Towards Conditional Inference under Disjunctive Rationality

Richard Booth[1] , **Ivan Varzinczak**[2,3]
[1]Cardiff University, United Kingdom
[2]CRIL, Univ. Artois & CNRS, France
[3]CAIR, Computer Science Division, Stellenbosch University, South Africa
boothr2@cardiff.ac.uk, varzinczak@cril.fr

Abstract

The question of *conditional inference*, i.e., of which conditional sentences of the form "if α then, normally, β" should follow from a set $\mathcal{KB}$ of such sentences, has been one of the classic questions of non-monotonic reasoning, with several well-known solutions proposed. Perhaps the most notable is the *rational closure* construction of Lehmann and Magidor, under which the set of inferred conditionals forms a rational consequence relation, i.e., satisfies all the rules of preferential reasoning, *plus* Rational Monotonicity. However, this last named rule is not universally accepted, and other researchers have advocated working within the larger class of *disjunctive* consequence relations, which satisfy the weaker requirement of *Disjunctive Rationality*. While there are convincing arguments that the rational closure forms the "simplest" rational consequence relation extending a given set of conditionals, the question of what is the simplest *disjunctive* consequence relation has not been explored. In this paper, we propose a solution to this question and explore some of its properties.

1 Introduction

The question of *conditional inference*, i.e., of which conditional sentences of the form "if α then, normally, β" should follow from a set $\mathcal{KB}$ of such sentences, has been one of the classic questions of non-monotonic reasoning, with several well-known solutions proposed. Since the work of Lehmann and colleagues in the early '90s, the so-called preferential approach to defeasible reasoning has established itself as one of the most elegant frameworks within which to answer this question. Central to the preferential approach is the notion of *rational closure* of a conditional knowledge base, under which the set of inferred conditionals forms a rational consequence relation, i.e., satisfies all the rules of preferential reasoning, *plus* Rational Monotonicity. One of the reasons for accepting rational closure is the fact it delivers a venturous notion of entailment that is conservative enough. Given that, rationality has for long been accepted as the core baseline for any appropriate form of non-monotonic entailment.

Very few have stood against this position, including Makinson (1994), who considered Rational Monotonicity too strong and has briefly advocated the weaker rule of Disjunctive Rationality instead. This rule is implied by Rational Monotonicity and may still be desirable in cases where the latter does not hold. Quite surprisingly, the debate did not catch on, and, for lack of rivals of the same stature, Rational Closure has since reigned alone as a role model in non-monotonic inference. That used to be the case until Rott (2014) reignited interest in Disjunctive Rationality by considering interval models in connection with belief contraction. Inspired by that, here we revisit disjunctive consequence relations and make the first steps in the quest for a suitable notion of disjunctive rational closure of a conditional knowledge base.

The plan of the paper is as follows. First, in Section 2, we give the usual summary of the formal background assumed in the following sections, in particular of the rational closure construction. Then, in Section 3, we make a case for weakening the rationality requirement and propose a semantics with an accompanying representation result for a weaker form of rationality enforcing the rule of Disjunctive Rationality. We move on, and in Section 4, we investigate a notion of closure of (or entailment from) a conditional knowledge base under Disjunctive Rationality. Our analysis is in terms of a set of postulates, all reasonable at first glance, that one can expect a suitable notion of closure to satisfy. Following that, in Section 5, we propose a specific construction for the Disjunctive Rational Closure of a conditional knowledge base and assess its suitability in the light of the postulates previously put forward (Section 6). We conclude with some remarks on future directions of investigation.

2 Formal preliminaries

In this section, we provide the required formal background for the remainder of the paper. In particular, we set up the notation and conventions that shall be followed in the upcoming sections. (The reader conversant with the KLM framework for non-monotonic reasoning can safely skip to Section 3.)

Let $\mathcal{P}$ be a finite set of propositional *atoms*. We use $p, q, \ldots$ as meta-variables for atoms. Propositional sentences are denoted by $\alpha, \beta, \ldots$, and are recursively defined in the usual way:

$$\alpha ::= \top \mid \bot \mid \mathcal{P} \mid \neg\alpha \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \alpha \rightarrow \alpha \mid \alpha \leftrightarrow \alpha$$

We use $\mathcal{L}$ to denote the set of all propositional sentences.

With $\mathcal{U} \stackrel{\text{def}}{=} \{0,1\}^{\mathcal{P}}$, we denote the set of all propositional *valuations*, where 1 represents truth and 0 falsity. We use

$v, u, \ldots$, possibly with primes, to denote valuations. Whenever it eases the presentation, we shall 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-barred 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 $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$, where the atoms stand for, respectively, "being a bird", "being a flying creature", and "being a penguin", the valuation in which b is true, f is false, and p is true will be represented as $\mathsf{b\bar{f}p}$.

Satisfaction of a sentence $\alpha \in \mathcal{L}$ by a valuation $v \in \mathcal{U}$ is defined in the usual truth-functional way and is denoted by $v \Vdash \alpha$. The set of *models* of a sentence α is defined as $[\![\alpha]\!] \stackrel{\text{def}}{=} \{v \in \mathcal{U} \mid v \Vdash \alpha\}$. This notion is extended to a set of sentences X in the usual way: $[\![X]\!] \stackrel{\text{def}}{=} \bigcap_{\alpha \in X} [\![\alpha]\!]$. We say a set of sentences X (classically) *entails* $\alpha \in \mathcal{L}$, denoted $X \models \alpha$, if $[\![X]\!] \subseteq [\![\alpha]\!]$. α is *valid*, denoted $\models \alpha$, if $[\![\alpha]\!] = \mathcal{U}$.

2.1 KLM-style rational defeasible consequence

Several approaches to non-monotonic reasoning have been proposed in the literature over the past 40 years. The *preferential approach*, initially put forward by Shoham (1988) and subsequently developed by Kraus et al. (1990) in much depth (the reason why it became known as the KLM-approach), has established itself as one of the main references in the area. This stems from at least three of its features: (*i*) its intuitive semantics and elegant proof-theoretic characterisation; (*ii*) its generality w.r.t. alternative approaches to non-monotonic reasoning such as circumscription (McCarthy 1980), default logic (Reiter 1980), and many others, and (*iii*) its formal links with AGM-style belief revision (Gärdenfors and Makinson 1994). The fruitfulness of the preferential approach is also witnessed by the great deal of recent work extending it to languages that are more expressive than that of propositional logic such as those of description logics (Bonatti 2019; Britz, Meyer, and Varzinczak 2011; Casini et al. 2015; Britz and Varzinczak 2017; Giordano et al. 2007; Giordano et al. 2015; Pensel and Turhan 2017; Varzinczak 2018), modal logics (Britz and Varzinczak 2018a; Britz and Varzinczak 2018b; Chafik et al. 2020), and others (Booth, Meyer, and Varzinczak 2012).

A *defeasible consequence relation* $\mid\!\sim$ is defined as a binary relation on sentences of our underlying propositional logic, i.e., $\mid\!\sim \subseteq \mathcal{L} \times \mathcal{L}$. We say that $\mid\!\sim$ is a *preferential* consequence relation (Kraus, Lehmann, and Magidor 1990) if it satisfies the following set of (Gentzen-style) rules:

$$\text{(Ref)} \quad \alpha \mid\!\sim \alpha \qquad\qquad \text{(LLE)} \quad \frac{\models \alpha \leftrightarrow \beta,\ \alpha \mid\!\sim \gamma}{\beta \mid\!\sim \gamma}$$

$$\text{(And)} \quad \frac{\alpha \mid\!\sim \beta,\ \alpha \mid\!\sim \gamma}{\alpha \mid\!\sim \beta \wedge \gamma} \qquad\qquad \text{(Or)} \quad \frac{\alpha \mid\!\sim \gamma,\ \beta \mid\!\sim \gamma}{\alpha \vee \beta \mid\!\sim \gamma}$$

$$\text{(RW)} \quad \frac{\alpha \mid\!\sim \beta,\ \models \beta \rightarrow \gamma}{\alpha \mid\!\sim \gamma} \qquad\qquad \text{(CM)} \quad \frac{\alpha \mid\!\sim \beta,\ \alpha \mid\!\sim \gamma}{\alpha \wedge \beta \mid\!\sim \gamma}$$

If, in addition to the preferential rules, the defeasible consequence relation $\mid\!\sim$ also satisfies the following Rational Monotonicity rule (Lehmann and Magidor 1992), it is said to be a *rational* consequence relation:

$$\text{(RM)} \quad \frac{\alpha \mid\!\sim \beta,\ \alpha \not\mid\!\sim \neg\gamma}{\alpha \wedge \gamma \mid\!\sim \beta}$$

Rational consequence relations can be given an intuitive semantics in terms of *ranked interpretations*.

Definition 1. *A* ***ranked interpretation*** $\mathscr{R}$ *is a function from* $\mathcal{U}$ *to* $\mathbb{N} \cup \{\infty\}$ *such that* $\mathscr{R}(v) = 0$ *for some* $v \in \mathcal{U}$*, and satisfying the following* ***convexity property****: for every* $i \in \mathbb{N}$*, if* $\mathscr{R}(u) = i$*, then, for every* j *s.t.* $0 \leq j < i$*, there is a* $u' \in \mathcal{U}$ *for which* $\mathscr{R}(u') = j$.

In a ranked interpretation, we call $\mathscr{R}(v)$ the *rank of* v w.r.t. $\mathscr{R}$. 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 'forbidden', e.g. by some background knowledge—see below. Given a ranked interpretation $\mathscr{R}$, we therefore partition the set $\mathcal{U}$ into the set of *plausible* valuations (those with finite rank), and that of *implausible* ones (with rank ∞).[1]

Figure 1 depicts an example of a ranked interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$. (In our graphical representations of ranked interpretations—and of interval-based interpretations later on—we shall plot the set of valuations in $\mathcal{U}$ on the y-axis so that the preference relation reads more naturally across the x-axis—from lower to higher. Moreover, plausible valuations are associated with the colour blue, whereas the implausible ones with red.)

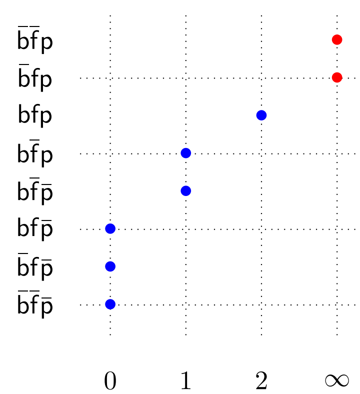

Figure 1: A ranked interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$.

