

Taking Defeasible Entailment beyond Rational Closure

Giovanni Casini¹[0000–0002–4267–4447], Thomas Meyer²[0000–0003–2204–6969], and
Ivan Varzinczak³[0000–0002–0025–9632]

¹ CSC, Université du Luxembourg, Luxembourg.

Email: giovanni.casini@uni.lu

² CAIR & University of Cape Town, South Africa.

Email: tmeyer@cs.uct.ac.za

³ CRIL, Univ. Artois & CNRS, France.

Email: varzinczak@cril.fr

Abstract. We present a systematic approach for extending the KLM framework for defeasible entailment. We first present a class of basic defeasible entailment relations, characterise it in three distinct ways and provide a high-level algorithm for computing it. This framework is then refined, with the refined version being characterised in a similar manner. We show that the two well-known forms of defeasible entailment, *rational closure* and *lexicographic closure*, fall within our refined framework, that rational closure is the most conservative of the defeasible entailment relations within the framework (with respect to subset inclusion), but that there are forms of defeasible entailment within our framework that are more “adventurous” than lexicographic closure.

1 Introduction

The approach by Kraus, Lehmann and Magidor [23] (a.k.a. KLM) is a well-known framework for defeasible reasoning. The KLM properties can be viewed as constraints on appropriate forms of defeasible entailment. We present what we believe to be the first systematic approach for extending the KLM framework. Our first proposal, *basic defeasible entailment*, strengthens the KLM framework by adding additional properties to it. We provide both a semantic characterisation in terms of a class of *ranked interpretations*, and a characterisation in terms of a class of functions that rank propositional (and defeasible) statements in a knowledge base according to their level of typicality. We also provide an algorithm for computing the framework. Next, we identify a crucial shortcoming in basic defeasible entailment, and propose a further strengthening, *rational defeasible entailment*, via an additional property. We prove that rational defeasible entailment can be characterised both semantically and in terms of ranks, and show that the algorithm is also applicable for computing rational defeasible entailment.

Currently there are two well-known forms of defeasible entailment satisfying those properties: rational closure (RC) [25] and lexicographic closure (LC) [24]. We show that both are rational (and basic) defeasible entailment relations, that RC is the most conservative form of rational defeasible entailment, but there are forms of rational defeasible entailment that are “bolder” than LC. We argue that the framework for rational defeasible entailment is reminiscent of the AGM framework for belief change [1].

In the next section we provide the relevant background material, after which we present our work on basic defeasible entailment, rational defeasible entailment, and a discussion on the relation between lexicographic closure and rational defeasible entailment. We conclude with a discussion of related work, a summary, and suggestions for future work.

2 Background

For \mathcal{P} being a finite set of propositional *atoms*, we use p, q, \dots as meta-variables for atoms. Propositional sentences are denoted by α, β, \dots , and are recursively defined in the usual way: $\alpha ::= \top \mid \perp \mid p \mid \neg\alpha \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \alpha \rightarrow \alpha \mid \alpha \leftrightarrow \alpha$. With \mathcal{L} we denote the set of all propositional sentences. With $\mathcal{U} \equiv_{\text{def}} \{0, 1\}^{\mathcal{P}}$ we denote the set of all propositional *valuations*, with 1 representing truth and 0 representing falsity. We use $u, v \dots$ to denote valuations. Sometimes we represent valuations as sequences of atoms (e.g., p) and barred atoms (e.g., \bar{p}), with the understanding that the presence of a non-bared atom indicates that the atom is true in the valuation, while the presence of a barred atom indicates that the atom is false in the valuation. 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 $v \Vdash \alpha$. The *models* of a set of sentences X is: $\llbracket X \rrbracket \equiv_{\text{def}} \{v \in \mathcal{U} \mid v \Vdash \alpha \text{ for every } \alpha \in X\}$.

2.1 KLM-style defeasible implication

In the logic proposed by Kraus et al. [23], often referred to as the *KLM approach*, we are interested in defeasible implications (or DIs) of the form $\alpha \vdash \beta$, read as “typically, if α , then β ”. The semantics of KLM-style rational defeasible implications is given by structures referred to as *ranked interpretations* [25]. In this work we adopt the following alternative representation thereof:

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

Given \mathcal{R} , we call $\mathcal{R}(v)$ the *rank* of v w.r.t. \mathcal{R} . Valuations with a lower rank are deemed more normal (or typical) than those with a higher rank, while those with an infinite rank are regarded as so atypical as to be impossible. With $\mathcal{U}^{\mathcal{R}} \equiv_{\text{def}} \{v \in \mathcal{U} \mid \mathcal{R}(v) < \infty\}$ we denote the *possible* valuations in \mathcal{R} . Given $\alpha \in \mathcal{L}$, we let $\llbracket \alpha \rrbracket^{\mathcal{R}} \equiv_{\text{def}} \{v \in \mathcal{U}^{\mathcal{R}} \mid v \Vdash \alpha\}$. \mathcal{R} satisfies (is a ranked model of) α (denoted $\mathcal{R} \Vdash \alpha$) if $\mathcal{U}^{\mathcal{R}} \subseteq \llbracket \alpha \rrbracket^{\mathcal{R}}$.

Note that \mathcal{R} generates a total preorder $\preceq_{\mathcal{R}}$ on \mathcal{U} as follows: $v \preceq_{\mathcal{R}} u$ iff $\mathcal{R}(v) \leq \mathcal{R}(u)$. Given any total preorder \preceq on $V \subseteq \mathcal{U}$, we can use its strict version \prec to generate a ranked interpretation as follows. Let the *height* $h(v)$ of $v \in V$ be the length of the \prec -path between any one of the \prec -minimal elements of V and v (the length of the \prec -path between any of the \prec -minimal elements and a \prec -minimal element is 0). For $V \subseteq \mathcal{U}$ and a total preorder \preceq on V , the ranked interpretation \mathcal{R}^{\preceq} generated from \preceq is defined as follows: for every $v \in \mathcal{U}$, $\mathcal{R}^{\preceq}(v) = h(v)$ if $v \in V$, and $\mathcal{R}^{\preceq}(v) = \infty$ otherwise.

Given a ranked interpretation \mathcal{R} and $\alpha, \beta \in \mathcal{L}$, \mathcal{R} satisfies (is a ranked model of) the conditional $\alpha \sim \beta$ (denoted $\mathcal{R} \Vdash \alpha \sim \beta$) if all the *possible* \prec -minimal α -valuations also satisfy β , i.e., if $\min_{\prec} \llbracket \alpha \rrbracket^{\mathcal{R}} \subseteq \llbracket \beta \rrbracket^{\mathcal{R}}$. \mathcal{R} satisfies a set of conditionals \mathcal{K} if $\mathcal{R} \Vdash \alpha \sim \beta$ for every $\alpha \sim \beta \in \mathcal{K}$.

