

Defeasible Entailment: from Rational Closure to Lexicographic Closure and Beyond

Giovanni Casini

CSC, Université du Luxembourg
Luxembourg
giovanni.casini@uni.lu

Thomas Meyer

CAIR & University of Cape Town
South Africa
tmeyer@cs.uct.ac.za

Ivan Varzinczak

CRIL, Univ. Artois & CNRS
France
varzinczak@cril.fr

Abstract

In this paper we present what we believe to be the first systematic approach for extending the framework for defeasible entailment first presented by Kraus, Lehmann, and Magidor—the so-called KLM approach. Drawing on the properties for KLM, we first propose a class of basic defeasible entailment relations. We characterise this basic framework in three ways: (i) semantically, (ii) in terms of a class of properties, and (iii) in terms of ranks on statements in a knowledge base. We also provide an algorithm for computing the basic framework. These results are proved through various representation results. We then refine this framework by defining the class of rational defeasible entailment relations. This refined framework is also characterised in three ways: semantically, in terms of a class of properties, and in terms of ranks on statements. We also provide an algorithm for computing the refined framework. Again, these results are proved through various representation results.

We argue that the class of rational defeasible entailment relations—a strengthening of basic defeasible entailment which is itself a strengthening of the original KLM proposal—is worthy of the term *rational* in the sense that all of them can be viewed as appropriate forms of defeasible entailment. We show that the two well-known forms of defeasible entailment, *rational closure* and *lexicographic closure*, fall within our rational defeasible framework. We show 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 (1990) (a.k.a. KLM, or System P plus Rational Monotony) is a well-known framework for defeasible reasoning. The KLM properties can be viewed as constraints on appropriate forms of defeasible entailment. Currently two are the most known forms of defeasible entailment satisfying those properties: rational closure (Lehmann and Magidor 1992) and lexicographic closure (Lehmann 1995). Both forms of defeasible entailment can be characterised in three ways: semantically, in terms of ranks, and algorithmically.

Here we present what we believe to be the first systematic approach for extending the framework for defeasible entail-

ment originally proposed by Kraus, Lehmann, and Magidor. Our first proposal for doing so is referred to as *basic defeasible entailment*. This framework can be obtained by a strengthening of the KLM properties: adding additional properties to those initially proposed by KLM. We then proceed to characterise basic defeasible entailment in two other ways. The first is a semantic characterisation in terms of a class of *ranked interpretations*. The second is 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 then provide an algorithm for computing the framework. Given a rank function, we present an algorithm for computing basic defeasible entailment. The algorithm is a generalisation of one that has been proposed for computing rational closure (Freund 1998).

Having defined basic defeasible entailment, we propose a further strengthening via an additional property; one that requires any defeasible entailment relation to extend rational closure. In doing so we ensure that rational closure is viewed as the most conservative form of defeasible entailment. We refer to the resulting class of defeasible entailment relations as *rational* and argue that the class as a whole is worthy of further investigation. Part of this justification is that both rational and lexicographic closure are rational defeasible entailment relations. But while rational closure is the most conservative form of rational defeasible entailment, it turns out there are forms of rational defeasible entailment that are “bolder” than lexicographic closure. We show that rational defeasible entailment can also be characterised in two other ways: a semantic characterisation in terms of a further restricted class of ranked interpretations, and a characterisation in terms of a further restricted class of ranks. We also provide an algorithm for computing rational defeasible entailment. Given one of the restricted rank functions, the algorithm computes rational defeasible entailment.

We argue that the framework for rational defeasible entailment is reminiscent of the AGM framework for belief change (Alchourrón, Gärdenfors, and Makinson 1985) when viewed from the semantic perspective. Rational closure is analogous to full-meet revision (and contraction); there are defeasible entailment relations analogous to (linear) max-choice revision (and contraction), and the semantic construction of the class of rational defeasible entailment relations bears a close resemblance to transitively relation partial-

meet revision (and contraction) (Gärdenfors 1988).

The remainder of the paper is structured as follows. Section 2 fixes the notation and terminology we use, and contains a summary of the necessary technical background: an introduction to KLM-style defeasible implication (Section 2.1), defeasible entailment for propositional languages enriched with a defeasible implication connective (Section 2.2), and an overview of rational closure (Section 2.3). Then, in Section 3 we present and discuss our notion of basic defeasible entailment. This is followed in Section 4 with an investigation into rational defeasible entailment. Section 5 is devoted to lexicographic closure and its relation to rational defeasible entailment, while Section 6 provides an overview of related work. Finally, in Section 7 we conclude and briefly point to future directions for this line of research.

2 Background

Let \mathcal{P} be 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 (finite) set of all propositional *valuations*, with 1 representing truth and 0 representing falsity. We use u, v, \dots to denote valuations. Whenever it eases the presentation, 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 (has the value 1) in the valuation, while the presence of a barred atom indicates that the atom is false (has the value 0) in the valuation. Thus, for the logic generated from $\mathcal{P} = \{p, q\}$, the valuation in which p is true and q is false will be represented as $p\bar{q}$. Satisfaction of a sentence $\alpha \in \mathcal{L}$ by $v \in \mathcal{U}$ is defined in the usual truth-functional way and is denoted by $v \Vdash \alpha$. The *models* of a set of sentences X is defined as: $\llbracket X \rrbracket \equiv_{\text{def}} \{v \in \mathcal{U} \mid v \Vdash \alpha \text{ for every } \alpha \in X\}$.