Given a ranked interpretation $\mathscr{R}$ and $\alpha \in \mathcal{L}$, with $[\![\alpha]\!]^{\mathscr{R}}$ we denote the set of plausible valuations satisfying α (α-*valuations* for short) in $\mathscr{R}$. If $[\![\alpha]\!]^{\mathscr{R}} = [\![\top]\!]^{\mathscr{R}}$, then we say α is *true* in $\mathscr{R}$ and denote it $\mathscr{R} \Vdash \alpha$. With $\mathscr{R}(\alpha) \stackrel{\text{def}}{=} \min\{\mathscr{R}(v) \mid$

[1] In the literature, it is customary to omit implausible valuations from ranked interpretations. Since they are not logically impossible, but rather judged as irrelevant on the grounds of contingent information (e.g. a knowledge base) which is prone to change, we shall include them in our semantic definitions. This does not mean that we do anything special with them in this paper; they are rather kept for future use.

$v \in [\![\alpha]\!]^{\mathscr{R}}\}$ we denote the *rank of* α *in* $\mathscr{R}$. By convention, if $[\![\alpha]\!]^{\mathscr{R}} = \emptyset$, we let $\mathscr{R}(\alpha) = \infty$. Defeasible consequence of the form $\alpha \mid\!\sim \beta$ is then given a semantics in terms of ranked interpretations in the following way: We say $\alpha \mid\!\sim \beta$ is *satisfied in* $\mathscr{R}$ (denoted $\mathscr{R} \Vdash \alpha \mid\!\sim \beta$) if $\mathscr{R}(\alpha) < \mathscr{R}(\alpha \land \neg\beta)$. (And here we adopt Jaeger's (1996) convention that $\infty < \infty$ always holds.) It is easy to see that for every $\alpha \in \mathcal{L}$, $\mathscr{R} \Vdash \alpha$ if and only if $\mathscr{R} \Vdash \neg\alpha \mid\!\sim \bot$. If $\mathscr{R} \Vdash \alpha \mid\!\sim \beta$, we say $\mathscr{R}$ is a *ranked model* of $\alpha \mid\!\sim \beta$. In the example in Figure 1, we have $\mathscr{R} \Vdash \mathsf{b} \mid\!\sim \mathsf{f}$, $\mathscr{R} \Vdash \mathsf{p} \to \mathsf{b}$ (and therefore $\mathscr{R} \Vdash \neg(\mathsf{p} \to \mathsf{b}) \mid\!\sim \bot$), $\mathscr{R} \Vdash \mathsf{p} \mid\!\sim \neg\mathsf{f}$, $\mathscr{R} \not\Vdash \mathsf{f} \mid\!\sim \mathsf{b}$, and $\mathscr{R} \Vdash \mathsf{p} \land \neg\mathsf{b} \mid\!\sim \mathsf{b}$, which are all according to the intuitive expectations.

That this semantic characterisation of rational defeasible consequence is appropriate is a consequence of a representation result linking the seven rationality rules above to precisely the class of ranked interpretations (Lehmann and Magidor 1992; Gärdenfors and Makinson 1994).

2.2 Rational closure

One can also view defeasible consequence as formalising some form of (defeasible) conditional and bring it down to the level of statements. Such was the stance adopted by Lehmann and Magidor (1992). A *conditional knowledge base* $\mathcal{KB}$ is thus a finite set of statements of the form $\alpha \mid\!\sim \beta$, with $\alpha, \beta \in \mathcal{L}$, and possibly containing classical statements. As an example, let $\mathcal{KB} = \{\mathsf{b} \mid\!\sim \mathsf{f}, \mathsf{p} \to \mathsf{b}, \mathsf{p} \mid\!\sim \neg\mathsf{f}\}$. Given a conditional knowledge base $\mathcal{KB}$, a *ranked model* of $\mathcal{KB}$ is a ranked interpretation satisfying all statements in $\mathcal{KB}$. As it turns out, the ranked interpretation in Figure 1 is a ranked model of the above $\mathcal{KB}$. It is not hard to see that, in every ranked model of $\mathcal{KB}$, the valuations $\bar{\mathsf{b}}\bar{\mathsf{f}}\mathsf{p}$ and $\bar{\mathsf{b}}\mathsf{f}\mathsf{p}$ are deemed implausible—note, however, that they are still *logically* possible, which is the reason why they feature in all ranked interpretations.

An important reasoning task in this setting is that of determining which conditionals follow from a conditional knowledge base. Of course, even when interpreted as a conditional in (and under) a given knowledge base $\mathcal{KB}$, $\mid\!\sim$ is expected to adhere to the rules of Section 2.1. Intuitively, that means whenever appropriate instantiations of the premises in a rule are sanctioned by $\mathcal{KB}$, so should the suitable instantiation of its conclusion.

To be more precise, we can take the defeasible conditionals in $\mathcal{KB}$ as the core elements of a defeasible consequence relation $\mid\!\sim^{\mathcal{KB}}$. By closing the latter under the preferential rules (in the sense of exhaustively applying them), we get a *preferential extension* of $\mid\!\sim^{\mathcal{KB}}$. Since there can be more than one such extension, the most cautious approach consists in taking their intersection. The resulting set, which also happens to be closed under the preferential rules, is the *preferential closure* of $\mid\!\sim^{\mathcal{KB}}$, which we denote by $\mid\!\sim^{\mathcal{KB}}_{PC}$. When interpreted again as a conditional knowledge base, the preferential closure of $\mid\!\sim^{\mathcal{KB}}$ contains all the conditionals entailed by $\mathcal{KB}$. (Hence, the notions of closure of and entailment from a conditional knowledge base are two sides of the same coin.) The same process and definitions carry over when one requires the defeasible consequence relations also to be closed under the rule RM, in which case we talk of *rational* extensions of $\mid\!\sim^{\mathcal{KB}}$. Nevertheless, as pointed out by Lehmann and Magidor (1992, Section 4.2), the intersection of all such rational extensions is not, in general, a rational consequence relation: it coincides with preferential closure and therefore may fail RM. Among other things, this means that the corresponding entailment relation, which is called *rank entailment* and defined as $\mathcal{KB} \models_{\mathscr{R}} \alpha \mid\!\sim \beta$ if every ranked model of $\mathcal{KB}$ also satisfies $\alpha \mid\!\sim \beta$, is *monotonic* and therefore it falls short of being a suitable form of entailment in a defeasible reasoning setting. As a result, several alternative notions of entailment from conditional knowledge bases have been explored in the literature on non-monotonic reasoning (Booth and Paris 1998; Booth et al. 2019; Casini, Meyer, and Varzinczak 2019; Giordano et al. 2012; Giordano et al. 2015; Lehmann 1995; Weydert 2003), with *rational closure* (Lehmann and Magidor 1992) commonly acknowledged as the gold standard in the matter.

Rational closure (RC) is a form of inferential closure extending the notion of rank entailment above. It formalises the principle of *presumption of typicality* (Lehmann 1995, p. 63), which, informally, specifies that a situation (in our case, a valuation) should be assumed to be as typical as possible (w.r.t. background information in a knowledge base).

Assume an ordering $\preceq_{\mathcal{KB}}$ on all ranked models of a knowledge base $\mathcal{KB}$, which is defined as follows: $\mathscr{R}_1 \preceq_{\mathcal{KB}} \mathscr{R}_2$, if, for every $v \in \mathcal{U}$, $\mathscr{R}_1(v) \leq \mathscr{R}_2(v)$. Intuitively, ranked models lower down in the ordering are more typical. It is easy to see that $\preceq_{\mathcal{KB}}$ is a weak partial order. Giordano et al. (2015) showed that there is a unique $\preceq_{\mathcal{KB}}$-minimal element. The rational closure of $\mathcal{KB}$ is defined in terms of this minimum ranked model of $\mathcal{KB}$.

Definition 2. *Let $\mathcal{KB}$ be a conditional knowledge base, and let $\mathscr{R}^{\mathcal{KB}}_{RC}$ be the minimum element of the ordering $\preceq_{\mathcal{KB}}$ on ranked models of $\mathcal{KB}$. The* ***rational closure*** *of $\mathcal{KB}$ is the defeasible consequence relation* $\mid\!\sim^{\mathcal{KB}}_{RC} \stackrel{\text{def}}{=} \{\alpha \mid\!\sim \beta \mid \mathscr{R}^{\mathcal{KB}}_{RC} \Vdash \alpha \mid\!\sim \beta\}$.

As an example, Figure 1 shows the minimum ranked model of $\mathcal{KB} = \{\mathsf{b} \mid\!\sim \mathsf{f}, \mathsf{p} \to \mathsf{b}, \mathsf{p} \mid\!\sim \neg\mathsf{f}\}$ w.r.t. $\preceq_{\mathcal{KB}}$. Hence we have that $\neg\mathsf{f} \mid\!\sim \neg\mathsf{b}$ is in the rational closure of $\mathcal{KB}$.

Observe that there are two levels of typicality at work for rational closure, namely *within* ranked models of $\mathcal{KB}$, where valuations lower down are viewed as more typical, but also *between* ranked models of $\mathcal{KB}$, where ranked models lower down in the ordering are viewed as more typical. The most typical ranked model $\mathscr{R}^{\mathcal{KB}}_{RC}$ is the one in which valuations are as typical as $\mathcal{KB}$ allows them to be (the principle of presumption of typicality we alluded to above).

Rational closure is commonly viewed as the *basic* (although certainly not the only acceptable) form of nonmonotonic entailment, on which other, more venturous forms of entailment can be and have been constructed (Booth et al. 2019; Casini et al. 2014; Casini, Meyer, and Varzinczak 2019; Lehmann 1995).

3 Disjunctive rationality and interval-based preferential semantics

One may argue that there are cases in which Rational Monotonicity is too strong a rule to enforce and for which a weaker defeasible consequence relation would suffice (Giordano et al. 2010; Makinson 1994). Nevertheless, doing away completely with rationality (i.e., sticking to the preferential rules only) is not particularly appropriate in a defeasible-reasoning context. Indeed, as widely known in the literature, preferential systems induce entailment relations that are monotonic. In that respect, here we are interested in defeasible consequence relations (or defeasible conditionals) that do not necessarily satisfy Rational Monotonicity while still encapsulating some form of rationality, i.e., a venturous passage from the premises to the conclusion. A case in point is that of the Disjunctive Rationality (DR) rule (Kraus, Lehmann, and Magidor 1990) below:

$$\text{(DR)} \quad \frac{\alpha \vee \beta \mathrel{|\!\sim} \gamma}{\alpha \mathrel{|\!\sim} \gamma \ \text{ or } \ \beta \mathrel{|\!\sim} \gamma}$$

Intuitively, DR says that if one may draw a conclusion from a disjunction of premises, then one should be able to draw this conclusion from at least one of these premises taken alone (Freund 1993). A preferential consequence relation is called *disjunctive* if it also satisfies DR.