Figure 1 depicts an example of a ranked interpretation for $\mathcal{P} = \{b, f, p\}$ satisfying $\mathcal{K} = \{p \rightarrow b, b \sim f, p \sim \neg f\}$. For brevity, we omit the valuations with rank ∞ in our graphical representations of ranked interpretations.

2	pbf
1	$\bar{p}b\bar{f}$ $p\bar{b}\bar{f}$
0	$\bar{p}\bar{b}\bar{f}$ $\bar{p}\bar{b}f$ $\bar{p}b\bar{f}$

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

Observe that *all* classical propositional sentences can be expressed as DIs: $\mathcal{R} \Vdash \alpha$ iff $\mathcal{R} \Vdash \neg\alpha \sim \perp$. The logic of defeasible implications can therefore be viewed as an extension of propositional logic.

2.2 Defeasible Entailment

Let a *knowledge base* \mathcal{K} be a finite set of defeasible implications. The main question in this paper is to analyse *defeasible entailment* (denoted by \approx): what it means for a defeasible implication to be entailed by a *fixed* knowledge \mathcal{K} . It is well-accepted that defeasible entailment (unlike classical entailment) is not unique. For example, Lehmann and Magidor [25] put forward *rational closure* as an appropriate form of defeasible entailment, while Lehmann [24] proposed *lexicographic closure* as an alternative. We consider both of these in more detail below. In studying different forms of defeasible entailment, the position advocated by Lehmann and Magidor [25], and one we adopt here as well, is to consider a number of *rationality properties*, referred to as the KLM properties, for defeasible entailment.

$$\begin{array}{l}
 \text{(Ref)} \quad \mathcal{K} \approx \alpha \sim \alpha \qquad \text{(LLE)} \quad \frac{\alpha \equiv \beta, \mathcal{K} \approx \alpha \sim \gamma}{\mathcal{K} \approx \beta \sim \gamma} \qquad \text{(RW)} \quad \frac{\mathcal{K} \approx \alpha \sim \beta, \beta \models \gamma}{\mathcal{K} \approx \alpha \sim \gamma} \\
 \text{(And)} \quad \frac{\mathcal{K} \approx \alpha \sim \beta, \mathcal{K} \approx \alpha \sim \gamma}{\mathcal{K} \approx \alpha \sim \beta \wedge \gamma} \qquad \text{(Or)} \quad \frac{\mathcal{K} \approx \alpha \sim \gamma, \mathcal{K} \approx \beta \sim \gamma}{\mathcal{K} \approx \alpha \vee \beta \sim \gamma} \qquad \text{(CM)} \quad \frac{\mathcal{K} \approx \alpha \sim \beta, \mathcal{K} \approx \alpha \sim \gamma}{\mathcal{K} \approx \alpha \wedge \beta \sim \gamma} \\
 \text{(RM)} \quad \frac{\mathcal{K} \approx \alpha \sim \gamma, \mathcal{K} \not\approx \alpha \sim \neg\beta}{\mathcal{K} \approx \alpha \wedge \beta \sim \gamma}
 \end{array}$$

Lehmann and Magidor argue that defeasible entailment ought to satisfy all the above KLM properties. We refer to this as *LM-rationality*.

Definition 2. A ranked interpretation \mathcal{R} is said to generate a defeasible \mathcal{K} -entailment relation $\approx_{\mathcal{R}}$ by setting $\mathcal{K} \approx_{\mathcal{R}} \alpha \sim \beta$ iff $\mathcal{R} \Vdash \alpha \sim \beta$. (If there isn't any ambiguity, we drop the subscript \mathcal{R} .)

Lehmann and Magidor proved the following useful result.

Observation 1 (Lehman & Magidor [25]) A defeasible entailment relation is LM-rational iff it can be generated from a ranked interpretation.

It is easy to see that rank entailment, defined next, is not LM-rational [25, Sect. 4.2].

Definition 3. A defeasible implication $\alpha \sim \beta$ is rank entailed by a knowledge base \mathcal{K} (denoted as $\mathcal{K} \approx_R \alpha \sim \beta$) if every ranked model of \mathcal{K} satisfies $\alpha \sim \beta$.

But rank entailment plays an important part in defining acceptable versions of defeasible entailment, since it can be viewed as the *monotonic core* of any appropriate form of defeasible entailment [16].

2.3 Rational Closure

The first version of defeasible entailment satisfying LM-rationality we consider is *rational closure* [25]. Consider the ordering $\preceq_{\mathcal{K}}$ on all ranked models of a knowledge base \mathcal{K} , which is defined as follows: $\mathcal{R}_1 \preceq_{\mathcal{K}} \mathcal{R}_2$ if for every $v \in \mathcal{U}$, $\mathcal{R}_1(v) \leq \mathcal{R}_2(v)$. Intuitively, ranked models lower down in the ordering are more typical. Giordano et al. [21] showed that there is a unique $\preceq_{\mathcal{K}}$ -minimal element.

Definition 4. Let $\mathcal{R}_{\mathcal{K}}^{RC}$ be the minimum element of the ordering $\preceq_{\mathcal{K}}$ on ranked models of \mathcal{K} . A defeasible implication $\alpha \sim \beta$ is in the rational closure of \mathcal{K} (denoted as $\mathcal{K} \approx_{RC} \alpha \sim \beta$) if $\mathcal{R}_{\mathcal{K}}^{RC} \Vdash \alpha \sim \beta$.

Observe that there are two levels of typicality at work for rational closure, namely *within* ranked models of \mathcal{K} , where valuations lower down are viewed as more typical, and *between* ranked models of \mathcal{K} , where ranked models lower down in the ordering are viewed as more typical. The most typical ranked model $\mathcal{R}_{\mathcal{K}}^{RC}$ is the one in which valuations are as typical as \mathcal{K} allows them to be.

Since rational closure can be defined in terms of a single ranked interpretation, it follows from Observation 1 that it is LM-rational (it satisfies all the KLM properties).

It will be useful to be able to refer to the possible valuations w.r.t. a knowledge base. We refer to $\mathcal{U}_R^{\mathcal{K}} \equiv_{\text{def}} \mathcal{U} \setminus \{u \in \llbracket \alpha \rrbracket \mid \mathcal{K} \approx_R \neg \alpha \sim \perp\}$ as the set of possible valuations w.r.t. \mathcal{K} . So $\mathcal{U}_R^{\mathcal{K}}$ refers to all the valuations not in conflict with rank entailment w.r.t. \mathcal{K} . From results by Lehmann and Magidor [25] (Lemmas 24 and 30) it follows that the possible valuations in the minimal model $\mathcal{R}_{\mathcal{K}}^{RC}$ are precisely the possible valuations w.r.t. \mathcal{K} : $\mathcal{U}_R^{\mathcal{K}} = \mathcal{U}^{\mathcal{R}_{\mathcal{K}}^{RC}}$.

