BKLM - An expressive logic for defeasible reasoning

Guy Paterson-Jones[2] , **Giovanni Casini**[1,2] , **Thomas Meyer**[2]
[1]ISTI-CNR, Italy
[2]CAIR and Univ. of Cape Town, South Africa
guy.paterson.jones@gmail.com , giovanni.casini@isti.cnr.it , tmeyer@cair.org.za

Abstract

Propositional KLM-style defeasible reasoning involves a core propositional logic capable of expressing defeasible (or conditional) implications. The semantics for this logic is based on Kripke-like structures known as ranked interpretations. KLM-style defeasible entailment is referred to as rational whenever the defeasible entailment relation under consideration generates a set of defeasible implications all satisfying a set of rationality postulates known as the KLM postulates. In a recent paper Booth et al. proposed PTL, a logic that is more expressive than the core KLM logic. They proved an impossibility result, showing that defeasible entailment for PTL fails to satisfy a set of rationality postulates similar in spirit to the KLM postulates. Their interpretation of the impossibility result is that defeasible entailment for PTL need not be unique.

In this paper we continue the line of research in which the expressivity of the core KLM logic is extended. We present the logic Boolean KLM (BKLM) in which we allow for disjunctions, conjunctions, and negations, but not nesting, of defeasible implications. Our contribution is twofold. Firstly, we show (perhaps surprisingly) that BKLM is more expressive than PTL. Our proof is based on the fact that BKLM can characterise all single ranked interpretations, whereas PTL cannot. Secondly, given that the PTL impossibility result also applies to BKLM, we adapt the different forms of PTL entailment proposed by Booth et al. to apply to BKLM.

1 Introduction

Non-monotonic reasoning has been extensively studied in the AI literature, as it provides a mechanism for making bold inferences that go beyond what classical methods can provide, while retaining the possibility of revising these inferences in light of new information. In their seminal paper, Kraus et al. (1990) consider a general framework for non-monotonic reasoning, phrased in terms of *defeasible, or conditional implications* of the form $\alpha \mid\!\sim \beta$, to be read as '*If* α *holds, then typically* β *holds*'. Importantly, they provide a set of *rationality conditions*, in the form of structural properties, that a reasonable form of entailment for these conditionals should satisfy, and characterise these semantically. Lehmann and Magidor (1992) also considered the question of which entailment relations definable in the KLM framework can be considered to be the *correct* ones for non-monotonic reasoning. In general, there is a large class of entailment relations for KLM-style logics (Casini, Meyer, and Varzinczak 2019), and it is widely agreed upon that there is no unique best answer. The options can be narrowed down, however, and Lehmann et al. propose *Rational Closure* (RC) as the minimally acceptable form of rational entailment. Rational closure is based on the principle of *presumption of typicality* (Lehmann 1995), which states that propositions should be considered typical unless there is reason to believe otherwise. For instance, if we know that birds typically fly, and all we know about a robin is that it is a bird, we should tentatively conclude that it flies, as there is no reason to believe it is atypical. While RC is not always appropriate, there is fairly general consensus that interesting forms of conditional reasoning should extend RC from an inferential perspective (Lehmann 1995; Casini, Meyer, and Varzinczak 2019).

Since KLM-style logics have limited conditional expressivity (see Section 2.1), there has been some work in extending the KLM constructions to more expressive logics. Perhaps the main question is whether entailment relations resembling RC can be defined also for more expressive logics. The first investigation in such a direction was proposed by Booth and Paris (1998) who consider an extension in which both positive ($\alpha \mid\!\sim \beta$) and negative ($\alpha \not\mid\!\sim \beta$) conditionals are allowed. Booth et al. (2013) introduce a significantly more expressive logic called *Propositional Typicality Logic* (PTL), in which propositional logic is extended with a modal-like typicality operator $\bullet$. This typicality operator can be used anywhere in a formula, in contrast to KLM-style logics, where typicality refers only to the antecedent of conditionals of the form $\alpha \mid\!\sim \beta$.

The price one pays for this expressiveness is that rational entailment becomes more difficult to pin down. This is shown by Booth et al. (2015), who prove that several desirable properties of rational closure are mutually inconsistent for PTL entailment. They interpret this as saying that the correct form of entailment for PTL is contextual, and depends on which properties are considered more important for the task at hand.

In this paper we consider a different extension of KLM-style logics, which we refer to as *Boolean KLM* (BKLM), and in which we allow negative conditionals, as well as arbitrary conjunctions and disjunctions of conditionals. We do not allow the nesting of conditionals, though. We show, perhaps surprisingly, that BKLM is strictly more expressive than PTL by exhibiting an explicit translation of PTL knowledge bases into BKLM. We also prove that BKLM entailment is more restrictive than PTL entailment, in the sense that a stronger class of entailment properties are inconsistent for BKLM. In particular, attempts to extend rational closure to BKLM in the manner of LM-entailment as defined by Booth et al. (2015), are shown to be untenable.

The rest of the paper is structured as follows. In section 2 we provide the relevant background on the KLM approach to defeasible reasoning, and discuss various forms of rational entailment. We then define Propositional Typicality Logic, and give a brief overview of the entailment problem for PTL. In section 3 we define the logic BKLM, an extension of KLM-style logics that allows for arbitrary boolean combinations of conditionals. We investigate the expressiveness of BKLM, and show that it is strictly more expressive PTL by exhibiting an explicit translation of PTL formulas into BKLM. In section 4 we turn to the entailment problem for BKLM, and show that BKLM suffers from stronger versions of the known impossibility results for PTL. Section 5 discusses some related work, while section 6 concludes and points out some future research directions.

2 Background

Let $\mathcal{P}$ be a set of propositional atoms, and let $p, q, \ldots$ be meta-variables for elements of $\mathcal{P}$. We write $\mathcal{L}^{\mathcal{P}}$ for the set of propositional formulas over $\mathcal{P}$, defined by $\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \top \mid \bot$. Other boolean connectives are defined as usual in terms of $\wedge, \neg, \rightarrow$, and $\leftrightarrow$. We write $\mathcal{U}^{\mathcal{P}}$ for the set of valuations of $\mathcal{P}$, which are functions $v : \mathcal{P} \rightarrow \{0, 1\}$. Valuations are extended to $\mathcal{L}^{\mathcal{P}}$ in the usual way, and satisfaction of a formula α will be denoted $v \Vdash \alpha$. For the remainder of this paper we will assume that $\mathcal{P}$ is finite and drop superscripts whenever there isn't any danger of ambiguity.

2.1 The Logic KLM

Kraus et al. (1990) study a conditional logic, which we refer to as KLM. It is defined by assertions of the form $\alpha \mid\!\sim \beta$, which are read "if α, then typically β". For example, if $\mathcal{P} = \{\mathsf{b}, \mathsf{f}\}$ refers to the properties of being a bird and flying respectively, then $\mathsf{b} \mid\!\sim \mathsf{f}$ states that birds typically fly. There are various possible semantic structures for this logic, but in this paper we are interested in the case of *rational* conditional assertions. The semantics for rational conditionals is given by *ranked interpretations* (Lehmann and Magidor 1992). The following is an alternative, but equivalent definition of such a class of interpretations.

Definition 1. *A ranked interpretation $\mathscr{R}$ is a function from $\mathcal{U}$ to $\mathbb{N} \cup \{\infty\}$ satisfying the following convexity condition: if $\mathscr{R}(u) < \infty$, then for every $0 \leq j < \mathscr{R}(u)$, there is some $v \in \mathcal{U}$ for which $\mathscr{R}(v) = j$.*

Given a ranked interpretation $\mathscr{R}$, we call $\mathscr{R}(u)$ the *rank* of u with respect to $\mathscr{R}$. Valuations with a lower rank are viewed as being more typical than those with a higher rank, whereas valuations with infinite rank are viewed as being impossibly atypical. We refer to the set of *possible valuations* as $\mathcal{U}^{\mathscr{R}} = \{u \in \mathcal{U} : \mathscr{R}(u) < \infty\}$, and for any $\alpha \in \mathcal{L}$ we define $[\![\alpha]\!]^{\mathscr{R}} = \{u \in \mathcal{U}^{\mathscr{R}} : u \Vdash \alpha\}$.