As it turns out, every rational consequence relation is also disjunctive, but not the other way round (Lehmann and Magidor 1992). Therefore, DR is a weaker form of rationality, as its name suggests. Given that, Disjunctive Rationality is indeed a suitable candidate for the type of investigation we have in mind.

A semantic characterisation of disjunctive consequence relations was given by Freund (1993) based on a filtering condition on the underlying ordering. Here, we provide an alternative semantics in terms of *interval-based interpretations*. (We conjecture Freund's semantic constructions and ours can be shown to be equivalent in the finite case.)

Definition 3. *An **interval-based interpretation** is a tuple $\mathscr{I} \stackrel{\text{def}}{=} \langle \mathscr{L}, \mathscr{U} \rangle$, where $\mathscr{L}$ and $\mathscr{U}$ are functions from $\mathcal{U}$ to $\mathbb{N} \cup \{\infty\}$ s.t. (i) $\mathscr{L}(v) = 0$, for some $v \in \mathcal{U}$; (ii) if $\mathscr{L}(u) = i$ or $\mathscr{U}(u) = i$, then for every $0 \leq j < i$, there is u' s.t. either $\mathscr{L}(u') = j$ or $\mathscr{U}(u') = j$, (iii) $\mathscr{L}(v) \leq \mathscr{U}(v)$, for every $v \in \mathcal{U}$, and (iv) $\mathscr{L}(u) = \infty$ iff $\mathscr{U}(u) = \infty$. Given $\mathscr{I} = \langle \mathscr{L}, \mathscr{U} \rangle$ and $v \in \mathcal{U}$, $\mathscr{L}(v)$ is the **lower rank of** v **in** $\mathscr{I}$, and $\mathscr{U}(v)$ is the **upper rank of** v **in** $\mathscr{I}$. Hence, for any v, the pair $(\mathscr{L}(v), \mathscr{U}(v))$ is the **interval of** v **in** $\mathscr{I}$. We say u is **more preferred than** v **in** $\mathscr{I}$, denoted $u \prec v$, if $\mathscr{U}(u) < \mathscr{L}(v)$.*

The preference order $\prec$ on $\mathcal{U}$ defined above via an interval-based interpretation forms an *interval order*, i.e., it is a strict partial order that additionally satisfies the *interval condition*: if $u \prec v$ and $u' \prec v'$, then either $u \prec v'$ or $u' \prec v$. Furthermore, every interval order can be defined from an interval-based interpretation in this way. See the work of Fishburn (1985) for a detailed treatise on interval orders.

Figure 2 illustrates an example of an interval-based interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$. In our depictions of interval-based interpretations, it will be convenient to see $\mathscr{I}$ as a function from $\mathcal{U}$ to intervals on the set $\mathbb{N} \cup \{\infty\}$. Whenever the intervals associated to valuations u and v overlap, the intuition is that both valuations are incomparable in $\mathscr{I}$; otherwise the leftmost interval is seen as more preferred than the rightmost one.

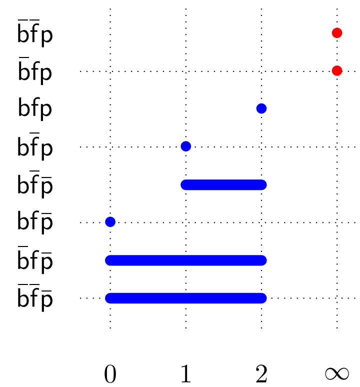

Figure 2: An interval-based interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$.

In Figure 2, the rationale behind the ordering is as follows: situations with flying birds are the most normal ones; situations with non-flying penguins are more normal than the flying-penguin ones, but both are incomparable to non-penguin situations; the situations with penguins that are not birds are the implausible ones; and finally those that are not about birds or penguins are so irrelevant as to be seen as incomparable with any of the plausible ones.

The notions of plausible and implausible valuations, as well as that of α-valuations, carry over to interval-based interpretations, only now the plausible valuations are the ones with finite lower ranks (and hence also finite upper ranks, by part *(iv)* of the previous definition). With $\mathscr{L}(\alpha) \stackrel{\text{def}}{=} \min\{\mathscr{L}(v) \mid v \in [\![\alpha]\!]^{\mathscr{I}}\}$ and $\mathscr{U}(\alpha) \stackrel{\text{def}}{=} \min\{\mathscr{U}(v) \mid v \in [\![\alpha]\!]^{\mathscr{I}}\}$ we denote, respectively, the *lower* and the *upper rank of* α *in* $\mathscr{I}$. By convention, if $[\![\alpha]\!]^{\mathscr{I}} = \emptyset$, we let $\mathscr{L}(\alpha) = \mathscr{U}(\alpha) = \infty$. We say $\alpha \mathrel{|\!\sim} \beta$ is *satisfied in* $\mathscr{I}$ (denoted $\mathscr{I} \Vdash \alpha \mathrel{|\!\sim} \beta$) if $\mathscr{U}(\alpha) < \mathscr{L}(\alpha \wedge \neg\beta)$. (Recall the convention that $\infty < \infty$.) As an example, in the interval-based interpretation of Figure 2, we have $\mathscr{I} \Vdash \mathsf{b} \mathrel{|\!\sim} \mathsf{f}$, $\mathscr{I} \Vdash \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}$, and $\mathscr{I} \not\Vdash \neg\mathsf{f} \mathrel{|\!\sim} \neg\mathsf{p}$ (contrary to the ranked interpretation $\mathscr{R}$ in Figure 1, which endorses the latter statement).

In the tradition of the KLM approach to defeasible reasoning, we define the defeasible consequence relation induced by an interval-based interpretation: $\mathrel{|\!\sim}_{\mathscr{I}} \stackrel{\text{def}}{=} \{\alpha \mathrel{|\!\sim} \beta \mid \mathscr{I} \Vdash \alpha \mathrel{|\!\sim} \beta\}$. We can now state a KLM-style representation result establishing that our interval-based semantics is suitable for characterising the class of disjunctive defeasible consequence relations, which is a variant of Freund's (1993) result:

Theorem 1. *A defeasible consequence relation is a disjunctive consequence relation if and only if it is defined by some* *interval-based interpretation, i.e., $\mid\!\sim$ is disjunctive if and only if there is $\mathscr{I}$ such that $\mid\!\sim \, = \, \mid\!\sim_{\mathscr{I}}$.*

4 Towards disjunctive rational closure

Given a conditional knowledge base $\mathcal{KB}$, the obvious definition of closure under Disjunctive Rationality consists in taking the intersection of all *disjunctive extensions* of $\mid\!\sim^{\mathcal{KB}}$ (cf. Section 2.2). Let us call it the *disjunctive closure* of $\mid\!\sim^{\mathcal{KB}}$, with *interval-based entailment*, defined as $\mathcal{KB} \models_{\mathscr{I}} \alpha \mid\!\sim \beta$ if every interval-based model of $\mathcal{KB}$ also satisfies $\alpha \mid\!\sim \beta$, being its semantic counterpart. The following result shows that the notion of disjunctive closure is stillborn, i.e., it does not even satisfy Disjunctive Rationality.

Theorem 2. *Given a conditional knowledge base $\mathcal{KB}$, (i) the disjunctive closure of $\mathcal{KB}$ coincides with its preferential closure $\mid\!\sim^{\mathcal{KB}}_{PC}$. (ii) There exists $\mathcal{KB}$ such that $\mid\!\sim^{\mathcal{KB}}_{PC}$ does not satisfy Disjunctive Rationality.*

For a simple counterexample showing that $\mid\!\sim^{\mathcal{KB}}_{PC}$ need not satisfy Disjunctive Rationality, consider $\mathcal{KB} = \{\top \mid\!\sim \mathsf{b}\}$. Clearly we have $\mathsf{p} \lor \neg\mathsf{p} \mid\!\sim^{\mathcal{KB}}_{PC} \mathsf{b}$, but one can easily construct interval-based interpretations $\mathscr{I}_1$, $\mathscr{I}_2$ whose corresponding consequence relations both satisfy $\mathcal{KB}$ but for which $\mathsf{p} \not\mid\!\sim_{\mathscr{I}_1} \mathsf{b}$ and $\neg\mathsf{p} \not\mid\!\sim_{\mathscr{I}_2} \mathsf{b}$.

This result suggests that the quest for a suitable definition of entailment under disjunctive rationality should follow the footprints in the road which led to the definition of rational closure. Such is our contention here, and our research question is now: 'Is there a single best disjunctive relation extending the one induced by a given conditional knowledge base $\mathcal{KB}$?'

Let us denote by $\mid\!\sim^{\mathcal{KB}}_{*}$ the special defeasible consequence relation that we are looking for. In the remainder of this section, we consider some desirable properties for the mapping from $\mathcal{KB}$ to $\mid\!\sim^{\mathcal{KB}}_{*}$, and consider some simple examples in order to build intuitions. In the following section, we will offer a concrete construction: the Disjunctive Rational Closure of $\mathcal{KB}$.

4.1 Basic postulates

Starting with our most basic requirements, we put forward the following two postulates:

Inclusion If $\alpha \mid\!\sim \beta \in \mathcal{KB}$, then $\alpha \mid\!\sim^{\mathcal{KB}}_{*} \beta$.

D-Rationality $\mid\!\sim^{\mathcal{KB}}_{*}$ is a disjunctive consequence relation.

Note that, given Theorem 1, D-Rationality is equivalent to saying there is an interval-based interpretation $\mathscr{I}$ such that $\mid\!\sim^{\mathcal{KB}}_{*} = \, \mid\!\sim_{\mathscr{I}}$. If we replace "disjunctive consequence" in the statement by "rational consequence", then that is the postulate that is usually considered in the area.

Another reasonable property to require from an induced consequence relation is for two equivalent knowledge bases to yield exactly the same set of inferences. This prompts the question of what it means to say that two conditional knowledge bases are equivalent. One weak notion of equivalence can be defined as follows.

