$ s.t. $v \Vdash \alpha \wedge \neg\beta$, then there is a valuation $v' \in \mathcal{R}^{\mathfrak{R}}$ s.t. $v' \Vdash \alpha \wedge \beta$ and $v' <_{\mathcal{R}^{\mathfrak{R}}} v$; hence $\mathcal{R}^{\mathfrak{R}}$ is a model of \mathcal{C} .

Now we prove that $\mathcal{R}^{\mathfrak{R}} \triangleleft_{\text{LM}} \mathcal{R}_i$ for every $\mathcal{R}_i \in \mathfrak{R}$.

Let $\mathcal{R}_i := (L_0, \dots, L_{n-1}, L_{\infty}^{\mathcal{R}_i})$ s.t. $\mathcal{R}_i \in \mathfrak{R}$ and $\mathcal{R}^{\mathfrak{R}} \not\triangleleft_{\text{LM}} \mathcal{R}_i$. That is, there is an i s.t. $L_i^{\mathcal{R}^{\mathfrak{R}}} \not\subseteq L_i$, while $L_j^{\mathcal{R}^{\mathfrak{R}}} = L_j$ for every $j < i$. That is, there is a $v \in L_i$ s.t. $v \notin L_i^{\mathcal{R}^{\mathfrak{R}}}$. By definition of $\mathcal{R}^{\mathfrak{R}}$, that implies that $v \in L_j^{\mathcal{R}^{\mathfrak{R}}}$ for some $j < i$, but that cannot be the case, since $L_j^{\mathcal{R}^{\mathfrak{R}}} = L_j$ for every $j < i$. Hence $\mathcal{R}^{\mathfrak{R}} \triangleleft_{\text{LM}} \mathcal{R}_i$ for every $\mathcal{R}_i \in \mathfrak{R}$.

We finish by proving that $\mathcal{R}_i \triangleleft_{\text{LM}} \mathcal{R}^{\mathfrak{R}}$ for every $\mathcal{R}_i \in (\mathfrak{R} \setminus \mathcal{R}^{\mathfrak{R}})$.

Let \mathcal{R}_i be a model in \mathcal{C} s.t. $\mathcal{R}_i \triangleleft_{\text{LM}} \mathcal{R}^{\mathfrak{R}}$. Since $\mathcal{R}^{\mathfrak{R}} \triangleleft_{\text{LM}} \mathcal{R}_i$, we must conclude that for every i , for every cell $L_i^{\mathcal{R}_i}$ composing \mathcal{R}_i and every cell $L_i^{\mathcal{R}^{\mathfrak{R}}}$ composing $\mathcal{R}^{\mathfrak{R}}$, $L_i^{\mathcal{R}_i} = L_i^{\mathcal{R}^{\mathfrak{R}}}$; that is, \mathcal{R}_i and $\mathcal{R}^{\mathfrak{R}}$ are exactly the same model. Hence $\mathcal{R}^{\mathfrak{R}} \triangleleft_{\text{LM}} \mathcal{R}_i$ for every $\mathcal{R}_i \in (\mathfrak{R} \setminus \mathcal{R}^{\mathfrak{R}})$. \square

Proposition A.2. Let \mathcal{C} be a satisfiable set of conditionals. Then the ranked union of the elements of $\text{Mod}(\mathcal{C})$ is the only $\triangleleft_{\text{LM}}$ -minimum element in $\text{Mod}(\mathcal{C})$.

Proof. It is an immediate consequence of Definition A.1 and Proposition A.1. \square

Proposition 2.1. Given a set of conditionals \mathcal{C} and a conditional $\alpha \sim \beta$, $\alpha \sim \beta$ is in the rational closure of \mathcal{C} iff:

1. \mathcal{C} is unsatisfiable; or
2. $\mathcal{R}^{\text{rc}}(\mathcal{C}) \Vdash \alpha \sim \beta$.