2	pbf
1	$\overline{\mathsf{p}}\mathsf{b}\overline{\mathsf{f}}$, $\mathsf{pb}\overline{\mathsf{f}}$
0	$\overline{\mathsf{p}}\overline{\mathsf{b}}\overline{\mathsf{f}}$, $\overline{\mathsf{p}}\overline{\mathsf{b}}\mathsf{f}$, $\overline{\mathsf{p}}\mathsf{bf}$

Figure 1: A ranked interpretation over $\mathcal{P} = \{\mathsf{p}, \mathsf{b}, \mathsf{f}\}$.

Every ranked interpretation $\mathscr{R}$ determines a total preorder on $\mathcal{U}$ in the obvious way, namely $u \leq_{\mathscr{R}} v$ iff $\mathscr{R}(u) \leq \mathscr{R}(v)$. Writing the strict version of this preorder as $\prec_{\mathscr{R}}$, it is straightforward to show that it is *modular*:

Proposition 1. *$\prec_{\mathscr{R}}$ is* modular*, i.e. for all $u, v, w \in \mathcal{U}$, $u \prec_{\mathscr{R}} v$ implies that either $w \prec_{\mathscr{R}} v$ or $u \prec_{\mathscr{R}} w$.*

Lehmann et al. (1992) define ranked interpretations in terms of modular orderings on $\mathcal{U}$. The following straightforward observation proves the equivalence of the two definitions:

Proposition 2. *Let $\mathscr{R}_1$ and $\mathscr{R}_2$ be ranked interpretations. Then $\mathscr{R}_1 = \mathscr{R}_2$ iff $\prec_{\mathscr{R}_1} = \prec_{\mathscr{R}_2}$.*

We define satisfaction with respect to ranked interpretations as follows. Given any $\alpha \in \mathcal{L}$, we say $\mathscr{R}$ *satisfies* α (written $\mathscr{R} \Vdash \alpha$) iff $[\![\alpha]\!]^{\mathscr{R}} = \mathcal{U}^{\mathscr{R}}$. Similarly, $\mathscr{R}$ satisfies a conditional assertion $\alpha \mid\!\sim \beta$ iff $\min_{\leq_{\mathscr{R}}} [\![\alpha]\!]^{\mathscr{R}} \subseteq [\![\beta]\!]^{\mathscr{R}}$, or in other words iff all of the $\leq_{\mathscr{R}}$-minimal valuations satisfying α also satisfy β.

Example 1. *Let $\mathscr{R}$ be the ranked interpretation in figure 1. Then $\mathscr{R}$ satisfies $\mathsf{p} \rightarrow \mathsf{b}$, $\mathsf{b} \mid\!\sim \mathsf{f}$ and $\mathsf{p} \mid\!\sim \neg\mathsf{f}$. Note that in our figures we omit rank ∞ for brevity, and we represent a valuation as a string of literals, with $\overline{\mathsf{p}}$ indicating the negation of the atom p.*

A useful simplification is the fact that classical statements (such as $\mathsf{p} \rightarrow \mathsf{b}$) can be viewed as special cases of conditional assertions:

Proposition 3. *(Kraus, Lehmann, and Magidor 1990, p.174) For all $\alpha \in \mathcal{L}$, $\mathscr{R} \Vdash \alpha$ iff $\mathscr{R} \Vdash \neg\alpha \mid\!\sim \bot$.*

In what follows we define a *knowledge base* as a finite set of conditional assertions. We sometimes abuse notation by including classical statements (of the form $\alpha \in \mathcal{L}$) in knowledge bases, but in the context of Proposition 3 this should be understood to be shorthand for the conditional assertion $\neg\alpha \mid\!\sim \bot$. For example, the knowledge base $\{\mathsf{p} \rightarrow \mathsf{b}, \mathsf{b} \mid\!\sim \mathsf{f}\}$ is shorthand for $\{\neg(\mathsf{p} \rightarrow \mathsf{b}) \mid\!\sim \bot, \mathsf{b} \mid\!\sim \mathsf{f}\}$.

We denote the set of all ranked interpretations over $\mathcal{P}$ by RI, and we write $\textsc{Mod}(\mathcal{K})$ for the set of ranked models of a knowledge base $\mathcal{K}$. For any $U \subseteq$ RI, we write $U \Vdash \alpha$ to mean $\mathscr{R} \Vdash \alpha$ for all $\mathscr{R} \in U$. Finally, we write sat($\mathscr{R}$) for the set of formulas satisfied by the ranked interpretation $\mathscr{R}$.

Even though KLM extends propositional logic, it is still quite restrictive, as it only permits positive conditional assertions. Booth et al. (1998) consider an extension allowing for *negative* conditionals, i.e. assertions of the form $\alpha \not\mid\!\sim \beta$. Such an assertion is satisfied by a ranked interpretation $\mathscr{R}$ if and only if $\mathscr{R} \not\Vdash \alpha \mid\!\sim \beta$.

(REFL) $\mathcal{K} \approx \alpha \mid\sim \alpha$

(AND) $\dfrac{\mathcal{K} \approx \alpha \mid\sim \beta,\ \mathcal{K} \approx \alpha \mid\sim \gamma}{\mathcal{K} \approx \alpha \mid\sim \beta \wedge \gamma}$

(LLE) $\dfrac{\models \alpha \leftrightarrow \beta,\ \mathcal{K} \approx \alpha \mid\sim \gamma}{\mathcal{K} \approx \beta \mid\sim \gamma}$

(OR) $\dfrac{\mathcal{K} \approx \alpha \mid\sim \gamma,\ \mathcal{K} \approx \beta \mid\sim \gamma}{\mathcal{K} \approx \alpha \vee \beta \mid\sim \gamma}$

(RW) $\dfrac{\models \beta \rightarrow \gamma,\ \mathcal{K} \approx \alpha \mid\sim \beta}{\mathcal{K} \approx \alpha \mid\sim \gamma}$

(CM) $\dfrac{\mathcal{K} \approx \alpha \mid\sim \beta,\ \mathcal{K} \approx \alpha \mid\sim \gamma}{\mathcal{K} \approx \alpha \wedge \beta \mid\sim \gamma}$

(RM) $\dfrac{\mathcal{K} \not\approx \alpha \wedge \beta \mid\sim \gamma,\ \mathcal{K} \not\approx \alpha \mid\sim \neg\beta}{\mathcal{K} \not\approx \alpha \mid\sim \gamma}$

Figure 2: Rationality properties for defeasible entailment.

2.2 Rank Entailment

A central question in non-monotonic reasoning is determining what forms of entailment are appropriate in a defeasible setting. Given a knowledge base $\mathcal{K}$, we write $\mathcal{K} \approx \alpha \mid\sim \beta$ to mean that $\mathcal{K}$ *defeasibly entails* $\alpha \mid\sim \beta$. In the literature, there are a plethora of options available for the entailment relation $\approx$, each with their own strengths and weaknesses (Casini, Meyer, and Varzinczak 2019). As such, it is useful to understand defeasible entailment relations in terms of their global properties. An obviously desirable property is *Inclusion*:

(*Inclusion*) $\mathcal{K} \approx \alpha \mid\sim \beta$ for all $\alpha \mid\sim \beta \in \mathcal{K}$

Kraus et al. (1990) argue that a defeasible entailment relation should satisfy each of the properties given in figure 2, known as the *rationality properties*. We will call such relations *rational*.

Rational properties are essentially intertwined with the class of ranked interpretations.

