Artificial Intelligence 348 (2025) 104418

Contents lists available at ScienceDirect

Artificial Intelligence

journal homepage: www.elsevier.com/locate/artint

On the disjunctive rational closure of a conditional knowledge base

Richard Booth [a,*], Ivan Varzinczak [b,c,d]

[a] *Cardiff University, United Kingdom*
[b] *Université Sorbonne Paris Nord, Inserm, Sorbonne Université, Limics, 93017 Bobigny, France*
[c] *CAIR, University of Cape Town, South Africa*
[d] *ISTI–CNR, Italy*

ARTICLE INFO

Keywords:
Knowledge representation and reasoning
Non-monotonic reasoning
Conditional inference
(Disjunctive) rationality

ABSTRACT

One of the most widely investigated decision problems in symbolic AI is that of which conditional sentences of the form “if α, then normally β” should follow from a knowledge base containing this type of statements. Probably, the most notable approach to this problem is the rational closure construction put forward by Lehmann and Magidor in the’90s, which has been adapted to logical languages of various expressive powers since then. At the core of rational closure is the Rational Monotonicity property, which allows one to retain existing (defeasible) conclusions whenever new information cannot be negated by existing conclusions. As it turns out, Rational Monotonicity is not universally accepted, with many researchers advocating the investigation of weaker versions thereof leading to a larger class of consequence relations. A case in point is that of the Disjunctive Rationality property, which states that if one may draw a (defeasible) conclusion from a disjunction of premises, then one should be able to draw this conclusion from at least one of the premises taken alone. 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 the simplest disjunctive consequence relation in this setting is has not been explored in depth. In this article, we do precisely that by motivating and proposing a concrete construction of the disjunctive rational closure of a conditional knowledge base, of which the properties and consequences of its adoption we also investigate in detail. (Previous versions of this work have been selected for presentation at the 18th International Workshop on Nonmonotonic Reasoning (NMR 2020) [1] and at the 35th AAAI Conference on Artificial Intelligence (AAAI 2021) [2]. The present submission extends and elaborates on both papers.)

1. Introduction

The question of *conditional inference*, i.e., of which conditionals of the form “if α, then normally β” should follow from a set $\mathcal{KB}$ of such statements, has been one of the classic questions of symbolic AI, with several well-known solutions being proposed over the past decades [3–7]. 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 postulates of preferential reasoning *plus* Rational Monotonicity. One of the reasons for

* Corresponding author.
E-mail address: boothr2@cardiff.ac.uk (R. Booth).

https://doi.org/10.1016/j.artint.2025.104418
Received 6 November 2023; Received in revised form 22 March 2025; Accepted 1 September 2025

Available online 12 September 2025

accepting rational closure is the fact it delivers a supra-classical, more venturous, notion of entailment that is conservative enough. Given that, rationality has long been accepted as the core baseline for any appropriate form of non-monotonic entailment from a set of conditionals.

Very few have stood against this position, including Stalnaker [8], with his example showing Rational Monotonicity can sometimes be counter-intuitive, and Makinson [9], who considered Rational Monotonicity too strong and has briefly advocated the weaker postulate of Disjunctive Rationality instead. The latter says that if one may draw a defeasible conclusion from a disjunction of premises, then one should (defeasibly) be able to draw this conclusion from at least one of these premises taken alone. This postulate holds for all rational consequence relations, and may still be desirable in cases where rationality does not hold. Quite surprisingly, the debate did not catch on at the time, and, for lack of rivals of the same stature, rational closure has since reigned alone as a role model in non-monotonic conditional inference. That used to be the case until Rott [10] reignited interest in Disjunctive Rationality by considering interval models in connection with belief contraction. Inspired by that, here we revisit disjunctive consequence relations and introduce 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 postulate of Disjunctive Rationality. We move on and investigate a notion of closure of (or entailment from) a conditional knowledge base under Disjunctive Rationality (Section 4). 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. We provide impossibility results showing incompatibilities between some of these postulates. Following that, in Section 5, we propose a specific construction for the disjunctive rational closure of a conditional knowledge base and discuss its computational complexity. Section 6 is devoted to an assessment of the suitability of our construction in light of the postulates previously put forward. In Section 7 we consider a slightly modified alternative to the disjunctive rational closure and assess the implications of its adoption. We conclude with some remarks on future directions of investigation. The more complex proofs for some of the stated results can be found in the appendix.

Previous versions of this work have been selected for presentation at the 18th International Workshop on Nonmonotonic Reasoning (NMR 2020) [1] and at the 35th AAAI Conference on Artificial Intelligence (AAAI 2021) [2]. The present paper extends and elaborates on both papers in the following way: (*i*) It contains the full proofs of all results stated in the above-mentioned versions (longer proofs appear in a dedicated appendix, whereas shorter proofs are interspersed within the text); (*ii*) the present version contains new results (Propositions 3 and 9—13, and Corollary 4) along with their proofs, and we also put forward and discuss some new postulates; (*iii*) the present paper contains a discussion on the computational complexity of the framework put forward in this work (end of Section 5); (*iv*) it fixes a conjecture made in the previous versions, which turned out to be incorrect, and expands on it, and (*v*) the present paper contains a new section (Section 7) devoted to an analysis of an alternative construction to the disjunctive rational closure we introduce here and how they compare.

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 $\mathcal{U} \overset{\text{def}}{=} \{0,1\}^{\mathcal{P}}$, 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., $\bar{p}$), in alphabetical order, 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 $\mathsf{b}\bar{\mathsf{f}}\mathsf{p}$.

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 $[\![\alpha]\!] \overset{\text{def}}{=} \{v \in \mathcal{U} \mid v \Vdash \alpha\}$. This notion is extended to a (possibly infinite) set of sentences $X \subseteq \mathcal{L}$ in the usual way: $[\![X]\!] \overset{\text{def}}{=} \bigcap_{\alpha \in X} [\![\alpha]\!]$. We say a set of sentences X (classically) *entails* $\alpha \in \mathcal{L}$, denoted $X \models \alpha$, if $[\![X]\!] \subseteq [\![\alpha]\!]$. Finally, α is a *tautology*, denoted $\models \alpha$, if $[\![\alpha]\!] = \mathcal{U}$.

2.1. KLM-style rational consequence

In this subsection, we briefly recall the so-called KLM approach to non-monotonic reasoning with a focus on a weaker form of monotonicity called Rational Monotonicity.

Several approaches to non-monotonic reasoning have been proposed in the literature over the past 40 years. Drawing on the semantics by Hansson [11], the *preferential approach*, initially put forward by Shoham [12] and subsequently developed by Kraus et al. [13] 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 [14], default logic [15], and many others, and (*iii*) its formal links with AGM-style belief revision [16]. The fruitfulness of the preferential approach is also witnessed by a great deal of recent work extending it to languages that are more expressive than that of propositional logic, such as those of description logics [17–21] and modal logics [22–24].

A *consequence relation* $\mid\!\sim$ is a binary relation on sentences of our underlying propositional language, i.e., $\mid\!\sim \subseteq \mathcal{L} \times \mathcal{L}$. We say that $\mid\!\sim$ is a *preferential* consequence relation [13] if it satisfies the following set of (Gentzen-style) 'rules', usually called *postulates* or *properties*, with the understanding that $\alpha \mid\!\sim \beta$ is a shorthand for $(\alpha, \beta) \in \mid\!\sim$:

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

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

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

Reflexivity (Ref) specifies the obvious requirement that every sentence is a (defeasible) consequence of itself. Left Logical Equivalence (LLE) captures a form of syntax independence regarding the antecedent of conditionals. The And postulate allows conjoining two defeasible conclusions from the same premise. The Or postulate specifies that if the same defeasible conclusion follows from two (possibly different) sentences, then it follows from either of them. Right Weakening (RW) formalises that every weaker sentence than a defeasible conclusion should follow defeasibly from the given antecedent. Finally, Cautious Monotonicity (CM) captures a weaker form of monotonicity, allowing us to strengthen the antecedent with any known defeasible conclusion from the premise under consideration.

If, in addition to the preferential postulates, the consequence relation $\mid\!\sim$ also satisfies the following Rational Monotonicity property [5], which is also a weak form of monotonicity, it is said to be a *rational* consequence relation:

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

For more details on all the postulates above, we refer the reader to the provided references [13,5].

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

Definition 1 *(Ranked Interpretation).* A **ranked interpretation** $\mathscr{R}$ is a function from $\mathcal{U}$ to $\mathbb{N} \cup \{\infty\}$ satisfying the following **convexity property**: for every $v \in \mathcal{U}$ and every $i \in \mathbb{N}$, if $\mathscr{R}(v) = i$, 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 valuations 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]

Example 1. Let $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$. An example of a ranked interpretation would be $\mathscr{R}$ such that $\mathscr{R}(\bar{\mathsf{b}}\bar{\mathsf{f}}\bar{\mathsf{p}}) = \mathscr{R}(\bar{\mathsf{b}}\mathsf{f}\bar{\mathsf{p}}) = \mathscr{R}(\mathsf{b}\mathsf{f}\bar{\mathsf{p}}) = 0$, $\mathscr{R}(\mathsf{b}\bar{\mathsf{f}}\bar{\mathsf{p}}) = \mathscr{R}(\mathsf{b}\bar{\mathsf{f}}\mathsf{p}) = 1$, $\mathscr{R}(\mathsf{b}\mathsf{f}\mathsf{p}) = 2$, and $\mathscr{R}(\bar{\mathsf{b}}\mathsf{f}\mathsf{p}) = \mathscr{R}(\bar{\mathsf{b}}\bar{\mathsf{f}}\mathsf{p}) = \infty$.

Fig. 1 depicts the ranked interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$ in Example 1. (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 are with red.)

Given a ranked interpretation $\mathscr{R}$ and $\alpha \in \mathcal{L}$, we say α is *true* in $\mathscr{R}$, denoted $\mathscr{R} \Vdash \alpha$, if all plausible valuations in $\mathscr{R}$ satisfy α. With $\mathscr{R}(\alpha) \stackrel{\text{def}}{=} \min\{\mathscr{R}(v) \mid v \Vdash \alpha \text{ and } \mathscr{R}(v) \neq \infty\}$, we denote the *rank of* α *in* $\mathscr{R}$. By convention, if none of the plausible valuations in $\mathscr{R}$ satisfies α, we let $\mathscr{R}(\alpha) = \infty$.

A 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 $\mathscr{R}(\alpha) < \mathscr{R}(\alpha \wedge \neg\beta)$ or, equivalently, $\mathscr{R}(\alpha \wedge \beta) < \mathscr{R}(\alpha \wedge \neg\beta)$. (And here we adopt Jaeger's [26] convention that $\infty < \infty$ always holds.) Equivalently, $\mathscr{R} \Vdash \alpha \mid\!\sim \beta$ if β holds in all the most normal α-valuations. 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 $\alpha \mid\!\sim \beta$.

[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 to ensure uniformity in examples. For an explicit use of implausible valuations in counterfactual reasoning, see the work of Casini et al. [25].

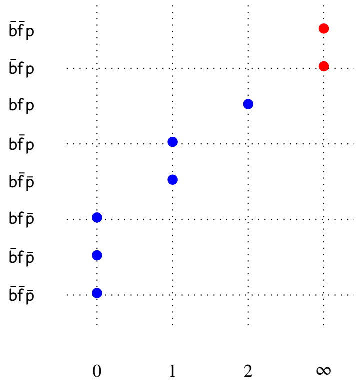

Fig. 1. A ranked interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

Example 2. Let $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$ and let $\mathscr{R}$ be as in Example 1 (depicted in Fig. 1). We have $\mathscr{R} \Vdash \mathsf{b} \mid\!\sim \mathsf{f}$ (birds usually fly), $\mathscr{R} \Vdash \mathsf{p} \to \mathsf{b}$ (penguins are birds), i.e., $\mathscr{R} \Vdash \neg(\mathsf{p} \to \mathsf{b}) \mid\!\sim \bot$, $\mathscr{R} \Vdash \mathsf{p} \mid\!\sim \neg\mathsf{f}$ (penguins usually do not fly), and $\mathscr{R} \nVdash \mathsf{f} \mid\!\sim \mathsf{b}$ (it is not the case that flying creatures are usually birds), which are all according to the intuitive expectations. Notice that we also have $\mathscr{R} \Vdash \mathsf{p} \land \neg\mathsf{b} \mid\!\sim \mathsf{b}$.

The fact this semantic characterisation of rational consequence is appropriate is a consequence of a representation result linking the seven rationality postulates above to precisely the class of ranked interpretations:

Theorem 1 *(Lehmann & Magidor, 1992). A consequence relation* $\mid\!\sim$ *is rational iff there is an* $\mathscr{R}$ *such that* $\alpha \mid\!\sim \beta$ *iff* $\mathscr{R} \Vdash \alpha \mid\!\sim \beta$.

2.2. Rational closure

In this subsection, we briefly recall Lehmann and Magidor's construction for reasoning defeasibly about conditionals.

One can view consequence relations of the type we are interested in here as formalising some form of (defeasible) conditional, so we can bring it down to the level of statements. Such was the stance adopted by Lehmann and Magidor [5]. A *conditional knowledge base* $\mathcal{KB}$ is thus a finite set of statements of the form $\alpha \mid\!\sim \beta$, with $\alpha, \beta \in \mathcal{L}$. In knowledge bases, we shall abbreviate $\neg\alpha \mid\!\sim \bot$ with α (see above).

Example 3. Let $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$. $\mathcal{KB} = \{\mathsf{b} \mid\!\sim \mathsf{f}, \mathsf{p} \to \mathsf{b}, \mathsf{p} \mid\!\sim \neg\mathsf{f}\}$ is a conditional knowledge base formalising a possible state of initial knowledge regarding penguins and birds.

Given a conditional knowledge base $\mathcal{KB}$, a *ranked model* of $\mathcal{KB}$ is a ranked interpretation satisfying all statements in $\mathcal{KB}$. With $[\![\mathcal{KB}]\!]_{\mathscr{R}}$ we shall denote the set of all ranked models of $\mathcal{KB}$. We use $[\![\alpha \mid\!\sim \beta]\!]_{\mathscr{R}}$ as an abbreviation for $[\![\{\alpha \mid\!\sim \beta\}]\!]_{\mathscr{R}}$.

Example 4. The ranked interpretation $\mathscr{R}$ from Example 1 (depicted in Fig. 1) is a ranked model of the conditional knowledge base $\mathcal{KB}$ in Example 3. 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.

From a knowledge representation perspective, 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 properties of Section 2.1. Intuitively, that means whenever appropriate instantiations of the premises in a property 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 consequence relation $\mid\!\sim^{\mathcal{KB}}$. By closing the latter under the preferential properties (in the sense of exhaustively applying them as 'inference rules'), 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 postulates, is the *preferential closure* of $\mid\!\sim^{\mathcal{KB}}$, which we denote by $\mid\!\sim^{\mathcal{KB}}_{PC}$.

The above process and definitions carry over when one requires the consequence relations also to be closed under the postulate RM, in which case we talk of *rational* extensions of $\mid\!\sim^{\mathcal{KB}}$. Nevertheless, as pointed out by Lehmann and Magidor [5, 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 $[\![\mathcal{KB}]\!]_{\mathscr{R}} \subseteq [\![\alpha \mid\!\sim \beta]\!]_{\mathscr{R}}$, i.e., 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 [4,27,7,28,20,29,30], with *rational closure* [5] 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* [4, 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 (see below). It is easy to see that $\preceq_{\mathcal{KB}}$ is a weak partial order. Giordano et al. [20] 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 *(Rational Closure).* Let $\mathcal{KB}$ be a conditional knowledge base, and let $\mathscr{R}^{\mathcal{KB}}_{RC}$ be the minimum element of $[\![\mathcal{KB}]\!]_{\mathscr{R}}$, i.e., the minimum of the ranked models of $\mathcal{KB}$ w.r.t. $\preceq_{\mathcal{KB}}$. The **rational closure** of $\mathcal{KB}$ is the consequence relation $\mid\!\sim^{\mathcal{KB}}_{RC} \stackrel{\text{def}}{=} \{\alpha \mid\!\sim \beta \mid \mathscr{R}^{\mathcal{KB}}_{RC} \Vdash \alpha \mid\!\sim \beta\}$.