Proof. It has been proved by Booth and Paris [2, Theorem 2] that the rational closure of a KB is determined by a model that is equivalent to the ranked union of $\text{Mod}(\mathcal{C})$. According to Proposition A.2, the ranked union of $\text{Mod}(\mathcal{C})$ is the only $\triangleleft_{\text{LM}}$ -minimum element in $\text{Mod}(\mathcal{C})$, that is, the model $\mathcal{R}^{\text{rc}}(\mathcal{C})$. \square

Appendix B. Proofs of Lemmas 5.1, 5.2 and 5.3

B.1. Proof of Lemma 5.1

Lemma 5.1. The following hold for each $i \geq 0$:

1. $S_i \subseteq S_{i+1}$, i.e., $[S_0 \subseteq S_1$ and, for all $i \geq 0$, $[[\mathcal{KB}]]^{\mathcal{R}_i} \subseteq [[\mathcal{KB}]]^{\mathcal{R}_{i+1}}$];
2. For all $v_1, v_2 \in \mathcal{U}$, if $\mathcal{R}_i(v_1) < \mathcal{R}_i(v_2)$, then $v_1 \in [[\mathcal{KB}]]^{\mathcal{R}_i}$;
3. \mathcal{R}_i is a ranked interpretation.

Proof. We show all three simultaneously by complete induction on i . So, assume all of Items 1, 2 and 3 hold for all $m < i$. We will show this implies all three hold also for i . We assume each $\alpha \in \mathcal{KB}$ is in normal form.

1. $[[\mathcal{KB}]]^{\mathcal{R}_i} \subseteq [[\mathcal{KB}]]^{\mathcal{R}_{i+1}}$.

Let $v \in \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i}$ and let $\alpha \in \mathcal{KB}$ with $\alpha = \bigwedge_{i \leq t} \bullet \theta_i \rightarrow (\phi \vee \bigvee_{i \leq s} \bullet \psi_i)$ (for some $s, t \geq 0$). We must show $v \in \llbracket \alpha \rrbracket^{\mathcal{R}_{i+1}}$. Since $v \in \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i}$ we know $v \in \llbracket \alpha \rrbracket^{\mathcal{R}_i}$. Hence we know that one of the following must hold:

- $v \notin \llbracket \bullet \theta_k \rrbracket^{\mathcal{R}_i}$ for some k : This means (since θ_k is propositional) v is not $\prec^{\mathcal{R}_i}$ -minimal in $\llbracket \theta_k \rrbracket^{\mathcal{R}_i} = \llbracket \theta_k \rrbracket^{\mathcal{R}_{i+1}}$. But then it is also not $\prec^{\mathcal{R}_{i+1}}$ -minimal since, by construction, if $\mathcal{R}_i(v) \leq \mathcal{R}_i(w)$ then $\mathcal{R}_{i+1}(v) \leq \mathcal{R}_{i+1}(w)$. Hence in this case $v \notin \llbracket \bullet \theta_k \rrbracket^{\mathcal{R}_{i+1}}$.
- $v \in \llbracket \phi \rrbracket^{\mathcal{R}_i}$: In this case also $v \in \llbracket \phi \rrbracket^{\mathcal{R}_{i+1}}$, since $\llbracket \phi \rrbracket^{\mathcal{R}_i} = \llbracket \phi \rrbracket^{\mathcal{R}_{i+1}}$ (because ϕ is purely propositional).
- $v \in \llbracket \bullet \psi_k \rrbracket^{\mathcal{R}_i}$ for some k : This means v is $\prec^{\mathcal{R}_i}$ -minimal in $\llbracket \psi_k \rrbracket^{\mathcal{R}_i}$. But then it is also $\prec^{\mathcal{R}_{i+1}}$ -minimal, since we assumed $v \in \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i} = S_{i+1}$, and so by construction of \mathcal{R}_{i+1} we have that $\mathcal{R}_{i+1}(w) < \mathcal{R}_{i+1}(v)$ if and only if $\mathcal{R}_i(w) < \mathcal{R}_i(v)$ for all $w \in \mathcal{U}$. Since $\llbracket \psi_k \rrbracket^{\mathcal{R}_i} = \llbracket \psi_k \rrbracket^{\mathcal{R}_{i+1}}$ (since ψ_k is purely propositional) we obtain that v is $\prec^{\mathcal{R}_{i+1}}$ -minimal in $\llbracket \psi_k \rrbracket^{\mathcal{R}_{i+1}}$, i.e., $v \in \llbracket \bullet \psi_k \rrbracket^{\mathcal{R}_{i+1}}$.