Definition 4. *Let $\alpha_1, \alpha_2, \beta_1, \beta_2 \in \mathcal{L}$. We say $\alpha_1 \mid\!\sim \beta_1$ is **equivalent** to $\alpha_2 \mid\!\sim \beta_2$ if $\models (\alpha_1 \leftrightarrow \alpha_2) \land (\beta_1 \leftrightarrow \beta_2)$. We say two knowledge bases $\mathcal{KB}_1$, $\mathcal{KB}_2$ are **equivalent**, written $\mathcal{KB}_1 \equiv \mathcal{KB}_2$, if there is a bijection $f : \mathcal{KB}_1 \longrightarrow \mathcal{KB}_2$ s.t. each $\alpha \mid\!\sim \beta \in \mathcal{KB}_1$ is equivalent to $f(\alpha \mid\!\sim \beta)$.*

Given this, we can express a weak form of syntax independence:

Equivalence If $\mathcal{KB}_1 \equiv \mathcal{KB}_2$, then $\mid\!\sim^{\mathcal{KB}_1}_{*} = \mid\!\sim^{\mathcal{KB}_2}_{*}$.

Finally, the last of our basic postulates requires rational closure to be the upper bound on how venturous our consequence relation should be.

Infra-Rationality $\mid\!\sim^{\mathcal{KB}}_{*} \subseteq \, \mid\!\sim^{\mathcal{KB}}_{RC}$.

4.2 Minimality postulates

Echoing a fundamental principle of reasoning in general and of non-monotonic reasoning in particular is a property requiring $\mid\!\sim^{\mathcal{KB}}_{*}$ to contain only conditionals whose inferences can be *justified* on the basis of $\mathcal{KB}$. The first idea to achieve this would be to set $\mid\!\sim^{\mathcal{KB}}_{*}$ to be a set-theoretically *minimal* disjunctive consequence relation that extends $\mathcal{KB}$.

Example 1. *Suppose the only knowledge we have is a single conditional saying "birds normally fly", i.e., $\mathcal{KB} = \{\mathsf{b} \mid\!\sim \mathsf{f}\}$. Assuming just two variables, we have a unique $\subseteq$-minimal disjunctive consequence relation extending this knowledge base, which is given by the interval-based interpretation $\mathscr{I}$ in Figure 3. Indeed, the conditional $\mathsf{b} \mid\!\sim \mathsf{f}$ is saying precisely that $\mathsf{bf} \prec \mathsf{b\bar{f}}$, but is telling us nothing with regard to the relative typicality of the other two possible valuations, so any pair of valuations other than this one is incomparable. For this reason, we do not have $\neg\mathsf{f} \mid\!\sim_{\mathscr{I}} \neg\mathsf{b}$ here. Note the rational closure in this example* does *endorse this latter conclusion, thus providing further evidence that the rational closure arguably gives some unwarranted conclusions.*

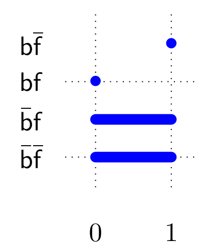

Figure 3: Interval-based model of $\mathcal{KB} = \{\mathsf{b} \mid\!\sim \mathsf{f}\}$.

The next example illustrates the fact that there might be more than one $\subseteq$-minimal extension of a $\mathcal{KB}$-induced consequence relation.

Example 2. *Assume a COVID-19 inspired scenario with only two propositions, m and s, standing for, respectively, "you wear a $\mathsf{m}$ask" and "you observe $\mathsf{s}$ocial distancing". Let $\mathcal{KB} = \{\mathsf{m} \mid\!\sim \mathsf{s}, \neg\mathsf{m} \mid\!\sim \mathsf{s}\}$. There are two $\subseteq$-minimal disjunctive consequence relations extending $\mid\!\sim^{\mathcal{KB}}$, corresponding to the two interval-based interpretations $\mathscr{I}_1$ and $\mathscr{I}_2$ (from left to right) in Figure 4. The first conditional is saying $\mathsf{ms} \prec \mathsf{m\bar{s}}$, while the second is saying $\mathsf{\bar{m}s} \prec \mathsf{\bar{m}\bar{s}}$. According* *to the interval condition (see the paragraph following Definition 3), we must then have either* $\mathsf{ms} \prec \bar{\mathsf{m}}\bar{\mathsf{s}}$ *or* $\bar{\mathsf{m}}\mathsf{s} \prec \mathsf{m}\bar{\mathsf{s}}$*. The choice of which gives rise to* $\mathscr{I}_1$ *and* $\mathscr{I}_2$*, respectively.*

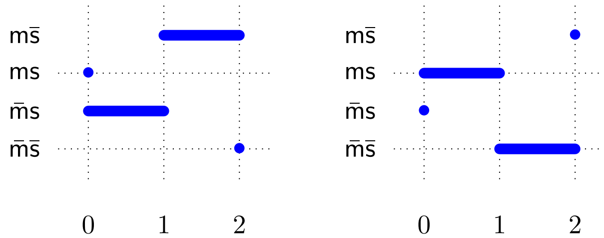

Figure 4: Interval-based models of the two $\subseteq$-minimal extensions of $\mid\!\sim^{\mathcal{KB}}$, for $\mathcal{KB} = \{\mathsf{m} \mid\!\sim \mathsf{s}, \neg\mathsf{m} \mid\!\sim \mathsf{s}\}$.

In the light of Example 2 above, a question that arises is what to do when one has more than a single $\subseteq$-minimal extension of $\mid\!\sim^{\mathcal{KB}}$. Theorem 2 already tells us we cannot, in general, take the obvious approach by taking their intersection. However, even though returning the disjunctive/preferential closure $\mid\!\sim^{\mathcal{KB}}_{PC}$ is not enough to ensure D-Rationality, we might still expect the following postulates as reasonable.

Vacuity If $\mid\!\sim^{\mathcal{KB}}_{PC}$ is a disjunctive consequence relation, then $\mid\!\sim^{\mathcal{KB}}_{*} = \mid\!\sim^{\mathcal{KB}}_{PC}$.

Preferential Extension $\mid\!\sim^{\mathcal{KB}}_{PC} \subseteq \mid\!\sim^{\mathcal{KB}}_{*}$.

(Note, given Theorem 2, the postulate above follows from Inclusion and D-Rationality.)

Justification If $\alpha \mid\!\sim^{\mathcal{KB}}_{*} \beta$, then $\alpha \mid\!\sim' \beta$ for at least one $\subseteq$-minimal disjunctive relation $\mid\!\sim'$ extending $\mid\!\sim^{\mathcal{KB}}$.

4.3 Representation independence postulates

Going back to Example 2, what should the expected output be in this case? Intuitively, faced with the choice of which of the pairs $\mathsf{ms} \prec \bar{\mathsf{m}}\bar{\mathsf{s}}$ or $\bar{\mathsf{m}}\mathsf{s} \prec \mathsf{m}\bar{\mathsf{s}}$ to include, and in the absence of any reason to prefer either one, it seems the right thing to do is to include both, and thereby let the interval-based interpretation depicted in Figure 5 yield the output. Notice that this will be the same as the rational closure in this case.

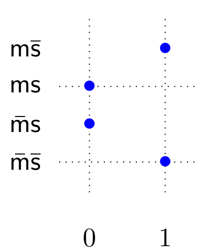

Figure 5: Interval-based models of the union of the two $\subseteq$-minimal extensions of $\mid\!\sim^{\mathcal{KB}}$, for $\mathcal{KB} = \{\mathsf{m} \mid\!\sim \mathsf{s}, \neg\mathsf{m} \mid\!\sim \mathsf{s}\}$.

We can express the desired symmetry requirement in a syntactic form, using the notion of *symbol translations* (Marquis and Schwind 2014). A symbol translation (on $\mathcal{P}$) is a function $\sigma : \mathcal{P} \longrightarrow \mathcal{L}$. A symbol translation can be extended to a function on $\mathcal{L}$ by setting, for each sentence α, $\sigma(\alpha)$ to be the sentence obtained from α by replacing each atom p occurring in α by its image $\sigma(p)$ throughout.[2] Similarly, given a conditional knowledge base $\mathcal{KB}$ and a symbol translation $\sigma(\cdot)$, we denote by $\sigma(\mathcal{KB})$ the knowledge base obtained by replacing each conditional $\alpha \mid\!\sim \beta$ in $\mathcal{KB}$ by $\sigma(\alpha) \mid\!\sim \sigma(\beta)$.

Representation Independence For any symbol translation $\sigma(\cdot)$, we have $\alpha \mid\!\sim^{\mathcal{KB}}_{*} \beta$ iff $\sigma(\alpha) \mid\!\sim^{\sigma(\mathcal{KB})}_{*} \sigma(\beta)$.

Note that Weydert (2003) also considers Representation Independence (RI) in the context of conditional inference, but in a slightly different framework. The idea behind it has also been explored by Jaeger (1996), who, in particular, looked at the property in relation to rational closure. As noted by Marquis and Schwind (2014), the property is a very demanding one that is likely hard to satisfy in its full, unrestricted form above. And indeed this is confirmed in our setting, since it can be shown that Representation Independence is jointly incompatible with two of our basic postulates, namely Inclusion and Infra-Rationality. This motivates the need to focus on specific families of symbol translation. Some examples are the following:

1. $\sigma(\cdot)$ is a permutation on $\mathcal{P}$, i.e., is just a *renaming* of the propositional variables;
2. $\sigma(p) \in \{p, \neg p\}$, for all $p \in \mathcal{P}$. Then, instead of using p to denote say "it's raining", we use it rather to denote "it's not raining". We call any symbol translation of this type a *negation-swapping* symbol translation.

Each special subfamily of symbol translations yields a corresponding weakening of RI that applies to just that kind of translation. In particular we have the following postulate:

Negated Representation Independence For any negation-swapping symbol translation $\sigma(\cdot)$, we have $\alpha \mid\!\sim^{\mathcal{KB}}_{*} \beta$ iff $\sigma(\alpha) \mid\!\sim^{\sigma(\mathcal{KB})}_{*} \sigma(\beta)$.