Rational closure can also be defined in terms of the *base rank* of a statement. A propositional sentence α is said to be *exceptional* w.r.t. \mathcal{K} if $\mathcal{K} \approx_R \top \sim \neg \alpha$ (i.e., α is false in all the most typical valuations in every ranked model of \mathcal{K}). Let $\varepsilon(\mathcal{K}) = \{\alpha \sim \beta \mid \mathcal{K} \approx_R \top \sim \neg \alpha\}$. Define a sequence of knowledge bases $\mathcal{E}_0^{\mathcal{K}}, \dots, \mathcal{E}_{\infty}^{\mathcal{K}}$ as follows:

$\mathcal{E}_0^{\mathcal{K}} \equiv_{\text{def}} \mathcal{K}$, $\mathcal{E}_i^{\mathcal{K}} \equiv_{\text{def}} \varepsilon(\mathcal{E}_{i-1}^{\mathcal{K}})$, for $0 < i < n$, and $\mathcal{E}_\infty \equiv_{\text{def}} \mathcal{E}_n^{\mathcal{K}}$, where n is the smallest k for which $\mathcal{E}_k^{\mathcal{K}} = \mathcal{E}_{k+1}^{\mathcal{K}}$ (since \mathcal{K} is finite, n must exist). The *base rank* $br_{\mathcal{K}}(\alpha)$ of a propositional statement α w.r.t. a knowledge base \mathcal{K} is defined to be the smallest r for which α is *not* exceptional w.r.t. $\mathcal{E}_r^{\mathcal{K}}$. $br_{\mathcal{K}}(\alpha) \equiv_{\text{def}} \min\{r \mid \mathcal{E}_r^{\mathcal{K}} \not\vdash_R \top \vdash \neg\alpha\}$.

Observation 2 (Giordano et al. [21]) $\mathcal{K} \approx_{RC} \alpha \vdash \beta$ iff $br_{\mathcal{K}}(\alpha) < br_{\mathcal{K}}(\alpha \wedge \neg\beta)$ or $br_{\mathcal{K}}(\alpha) = \infty$.

There is a fundamental connection between the base ranks of propositional statements w.r.t. \mathcal{K} and the ranks of valuations in the minimal ranked model $\mathcal{R}_{\mathcal{K}}^{RC}$.

Observation 3 (Giordano et al. [21]) For every knowledge base \mathcal{K} and $\alpha \in \mathcal{L}$, $br_{\mathcal{K}}(\alpha) = \min\{i \mid \text{there is a } v \in \llbracket \alpha \rrbracket \text{ s.t. } \mathcal{R}_{\mathcal{K}}^{RC}(v) = i\}$.

From Observation 3 it follows that a classical statement α (or its defeasible representation $\neg\alpha \vdash \perp$) is in the rational closure of \mathcal{K} iff the base rank of $\neg\alpha$ w.r.t. \mathcal{K} is ∞ . The definition of base rank can be extended to defeasible implications as follows: $br_{\mathcal{K}}(\alpha \vdash \beta) \equiv_{\text{def}} br_{\mathcal{K}}(\alpha)$. Assigning base ranks to defeasible implications in this way forms the basis of an algorithm for computing rational closure; an algorithm that can be reduced to a number of classical entailment checks. Define the *materialisation* of a knowledge base \mathcal{K} as $\vec{\mathcal{K}} \equiv_{\text{def}} \{\alpha \rightarrow \beta \mid \alpha \vdash \beta \in \mathcal{K}\}$. It can be shown [25] that a sentence α is exceptional w.r.t. \mathcal{K} iff $\vec{\mathcal{K}} \models \neg\alpha$. From this we can define a procedure `BaseRank` which partitions the materialisation of \mathcal{K} into $n + 1$ classes according to base rank: $i = 0, \dots, n - 1, \infty$, $R_i \equiv_{\text{def}} \{\alpha \rightarrow \beta \mid \alpha \vdash \beta \in \mathcal{K}, br_{\mathcal{K}}(\alpha) = i\}$.

Algorithm 1: BaseRank

Input: A knowledge base \mathcal{K}
Output: An ordered tuple $(R_0, \dots, R_{n-1}, R_\infty, n)$

- 1 $i := 0$;
- 2 $E_0 := \vec{\mathcal{K}}$;
- 3 **repeat**
- 4 $E_{i+1} := \{\alpha \rightarrow \beta \in E_i \mid E_i \models \neg\alpha\}$;
- 5 $R_i := E_i \setminus E_{i+1}$;
- 6 $i := i + 1$;
- 7 **until** $E_{i-1} = E_i$;
- 8 $R_\infty := E_{i-1}$;
- 9 **if** $E_{i-1} = \emptyset$ **then**
- 10 $n := i - 1$;
- 11 **else**
- 12 $n := i$;
- 13 **return** $(R_0, \dots, R_{n-1}, R_\infty, n)$

We use `BaseRank` to describe an algorithm originally proposed by Freund [18] for computing rational closure. It takes as input \mathcal{K} and $\alpha \vdash \beta$, and returns **true** iff $\alpha \vdash \beta$ is in the rational closure of \mathcal{K} .

Algorithm 2: RationalClosure

Input: A knowledge base \mathcal{K} and a DI $\alpha \sim \beta$
Output: **true**, if $\mathcal{K} \approx \alpha \sim \beta$, and **false**, otherwise

- 1 $(R_0, \dots, R_{n-1}, R_\infty, n) := \text{BaseRank}(\mathcal{K});$
- 2 $i := 0;$
- 3 $R := \bigcup_{i=0}^{j < n} R_j;$
- 4 **while** $R_\infty \cup R \models \neg \alpha$ **and** $R \neq \emptyset$ **do**
- 5 $R := R \setminus R_i;$
- 6 $i := i + 1;$
- 7 **return** $R_\infty \cup R \models \alpha \rightarrow \beta;$

The algorithm keeps on removing (materialisations of) defeasible implications from (the materialisation of) \mathcal{K} , starting with the lowest base rank, and proceeding base rank by base rank, until it finds the first R which is classically consistent with α (and therefore α is not exceptional w.r.t. the defeasible version of R). $\alpha \sim \beta$ is then taken to be in the rational closure of \mathcal{K} iff R classically entails the materialisation of $\alpha \sim \beta$.

Observation 4 (Freund [18]) *Given \mathcal{K} and $\alpha \sim \beta$, RationalClosure returns true iff $\mathcal{K} \approx_{RC} \alpha \sim \beta$.*

Observe that RationalClosure involves a number of calls to a classical-entailment checker that is polynomial in the size of \mathcal{K} . Computing rational closure is therefore no harder than checking classical entailment.

3 Basic Defeasible Entailment

Our departure point for defining defeasible entailment is that it ought to be LM-rational. The central question we address in this paper is whether LM-rationality is sufficient. The immediate answer is that it is not. For starters, we also require \approx to satisfy Inclusion (all elements of \mathcal{K} should be defeasibly entailed by \mathcal{K}):

(Inclusion) $\mathcal{K} \approx \alpha \sim \beta$ for every $\alpha \sim \beta \in \mathcal{K}$

and Classic Preservation—the classical defeasible implications (those corresponding to classical sentences) defeasibly entailed by \mathcal{K} should correspond exactly to those in the monotonic core of \mathcal{K} (i.e., those that are rank entailed by \mathcal{K}):

(Classic Preservation) $\mathcal{K} \approx \alpha \sim \perp$ iff $\mathcal{K} \approx_R \alpha \sim \perp$

An easy corollary of Classic Preservation is Classic Consistency, requiring that a knowledge base is consistent iff it is consistent w.r.t. rank entailment.