Example 5. The ranked interpretation $\mathscr{R}$ from Example 1 (depicted in Fig. 1) is the minimum ranked model of the conditional knowledge base $\mathcal{KB}$ in Example 3 w.r.t. $\preceq_{\mathcal{KB}}$. Hence we have both $\neg \mathsf{f} \mid\!\sim^{\mathcal{KB}}_{RC} \neg \mathsf{b}$ and $\top \mid\!\sim^{\mathcal{KB}}_{RC} \neg \mathsf{p}$.

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 descriptions of how valuations compare to each other. 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 non-monotonic entailment, on which other, more venturous forms of entailment can be and have been constructed [4,31,32,29,30,25].

3. Disjunctive rationality and interval-based preferential semantics

In this section, we motivate the need for a weaker form of rationality and propose a semantics which we show to be suitable for its characterisation.

One may argue that there are cases in which Rational Monotonicity is too strong a postulate to enforce and for which a weaker consequence relation would suffice [9,33]. Nevertheless, as pointed out above, doing away completely with rationality (i.e., sticking to the preferential postulates only) is not particularly suitable in a broader defeasible-reasoning context. Indeed, as widely known in the literature, preferential systems induce entailment relations that are monotonic [5]. In that respect, here we are interested in 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) postulate [13] below:

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

Intuitively, DR says that if one may draw a (defeasible) conclusion from a disjunction of premises, then one should be able to draw this conclusion from at least one of these premises taken alone. Kraus et al. [13] offered the following example to illustrate the plausibility of DR:

> “If we do not hold that if Peter comes to the party, it will be great and do not hold that if Cathy comes to the party, it will be great, how could we hold that if at least one of Peter or Cathy comes, the party will be great?”

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 [5]. Therefore, DR is a weaker form of rationality, as its name already 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 [34] 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, but we shall not investigate this point here.)

Definition 3 *(Interval-Based Interpretation).* An **interval-based interpretation** is a pair $\mathscr{I} \stackrel{\text{def}}{=} \langle \mathscr{L}, \mathscr{U} \rangle$, where $\mathscr{L}$ and $\mathscr{U}$ are functions from $\mathcal{U}$ to $\mathbb{N} \cup \{\infty\}$ s.t. for all $v \in \mathcal{U}$, (a) $\mathscr{L}(v) \leq \mathscr{U}(v)$; (b) if $\mathscr{L}(v) = i$ or $\mathscr{U}(v) = i$, then for every $0 \leq j < i$, there is u s.t. $\mathscr{L}(u) = j$ or $\mathscr{U}(u) = j$, and (c) $\mathscr{L}(v) = \infty$ iff $\mathscr{U}(v) = \infty$. 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, the pair $(\mathscr{L}(v), \mathscr{U}(v))$ is the **interval of** v **in** $\mathscr{I}$. We say v is **more preferred than** u **in** $\mathscr{I}$, denoted $v \prec u$, if $\mathscr{U}(v) < \mathscr{L}(u)$.

The order $\prec$ on $\mathcal{U}$ defined above via an interval-based interpretation forms an *interval order* over the set of valuations of finite lower or upper rank, i.e., it is a strict partial order that additionally satisfies the *interval condition* below:

$$\text{If } v \prec u \text{ and } v' \prec u', \text{ then } v \prec u' \text{ or } v' \prec u \qquad (1)$$

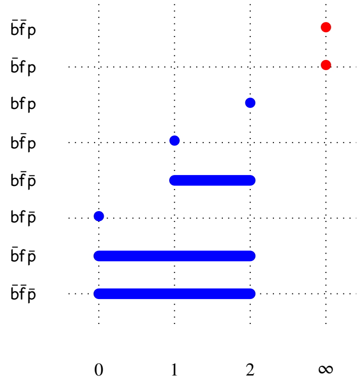

Fig. 2. An interval-based interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$.

Furthermore, every interval order over any subset of $\mathcal{U}$ can be defined from an interval-based interpretation in this way. See the work of Fishburn [35] for more details, and also that of Rott [10], who more recently explored interval orders in the context of belief contraction. See also Appendix A for a bit more detail on the relationship between interval orders and interval-based interpretations.

Example 6. Let $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$. An example of an interval-based interpretation would be $\mathscr{I} = \langle \mathscr{L}, \mathscr{U} \rangle$, where $\mathscr{L}(\bar{\mathsf{b}}\bar{\mathsf{f}}\bar{\mathsf{p}}) = 0$, $\mathscr{U}(\bar{\mathsf{b}}\bar{\mathsf{f}}\bar{\mathsf{p}}) = 2$, $\mathscr{L}(\bar{\mathsf{b}}\mathsf{f}\bar{\mathsf{p}}) = 0$, $\mathscr{U}(\bar{\mathsf{b}}\mathsf{f}\bar{\mathsf{p}}) = 2$, $\mathscr{L}(\mathsf{b}\mathsf{f}\bar{\mathsf{p}}) = \mathscr{U}(\mathsf{b}\mathsf{f}\bar{\mathsf{p}}) = 0$, $\mathscr{L}(\mathsf{b}\bar{\mathsf{f}}\bar{\mathsf{p}}) = 1$, $\mathscr{U}(\mathsf{b}\bar{\mathsf{f}}\bar{\mathsf{p}}) = 2$, $\mathscr{L}(\mathsf{b}\bar{\mathsf{f}}\mathsf{p}) = \mathscr{U}(\mathsf{b}\bar{\mathsf{f}}\mathsf{p}) = 1$, $\mathscr{L}(\mathsf{b}\mathsf{f}\mathsf{p}) = \mathscr{U}(\mathsf{b}\mathsf{f}\mathsf{p}) = 2$, and $\mathscr{L}(\bar{\mathsf{b}}\mathsf{f}\mathsf{p}) = \mathscr{U}(\bar{\mathsf{b}}\mathsf{f}\mathsf{p}) = \mathscr{L}(\bar{\mathsf{b}}\bar{\mathsf{f}}\mathsf{p}) = \mathscr{U}(\bar{\mathsf{b}}\bar{\mathsf{f}}\mathsf{p}) = \infty$.

Fig. 2 depicts the interval-based interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$ in Example 6. 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 with valuations v and u overlap, the intuition is that both valuations are incomparable in $\mathscr{I}$, i.e., neither is preferred to the other; otherwise, the valuation assigned the leftmost interval is seen as more preferred than the valuation which gets assigned the rightmost one.

In Fig. 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 (in red); 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 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 *(c)* of the previous definition). With $\mathscr{L}(\alpha) \stackrel{\text{def}}{=} \min\{\mathscr{L}(v) \mid v \Vdash \alpha \text{ and } \mathscr{L}(v) \neq \infty\}$ and $\mathscr{U}(\alpha) \stackrel{\text{def}}{=} \min\{\mathscr{U}(v) \mid v \Vdash \alpha \text{ and } \mathscr{L}(v) \neq \infty\}$ we denote, respectively, the *lower* and the *upper rank of α in $\mathscr{I}$*. By convention, if α is satisfied by none of the plausible valuations in $\mathscr{I}$, 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 $\mathscr{U}(\alpha) < \mathscr{L}(\alpha \wedge \neg\beta)$, or equivalently (using the facts that $\mathscr{U}(\alpha) = \min\{\mathscr{U}(\alpha \wedge \beta), \mathscr{U}(\alpha \wedge \neg\beta)\}$ and $\mathscr{L}(\alpha \wedge \neg\beta) \leq \mathscr{U}(\alpha \wedge \neg\beta)$) $\mathscr{U}(\alpha \wedge \beta) < \mathscr{L}(\alpha \wedge \neg\beta)$. (Recall the convention that $\infty < \infty$.)

Example 7. Let $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$ and let $\mathscr{I}$ be as in Example 6 (depicted in Fig. 2). We have $\mathscr{I} \Vdash \mathsf{b} \mathrel{|\!\sim} \mathsf{f}$, $\mathscr{I} \Vdash \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}$, and $\mathscr{I} \nVdash \neg\mathsf{f} \mathrel{|\!\sim} \neg\mathsf{b}$ (contrary to the ranked interpretation $\mathscr{R}$ in Example 1, which endorses the latter).

Given a conditional knowledge base $\mathcal{KB}$, an *interval-based model* of $\mathcal{KB}$ is an interval-based interpretation satisfying all statements in $\mathcal{KB}$. With $[\![\mathcal{KB}]\!]_{\mathscr{I}}$ we shall denote the set of all interval-based models of $\mathcal{KB}$. As before, we use $[\![\alpha \mathrel{|\!\sim} \beta]\!]_{\mathscr{I}}$ as an abbreviation for $[\![\{\alpha \mathrel{|\!\sim} \beta\}]\!]_{\mathscr{I}}$.

In the tradition of the KLM approach to defeasible reasoning, we define the consequence relation induced by an interval-based interpretation: given $\mathscr{I}$, $\mathrel{|\!\sim}_{\mathscr{I}} \stackrel{\text{def}}{=} \{\alpha \mathrel{|\!\sim} \beta \mid \mathscr{I} \Vdash \alpha \mathrel{|\!\sim} \beta\}$. We can now state a KLM-style representation result establishing that our interval-based semantics is suitable for characterising the class of disjunctive consequence relations, which is a variant of Freund's result [34]:

Theorem 2. *A consequence relation is a disjunctive consequence relation if and only if it is defined by some interval-based interpretation, i.e., $\mathrel{|\!\sim}$ is disjunctive if and only if there is $\mathscr{I}$ such that $\mathrel{|\!\sim} = \mathrel{|\!\sim}_{\mathscr{I}}$.*

The proof for the theorem above is available in the appendix.

4. Towards disjunctive rational closure

In this section, we address the question of what a suitable definition of entailment under disjunctive rationality is. We do so by stating and assessing a set of postulates that such a definition could satisfy. Before that, we make explicit the notion of disjunctive extension and show a result which serves as motivation for the remaining of the section.

Given a conditional knowledge base $\mathcal{KB}$, the definition of a *disjunctive extension* of $\mid\!\sim^{\mathcal{KB}}$ is analogous to that of a preferential one, namely closing $\mid\!\sim^{\mathcal{KB}}$ under the preferential properties plus disjunctive rationality (cf. Section 2.2). In that sense, the obvious definition of closure under Disjunctive Rationality consists in taking the intersection of all *disjunctive extensions* of $\mid\!\sim^{\mathcal{KB}}$. 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 $[\![\mathcal{KB}]\!]_{\mathscr{I}} \subseteq [\![\alpha \mid\!\sim \beta]\!]_{\mathscr{I}}$, i.e., 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.

Proposition 1. *Given a conditional knowledge base $\mathcal{KB}$, we have: (i) the disjunctive closure of $\mathcal{KB}$ coincides with its preferential closure $\mid\!\sim^{\mathcal{KB}}_{PC}$, and (ii) there exists $\mathcal{KB}$ such that $\mid\!\sim^{\mathcal{KB}}_{PC}$ does not satisfy DR.*

Proof. Showing (*i*): Since every rational consequence relation is disjunctive, and every disjunctive consequence relation is preferential, we know that the disjunctive closure must contain the preferential closure of $\mathcal{KB}$ and be contained in the intersection of all the rational extensions of $\mathcal{KB}$. But we already know that the latter intersection is equal to the preferential closure, from which the result immediately follows.

Showing (*ii*): 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 $\mathsf{p} \vee \neg\mathsf{p} \mid\!\sim^{\mathcal{KB}}_{PC} \mathsf{b}$, 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 $\neg\mathsf{p} \not\mid\!\sim_{\mathscr{I}_2} \mathsf{b}$. □

The result above 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 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

In this subsection, we propose the first four postulates that a suitable notion of entailment under disjunctive rationality should satisfy. 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 2, D-Rationality is equivalent to saying that 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 with "rational consequence", then that is the postulate that is usually considered in the area.

We point out that, given Proposition 1, the postulate below follows from Inclusion and D-Rationality:

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

Finally, in line with the motivation laid out in the Introduction and in the beginning of Section 3, the last of our basic postulates requires rational closure to be the upper bound on how venturous our consequence relation should be.

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

Therefore, given the last two postulates, the consequence relation we are looking for here should be one lying between the preferential closure and the rational closure of a conditional knowledge base.

4.2. Equivalence postulates

In this subsection, we turn our attention to postulates capturing different forms of syntax independence.

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. We present a series of notions of equivalence of decreasing strength leading to correspondingly progressively stronger requirements of syntax independence. One strong albeit natural notion of equivalence between knowledge bases can be defined as follows.

Definition 4 *(Piecewise Equivalence).* Let $\alpha, \beta, \gamma, \delta \in \mathcal{L}$. We say $\alpha \mid\!\sim \beta$ is **piecewise equivalent** to $\gamma \mid\!\sim \delta$ if $\models (\alpha \leftrightarrow \gamma) \wedge (\beta \leftrightarrow \delta)$. We say that knowledge bases $\mathcal{KB}$ and $\mathcal{KB}'$ are **piecewise equivalent**, written $\mathcal{KB} \equiv_{pw} \mathcal{KB}'$, if there is a bijection $f : \mathcal{KB} \longrightarrow \mathcal{KB}'$ s.t. each $\alpha \mid\!\sim \beta \in \mathcal{KB}$ is piecewise equivalent to $f(\alpha \mid\!\sim \beta)$.

Given this, we can state a (weak) requirement of syntax independence:

Piecewise Equivalence If $\mathcal{KB}_1 \equiv_{pw} \mathcal{KB}_2$, then $\mid\!\sim_*^{\mathcal{KB}_1} = \mid\!\sim_*^{\mathcal{KB}_2}$.

Weaker notions of equivalence between knowledge bases are also possible. For example, one can adopt the variant which substitutes the following definition of equivalence between conditionals into Definition 4 (see for example the work of Kern-Isberner [31] and of Calabrese [36]).

Definition 5 *(Pairwise Equivalence).* Let $\alpha, \beta, \gamma, \delta \in \mathcal{L}$. We say $\alpha \mid\!\sim \beta$ is **weakly equivalent** to $\gamma \mid\!\sim \delta$ if $\models (\alpha \leftrightarrow \gamma) \wedge ((\alpha \wedge \beta) \leftrightarrow (\gamma \wedge \delta))$. We say that knowledge bases $\mathcal{KB}$ and $\mathcal{KB}'$ are **pairwise equivalent**, written $\mathcal{KB} \equiv_p \mathcal{KB}'$, if there is a bijection $f : \mathcal{KB} \longrightarrow \mathcal{KB}'$ s.t. each $\alpha \mid\!\sim \beta \in \mathcal{KB}$ is weakly equivalent to $f(\alpha \mid\!\sim \beta)$.

This yields a correspondingly stronger syntax-independence requirement:

Pairwise Equivalence If $\mathcal{KB}_1 \equiv_p \mathcal{KB}_2$, then $\mid\!\sim_*^{\mathcal{KB}_1} = \mid\!\sim_*^{\mathcal{KB}_2}$.

Instead of requiring two knowledge bases to be equivalent when their constituent conditionals are individually equivalent on a one-to-one basis, we can adopt a more *global*, more *semantic*, perspective and say that $\mathcal{KB}_1$ and $\mathcal{KB}_2$ are equivalent if they collectively place the same constraints on the set of interval-based interpretations, i.e., they have the same interval-based models. The following proposition says that this is the same as saying they have the same *ranked* models.