Example 3. *Going back to Example 2, when modelling the scenario, instead of using propositional atom* m *to denote "you wear a mask" we could equally well have used it to denote "you do not wear a mask". Then the statement "if you wear a mask then, normally, you do social distancing" would be modelled by* $\neg\mathsf{m} \mid\!\sim \mathsf{s}$*, etc. This boils down to taking a negation-swapping symbol translation such that* $\sigma(\mathsf{m}) = \neg\mathsf{m}$ *and* $\sigma(\mathsf{s}) = \mathsf{s}$*. Then* $\sigma(\mathcal{KB}) = \{\neg\mathsf{m} \mid\!\sim \mathsf{s}, \neg\neg\mathsf{m} \mid\!\sim \mathsf{s}\}$*, and if we inferred, say,* $\mathsf{m} \leftrightarrow \mathsf{s} \mid\!\sim \mathsf{s}$ *from* $\mathcal{KB}$ *then we would expect to infer* $\neg\mathsf{m} \leftrightarrow \mathsf{s} \mid\!\sim \mathsf{s}$ *from* $\sigma(\mathcal{KB})$*.*

4.4 Cumulativity postulates

The idea behind a notion of Cumulativity in our setting is that adding a conditional to the knowledge base that was already inferred should not change anything in terms of its consequences. We can split this into two 'halves'.

Cautious Monotonicity If $\alpha \mid\!\sim^{\mathcal{KB}}_{*} \beta$ and $\mathcal{KB}' = \mathcal{KB} \cup \{\alpha \mid\!\sim \beta\}$, then $\mid\!\sim^{\mathcal{KB}}_{*} \subseteq \mid\!\sim^{\mathcal{KB}'}_{*}$.

[2] Marquis and Schwind (2014) consider much more general settings, but this is all we need in the present paper.

Cut If $\alpha \mid\!\sim_*^{\mathcal{KB}} \beta$ and $\mathcal{KB}' = \mathcal{KB} \cup \{\alpha \mid\!\sim \beta\}$, then $\mid\!\sim_*^{\mathcal{KB}'} \subseteq \mid\!\sim_*^{\mathcal{KB}}$.

We conclude this section with an impossibility result concerning a subset of the postulates we have mentioned so far.

Theorem 3. *There is no method $*$ simultaneously satisfying all of Inclusion, D-Rationality, Equivalence, Vacuity, Cautious Monotonicity and Negated Representation Independence.*

Proof. Assume, for contradiction, that $*$ satisfies all the listed properties. Suppose $\mathcal{P} = \{\mathsf{m}, \mathsf{s}\}$ and let $\mathcal{KB}$ be the knowledge base from Example 2, i.e., $\{\mathsf{m} \mid\!\sim \mathsf{s}, \neg\mathsf{m} \mid\!\sim \mathsf{s}\}$. By Inclusion, $\mathsf{m} \mid\!\sim_*^{\mathcal{KB}} \mathsf{s}$ and $\neg\mathsf{m} \mid\!\sim_*^{\mathcal{KB}} \mathsf{s}$. By D-Rationality, we know $\mid\!\sim_*^{\mathcal{KB}}$ satisfies the Or rule, so, from these two, we get $\mathsf{m} \vee \neg\mathsf{m} \mid\!\sim_*^{\mathcal{KB}} \mathsf{s}$ which, in turn, yields $(\mathsf{m} \leftrightarrow \mathsf{s}) \vee (\neg\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s}$, by LLE. Applying DR to this means we have:

$$(\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s} \text{ or } (\neg\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s} \qquad (1)$$

Now, let $\sigma(\cdot)$ be the negation-swapping symbol translation mentioned in Example 3, i.e., $\sigma(\mathsf{m}) = \neg\mathsf{m}$, $\sigma(\mathsf{s}) = \mathsf{s}$, so $\sigma(\mathcal{KB}) = \{\neg\mathsf{m} \mid\!\sim \mathsf{s}, \neg\neg\mathsf{m} \mid\!\sim \mathsf{s}\}$. Then, by Negated Representation Independence, we have $(\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s}$ iff $(\neg\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\sigma(\mathcal{KB})} \mathsf{s}$. But clearly we have $\mathcal{KB} \equiv \sigma(\mathcal{KB})$, so, by Equivalence, we obtain from this:

$$(\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s} \text{ iff } (\neg\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s} \qquad (2)$$

Putting (1) and (2) together gives us both $(\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s}$ and $(\neg\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s}$. Now, let $\mathcal{KB}' = \mathcal{KB} \cup \{(\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim \mathsf{s}\}$. By Cautious Monotonicity, $\mid\!\sim_*^{\mathcal{KB}} \subseteq \mid\!\sim_*^{\mathcal{KB}'}$. In particular, $(\neg\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}'} \mathsf{s}$. It can be checked that the disjunctive/preferential closure of $\mathcal{KB}'$ is itself a disjunctive consequence relation. In fact, it corresponds to the interval-based interpretation on the left of Figure 4. Hence, by Vacuity, this particular interval-based interpretation corresponds also to $\mid\!\sim_*^{\mathcal{KB}'}$. But, by inspecting this picture, we see $(\neg\mathsf{m} \leftrightarrow \mathsf{s}) \not\mid\!\sim_*^{\mathcal{KB}'} \mathsf{s}$, which leads to a contradiction. □

Theorem 3 is both surprising and disappointing, since all of the properties mentioned seem to be rather intuitive and desirable. Note that a close inspection of the proof shows that even just Vacuity and Cautious Monotonicity together place some quite severe restrictions on the behaviour of $*$.

Corollary 1. *Let $\mathcal{P} = \{p, q\}$ and $\mathcal{KB} = \{p \mid\!\sim q, \neg p \mid\!\sim q\}$. There is no operator $*$ satisfying Vacuity and Cautious Monotonicity that infers both $(p \leftrightarrow q) \mid\!\sim_*^{\mathcal{KB}} q$ and $(\neg p \leftrightarrow q) \mid\!\sim_*^{\mathcal{KB}} q$.*

What can we do in the face of these results? Our strategy will be to seek to construct a method that can satisfy as many of these properties as possible. We now provide our candidate for such a method - the disjunctive rational closure.

5 A construction for disjunctive rational closure

In order to satisfy D-Rationality, we can focus on constructing a special interval-based interpretation from $\mathcal{KB}$ and then take all conditionals holding in this interpretation as the consequences of $\mathcal{KB}$. In this section, we give our construction of the interpretation $\mathscr{I}_{DC}^{\mathcal{KB}}$ that gives us the *disjunctive rational closure* of a conditional knowledge base.

To specify $\mathscr{I}_{DC}^{\mathcal{KB}}$, we will construct the pair $\langle \mathscr{L}_{DC}^{\mathcal{KB}}, \mathscr{U}_{DC}^{\mathcal{KB}} \rangle$ of functions specifying the *lower* and *upper ranks* for each valuation. Since we aim to satisfy Infra-Rationality, our construction method takes the rational closure $\mathscr{R}_{RC}^{\mathcal{KB}}$ of $\mathcal{KB}$ as a point of departure. Starting with the lower ranks, we simply set, for all $v \in \mathcal{U}$:

$$\mathscr{L}_{DC}^{\mathcal{KB}}(v) \stackrel{\text{def}}{=} \mathscr{R}_{RC}^{\mathcal{KB}}(v).$$

That is, the lower ranks are given by the rational closure.

For the upper ranks $\mathscr{U}_{DC}^{\mathcal{KB}}$, if we happen to have $\mathscr{L}_{DC}^{\mathcal{KB}}(v) = \mathscr{R}_{RC}^{\mathcal{KB}}(v) = \infty$, then, to conform with the definition of interval-based interpretation, it is clear that we must set $\mathscr{U}_{DC}^{\mathcal{KB}}(v) = \infty$ also. If $\mathscr{L}_{DC}^{\mathcal{KB}}(v) \neq \infty$, then the construction of $\mathscr{U}_{DC}^{\mathcal{KB}}(v)$ becomes a little more involved. We require first the following definition.

Definition 5. *Given a ranked interpretation $\mathscr{R}$ and a conditional $\alpha \mid\!\sim \beta$ such that $\mathscr{R} \Vdash \alpha \mid\!\sim \beta$, we say a valuation v* ***verifies*** *$\alpha \mid\!\sim \beta$* ***in*** *$\mathscr{R}$ if $\mathscr{R}(v) = \mathscr{R}(\alpha)$.*

Now, assuming $\mathscr{L}_{DC}^{\mathcal{KB}}(v) \neq \infty$, our construction of $\mathscr{U}_{DC}^{\mathcal{KB}}(v)$ splits into two cases, according to whether v verifies any of the conditionals from $\mathcal{KB}$ in $\mathscr{R}_{RC}^{\mathcal{KB}}$ or not.

Case 1: v does not verify any of the conditionals in $\mathcal{KB}$ in $\mathscr{R}_{RC}^{\mathcal{KB}}$. In this case, we set:

$$\mathscr{U}_{DC}^{\mathcal{KB}}(v) \stackrel{\text{def}}{=} \max\{\mathscr{R}_{RC}^{\mathcal{KB}}(u) \mid \mathscr{R}_{RC}^{\mathcal{KB}}(u) \neq \infty\}$$

Case 2: v verifies at least one conditional from $\mathcal{KB}$ in $\mathscr{R}_{RC}^{\mathcal{KB}}$. In this case, the idea is to extend the upper rank of v as much as possible while still ensuring the constraints represented by $\mathcal{KB}$ are respected in the resulting $\mathscr{I}_{DC}^{\mathcal{KB}}$. If v verifies $\alpha \mid\!\sim \beta$ in $\mathscr{R}_{RC}^{\mathcal{KB}}$, then this is achieved by setting $\mathscr{U}_{DC}^{\mathcal{KB}}(v) = \mathscr{R}_{RC}^{\mathcal{KB}}(\alpha \wedge \neg\beta) - 1$; or, if $\mathscr{R}(\alpha \wedge \neg\beta) = \infty$, then again just set $\mathscr{U}_{DC}^{\mathcal{KB}}(v) = \max\{\mathscr{R}_{RC}^{\mathcal{KB}}(u) \mid \mathscr{R}_{RC}^{\mathcal{KB}}(u) \neq \infty\}$, as in Case 1. (This takes care of 'redundant' conditionals that might occur in $\mathcal{KB}$, like $\alpha \mid\!\sim \alpha$). We introduce now the following notation. Given sentences α, β:

$$t_{RC}^{\mathcal{KB}}(\alpha, \beta) \stackrel{\text{def}}{=} \begin{cases} \mathscr{R}_{RC}^{\mathcal{KB}}(\alpha \wedge \neg\beta) - 1, & \text{if } \mathscr{R}_{RC}^{\mathcal{KB}}(\alpha \wedge \neg\beta) \neq \infty \\ \max\{\mathscr{R}_{RC}^{\mathcal{KB}}(u) \mid \mathscr{R}_{RC}^{\mathcal{KB}}(u) \neq \infty\}, & \text{otherwise.} \end{cases}$$