Proposition 4 (Lehmann et al. (1992))**.** *A defeasible entailment relation $\approx$ is rational iff for each knowledge base $\mathcal{K}$, there is a ranked interpretation $\mathscr{R}_{\mathcal{K}}$ such that $\mathcal{K} \approx \alpha \mid\sim \beta$ iff $\mathscr{R}_{\mathcal{K}} \Vdash \alpha \mid\sim \beta$.*

The following natural form of entailment, called *rank entailment*, is *not* rational in general, as it fails to satisfy the property of rational monotonicity (RM):

Definition 2. *A conditional $\alpha \mid\sim \beta$ is* rank entailed *by a knowledge base $\mathcal{K}$ (written $\mathcal{K} \approx_R \alpha \mid\sim \beta$) iff $\mathscr{R} \Vdash \alpha \mid\sim \beta$ for every ranked model $\mathscr{R}$ of $\mathcal{K}$.*

Despite failing to be rational, rank entailment is important as it can be viewed as the *monotonic core* of an appropriate defeasible entailment relation. In other words, the following property is desirable:

(*KLM-Ampliativity*) $\mathcal{K} \approx \alpha \mid\sim \beta$ whenever $\mathcal{K} \approx_R \alpha \mid\sim \beta$

Note that a rational entailment relation satisfying Inclusion also satisfies KLM-Ampliativity by proposition 4.

2.3 Rational Closure

A well-known form of rational entailment for KLM is *rational closure*. Lehmann et al. (1992) propose rational closure as the minimum acceptable form of rational defeasible entailment, and give a syntactic characterisation of rational closure in terms of an ordering on KLM knowledge bases. Here we refer to the semantic approach (Giordano et al. 2015) and define rational closure in terms of an ordering on ranked interpretations:

Definition 3. *(Giordano et al. 2015, Definition 7) Given two ranked interpretations $\mathscr{R}_1$ and $\mathscr{R}_2$, we write $\mathscr{R}_1 <_G \mathscr{R}_2$, that is, $\mathscr{R}_1$ is preferred to $\mathscr{R}_2$, iff $\mathscr{R}_1(u) \leq_G \mathscr{R}_2(u)$ for every $u \in \mathcal{U}$, and there is some $v \in \mathcal{U}$ s.t. $\mathscr{R}_1(v) <_G \mathscr{R}_2(v)$.*

Consider the set of models of a KLM knowledge base $\mathcal{K}$. Intuitively, the lower a model is with respect to the ordering $\leq_G$, the fewer exceptional valuations it has modulo the constraints of $\mathcal{K}$. Thus the $\leq_G$-minimal models can be thought of as the semantic counterpart to the *principle of typicality* seen above. This idea of making valuations as typical as possible has first been presented by Booth et al. (1998) for the case of KLM knowledge bases with both positive and negative conditionals. For these knowledge bases, it turns out that there is always a unique minimal model:

Proposition 5. *Let $\mathcal{K} \subseteq \mathcal{L}^{\mid\sim}$ be a knowledge base. Then if $\mathcal{K}$ is consistent,* MOD$(\mathcal{K})$ *has a unique $\leq_G$-minimal element, denoted $\mathscr{R}^{RC}_{\mathcal{K}}$.*

The rational closure of a knowledge base can be characterised as the set of formulas satisfied by this minimal model:

Proposition 6. *(Giordano et al. 2015, Theorem 2) A conditional $\alpha \mid\sim \beta$ is in the rational closure of a knowledge base $\mathcal{K} \subseteq \mathcal{L}^{\mid\sim}$ (written $\mathcal{K} \approx_{RC} \alpha \mid\sim \beta$) iff $\mathscr{R}^{RC}_{\mathcal{K}} \Vdash \alpha \mid\sim \beta$.*

A well-known behaviour of rational closure is the so-called *drowning effect*. To make this concrete, consider the knowledge base $\mathcal{K} = \{\mathsf{b} \mid\sim \mathsf{f}, \mathsf{b} \mid\sim \mathsf{w}, \mathsf{r} \rightarrow \mathsf{b}, \mathsf{p} \rightarrow \mathsf{b}, \mathsf{p} \mid\sim \neg\mathsf{f}, \}$. This states that birds have wings and typically fly, that robins are birds, and that penguins are birds that typically don't fly. Intuitively one would expect to be able to conclude from this that robins typically have wings ($\mathsf{r} \mid\sim \mathsf{w}$), since robins are not exceptional birds. More generally, every subclass that does not show any exceptional behaviour should inherit all the typical properties of a class by default. This is the principle of the Presumption of Typicality mentioned earlier, to which rational closure obeys. But what happens with subclasses that are exceptional with respect to some property?

In the above example, since penguins are exceptional only with respect to their ability to fly, the question is whether penguins should inherit the other typical properties of birds, such as having wings ($\mathsf{p} \mid\sim \mathsf{w}$). Rational closure does *not* sanction this type of conclusion. That is, subclasses that are exceptional with respect to a typical property of a class do not inherit the other typical properties of the class. This is the drowning effect which, while being a desirable form of reasoning in some contexts, is considered a limitation if we are interested in modelling some form of Presumption of Independence (Lehmann 1995), in which a subclass inherits all the typical properties of a class, unless there is explicit information to the contrary. So, even though penguins are exceptional birds in the sense of typically not being able to fly, the Presumption of Independence requires us to conclude that penguins typically have wings.

There are several refinements of rational closure, such as lexicographic closure (Lehmann 1995), relevant closure (Casini et al. 2014) and inheritance-based closure (Casini and Straccia 2013), that satisfy both the Presumption of Typicality and the Presumption of Independence. Unlike rational closure, lexicographic closure formalises the *presumptive* reading of $\alpha \mid\sim \beta$, which states that "α implies β unless there is reason to believe otherwise" (Lehmann 1995; Casini, Meyer, and Varzinczak 2019).

2.4 Propositional Typicality Logic

The present paper investigates whether the notion of rational closure can be extended to more expressive logics. The first investigation in such a direction was proposed by Booth and Paris (1998), who consider an extension of KLM in which both positive ($\alpha \mid\sim \beta$) and negative ($\alpha \not\mid\sim \beta$) conditionals are allowed. This additional expressiveness introduces some technical issues, as not every such knowledge base has a model (consider $\mathcal{K} = \{\alpha \mid\sim \beta, \alpha \not\mid\sim \beta\}$, for instance). Nevertheless, Booth and Paris show that this is the only limit in the validity of Proposition 5: every consistent knowledge base in this extension has a rational closure.

Another investigated logic that extends KLM is Propositional Typicality Logic (PTL), a logic for defeasible reasoning proposed by Booth et al. (2015), in which propositional logic is enriched with a modal *typicality operator* (denoted $\bullet$). Formulas for PTL are defined by $\alpha ::= \top \mid \bot \mid p \mid \bullet\alpha \mid \neg\alpha \mid \alpha \wedge \alpha$, where p is any propositional atom. As before, other boolean connectives are defined in terms of $\neg, \wedge, \rightarrow, \leftrightarrow$. The intuition behind a formula $\bullet\alpha$ is that it is true for *typical* instances of α. Note that the typicality operator can be nested, so α may itself contain some $\bullet\beta$ as a subformula. The set of all PTL formulas is denoted $\mathcal{L}^\bullet$.

Satisfaction for PTL is defined with respect to a ranked interpretation $\mathscr{R}$. Given a valuation $u \in \mathcal{U}$ and formula $\alpha \in \mathcal{L}^\bullet$, we define $u \Vdash_{\mathscr{R}} \alpha$ inductively in the same way as propositional logic, with an additional rule for the typicality operator: $u \Vdash_{\mathscr{R}} \bullet\alpha$ if and only if $u \Vdash_{\mathscr{R}} \alpha$ and there is no $v \prec_{\mathscr{R}} u$ such that $v \Vdash_{\mathscr{R}} \alpha$. We then say that $\mathscr{R}$ satisfies the formula α, written $\mathscr{R} \Vdash \alpha$, iff $u \Vdash_{\mathscr{R}} \alpha$ for all $u \in \mathcal{U}^{\mathscr{R}}$. Given that the typicality operator can be nested and used anywhere within a PTL formula, one would intuitively expect PTL to be at least as expressive as KLM. The following lemma shows that this is indeed the case:

Proposition 7 (Booth et al. (2013))**.** *A ranked interpretation $\mathscr{R}$ satisfies the* KLM *formula $\alpha \mid\sim \beta$ if and only if it satisfies the* PTL *formula $\bullet\alpha \rightarrow \beta$.*