Proposition 2. *For any two conditional knowledge bases $\mathcal{KB}_1$ and $\mathcal{KB}_2$, we have $[\![\mathcal{KB}_1]\!]_{\mathscr{I}} = [\![\mathcal{KB}_2]\!]_{\mathscr{I}}$ iff $[\![\mathcal{KB}_1]\!]_{\mathscr{R}} = [\![\mathcal{KB}_2]\!]_{\mathscr{R}}$.*

Proof. The 'only-if' direction is clear since every ranked interpretation is also an interval-based interpretation. For the 'if' direction, suppose $[\![\mathcal{KB}_1]\!]_{\mathscr{I}} \neq [\![\mathcal{KB}_2]\!]_{\mathscr{I}}$ and assume, without loss of generality, that $\mathscr{I} = \langle \mathscr{L}, \mathscr{U} \rangle \in [\![\mathcal{KB}_1]\!]_{\mathscr{I}} \setminus [\![\mathcal{KB}_2]\!]_{\mathscr{I}}$. Choose a conditional $\alpha \mid\!\sim \beta \in \mathcal{KB}_2$ such that $\alpha \not\mid\!\sim_{\mathscr{I}} \beta$. We will construct from $\mathscr{I}$ a ranked interpretation $\mathscr{R} \in [\![\mathcal{KB}_1]\!]_{\mathscr{R}}$ such that $\alpha \not\mid\!\sim_{\mathscr{R}} \beta$. Define $\mathscr{R}$ as follows, for any $v \in \mathcal{U}$:

$$\mathscr{R}(v) = \begin{cases} \mathscr{U}(v), & \text{if } v \in [\![\alpha \wedge \beta]\!]; \\ \mathscr{L}(v), & \text{otherwise.} \end{cases}$$

(Strictly speaking, $\mathscr{R}$ is not a ranked interpretation since it will not, in general, satisfy the convexity property mentioned in Definition 1. However, it should be clear that a non-convex ranked interpretation can always be 'normalised' into an equivalent convex one by essentially ignoring any 'gaps' that appear in the ranking.) Notice that $\mathscr{L}(v) \leq \mathscr{R}(v) \leq \mathscr{U}(v)$ for all $v \in \mathcal{U}$. From this we can show $\mid\!\sim_{\mathscr{I}} \subseteq \mid\!\sim_{\mathscr{R}}$ and so, since $\mathscr{I}$ satisfies $\mathcal{KB}_1$, we have $\mathscr{R} \in [\![\mathcal{KB}_1]\!]_{\mathscr{R}}$. It thus remains to show $\alpha \not\mid\!\sim_{\mathscr{R}} \beta$. Now, for any interval-based interpretation $\mathscr{I}' = \langle \mathscr{L}', \mathscr{U}' \rangle$, and any sentences γ, δ, recall that $\gamma \mid\!\sim_{\mathscr{I}'} \delta$ iff $\mathscr{U}'(\gamma \wedge \delta) < \mathscr{L}'(\gamma \wedge \neg\delta)$. Similarly, for any ranked interpretation $\mathscr{R}'$ we have $\gamma \mid\!\sim_{\mathscr{R}'} \delta$ iff $\mathscr{R}'(\gamma \wedge \delta) < \mathscr{R}'(\gamma \wedge \neg\delta)$. By construction of $\mathscr{R}$ we have $\mathscr{R}(\alpha \wedge \beta) = \mathscr{U}(\alpha \wedge \beta)$ and $\mathscr{R}(\alpha \wedge \neg\beta) = \mathscr{L}(\alpha \wedge \neg\beta)$. Since $\alpha \not\mid\!\sim_{\mathscr{I}} \beta$, we have $\mathscr{U}(\alpha \wedge \beta) \not< \mathscr{L}(\alpha \wedge \neg\beta)$, so $\mathscr{R}(\alpha \wedge \beta) \not< \mathscr{R}(\alpha \wedge \neg\beta)$, i.e., $\alpha \not\mid\!\sim_{\mathscr{R}} \beta$ as required. □

It is straightforward to show that, for any knowledge bases $\mathcal{KB}_1$ and $\mathcal{KB}_2$, we have that $\mathcal{KB}_1 \equiv_p \mathcal{KB}_2$ implies $[\![\mathcal{KB}_1]\!]_{\mathscr{R}} = [\![\mathcal{KB}_2]\!]_{\mathscr{R}}$ (but not conversely). This semantic notion of equivalence between knowledge bases thus gives rise to the following equivalence postulate, which is the strongest one that we consider in this paper.

Global Equivalence If $[\![\mathcal{KB}_1]\!]_{\mathscr{R}} = [\![\mathcal{KB}_2]\!]_{\mathscr{R}}$, then $\mid\!\sim_*^{\mathcal{KB}_1} = \mid\!\sim_*^{\mathcal{KB}_2}$.

Having agreed on a set of reasonable basic postulates and some equivalence postulates, in what follows we assess the formalisation of notions of 'caution' in conclusion drawing.

4.3. Conservativeness 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 8. 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

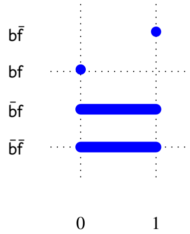

Fig. 3. Interval-based model of $\mathcal{KB} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}\}$.

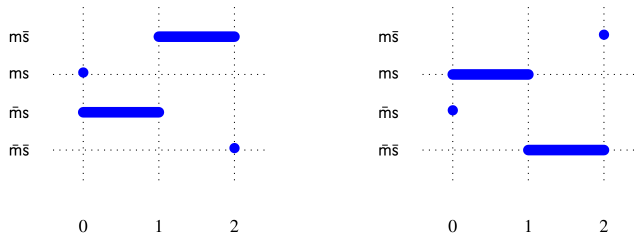

Fig. 4. Interval-based models of the two $\subseteq$-minimal extensions of $\mathrel{|\!\sim}^{\mathcal{KB}}$, for $\mathcal{KB} = \{\mathsf{m} \mathrel{|\!\sim} \mathsf{s}, \neg\mathsf{m} \mathrel{|\!\sim} \mathsf{s}\}$.

interval-based interpretation $\mathscr{I}$ in Fig. 3. Indeed, the conditional $\mathsf{b} \mathrel{|\!\sim} \mathsf{f}$ is saying precisely that $\mathsf{bf} \prec \mathsf{b}\bar{\mathsf{f}}$, 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 $\neg\mathsf{f} \mathrel{|\!\sim}_{\mathscr{I}} \neg\mathsf{b}$ here. Notice that the rational closure in this example *does* endorse this latter conclusion, thus providing further evidence that the rational closure arguably gives some unwarranted conclusions.

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

Example 9. Assume a COVID-19 inspired scenario with only two propositions, m and s, standing for, respectively, "you wear a face mask" and "you observe social distancing". Let $\mathcal{KB} = \{\mathsf{m} \mathrel{|\!\sim} \mathsf{s}, \neg\mathsf{m} \mathrel{|\!\sim} \mathsf{s}\}$. There are two $\subseteq$-minimal disjunctive consequence relations extending $\mathrel{|\!\sim}^{\mathcal{KB}}$, corresponding to the two interval-based interpretations $\mathscr{I}_1$ and $\mathscr{I}_2$ (from left to right) in Fig. 4. The first conditional is saying $\mathsf{ms} \prec \mathsf{m}\bar{\mathsf{s}}$, while the second is saying $\bar{\mathsf{m}}\mathsf{s} \prec \bar{\mathsf{m}}\bar{\mathsf{s}}$. According to the interval condition (see (1) in 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.

In the light of Example 9 above, a question that arises is what to do when one has more than a single $\subseteq$-minimal extension of $\mathrel{|\!\sim}^{\mathcal{KB}}$. Proposition 1 already tells us we cannot, in general, take the obvious approach by taking their intersection. However, even though returning the disjunctive/preferential closure $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ is not enough to ensure D-Rationality, we might still expect the following postulate as reasonable.

Vacuity If $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ is disjunctive, then $\mathrel{|\!\sim}^{\mathcal{KB}}_{*} = \mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$.

One property we obviously do not want to enforce on $\mathrel{|\!\sim}^{\mathcal{KB}}_{*}$ is Monotonicity, i.e., the requirement that $\alpha \wedge \gamma \mathrel{|\!\sim}^{\mathcal{KB}}_{*} \beta$, for any γ, whenever $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{*} \beta$. But are there any weaker versions that are justifiably worth retaining? Under the assumption of the basic postulate D-Rationality, we already obtain one such weakening, viz. the KLM property CM from Section 2.1. We now suggest another restricted version of Monotonicity — one that explicitly relates to the knowledge base $\mathcal{KB}$ under consideration. We first need some more notation: Given a conditional knowledge base $\mathcal{KB}$, let us denote by $\mathit{ante}(\mathcal{KB})$ the set of all antecedents of the conditionals in $\mathcal{KB}$. We put forward the following postulate:

Knowledge-based Monotonicity If $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{*} \beta$, then $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mathrel{|\!\sim}^{\mathcal{KB}}_{*} \beta$.

For an informal justification of this postulate, suppose we have $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{*} \beta$. Then that means that given α, we tentatively infer *on the basis of* $\mathcal{KB}$ that β holds. If we learn that, in addition to α, at least one (but without knowing which) of the antecedents of the conditionals in $\mathcal{KB}$ holds, i.e., that at least one of the conditionals in $\mathcal{KB}$ "fires", then this information should not invalidate the conclusion β.

Knowledge-based Monotonicity (KBM) above expresses a specific kind of cautiousness in the inferences we make from $\mathcal{KB}$: Given α, we should not draw any conclusion β that can be "defeated" by adding $\bigvee \mathit{ante}(\mathcal{KB})$. This behaviour is in keeping with how the preferential closure works.

Proposition 3. *The preferential closure satisfies Knowledge-based Monotonicity.*

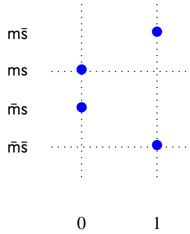

Fig. 5. Interval-based model of the union of the two $\subseteq$-minimal extensions of $\mid\!\sim^{\mathcal{KB}}$, for $\mathcal{KB} = \{\mathsf{m} \mid\!\sim \mathsf{s}, \neg\mathsf{m} \mid\!\sim \mathsf{s}\}$.

Proof. We make use of the known characterisation [5] of preferential closure of a knowledge base $\mathcal{KB}$ in terms of derivations from $\mathcal{KB}$ using the inference rules of preferential reasoning given in Section 2.1. Specifically, we know that for any $\mathcal{KB}$ and sentences α, β we have that $\alpha \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ iff there exists a derivation of $\alpha \mid\!\sim \beta$ from $\mathcal{KB}$ using the rules Ref, And, RW, LLE, Or, and CM. We show that $\mid\!\sim^{\mathcal{KB}}_{PC}$ satisfies KBM by induction on the length i of the derivation of $\alpha \mid\!\sim \beta$.

For the base case, assume $i = 1$. Then either $\alpha \mid\!\sim \beta$ is an instance of Ref $\alpha \mid\!\sim \alpha$, or we have $\alpha \mid\!\sim \beta \in \mathcal{KB}$ (and hence $\alpha \in \mathit{ante}(\mathcal{KB})$). In the first case, we get $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \alpha$ by Ref and RW for $\mid\!\sim^{\mathcal{KB}}_{PC}$, while in the second case we obtain $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ by LLE from $\alpha \mid\!\sim^{\mathcal{KB}}_{PC} \beta$, since $\models (\alpha \leftrightarrow (\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB})))$.

Now for the inductive case let $i > 1$ and assume that, for all $j < i$ and all $\eta, \rho \in \mathcal{L}$, whenever there exists a derivation of length j of $\eta \mid\!\sim \rho$ from $\mathcal{KB}$, we have $\eta \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim \rho$. Suppose we have a derivation of length i of $\alpha \mid\!\sim \beta$ from $\mathcal{KB}$. We go through each of the possible last rules applied in the derivation.

And. Then $\beta = \gamma \wedge \delta$ and we must have derivations of $\alpha \mid\!\sim \gamma$ and $\alpha \mid\!\sim \delta$, both of length less than i. By induction we know $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \gamma$ and $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \delta$ and so we obtain the required $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \gamma \wedge \delta$ by And for $\mid\!\sim^{\mathcal{KB}}_{PC}$.

RW. Then we must have a derivation of $\alpha \mid\!\sim \gamma$ of length $i-1$, for some γ such that $\models \gamma \rightarrow \beta$. By induction we know $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \gamma$ and so we obtain $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ by RW for $\mid\!\sim^{\mathcal{KB}}_{PC}$.

LLE. Then we must have a derivation of $\gamma \mid\!\sim \beta$ of length $i-1$, for some γ such that $\models \alpha \leftrightarrow \gamma$. By induction we know $\gamma \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ and so we obtain $\alpha \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ by LLE for $\mid\!\sim^{\mathcal{KB}}_{PC}$.

Or. Then $\alpha = \gamma \vee \delta$ and we must have derivations of $\gamma \mid\!\sim \beta$ and $\delta \mid\!\sim \beta$, both of length less than i. By induction we know $\gamma \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ and $\delta \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ and so we obtain the required $(\gamma \vee \delta) \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ using Or and LLE for $\mid\!\sim^{\mathcal{KB}}_{PC}$.

CM. Then $\alpha = \gamma \wedge \delta$ and we must have derivations of $\gamma \mid\!\sim \delta$ and $\gamma \mid\!\sim \beta$, both of length less than i. By induction we know $\gamma \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \delta$ and $\gamma \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ and so we obtain the required $(\gamma \wedge \delta) \wedge \bigvee \mathit{ante}(\mathcal{KB}) \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ using CM and LLE for $\mid\!\sim^{\mathcal{KB}}_{PC}$. $\square$

Returning to Example 8 above, KBM goes some way to blocking the inference $\neg\mathsf{f} \mid\!\sim \neg\mathsf{b}$ from $\mathcal{KB} = \{\mathsf{b} \mid\!\sim \mathsf{f}\}$, since in this case if we had $\neg\mathsf{f} \mid\!\sim^{\mathcal{KB}}_{*} \neg\mathsf{b}$, then KBM would force $\neg\mathsf{f} \wedge \mathsf{b} \mid\!\sim^{\mathcal{KB}}_{*} \neg\mathsf{b}$ and so $\neg\mathsf{f} \wedge \mathsf{b} \mid\!\sim^{\mathcal{KB}}_{*} \bot$ by D-Rationality, i.e., non-flying birds are not just exceptional, but *impossible*, which seems too strong an inference to make from this $\mathcal{KB}$. The rational closure, while endorsing the inference of $\neg\mathsf{f} \mid\!\sim \neg\mathsf{b}$ from this $\mathcal{KB}$, does *not* endorse $\neg\mathsf{f} \wedge \mathsf{b} \mid\!\sim^{\mathcal{KB}}_{RC} \bot$. Thus rational closure does *not* satisfy KBM. Indeed this example shows it does not even satisfy the special case in which $\mathcal{KB}$ consists of a single conditional. Despite this we will later see that KBM is still attainable for inference operators $*$ satisfying D-Rationality.

4.4. Representation independence postulates

In this subsection, we recast notions of representation independence as recently studied in the literature in the form of postulates.

Going back to Example 9, 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 that the right thing to do is to include both, and thereby let the interval-based interpretation depicted in Fig. 5 yield the output. Notice that this will be the same as the rational closure in this case.

We can express the desired symmetry requirement in a syntactic form, using the notion of *symbol translations* [37]. 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 σ, 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 σ, we have $\alpha \mid\!\sim^{\mathcal{KB}}_{*} \beta$ iff $\sigma(\alpha) \mid\!\sim^{\sigma(\mathcal{KB})}_{*} \sigma(\beta)$.

[2] Marquis and Schwind [37] consider much more general settings, but this is all we need in the present paper.

Note that Weydert [7] 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 [26], who, in particular, looked at the property in relation to rational closure. As noted by Marquis and Schwind [37], 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 by the following result showing RI is jointly incompatible with two of our basic postulates, namely Inclusion and Infra-Rationality.