(Classic Consistency) $\mathcal{K} \approx \top \sim \perp$ iff $\mathcal{K} \approx_R \top \sim \perp$

We refer to a defeasible entailment relation satisfying LM-rationality, Inclusion, and Classic Preservation as a *basic defeasible entailment relation*.

We shall see below (using Theorem 1) that rational closure is a basic defeasible entailment relation. However, since ranked entailment does not satisfy RM, it is not LM-rational, and is therefore not a basic defeasible entailment relation.

Definition 5. A ranked model \mathcal{R} of \mathcal{K} is said to be \mathcal{K} -faithful if the possible valuations in \mathcal{R} are precisely the possible valuations w.r.t. \mathcal{K} : $\mathcal{U}^{\mathcal{R}} = \mathcal{U}_{\mathcal{R}}^{\mathcal{K}}$.

Note that the minimal model $\mathcal{R}_{\mathcal{K}}^{RC}$ is \mathcal{K} -faithful.

Our first fundamental result (using points 1 and 2b of Theorem 1 below) is a semantic characterisation of basic defeasible entailment in terms of the \mathcal{K} -faithful ranked models. From this it also follows immediately that basic defeasible entailment satisfies the following property.

(Rank Extension) If $\mathcal{K} \approx_R \alpha \sim \beta$, then $\mathcal{K} \approx \alpha \sim \beta$

Rank Extension requires \approx to extend its monotonic core.

We can also characterise basic defeasible entailment by generalising the notion of base rank.

Definition 6. Let $r : \mathcal{L} \rightarrow \mathbb{N} \cup \{\infty\}$ be a rank function s.t. $r(\top) = 0$, satisfying the following convexity property: for every $i \in \mathbb{N}$, if $r(\alpha) = i$ then, for every j such that $0 \leq j < i$, there is a $\beta \in \mathcal{L}$ for which $r(\beta) = j$. r is entailment preserving if $\alpha \models \beta$ implies $r(\alpha) \geq r(\beta)$. r is \mathcal{K} -faithful if (i) it is entailment preserving; (ii) $r(\alpha) < r(\alpha \wedge \neg\beta)$ or $r(\alpha) = \infty$, for every $\alpha \sim \beta \in \mathcal{K}$, and (iii) $r(\alpha) = \infty$ iff $\mathcal{K} \approx_R \alpha \sim \perp$.

Observe that the base rank $br_{\mathcal{K}}(\cdot)$ is \mathcal{K} -faithful.

Definition 7. A rank function r generates a defeasible entailment relation \approx whenever $\mathcal{K} \approx \alpha \sim \beta$ if $r(\alpha) < r(\alpha \wedge \neg\beta)$ or $r(\alpha) = \infty$.

It follows (using points 1 and 2c of Theorem 1 below), that basic defeasible entailment can be characterised using the \mathcal{K} -faithful rank functions.

Next, we present an algorithm that computes the defeasible entailment relation generated by a \mathcal{K} -faithful rank function. It is a modified version of `RationalClosure`, differing from it in that the call to `BaseRank` is replaced with a call to the `Rank` algorithm described below. It receives as input a knowledge base \mathcal{K} and a \mathcal{K} -faithful rank function r . It produces as output a sequence $(R_0, \dots, R_{n-1}, R_\infty, n)$ where the R_i s are *sentences*, unlike `BaseRank`, which produces sets of sentences. `DefeasibleEntailment` is then adjusted accordingly.

`DefeasibleEntailment` removes statements, starting with the lowest rank, and proceeding rank by rank, until it finds the first R which is classically consistent with α . $\alpha \sim \beta$ is then taken to be defeasibly entailed by \mathcal{K} iff R classically entails the materialisation of $\alpha \sim \beta$. The R_i s correspond to classical representations of defeasible information, with different R_i s representing information with different levels of typicality, and with R_∞ corresponding to information that is classical. In fact, the set containing all the R_i s is equivalent to the materialisation of \mathcal{K} .

Algorithm 3: DefeasibleEntailment

Input: A knowledge base \mathcal{K} , a \mathcal{K} -faithful rank function r , and a DI $\alpha \sim \beta$
Output: **true**, if $\mathcal{K} \approx \alpha \sim \beta$, and **false**, otherwise

- 1 $(R_0, \dots, R_{n-1}, R_\infty, n) := \text{Rank}(\mathcal{K}, r)$;
- 2 $i := 0$;
- 3 $R := \bigcup_{i=0}^{j < n} \{R_j\}$;
- 4 **while** $\{R_\infty\} \cup R \models \neg \alpha$ **and** $R \neq \emptyset$ **do**
- 5 $R := R \setminus \{R_i\}$;
- 6 $i := i + 1$;
- 7 **return** $\{R_\infty\} \cup R \models \alpha \rightarrow \beta$;

Algorithm 4: Rank

Input: A knowledge base \mathcal{K} and a \mathcal{K} -faithful rank function r
Output: An ordered tuple $(R_0, \dots, R_{n-1}, R_\infty, n)$

- 1 $R_\infty := \neg \left(\bigvee_{r([\alpha]=\infty} [\alpha] \right)$;
- 2 $n := \max\{i \in \mathbb{N} \mid \text{there is an } \alpha \in \mathcal{L} \text{ s.t. } r(\alpha) = i\}$;
- 3 **if** $n = 0$ **then**
- 4 $R_0 := \top$; $n := 1$;
- 5 **else**
- 6 **for** $i := 0$ **to** $n - 1$ **do**
- 7 $R_i \equiv_{\text{def}} \neg \left(\bigvee_{r([\alpha]=i+1} [\alpha] \right)$
- 8 **return** $(R_0, \dots, R_{n-1}, R_\infty, n)$

For $\alpha \in \mathcal{L}$, let $[\alpha]$ be a canonical representative of the set $\{\beta \mid \beta \equiv \alpha\}$. Rank receives as input a knowledge base \mathcal{K} and a \mathcal{K} -faithful rank function r and, as mentioned above, produces as output an ordered tuple of *sentences* $(R_0, \dots, R_{n-1}, R_\infty, n)$.

If there is no α such that $r(\alpha) = \infty$, then R_∞ will be set to \top . This corresponds to the case where all information is defeasible. If $n = 0$, it corresponds to the case where there is no defeasible information. In this case we set n to 1 and set R_0 to \top .