Given two knowledge bases $\mathcal{K}_1$ and $\mathcal{K}_2$, we say they are *equivalent* if they have exactly the same set of ranked models, i.e. if $\text{MOD}(\mathcal{K}_1) = \text{MOD}(\mathcal{K}_2)$. Proposition 7 can be rephrased as saying that every KLM knowledge base has an equivalent PTL knowledge base. Note that the converse doesn't hold; there are PTL knowledge bases with no equivalent in KLM:

Proposition 8 (Booth et al. (2013))**.** *For any $p \in \mathcal{P}$, the knowledge base $\mathcal{K} = \{\bullet p\}$ has no equivalent* KLM *knowledge base.*

The obvious form of entailment for a PTL knowledge base $\mathcal{K}$ is *rank entailment* (denoted $\approx_R$), presented earlier in definition 2. As noted before, rank entailment is monotonic and therefore inappropriate in many contexts. To pin down better forms of PTL entailment, Booth et al. (2015) consider the following properties, modelled after properties of rational closure, where $\approx_?$ is a PTL entailment relation and $\text{Cn}_?(\mathcal{K}) = \{\alpha \in \mathcal{L}^\bullet : \mathcal{K} \approx_? \alpha\}$ is its associated consequence operator:

(*Cumulativity*) For all $\mathcal{K}_1, \mathcal{K}_2 \subseteq \mathcal{L}^\bullet$, if $\mathcal{K}_1 \subseteq \mathcal{K}_2 \subseteq \text{Cn}_?(\mathcal{K}_1)$, then $\text{Cn}_?(\mathcal{K}_1) = \text{Cn}_?(\mathcal{K}_2)$.

(*Ampliativity*) For all $\mathcal{K} \subseteq \mathcal{L}^\bullet$, $\text{Cn}_R(\mathcal{K}) \subseteq \text{Cn}_?(\mathcal{K})$.

(*Strict Entailment*) For all $\mathcal{K} \subseteq \mathcal{L}^\bullet$ and $\alpha \in \mathcal{L}$, $\alpha \in \text{Cn}_?(\mathcal{K})$ iff $\alpha \in \text{Cn}_R(\mathcal{K})$.

(*Typical Entailment*) For all $\mathcal{K} \subseteq \mathcal{L}^\bullet$ and $\alpha \in \mathcal{L}$, $\bullet\top \rightarrow \alpha \in \text{Cn}_?(\mathcal{K})$ iff $\bullet\top \rightarrow \alpha \in \text{Cn}_R(\mathcal{K})$.

(*Single Model*) For all $\mathcal{K} \subseteq \mathcal{L}^\bullet$, there's some $\mathscr{R} \in \text{MOD}(\mathcal{K})$ such that for all $\alpha \in \mathcal{L}^\bullet$, $\alpha \in \text{Cn}_?(\mathcal{K})$ iff $\mathscr{R} \Vdash \alpha$.

Surprisingly, it turns out that an entailment relation cannot satisfy all of these properties simultaneously:

Proposition 9 (Booth et al. (2015))**.** *There is no* PTL *entailment relation $\approx_?$ satisfying Cumulativity, Ampliativity, Strict entailment, Typical entailment and the Single Model property.*

Booth et al. suggest that this is best interpreted as an argument for developing more than one form of PTL entailment, which can be compared to the divide between presumptive and prototypical readings for KLM entailment. An example of PTL entailment is *LM-entailment*, which is based on the following adaption of proposition 5:

Proposition 10 (Booth et al. (2019))**.** *Let $\mathcal{K} \subseteq \mathcal{L}^\bullet$ be a consistent knowledge base. Then* $\text{MOD}(\mathcal{K})$ *has a unique $\leq_G$-minimal element, denoted $\mathscr{R}^{LM}_{\mathcal{K}}$.*

Given a knowledge base $\mathcal{K} \subseteq \mathcal{L}^\bullet$, we define LM-entailment by writing $\mathcal{K} \approx_{LM} \alpha$ iff either $\mathcal{K}$ is inconsistent or $\mathscr{R}^{LM}_{\mathcal{K}} \Vdash \alpha$. Booth et al. prove that LM-entailment satisfies all of the above properties except for Strict Entailment, and hence in general there may be classical statements that are LM-entailed by $\mathcal{K}$ but not rank-entailed by it. Other forms of entailment, such as *PT-entailment*, can be shown to satisfy Strict Entailment but fail both Typical Entailment and the Single Model property.

3 Boolean KLM

In section 2.1, we noted that the logic KLM is quite restrictive, as it allows only for positive conditional assertions. As mentioned there, Booth and Paris (1998) consider an extension allowing for *negative* conditionals, i.e. assertions of the form $\alpha \not\mid\sim \beta$. Here we take that extension further, and propose *Boolean KLM* (BKLM), which allows for arbitrary boolean combinations of conditionals, but not for nested conditionals. BKLM formulas are defined by $A ::= \alpha \mid\sim \beta \mid \neg A \mid A \wedge A$, with other boolean connectives defined as usual. Following Booth and Paris, we will write $\alpha \not\mid\sim \beta$ as a synonym for $\neg(\alpha \mid\sim \beta)$ where convenient, and we denote the set of all BKLM formulas by $\mathcal{L}^b$. So, for example, $(\alpha \mid\sim \beta) \wedge (\gamma \not\mid\sim \delta)$ and $\neg((\alpha \not\mid\sim \beta) \vee (\gamma \mid\sim \delta))$ are BKLM formulas, but $\alpha \mid\sim (\beta \mid\sim \gamma)$ is not.

1	$\mathsf{p}\overline{\mathsf{q}}$
0	$\overline{\mathsf{p}}\mathsf{q}$

Figure 3: A ranked interpretation illustrating the difference between classical disjunction and BKLM disjunction.

Satisfaction for BKLM is defined in terms of ranked interpretations, by extending KLM satisfaction to boolean combinations of conditionals in the obvious fashion, namely $\mathscr{R} \Vdash \neg A$ iff $\mathscr{R} \nVdash A$ and $\mathscr{R} \Vdash A \wedge B$ iff $\mathscr{R} \Vdash A$ and $\mathscr{R} \Vdash B$. This leads to some subtle differences between BKLM satisfaction and the other logics. For instance, care must be taken to apply proposition 3 correctly when translating between propositional formulas and BKLM formulas. The propositional formula $\mathsf{p} \vee \mathsf{q}$ translates to the BKLM formula $\neg(\mathsf{p} \vee \mathsf{q}) \mid\!\sim \bot$, and *not* to the BKLM formula $(\neg\mathsf{p} \mid\!\sim \bot) \vee (\neg\mathsf{q} \mid\!\sim \bot)$, as the following example illustrates:

Example 2. *Consider the propositional formula $A = \mathsf{p} \vee \mathsf{q}$ and the* BKLM *formula $B = (\neg\mathsf{p} \mid\!\sim \bot) \vee (\neg\mathsf{q} \mid\!\sim \bot)$. If $\mathscr{R}$ is the ranked interpretation in figure 3, then $\mathscr{R}$ satisfies A but not B, as neither clause of the disjunction is satisfied.*

To prevent possible confusion, we will avoid mixing classical and defeasible assertions in a BKLM knowledge base. For similar reasons, it's also worth noting the difference between boolean connectives in PTL and the corresponding connectives in BKLM. By proposition 7, one might expect a BKLM formula such as $\neg(\mathsf{p} \mid\!\sim \mathsf{q})$ to translate into the PTL formula $\neg(\bullet\mathsf{p} \rightarrow \mathsf{q})$. The following example shows that this naïve approach fails:

Example 3. *Consider the formulas $A = \neg(\bullet\mathsf{p} \rightarrow \mathsf{q})$ and $B = \neg(\mathsf{p} \mid\!\sim \mathsf{q})$, and let $\mathscr{R}$ be the ranked interpretation in figure 3. Then A is equivalent to $\bullet\mathsf{p} \wedge \neg\mathsf{q}$, which isn't satisfied by $\mathscr{R}$. On the other hand, $\mathscr{R}$ satisfies B.*

One might ask whether there is a more nuanced way of translating BKLM knowledge bases into PTL. In the next section we answer this question in the negative, by showing that BKLM is in fact strictly more expressive than PTL.