Theorem 3. *There is no method $*$ satisfying all of Inclusion, Infra-Rationality and Representation Independence.*

Proof. Assume for contradiction that $*$ satisfies all three properties. Let $\mathcal{KB} = \{\top \mid\!\sim q\}$. By Inclusion, $\top \mid\!\sim_*^{\mathcal{KB}} q$. Now consider a symbol translation σ_1 such that $\sigma_1(p) = \top$ and $\sigma_1(q) = q$. Clearly $\sigma_1(\mathcal{KB}) = \mathcal{KB}$, so from $\top \mid\!\sim_*^{\mathcal{KB}} q$ we get $\top \mid\!\sim_*^{\sigma_1(\mathcal{KB})} q$, i.e., $\sigma_1(p) \mid\!\sim_*^{\sigma_1(\mathcal{KB})} \sigma_1(q)$. Then, from this and Representation Independence, we get $p \mid\!\sim_*^{\mathcal{KB}} q$. Now consider another symbol translation σ_2 such that $\sigma_2(p) = p$ and $\sigma_2(q) = \neg p$. Note that $\sigma_2(\mathcal{KB}) = \{\top \mid\!\sim \neg p\}$. From $p \mid\!\sim_*^{\mathcal{KB}} q$ and Representation Independence we have $\sigma_2(p) \mid\!\sim_*^{\sigma_2(\mathcal{KB})} \sigma_2(q)$, i.e., $p \mid\!\sim_*^{\sigma_2(\mathcal{KB})} \neg p$, and so, by Infra-Rationality, $p \mid\!\sim_{RC}^{\sigma_2(\mathcal{KB})} \neg p$. But $p \mid\!\sim \neg p$ is not inferred from $\sigma_2(\mathcal{KB}) = \{\top \mid\!\sim \neg p\}$ by the rational closure, which gives a contradiction. □

We remark that, since rational closure satisfies Inclusion and Infra-Rationality, the above result shows, as a corollary, that rational closure itself does not satisfy Representation Independence.

The preceding theorem motivates the need to focus on specific families of symbol translation. Some examples are the following:

1. σ is a permutation on $\mathcal{P}$, i.e., is just a *renaming* of the propositional variables[3];
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.
3. A family of symbol translations that includes both of the above as special cases is as follows. Call σ *belief amount preserving (BAP)* [37] if there exists a permutation R_σ on the set $\mathcal{U}$ of valuations such that, for all $\alpha \in \mathcal{L}$, $[\![\sigma(\alpha)]\!] = \{R_\sigma(v) \mid v \in [\![\alpha]\!]\}$. If σ is a permutation on $\mathcal{P}$ (as in the first case above) then $R_\sigma(v)$ is the valuation such that $[R_\sigma(v)](p) = v(\sigma(p))$, while if σ is a negation-swapping symbol translation (as in the second case) then

$$[R_\sigma(v)](p) = \begin{cases} v(p), & \text{if } \sigma(p) = p; \\ v(\neg p), & \text{if } \sigma(p) = \neg p, \end{cases}$$

i.e., $R_\sigma(v)$ 'flips' the truth-values in v of all those p such that $\sigma(p) = \neg p$. An example of a BAP symbol translation that does not fall under the above two cases would be one such that (assuming $\mathcal{P} = \{p, q\}$) $\sigma(p) = p$ and $\sigma(q) = p \leftrightarrow q$.

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 two postulates:

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

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

Since every negation-swapping symbol translation is a BAP symbol translation, we have that BAP Representation Independence implies Negated Representation Independence. A property in the same spirit as BAP Representation Independence has been studied by Paris and Vencovská [38] in the setting of inference from probabilistic knowledge bases. The property, which went under the name *Renaming Principle*, was shown to be a characteristic property of the *maximum entropy* inference process [39].

Example 10. Going back to Example 9, 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 $\sigma(\mathcal{KB})$.

4.5. Cumulativity postulates

In this subsection, we state our last two postulates and conclude the section with an impossibility result concerning a subset of the postulates we have motivated so far.

[3] Symbol translations of this type were already considered by Lehmann and Magidor, who showed [5, Lemma 5.10] that rational closure satisfies RI when we restrict σ to be a renaming.

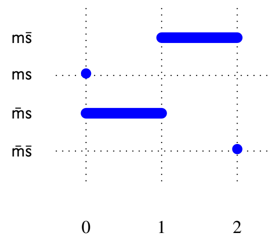

Fig. 6. Disjunctive/preferential closure of $\mid\!\sim^{\mathcal{KB}}$, for $\mathcal{KB} = \{\mathsf{m} \mid\!\sim \mathsf{s}, \neg\mathsf{m} \mid\!\sim \mathsf{s}\}$.

The idea behind the 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', formalised as the following postulates:

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

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

Theorem 4 *(Impossibility Result). There is no method $*$ satisfying all of Inclusion, D-Rationality, Piecewise 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 9, 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 that $\mid\!\sim_*^{\mathcal{KB}}$ satisfies the Or postulate, 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:

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

Now, let σ be the negation-swapping symbol translation mentioned in Example 10, 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 $\mathcal{KB} \equiv_{pw} \sigma(\mathcal{KB})$, so, by Piecewise Equivalence, we obtain from this:

$$(\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s} \ \text{ iff } \ (\neg\mathsf{m} \leftrightarrow \mathsf{s}) \mid\!\sim_*^{\mathcal{KB}} \mathsf{s} \tag{3}$$

Putting (2) and (3) 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 Fig. 4, reproduced in Fig. 6 for the reader's convenience. Hence, by Vacuity, this particular interval-based interpretation corresponds also to $\mid\!\sim_*^{\mathcal{KB}'}$. However, 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 4 is both surprising and disappointing, since all of the properties mentioned seem to be rather intuitive and desirable. Piecewise Equivalence and Negated Representation Independence are even both relatively weak formulations of equivalence and representation independence, respectively, as the discussion in Sections 4.2 and 4.4 indicates. Also, 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 construct a method that can satisfy as many of these properties as possible. In the next section, we shall provide our candidate for such a method — the disjunctive rational closure.

5. A construction for disjunctive rational closure

In this section, we present a constructive definition of the disjunctive rational closure of a conditional knowledge base and illustrate its behaviour with the scenario examples we have seen earlier.

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 what we shall call the *disjunctive rational closure* of a conditional knowledge base.

To characterise $\mathscr{I}_{DC}^{\mathcal{KB}}$, 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 every valuation $v \in \mathcal{U}$:

$$\mathscr{L}^{\mathcal{KB}}_{DC}(v) \stackrel{\text{def}}{=} \mathscr{R}^{\mathcal{KB}}_{RC}(v) \tag{4}$$

That is, the lower ranks are given by the rational closure of the underlying conditional knowledge base.

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

Definition 6 *(Min-verification of a Conditional).* Given a ranked interpretation $\mathscr{R}$ and a conditional $\alpha \mathrel{|\!\sim} \beta$ such that $\mathscr{R} \Vdash \alpha \mathrel{|\!\sim} \beta$, we say a valuation v **min-verifies** $\alpha \mathrel{|\!\sim} \beta$ **in** $\mathscr{R}$ if $v \Vdash \alpha$ and $\mathscr{R}(v) = \mathscr{R}(\alpha)$.

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

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

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

Case 2: The valuation v min-verifies at least one conditional from $\mathcal{KB}$ in $\mathscr{R}^{\mathcal{KB}}_{RC}$. 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 $\mathscr{I}^{\mathcal{KB}}_{DC}$. If v min-verifies $\alpha \mathrel{|\!\sim} \beta$ in $\mathscr{R}^{\mathcal{KB}}_{RC}$, then this is achieved by setting $\mathscr{U}^{\mathcal{KB}}_{DC}(v) = \mathscr{R}^{\mathcal{KB}}_{RC}(\alpha \wedge \neg\beta) - 1$; or, if $\mathscr{R}(\alpha \wedge \neg\beta) = \infty$, then again just set $\mathscr{U}^{\mathcal{KB}}_{DC}(v) = \max\{\mathscr{R}^{\mathcal{KB}}_{RC}(u) \mid \mathscr{R}^{\mathcal{KB}}_{RC}(u) \neq \infty\}$, as in Case 1. We introduce now the following notation. Given sentences α, β:

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

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

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

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

$$\mathscr{U}^{\mathcal{KB}}_{DC}(v) \stackrel{\text{def}}{=} \begin{cases} \min\{t^{\mathcal{KB}}_{RC}(\alpha,\beta) \mid \alpha \mathrel{|\!\sim} \beta \in \mathcal{KB} \text{ and } v \text{ min-verifies } \alpha \mathrel{|\!\sim} \beta \text{ in } \mathscr{R}^{\mathcal{KB}}_{RC}\}, \\ \qquad \text{if } v \text{ min-verifies at least one conditional from } \mathcal{KB} \text{ in } \mathscr{R}^{\mathcal{KB}}_{RC}; \\ \max\{\mathscr{R}^{\mathcal{KB}}_{RC}(u) \mid \mathscr{R}^{\mathcal{KB}}_{RC}(u) \neq \infty\}, \text{otherwise.} \end{cases} \tag{5}$$

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

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

Definition 7 *(Disjunctive Rational Closure).* Let $\mathscr{I}^{\mathcal{KB}}_{DC} \stackrel{\text{def}}{=} \langle \mathscr{L}^{\mathcal{KB}}_{DC}, \mathscr{U}^{\mathcal{KB}}_{DC} \rangle$ be the interval-based interpretation specified by $\mathscr{L}^{\mathcal{KB}}_{DC}$ and $\mathscr{U}^{\mathcal{KB}}_{DC}$ as in (4) and (5) above, respectively. The **disjunctive rational closure** (hereafter DRC) of $\mathcal{KB}$ is the consequence relation $\mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \stackrel{\text{def}}{=} \{\alpha \mathrel{|\!\sim} \beta \mid \mathscr{I}^{\mathcal{KB}}_{DC} \Vdash \alpha \mathrel{|\!\sim} \beta\}$.

Next, we revisit the examples we have seen throughout the paper, to see what answer the disjunctive rational closure gives.

Example 11. Going back to Example 8, with $\mathcal{KB} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}\}$, the rational closure yields $\mathscr{R}^{\mathcal{KB}}_{RC}(\mathsf{bf}) = \mathscr{R}^{\mathcal{KB}}_{RC}(\bar{\mathsf{b}}\mathsf{f}) = \mathscr{R}^{\mathcal{KB}}_{RC}(\bar{\mathsf{b}}\bar{\mathsf{f}}) = 0$ and $\mathscr{R}^{\mathcal{KB}}_{RC}(\mathsf{b}\bar{\mathsf{f}}) = 1$. Since $\mathscr{L}^{\mathcal{KB}}_{DC} = \mathscr{R}^{\mathcal{KB}}_{RC}$, this gives us the lower ranks for each valuation in $\mathscr{I}^{\mathcal{KB}}_{DC}$. Turning to the upper ranks, the only valuation that min-verifies the single conditional $\mathsf{b} \mathrel{|\!\sim} \mathsf{f}$ in $\mathcal{KB}$ is bf, thus $\mathscr{U}^{\mathcal{KB}}_{DC}(\mathsf{bf}) = t^{\mathcal{KB}}_{RC}(\mathsf{b},\mathsf{f}) = \mathscr{R}^{\mathcal{KB}}_{RC}(\mathsf{b} \wedge \neg\mathsf{f}) - 1 = 1 - 1 = 0$, 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 $\mathscr{R}^{\mathcal{KB}}_{RC}$, 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 $\mathscr{I}^{\mathcal{KB}}_{DC}$ outputs exactly the same interval-based interpretation depicted in Fig. 3 (reproduced in Fig. 7) which, recall, gives the unique $\subseteq$-minimal disjunctive consequence relation extending $\mathcal{KB}$ in this case.

Example 12. Returning to Example 9, with $\mathcal{KB} = \{\mathsf{m} \mathrel{|\!\sim} \mathsf{s}, \neg\mathsf{m} \mathrel{|\!\sim} \mathsf{s}\}$, the rational closure yields $\mathscr{R}^{\mathcal{KB}}_{RC}(\mathsf{ms}) = \mathscr{R}^{\mathcal{KB}}_{RC}(\bar{\mathsf{m}}\mathsf{s}) = 0$ and $\mathscr{R}^{\mathcal{KB}}_{RC}(\mathsf{m}\bar{\mathsf{s}}) = \mathscr{R}^{\mathcal{KB}}_{RC}(\bar{\mathsf{m}}\bar{\mathsf{s}}) = 1$, which gives us the lower ranks. The valuation ms min-verifies only the conditional $\mathsf{m} \mathrel{|\!\sim} \mathsf{s}$, and so $\mathscr{U}^{\mathcal{KB}}_{DC}(\mathsf{ms}) = t^{\mathcal{KB}}_{RC}(\mathsf{m},\mathsf{s}) = \mathscr{R}^{\mathcal{KB}}_{RC}(\mathsf{m} \wedge \neg\mathsf{s}) - 1 = 1 - 1 = 0$. Similarly, the valuation $\bar{\mathsf{m}}\mathsf{s}$ min-verifies only the conditional $\neg\mathsf{m} \mathrel{|\!\sim} \mathsf{s}$ and so, by analogous reasoning, $\mathscr{U}^{\mathcal{KB}}_{DC}(\bar{\mathsf{m}}\mathsf{s}) = t^{\mathcal{KB}}_{RC}(\neg\mathsf{m},\mathsf{s}) = 0$. So both of these valuations are assigned the interval $(0,0)$ by $\mathscr{I}^{\mathcal{KB}}_{DC}$. The other two

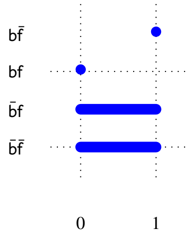

Fig. 7. Interval-based model of $\mathcal{KB} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}\}$.

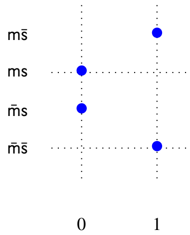

Fig. 8. Rational closure of $\mathcal{KB} = \{\mathsf{m} \mathrel{|\!\sim} \mathsf{s}, \neg\mathsf{m} \mathrel{|\!\sim} \mathsf{s}\}$.

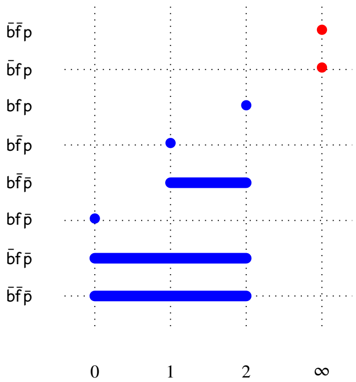

Fig. 9. Disjunctive rational closure of $\mathcal{KB} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \rightarrow \mathsf{b}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}\}$.

valuations, which min-verify neither conditional in $\mathcal{KB}$, are assigned $(1, 1)$. Thus, in this case, $\mathscr{I}^{\mathcal{KB}}_{DC}$ returns just the rational closure of $\mathcal{KB}$, as pictured in Fig. 8.

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

Example 13. 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 $\mathscr{R}^{\mathcal{KB}}_{RC}$ for this $\mathcal{KB}$ is depicted in Fig. 1. Since both of the valuations in $[\![\mathsf{p} \land \neg\mathsf{b}]\!]$ (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 min-verifying $\mathsf{b} \mathrel{|\!\sim} \mathsf{f}$ in $\mathscr{R}^{\mathcal{KB}}_{RC}$ is $\mathsf{b}\mathsf{f}\bar{\mathsf{p}}$ (which min-verifies no other conditional in $\mathcal{KB}$), so $\mathscr{U}^{\mathcal{KB}}_{DC}(\mathsf{b}\mathsf{f}\bar{\mathsf{p}}) = \mathscr{R}^{\mathcal{KB}}_{RC}(\mathsf{b} \land \neg\mathsf{f}) - 1 = 1 - 1 = 0$. The only valuation min-verifying $\mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}$ is $\mathsf{b}\bar{\mathsf{f}}\mathsf{p}$, so $\mathscr{U}^{\mathcal{KB}}_{DC}(\mathsf{b}\bar{\mathsf{f}}\mathsf{p}) = \mathscr{R}^{\mathcal{KB}}_{RC}(\mathsf{p} \land \mathsf{f}) - 1 = 2 - 1 = 1$. All other plausible valuations get assigned as their upper rank the maximum finite rank, which is 2. The resulting $\mathscr{I}^{\mathcal{KB}}_{DC}$ is the interval-based interpretation depicted in Fig. 9.