2.1 KLM-style defeasible implication

In the logic proposed by Kraus et al. (1990), often referred to as the *KLM approach*, we are interested in defeasible implications (or DIs) of the form $\alpha \sim \beta$, read as “typically, if α , then β ”. For instance, if $\mathcal{P} = \{b, f, p\}$, where b, f and p stand for, respectively, “being a bird”, “being able to fly”, and “being a penguin”, the following are examples of defeasible implications: $b \sim f$ (birds typically fly), $p \wedge b \sim \neg f$ (penguins that are birds typically do not fly).

The semantics of KLM-style rational defeasible implications is given by structures referred to as *ranked interpretations* (Lehmann and Magidor 1992). In this work we adopt the following alternative representation thereof:

Definition 1 A ranked interpretation \mathcal{R} is a function from \mathcal{U} to $\mathbb{N} \cup \{\infty\}$ such that $\mathcal{R}(u) = 0$ 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 $\mathcal{R}(u) = j$.

In a ranked interpretation, we call $\mathcal{R}(v)$ the *rank* of v w.r.t. \mathcal{R} . The intuition is that 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 (a connected, transitive ordering) $\preceq_{\mathcal{R}}$ on \mathcal{U} as follows: $v \preceq_{\mathcal{R}} u$ if and only if $\mathcal{R}(v) \leq \mathcal{R}(u)$. Moreover, given any total preorder \preceq on $V \subseteq \mathcal{U}$, we can use its strict version \prec to generate a ranked interpretation. To see how, first 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 (where the length of the \prec -path between any of the \prec -minimal elements and a \prec -minimal element is 0).

Definition 2 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 $\alpha, \beta \in \mathcal{L}$, we say \mathcal{R} satisfies (is a ranked model of) the conditional $\alpha \sim \beta$ (denoted $\mathcal{R} \Vdash \alpha \sim \beta$) if all the \prec -minimal α -valuations also satisfy β , i.e., if $\min_{\prec} \llbracket \alpha \rrbracket^{\mathcal{R}} \subseteq \llbracket \beta \rrbracket^{\mathcal{R}}$. We say \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 $\mathcal{K} = \{p \rightarrow b, b \sim f, p \sim \neg f\}$.¹

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

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

An important property of ranked interpretations is that *all* classical propositional sentences can be expressed as DIs. More precisely, we have the following result: for every \mathcal{R} and every $\alpha \in \mathcal{L}$, $\mathcal{R} \Vdash \alpha$ if and only if $\mathcal{R} \Vdash \neg\alpha \sim \perp$. 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. One of the central questions is to determine what *entailment* means in this context. That is, we aim to specify what it means for a defeasible implication to be entailed by a fixed knowledge \mathcal{K} . This is the main question with which we concern ourselves in this paper. We refer to this type of reasoning as *defeasible entailment* and denote it by \approx . It is

¹For brevity, we shall omit the valuations with rank ∞ in our graphical representations of ranked interpretations.

important to note that, for the purposes of this paper, we assume \mathcal{K} to be fixed.² Strictly speaking, we should therefore refer to defeasible \mathcal{K} -entailment. However, where there is no ambiguity, we frequently drop the \mathcal{K} prefix, and just refer to defeasible entailment.

It is well-accepted that defeasible entailment (unlike classical entailment) is not unique. Lehmann and Magidor (1992) put forward *rational closure* as an appropriate form of defeasible entailment, while Lehmann (1995) proposed *lexicographic closure* as an alternative. We consider both of these in more detail below.

More generally, in studying different forms of defeasible entailment, the position advocated by Lehmann and Magidor (1992), and one we adopt here as well, is to consider a number of *rationality properties*, referred to as the KLM properties, for defeasible entailment.

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

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

We can refer to defeasible entailment as being generated from single ranked interpretations.

Definition 3 A ranked interpretation \mathcal{R} is said to generate a defeasible \mathcal{K} -entailment relation $\approx_{\mathcal{R}}$ by setting $\mathcal{K} \approx_{\mathcal{R}} \alpha \vdash \beta$ if and only if $\mathcal{R} \Vdash \alpha \vdash \beta$. (If there isn't any ambiguity, we drop the subscript \mathcal{R} .)

Lehmann and Magidor proved the following useful result.

Observation 1 (Lehman & Magidor (1992)) *A defeasible entailment relation is LM-rational if and only if it can be generated from a ranked interpretation.*³

Using LM-rationality as a starting point, it is easy to see that the most obvious attempt at defining defeasible entailment does not make the grade (Lehmann and Magidor 1992, Sect. 4.2).

²For an investigation of the case where \mathcal{K} may vary, the reader is invited to consult the work of Casini and Meyer (2017).

³This result was originally phrased as a representation result about non-monotonic consequence relations on propositional statements (Kraus, Lehmann, and Magidor 1990), but were subsequently applied to defeasible entailment relations for propositional logic enriched with defeasible implication. Section 6 explains this in more detail.

Definition 4 A defeasible implication $\alpha \vdash \beta$ is rank entailed by a knowledge base \mathcal{K} (denoted as $\mathcal{K} \approx_R \alpha \vdash \beta$) if every ranked model of \mathcal{K} is also a ranked model of $\alpha \vdash \beta$.