3.1 Expressiveness of BKLM for Ranked Interpretations

So far we have been rather vague about what we mean by the *expressiveness* of a logic. All of the logics we consider in this paper share the same semantic structures, which provides us with a handy definition. We say that a logic can *characterise* a set of ranked interpretations $U \subseteq \mathrm{RI}$ if there is some knowledge base $\mathcal{K}$ with U as its set of ranked models. Given this, we say that a logic is *more expressive* than another logic if it can characterise at least as many sets of interpretations.

Example 4. *Let $\mathcal{K} \subseteq \mathcal{L}^{\mid\sim}$ be a* KLM *knowledge base. Then its* PTL *translation $\mathcal{K}' = \{\bullet\alpha \rightarrow \beta : \alpha \mid\!\sim \beta \in \mathcal{K}\}$ has exactly the same ranked models by proposition 7, and hence* PTL *is at least as expressive as* KLM. *Proposition 8 shows that this comparison is strict.*

In this section we show that BKLM is maximally expressive, in the sense that it can characterise *any* set of ranked interpretations. For a valuation $u \in \mathcal{U}$, we write $\hat{u}$ to mean any *characteristic formula* of u, namely any propositional formula such that $v \Vdash \hat{u}$ iff $v = u$. It is easy to see that these always exist, as $\mathcal{P}$ is finite, and that all characteristic formulas of u are logically equivalent.

Lemma 1. *For any ranked interpretation $\mathscr{R}$ and valuations $u, v \in \mathcal{U}$, it is straightforward to check that:*

1. *$\mathscr{R} \Vdash \top \not\mid\!\sim \neg\hat{u}$ iff $\mathscr{R}(u) = 0$.*
2. *$\mathscr{R} \Vdash \hat{u} \mid\!\sim \bot$ iff $\mathscr{R}(u) = \infty$.*
3. *$\mathscr{R} \Vdash \hat{u} \vee \hat{v} \mid\!\sim \neg\hat{v}$ iff $u \prec_{\mathscr{R}} v$ or $\mathscr{R}(u) = \mathscr{R}(v) = \infty$.*

Note that this lemma holds even in the vacuous case where $\mathscr{R}(u) = \infty$ for all $u \in \mathcal{U}$. Following Lehmann et al. (1992), we write $\alpha < \beta$ as shorthand for the defeasible implication $\alpha \vee \beta \mid\!\sim \neg\beta$. We now show that the concept of characteristic formulas can be applied to ranked interpretations as well:

Lemma 2. *Let $\mathscr{R}$ be any ranked interpretation. Then there exists a formula $ch(\mathscr{R}) \in \mathcal{L}^b$ with $\mathscr{R}$ as its unique model.*

Proof. Consider the following knowledge bases.

1. $\mathcal{K}_{\prec} = \{\hat{u} < \hat{v} : u \prec_{\mathscr{R}} v\} \cup \{\hat{u} \not< \hat{v} : u \not\prec_{\mathscr{R}} v\}$
2. $\mathcal{K}_{\infty} = \{\hat{u} \mid\!\sim \bot : \mathscr{R}(u) = \infty\} \cup \{\hat{u} \not\mid\!\sim \bot : \mathscr{R}(u) < \infty\}$

By lemma 1, $\mathscr{R}$ satisfies $\mathcal{K} = \mathcal{K}_{\prec} \cup \mathcal{K}_{\infty}$. To show that it is the unique model of $\mathcal{K}$, consider any $\mathscr{R}^* \in \mathrm{MOD}(\mathcal{K})$. Since $\mathscr{R}^*$ satisfies $\mathcal{K}_{\infty}$, $\mathscr{R}^*(u) = \infty$ iff $\mathscr{R}(u) = \infty$ for any $u \in \mathcal{U}$. Now consider any $u, v \in \mathcal{U}$, and suppose that $\mathscr{R}(u) < \infty$. Then $u \prec_{\mathscr{R}} v$ iff $\mathcal{K}_{\prec}$ contains $\hat{u} < \hat{v}$. But $\mathscr{R}^*$ satisfies $\mathcal{K}_{\prec}$, so this is true iff $u \prec_{\mathscr{R}^*} v$ as $\mathscr{R}^*(u) < \infty$. On the other hand, if $\mathscr{R}(u) = \infty$, then $u \not\prec_{\mathscr{R}} v$ and $u \not\prec_{\mathscr{R}^*} v$. Hence $\prec_{\mathscr{R}} = \prec_{\mathscr{R}^*}$, which implies that $\mathscr{R} = \mathscr{R}^*$ by proposition 2. We conclude the proof by letting $\mathrm{ch}(\mathscr{R}) = \bigwedge_{\alpha \in \mathcal{K}} \alpha$. □

We refer to $\mathrm{ch}(\mathscr{R})$ as the *characteristic formula* of $\mathscr{R}$. A simple application of disjunction allows us to prove the following more general corollary:

Corollary 1. *Let $U \subseteq$ RI be a set of ranked interpretations. Then there exists a formula $ch(U) \in \mathcal{L}^b$ with U as its set of models.*

This proves that BKLM is at least as expressive as PTL since, in principle, for every PTL knowledge base there is some BKLM knowledge base with the same set of models. It is not clear, however, whether there is a more natural description of this knowledge base than that provided by characteristic formulas. In the next section we will address this shortcoming by describing an explicit translation from PTL to BKLM knowledge bases.

In fact, BKLM is *strictly* more expressive than PTL. This is illustrated by the knowledge base $\mathcal{K} = \{(\top \mid\!\sim \mathsf{p}) \vee (\top \mid\!\sim \neg\mathsf{p})\}$, which expresses the "excluded-middle" statement that typically one of p or $\neg\mathsf{p}$ is true. There are two distinct $\leq_G$-minimal ranked models of $\mathcal{K}$, given by $\mathscr{R}_1$ and $\mathscr{R}_2$ in figure 4, and hence $\mathcal{K}$ cannot have an equivalent PTL knowledge base by proposition 10.

3.2 Translating PTL Into BKLM

In section 2.4, satisfaction for PTL formulas with respect to a ranked interpretation $\mathscr{R}$ was defined in terms of the possible valuations of $\mathscr{R}$. In order to define a translation operator between PTL and BKLM, our main idea is to *encode* satisfaction with respect to a valuation $u \in \mathcal{U}$ in terms of an appropriate BKLM formula. In other words, we will define an operator $\text{tr}_u : \mathcal{L}^\bullet \to \mathcal{L}^b$ such that for each $u \in \mathcal{U}^{\mathscr{R}}$, $\mathscr{R} \Vdash \text{tr}_u(\alpha)$ iff $u \Vdash_{\mathscr{R}} \alpha$.

Definition 4. *Given $\alpha, \beta \in \mathcal{L}^\bullet$, $p \in \mathcal{P}$ and $u \in \mathcal{U}$, we define tr_u by structural induction as follows:*

1. $tr_u(p) \stackrel{\text{def}}{=} \hat{u} \mid\sim p$
2. $tr_u(\top) \stackrel{\text{def}}{=} \hat{u} \mid\sim \top$
3. $tr_u(\bot) \stackrel{\text{def}}{=} \hat{u} \mid\sim \bot$
4. $tr_u(\neg\alpha) \stackrel{\text{def}}{=} \neg tr_u(\alpha)$
5. $tr_u(\alpha \wedge \beta) \stackrel{\text{def}}{=} tr_u(\alpha) \wedge tr_u(\beta)$
6. $tr_u(\bullet\alpha) \stackrel{\text{def}}{=} tr_u(\alpha) \wedge \bigwedge_{v \in \mathcal{U}} \left[(\hat{v} < \hat{u}) \to \neg tr_v(\alpha)\right]$

Note that this is well-defined, as each case is defined in terms of the translation of strict subformulas. The translations can be viewed as formal version of the definition of PTL satisfaction - case 6 states that $\bullet\alpha$ is satisfied by a possible valuation u iff u is a minimal valuation satisfying α, for instance.