We conclude this section by considering our construction from the computational complexity standpoint.

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 min-verifies at least one conditional $\alpha \mathrel{|\!\sim} \beta$ in $\mathcal{KB}$ (cf. Definition 6). 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 $\mathscr{R}^{\mathcal{KB}}_{RC}$ 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 min-verifies a conditional $\alpha \mathrel{|\!\sim} \beta$, the computation of $t^{\mathcal{KB}}_{RC}(\cdot)$ also requires that of $\mathscr{R}^{\mathcal{KB}}_{RC}(\alpha \wedge \neg\beta)$. In the worst case, the latter requires $2^{|\mathcal{P}|}$ propositional verifications. So, the computation of $t^{\mathcal{KB}}_{RC}(\cdot)$ 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 $\mathscr{I}^{\mathcal{KB}}_{DC}$. This task amounts to computing $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha)$ and $\mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg\beta)$ 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 4. *DRC satisfies Inclusion, D-Rationality and Infra-Rationality.*

Proof. D-Rationality is immediate since we construct an interval-based interpretation. For Infra-Rationality, first recall that $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \beta$ iff $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha) < \mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg\beta)$. Since $\mathscr{L}^{\mathcal{KB}}_{DC}(\alpha) \leq \mathscr{U}^{\mathcal{KB}}_{DC}(\alpha)$ (follows by definition of interval-based interpretation) and $\mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg\beta) = \mathscr{R}^{\mathcal{KB}}_{RC}(\alpha \wedge \neg\beta)$ (by construction), we have $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha) < \mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg\beta)$ implies $\mathscr{R}^{\mathcal{KB}}_{RC}(\alpha) = \mathscr{L}^{\mathcal{KB}}_{DC}(\alpha) < \mathscr{R}^{\mathcal{KB}}_{RC}(\alpha \wedge \neg\beta)$, giving $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{RC} \beta$, 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' min-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, DRC also satisfies Preferential Extension.

Regarding the equivalence postulates, we can show that DRC satisfies Pairwise (and hence also Piecewise) Equivalence.

Proposition 5. *DRC satisfies Pairwise Equivalence.*

Proof. Suppose $\mathcal{KB}_1 = \{\alpha_i \mathrel{|\!\sim} \beta_i \mid i = 1, \ldots, n\}$ and let $\mathcal{KB}_2$ be such that $\mathcal{KB}_1 \equiv_p \mathcal{KB}_2$. Then $\mathcal{KB}_2 = \{\gamma_i \mathrel{|\!\sim} \delta_i \mid i = 1, \ldots, n\}$ where, for each $i = 1, \ldots, n$ we have $\models (\alpha_i \leftrightarrow \gamma_i) \wedge ((\alpha_i \wedge \beta_i) \leftrightarrow (\gamma_i \wedge \delta_i))$. It suffices to show, for all $v \in \mathcal{U}$, (*i*) $\mathscr{L}^{\mathcal{KB}_1}_{DC}(v) = \mathscr{L}^{\mathcal{KB}_2}_{DC}(v)$, and (*ii*) $\mathscr{U}^{\mathcal{KB}_1}_{DC}(v) = \mathscr{U}^{\mathcal{KB}_2}_{DC}(v)$.

(*i*) $\mathscr{L}^{\mathcal{KB}_1}_{DC}(v) = \mathscr{L}^{\mathcal{KB}_2}_{DC}(v)$. By construction this is the same as $\mathscr{R}^{\mathcal{KB}_1}_{RC}(v) = \mathscr{R}^{\mathcal{KB}_2}_{RC}(v)$, which holds since the rational closure clearly satisfies Global Equivalence, and hence Pairwise Equivalence.

(*ii*) $\mathscr{U}^{\mathcal{KB}_1}_{DC}(v) = \mathscr{U}^{\mathcal{KB}_2}_{DC}(v)$. For each $i = 1, \ldots, n$, since $\mathscr{R}^{\mathcal{KB}_1}_{RC} = \mathscr{R}^{\mathcal{KB}_2}_{RC}$ and $\models (\alpha_i \leftrightarrow \gamma_i)$ we have that v min-verifies $\alpha_i \mathrel{|\!\sim} \beta_i$ in $\mathscr{R}^{\mathcal{KB}_1}_{RC}$ iff v min-verifies $\gamma_i \mathrel{|\!\sim} \delta_i$ in $\mathscr{R}^{\mathcal{KB}_2}_{RC}$. Thus, if v min-verifies no conditional from $\mathcal{KB}_1$ in $\mathscr{R}^{\mathcal{KB}_1}_{RC}$ then it also min-verifies no conditional from $\mathcal{KB}_2$ in $\mathscr{R}^{\mathcal{KB}_2}_{RC}$ and in this case

$$\begin{aligned}
\mathscr{U}^{\mathcal{KB}_1}_{DC}(v) &= \max\{\mathscr{R}^{\mathcal{KB}_1}_{RC}(u) \mid \mathscr{R}^{\mathcal{KB}_1}_{RC}(u) \neq \infty\} \\
&= \max\{\mathscr{R}^{\mathcal{KB}_2}_{RC}(u) \mid \mathscr{R}^{\mathcal{KB}_2}_{RC}(u) \neq \infty\} \\
&= \mathscr{U}^{\mathcal{KB}_2}_{DC}(v).
\end{aligned}$$

Otherwise

$$\begin{aligned}
\mathscr{U}^{\mathcal{KB}_1}_{DC}(v) &= \min\{t^{\mathcal{KB}_1}_{RC}(\alpha_i, \beta_i) \mid v \text{ min-verifies } \alpha_i \mathrel{|\!\sim} \beta_i \text{ in } \mathscr{R}^{\mathcal{KB}_1}_{RC}\} \\
&= \min\{t^{\mathcal{KB}_1}_{RC}(\alpha_i, \beta_i) \mid v \text{ min-verifies } \gamma_i \mathrel{|\!\sim} \delta_i \text{ in } \mathscr{R}^{\mathcal{KB}_2}_{RC}\}.
\end{aligned}$$

But, for each $i = 1, \ldots, n$ we have $\models ((\alpha_i \wedge \neg\beta_i) \leftrightarrow (\gamma_i \wedge \neg\delta_i))$ and thus $\mathscr{R}^{\mathcal{KB}_1}_{RC}(\alpha_i \wedge \neg\beta_i) = \mathscr{R}^{\mathcal{KB}_2}_{RC}(\gamma_i \wedge \neg\delta_i)$. This is enough to show $t^{\mathcal{KB}_1}_{RC}(\alpha_i, \beta_i) = t^{\mathcal{KB}_2}_{RC}(\gamma_i, \delta_i)$ and so

$$\begin{aligned}
\mathscr{U}^{\mathcal{KB}_1}_{DC}(v) &= \min\{t^{\mathcal{KB}_1}_{RC}(\alpha_i, \beta_i) \mid v \text{ min-verifies } \gamma_i \mathrel{|\!\sim} \delta_i \text{ in } \mathscr{R}^{\mathcal{KB}_2}_{RC}\} \\
&= \min\{t^{\mathcal{KB}_2}_{RC}(\gamma_i, \delta_i) \mid v \text{ min-verifies } \gamma_i \mathrel{|\!\sim} \delta_i \text{ in } \mathscr{R}^{\mathcal{KB}_2}_{RC}\} \\
&= \mathscr{U}^{\mathcal{KB}_2}_{DC}(v)
\end{aligned}$$

as required. □

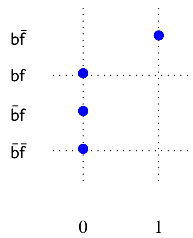

Fig. 10. Output for $\mathcal{KB}' = \{\mathsf{b} \mid\!\sim \mathsf{f}, \top \mid\!\sim (\mathsf{b} \to \mathsf{f})\}$.

DRC does *not*, however, satisfy our strongest equivalence postulate, namely Global Equivalence. This will be confirmed below when we talk about Cut.

We can also confirm that DRC conforms with BAP Representation Independence, and hence also Negated Representation Independence (the proof of this result is a little more involved and is deferred to Appendix B).

Proposition 6. *DRC satisfies BAP Representation Independence.*

DRC essentially inherits this property from rational closure, which also satisfies it, as shown in the proof of the above result in the appendix. Although Jaeger [26] showed that rational closure conforms with his version of Representation Independence, the relationship between his version and ours remains to be explored.

Now we look at the Cumulativity properties. It is known from the work by Lehmann and Magidor [5] that rational closure satisfies both Cautious Monotonicity and Cut, and, in fact, if $\alpha \mid\!\sim^{\mathcal{KB}}_{RC} \beta$ and $\mathcal{KB}' = \mathcal{KB} \cup \{\alpha \mid\!\sim \beta\}$, then $\mathscr{R}^{\mathcal{KB}}_{RC} = \mathscr{R}^{\mathcal{KB}'}_{RC}$. We can show the following for DRC.

Proposition 7. *DRC satisfies Cautious Monotonicity, but does not satisfy Cut.*

Proof. Turning first to Cut, assume $\mathcal{P} = \{\mathsf{b}, \mathsf{f}\}$, and $\mathcal{KB}$ is again the knowledge base from Example 1, i.e., $\{\mathsf{b} \mid\!\sim \mathsf{f}\}$. We have seen in Example 11 that $\mathscr{I}^{\mathcal{KB}}_{DC}$ is given by the interval-based interpretation depicted in Fig. 3. By inspecting this picture, we see $\mathscr{I}^{\mathcal{KB}}_{DC} \Vdash \top \mid\!\sim (\mathsf{b} \to \mathsf{f})$. Now let $\mathcal{KB}' = \mathcal{KB} \cup \{\top \mid\!\sim (\mathsf{b} \to \mathsf{f})\}$. Then $\mathscr{I}^{\mathcal{KB}'}_{DC}$ is given by the model in Fig. 10. We now have $\mathscr{I}^{\mathcal{KB}'}_{DC} \Vdash \neg\mathsf{f} \mid\!\sim \neg\mathsf{b}$, whereas before we had $\mathscr{I}^{\mathcal{KB}}_{DC} \nVdash \neg\mathsf{f} \mid\!\sim \neg\mathsf{b}$.

Regarding Cautious Monotonicity, it is easy to see that the addition of an inferred conditional to a knowledge base will never lead to an *increase* in the upper ranks, which means DRC *does* satisfy Cautious Monotonicity. □

The reason for the failure of Cut is that by adding a new conditional $\alpha \mid\!\sim \beta$ to $\mathcal{KB}$, even when that conditional is already inferred by DRC, we give certain valuations (i.e., those in $[\![\alpha]\!]$) opportunity to min-verify one more conditional from the knowledge base in $\mathscr{R}^{\mathcal{KB}'}_{RC}$. This leads, potentially, to a corresponding decrease in their upper ranks $\mathscr{U}^{\mathcal{KB}}_{DC}$, leading in turn to more inferences being made available. This behaviour reveals that DRC 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, as mentioned in the above proof, adding an inferred conditional will never lead to an increase in the upper ranks, and so DRC does satisfy Cautious Monotonicity.

In the counterexample to Cut given in the proof of Proposition 7, the conditional $\top \mid\!\sim (\mathsf{b} \to \mathsf{f})$ being added to $\mathcal{KB}$ isn't just in the DRC of $\mathcal{KB}$, it is in the *preferential closure* of $\mathcal{KB}$. Thus our proof actually shows the following variant of Cut fails, which is weaker in the presence of Preferential Extension (which, as earlier remarked, itself follows from D-Rationality and Inclusion):

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

As well as being a weakening of Cut, Preferential Cut doubles as a weakening of Global Equivalence, as the following result shows.

Proposition 8. *Every method $*$ satisfying Global Equivalence also satisfies Preferential Cut.*

Proof. Suppose $\alpha \mid\!\sim^{\mathcal{KB}}_{PC} \beta$ and let $\mathcal{KB}' = \mathcal{KB} \cup \{\alpha \mid\!\sim \beta\}$. Since $\alpha \mid\!\sim^{\mathcal{KB}}_{PC} \beta$, every ranked interpretation satisfying $\mathcal{KB}$ also satisfies $\alpha \mid\!\sim \beta$, and so $[\![\mathcal{KB}]\!]_{\mathscr{R}} = [\![\mathcal{KB}']\!]_{\mathscr{R}}$. Hence, by Global Equivalence, $\mid\!\sim^{\mathcal{KB}'}_{*} = \mid\!\sim^{\mathcal{KB}}_{*}$. □

We thus obtain the following corollary to the proof of Proposition 7.

Corollary 2. *DRC does not satisfy Global Equivalence.*

In a previous version of this work [2], we conjectured that another weaker version of Cut holds, according to which the new conditional added must be such that its antecedent already appears as an antecedent of another conditional present in $\mathcal{KB}$.

Cut′ If $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{*} \beta$, $\mathcal{KB}' = \mathcal{KB} \cup \{\alpha \mathrel{|\!\sim} \beta\}$, and there is γ s.t. $\alpha \mathrel{|\!\sim} \gamma \in \mathcal{KB}$, then $\mathrel{|\!\sim}^{\mathcal{KB}'}_{*} \subseteq \mathrel{|\!\sim}^{\mathcal{KB}}_{*}$.

As it turns out, this conjecture is also false. To witness, consider the slightly different albeit equivalent version of the counterexample in Proposition 7 with $\mathcal{KB} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \top \mathrel{|\!\sim} \top\}$. We still get that the addition of the implied $\top \mathrel{|\!\sim} (\mathsf{b} \rightarrow \mathsf{f})$ results in the extra inference of $\neg\mathsf{f} \mathrel{|\!\sim} \neg\mathsf{b}$ and the violation of Cut′.

As we have seen in Corollary 1 in Section 4.5, the satisfaction of Cautious Monotonicity, plus the seemingly very reasonable behaviour displayed by disjunctive rational closure in Example 12, 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.

Corollary 3. *DRC does not satisfy Vacuity.*

Proof. Follows immediately from the impossibility result in Theorem 4, since Propositions 4, 6 and 7 tell us that DRC satisfies all the postulates listed in Theorem 4 other than Vacuity. □

However, we end this section by showing that DRC *does* satisfy our other major conservativeness postulate from Section 4.3 which, recall, was shown to fail for rational closure.

Proposition 9. *DRC satisfies Knowledge-based Monotonicity.*

Proof. Suppose $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{DC} \beta$, i.e., $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha) < \mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg\beta)$. We must show

$$\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})) < \mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB}) \wedge \neg\beta).$$

If $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})) = \infty$, then also $\mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})) = \infty$, and so $\mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB}) \wedge \neg\beta) = \infty$, giving the required conclusion. So, assume $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})) \neq \infty$. Clearly $\mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg\beta) \leq \mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB}) \wedge \neg\beta)$, so, using this with our assumption $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha) < \mathscr{L}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg\beta)$, it suffices to show $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})) = \mathscr{U}^{\mathcal{KB}}_{DC}(\alpha)$. We know that $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha) = \min\{\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})), \mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg \bigvee ante(\mathcal{KB}))\}$, so it suffices, in turn, to show that $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})) \leq \mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg \bigvee ante(\mathcal{KB}))$. Now, by construction of $\mathscr{I}^{\mathcal{KB}}_{DC}$, for all $v \in [\![\neg \bigvee ante(\mathcal{KB})]\!]$ we know that $\mathscr{L}^{\mathcal{KB}}_{DC}(v) = \mathscr{R}^{\mathcal{KB}}_{RC}(v) = 0$ (by definition of rational closure) and $\mathscr{U}^{\mathcal{KB}}_{DC}(v) = \max\{\mathscr{R}^{\mathcal{KB}}_{RC}(u) \mid \mathscr{R}^{\mathcal{KB}}_{RC}(u) \neq \infty\}$ (since v does not min-verify any of the conditionals in $\mathcal{KB}$). Hence, in particular from the latter, we get $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg \bigvee ante(\mathcal{KB})) = \max\{\mathscr{R}^{\mathcal{KB}}_{RC}(u) \mid \mathscr{R}^{\mathcal{KB}}_{RC}(u) \neq \infty\}$, which means (since we are assuming $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})) \neq \infty$) that we obtain $\mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \bigvee ante(\mathcal{KB})) \leq \mathscr{U}^{\mathcal{KB}}_{DC}(\alpha \wedge \neg \bigvee ante(\mathcal{KB}))$, as required. □

In terms of the interval order $\prec$ corresponding to its interval-based interpretation $\mathscr{I}^{\mathcal{KB}}_{DC}$, the essential reason that DRC satisfies KBM is that there is no plausible $\neg \bigvee ante(\mathcal{KB})$-valuation v and plausible $\bigvee ante(\mathcal{KB})$-valuation u such that $v \prec u$. This is ensured in the construction by the fact that $\mathscr{U}^{\mathcal{KB}}_{DC}(v) = \max\{\mathscr{R}^{\mathcal{KB}}_{RC}(u) \mid \mathscr{R}^{\mathcal{KB}}_{RC}(u) \neq \infty\}$ for all $v \in [\![\neg \bigvee ante(\mathcal{KB})]\!]$.