Rank entailment is an application of the classical Tarskian tradition of entailment to ranked interpretations, but it does not satisfy rational monotonicity (RM) and is therefore not rational (Lehmann and Magidor 1992, Th. 4.2). Despite this, 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 (Casini and Meyer 2017).

2.3 Rational Closure

The first version of defeasible entailment satisfying LM-rationality that we consider is *rational closure* (Lehmann and Magidor 1992). Given a knowledge base \mathcal{K} , consider the ordering $\preceq_{\mathcal{K}}$ on all ranked models of \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. It is easy to see that $\preceq_{\mathcal{K}}$ is a weak partial order. Giordano et al. (2015) showed that there is a unique $\preceq_{\mathcal{K}}$ -minimal element. The rational closure of \mathcal{K} is defined in terms of this minimum ranked model (Giordano et al. 2015, Theorem 2).

Definition 5 Consider a knowledge base \mathcal{K} , and 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 $\alpha \vdash \beta$ is in the rational closure of \mathcal{K} (denoted as $\mathcal{K} \approx_{RC} \alpha \vdash \beta$) if $\mathcal{R}_{\mathcal{K}}^{RC} \Vdash \alpha \vdash \beta$.

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. Essentially, 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 is 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 prove useful to be able to refer to the possible valuations w.r.t. a knowledge base.

Definition 6 We refer to $\mathcal{U}_R^{\mathcal{K}} \equiv_{\text{def}} \mathcal{U} \setminus \{u \in \mathcal{U} \mid u \in \llbracket \alpha \rrbracket\}$ for some α s.t. $\mathcal{K} \approx_R \neg\alpha \vdash \perp$ as the set of possible valuations w.r.t. \mathcal{K} .

Informally $\mathcal{U}_R^{\mathcal{K}}$ refers to all the valuations not in conflict with rank entailment w.r.t. \mathcal{K} . From results by Lehmann and Magidor (1992) (Lemmas 24 and 30, to be precise) 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 w.r.t. \mathcal{K} .⁴ Given a knowledge base \mathcal{K} , a propositional sentence α is said to be *exceptional* w.r.t. \mathcal{K} if $\mathcal{K} \approx_R \top \vdash \neg\alpha$ (i.e., α is false in all the most typical valuations in every ranked model of \mathcal{K}). Let $\varepsilon(\mathcal{K}) =$

⁴In the literature, the base rank of a sentence is just referred to as its rank, but for reasons that will become clear we have opted for the term *base rank* here.

$\{\alpha \sim \beta \mid \mathcal{K} \approx_R \top \sim \neg\alpha\}$. Now define a sequence of knowledge bases $\mathcal{E}_0^K, \dots, \mathcal{E}_\infty^K$ as follows: $\mathcal{E}_0^K \equiv_{\text{def}} \mathcal{K}$, $\mathcal{E}_i^K \equiv_{\text{def}} \varepsilon(\mathcal{E}_{i-1}^K)$, for $0 < i < n$, and $\mathcal{E}_\infty^K \equiv_{\text{def}} \mathcal{E}_n^K$, where n is the smallest k for which $\mathcal{E}_k^K = \mathcal{E}_{k+1}^K$ (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^K . $br_{\mathcal{K}}(\alpha) \equiv_{\text{def}} \min\{r \mid \mathcal{E}_r^K \not\approx_R \top \sim \neg\alpha\}$.

Observation 2 (Lehmann and Magidor (1992)) $\mathcal{K} \approx_{RC} \alpha \sim \beta$ if and only if $br_{\mathcal{K}}(\alpha) < br_{\mathcal{K}}(\alpha \wedge \neg\beta)$ or $br_{\mathcal{K}}(\alpha) = \infty$.

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

Observation 3 (Lehmann and Magidor (1992)) 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 also follows immediately that a classical statement α (or its defeasible representation $\neg\alpha \sim \perp$) is in the rational closure of \mathcal{K} if and only if the base rank of $\neg\alpha$ w.r.t. \mathcal{K} is ∞ , as intuitively expected.

The definition of base rank can be extended to defeasible implications as follows: $br_{\mathcal{K}}(\alpha \sim \beta) \equiv_{\text{def}} br_{\mathcal{K}}(\alpha)$.

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 $\vec{\mathcal{K}} \equiv_{\text{def}} \{\alpha \rightarrow \beta \mid \alpha \sim \beta \in \mathcal{K}\}$. It can be shown (Lehmann and Magidor 1992) that a sentence α is exceptional w.r.t. \mathcal{K} if and only if $\vec{\mathcal{K}} \models \neg\alpha$. From this we can define a procedure `BaseRank` which partitions the materialisation of \mathcal{K} into $n + 1$ equivalence classes according to base rank: $i = 0, \dots, n-1, \infty$, $R_i \equiv_{\text{def}} \{\alpha \rightarrow \beta \mid \alpha \sim \beta \in \mathcal{K}, br_{\mathcal{K}}(\alpha) = i\}$.

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 can use the `BaseRank` procedure to define an algorithm for computing rational closure. It takes as input a

knowledge base \mathcal{K} and a DI $\alpha \sim \beta$, and returns **true** if and only if $\alpha \sim \beta$ 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 $R := \bigcup_{j=0}^{j < n} R_j$;
- 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** $R_\infty \cup R \models \alpha \rightarrow \beta$;

Informally, 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} if and only if R classically entails the materialisation of $\alpha \sim \beta$.