Proposition 1. *Let $(R_0, \dots, R_{n-1}, R_\infty, n)$ be the output obtained from the Rank algorithm, given a knowledge base \mathcal{K} and a \mathcal{K} -faithful ranking function r . Then $\{R_\infty\} \cup \bigcup_{i=0}^{j < n} \{R_j\} \equiv \vec{\mathcal{K}}$.*

Example 1. Let $\mathcal{K} = \{\mathbf{p} \rightarrow \mathbf{b}, \mathbf{b} \sim \mathbf{f}, \mathbf{p} \sim \neg \mathbf{f}\}$. One can see there is only one ranking function r for which $r((\mathbf{b} \rightarrow \mathbf{f}) \rightarrow \mathbf{p}) = 1$, $r(\mathbf{p} \wedge (\mathbf{b} \rightarrow \mathbf{f})) = 2$, and $r(\neg(\mathbf{p} \rightarrow \mathbf{b})) = \infty$. Moreover, for every $\alpha \in \mathcal{L}$, $r(\alpha) = \infty$ or $r(\alpha) \leq 2$. Given \mathcal{K} and r , Rank will output the ordered tuple $(R_0, R_1, R_\infty, 2)$, where $R_\infty \equiv \mathbf{p} \rightarrow \mathbf{b}$, $R_1 \equiv \neg(\mathbf{p} \wedge (\mathbf{b} \rightarrow \mathbf{f})) \equiv \mathbf{p} \rightarrow (\mathbf{b} \wedge \neg \mathbf{f})$, and $R_0 \equiv \neg((\mathbf{b} \rightarrow \mathbf{f}) \rightarrow \mathbf{p}) \equiv (\mathbf{b} \rightarrow \mathbf{f}) \wedge \neg \mathbf{p}$. Given \mathcal{K} , r , and $(\mathbf{p} \leftrightarrow \mathbf{b}) \wedge (\mathbf{b} \leftrightarrow \mathbf{f}) \sim \neg \mathbf{f}$, DefeasibleEntailment will return **true**. It will do so by first verifying that $\{R_0, R_1, R_\infty\} \not\models \neg((\mathbf{p} \leftrightarrow \mathbf{b}) \wedge (\mathbf{b} \leftrightarrow \mathbf{f}))$ and then checking whether $\{R_0, R_1, R_\infty\} \models ((\mathbf{p} \leftrightarrow \mathbf{b}) \wedge (\mathbf{b} \leftrightarrow \mathbf{f})) \rightarrow \neg \mathbf{f}$ (which it does). Note that, given this r , DefeasibleEntailment computes the rational closure of \mathcal{K} . \square

Example 2. Let \mathcal{K} be as in Example 1. It can be shown that there is only one ranking function r s.t. $r(f \rightarrow p) = 1$, $r((b \vee f) \rightarrow (p \wedge f)) = 2$, and $r(\neg(p \rightarrow b)) = \infty$, and that r is \mathcal{K} -faithful. Moreover, for r it will be the case that for every $\alpha \in \mathcal{L}$, $r(\alpha) = \infty$ or $r(\alpha) \leq 2$. Given \mathcal{K} and r , the Rank algorithm will output the ordered tuple $(R_0, R_1, R_\infty, 2)$ where $R_\infty \equiv p \rightarrow b$, $R_1 \equiv \neg((b \vee f) \rightarrow (p \wedge f)) \equiv (\neg b \rightarrow f) \wedge (p \rightarrow \neg f)$, and $R_0 \equiv \neg(f \rightarrow p) \equiv f \wedge \neg p$. Given \mathcal{K} , r , and the DI $(p \leftrightarrow b) \wedge (b \leftrightarrow f) \vdash \neg f$, algorithm `DefeasibleEntailment` will return **false**. It will do so by first removing R_0 (since $\{R_0, R_1, R_\infty\} \models \neg((p \leftrightarrow b) \wedge (b \leftrightarrow f))$), then removing R_1 (since $\{R_1, R_\infty\} \models \neg((p \leftrightarrow b) \wedge (b \leftrightarrow f))$), and then, since $\{R_\infty\} \not\models \neg((p \leftrightarrow b) \wedge (b \leftrightarrow f))$, it will check whether $\{R_\infty\} \models ((p \leftrightarrow b) \wedge (b \leftrightarrow f)) \rightarrow \neg f$ (which it does not). \square

Definition 8. *DefeasibleEntailment* computes a defeasible entailment relation \approx for a knowledge base \mathcal{K} and a rank function r if *DefeasibleEntailment*, when presented with \mathcal{K} , r , and $\alpha \vdash \beta$, returns **true** if and only if $\mathcal{K} \approx \alpha \vdash \beta$.

It follows (using points 1 and 2d of Theorem 1) that `DefeasibleEntailment` computes exactly basic defeasible entailment.

Theorem 1. *The following statements are equivalent: (1) \approx is a basic defeasible \mathcal{K} -entailment relation, and (2) there is a \mathcal{K} -faithful ranked model \mathcal{R} and a \mathcal{K} -faithful rank function r such that:*

- a. $r(\alpha) = \min\{i \mid \text{there is a } v \in \llbracket \alpha \rrbracket \text{ s.t. } \mathcal{R}(v) = i\}$;
- b. \approx can be generated from \mathcal{R} ;
- c. \approx can be generated from r ;
- d. \approx can be computed by *Defeasible Entailment*, given \mathcal{K} and r as input.

Note that points 1 and 2 in Theorem 1 establish a connection between \mathcal{R} and r via a result that is a generalisation of Observation 3. And observe that `DefeasibleEntailment` involves a number of calls to a classic entailment checker that is linear in n times the size of \mathcal{K} (where n is the number returned by the Rank algorithm). But note also that n may be exponential in the size of \mathcal{K} .

4 Rational Defeasible Entailment

We now proceed by suggesting that basic defeasible entailment is too permissive. We first show that it does not satisfy RC Extension:

(RC Extension) If $\mathcal{K} \approx_{RC} \alpha \vdash \beta$, then $\mathcal{K} \approx \alpha \vdash \beta$

To see that basic defeasible entailment does not satisfy RC Extension, consider the following example.

Example 3. Figure 2(a) depicts the (\mathcal{K} -faithful) minimal ranked model \mathcal{R}_K^{RC} of $\mathcal{K} = \{p \rightarrow b, b \vdash f, p \vdash \neg f\}$. Note that $\mathcal{R}_K^{RC} \models \neg p \wedge \neg f \vdash \neg b$ and (from Definition 4) that $\mathcal{K} \approx_{RC} \neg p \wedge \neg f \vdash \neg b$. But for the \mathcal{K} -faithful ranked model \mathcal{R} in Figure 2(b) below it follows that $\mathcal{R} \not\models \neg p \wedge \neg f \vdash \neg b$. And from Theorem 1 it follows that for the basic defeasible \mathcal{K} -entailment relation \approx generated from \mathcal{R} , $\mathcal{K} \not\approx \neg p \wedge \neg f \vdash \neg b$. So RC Extension does not hold. \square

2	pbf
1	$\overline{\text{pbf}} \text{pbf}$
0	$\overline{\text{pbf}} \overline{\text{pbf}} \text{pbf}$

(a)

2	pbf
1	$\overline{\text{pbf}} \overline{\text{pbf}} \text{pbf}$
0	$\overline{\text{pbf}} \text{pbf}$

(b)

Fig. 2. Ranked models of the knowledge base in Example 3. 2(a) shows the minimal \mathcal{K} -faithful ranked model $\mathcal{R}_{\mathcal{K}}^{\text{RC}}$, while 2(b) depicts the \mathcal{K} -faithful ranked model \mathcal{R} .