7. Forcing the Vacuity postulate

Regarding the properties of disjunctive rational closure that we have seen, the news is somewhat mixed, with several basic postulates satisfied, as well as Cautious Monotonicity and Knowledge-based Monotonicity, but with neither Cut nor Vacuity holding in general.

Regarding Vacuity, our impossibility result tells us that its failure is unavoidable given the other, reasonable, behaviour that we have shown DRC to exhibit (see Section 6). 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 DRC favouring the latter (Proposition 7) at the expense of the former (Corollary 3). 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 DRC otherwise. In what follows, we assess the potential ripple effects on the other properties of $\mathrel{|\!\sim}^{\mathcal{KB}}_{DC}$ of making this manoeuvre.

Let us assume a slightly modified version of our construction as motivated above which outputs the preferential closure of a conditional knowledge base $\mathcal{KB}$ when it is already a disjunctive relation, and let us call it DCV, the DRC with Vacuity enforced.

$$\mathrel{|\!\sim}^{\mathcal{KB}}_{DCV} \stackrel{\text{def}}{=} \begin{cases} \mathrel{|\!\sim}^{\mathcal{KB}}_{PC}, & \text{if } \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \text{ is a disjunctive relation;} \\ \mathrel{|\!\sim}^{\mathcal{KB}}_{DC}, & \text{otherwise.} \end{cases}$$

That the DCV of a conditional knowledge base satisfies Vacuity follows immediately from the definition above. As it can easily be checked, both Inclusion and D-Rationality are obviously satisfied by $\mathrel{|\!\sim}^{\mathcal{KB}}_{DCV}$. The same can be said of Infra-Rationality and Preferential Extension, which are satisfied by both the preferential closure of a knowledge base and its DRC. It remains to check the other, less obvious, properties, starting with Pairwise Equivalence.

Proposition 10. *DCV satisfies Pairwise Equivalence.*

Proof. Suppose $\mathcal{KB}_1 \equiv_p \mathcal{KB}_2$. Then $[\![\mathcal{KB}_1]\!]_{\mathscr{R}} = [\![\mathcal{KB}_2]\!]_{\mathscr{R}}$. We have two cases to consider:

Case 1: $\mathrel{|\!\sim}^{\mathcal{KB}_1}_{PC}$ is disjunctive. Then so is $\mathrel{|\!\sim}^{\mathcal{KB}_2}_{PC}$, from the fact that the preferential closure satisfies Global Equivalence. Hence we have $\mathrel{|\!\sim}^{\mathcal{KB}_1}_{DCV} = \mathrel{|\!\sim}^{\mathcal{KB}_1}_{PC} = \mathrel{|\!\sim}^{\mathcal{KB}_2}_{PC} = \mathrel{|\!\sim}^{\mathcal{KB}_2}_{DCV}$.

Case 2: $\mathrel{|\!\sim}^{\mathcal{KB}_1}_{PC}$ is not disjunctive. Then neither is $\mathrel{|\!\sim}^{\mathcal{KB}_2}_{PC}$, from the fact that the preferential closure satisfies Global Equivalence. Hence we have $\mathrel{|\!\sim}^{\mathcal{KB}_1}_{DCV} = \mathrel{|\!\sim}^{\mathcal{KB}_1}_{DC} = \mathrel{|\!\sim}^{\mathcal{KB}_2}_{DC} = \mathrel{|\!\sim}^{\mathcal{KB}_2}_{DCV}$, as DRC has been shown to satisfy Pairwise Equivalence (Proposition 5). □

The following result establishes that Knowledge-based Monotonicity is also satisfied.

Proposition 11. *DCV satisfies Knowledge-based Monotonicity.*

Proof. By definition, $\mathrel{|\!\sim}^{\mathcal{KB}}_{DCV}$ is equal to either the preferential closure or the disjunctive rational closure of $\mathcal{KB}$. So, satisfaction of KBM follows from the facts that both of these satisfy it (Propositions 3 and 9). □

Turning to BAP Representation Independence, we first establish a couple of useful facts concerning the behaviour of the preferential closure under BAP symbol translations.

Lemma 1. (i). *The preferential closure satisfies BAP Representation Independence.* (ii). *Let $\mathcal{KB}$ be a conditional knowledge base and σ a BAP symbol translation. Then $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ is disjunctive iff $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC}$ is disjunctive.*

Proof. *(i).* Let $\mathcal{KB}$ be a knowledge base, σ a BAP symbol translation and $\alpha, \beta \in \mathcal{L}$. We must show $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \beta$ iff $\sigma(\alpha) \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\beta)$. As stated by Lehmann and Magidor [5, p.41] (observation credited there to Dix), for any $\delta, \gamma \in \mathcal{L}$ and any knowledge base $\mathcal{KB}'$ we have $\delta \mathrel{|\!\sim}^{\mathcal{KB}'}_{PC} \gamma$ iff $\delta \mathrel{|\!\sim}^{\mathcal{KB}''}_{RC} \bot$, where $\mathcal{KB}'' = \mathcal{KB}' \cup \{\delta \mathrel{|\!\sim} \neg\gamma\}$. Thus $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \beta$ iff $\alpha \mathrel{|\!\sim}^{\mathcal{KB}_1}_{RC} \bot$, where $\mathcal{KB}_1 = \mathcal{KB} \cup \{\alpha \mathrel{|\!\sim} \neg\beta\}$, while $\sigma(\alpha) \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\beta)$ iff $\sigma(\alpha) \mathrel{|\!\sim}^{\mathcal{KB}_2}_{RC} \bot$, where $\mathcal{KB}_2 = \sigma(\mathcal{KB}) \cup \{\sigma(\alpha) \mathrel{|\!\sim} \neg\sigma(\beta)\}$. We see that $\mathcal{KB}_2 = \sigma(\mathcal{KB}_1)$. Thus, using the fact that $\sigma(\bot) = \bot$, we can conclude the desired $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \beta$ iff $\sigma(\alpha) \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\beta)$ from the fact (established in the proof of Proposition 6 in the appendix) that rational closure satisfies BAP Representation Independence.
(ii). Since the preferential closure of any conditional knowledge base is always a preferential relation, what needs to be shown is that $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ satisfies DR iff $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC}$ does. For the 'if' direction assume $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC}$ satisfies DR and let $\alpha, \beta, \gamma \in \mathcal{L}$ be such that $\alpha \vee \beta \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \gamma$. By part *(i)* just proved above this gives $\sigma(\alpha \vee \beta) \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\gamma)$, i.e., $\sigma(\alpha) \vee \sigma(\beta) \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\gamma)$. Since $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC}$ satisfies DR we get either $\sigma(\alpha) \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\gamma)$ or $\sigma(\beta) \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\gamma)$, equivalently (again using part *(i)*) $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \gamma$ or $\beta \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \gamma$ as required.

For the 'only if' direction, suppose $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ satisfies DR and let $\alpha, \beta, \gamma \in \mathcal{L}$ be such that $\alpha \vee \beta \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \gamma$. Since σ is a BAP symbol translation we know there exists α' such that $\models (\alpha \leftrightarrow \sigma(\alpha'))$ (just take any α' such that $[\![\alpha']\!] = \{R^{-1}_{\sigma}(v) \mid v \in [\![\alpha]\!]\}$, where R_{σ} is the permutation on $\mathcal{U}$ associated with σ), and similarly for β and γ. Thus from $\alpha \vee \beta \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \gamma$, and since we know $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC}$ is a preferential relation (and thus satisfies LLE and RWE) we obtain $\sigma(\alpha') \vee \sigma(\beta') \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\gamma')$, i.e., $\sigma(\alpha' \vee \beta') \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\gamma')$ which, using part *(i)* proved above, is equivalent to $\alpha' \vee \beta' \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \gamma'$. The assumption that $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ satisfies DR then gives either $\alpha' \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \gamma'$ or $\beta' \mathrel{|\!\sim}^{\mathcal{KB}}_{PC} \gamma'$, i.e., (using *(i)* again) either $\sigma(\alpha') \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\gamma')$ or $\sigma(\beta') \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \sigma(\gamma')$. This, in turn, is equivalent to $\alpha \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \gamma$ or $\beta \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC} \gamma'$ as required. □

Given the preceding facts, we can now show that DCV essentially inherits BAP Representation Independence from the preferential closure and the disjunctive rational closure.

Proposition 12. *DCV satisfies BAP Representation Independence.*

Proof. Let σ be a BAP symbol translation and let $\alpha, \beta \in \mathcal{L}$. We must show $\alpha \mathrel{|\!\sim}^{\mathcal{KB}}_{DCV} \beta$ iff $\sigma(\alpha) \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{DCV} \sigma(\beta)$.

Case 1: $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ is disjunctive. Then, by Lemma 1*(ii)*, so is $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC}$ and so we have $\mathrel{|\!\sim}^{\mathcal{KB}}_{DCV} = \mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ and $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{DCV} = \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC}$ and we can conclude using the fact that PC satisfies BAP Representation Independence (Lemma 1*(i)*).

Case 2: $\mathrel{|\!\sim}^{\mathcal{KB}}_{PC}$ is not disjunctive. Then, by Lemma 1*(ii)*, neither is $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{PC}$ and so we have $\mathrel{|\!\sim}^{\mathcal{KB}}_{DCV} = \mathrel{|\!\sim}^{\mathcal{KB}}_{DC}$ and $\mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{DCV} = \mathrel{|\!\sim}^{\sigma(\mathcal{KB})}_{DC}$ and we can conclude using the fact that DRC satisfies BAP Representation Independence (Proposition 6). □

The following is a direct consequence of the results above together with the impossibility result (Theorem 4) and links back to our comment in the second paragraph of the present section on the fact a choice has to be made between Vacuity and Cautious Monotonicity:

Corollary 4. *DCV fails Cautious Monotonicity.*

We close the discussion on the DCV of a conditional knowledge base with the following result:

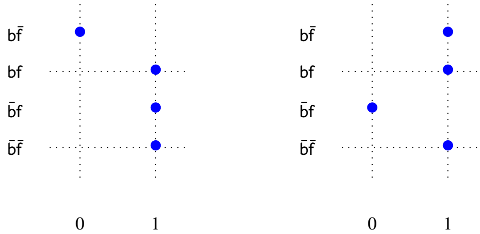

Fig. 11. Two ranked models of $\mathcal{KB} = \{\neg(\mathsf{b} \land \mathsf{f}) \mid\!\sim \mathsf{b} \lor \mathsf{f}\}$.

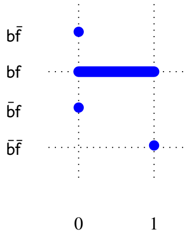

Fig. 12. The DRC of $\mathcal{KB} = \{\neg(\mathsf{b} \land \mathsf{f}) \mid\!\sim \mathsf{b} \lor \mathsf{f}\}$.

Proposition 13. *DCV fails Preferential Cut.*

Proof. Assume $\mathcal{P} = \{\mathsf{b}, \mathsf{f}\}$ and consider the conditional knowledge base $\mathcal{KB} = \{\neg(\mathsf{b} \land \mathsf{f}) \mid\!\sim \mathsf{b} \lor \mathsf{f}\}$. It is easy to check (either semantically via ranked models or syntactically via preferential reasoning) that $\top \mid\!\sim^{\mathcal{KB}}_{PC} \mathsf{b} \lor \mathsf{f}$. We will show that $\mid\!\sim^{\mathcal{KB}'}_{DCV} \not\subseteq \mid\!\sim^{\mathcal{KB}}_{DCV}$, where $\mathcal{KB}' = \mathcal{KB} \cup \{\top \mid\!\sim \mathsf{b} \lor \mathsf{f}\}$, thus violating Preferential Cut.

First note that, since $\top \mid\!\sim^{\mathcal{KB}}_{PC} \mathsf{b} \lor \mathsf{f}$ and $\models (\mathsf{b} \to \mathsf{f}) \lor (\mathsf{f} \to \mathsf{b})$, we have $(\mathsf{b} \to \mathsf{f}) \lor (\mathsf{f} \to \mathsf{b}) \mid\!\sim^{\mathcal{KB}}_{PC} \mathsf{b} \lor \mathsf{f}$. We have that $\mid\!\sim^{\mathcal{KB}}_{PC}$ is not disjunctive, since $\mathsf{b} \to \mathsf{f} \not\mid\!\sim^{\mathcal{KB}}_{PC} \mathsf{b} \lor \mathsf{f}$ and $\mathsf{f} \to \mathsf{b} \not\mid\!\sim^{\mathcal{KB}}_{PC} \mathsf{b} \lor \mathsf{f}$, as witnessed by the two ranked (and hence preferential) models of $\mathcal{KB}$ depicted in Fig. 11.

Hence $\mid\!\sim^{\mathcal{KB}}_{DCV} = \mid\!\sim^{\mathcal{KB}}_{DC}$, which is depicted in Fig. 12.

Note that we have $\mathsf{b} \leftrightarrow \mathsf{f} \not\mid\!\sim^{\mathcal{KB}}_{DC} \mathsf{b}$. Now, we know also $\mid\!\sim^{\mathcal{KB}'}_{PC}$ is not disjunctive (for the same reasons as $\mid\!\sim^{\mathcal{KB}}_{PC}$) and hence $\mid\!\sim^{\mathcal{KB}'}_{DCV} = \mid\!\sim^{\mathcal{KB}'}_{DC}$. In this case, the DRC of $\mathcal{KB}'$ corresponds to its rational closure. In particular, notice that $\mathsf{b} \leftrightarrow \mathsf{f} \mid\!\sim \mathsf{b}$ is in the RC of $\mathcal{KB}'$ and is therefore gained when constructing the DCV of $\mathcal{KB}'$, as required. □

Since, as we have seen, Preferential Cut is implied by either Global Equivalence (see Proposition 8) or Cut (in the presence of D-Rationality and Inclusion, both of which hold for DCV), the above result shows that also those properties fail for DCV.

To sum it up, DRC and DCV both fail Preferential Cut but satisfy the same set of postulates, except for the fact that the former fails Vacuity whereas the latter fails Cautious Monotonicity. We conclude that the two constructions are objectively incomparable. As already alluded to at the beginning of this section, the choice of which one to adopt is up to the system designer based on which of Cautious Monotonicity and Vacuity is preferable in specific cases.

8. Conclusion

In this paper, we have set ourselves the task of reviving 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 and its alter ego, the DCV, of a conditional knowledge base. Both allow for inferring a set of conditionals intermediate between the preferential closure and the rational closure. Finally, we have shown the properties satisfied by our constructions.

Table 1 shows a comparison between DRC, DCV and RC w.r.t. the postulates we have discussed throughout this work. As can be seen, DRC and DCV differ only in which of the two incompatible postulates Vacuity or Cautious Monotonicity are satisfied, with DRC satisfying the latter and DCV the former. The main postulate on which they both differ from RC is Knowledge-based Monotonicity, which is satisfied by both DRC and DCV, but which fails for RC.