Thus in all three possible cases we obtain $v \in \llbracket \alpha \rrbracket^{\mathcal{R}_{i+1}}$ as required.

2. $\mathcal{R}_i(v_1) < \mathcal{R}_i(v_2)$ implies $v_1 \in \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i}$.

Suppose $\mathcal{R}_i(v_1) < \mathcal{R}_i(v_2)$. Observe that, by construction, this can only be the case if $i > 0$. Then either $\mathcal{R}_{i-1}(v_1) < \mathcal{R}_{i-1}(v_2)$ or $v_2 \notin S_i$. If $\mathcal{R}_{i-1}(v_1) < \mathcal{R}_{i-1}(v_2)$ then, by the inductive hypothesis, $v_1 \in \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_{i-1}}$, while if $v_2 \notin S_i$, then $v_1 \in S_i = \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_{i-1}}$. So either way we get $v_1 \in \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_{i-1}}$ and so we get the desired conclusion by applying $\llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_{i-1}} \subseteq \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_i}$ which was just proved in Item 1 above.

3. \mathcal{R}_i is a ranked interpretation.

By construction it immediately follows that \mathcal{R}_i is a function from \mathcal{U} to $\mathbb{N} \cup \{\infty\}$. We need to show the convexity property: if $\mathcal{R}_i(u) = j$ then, for every k such that $0 \leq k < j$, there is a $v \in \mathcal{U}$ for which $\mathcal{R}_i(v) = k$. If $i = 0$, this follows immediately (since $\mathcal{R}_0(u) = 0$ for all $u \in \mathcal{U}$). Otherwise we have by the inductive hypothesis that \mathcal{R}_{i-1} is a ranked interpretation. We have two cases. (1) $S_i = S_{i-1}$: Then $\mathcal{R}_i = (\mathcal{R}_{i-1})_{S_i}^\infty$ from which convexity follows immediately. (2) $S_i \neq S_{i-1}$: Then $\mathcal{R}_i = (\mathcal{R}_{i-1})_{S_i}^1$ from which convexity also follows immediately. \square

B.2. Proof of Lemma 5.2

Lemma 5.2. For every $i > 0$, $(\mathcal{R}_i)_{S_i}^\infty$ is a ranked model of \mathcal{KB} .

Proof. Let \mathcal{R} denote $(\mathcal{R}_i)_{S_i}^\infty$. We need to show that for every valuation $v \in \mathcal{U}^{\mathcal{R}}$, i.e., every $v \in S_i = \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_{i-1}}$, and for every $\alpha \in \mathcal{KB}$, we have $v \in \llbracket \alpha \rrbracket^{\mathcal{R}}$. Since $v \in \llbracket \alpha \rrbracket^{\mathcal{R}_{i-1}}$ we know one of the following must hold (recalling that α is expressed in normal form $\bigwedge_{i \leq t} \bullet \theta_i \rightarrow (\phi \vee \bigvee_{i \leq s} \bullet \psi_i)$):

- $v \notin \llbracket \bullet \theta_k \rrbracket^{\mathcal{R}_{i-1}}$ for some k : This means v is not $\prec^{\mathcal{R}_{i-1}}$ -minimal in $\llbracket \theta_k \rrbracket^{\mathcal{R}_{i-1}}$. But then it is also not $\prec^{\mathcal{R}}$ -minimal in $\llbracket \theta_k \rrbracket^{\mathcal{R}} = \llbracket \theta_k \rrbracket^{\mathcal{R}_{i-1}} \cap S_i$, since if $w \prec^{\mathcal{R}_{i-1}} v$ and $w \in \llbracket \theta_k \rrbracket^{\mathcal{R}_{i-1}}$, then from the former we know $w \in S_i$ by Item 2 of the previous lemma. Hence in this case $v \notin \llbracket \bullet \theta_k \rrbracket^{\mathcal{R}}$.
- $v \in \llbracket \phi \rrbracket^{\mathcal{R}_{i-1}}$: In this case also $v \in \llbracket \phi \rrbracket^{\mathcal{R}}$, since $\llbracket \phi \rrbracket^{\mathcal{R}} = \llbracket \phi \rrbracket^{\mathcal{R}_{i-1}} \cap S_i$ (because ϕ is purely propositional).
- $v \in \llbracket \bullet \psi_k \rrbracket^{\mathcal{R}_{i-1}}$ for some k : This means v is $\prec^{\mathcal{R}_{i-1}}$ -minimal in $\llbracket \psi_k \rrbracket^{\mathcal{R}_{i-1}}$. But then it is also $\prec^{\mathcal{R}}$ -minimal in $\llbracket \psi_k \rrbracket^{\mathcal{R}} = \llbracket \psi_k \rrbracket^{\mathcal{R}_{i-1}} \cap S_i$, since $\prec^{\mathcal{R}_{i-1}} \subseteq \prec^{\mathcal{R}}$. Hence $v \in \llbracket \bullet \psi_k \rrbracket^{\mathcal{R}}$.