Observation 4 (Freund (1998)) Given a knowledge base \mathcal{K} and a defeasible implication $\alpha \sim \beta$, the algorithm `RationalClosure` returns **true** if and only if $\mathcal{K} \approx_{RC} \alpha \sim \beta$.

To conclude this section, we observe that algorithm `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

As discussed in the previous section, the 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. That is, is it justifiable to regard any form of defeasible entailment that is LM-rational as appropriate? The immediate answer to this question is that it is not. For starters, we require \approx to satisfy Inclusion and Classic Preservation.

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

(Classic Preservation) $\mathcal{K} \approx \alpha \sim \perp$ if and only if $\mathcal{K} \approx_R \alpha \sim \perp$

Inclusion simply requires all elements of \mathcal{K} to be defeasibly entailed by \mathcal{K} . Classic Preservation states that 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}). An easy corollary of Classic Preservation is Classic Consistency, requiring that a knowledge base is consistent if and only if it is consistent w.r.t. rank entailment.

(Classic Consistency) $\mathcal{K} \approx \top \sim \perp$ if and only if $\mathcal{K} \approx_R \top \sim \perp$

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.

We now proceed with presenting our first fundamental result, a semantic characterisation of basic defeasible entailment relations in which we consider a class of ranked models we refer to as \mathcal{K} -faithful ranked.

Definition 7 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} : $\mathcal{U}^{\mathcal{R}} = \mathcal{U}^{\mathcal{K}}$.

Note that the minimal model $\mathcal{R}_{\mathcal{K}}^{RC}$ is \mathcal{K} -faithful. This brings us to our first representation result.

Theorem 1 Every basic defeasible \mathcal{K} -entailment relation can be generated from a \mathcal{K} -faithful ranked model. Conversely, every \mathcal{K} -faithful ranked model generates a defeasible \mathcal{K} -entailment relation.

From this it follows immediately that basic defeasible entailment satisfies the following property.

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

To see why, note that if $\mathcal{K} \approx_R \alpha \vdash \beta$ then $\alpha \vdash \beta$ is satisfied by every ranked model of \mathcal{K} , and in particular, by the ranked model used to generate \approx .

Rank Extension requires \approx to extend its monotonic core (i.e., it is required to extend the rank entailment of \mathcal{K}).

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

Definition 8 Let $r : \mathcal{L} \rightarrow \mathbb{N} \cup \{\infty\}$ be a rank function such that $r(\top) = 0$ and 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 referred to as entailment preserving if $\alpha \models \beta$ implies that $r(\alpha) \geq r(\beta)$. Given a knowledge base \mathcal{K} , r is said to be \mathcal{K} -faithful if (i) it is entailment preserving; (ii) $r(\alpha) < r(\alpha \wedge \neg\beta)$ or $r(\alpha) = \infty$, for every $\alpha \vdash \beta \in \mathcal{K}$, and (iii) $r(\alpha) = \infty$ if and only if $\mathcal{K} \approx_R \alpha \vdash \perp$.

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

Definition 9 A rank function r generates a defeasible entailment relation \approx whenever $\mathcal{K} \approx \alpha \vdash \beta$ if $r(\alpha) < r(\alpha \wedge \neg\beta)$ or $r(\alpha) = \infty$.

We can now present our second representation result.

Theorem 2 Every basic defeasible \mathcal{K} -entailment relation can be generated by a \mathcal{K} -faithful rank function r . Conversely, every \mathcal{K} -faithful rank function r generates a basic defeasible \mathcal{K} -entailment relation.

Next, we present algorithm `DefeasibleEntailment` that computes the defeasible entailment relation generated by a \mathcal{K} -faithful rank function. It is a modified version of the `RationalClosure` algorithm presented earlier, differing from that algorithm in that the call to the `BaseRank` algorithm is replaced with a call to the `Rank` algorithm described below. As for the `Rank` algorithm, it receives as

input, not just a knowledge base \mathcal{K} as `BaseRank` does, but also a \mathcal{K} -faithful rank function r , and then produces as output a sequence $(R_0, \dots, R_{n-1}, R_{\infty}, n)$ where the R_i s are *sentences*, unlike the `BaseRank` algorithm, which produces sets of sentences. `DefeasibleEntailment` is then adjusted accordingly.

Algorithm 3: `DefeasibleEntailment`

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

- 1 $(R_0, \dots, R_{n-1}, R_{\infty}, n) := \text{Rank}(\mathcal{K}, r)$;
- 2 $i := 0$;
- 3 $R := \bigcup_{j=0}^{i-1} \{R_j\}$;
- 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** $\{R_{\infty}\} \cup R \models \alpha \rightarrow \beta$;

Like the `RationalClosure` algorithm, the `DefeasibleEntailment` algorithm keeps on removing statements, starting with the lowest rank, and proceeding rank by rank, until it finds the first R which is classically consistent with α . $\alpha \vdash \beta$ is then taken to be defeasibly entailed by \mathcal{K} if and only if R classically entails the materialisation of $\alpha \vdash \beta$. Intuitively, 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, rather than defeasible. In fact, the set containing all the R_i s is equivalent to the materialisation of \mathcal{K} .

For $\alpha \in \mathcal{L}$, let $[\alpha]$ be a canonical representative of the set $\{\beta \mid \beta \equiv \alpha\}$. The `Rank` algorithm 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* $(R_0, \dots, R_{n-1}, R_{\infty}, n)$.