As for the next steps, we plan to investigate suitable definitions of a preference relation on the set of interval-based interpretations. A promising starting point is the work of Disanto et al. [40]. We hope the construction we have introduced here can be shown to

Table 1
Postulates satisfied by each of DRC, DCV and RC.

Postulate	DRC	DCV	RC
Inclusion	✓	✓	✓
D-Rationality	✓	✓	✓
Preferential Extension	✓	✓	✓
Infra-Rationality	✓	✓	✓
Piecewise Equivalence	✓	✓	✓
Pairwise Equivalence	✓	✓	✓
Global Equivalence	✗	✗	✓
Vacuity	✗	✓	✗
Knowledge-based Monotonicity	✓	✓	✗
Representation Independence	✗	✗	✗
Negated Representation Independence	✓	✓	✓
BAP Representation Independence	✓	✓	✓
Cautious Monotonicity	✓	✗	✓
Cut	✗	✗	✓
Preferential Cut	✗	✗	✓

be the most preferred extension of the knowledge base according to some intuitively defined preference relation, along the lines of well-known results in the rational case.

The DRC construction we proposed in this work relies on the *offline* pre-computation of the rational closure of a (propositional) conditional knowledge base. This is usually done in the field, either semantically (as suitable for us here) or syntactically (via an algorithm which partitions the knowledge base according to the exceptionality of the conditionals therein). This minimises, to some extent, the impact of computing the DRC of a conditional knowledge base, which, as we have seen in Section 5, runs in time that grows (singly) exponentially with the size of the input.

We conjecture that another, more syntactic, construction for the DRC of a conditional knowledge base is possible. This should take the shape of an algorithm along the lines of Freund's or Casini & Straccia's but which would stratify the $\mid\!\sim$-statements in the KB rather into (possibly overlapping) sequences of ranks, somehow mimicking the intervals in the semantics, of which the complexity should be the same as that of the above-mentioned ranking algorithms. This conjecture is left for future work.

In this work, we required the postulate of Infra-Rationality. As a result, our construction of disjunctive rational closure took the rational closure of a conditional knowledge base 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. Examples of these are the lexicographic closure or System JLZ [7] or those based on c-revisions [31] and others [30]. 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 (DLs) [20,41,42,18,43], an investigation of our approach beyond the propositional case for low-complexity DLs may also be envisaged.

CRediT authorship contribution statement

Richard Booth: Investigation, Formal analysis, Methodology, Writing – review & editing, Conceptualization, Writing – original draft. **Ivan Varzinczak:** Methodology, Formal analysis, Writing – review & editing, Investigation, Validation, Writing – original draft, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work has been partially supported by the CNRS/Royal Society project *Non-Classical Reasoning for Enhanced Ontology-based Semantic Technologies*, and has also benefited from the support of the project AGGREEY ANR-22-CE23-0005 of the French National Research Agency (ANR). This work was partially supported by the ANR Chaire IA BE4musIA: BElief change FOR better MUlti-Source Information Analysis. The presentation of the equivalence postulates benefited from some discussions with Jesse Heyninck, Thomas Meyer, Giovanni Casini and Umberto Straccia. Thanks also to the anonymous reviewers of this journal for their insightful suggestions.

Appendix A. Proof of Theorem 2

Theorem 2. *A 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 $\mid\!\sim \, = \, \mid\!\sim_{\mathscr{I}}$.*

The 'if' direction (soundness of the properties) is relatively straightforward. For the only-if direction, we will utilise an equivalent alternative, purely qualitative, semantic structure, which we here call *interval models*.

Definition 8. An **interval model for $\mathcal{L}$** is a pair $\mathbb{D} = \langle V, \prec \rangle$, where $V \subseteq \mathcal{U}$ and $\prec$ is an interval order over V, i.e., a relation that is irreflexive, transitive and satisfies the interval condition: if $u \prec v$ and $s \prec t$ then $u \prec t$ or $s \prec v$.

Each interval model $\mathbb{D} = \langle V, \prec \rangle$ yields a consequence relation $\mid\!\sim_{\mathbb{D}}$ by setting, for $\alpha, \beta \in \mathcal{L}$, $\alpha \mid\!\sim_{\mathbb{D}} \beta$ iff either (*i*) $V \cap [\![\alpha]\!] = \emptyset$, or (*ii*) $V \cap [\![\alpha]\!] \neq \emptyset$ and $\min_{\prec}([\![\alpha]\!]) \subseteq [\![\beta]\!]$.

We give now a proof that disjunctive consequence relations can be characterised in terms of interval models. From well-known established results from the theory of interval orders (see [35,44]), we know that we can translate such a model $\mathbb{D}$ into an equivalent interval-based interpretation $\mathscr{I}$ that yields the same consequence relation. To do this, first of all if $u \notin V$ then we just set $\mathscr{L}(u) = \mathscr{U}(u) = \infty$. Otherwise, for each $u \in V$ we can define the *down set* $\mathrm{down}_{\prec}(u)$, respectively the *up set* $\mathrm{up}_{\prec}(u)$, of u with respect to $\prec$ to be the set of elements $v \in V$ such that $v \prec u$, respectively $u \prec v$. By virtue of being an interval model, the collection of all down sets for each $u \in V$ is linearly ordered by set inclusion, and similarly for the collection of all up sets. Let $P_1 \subset P_2 \subset \cdots \subset P_n$ be the chain of all down sets, and $S_n \subset S_{n-1} \subset \cdots \subset S_1$ be the chain of all up sets (both chains will be of the same length n). Then we can set $\mathscr{L}(u) = i$, where i is such that $P_i = \mathrm{down}_{\prec}(u)$ and $\mathscr{U}(u) = j$ where j is such that $S_j = \mathrm{up}_{\prec}(u)$. Given the existence of this translation, the following result is enough to give us Theorem 2.

Lemma 2. *A consequence relation $\mid\!\sim$ is disjunctive if and only if there is $\mathbb{D}$ such that $\mid\!\sim \, = \, \mid\!\sim_{\mathbb{D}}$.*

Proof. The 'if' direction (soundness of the properties) is relatively straightforward. We prove here the 'only-if' direction.

So, let $\mid\!\sim$ be a disjunctive relation. Let $\mathbb{D} = \langle V, \prec \rangle$ be the interval model specified by setting $V = \{u \mid u \not\mid\!\sim \bot\}$ and then setting, for all $u, v \in V$, $u \prec v$ iff $u \vee v \mid\!\sim \neg v$.[4] We must now show two things: (*a*) $\prec$ is an interval order over V, and (*b*) $\mid\!\sim \, = \, \mid\!\sim_{\mathbb{D}}$.

(*a*) $\prec$ is an interval order over V. For irreflexivity, we need to show $u \vee u \not\mid\!\sim \neg u$ for all $u \in V$. But if $u \vee u \mid\!\sim \neg u$, then $u \mid\!\sim \bot$ by LLE, Ref, RW, And, contradicting $u \in V$. Hence $u \vee u \not\mid\!\sim \neg u$ as required.

For transitivity, let $u, v, w \in V$ such that $u \vee v \mid\!\sim \neg v$ and $v \vee w \mid\!\sim \neg w$. We must show $u \vee w \mid\!\sim \neg w$. We know $u \neq w$, since otherwise we would have both $v \vee w \mid\!\sim \neg v$ (from $u \vee v \mid\!\sim \neg v$ and LLE) and $v \vee w \mid\!\sim \neg w$ and so $v \vee w \mid\!\sim \neg v \wedge \neg w$ (And) and so $v \vee w \mid\!\sim \bot$ (Ref, And, RW) giving either $v \mid\!\sim \bot$ or $w \mid\!\sim \bot$ by DR, which contradicts $v, w \in V$. So, given $u \neq w$, we now know $\models u \rightarrow \neg w$, and so to show $u \vee w \mid\!\sim \neg w$ it suffices to show $u \vee w \mid\!\sim u$ by RW. Now, we know $u \neq v$ by the already proved irreflexivity, so $\models v \rightarrow \neg u$. Hence $v \mid\!\sim \neg u$ by the following postulate, which is derivable from Ref, RW:

$$(\mathrm{SCL}) \quad \frac{\models \alpha \rightarrow \beta}{\alpha \mid\!\sim \beta}.$$

This gives $v \not\mid\!\sim u$ (by And, RW and the assumption $v \in V$). Thus, to show $u \vee w \mid\!\sim u$ it suffices, by DR, to show $u \vee w \vee v \mid\!\sim u$. To show this, in turn, it suffices to show $u \vee w \vee v \mid\!\sim \neg v \wedge \neg w$ (Ref, And, RW). Since $w \neq v$ (by irreflexivity), we know $w \mid\!\sim \neg v$ (SCL). Using this together with $u \vee v \mid\!\sim \neg v$ gives $u \vee v \vee w \mid\!\sim \neg v$ (by Or). Similarly, from $u \mid\!\sim \neg w$ (SCL) and $v \vee w \mid\!\sim \neg w$ we get $u \vee v \vee w \mid\!\sim \neg w$. Hence $u \vee w \vee v \mid\!\sim \neg v \wedge \neg w$ by And, as required.

For the interval condition, let $u, v, s, t \in V$ such that $u \prec v$ and $s \prec t$. We must show $u \prec t$ or $s \prec v$. If $v = s$, then we have $u \prec s$ and $s \prec t$ and we conclude $u \prec t$ as required by the just-proved transitivity, and similarly if $u = t$. Also, if $v = t$ then the required conclusion is immediate. So we may assume $v \neq s$, $v \neq t$ and $u \neq t$. Under this assumption, we know $s \vee t \mid\!\sim \neg v$ by SCL and so from this and $u \vee v \mid\!\sim \neg v$ we obtain $(u \vee v) \vee (s \vee t) \mid\!\sim \neg v$ by Or. We know $u \vee v \mid\!\sim \neg t$ from the assumption $u \neq t \neq v$ and SCL. Hence from this and $s \vee t \mid\!\sim \neg t$ we obtain $(u \vee v) \vee (s \vee t) \mid\!\sim \neg t$ (Or). From this together with $(u \vee v) \vee (s \vee t) \mid\!\sim \neg v$ we have, by And,

$$(u \vee v) \vee (s \vee t) \mid\!\sim \neg t \wedge \neg v \tag{A.1}$$

Now, assume for contradiction both $u \not\prec t$ and $s \not\prec v$, i.e., $u \vee t \not\mid\!\sim \neg t$ and $s \vee v \not\mid\!\sim \neg v$. From the former we know $u \vee t \not\mid\!\sim \neg t \wedge \neg v$ (RW) and, similarly, from the latter $s \vee v \not\mid\!\sim \neg t \wedge \neg v$. Then, from DR, $(u \vee t) \vee (s \vee v) \not\mid\!\sim \neg t \wedge \neg v$ and so, by LLE, $(u \vee v) \vee (s \vee t) \not\mid\!\sim \neg t \wedge \neg v$. But this contradicts (A.1). Hence either $u \prec t$ or $s \prec v$, as required.

(*b*) $\mid\!\sim \, = \, \mid\!\sim_{\mathbb{D}}$. Turning first to the $\subseteq$-inclusion, suppose $\alpha \mid\!\sim \beta$. If $V \cap [\![\alpha]\!] = \emptyset$, then $\alpha \mid\!\sim_{\mathbb{D}} \beta$, as required. So assume $V \cap [\![\alpha]\!] \neq \emptyset$. We must show $\min_{\prec}([\![\alpha]\!]) \subseteq [\![\beta]\!]$. Assume for contradiction $v \in \min_{\prec}([\![\alpha]\!]) \cap [\![\neg\beta]\!]$. Note from $v \in [\![\neg\beta]\!]$ and $v \in V$ we know $v \not\mid\!\sim \beta$ (since o.w. $v \mid\!\sim \bot$ by SCL, And, RW). We now claim that, for each $u \in [\![\alpha]\!]$, $u \vee v \not\mid\!\sim \beta$. This will suffice to reach contradiction with $\alpha \mid\!\sim \beta$,

[4] Whenever a valuation u appears within the scope of a propositional connective, it should be understood as denoting any sentence γ such that $[\![\gamma]\!] = \{u\}$. By LLE, RW, the precise choice of γ is irrelevant.

since it would give us $\bigvee_{u\in[\![\alpha]\!]}(u \vee v) \not\mid\!\sim \beta$ by DR, and so $\alpha \not\mid\!\sim \beta$ by LLE. To show the claim, let $u \in [\![\alpha]\!]$. If $u \notin V$, i.e., $u \mid\!\sim \bot$, then from this with $v \not\mid\!\sim \beta$ we obtain $u \vee v \not\mid\!\sim \beta$ as required by applying the following postulate (which follows from CM, RW, LLE, Ref, Or)[5]:

$$\frac{\alpha \mid\!\sim \bot,\ \ \alpha \vee \beta \mid\!\sim \gamma}{\beta \mid\!\sim \gamma}$$

So now assume instead $u \in V$. Then by minimality of v, $u \not\prec v$, i.e., $u \vee v \not\mid\!\sim \neg v$. But then $u \vee v \not\mid\!\sim \beta$ again, as required, by RW (since $v \in [\![\neg\beta]\!]$).

Turning to the $\supseteq$-direction, if $u \mid\!\sim \bot$ for all $u \in [\![\alpha]\!]$, then $\bigvee_{u\in[\![\alpha]\!]} u \mid\!\sim \bot$ by Or and so $\alpha \mid\!\sim \beta$ by LLE, RW, as required. So assume $V \cap [\![\alpha]\!] \neq \emptyset$. Let $\min_{\prec}([\![\alpha]\!]) = \{u_1, \ldots, u_m\}$. We know $\min_{\prec}([\![\alpha]\!]) \subseteq [\![\beta]\!]$ and so $\bigvee_{i=1}^{m} u_i \mid\!\sim \beta$ by SCL. For each $u \in (V \cap [\![\alpha]\!]) \setminus \min_{\prec}([\![\alpha]\!])$, we have $u_j \vee u \mid\!\sim \neg u$, for some j, so $u_j \vee u \mid\!\sim u_j$ (Ref, And, RW) and so $u_j \vee u \mid\!\sim \beta$ (RW). Combining $u_j \vee u \mid\!\sim \beta$ for each u with $\bigvee_{i=1}^{m} u_i \mid\!\sim \beta$ using Or, LLE then gives $\bigvee_{u\in V\cap[\![\alpha]\!]} u \mid\!\sim \beta$. We also have $\bigvee_{u\in V^c\cap[\![\alpha]\!]} u \mid\!\sim \bot$ using Or (where V^c denotes the complement of V), so $\bigvee_{u\in V^c\cap[\![\alpha]\!]} u \mid\!\sim \beta$ (RW). Using this together with $\bigvee_{u\in V\cap[\![\alpha]\!]} u \mid\!\sim \beta$ and Or, LLE gives us the required $\alpha \mid\!\sim \beta$. □

Appendix B. Proof of Proposition 6

Our proof that DRC satisfies BAP Representation Independence will make use of an alternative construction method for rational closure (see e.g. the work of Lehmann and Magidor [5] or the one by Freund [45]). Given any conditional knowledge base $\mathcal{KB}$, we denote by $\widetilde{\mathcal{KB}}$ its set of *material counterparts*, where the material counterpart of a conditional $\alpha \mid\!\sim \beta$ is just the sentence $\alpha \to \beta$. A sentence γ is *exceptional* for $\mathcal{KB}$ if $\widetilde{\mathcal{KB}} \vDash \neg\gamma$, and a conditional $\gamma \mid\!\sim \delta$ is exceptional for $\mathcal{KB}$ if its antecedent γ is. To build the ranked model $\mathscr{R}^{\mathcal{KB}}_{RC}$ corresponding to the rational closure of $\mathcal{KB}$, we define a decreasing set of conditionals $\mathcal{KB} = A_0 \supseteq A_1 \supseteq \cdots \supseteq A_n$, where, for each $i > 0$, A_i is the set of conditionals in A_{i-1} that are exceptional for A_{i-1}, and n is minimal such that $A_n = A_{n+1}$. Then $\mathscr{R}^{\mathcal{KB}}_{RC}$ is given by:

$$\mathscr{R}^{\mathcal{KB}}_{RC}(v) = \begin{cases} \min(\{i \mid v \in [\![\widetilde{A_i}]\!]\}), & \text{if } v \in [\![\widetilde{A_n}]\!]; \\ \infty, & \text{otherwise.} \end{cases}$$

Given a BAP symbol translation σ with corresponding permutation R_σ on $\mathcal{U}$, for any set V of valuations we will denote by $R_\sigma(V)$ the set $\{R_\sigma(v) \mid v \in V\}$. In what follows we make free use of the following facts, the second of which follows from the substitutivity theorem of propositional logic.[6]

Lemma 3. *Let σ be a BAP symbol translation, and let R_σ be its corresponding permutation on $\mathcal{U}$. Then, for any $X \cup \{\alpha\} \subseteq \mathcal{L}$, we have* (i) $[\![\sigma(X)]\!] = R_\sigma([\![X]\!])$, *and* (ii) $X \vDash \alpha$ *iff* $\sigma(X) \vDash \sigma(\alpha)$.