But we need to take care of the situation in which v possibly verifies more than one conditional from $\mathcal{KB}$ in $\mathscr{R}_{RC}^{\mathcal{KB}}$. In order to ensure that *all* conditionals in $\mathcal{KB}$ will still be satisfied, we need to take:

$$\mathscr{U}_{DC}^{\mathcal{KB}}(v) \stackrel{\text{def}}{=} \min\{t_{RC}^{\mathcal{KB}}(\alpha, \beta) \mid (\alpha \mid\!\sim \beta) \in \mathcal{KB} \text{ and } v \text{ verifies } \alpha \mid\!\sim \beta \text{ in } \mathscr{R}_{RC}^{\mathcal{KB}}\}$$

So, summarising the two cases, we arrive at our final definition of $\mathscr{U}_{DC}^{\mathcal{KB}}$:

$$\mathscr{U}_{DC}^{\mathcal{KB}}(v) \stackrel{\text{def}}{=} \begin{cases} \min\{t_{RC}^{\mathcal{KB}}(\alpha, \beta) \mid \alpha \mid\!\sim \beta \in \mathcal{KB} \text{ and } v \text{ verifies } \alpha \mid\!\sim \beta \text{ in } \mathscr{R}_{RC}^{\mathcal{KB}}\}, & \text{if } v \text{ verifies at least one conditional from } \mathcal{KB} \text{ in } \mathscr{R}_{RC}^{\mathcal{KB}} \\ \max\{\mathscr{R}_{RC}^{\mathcal{KB}}(u) \mid \mathscr{R}_{RC}^{\mathcal{KB}}(u) \neq \infty\}, & \text{otherwise.} \end{cases}$$

Note that if v verifies $\alpha \mathrel{|\!\sim} \beta \in \mathcal{KB}$ in $\mathscr{R}_{RC}^{\mathcal{KB}}$, then $\mathscr{R}_{RC}^{\mathcal{KB}}(v) = \mathscr{R}_{RC}^{\mathcal{KB}}(\alpha) \leq \mathscr{R}_{RC}^{\mathcal{KB}}(\alpha \wedge \neg\beta) - 1 = t_{RC}^{\mathcal{KB}}(\alpha, \beta)$. Thus, in both cases above, we have $\mathscr{L}_{DC}^{\mathcal{KB}}(v) \leq \mathscr{U}_{DC}^{\mathcal{KB}}(v)$ and so the pair $\mathscr{L}_{DC}^{\mathcal{KB}}$ and $\mathscr{U}_{DC}^{\mathcal{KB}}$ form a legitimate interval-based interpretation.

We thus arrive at our final definition of the disjunctive rational closure of a conditional knowledge base.

Definition 6. *Let $\mathscr{I}_{DC}^{\mathcal{KB}} \stackrel{\text{def}}{=} \langle \mathscr{L}_{DC}^{\mathcal{KB}}, \mathscr{U}_{DC}^{\mathcal{KB}} \rangle$ be the interval-based interpretation specified by $\mathscr{L}_{DC}^{\mathcal{KB}}$ and $\mathscr{U}_{DC}^{\mathcal{KB}}$ as above. The **disjunctive rational closure** of $\mathcal{KB}$ is the defeasible consequence relation $\mathrel{|\!\sim}_{DC}^{\mathcal{KB}} \stackrel{\text{def}}{=} \{\alpha \mathrel{|\!\sim} \beta \mid \mathscr{I}_{DC}^{\mathcal{KB}} \Vdash \alpha \mathrel{|\!\sim} \beta\}$.*

In the remainder of this section, we revisit the examples we have seen throughout the paper, to see what answer the disjunctive rational closure gives.

Example 4. *Going back to Example 1, with $\mathcal{KB} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}\}$, the rational closure yields $\mathscr{R}_{RC}^{\mathcal{KB}}(\mathsf{bf}) = \mathscr{R}_{RC}^{\mathcal{KB}}(\bar{\mathsf{b}}\mathsf{f}) = \mathscr{R}_{RC}^{\mathcal{KB}}(\bar{\mathsf{b}}\bar{\mathsf{f}}) = 0$ and $\mathscr{R}_{RC}^{\mathcal{KB}}(\mathsf{b}\bar{\mathsf{f}}) = 1$. Since $\mathscr{L}_{DC}^{\mathcal{KB}} = \mathscr{R}_{RC}^{\mathcal{KB}}$, this gives us the lower ranks for each valuation in $\mathscr{I}_{DC}^{\mathcal{KB}}$. Turning to the upper ranks, the only valuation that verifies the single conditional $\mathsf{b} \mathrel{|\!\sim} \mathsf{f}$ in $\mathcal{KB}$ is bf, thus $\mathscr{U}_{DC}^{\mathcal{KB}}(\mathsf{bf}) = t_{RC}^{\mathcal{KB}}(\mathsf{b}, \mathsf{f}) = \mathscr{R}_{RC}^{\mathcal{KB}}(\mathsf{b} \wedge \neg\mathsf{f}) - 1 = 1 - 1 = 0$, meaning that the interval assigned to bf is $(0, 0)$. The other three valuations all get assigned the same upper rank, which is just the maximum finite rank occurring in $\mathscr{R}_{RC}^{\mathcal{KB}}$, which is 1. Thus the interval assigned to $\mathsf{b}\bar{\mathsf{f}}$ is $(1, 1)$, while both the valuations in $[\![\neg\mathsf{b}]\!]$ are assigned $(0, 1)$. So $\mathscr{I}_{DC}^{\mathcal{KB}}$ outputs exactly the same interval-based interpretation depicted in Figure 3 which, recall, gives the unique $\subseteq$-minimal disjunctive consequence relation extending $\mathcal{KB}$ in this case.*

Example 5. *Returning to Example 2, with $\mathcal{KB} = \{\mathsf{m} \mathrel{|\!\sim} \mathsf{s}, \neg\mathsf{m} \mathrel{|\!\sim} \mathsf{s}\}$, the rational closure yields $\mathscr{R}_{RC}^{\mathcal{KB}}(\mathsf{ms}) = \mathscr{R}_{RC}^{\mathcal{KB}}(\bar{\mathsf{m}}\mathsf{s}) = 0$ and $\mathscr{R}_{RC}^{\mathcal{KB}}(\mathsf{m}\bar{\mathsf{s}}) = \mathscr{R}_{RC}^{\mathcal{KB}}(\bar{\mathsf{m}}\bar{\mathsf{s}}) = 1$, which gives us the lower ranks. The valuation ms verifies only the conditional $\mathsf{m} \mathrel{|\!\sim} \mathsf{s}$, and so $\mathscr{U}_{DC}^{\mathcal{KB}}(\mathsf{ms}) = t_{RC}^{\mathcal{KB}}(\mathsf{m}, \mathsf{s}) = \mathscr{R}_{RC}^{\mathcal{KB}}(\mathsf{m} \wedge \neg\mathsf{s}) - 1 = 1 - 1 = 0$. Similarly, the valuation $\bar{\mathsf{m}}\mathsf{s}$ verifies only the conditional $\neg\mathsf{m} \mathrel{|\!\sim} \mathsf{s}$ and so, by analogous reasoning, $\mathscr{U}_{DC}^{\mathcal{KB}}(\bar{\mathsf{m}}\mathsf{s}) = t_{RC}^{\mathcal{KB}}(\neg\mathsf{m}, \mathsf{s}) = 0$. So both of these valuations are assigned the interval $(0, 0)$ by $\mathscr{I}_{DC}^{\mathcal{KB}}$. The other two valuations, which verify neither conditional in $\mathcal{KB}$, are assigned $(1, 1)$. Thus, in this case, $\mathscr{I}_{DC}^{\mathcal{KB}}$ returns just the rational closure of $\mathcal{KB}$, as pictured in Figure 5.*

In both the above examples, the disjunctive rational closure returns arguably the right answers.

Example 6. *Consider $\mathcal{KB} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \rightarrow \mathsf{b}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}\}$. As previously mentioned, the rational closure $\mathscr{R}_{RC}^{\mathcal{KB}}$ for this $\mathcal{KB}$ is depicted in Figure 1. Since both of the valuations in $[\![\mathsf{p} \wedge \neg\mathsf{b}]\!]$ (in red at the top of the picture) are deemed implausible (i.e., have rank ∞), they are both assigned interval (∞, ∞). Focusing then on just the plausible valuations, the only valuation verifying $\mathsf{b} \mathrel{|\!\sim} \mathsf{f}$ in $\mathscr{R}_{RC}^{\mathcal{KB}}$ is $\mathsf{bf}\bar{\mathsf{p}}$ (which verifies no other conditional in $\mathcal{KB}$), so $\mathscr{U}_{DC}^{\mathcal{KB}}(\mathsf{bf}\bar{\mathsf{p}}) = \mathscr{R}_{RC}^{\mathcal{KB}}(\mathsf{b} \wedge \neg\mathsf{f}) - 1 = 1 - 1 = 0$. The only valuation verifying $\mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}$ is $\mathsf{b}\bar{\mathsf{f}}\mathsf{p}$, so $\mathscr{U}_{DC}^{\mathcal{KB}}(\mathsf{b}\bar{\mathsf{f}}\mathsf{p}) = \mathscr{R}_{RC}^{\mathcal{KB}}(\mathsf{p} \wedge \mathsf{f}) - 1 = 2 - 1 = 1$. All other plausible valuations get assigned as their upper rank the maximum finite rank, which is 2. The resulting $\mathscr{I}_{DC}^{\mathcal{KB}}$ is the interval-based interpretation depicted in Figure 2.*

We end this section by considering our construction from the standpoint of complexity. The construction method above runs in time that grows (singly) exponentially with the size of the input, even if the rational closure of the knowledge base has been computed offline. To see why, let the input be a set of propositional atoms $\mathcal{P}$ together with a conditional knowledge base $\mathcal{KB}$, and let $|\mathcal{KB}| = n$. (For simplicity, we assume the size of $\mathcal{KB}$ to be the number of conditionals therein.) We know that $|\mathcal{U}| = 2^{|\mathcal{P}|}$. Now, for each valuation $v \in \mathcal{U}$, one has to check whether v verifies at least one conditional $\alpha \mathrel{|\!\sim} \beta$ in $\mathcal{KB}$ (cf. Definition 5). In the worst case, we have (*i*) all conditionals in $\mathcal{KB}$ will be checked against v, i.e., we will have n checks per valuation. Each of such checks amounts to comparing $\mathscr{R}(v)$ with $\mathscr{R}(\alpha)$, where α is the antecedent of the conditional under inspection. While $\mathscr{R}(v)$ is already known, $\mathscr{R}(\alpha)$ has to be computed (unless, of course, we also assume it has been done offline in the computation of the rational closure). Computing $\mathscr{R}(\alpha)$ is done by searching for the lowest valuations in $\mathscr{R}_{RC}^{\mathcal{KB}}$ satisfying α. In the worst case, we have that (*ii*) $2^{|\mathcal{P}|}$ valuations have to be inspected. Each such inspection amounts to a propositional verification, which is a polynomial-time task. Every time v verifies a conditional $\alpha \mathrel{|\!\sim} \beta$, the computation of $t_{RC}^{\mathcal{KB}}(\cdot)$ also requires that of $\mathscr{R}_{RC}^{\mathcal{KB}}(\alpha \wedge \neg\beta)$. In the worst case, the latter requires $2^{|\mathcal{P}|}$ propositional verifications. So, the computation of $t_{RC}^{\mathcal{KB}}(\cdot)$ takes at most (*iii*) $n \times 2^{|\mathcal{P}|}$ checks. From (*i*), (*ii*) and (*iii*), it follows that $n^2 \times 2^{2|\mathcal{P}|}$ propositional verifications are required. This has to be done for each of the $2^{|\mathcal{P}|}$ valuations, and therefore we have a total of $n^2 \times 2^{3|\mathcal{P}|}$ verifications in the worst case, from which the result follows.