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** $(R_0, \dots, R_{n-1}, R_{\infty}, n)$

Note that 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. Note also that 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 . Also, as mentioned above, the set consisting of all the R_i s is equivalent to \mathcal{K} .

Lemma 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 $\{R_\infty\} \cup \bigcup_{j=0}^{j < n} \{R_j\} \equiv \vec{\mathcal{K}}$.

To get a sense of how the algorithm works, consider the following examples.

Example 1 Let $\mathcal{K} = \{p \rightarrow b, b \vdash f, p \vdash \neg f\}$. It can be shown that there is only one ranking function r for which $r((b \rightarrow f) \rightarrow p) = 1$, $r(p \wedge (b \rightarrow f)) = 2$, and $r(\neg(p \rightarrow b)) = \infty$. 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 $R_\infty \equiv p \rightarrow b$,

$$R_1 \equiv \neg(p \wedge (b \rightarrow f)) \equiv p \rightarrow (b \wedge \neg f), \text{ and}$$

$$R_0 \equiv \neg((b \rightarrow f) \rightarrow p) \equiv (b \rightarrow f) \wedge \neg p.$$

Given \mathcal{K} , r , and $(p \leftrightarrow b) \wedge (b \leftrightarrow f) \vdash \neg f$, DefeasibleEntailment will return **true**. It will do so by first verifying that $\{R_0, R_1, R_\infty\} \not\models \neg((p \leftrightarrow b) \wedge (b \leftrightarrow f))$ and then checking whether $\{R_0, R_1, R_\infty\} \models ((p \leftrightarrow b) \wedge (b \leftrightarrow f)) \rightarrow \neg f$ (which it does). It is worth noting that, given this r , algorithm DefeasibleEntailment computes the rational closure of \mathcal{K} . ■

Example 2 Consider again $\mathcal{K} = \{p \rightarrow b, b \vdash f, p \vdash \neg f\}$. 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 $R_\infty \equiv p \rightarrow b$,

$$R_1 \equiv \neg((b \vee f) \rightarrow (p \wedge f)) \equiv (\neg b \rightarrow f) \wedge (p \rightarrow \neg f), \text{ and}$$

$$R_0 \equiv \neg(f \rightarrow p) \equiv f \wedge \neg p.$$

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 $\{R_\infty\} \not\models \neg((p \leftrightarrow b) \wedge (b \leftrightarrow f)) \rightarrow \neg f$ (which it does not). ■

Definition 10 Algorithm DefeasibleEntailment is said to compute a defeasible entailment relation \approx for a knowledge base \mathcal{K} and a rank function r whenever $\mathcal{K} \approx \alpha \vdash \beta$ if DefeasibleEntailment, when presented with \mathcal{K} , r , and $\alpha \vdash \beta$, returns **true**.

This provides us with the material for our third representation result.

Theorem 3 Given a \mathcal{K} -faithful rank function r , the basic defeasible entailment relation generated by r is exactly the defeasible entailment relation computed by the DefeasibleEntailment algorithm when given \mathcal{K} and r as input.

The results obtained for basic defeasible entailment can therefore be summarised in the following theorem.

Theorem 4 The following statements are equivalent.

- \approx is a basic defeasible \mathcal{K} -entailment relation.
- There is a \mathcal{K} -faithful ranked model \mathcal{R} and a \mathcal{K} -faithful rank function r such that:
 1. $r(\alpha) = \min\{i \mid \text{there is a } v \in \llbracket \alpha \rrbracket \text{ s.t. } \mathcal{R}(v) = i\}$;
 2. \approx can be generated from \mathcal{R} ;
 3. \approx can be generated from r ;
 4. \approx can be computed by algorithm DefeasibleEntailment, given \mathcal{K} and r as input.

Note that point 1 in Theorem 4 establishes a connection between \mathcal{R} and r via a result that is a generalisation of Observation 3.

Finally, 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

Having analysed basic defeasible entailment in the previous section, we now proceed by contending that it is too permissive. In particular, we first show that it does not satisfy the following property.

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

RC Extension requires of \approx to extend the rational closure of \mathcal{K} . To see that basic defeasible entailment does not satisfy RC Extension, consider the following example.

Example 3 Figure 2 depicts the (\mathcal{K} -faithful) minimal ranked model $\mathcal{R}_{\mathcal{K}}^{RC}$ of $\mathcal{K} = \{p \rightarrow b, b \vdash f, p \vdash \neg f\}$. Note that $\mathcal{R}_{\mathcal{K}}^{RC} \models \neg p \wedge \neg f \vdash \neg b$. From Definition 5 it then follows that $\mathcal{K} \approx_{RC} \neg p \wedge \neg f \vdash \neg b$. But also note that for the \mathcal{K} -faithful ranked model \mathcal{R} in Figure 3 below it follows that $\mathcal{R} \not\models \neg p \wedge \neg f \vdash \neg b$. And from Theorem 4 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. ■

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

Figure 2: The minimal \mathcal{K} -faithful ranked model $\mathcal{R}_{\mathcal{K}}^{RC}$

If a basic defeasible entailment relation satisfies RC Extension as well, 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 that rational defeasible entailment can be characterised semantically in terms of a subset of the \mathcal{K} -faithful ranked models.

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

Figure 3: The \mathcal{K} -faithful ranked model \mathcal{R}

Definition 11 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 $\mathcal{R}_K^{RC}(v) < \mathcal{R}_K^{RC}(u)$, 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}_K^{RC} of \mathcal{K} .