If a basic defeasible entailment relation also satisfies RC Extension, we refer to it as *rational* defeasible entailment. We propose the class of rational defeasible entailment relations as those worthy of the term *rational* and analyse them further in the remainder of this section. We start by showing (points 1 and 2b of Theorem 2) that rational defeasible entailment can be characterised in terms of a subset of the \mathcal{K} -faithful ranked models, referred to as *rank preserving*.

Definition 9. A \mathcal{K} -faithful ranked model \mathcal{R} is said to be *rank preserving* if the following condition holds: for all $v, u \in \mathcal{U}$, if $\mathcal{R}_{\mathcal{K}}^{\text{RC}}(v) < \mathcal{R}_{\mathcal{K}}^{\text{RC}}(u)$, then $\mathcal{R}(v) < \mathcal{R}(u)$.

Informally, rank preservation requires the total preorder $\preceq_{\mathcal{R}}$ generated from \mathcal{R} to respect the relative positions assigned to valuations in the minimal model $\mathcal{R}_{\mathcal{K}}^{\text{RC}}$ of \mathcal{K} .

We can also characterise rational defeasible entailment (points 1 and 2c of Theorem 2) using a subclass of \mathcal{K} -faithful rank functions referred to as *base rank preserving*.

Definition 10. A \mathcal{K} -faithful rank function r is said to be *base rank preserving* if the following condition holds: for all $\alpha, \beta \in \mathcal{L}$, if $br_{\mathcal{K}}(\alpha) < br_{\mathcal{K}}(\beta)$, then $r(\alpha) < r(\beta)$.

Base rank preserving rank functions (or, the relations $<$ derivable from base rank preserving rank functions) respect the base rank (or rather, the relation $<$ derivable from the base rank). We show (points 1 and 2d of Theorem 2) that `DefeasibleEntailment` described in the previous section can also be used to compute rational defeasible entailment, provided it receives base rank preserving rank functions as input.

Theorem 2. The following statements are equivalent: (1) \models is a rational defeasible \mathcal{K} -entailment relation, and (2) there is a rank preserving \mathcal{K} -faithful ranked model \mathcal{R} and a \mathcal{K} -faithful base rank preserving rank function r s.t.:

- a. $r(\alpha) = \min\{i \mid v \in \llbracket \alpha \rrbracket \text{ and } \mathcal{R}(v) = i\}$;
- b. \models can be generated from \mathcal{R} ;
- c. \models can be generated from r ;
- d. \models can be computed from `Defeasible Entailment`, given \mathcal{K} and r as input.

Analogous to basic defeasible entailment, Points 1 and 2 of Theorem 2 establish a connection between \mathcal{R} and r via a result that is a generalisation of Observation 3.

5 Lexicographic Closure

We now turn our attention to *lexicographic closure*, a second form of defeasible entailment that has been studied in the literature [24]. Our central result is that lexicographic closure is a rational defeasible entailment relation. We also show that lexicographic closure can be characterised in three different ways: semantically via a rank preserving \mathcal{K} -faithful ranked model, in terms of a base preserving \mathcal{K} -faithful rank function r , and via `DefeasibleEntailment` when it is presented with r (and a knowledge base \mathcal{K}) as input. While the semantic construction of lexicographic closure is known [24], the other two constructions are new. We also show that there are rational defeasible entailment relations that extend lexicographic closure, which means that lexicographic closure is not the “boldest” form of rational defeasible entailment, as has been the conjecture in the literature.

Let $\mathcal{C}^{\mathcal{K}}$ be a function from \mathcal{U} to \mathbb{N} s.t. $\mathcal{C}^{\mathcal{K}}(v) = \#\{\alpha \sim \beta \in \mathcal{K} \mid v \Vdash \alpha \rightarrow \beta\}$ (where $\#X$ denotes the cardinality of the set X). The goal is to refine the ordering on \mathcal{U} obtained from the minimal model $\mathcal{R}_{\mathcal{K}}^{RC}$ with $\mathcal{C}^{\mathcal{K}}$: in comparing two valuations with the same rank w.r.t. $\mathcal{R}_{\mathcal{K}}^{RC}$, the one with a higher number will be viewed as more typical.

We define an ordering $\preceq_{LC}^{\mathcal{K}}$ on \mathcal{U} : $v \preceq_{LC}^{\mathcal{K}} u$ if $\mathcal{R}_{\mathcal{K}}^{RC}(u) = \infty$, or $\mathcal{R}_{\mathcal{K}}^{RC}(v) < \mathcal{R}_{\mathcal{K}}^{RC}(u)$, or $\mathcal{R}_{\mathcal{K}}^{RC}(v) = \mathcal{R}_{\mathcal{K}}^{RC}(u)$ and $\mathcal{C}^{\mathcal{K}}(v) \geq \mathcal{C}^{\mathcal{K}}(u)$. Then let $\mathcal{R}_{\mathcal{K}}^{LC}$ be the ranked interpretation obtained from $\preceq_{LC}^{\mathcal{K}}$, which we call the *lexicographic ranked model* of \mathcal{K} .

Definition 11. *The lexicographic closure \approx_{LC} of \mathcal{K} is defined as follows: $\mathcal{K} \approx_{LC} \alpha \sim \beta$ if $\mathcal{R}_{\mathcal{K}}^{LC} \Vdash \alpha \sim \beta$.*

Proposition 2. *$\mathcal{R}_{\mathcal{K}}^{LC}$ is a \mathcal{K} -faithful and rank preserving ranked model.*

From this result it follows from Theorems 2 and 1 that lexicographic closure is a rational and basic defeasible entailment relation. Lehmann [24, Theorem 3] already showed that lexicographic closure satisfies RC Extension.