Proposition 6. *DRC satisfies BAP Representation Independence.*

Proof. Let σ be a BAP symbol translation, $\mathcal{KB}$ be a knowledge base, and $\alpha, \beta \in \mathcal{L}$. We must show $\alpha \mid\!\sim^{\mathcal{KB}}_{DC} \beta$ iff $\sigma(\alpha) \mid\!\sim^{\sigma(\mathcal{KB})}_{DC} \sigma(\beta)$. It suffices to show that, for all $\gamma \in \mathcal{L}$, (*i*) $\mathscr{L}^{\mathcal{KB}}_{DC}(\gamma) = \mathscr{L}^{\sigma(\mathcal{KB})}_{DC}(\sigma(\gamma))$, and (*ii*) $\mathscr{U}^{\mathcal{KB}}_{DC}(\gamma) = \mathscr{U}^{\sigma(\mathcal{KB})}_{DC}(\sigma(\gamma))$.

(*i*) $\mathscr{L}^{\mathcal{KB}}_{DC}(\gamma) = \mathscr{L}^{\sigma(\mathcal{KB})}_{DC}(\sigma(\gamma))$. By construction of DRC, this is the same as showing $\mathscr{R}^{\mathcal{KB}}_{RC}(\gamma) = \mathscr{R}^{\sigma(\mathcal{KB})}_{RC}(\sigma(\gamma))$. It suffices to show, for all valuations v, that $\mathscr{R}^{\mathcal{KB}}_{RC}(v) = \mathscr{R}^{\sigma(\mathcal{KB})}_{RC}(R_\sigma(v))$. (Indeed this part is essentially a proof that rational closure satisfies BAP Representation Independence.) Let the sequences of conditionals constructed from $\mathcal{KB}$, $\sigma(\mathcal{KB})$ in the above construction method for rational closure be, respectively,

$$\mathcal{KB} = A_0 \supseteq A_1 \supseteq \cdots \supseteq A_n$$

$$\sigma(\mathcal{KB}) = B_0 \supseteq B_1 \supseteq \cdots \supseteq B_m$$

We prove, by induction on i, that $B_i = \sigma(A_i)$ for $i = 0, 1, \ldots$ (which also implies $n = m$). This statement clearly holds for $i = 0$, so now fix $i > 0$ and assume $B_{i-1} = \sigma(A_{i-1})$. We must show that, for any conditional $\lambda \mid\!\sim \delta$, we have

$$(\lambda \mid\!\sim \delta) \in B_i \quad \text{iff} \quad \begin{array}{l}(\lambda \mid\!\sim \delta) = \sigma(\alpha) \mid\!\sim \sigma(\beta) \\ \text{for some } (\alpha \mid\!\sim \beta) \in A_i.\end{array}$$

Taking the only-if direction first, since $(\lambda \mid\!\sim \delta) \in B_i \subseteq B_0 = \sigma(\mathcal{KB})$ we know $\lambda \mid\!\sim \delta = \sigma(\alpha) \mid\!\sim \sigma(\beta)$ for *some* $(\alpha \mid\!\sim \beta) \in \mathcal{KB}$. It remains to show $(\alpha \mid\!\sim \beta) \in A_i$. Since $(\lambda \mid\!\sim \delta) \in B_i$ we know $\lambda = \sigma(\alpha)$ is exceptional for B_{i-1}, i.e., via the inductive hypothesis, $\widetilde{\sigma(A_{i-1})} \vDash \neg\sigma(\alpha)$. This is the same as $\sigma(\widetilde{A_{i-1}}) \vDash \sigma(\neg\alpha)$ which, in turn, is equivalent to $\widetilde{A_{i-1}} \vDash \neg\alpha$ by Lemma 3–(*ii*). Thus α is exceptional for A_{i-1} and so $(\alpha \mid\!\sim \beta) \in A_i$ as required. The converse direction is proved similarly.

We know that, for any valuation v, $v \in [\![\widetilde{A_i}]\!]$ iff $R_\sigma(v) \in [\![\sigma(\widetilde{A_i})]\!]$. Thus from the above proved $B_i = \sigma(A_i)$ for $i = 0, 1, \ldots$ we have $v \in [\![\widetilde{A_i}]\!]$ iff $R_\sigma(v) \in [\![\widetilde{B_i}]\!]$ and so $\mathscr{R}^{\mathcal{KB}}_{RC}(v) = \mathscr{R}^{\sigma(\mathcal{KB})}_{RC}(R_\sigma(v))$ as required.

[5] Note this is the only place in the entire proof which uses CM.

[6] Note that, for any set X of sentences and any symbol translation σ, we use $\sigma(X)$ to denote the set $\{\sigma(\alpha) \mid \alpha \in X\}$.

(*ii*) $\mathcal{U}^{\mathcal{KB}}_{DC}(\gamma) = \mathcal{U}^{\sigma(\mathcal{KB})}_{DC}(\sigma(\gamma))$. This time it suffices to show, for all valuations v, that $\mathcal{U}^{\mathcal{KB}}_{DC}(v) = \mathcal{U}^{\sigma(\mathcal{KB})}_{DC}(R_\sigma(v))$. If $\mathcal{R}^{\mathcal{KB}}_{RC}(v) = \mathcal{R}^{\sigma(\mathcal{KB})}_{RC}(R_\sigma(v)) = \infty$, then $\mathcal{U}^{\mathcal{KB}}_{DC}(v) = \mathcal{U}^{\sigma(\mathcal{KB})}_{DC}(R_\sigma(v)) = \infty$ by construction and so the result holds. So assume $\mathcal{R}^{\mathcal{KB}}_{RC}(v) \neq \infty \neq \mathcal{R}^{\sigma(\mathcal{KB})}_{RC}(R_\sigma(v))$.

We know from Lemma 3–(*i*) that $v \Vdash \alpha$ iff $R_\sigma(v) \Vdash \sigma(\alpha)$, while from part (*i*) just proved above, we know $\mathcal{R}^{\mathcal{KB}}_{RC}(v) = \mathcal{R}^{\sigma(\mathcal{KB})}_{RC}(R_\sigma(v))$ and $\mathcal{R}^{\mathcal{KB}}_{RC}(\alpha) = \mathcal{R}^{\sigma(\mathcal{KB})}_{RC}(\sigma(\alpha))$. Thus (1) v min-verifies $\alpha \mid\!\sim \beta$ in $\mathcal{R}^{\mathcal{KB}}_{RC}$ iff $R_\sigma(v)$ min-verifies $\sigma(\alpha) \mid\!\sim \sigma(\beta)$ in $\mathcal{R}^{\sigma(\mathcal{KB})}_{RC}$.

From part (*i*) proved above we also have $\mathcal{R}^{\mathcal{KB}}_{RC}(\alpha \wedge \neg\beta) = \mathcal{R}^{\sigma(\mathcal{KB})}_{RC}(\sigma(\alpha) \wedge \neg\sigma(\beta))$, while the maximal finite rank in both $\mathcal{R}^{\mathcal{KB}}_{RC}$ and $\mathcal{R}^{\sigma(\mathcal{KB})}_{RC}$ is equal to n. These two facts give us (2) $t^{\mathcal{KB}}_{RC}(\alpha, \beta) = t^{\sigma(\mathcal{KB})}_{RC}(\sigma(\alpha), \sigma(\beta))$. According to the construction, (1) and (2) are enough to give us $\mathcal{U}^{\mathcal{KB}}_{DC}(v) = \mathcal{U}^{\sigma(\mathcal{KB})}_{DC}(R_\sigma(v))$ as required. □

Data availability

No data was used for the research described in the article.

References

[1] R. Booth, I. Varzinczak, Towards conditional inference under disjunctive rationality, in: 18th International Workshop on Nonmonotonic Reasoning (NMR), 2020.
[2] R. Booth, I. Varzinczak, Conditional inference under disjunctive rationality, in: K. Leyton-Brown Mausam (Ed.), Proceedings of the 35th AAAI Conference on Artificial Intelligence, AAAI, 2021, pp. 6227–6234.
[3] M. Goldszmidt, P. Morris, J. Pearl, A maximum entropy approach to nonmonotonic reasoning, IEEE Trans. Pattern Anal. Mach. Intell. 15 (3) (1993) 220–232.
[4] D. Lehmann, Another perspective on default reasoning, Ann. Math. Artif. Intell. 15 (1) (1995) 61–82.
[5] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artif. Intell. 55 (1992) 1–60.
[6] J. Pearl, System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: Proc. of the 3rd Conference on Theoretical Aspects of Rationality and Knowledge (TARK), 1990, pp. 121–135.
[7] E. Weydert, System JLZ - rational default reasoning by minimal ranking constructions, J. Appl. Log. 1 (3–4) (2003) 273–308.
[8] R. Stalnaker, What is a nonmonotonic consequence relation?, Fundam. Inform. 21 (1) (1994) 7–21.
[9] D. Makinson, General patterns in nonmonotonic reasoning, in: D. Gabbay, C. Hogger, J. Robinson (Eds.), Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3, Oxford University Press, 1994, pp. 35–110.
[10] H. Rott, Four floors for the theory of theory change: the case of imperfect discrimination, in: Proc. of the 14th European Conference on Logics in Artificial Intelligence (JELIA), in: LNCS, vol. 8761, Springer, 2014, pp. 368–382.
[11] B. Hansson, An analysis of some deontic logics, Noûs 3 (1969) 373–398.
[12] Y. Shoham, Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence, MIT Press, 1988.
[13] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artif. Intell. 44 (1990) 167–207.
[14] J. McCarthy, Circumscription, a form of nonmonotonic reasoning, Artif. Intell. 13 (1–2) (1980) 27–39.
[15] R. Reiter, A logic for default reasoning, Artif. Intell. 13 (1–2) (1980) 81–132.
[16] P. Gärdenfors, D. Makinson, Nonmonotonic inference based on expectations, Artif. Intell. 65 (2) (1994) 197–245.
[17] P. Bonatti, Rational closure for all description logics, Artif. Intell. 274 (2019) 197–223.
[18] K. Britz, G. Casini, T. Meyer, K. Moodley, U. Sattler, I. Varzinczak, Principles of KLM-style defeasible description logics, ACM Trans. Comput. Log. 22 (1) (2021) 1.
[19] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, Preferential description logics, in: Logic for Programming, Artificial Intelligence, and Reasoning (LPAR), in: LNAI, vol. 4790, Springer, 2007, pp. 257–272.
[20] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, Semantic characterization of rational closure: from propositional logic to description logics, Artif. Intell. 226 (2015) 1–33.
[21] M. Pensel, A.-Y. Turhan, Including quantification in defeasible reasoning for the description logic $\mathcal{EL}_\perp$, in: M. Balduccini, T. Janhunen (Eds.), Proceedings of the 14th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR), in: LNCS, vol. 10377, Springer, 2017, pp. 78–84.
[22] K. Britz, T. Meyer, I. Varzinczak, Preferential reasoning for modal logics, in: Proceedings of the 7th Workshop on Methods for Modalities (M4M'2011), Electron. Notes Theor. Comput. Sci. 278 (2011) 55–69.
[23] K. Britz, I. Varzinczak, From KLM-style conditionals to defeasible modalities, and back, J. Appl. Non-Class. Log. 28 (1) (2018) 92–121.
[24] A. Chafik, F. Cheikh, J.-F. Condotta, I. Varzinczak, On the decidability of a fragment of preferential LTL, in: Proceedings of the 27th International Symposium on Temporal Representation and Reasoning (TIME), 2020.
[25] G. Casini, T. Meyer, I. Varzinczak, Situated conditional reasoning, Artif. Intell. 319 (2023) 103917.
[26] M. Jaeger, Representation independence of nonmonotonic inference relations, in: Proc. of the 5th International Conference on Principles of Knowledge Representation and Reasoning (KR), Morgan Kaufmann Publishers, 1996, pp. 461–472.
[27] R. Booth, J. Paris, A note on the rational closure of knowledge bases with both positive and negative knowledge, J. Log. Lang. Inf. 7 (2) (1998) 165–190.
[28] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, A minimal model semantics for nonmonotonic reasoning, in: Proc. of the 13th European Conference on Logics in Artificial Intelligence (JELIA), in: LNCS, vol. 7519, Springer, 2012, pp. 228–241.
[29] R. Booth, G. Casini, T. Meyer, I. Varzinczak, On rational entailment for propositional typicality logic, Artif. Intell. 277.
[30] G. Casini, T. Meyer, I. Varzinczak, Taking defeasible entailment beyond rational closure, in: Proc. of the 16th European Conference on Logics in Artificial Intelligence (JELIA), in: LNCS, vol. 11468, Springer, 2019, pp. 182–197.
[31] G. Kern-Isberner, Conditionals in Nonmonotonic Reasoning and Belief Revision, Lecture Notes in Artificial Intelligence, Springer, 2001.
[32] G. Casini, T. Meyer, K. Moodley, R. Nortjé, Relevant closure: a new form of defeasible reasoning for description logics, in: Proc. of the 14th European Conference on Logics in Artificial Intelligence (JELIA), in: LNCS, vol. 8761, Springer, 2014, pp. 92–106.
[33] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, Preferential vs rational description logics: which one for reasoning about typicality?, in: Proc. of the European Conference on Artificial Intelligence (ECAI), 2010, pp. 1069–1070.
[34] M. Freund, Injective models and disjunctive relations, J. Log. Comput. 3 (3) (1993) 231–247.
[35] P. Fishburn, Interval Orders and Interval Graphs: A Study of Partially Ordered Sets, Wiley, 1985.
[36] P. Calabrese, Deduction and inference using conditional logic and probability, in: Conditional Logic in Expert Systems, Elsevier, North Holland, 1991.
[37] P. Marquis, N. Schwind, Lost in translation: language independence in propositional logic – application to belief change, Artif. Intell. 206 (2014) 1–24.
[38] J. Paris, A. Vencovská, A note on the inevitability of maximum entropy, Int. J. Approx. Reason. 4 (3) (1990) 183–224.
[39] J. Paris, The Uncertain Reasoner's Companion - a Mathematical Perspective, Cambridge University Press, 1994.

[40] F. Disanto, L. Ferrari, S. Rinaldi, A partial order structure on interval orders, Util. Math. 102 (2017) 135–147.
[41] P. Bonatti, M. Faella, I. Petrova, L. Sauro, A new semantics for overriding in description logics, Artif. Intell. 222 (2015) 1–48.
[42] P. Bonatti, L. Sauro, On the logical properties of the nonmonotonic description logic DL^N, Artif. Intell. 248 (2017) 85–111.
[43] M. Pensel, A.-Y. Turhan, Reasoning in the defeasible description logic $\mathcal{EL}_\perp$ – computing standard inferences under rational and relevant semantics, Int. J. Approx. Reason. 112 (2018) 28–70.
[44] T.L. Greenaugh, Representation and enumeration of interval orders and semiorders, Ph.D. thesis, Dartmouth College, 1976.
[45] M. Freund, On the revision of preferences and rational inference processes, Artif. Intell. 152 (1) (2004) 105–137.