Theorem 5 Every rational defeasible \mathcal{K} -entailment relation can be generated by a rank preserving \mathcal{K} -faithful model. Conversely, every rank preserving \mathcal{K} -faithful model generates a rational defeasible \mathcal{K} -entailment relation.

We can also characterise rational defeasible entailment using a subclass of \mathcal{K} -faithful rank functions.

Definition 12 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 $r(\alpha) < r(\beta)$.

As the name indicates, base rank preserving rank functions (or rather, the relations $<$ derivable from base rank preserving rank functions) respect the base rank (or rather, the relation $<$ derivable from the base rank).

Theorem 6 Every rational defeasible \mathcal{K} -entailment relation can be generated by a \mathcal{K} -faithful base rank preserving rank function. Conversely, every \mathcal{K} -faithful base rank preserving rank function generates a rational defeasible \mathcal{K} -entailment relation.

The following result shows that the algorithm `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 7 The defeasible entailment relation computed from algorithm `DefeasibleEntailment`, given a knowledge base \mathcal{K} and a \mathcal{K} -faithful base rank preserving rank function, is a rational defeasible \mathcal{K} -entailment relation. Conversely, every rational defeasible \mathcal{K} -entailment relation can be computed from algorithm `DefeasibleEntailment` when given \mathcal{K} and a \mathcal{K} -faithful base rank preserving rank function as input.

The results obtained for rational defeasible entailment can therefore be summarised in the following theorem.

Theorem 8 The following statements are equivalent.

- \approx is a rational defeasible \mathcal{K} -entailment relation.
- 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.:
 1. $r(\alpha) = \min\{i \mid v \in \llbracket \alpha \rrbracket \text{ and } \mathcal{R}(v) = i\}$;
 2. \approx can be generated from \mathcal{R} ;

3. \approx can be generated from r ;
4. \approx can be computed from algorithm `DefeasibleEntailment`, given \mathcal{K} and r as input.

Analogous to the case for basic defeasible entailment, Point 1 of Theorem 8 establishes a connection between \mathcal{R} and r via a result that is a generalisation of Observation 3.

5 Lexicographic Closure

In this section we turn our attention to *lexicographic closure*, a second form of defeasible entailment, other than rational closure, that has been studied in the literature (Lehmann 1995). Our central result is that lexicographic closure is a rational defeasible entailment relation, confirming our contention that rational defeasible entailment is a class of defeasible relations worth investigating. 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 the `DefeasibleEntailment` algorithm when it is presented with r (and a knowledge base \mathcal{K}) as input. While the semantic construction of lexicographic closure is known (Lehmann 1995), the other two constructions are new. Finally, we 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.

For a knowledge base \mathcal{K} , let $\mathcal{C}^{\mathcal{K}}$ be a function from \mathcal{U} to \mathbb{N} s.t. $\mathcal{C}^{\mathcal{K}}(v) = \#\{\alpha \sim \beta \in \mathcal{K} \mid v \Vdash \alpha \rightarrow \beta\}$.⁵ So $\mathcal{C}^{\mathcal{K}}(v)$ is the number of DIs in \mathcal{K} whose materialisations are satisfied by v . In defining lexicographic closure, the goal is to refine the ordering on \mathcal{U} obtained from the minimal model \mathcal{R}_K^{RC} with $\mathcal{C}^{\mathcal{K}}$: in comparing two valuations with the same rank w.r.t. \mathcal{R}_K^{RC} , the one with a higher number will be viewed as more typical.

Given a knowledge base \mathcal{K} , we define an ordering $\preceq_{LC}^{\mathcal{K}}$ on \mathcal{U} as follows: $v \preceq_{LC}^{\mathcal{K}} u$ if $\mathcal{R}_K^{RC}(u) = \infty$, or $\mathcal{R}_K^{RC}(v) < \mathcal{R}_K^{RC}(u)$, or $\mathcal{R}_K^{RC}(v) = \mathcal{R}_K^{RC}(u)$ and $\mathcal{C}^{\mathcal{K}}(v) \geq \mathcal{C}^{\mathcal{K}}(u)$. Then we let \mathcal{R}_K^{LC} be the ranked interpretation generated from $\preceq_{LC}^{\mathcal{K}}$. We refer to \mathcal{R}_K^{LC} as the lexicographic ranked model of \mathcal{K} .

Definition 13 The lexicographic closure \approx_{LC} of \mathcal{K} is defined as follows: $\mathcal{K} \approx_{LC} \alpha \sim \beta$ if $\mathcal{R}_K^{LC} \Vdash \alpha \sim \beta$.

The next result shows that the lexicographic ranked model of \mathcal{K} is \mathcal{K} -faithful and rank preserving.

Proposition 1 \mathcal{R}_K^{LC} is a \mathcal{K} -faithful and rank preserving ranked model.

From this result it follows immediately from Theorems 8 and 4 that lexicographic closure is a rational and basic defeasible entailment relation. In fact, Lehmann (1995, Theorem 3) already showed that lexicographic closure satisfies RC Extension.

To appreciate some of the differences between rational and lexicographic closure, consider the following example.

⁵ $\#X$ denotes the cardinality of the set X .