Let us now take a look at the complexity of entailment checking, i.e., that of checking whether a conditional $\alpha \mathrel{|\!\sim} \beta$ is satisfied by $\mathscr{I}_{DC}^{\mathcal{KB}}$. This task amounts to computing $\mathscr{U}_{DC}^{\mathcal{KB}}(\alpha)$ and $\mathscr{L}_{DC}^{\mathcal{KB}}(\alpha \wedge \neg\beta)$ and comparing them. It is easy to see that in the worst-case scenario both require $2^{|\mathcal{P}|}$ propositional verifications.

6 Properties of the Disjunctive Rational Closure

We now turn to the question of which of the postulates from Section 4 are satisfied by the disjunctive rational closure. We start by observing that we obtain all of the basic postulates proposed in Section 4.1:

Proposition 1. *The disjunctive rational closure satisfies Inclusion, D-Rationality, Equivalence and Infra-Rationality.*

Proof. (Outline) D-Rationality is immediate since we construct an interval-based interpretation. Equivalence is also straightforward. For Infra-Rationality, first recall that $\alpha \mathrel{|\!\sim}_{DC}^{\mathcal{KB}} \beta$ iff $\mathscr{U}_{DC}^{\mathcal{KB}}(\alpha) < \mathscr{L}_{DC}^{\mathcal{KB}}(\alpha \wedge \neg\beta)$. Since $\mathscr{L}_{DC}^{\mathcal{KB}}(\alpha) \leq \mathscr{U}_{DC}^{\mathcal{KB}}(\alpha)$ (follows by definition of interval-based interpretation) and $\mathscr{L}_{DC}^{\mathcal{KB}}(\alpha \wedge \neg\beta) = \mathscr{R}_{RC}^{\mathcal{KB}}(\alpha \wedge \neg\beta)$ (by construction), we have $\mathscr{U}_{DC}^{\mathcal{KB}}(\alpha) < \mathscr{L}_{DC}^{\mathcal{KB}}(\alpha \wedge \neg\beta)$ implies $\mathscr{R}_{RC}^{\mathcal{KB}}(\alpha) = \mathscr{L}_{DC}^{\mathcal{KB}}(\alpha) < \mathscr{R}_{RC}^{\mathcal{KB}}(\alpha \wedge \neg\beta)$, giving $\alpha \mathrel{|\!\sim}_{RC}^{\mathcal{KB}} \beta$, as required for Infra-Rationality. For Inclusion, suppose $\alpha \mathrel{|\!\sim} \beta \in \mathcal{KB}$.

If $\mathscr{R}^{\mathcal{KB}}_{RC}(\alpha) = \infty$, then $\mathscr{L}^{\mathcal{KB}}_{DC}(\alpha) = \mathscr{U}^{\mathcal{KB}}_{DC}(\alpha) = \infty$ by construction and so $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \beta$. So assume $\mathscr{R}^{\mathcal{KB}}_{RC}(\alpha) \neq \infty$. Then, to show $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \beta$, it suffices to show $\mathscr{U}^{\mathcal{KB}}_{DC}(v) < \mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg\beta) = \mathscr{R}^{\mathcal{KB}}_{RC}(\alpha \wedge \neg\beta)$ for at least one $v \in [\![\alpha]\!]$. Since rational closure satisfies inclusion, we know $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{RC} \beta$ and so, since $\mathscr{R}^{\mathcal{KB}}_{RC}(\alpha) \neq \infty$, there must exist at least one v' verifying $\alpha \mathrel{|\!\sim} \beta$ in $\mathscr{R}^{\mathcal{KB}}_{RC}$. By construction of $\mathscr{U}^{\mathcal{KB}}_{DC}$, we have $\mathscr{U}^{\mathcal{KB}}_{DC}(v') \leq t^{\mathcal{KB}}_{RC}(\alpha, \beta) = \mathscr{R}^{\mathcal{KB}}_{RC}(\alpha \wedge \neg\beta) - 1$ as required. □

We remind the reader that, since Inclusion and D-Rationality hold, disjunctive rational closure also satisfies Preferential Extension.

Now we look at the Cumulativity properties. It is known from the work by Lehmann and Magidor (1992) that rational closure satisfies both Cautious Monotonicity and Cut, and, in fact, if $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{RC} \beta$ and $\mathcal{KB}' = \mathcal{KB} \cup \{\alpha \mathrel{|\!\sim} \beta\}$, then $\mathscr{R}^{\mathcal{KB}}_{RC} = \mathscr{R}^{\mathcal{KB}'}_{RC}$. We can show the following for disjunctive rational closure.

Proposition 2. *The disjunctive rational closure does not satisfy Cut.*

Proof. Assume $\mathcal{P} = \{\mathsf{b}, \mathsf{f}\}$, and $\mathcal{KB}$ is again the knowledge base from Example 1, i.e., $\{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}\}$. We have seen in Example 4 that $\mathscr{I}^{\mathcal{KB}}_{DC}$ is given by the interval-based interpretation depicted in Figure 3. By inspecting this picture, we see $\mathscr{I}^{\mathcal{KB}}_{DC} \Vdash \top \mathrel{|\!\sim} (\mathsf{b} \rightarrow \mathsf{f})$. Now let $\mathcal{KB}' = \mathcal{KB} \cup \{\top \mathrel{|\!\sim} (\mathsf{b} \rightarrow \mathsf{f})\}$. Then $\mathscr{I}^{\mathcal{KB}'}_{DC}$ is given by the model in Figure 6. We now have $\mathscr{I}^{\mathcal{KB}'}_{DC} \Vdash \neg\mathsf{f} \mathrel{|\!\sim} \neg\mathsf{b}$, whereas before we had $\mathscr{I}^{\mathcal{KB}}_{DC} \nVdash \neg\mathsf{f} \mathrel{|\!\sim} \neg\mathsf{b}$. □

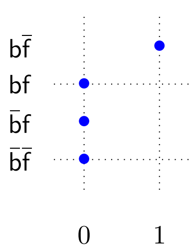

Figure 6: Output for $\mathcal{KB}' = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \top \mathrel{|\!\sim} (\mathsf{b} \rightarrow \mathsf{f})\}$.

Essentially, the reason for the failure of Cut is that by adding a new conditional $\alpha \mathrel{|\!\sim} \beta$ to the knowledge base, even when that conditional is already inferred by the disjunctive rational closure, we give certain valuations (namely those in $[\![\alpha]\!]$) the opportunity to verify one more conditional from the knowledge base in $\mathscr{R}^{\mathcal{KB}'}_{RC}$. (See, e.g. the two valuations in $[\![\neg\mathsf{b}]\!]$ in the above counterexample.) This leads, potentially, to a corresponding decrease in their upper ranks $\mathscr{U}^{\mathcal{KB}}_{*}$, leading in turn to more inferences being made available. This behaviour reveals that disjunctive rational closure can be termed a *base-driven* approach, since the conditionals that are included explicitly in the knowledge base have more influence compared to those that are merely derived. However, adding an inferred conditional will never lead to an *increase* in the upper ranks, which means the disjunctive rational closure *does* satisfy Cautious Monotonicity.

Proposition 3. *The disjunctive rational closure satisfies Cautious Monotonicity.*

Proof. (Outline) Suppose $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \beta$ and let $\mathcal{KB}' = \mathcal{KB} \cup \{\alpha \mathrel{|\!\sim} \beta\}$. Since disjunctive rational closure satisfies Infra-Rationality, we know $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{RC} \beta$, and so, since rational closure satisfies Cautious Monotonicity, $\mathscr{R}^{\mathcal{KB}}_{RC} = \mathscr{R}^{\mathcal{KB}'}_{RC}$, i.e., the lower ranks of all valuations in $\mathscr{I}^{\mathcal{KB}'}_{DC}$ are unchanged from $\mathscr{I}^{\mathcal{KB}}_{DC}$. To show $\mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \subseteq \mathrel{|\!\sim}^{\mathcal{KB}'}_{*}$, it thus suffices to show $\mathscr{U}^{\mathcal{KB}'}_{DC}(v) \leq \mathscr{U}^{\mathcal{KB}}_{DC}(v)$ for all valuations v. If v does not verify $\alpha \mathrel{|\!\sim} \beta$ in $\mathscr{R}^{\mathcal{KB}'}_{RC}$, then $\mathscr{U}^{\mathcal{KB}'}_{DC}(v) = \mathscr{U}^{\mathcal{KB}}_{DC}(v)$ (since all terms and cases in the definition of $\mathscr{U}^{\mathcal{KB}'}_{DC}$ depend only on $\mathscr{R}^{\mathcal{KB}'}_{RC} = \mathscr{R}^{\mathcal{KB}}_{RC}$), while if v does verify $\alpha \mathrel{|\!\sim} \beta$ in $\mathscr{R}^{\mathcal{KB}'}_{RC}$, then $\mathscr{U}^{\mathcal{KB}'}_{DC}(v) = \min\{\mathscr{U}^{\mathcal{KB}}_{DC}(v), t^{\mathcal{KB}'}_{RC}(\alpha, \beta)\} \leq \mathscr{U}^{\mathcal{KB}}_{DC}(v)$, as required. □

As we have seen in Corollary 1 in Section 4.4, the satisfaction of Cautious Monotonicity, plus the seemingly very reasonable behaviour displayed by disjunctive rational closure in Example 5, come at the cost of Vacuity, i.e., even if the preferential closure happens to be a disjunctive relation, the output may sanction extra conclusions.