Example 4. Figure 3(a) depicts the minimal ranked model $\mathcal{R}_{\mathcal{K}}^{RC}$ of $\mathcal{K} = \{p \rightarrow b, b \sim f, p \sim \neg f, b \sim w\}$, while Figure 3(b) depicts the lexicographic ranked model $\mathcal{R}_{\mathcal{K}}^{LC}$ of \mathcal{K} . From these two models we can see that $p \sim w$ (penguins usually have wings) is not in the rational closure of \mathcal{K} , but is in the lexicographic closure of \mathcal{K} . This is indicative of the difference between, what Lehmann refers to as Prototypical Reasoning and Presumptive Reasoning [24]. Presumptive Reasoning states that properties of a class are presumed to hold for all members of that class unless we have knowledge to the contrary. Because birds usually have wings we assume that penguins, being birds, usually have wings as well. Contrast this with Prototypical Reasoning which states that, while typical members of a class are presumed to inherit the properties of that class, the same does not hold for atypical members. According to Prototypical Reasoning, since penguins are atypical members of the class of birds, they do not inherit the property of having wings. Rational closure operates according to Prototypical Reasoning, while lexicographic closure adheres to Presumptive Reasoning. \square

We have seen that lexicographic closure (\approx_{LC}) can be generated from a \mathcal{K} -faithful rank preserving model. From Theorem 2 it then follows that there is a \mathcal{K} -faithful base

2	pbf \bar{w} pbfw
1	$\bar{p}\bar{b}\bar{f}\bar{w}$ $\bar{p}\bar{b}\bar{f}w$ $\bar{p}\bar{b}\bar{f}\bar{w}$ pbfw
0	$\bar{p}\bar{b}\bar{f}\bar{w}$ $\bar{p}\bar{b}\bar{f}w$ $\bar{p}\bar{b}\bar{f}\bar{w}$ $\bar{p}\bar{b}\bar{f}w$ $\bar{p}\bar{b}\bar{f}\bar{w}$ pbfw

(a)

5	pbf \bar{w}
4	pbfw
3	pbf \bar{w} $\bar{p}\bar{b}\bar{f}\bar{w}$
2	pbf \bar{w} $\bar{p}\bar{b}\bar{f}w$
1	pbfw
0	$\bar{p}\bar{b}\bar{f}\bar{w}$ $\bar{p}\bar{b}\bar{f}w$ $\bar{p}\bar{b}\bar{f}\bar{w}$ $\bar{p}\bar{b}\bar{f}w$ pbfw

(b)

Fig. 3. Ranked models of the knowledge base in Example 4. 3(a) shows the minimal model \mathcal{K} -faithful ranked model $\mathcal{R}_{\mathcal{K}}^{RC}$, while 3(b) depicts the lexicographic ranked model $\mathcal{R}_{\mathcal{K}}^{LC}$.

rank preserving rank function r from which \approx_{LC} can be generated. Furthermore, it can be generated by `DefeasibleEntailment`, given \mathcal{K} and r as input. We now show how to construct the \mathcal{K} -faithful base rank preserving rank function r mentioned above.

Definition 12. *The lexicographic rank w.r.t. a knowledge base \mathcal{K} is defined as $r_{\mathcal{K}}^{LC}(\alpha) \equiv_{def} \min\{\mathcal{R}_{\mathcal{K}}^{LC}(v) \mid v \in \llbracket \alpha \rrbracket\}$.*

Proposition 3. *The lexicographic rank $r_{\mathcal{K}}^{LC}$ w.r.t. a knowledge base \mathcal{K} is \mathcal{K} -faithful and base rank preserving.*

Now we show $r_{\mathcal{K}}^{LC}$ generates the same rational defeasible entailment relation as $\mathcal{R}_{\mathcal{K}}^{LC}$.

Proposition 4. $\mathcal{R}_{\mathcal{K}}^{LC} \Vdash \alpha \sim \beta$ iff $r_{\mathcal{K}}^{LC}(\alpha) < r_{\mathcal{K}}^{LC}(\alpha \wedge \neg\beta)$ or $r_{\mathcal{K}}^{LC}(\alpha) = \infty$.

Finally, `DefeasibleEntailment` computes the same (rational) defeasible entailment relation as $\mathcal{R}_{\mathcal{K}}^{LC}$ does when given the input \mathcal{K} and $r_{\mathcal{K}}^{LC}$.

Proposition 5. *`DefeasibleEntailment` returns **true** when given the input \mathcal{K} , $r_{\mathcal{K}}^{LC}$, and $\alpha \sim \beta$ iff $r_{\mathcal{K}}^{LC}(\alpha) < r_{\mathcal{K}}^{LC}(\alpha \wedge \neg\beta)$, or $r_{\mathcal{K}}^{LC}(\alpha) = \infty$.*

We conclude this section with an example which shows that lexicographic closure is not (always) the “boldest” form of rational defeasible entailment.

Example 5. Consider the knowledge base \mathcal{K} in Example 4 and let a \mathcal{K} -faithful ranked model \mathcal{R} be as depicted in Figure 5 below. \mathcal{R} is a refinement of the lexicographic ranked model $\mathcal{R}_{\mathcal{K}}^{LC}$ in Figure 4. It can be shown that \mathcal{R} is rank base preserving, and therefore it generates a rational defeasible \mathcal{K} -entailment relation \approx , and that \approx strictly extends lexicographic closure: If $\mathcal{K} \approx_{LC} \alpha \sim \beta$, then $\mathcal{K} \approx \alpha \sim \beta$, and there is at least one defeasible implication $\alpha \sim \beta$ such that $\mathcal{K} \approx \alpha \sim \beta$, but $\mathcal{K} \not\approx_{LC} \alpha \sim \beta$. For example, observe that $\mathcal{K} \approx b \wedge \neg f \wedge w \sim \neg p$, but $\mathcal{K} \not\approx_{LC} b \wedge \neg f \wedge w \sim \neg p$ \square

7	$\text{pb}\bar{\text{f}}\bar{\text{w}}$
6	$\text{pb}\bar{\text{f}}\text{w}$
5	$\text{pb}\bar{\text{f}}\bar{\text{w}}$
4	$\bar{\text{p}}\bar{\text{b}}\bar{\text{f}}\bar{\text{w}}$
3	$\text{pb}\bar{\text{f}}\text{w}$
2	$\bar{\text{p}}\bar{\text{b}}\bar{\text{f}}\text{w}$
1	$\bar{\text{p}}\bar{\text{b}}\text{f}\bar{\text{w}}$
0	$\bar{\text{p}}\bar{\text{b}}\bar{\text{f}}\bar{\text{w}} \bar{\text{p}}\bar{\text{b}}\bar{\text{f}}\text{w} \bar{\text{p}}\bar{\text{b}}\bar{\text{f}}\text{w} \bar{\text{p}}\bar{\text{b}}\text{f}\bar{\text{w}} \bar{\text{p}}\bar{\text{b}}\text{f}\text{w}$

Fig. 4. The ranked model \mathcal{R} of Example 5.

6 Related Work

The original work in the KLM style [23] was inspired by the work of Shoham [28], and investigated a class of non-monotonic consequence relations, where defeasible implication was viewed as a (non-monotonic) form of entailment. This approach was subsequently adapted by Lehmann and Magidor [25] to the case where \vdash is viewed as an object-level connective for defeasible implication, and where the focus then shifts to defeasible entailment (i.e., \approx) for a logic language that extends propositional logic with the defeasible implication connective \vdash .

We are aware of four instances of defeasible entailment that have been studied: ranked entailment [25] which is not LM-rational, rational closure [25, 4, 7, 21], and lexicographic closure [24] which are both regarded as appropriate forms of defeasible entailment, and relevant closure [17] which is also not LM-rational.