Example 4 Figure 4 depicts the minimal ranked model \mathcal{R}_K^{RC} of $\mathcal{K} = \{p \rightarrow b, b \sim f, p \sim \neg f, b \sim w\}$, while Figure 5 depicts the lexicographic ranked model \mathcal{R}_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 (1995). 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. So, because birds usually have wings we assume that penguins, being birds, usually have wings as well, since we don't have information to the contrary. 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 usually don't fly), they do not inherit the property of having wings. Rational closure operates according to Prototypical Reasoning, while lexicographic closure adheres to Presumptive Reasoning. ■

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

Figure 4: The minimal model \mathcal{K} -faithful ranked model \mathcal{R}_K^{RC}

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

Figure 5: The lexicographic ranked model \mathcal{R}_K^{LC}

We have already seen that lexicographic closure (\approx_{LC}) can be generated from a \mathcal{K} -faithful rank preserving model. From Theorem 8 it then follows that there is a \mathcal{K} -faithful base rank preserving rank function r from which \approx_{LC} can be generated. Furthermore, it can be generated by algorithm `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 14 The lexicographic rank w.r.t. a knowledge base \mathcal{K} is defined as $r_K^{LC}(\alpha) \equiv_{\text{def}} \min\{\mathcal{R}_K^{LC}(v) \mid v \in \llbracket \alpha \rrbracket\}$.

First we show that r_K^{LC} is \mathcal{K} -faithful and base rank preserving.

Proposition 2 The lexicographic rank r_K^{LC} w.r.t. a knowledge base \mathcal{K} is \mathcal{K} -faithful and base rank preserving.

Next, we show that r_K^{LC} generates the same rational defeasible entailment relation as \mathcal{R}_K^{LC} .

Proposition 3 $\mathcal{R}_K^{LC} \Vdash \alpha \sim \beta$ if and only if $r_K^{LC}(\alpha) < r_K^{LC}(\alpha \wedge \neg\beta)$ or $r_K^{LC}(\alpha) = \infty$.

Finally, we show that the `DefeasibleEntailment` algorithm computes the same (rational) defeasible entailment relation as \mathcal{R}_K^{LC} does when given the input \mathcal{K} and r_K^{LC} .

Proposition 4 Given a knowledge base \mathcal{K} and a defeasible implication $\alpha \sim \beta$, the `DefeasibleEntailment` algorithm returns **true** when given the input \mathcal{K} , r_K^{LC} , and $\alpha \sim \beta$ if and only if $r_K^{LC}(\alpha) < r_K^{LC}(\alpha \wedge \neg\beta)$, or $r_K^{LC}(\alpha) = \infty$.

We conclude this section by showing that, while lexicographic closure extends rational closure, it is not (always) the “boldest” form of rational defeasible entailment. To do so, we give an example of a knowledge base for which there is a rational defeasible entailment relation that extends lexicographic closure.

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. It is easy to see that \mathcal{R} is a refinement of the lexicographic ranked model \mathcal{R}_K^{LC} in Figure 6. 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 $\mathcal{K} \not\approx_{LC} b \wedge \neg f \wedge w \sim \neg p$. ■

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

Figure 6: The ranked model \mathcal{R} of Example 5.

6 Related Work

The original work in the KLM style (Kraus, Lehmann, and Magidor 1990) was inspired by the work of Shoham (1988),

and investigated a class of non-monotonic consequence relations, where the defeasible implication \vdash was viewed as the (non-monotonic) form of entailment. This approach was subsequently adapted by Lehmann and Magidor (1992) 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 .

When it comes to defeasible entailment, while there has been some work in this regard (Goldszmidt and Pearl 1996; Bezzazi, Makinson, and Perez 1997; Casini et al. 2014; Beierle et al. 2016; Kern-Isberner 2018), we are aware of three instances that have been studied in detail: ranked entailment (Lehmann and Magidor 1992) which is not LM-rational and is judged to be too weak, rational closure (Lehmann and Magidor 1992) and lexicographic closure (Lehmann 1995) which are both regarded as appropriate forms of defeasible entailment. Of these, rational closure is by far the best studied form of defeasible entailment (Booth et al. 2015; Booth and Paris 1998; Giordano et al. 2015).

Despite the work mentioned above, the present paper is, to the best of our knowledge, the first systematic attempt to characterise appropriate classes of defeasible entailment relations (to be distinguished from the original work by Kraus et al. (1990), which is a study of \vdash as a form of non-monotonic consequence).

Our work is reminiscent of the AGM framework for belief change (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1988), where belief change operators are studied. Taking the analogy further, 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 we regard as appropriate. Taking it even further, the boldest forms of rational defeasible entailment can be regarded as analogous to *maxichoice* belief contraction and revision. To see this, observe that maxichoice operators are obtained by imposing a linear ordering on the propositional valuations that are counter-models of a *belief set*. Similarly, 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 have also started to move beyond the propositional case, and now includes cases involving versions of defeasible implication in more expressive languages, most notably description logics (Bonatti et al. 2015; Bonatti and Sauro 2017; Britz, Meyer, and Varzinczak 2011b; Britz and Varzinczak 2018b; Casini and Straccia 2013; Giordano et al. 2013; Quantz and Royer 1992; Pensel and Turhan 2018; Casini, Straccia, and Meyer 2018) and modal logics (Boutilier 1994; Britz, Meyer, and Varzinczak 2011a; 2012). A slightly different type of extension is one in which defeasible implication is enriched by either introducing an explicit notion of typicality in propositional logic (Booth, Meyer, and Varzinczak 2012; 2013;

Booth et al. 2015) or a notion of defeasible modality (Britz and Varzinczak 2017; 2018a).