Lemma 3. *Let $\mathscr{R}$ be a ranked interpretation, and $u \in \mathcal{U}^{\mathscr{R}}$ a valuation with $\mathscr{R}(u) < \infty$. Then for all $\alpha \in \mathcal{L}^\bullet$ we have $\mathscr{R} \Vdash tr_u(\alpha)$ if and only if $u \Vdash_{\mathscr{R}} \alpha$.*

Proof. We will prove the result by structural induction on the cases in definition 4:

1. Suppose that $\mathscr{R} \Vdash \text{tr}_u(p)$, i.e. $\mathscr{R} \Vdash \hat{u} \mid\sim p$. This is true iff $u \models p$, which is equivalent by definition to $u \Vdash_{\mathscr{R}} p$. Cases 2 and 3 are similar.
4. Suppose that $\mathscr{R} \Vdash \text{tr}_u(\neg\alpha)$, i.e. $\mathscr{R} \Vdash \neg\text{tr}_u(\alpha)$. This is true iff $\mathscr{R} \nVdash \text{tr}_u(\alpha)$, which by the induction hypothesis is equivalent to $u \nVdash_{\mathscr{R}} \alpha$. But this is equivalent to $u \Vdash_{\mathscr{R}} \neg\alpha$ by definition. Case 5 is similar.
6. Suppose there exists an $\alpha \in \mathcal{L}^\bullet$ such that $\mathscr{R} \Vdash \text{tr}_u(\bullet\alpha)$ but $u \nVdash_{\mathscr{R}} \bullet\alpha$. Then either $u \nVdash_{\mathscr{R}} \alpha$, which by the induction hypothesis is a contradiction since $\mathscr{R} \Vdash \text{tr}_u(\alpha)$, or there is some $v \in \mathcal{U}$ with $v \prec_{\mathscr{R}} u$ such that $v \Vdash_{\mathscr{R}} \alpha$. But by lemma 1, $v \prec_{\mathscr{R}} u$ is true only if $\mathscr{R} \Vdash \hat{v} < \hat{u}$. We also have, by the induction hypothesis, that $\mathscr{R} \Vdash \text{tr}_v(\alpha)$ since $v \Vdash_{\mathscr{R}} \alpha$. Hence $\mathscr{R} \Vdash (\hat{v} < \hat{u}) \wedge \text{tr}_v(\alpha)$, which implies that one of the clauses in $\text{tr}_u(\bullet\alpha)$ is false. This is a contradiction, so we conclude that $\mathscr{R} \Vdash \text{tr}_u(\bullet\alpha)$ implies $u \Vdash_{\mathscr{R}} \bullet\alpha$.
   Conversely, suppose that $u \Vdash_{\mathscr{R}} \bullet\alpha$. Then $u \Vdash_{\mathscr{R}} \alpha$, and hence $\mathscr{R} \Vdash \text{tr}_u(\alpha)$ by the induction hypothesis. We also have that if $v \prec_{\mathscr{R}} u$ then $v \nVdash_{\mathscr{R}} \alpha$, which is equivalent to $\mathscr{R} \Vdash \neg\text{tr}_v(\alpha)$ by the induction hypothesis. But by lemma 1, $v \prec_{\mathscr{R}} u$ iff $\mathscr{R} \Vdash \hat{v} < \hat{u}$. We conclude that $\mathscr{R} \Vdash (\hat{v} < \hat{u}) \to \neg\text{tr}_v(\alpha)$ for all $v \in \mathcal{U}$, and hence $\mathscr{R} \Vdash \text{tr}_u(\bullet\alpha)$.

□

A formula $\alpha \in \mathcal{L}^\bullet$ is satisfied by a ranked interpretation $\mathscr{R}$ iff it is satisfied by every possible valuation of $\mathscr{R}$. We can combine the translation operators of definition 4 to formalise this statement as follows:

Definition 5. $tr(\alpha) \stackrel{\text{def}}{=} \bigwedge_{u \in \mathcal{U}} \Big((\hat{u} \not\mid\sim \bot) \to tr_u(\alpha)\Big)$

All that remains is to check that this formula correctly encodes PTL satisfaction:

Lemma 4. *For all $\alpha \in \mathcal{L}^\bullet$, a ranked model $\mathscr{R}$ satisfies α if and only if it satisfies $tr(\alpha)$.*

Proof. Suppose $\mathscr{R} \Vdash \alpha$. Then for all $u \in \mathcal{U}$, either $\mathscr{R}(u) = \infty$ or $u \Vdash_{\mathscr{R}} \alpha$. The former implies $\mathscr{R} \Vdash \hat{u} \mid\sim \bot$ by lemma 1, and the latter implies $\mathscr{R} \Vdash \text{tr}_u(\alpha)$ by lemma 3. Thus $\mathscr{R} \Vdash (\hat{u} \not\mid\sim \bot) \to \text{tr}_u(\alpha)$ for all $u \in \mathcal{U}$, which proves $\mathscr{R} \Vdash \text{tr}(\alpha)$ as required.

Conversely, suppose $\mathscr{R} \Vdash \text{tr}(\alpha)$. Then for any $u \in \mathcal{U}$, either $\mathscr{R} \Vdash \hat{u} \mid\sim \bot$ and hence $\mathscr{R}(u) = \infty$ by lemma 1, or $\mathscr{R} \Vdash \hat{u} \not\mid\sim \bot$ and hence $\mathscr{R} \Vdash \text{tr}_u(\alpha)$ by hypothesis. But then $\mathscr{R} \Vdash \alpha$ by lemma 3. □

4 Entailment Results for BKLM

We now turn to the question of defeasible entailment for BKLM knowledge bases. As in previous cases, an obvious approach to this is *rank entailment*, which we define in the usual fashion:

Definition 6. *Given any $\mathcal{K} \subseteq \mathcal{L}^b$ and $A \in \mathcal{L}^b$, we say $\mathcal{K}$* rank entails *A (written $\mathcal{K} \approx_R A$) iff $\mathscr{R} \Vdash A$ for all $\mathscr{R} \in$* MOD$(\mathcal{K})$.

Being monotonic, rank entailment serves as a useful lower bound for defeasible BKLM entailment, but cannot be considered a good solution in its own right. Letting $\approx_?$ be an entailment relation and $\text{Cn}_?$ its associated consequence operator, consider the entailment properties in section 2.4 in the context of BKLM. Our first observation is that the premises of proposition 9 can be weakened as a consequence of global disjunction:

Lemma 5. *There is no* BKLM *entailment relation $\approx_?$ satisfying Ampliativity, Typical Entailment and the Single Model property.*

Proof. Suppose that $\approx_?$ is such an entailment relation, and consider the knowledge base $\mathcal{K} = \{(\top \mid\sim \mathsf{p}) \vee (\top \mid\sim \neg\mathsf{p})\}$. Both interpretations in figure 4, $\mathscr{R}_1$ and $\mathscr{R}_2$, are models of $\mathcal{K}$. $\mathscr{R}_1$ satisfies $\top \mid\sim \mathsf{p}$ and not $\top \mid\sim \neg\mathsf{p}$, whereas $\mathscr{R}_2$ satisfies $\top \mid\sim \neg\mathsf{p}$ and not $\top \mid\sim \mathsf{p}$. Thus, by the Typical Entailment property, $\mathcal{K} \not\approx_? \top \mid\sim \mathsf{p}$ and $\mathcal{K} \not\approx_? \top \mid\sim \neg\mathsf{p}$. On the other hand, by Ampliativity we get $\mathcal{K} \approx_? (\top \mid\sim \mathsf{p}) \vee (\top \mid\sim \neg\mathsf{p})$. A single ranked interpretation cannot satisfy all three of these assertions, however, and hence no such entailment relation can exist. □

In the PTL context, LM-entailment satisfies Ampliativity, Typical Entailment and the Single Model property. Thus lemma 5 is a concrete sense in which BKLM entailment is more constrained than PTL entailment. This raises an interesting question - can we nevertheless define a notion of entailment for BKLM, in the same spirit as rational closure

1	$\overline{\mathsf{p}}$
0	p

(a) $\mathscr{R}_1$

1	p
0	$\overline{\mathsf{p}}$

(b) $\mathscr{R}_2$

Figure 4: Ranked models of $\mathcal{K} = \{(\top \mid\!\sim \mathsf{p}) \vee (\top \mid\!\sim \neg\mathsf{p})\}$.

and LM-entailment, by giving up one of the above properties? In order to guarantee a rational entailment relation, it is desirable to keep the Single Model property in view of proposition 4. For the rest of this section we will investigate the consequences of this choice, and show that while it is possible to satisfy the Single Model property for BKLM entailment, the resulting entailment relations are heavily restricted.