Thus in all three possible cases we obtain $v \in \llbracket \alpha \rrbracket^{\mathcal{R}}$ as required. \square

B.3. Proof of Lemma 5.3

Lemma 5.3. Let $\mathcal{R}_{\mathcal{KB}}^* := (L_0, \dots, L_{n-1}, L_\infty)$ and let $\mathcal{R} := (M_0, \dots, M_{n-1}, M_\infty)$ be any other ranked model of \mathcal{KB} . Let $i \in \{0, \dots, n-1\}$. If $L_j = M_j$ for all $j < i$, then $M_i \subseteq L_i$.

Proof. Let $v \in M_i$. By construction, $S_i = \llbracket \mathcal{KB} \rrbracket^{\mathcal{R}_{i-1}}$ where $\mathcal{R}_{i-1} = (L_0, \dots, L_{i-1}, (\mathcal{U} \setminus \bigcup_{j < i} L_j), \emptyset)$. Let $\alpha \in \mathcal{KB}$, with $\alpha = \bigwedge_{i \leq t} \bullet \theta_i \rightarrow (\phi \vee \bigvee_{i \leq s} \bullet \psi_i)$ (for some $s, t \geq 0$). We must show v satisfies α in \mathcal{R}_{i-1} , so assume v satisfies $\neg \phi$ and is a minimal θ_k -state in \mathcal{R}_{i-1} for all k . We must show v is a minimal ψ_k -state in \mathcal{R}_{i-1} for at least one k . Since we assume $M_j = L_j$ for all $j < i$, we have $\mathcal{R}_{i-1} = (M_0, \dots, M_{i-1}, (\mathcal{U} \setminus \bigcup_{j < i} M_j), \emptyset)$. Since $v \in M_i$, we can show that, for any propositional sentence λ , we have that v is a minimal λ -state in $(M_0, \dots, M_{i-1}, (\mathcal{U} \setminus \bigcup_{j < i} M_j), \emptyset)$ if and only if it is a minimal λ -state in $(M_0, \dots, M_i, \emptyset)$. Thus, from the fact that $(M_0, \dots, M_i, \emptyset)$ is a ranked model of \mathcal{KB} , we obtain our conclusion. \square

References

- [1] P. Bonatti, M. Faella, I. Petrova, L. Sauro, A new semantics for overriding in description logics, *Artif. Intell.* 222 (2015) 1–48.
- [2] R. Booth, J. Paris, A note on the rational closure of knowledge bases with both positive and negative knowledge, *J. Log. Lang. Inf.* 7 (2) (1998) 165–190.
- [3] R. Booth, T. Meyer, I. Varzinczak, PTL: a propositional typicality logic, in: L. Fariñas del Cerro, A. Herzig, J. Mengin (Eds.), *Proceedings of the 13th European Conference on Logics in Artificial Intelligence, JELIA*, in: LNCS, vol. 7519, Springer, 2012, pp. 107–119.