Our investigation here is reminiscent of the AGM framework for belief change [1, 19], where classes of belief change operators are studied. Rational closure can be viewed as the defeasible entailment equivalent of *full-meet belief contraction or revision* since, by virtue of the property of RC Extension, it is the most conservative of those defeasible entailment relations regarded as appropriate. The boldest forms of rational defeasible entailment can be seen as analogous to *maxichoice* belief contraction and revision: maxichoice operators are obtained by imposing a linear ordering on the propositional valuations that are counter-models of a *belief set*, while the boldest forms of rational defeasible entailment are obtained by imposing a linear ordering on \mathcal{U}_R^K , the set of possible valuations w.r.t. a knowledge base \mathcal{K} and then considering the defeasible entailment relations generated from the base rank preserving \mathcal{K} -faithful ranked models obtained from such linear orderings.

Studies of defeasible entailment beyond the propositional case include versions of defeasible implication in more expressive languages, most notably description logics [2, 3, 10, 14, 15, 20, 27, 26] and modal logics [8, 9, 11]. A different type of extension is one in which defeasible implication is enriched by either introducing a notion of typicality in propositional logic [5, 6, 4] or a notion of defeasible modality [12, 13].

7 Conclusion

The central focus of this paper is the question of determining what (defeasible) entailment means for propositional logic enriched with a defeasible implication connective. The short answer is that a defeasible entailment relation needs to be *rational* in a technical sense provided above. In arriving at this conclusion we have made a detour through the more permissive class of *basic* defeasible entailment relations.

There are at least three lines of research to which the work in this paper can lead. First is an analysis of concrete forms of rational defeasible entailment other than rational and lexicographic closure. Secondly, both basic and rational defeasible entailment is on the *knowledge level* [19] in the sense that the syntactic form of knowledge bases are, for the most part, irrelevant. But there is a strong case to be made for defining defeasible implication where syntax matters. This is analogous to the distinction between belief change on sets closed under classical consequence and base change [22], where the structure of the set of beliefs of an agent plays a role. And although lexicographic closure is an instance of rational defeasible entailment, it is an example of a form of entailment where the structure of the knowledge base matters. We conjecture that a syntax-based class of defeasible entailment will form a strict subclass of the class of rational defeasible entailment relations, and that lexicographic closure will be the strongest form of syntax-based rational defeasible entailment. Finally, we have presented an algorithm for computing any rational defeasible entailment relation, but the algorithm depends on the provision of a knowledge base \mathcal{K} , as well as a function that ranks all statements. With a syntax-based approach, it may be possible to use the structure of \mathcal{K} to rank statements, in the way that the `BaseRank` algorithm does in the process of computing rational closure.

Acknowledgments

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

References

1. Alchourrón, C., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* **50**, 510–530 (1985)
2. Bonatti, P., Faella, M., Petrova, I., Sauro, L.: A new semantics for overriding in description logics. *Artificial Intelligence* **222**, 1–48 (2015)
3. Bonatti, P., Sauro, L.: On the logical properties of the nonmonotonic description logic DL^N . *Artificial Intelligence* **248**, 85–111 (2017)
4. Booth, R., Casini, G., Meyer, T., Varzinczak, I.: On the entailment problem for a logic of typicality. In: *IJCAI 2015*. pp. 2805–2811 (2015)
5. Booth, R., Meyer, T., Varzinczak, I.: PTL: A propositional typicality logic. In: *Proceedings of the 13th European Conference on Logics in Artificial Intelligence (JELIA)*. pp. 107–119. No. 7519 in LNCS, Springer (2012)

6. Booth, R., Meyer, T., Varzinczak, I.: A propositional typicality logic for extending rational consequence. In: Fermé, E., Gabbay, D., Simari, G. (eds.) *Trends in Belief Revision and Argumentation Dynamics, Studies in Logic – Logic and Cognitive Systems*, vol. 48, pp. 123–154. King’s College Publications (2013)
7. Booth, R., Paris, J.: 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 (1998)
8. Boutilier, C.: Conditional logics of normality: A modal approach. *Artificial Intelligence* 68(1), 87–154 (1994)
9. Britz, K., Meyer, T., Varzinczak, I.: Preferential reasoning for modal logics. *Electronic Notes in Theoretical Computer Science* 278, 55–69 (2011)
10. Britz, K., Meyer, T., Varzinczak, I.: Semantic foundation for preferential description logics. In: Wang, D., Reynolds, M. (eds.) *AI 2011*. pp. 491–500. No. 7106 in *LNAI*, Springer (2011)
11. Britz, K., Meyer, T., Varzinczak, I.: Normal modal preferential consequence. In: Thielscher, M., Zhang, D. (eds.) *AI 2012*. pp. 505–516. No. 7691 in *LNAI*, Springer (2012)
12. Britz, K., Varzinczak, I.: From KLM-style conditionals to defeasible modalities, and back. *Journal of Applied Non-Classical Logics (JANCL)* (2017)
13. Britz, K., Varzinczak, I.: Preferential accessibility and preferred worlds. *Journal of Logic, Language and Information (JoLLI)* 27(2), 133–155 (2018)
14. Britz, K., Varzinczak, I.: Rationality and context in defeasible subsumption. In: Woltran, S., Ferrarotti, F. (eds.) *FoIKS 2018. LNCS*, Springer (2018)
15. Casini, G., Straccia, U.: Defeasible inheritance-based description logics. *JAIR* 48, 415–473 (2013)
16. Casini, G., Meyer, T.: Belief change in a preferential non-monotonic framework. In: *IJCAI 2017*. pp. 929–935 (2017)
17. Casini, G., Meyer, T., Moodley, K., Nortje, R.: Relevant closure: A new form of defeasible reasoning for description logics. In: *JELIA 2014*. pp. 92–106 (2014)
18. Freund, M.: Preferential reasoning in the perspective of Poole default logic. *Artificial Intelligence* 98, 209–235 (1998), <http://www.sciencedirect.com/science/article/pii/S0004370297000532>
19. Gärdenfors, P.: *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press (1988)
20. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.: A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence* 195, 165–202 (2013)
21. Giordano, L., Gliozzi, V., Olivetti, N., Pozzato, G.: Semantic characterization of rational closure: From propositional logic to description logics. *Art. Int.* 226, 1–33 (2015)
22. Hansson, S.: *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer (1999)
23. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44, 167–207 (1990)
24. Lehmann, D.: Another perspective on default reasoning. *Annals of Math. and Art. Int.* 15(1), 61–82 (1995)
25. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? *Art. Int.* 55, 1–60 (1992)
26. Pensel, M., Turhan, A.Y.: Including quantification in defeasible reasoning for the description logic \mathcal{EL}_\perp . In: Balduccini, M., Janhunen, T. (eds.) *LPNMR 2017*. pp. 78–84. No. 10377 in *LNCS*, Springer (2017)
27. Quantz, J., Royer, V.: A preference semantics for defaults in terminological logics. In: *KR 1992*. pp. 294–305 (1992)
28. Shoham, Y.: *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press (1988)