4.1 Order Entailments

As we have seen in Section 2.3, rational closure can be modeled as a form of *minimal model entailment*. In other words, given a knowledge base $\mathcal{K}$, we can construct the rational closure of $\mathcal{K}$ by placing an appropriate ordering on its set of ranked models (in this case $\leq_G$), and picking out the consequences of the minimal ones. In this section we formalise this notion of entailment, with a view towards understanding the Single Model property for BKLM.

Definition 7. *Let $<$ be a strict partial order on* RI. *Then for all knowledge bases $\mathcal{K}$ and formulas α, we define $\mathcal{K} \mid\!\approx_< \alpha$ iff $\mathscr{R} \Vdash \alpha$ for all $<$-minimal elements of* $\text{Mod}(\mathcal{K})$.

The relation $\mid\!\approx_<$ will be referred to as the *order entailment relation* of $<$. Note that we have been deliberately vague about which logic we are dealing with, as the construction works identically for KLM, PTL and BKLM. It is also worth mentioning that the set of models of a consistent knowledge base always has $<$-minimal elements, as we have assumed finiteness of $\mathcal{P}$, that implies a finite set of ranked interpretations.

Example 5. *By definition 6, the rational closure of any* KLM *knowledge base $\mathcal{K}$ is the set of formulas satisfied by the (unique) $<_G$-minimal element of* $\text{Mod}(\mathcal{K})$. *Thus rational closure is the order entailment relation of $<_G$ over* KLM.

In general, order entailment relations satisfy all of the rationality properties in figure 2 except for rational monotonicity (RM). Rational monotonicity holds, for instance, if $\text{Mod}(\mathcal{K})$ has a unique $<$-minimal model for every knowledge base $\mathcal{K}$. This is the case for rational closure and LM-entailment, which both satisfy the Single Model property. The following proposition follows easily from the definitions, and shows that this is typical:

Proposition 11. *An order entailment relation $\mid\!\approx_<$ satisfies the Single Model property iff* $\text{Mod}(\mathcal{K})$ *has a unique $<$-minimal model for every knowledge base $\mathcal{K}$.*

A class of order entailment relations for which the Single Model property always holds are the *total* order entailment relations, i.e. those $\mid\!\approx_<$ corresponding to a total order $<$. Intuitively, this is a strong restriction, as an a priori total ordering over all possible ranked interpretations is unnatural in the context of an agent's knowledge. For BKLM entailment, it turns out that there is a partial converse to this discussion, which we will prove in the next section.

4.2 The Single Model Property

In this section we prove that, under some mild assumptions, a BKLM entailment relation satisfying the Single Model property is always equivalent to a total order entailment relation.

Theorem 1. *Suppose $\mid\!\approx_?$ is a* BKLM *entailment relation satisfying Cumulativity, Ampliativity and the Single Model property. Then $\mid\!\approx_? = \mid\!\approx_<$, where $\mid\!\approx_<$ is a total order entailment relation.*

For the remainder of the section, consider a fixed BKLM entailment relation $\mid\!\approx_?$ (with associated consequence operator $\text{Cn}_?$), and suppose that $\mid\!\approx_?$ satisfies Cumulativity, Ampliativity and the Single Model property. In what follows, we will move between the entailment relation and consequence operator notations freely as convenient. To begin with, we note the following straightforward lemma:

Lemma 6. *For any knowledge base $\mathcal{K} \subseteq \mathcal{L}^b$, $Cn_?(\mathcal{K}) = Cn_R(Cn_?(\mathcal{K})) = Cn_?(Cn_R(\mathcal{K}))$.*

Our approach to proving theorem 1 is to assign a unique index $\text{ind}(\mathscr{R}) \in \mathbb{N}$ to each ranked interpretation $\mathscr{R} \in \text{RI}$, and then show that $\text{Cn}_?(\mathcal{K})$ corresponds to minimisation of index in $\text{Mod}(\mathcal{K})$. To construct this indexing scheme, consider the following algorithm:

1. Set $M_0 := \text{RI}$, $i := 0$.
2. If $M_i = \emptyset$, terminate.
3. By corollary 1, there is some knowledge base $\mathcal{K}_i \subseteq \mathcal{L}^b$ such that $\text{Mod}(\mathcal{K}_i) = M_i$.
4. By the single model property, there is some $\mathscr{R}_i \in M_i$ such that $\text{Cn}_?(\mathcal{K}_i) = \text{sat}(\mathscr{R}_i)$.
5. Set $M_{i+1} := M_i \setminus \{\mathscr{R}_i\}$, $i := i + 1$.
6. Go to step 2, and iterate until termination.

This algorithm is guaranteed to terminate as M_0 is finite and $0 \leq |M_{i+1}| < |M_i|$. Note that once the algorithm terminates, for each $\mathscr{R} \in \text{RI}$ there will have been a unique $i \in \mathbb{N}$ such that $\mathscr{R} = \mathscr{R}_i$. We will call this i the *index* of $\mathscr{R}$, and denote it by $\text{ind}(\mathscr{R})$. Given a knowledge base $\mathcal{K}$, we define $\text{ind}(\mathcal{K}) = \min\{\text{ind}(\mathscr{R}) : \mathscr{R} \in \text{Mod}(\mathcal{K})\}$ to be the *minimum index* of the knowledge base.

For clarity, when we write $\mathscr{R}_n$, $\mathcal{K}_n$ and M_n in the following lemmas, we mean the ranked interpretations, knowledge bases and sets of models constructed in steps 3 to 5 of the algorithm when $i = n$:

Lemma 7. *Given any knowledge base $\mathcal{K} \subseteq \mathcal{L}^b$,* $\text{Mod}(\mathcal{K}) \subseteq$ *M_n, where $n = ind(\mathcal{K})$.*

Proof. An easy induction on step 5 of the algorithm proves that $M_n = \{\mathscr{R} \in \text{RI} : \text{ind}(\mathscr{R}) \geq n\}$. By hypothesis, $\text{ind}(\mathscr{R}) \geq n$ for all $\mathscr{R} \in \text{Mod}(\mathcal{K})$, and hence $\text{Mod}(\mathcal{K}) \subseteq M_n$. □

The following lemma proves that entailment under $\mid\!\approx_?$ corresponds to minimisation of index:

Lemma 8. *Given any knowledge base $\mathcal{K} \subseteq \mathcal{L}^b$, $Cn_?(\mathcal{K}) = sat(\mathscr{R}_n)$, where $n = ind(\mathcal{K})$.*

Proof. For all A, $\mathcal{K}_n \approx_R A$ iff $\mathscr{R} \Vdash A$ for all $\mathscr{R} \in$ MOD$(\mathcal{K}_n) = M_n$. But by lemma 7, MOD$(\mathcal{K}) \subseteq M_n$ and hence $\text{Cn}_R(\mathcal{K}_n) \subseteq \text{Cn}_R(\mathcal{K})$. On the other hand, $\mathscr{R}_n \in$ MOD$(\mathcal{K})$ by hypothesis and hence $\mathscr{R}_n \Vdash A$ for all $A \in \mathcal{K}$. By the definition of step 4 of the algorithm we have $\text{sat}(\mathscr{R}_n) = \text{Cn}_?(\mathcal{K}_n)$, and thus $\mathcal{K} \subseteq \text{Cn}_?(\mathcal{K}_n)$. Applying Cn_R to each side of this inclusion (using the monotonicity of rank entailment), we get $\text{Cn}_R(\mathcal{K}) \subseteq \text{Cn}_R(\text{Cn}_?(\mathcal{K}_n)) = \text{Cn}_?(\mathcal{K}_n)$, with the last equality following from lemma 6. Putting it all together, we have $\text{Cn}_R(\mathcal{K}_n) \subseteq \text{Cn}_R(\mathcal{K}) \subseteq \text{Cn}_?(\mathcal{K}_n)$, and hence by Cumulativity we conclude $\text{Cn}_?(\mathcal{K}) = \text{Cn}_?(\mathcal{K}_n) = \text{sat}(\mathscr{R}_n)$. □

Consider the strict partial order on RI defined by $\mathscr{R}_1 < \mathscr{R}_2$ iff $\text{ind}(\mathscr{R}_1) < \text{ind}(\mathscr{R}_2)$. By construction, the index of a ranked interpretation is unique, and hence $<$ is total. It follows from lemma 8 that $\approx_? = \approx_<$, and hence $\approx_?$ is equivalent to a total order entailment relation. This completes the proof of theorem 1.