- [4] R. Booth, T. Meyer, I. Varzinczak, A propositional typicality logic for extending rational consequence, in: E. Fermé, D. Gabbay, G. Simari (Eds.), Trends in Belief Revision and Argumentation Dynamics, in: Studies in Logic – Logic and Cognitive Systems, vol. 48, King's College Publications, 2013, pp. 123–154.
- [5] R. Booth, G. Casini, T. Meyer, I. Varzinczak, On the entailment problem for a logic of typicality, in: Proceedings of the 24th International Joint Conference on Artificial Intelligence, IJCAI, 2015, pp. 2805–2811.
- [6] C. Boutilier, Conditional logics of normality: a modal approach, *Artif. Intell.* 68 (1) (1994) 87–154.
- [7] K. Britz, I. Varzinczak, Defeasible modalities, in: Proceedings of the 14th Conference on Theoretical Aspects of Rationality and Knowledge, TARK, 2013, pp. 49–60.
- [8] K. Britz, I. Varzinczak, Introducing role defeasibility in description logics, in: L. Michael, A. Kakas (Eds.), Proceedings of the 15th European Conference on Logics in Artificial Intelligence, JELIA, in: LNCS, vol. 10021, Springer, 2016, pp. 174–189.
- [9] K. Britz, I. Varzinczak, Toward defeasible \mathcal{SRLQ} , in: Proceedings of the 30th International Workshop on Description Logics, 2017.
- [10] K. Britz, I. Varzinczak, From KLM-style conditionals to defeasible modalities, and back, *J. Appl. Non-Class. Log.* (2017).
- [11] K. Britz, I. Varzinczak, From KLM-style conditionals to defeasible modalities, and back, *J. Appl. Non-Class. Log.* 28 (1) (2018) 92–121.
- [12] K. Britz, I. Varzinczak, Preferential accessibility and preferred worlds, *J. Log. Lang. Inf.* 27 (2) (2018) 133–155.
- [13] K. Britz, J. Heidema, T. Meyer, Modelling object typicality in description logics, in: A. Nicholson, X. Li (Eds.), Proceedings of the 22nd Australasian Joint Conference on Artificial Intelligence, in: LNAI, vol. 5866, Springer, 2009, pp. 506–516.
- [14] K. Britz, T. Meyer, I. Varzinczak, Preferential reasoning for modal logics, in: Proceedings of the 7th Workshop on Methods for Modalities, M4M'2011, in: Electronic Notes in Theoretical Computer Science, vol. 278, 2011, pp. 55–69.
- [15] K. Britz, T. Meyer, I. Varzinczak, Semantic foundation for preferential description logics, in: D. Wang, M. Reynolds (Eds.), Proceedings of the 24th Australasian Joint Conference on Artificial Intelligence, in: LNAI, vol. 7106, Springer, 2011, pp. 491–500.
- [16] K. Britz, T. Meyer, I. Varzinczak, Normal modal preferential consequence, in: M. Thielscher, D. Zhang (Eds.), Proceedings of the 25th Australasian Joint Conference on Artificial Intelligence, in: LNAI, vol. 7691, Springer, 2012, pp. 505–516.
- [17] K. Britz, G. Casini, T. Meyer, K. Moodley, U. Sattler, I. Varzinczak, Theoretical Foundations of Defeasible Description Logics, Technical report, 2019, <http://arxiv.org/abs/1904.07559>.
- [18] K. Britz, G. Casini, T. Meyer, I. Varzinczak, A KLM perspective on defeasible reasoning for description logics, in: Description Logic, Theory Combination, and All That – Essays Dedicated to Franz Baader on the Occasion of His 60th Birthday, 2019, pp. 147–173.
- [19] G. Casini, U. Straccia, Rational closure for defeasible description logics, in: T. Janhunen, I. Niemelä (Eds.), Proceedings of the 12th European Conference on Logics in Artificial Intelligence, JELIA, in: LNCS, vol. 6341, Springer-Verlag, 2010, pp. 77–90.
- [20] G. Casini, U. Straccia, Defeasible inheritance-based description logics, in: Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI 2011, Barcelona, Catalonia, Spain, July 16–22, 2011, 2011, pp. 813–818.
- [21] G. Casini, U. Straccia, Defeasible inheritance-based description logics, *J. Artif. Intell. Res.* 48 (2013) 415–473.
- [22] G. Casini, U. Straccia, T. Meyer, A polynomial time subsumption algorithm for nominal safe \mathcal{ELC}_{\perp} under rational closure, *Inf. Sci.* 50 (2018) 588–620, <https://doi.org/10.1016/j.ins.2018.09.037>.