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 to this question provided here is that a defeasible entailment relation needs to be *rational* in the technical sense described in Section 4. In arriving at this conclusion we have made a detour through the more permissive class of *basic* defeasible entailment relations defined in Section 3. Both basic and rational defeasible entailment are characterised in four different ways, through sets of properties, semantically via ranked interpretations, in terms of ranks assigned to (propositional and defeasible) statements, and algorithmically. While basic defeasible entailment tightens the requirements imposed by KLM-style defeasible entailment somewhat, rational defeasible entailment goes further by requiring that the form of defeasible entailment known as rational closure ought to be viewed as the most basic form of defeasible entailment. Part of the argument in favour of rational defeasible entailment is that, as is the case for rational closure, lexicographic closure (the other well-known form of defeasible entailment) is also rational.

There are at least three important lines of research to which the work in this paper can lead. First on the agenda is an analysis of concrete forms of rational defeasible entailment other than rational and lexicographic closure.

Secondly, the description of both basic and rational defeasible entailment in this paper can be viewed as being on the *knowledge level* (Gärdenfors 1988) 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. Roughly speaking, this is analogous to the distinction between belief change on belief sets (sets closed under classical consequence) and base change (Hansson 1999), where the structure of the set of beliefs of an agent plays a role in determining how change ought to occur. In fact, 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. Our current conjecture is 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, syntax-based defeasible entailment opens the door for studying the computation of defeasible entailment in more detail. 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 propositional (and therefore all defeasible implication) statements. With a syntax-based approach, it is possible to use the structure of \mathcal{K} to rank statements, in the way that the `BaseRank` algorithm in Section 2.3 does in the process of computing rational closure.

Acknowledgements

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

References

- Alchourrón, C.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50:510–530.
- Beierle, C.; Eichhorn, C.; Kern-Isberner, G.; and Kutsch, S. 2016. Skeptical, weakly skeptical, and credulous inference based on preferred ranking functions. In *ECAI 2016*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, 1149–1157. IOS Press.
- Bezzazi, H.; Makinson, D.; and Perez, R. P. 1997. Beyond rational monotony: Some strong non-horn rules for nonmonotonic inference relations. *J. Logic Computat.* 7(5):605–631.
- Bonatti, P., and Sauro, L. 2017. On the logical properties of the nonmonotonic description logic DL^N . *Artificial Intelligence* 248:85–111.
- Bonatti, P.; Faella, M.; Petrova, I.; and Sauro, L. 2015. A new semantics for overriding in description logics. *Artificial Intelligence* 222:1–48.
- 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.; Meyer, T.; and Varzinczak, I. 2012. PTL: A propositional typicality logic. In *Proceedings of the 13th European Conference on Logics in Artificial Intelligence (JELIA)*, number 7519 in LNCS, 107–119. Springer.
- 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.
- Britz, K., and Varzinczak, I. 2017. From KLM-style conditionals to defeasible modalities, and back. *Journal of Applied Non-Classical Logics (JANCL)*.
- Britz, K., and Varzinczak, I. 2018a. Preferential accessibility and preferred worlds. *Journal of Logic, Language and Information (JoLLI)* 27(2):133–155.
- Britz, K., and Varzinczak, I. 2018b. Rationality and context in defeasible subsumption. In Woltran, S., and Ferrarotti, F., eds., *FoIKS 2018*, LNCS. Springer.
- Britz, K.; Meyer, T.; and Varzinczak, I. 2011a. Preferential reasoning for modal logics. *Electronic Notes in Theoretical Computer Science* 278:55–69.
- Britz, K.; Meyer, T.; and Varzinczak, I. 2011b. Semantic foundation for preferential description logics. In Wang, D., and Reynolds, M., eds., *AI 2011*, number 7106 in LNAI, 491–500. Springer.
- Britz, K.; Meyer, T.; and Varzinczak, I. 2012. Normal modal preferential consequence. In Thielscher, M., and Zhang, D., eds., *AI 2012*, number 7691 in LNAI, 505–516. Springer.
- Casini, G., and Meyer, T. 2017. Belief change in a preferential non-monotonic framework. In *IJCAI 2017*, 929–935.
- 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.; Straccia, U.; and Meyer, T. 2018. A polynomial time subsumption algorithm for nominal safe $\mathcal{EL}\mathcal{O}_\perp$ under rational closure. *Information Sciences*, <https://doi.org/10.1016/j.ins.2018.09.037>.
- Freund, M. 1998. Preferential reasoning in the perspective of Poole default logic. *Artificial Intelligence* 98:209–235.
- Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press.
- Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2013. A non-monotonic description logic for reasoning about typicality. *Artificial Intelligence* 195:165–202.
- 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.
- Goldszmidt, M., and Pearl, J. 1996. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence* 85:57–112.
- Hansson, S. 1999. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer.
- Kern-Isberner, G. 2018. Axiomatizing a qualitative principle of conditional preservation for iterated belief change. In *KR 2018*, accepted.
- 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.
- Pensel, M., and Turhan, A.-Y. 2018. Reasoning in the defeasible description logic \mathcal{EL}_\perp - computing standard inferences under rational and relevant semantics. *International Journal of Approximate Reasoning* 103:28 – 70.
- Quantz, J., and Royer, V. 1992. A preference semantics for defaults in terminological logics. In *KR 1992*, 294–305.
- Shoham, Y. 1988. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press.