5 Related Work

The most relevant work w.r.t. the present paper is that of Booth and Paris (1998) in which they define rational closure for the extended version of KLM for which negated conditionals are allowed, and the work on PTL (Booth et al. 2015; Booth et al. 2019). The relation this work has with BKLM was investigated in detail throughout the paper.

Delgrande (1987) proposes a logic that is as expressive as BKLM. The entailment relation he proposes is different from the minimal entailment relations we consider here and, given the strong links between our constructions and the KLM approach, the remarks in the comparison made by Lehmann and Magidor (1992, Section 3.7) are also applicable here.

Boutilier (1994) defines a family of conditional logics using preferential and ranked interpretations. His logic is closer to ours and even more expressive, since nesting of conditionals is allowed, but he too does not consider minimal constructions. That is, both Delgrande and Boutilier's approaches adopt a Tarskian-style notion of consequence, in line with rank entailment. The move towards a non-monotonic notion of defeasible entailment was precisely our motivation in the present work.

Giordano et al. (2010) propose the system P_{min} which is based on a language that is as expressive as PTL. 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*). Also, the system P_{min} relies on preferential models and a notion of minimality that is closer to circumscription (McCarthy 1980).

In the context of description logics, Giordano et al. (2007; 2015) propose to extend the conditional language with an explicit typicality operator $T(\cdot)$, with a meaning that is closely related to the PTL operator $\bullet$. 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 $T(\cdot)$ that does not go beyond the expressivity of a KLM-style conditional language, but revised, of course, for the expressivity of description logics.

In the context of adaptive logics, Straßer (2014) defines the logic R^+ as an extension of KLM in which arbitrary boolean combinations of defeasible implications are allowed, and the set of propositional atoms has been extended to include the symbols $\{l_i : i \in \mathbb{N}\}$. Semantically, these symbols encode rank in the object language, in the sense that $u \Vdash l_i$ in a ranked interpretation $\mathscr{R}$ iff $\mathscr{R}(u) \geq i$. Straßer's interest in R^+ is to define an adaptive logic ALC^S that provides a dynamic proof theory for rational closure, whereas our interest in BKLM is to generalise rational closure to more expressive extensions of KLM. Nevertheless, the Minimal Abnormality Strategy (see the work of Batens (2007), for instance) for ALC^S is closely related to LM-entailment as defined in this paper.

6 Conclusion

The main focus of this paper is exploring the connection between expressiveness and entailment for extensions of the core logic KLM. Accordingly, we introduce the logic BKLM, an extension of KLM that allows for arbitrary boolean combinations of defeasible implications. We take an abstract approach to the analysis of BKLM, and show that it is strictly more expressive than existing extensions of KLM such as PTL (Booth, Meyer, and Varzinczak 2013) and KLM with negation (Booth and Paris 1998). Our primary conclusion is that a logic as expressive as BKLM has to give up several desirable properties for defeasible entailment, most notably the Single Model property, and thus appealing forms of entailment for PTL such as LM-entailment (Booth et al. 2015) cannot be lifted to the BKLM case.

For future work, an obvious question is what forms of defeasible entailment *are* appropriate for BKLM. For instance, is it possible to skirt the impossibility results proven in this paper while still retaining the KLM rationality properties? Other forms of entailment for PTL, such as PT-entailment, have also yet to be analysed in the context of BKLM and may be better suited to such an expressive logic.

Another line of research to be explored is whether there is a more natural translation of PTL formulas into BKLM than that defined in this paper. Our translation is based on a direct encoding of PTL semantics, and consequently results in an exponential blow-up in the size of the formulas being translated. It is clear that there are much more efficient ways to translate *specific* PTL formulas, but we leave it as an open problem whether this can be done in general. In a similar vein, it is interesting to ask how PTL could be extended in order to make it equiexpressive with BKLM.

Finally, it may be interesting to compare BKLM with an extension of KLM that allows for nested defeasible implications, i.e. formulas such as $\alpha \mid\sim (\beta \mid\sim \gamma)$. While such an extension cannot be more expressive than BKLM, at least for a semantics given by ranked interpretations, it may provide more natural encodings of various kinds of typicality, and thus be easier to work with from a pragmatic point of view.

References

Batens, D. 2007. A universal logic approach to adaptive logics. *Logica Universalis* 1:221–242.

Booth, R., and Paris, J. 1998. A note on the rational closure of knowledge bases with both positive and negative knowledge. *Journal of Logic, Language and Information* 7(2):165–190.

Booth, R.; Casini, G.; Meyer, T.; and Varzinczak, I. 2015. On the entailment problem for a logic of typicality. In *IJCAI 2015*, 2805–2811.

Booth, R.; Casini, G.; Meyer, T.; and Varzinczak, I. 2019. On rational entailment for propositional typicality logic. *Artificial Intelligence* 277.

Booth, R.; Meyer, T.; and Varzinczak, I. 2013. A propositional typicality logic for extending rational consequence. In Fermé, E.; Gabbay, D.; and Simari, G., eds., *Trends in Belief Revision and Argumentation Dynamics*, volume 48 of *Studies in Logic – Logic and Cognitive Systems*. King's College Publications. 123–154.

Boutilier, C. 1994. Conditional logics of normality: A modal approach. *Artificial Intelligence* 68(1):87–154.

Casini, G., and Straccia, U. 2013. Defeasible inheritance-based description logics. *JAIR* 48:415–473.

Casini, G.; Meyer, T.; Moodley, K.; and Nortje, R. 2014. Relevant closure: A new form of defeasible reasoning for description logics. In *JELIA 2014*, 92–106.

Casini, G.; Meyer, T.; and Varzinczak, I. 2019. Taking defeasible entailment beyond rational closure. In Calimeri, F.; Leone, N.; and Manna, M., eds., *Logics in Artificial Intelligence - 16th European Conference, JELIA 2019, Rende, Italy, May 7-11, 2019, Proceedings*, volume 11468 of *Lecture Notes in Computer Science*, 182–197. Springer.

Delgrande, J. 1987. A first-order logic for prototypical properties. *Artificial Intelligence* 33:105–130.

Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2007. Preferential description logics. In Dershowitz, N., and Voronkov, A., eds., *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, number 4790 in LNAI, 257–272. Springer.

Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. L. 2010. A nonmonotonic extension of KLM preferential logic P. In *Logic for Programming, Artificial Intelligence, and Reasoning - 17th International Conference, LPAR-17, Yogyakarta, Indonesia, October 10-15, 2010. Proceedings*, 317–332.

Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2015. Semantic characterization of rational closure: From propositional logic to description logics. *Art. Int.* 226:1–33.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167–207.

Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Art. Int.* 55:1–60.

Lehmann, D. 1995. Another perspective on default reasoning. *Annals of Math. and Art. Int.* 15(1):61–82.

McCarthy, J. 1980. Circumscription, a form of nonmonotonic reasoning. *Art. Int.* 13(1-2):27–39.

Straßer, C. 2014. An adaptive logic for rational closure. In *Adaptive Logics For Defeasible Reasoning*, volume 38 of *Trends in Logic*. Springer International Publishing. 181–206.