- [23] G. Crocco, P. Lamarre, On the connections between nonmonotonic inference systems and conditional logics, in: R. Nebel, C. Rich, W. Swartout (Eds.), Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning, KR, Morgan Kaufmann Publishers, 1992, pp. 565–571.
- [24] J. Delgrande, A first-order logic for prototypical properties, *Artif. Intell.* 33 (1987) 105–130.
- [25] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, Preferential description logics, in: N. Dershowitz, A. Voronkov (Eds.), Logic for Programming, Artificial Intelligence, and Reasoning, LPAR, in: LNAI, vol. 4790, Springer, 2007, pp. 257–272.
- [26] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, Reasoning about typicality in preferential description logics, in: S. Hölldobler, C. Lutz, H. Wansing (Eds.), Proceedings of the 11th European Conference on Logics in Artificial Intelligence, JELIA, in: LNAI, vol. 5293, Springer, 2008, pp. 192–205.
- [27] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, Analytic tableaux calculi for KLM logics of nonmonotonic reasoning, *ACM Trans. Comput. Log.* 10 (3) (2009) 18:1–18:47.
- [28] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, $\mathcal{ALC} + T$: a preferential extension of description logics, *Fundam. Inform.* 96 (3) (2009) 341–372.
- [29] L. Giordano, V. Gliozzi, N. Olivetti, G.L. Pozzato, A nonmonotonic extension of KLM preferential logic P, in: Logic for Programming, Artificial Intelligence, and Reasoning – Proceedings of 17th International Conference, LPAR-17, Yogyakarta, Indonesia, October 10–15, 2010, 2010, pp. 317–332.
- [30] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, A minimal model semantics for nonmonotonic reasoning, in: L. Fariñas del Cerro, A. Herzig, J. Mengin (Eds.), Proceedings of the 13th European Conference on Logics in Artificial Intelligence, JELIA, in: LNCS, vol. 7519, Springer, 2012, pp. 228–241.
- [31] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, A non-monotonic description logic for reasoning about typicality, *Artif. Intell.* 195 (2013) 165–202.
- [32] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, Semantic characterization of rational closure: from propositional logic to description logics, *Artif. Intell.* 226 (2015) 1–33.
- [33] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, *Artif. Intell.* 44 (1990) 167–207.
- [34] D. Lehmann, Another perspective on default reasoning, *Ann. Math. Artif. Intell.* 15 (1) (1995) 61–82.
- [35] D. Lehmann, M. Magidor, What does a conditional knowledge base entail? *Artif. Intell.* 55 (1992) 1–60.
- [36] D. Makinson, General patterns in nonmonotonic reasoning, in: Handbook of Logic in Artificial Intelligence and Logic Programming, vol. III, Clarendon Press, Oxford, 1994, pp. 35–110.
- [37] D. Makinson, Bridges from Classical to Nonmonotonic Logic, Texts in Computing, vol. 5, King's College Publications, 2005.
- [38] J. McCarthy, Circumscription, a form of nonmonotonic reasoning, *Artif. Intell.* 13 (1–2) (1980) 27–39.
- [39] J. Pearl, System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: Proceedings of the 3rd Conference on Theoretical Aspects of Rationality and Knowledge, TARK, 1990.
- [40] J. Pearl, M. Goldszmidt, On the relation between rational closure and system-z, in: Proceedings of the 3rd International Workshop on Non-Monotonic Reasoning, 1990.
- [41] J. Quantz, V. Royer, A preference semantics for defaults in terminological logics, in: Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning, KR, 1992, pp. 294–305.
- [42] Y. Shoham, Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence, MIT Press, 1988.
- [43] A. Tarski, On Some Fundamental Concepts of Metamathematics, [1930] Logic, Semantics, Metamathematics. Papers from 1923 to 1938, translated by J.H. Woodger, Clarendon Press, 1956, pp. 30–36.
- [44] I. Varzinczak, A note on a description logic of concept and role typicality for defeasible reasoning over ontologies, *Log. Univers.* 12 (3–4) (2018) 297–325.