Proposition 4. *The disjunctive rational closure does not satisfy Vacuity.*

Proof. By Corollary 1, there can be no operator $*$ satisfying Cautious Monotonicity and Vacuity that infers both $(\neg\mathsf{m} \leftrightarrow \mathsf{s}) \mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \mathsf{s}$ and $(\mathsf{m} \leftrightarrow \mathsf{s}) \mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \mathsf{s}$. We saw in Example 5 that the disjunctive rational closure returns the rational closure for this $\mathcal{KB}$, and so yields both these conditional inferences. We have also just seen that disjunctive rational closure satisfies Cautious Monotonicity. Hence we deduce that disjunctive rational closure cannot satisfy Vacuity. □

What about the Representation Independence postulates? Concerning full Representation Independence, we have remarked earlier that this postulate is not compatible with the basic postulates, and so Proposition 1 already tells us that disjunctive rational closure fails it. However, we conjecture that Negated Representation Independence is satisfied, since we can show that if rational closure satisfies it, then the disjunctive rational closure will inherit the property. Although Jaeger (1996) showed that rational closure does indeed conform with his version of Representation Independence, it remains to be proved that his notion coincides precisely with ours.

7 Concluding remarks

In this paper, we have set ourselves the task to revive interest in weaker alternatives to Rational Monotonicity when reasoning with conditional knowledge bases. We have studied the case of Disjunctive Rationality, a property already known by the community from the work of Kraus et al. and Freund in the early ’90s, which we have then coupled with a semantics in terms of interval orders borrowed from a more recent work by Rott in belief revision.

In our quest for a suitable form of entailment ensuring Disjunctive Rationality, we started by putting forward a set of postulates, all reasonable at first glance, characterising its expected behaviour. As it turns out, not all of them can be satisfied simultaneously, which suggests there might be more than one answer to our research question. We have then provided a construction of the disjunctive rational closure of a conditional knowledge base, which infers a set of conditionals intermediate between the preferential closure and the rational closure.

Regarding the properties of disjunctive rational closure, the news is somewhat mixed, with several basic postulates satisfied, as well as Cautious Monotonicity, but with neither Cut nor Vacuity holding in general. Regarding Cut, the reason for its failure seems tied to the fact that disjunctive rational closure places special importance on the conditionals that are explicitly written as part of the knowledge base. In this regard it shares commonalities with other base-driven approaches to defeasible inference such as the lexicographic closure (Lehmann 1995). We conjecture that a weaker version of Cut will still hold for our approach, according to which the new conditional added $\alpha \mathrel{|\!\sim} \beta$ is such that α already appears as an antecedent of another conditional already in $\mathcal{KB}$.

Regarding Vacuity, our impossibility result and surrounding discussion tells us that its failure is unavoidable given the other, reasonable, behaviour that we have shown disjunctive rational closure to exhibit. Essentially, when trying to devise a method for conditional inference under Disjunctive Rationality, we are faced with a choice between Vacuity and Cautious Monotonicity, with disjunctive rational closure favouring the latter at the expense of the former. It is possible, of course, to tweak the current approach by treating the case when $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ happens to be a disjunctive relation separately, outputting the preferential closure in this case, while returning the disjunctive rational closure otherwise. However the full ripple effects on the other properties of $\mathrel{|\!\sim}^{\mathcal{KB}}_{DC}$ of making this manoeuvre remain to be worked out.

As for future work, we plan to start by checking whether disjunctive rational closure satisfies Negated Representation Independence, as well as the Justification postulate. We also plan to investigate suitable definitions of a preference relation on the set of interval-based interpretations. We hope our construction can be shown to be the most preferred extension of the knowledge base according to some intuitively defined preference relation, as has been done in the rational case.

In this work we required the postulate of Infra-Rationality. As a result our construction of disjunctive rational closure took the rational closure as a starting point and then performed a particular modification to it to obtain a special 'privileged' subset of it that extends the input knowledge base and forms a disjunctive consequence relation. However it is clear that this modification could just as well be applied to any of the other conditional inference methods that have been suggested in the literature and that output a rational consequence relation, such as the lexicographic closure or System JLZ (Weydert 2003) or those based on c-revisions (Kern-Isberner 2001). It will be interesting to see what kind of properties will be gained or lost in these cases.

Finally, given the recent trend in applying defeasible reasoning to formal ontologies in Description Logics (Bonatti et al. 2015; Bonatti and Sauro 2017; Britz, Meyer, and Varzinczak 2011; Britz and Varzinczak 2019; Giordano et al. 2015; Pensel and Turhan 2018), an investigation of our approach beyond the propositional case is also envisaged.

Acknowledgments

This work is based upon research supported in part by the "Programme de Recherche Commun" *Non-Classical Reasoning for Enhanced Ontology-based Semantic Technologies* between the CNRS and the Royal Society. Thanks to the anonymous NMR reviewers for some helpful suggestions.

References

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.

Bonatti, P. 2019. Rational closure for all description logics. *Artificial Intelligence* 274:197–223.

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. 2019. On rational entailment for propositional typicality logic. *Artificial Intelligence* 277.

Booth, R.; Meyer, T.; and Varzinczak, I. 2012. PTL: A propositional typicality logic. In Fariñas del Cerro, L.; Herzig, A.; and Mengin, J., eds., *Proceedings of the 13th European Conference on Logics in Artificial Intelligence (JELIA)*, number 7519 in LNCS, 107–119. Springer.

Britz, K., and Varzinczak, I. 2017. Toward defeasible $\mathcal{SROIQ}$. In *Proceedings of the 30th International Workshop on Description Logics*.

Britz, K., and Varzinczak, I. 2018a. From KLM-style conditionals to defeasible modalities, and back. *Journal of Applied Non-Classical Logics (JANCL)* 28(1):92–121.

Britz, K., and Varzinczak, I. 2018b. Preferential accessibility and preferred worlds. *Journal of Logic, Language and Information (JoLLI)* 27(2):133–155.

Britz, K., and Varzinczak, I. 2019. Contextual rational closure for defeasible $\mathcal{ALC}$. *Annals of Mathematics and Artificial Intelligence* 87(1-2):83–108.

Britz, K.; Meyer, T.; and Varzinczak, I. 2011. Semantic foundation for preferential description logics. In Wang, D., and Reynolds, M., eds., *Proceedings of the 24th Australasian Joint Conference on Artificial Intelligence*, number 7106 in LNAI, 491–500. Springer.

Casini, G.; Meyer, T.; Moodley, K.; and Nortjé, R. 2014. Relevant closure: A new form of defeasible reasoning for description logics. In Fermé, E., and Leite, J., eds., *Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA)*, number 8761 in LNCS, 92–106. Springer.

Casini, G.; Meyer, T.; Moodley, K.; Sattler, U.; and Varzinczak, I. 2015. Introducing defeasibility into OWL ontologies. In Arenas, M.; Corcho, O.; Simperl, E.; Strohmaier, M.; d'Aquin, M.; Srinivas, K.; Groth, P.; Dumontier, M.; Heflin, J.; Thirunarayan, K.; and Staab, S., eds., *Proceedings of the 14th International Semantic Web Conference (ISWC)*, number 9367 in LNCS, 409–426. Springer.

Casini, G.; Meyer, T.; and Varzinczak, I. 2019. Taking defeasible entailment beyond rational closure. In Calimeri, F.; Leone, N.; and Manna, M., eds., *Proceedings of the 16th European Conference on Logics in Artificial Intelligence (JELIA)*, number 11468 in LNCS, 182–197. Springer.

Chafik, A.; Cheikh, F.; Condotta, J.-F.; and Varzinczak, I. 2020. On the decidability of a fragment of preferential LTL. In *Proceedings of the 27th International Symposium on Temporal Representation and Reasoning (TIME)*.

Fishburn, P. 1985. *Interval Orders and Interval Graphs: A Study of Partially Ordered Sets*. Wiley.

Freund, M. 1993. Injective models and disjunctive relations. *Journal of Logic and Computation* 3(3):231–247.

Gärdenfors, P., and Makinson, D. 1994. Nonmonotonic inference based on expectations. *Artificial Intelligence* 65(2):197–245.

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

Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2010. Preferential vs rational description logics: which one for reasoning about typicality? In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*, 1069–1070.

Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2012. A minimal model semantics for nonmonotonic reasoning. In Fariñas del Cerro, L.; Herzig, A.; and Mengin, J., eds., *Proceedings of the 13th European Conference on Logics in Artificial Intelligence (JELIA)*, number 7519 in LNCS, 228–241. Springer.

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

Jaeger, M. 1996. Representation independence of nonmonotonic inference relations. In Aiello, L.; Doyle, J.; and Shapiro, S., eds., *Proceedings of the 5th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 461–472. Morgan Kaufmann Publishers.

Kern-Isberner, G. 2001. *Conditionals in Nonmonotonic Reasoning and Belief Revision*. Springer, Lecture Notes in Artificial Intelligence.

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? *Artificial Intelligence* 55:1–60.

Lehmann, D. 1995. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence* 15(1):61–82.

Makinson, D. 1994. General patterns in nonmonotonic reasoning. In Gabbay, D.; Hogger, C.; and Robinson, J., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3. Oxford University Press. 35–110.

Marquis, P., and Schwind, N. 2014. Lost in translation: Language independence in propositional logic – application to belief change. *Artificial Intelligence* 206:1–24.

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

Pensel, M., and Turhan, A.-Y. 2017. Including quantification in defeasible reasoning for the description logic $\mathcal{EL}_{\perp}$. In Balduccini, M., and Janhunen, T., eds., *Proceedings of the 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, number 10377 in LNCS, 78–84. Springer.

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* 112:28–70.

Reiter, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13(1-2):81–132.

Rott, H. 2014. Four floors for the theory of theory change: The case of imperfect discrimination. In Fermé, E., and Leite, J., eds., *Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA)*, number 8761 in LNCS, 368–382. Springer.

Shoham, Y. 1988. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press.

Varzinczak, I. 2018. A note on a description logic of concept and role typicality for defeasible reasoning over ontologies. *Logica Universalis* 12(3-4):297–325.

Weydert, E. 2003. System JLZ - rational default reasoning by minimal ranking constructions. *Journal of Applied Logic* 1(3-4):273–308.