

Conditional Inference under Disjunctive Rationality

Richard Booth¹ Ivan Varzinczak^{2,3,4}

¹ Cardiff University, United Kingdom

² CRIL, Univ. Artois & CNRS, France

³ CAIR, Computer Science Division, Stellenbosch University, South Africa

⁴ ISTI-CNR, Italy

boothr2@cardiff.ac.uk, varzinczak@cril.fr

Abstract

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

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 sentences, has been one of the classic questions of AI, with several well-known solutions proposed (Goldszmidt, Morris, and Pearl 1993; Lehmann 1995; Lehmann and Magidor 1992; Pearl 1990; Weydert 2003). Since the work of Lehmann and colleagues in the early ’90s, the so-called preferential approach to defeasible reasoning has established itself as one of the most elegant frameworks within which to answer this question. Central to the preferential approach is the notion of *rational closure* of a conditional knowledge base, under which the set of inferred conditionals forms a rational consequence relation, i.e., satisfies all the rules of preferential reasoning, *plus* Rational Monotonicity. One of the reasons for accepting rational closure is the fact it delivers a venturous notion of entailment that is conservative enough. Given that, rationality has for long been accepted as the core baseline for any appropriate form of non-monotonic entailment.

Very few have stood against this position, including Makinson (1994), who considered Rational Monotonic-

ity too strong and has briefly advocated the weaker rule of Disjunctive Rationality instead, which says that if one may draw a conclusion from a disjunction of premises, then one should be able to draw this conclusion from at least one of these premises taken alone. This rule is implied by Rational Monotonicity and may still be desirable in cases where the latter does not hold. Quite surprisingly, the debate did not catch on, and, for lack of rivals of the same stature, rational closure has since reigned alone as a role model in non-monotonic inference. That used to be the case until Rott (2014) reignited interest in Disjunctive Rationality by considering interval models in connection with belief contraction. Inspired by that, here we revisit disjunctive consequence relations and introduce a suitable notion of *disjunctive rational closure* of a conditional knowledge base.

We start by giving a summary of the formal background and of the rational closure construction. Then, in the following section, we make a case for weakening the rationality requirement and propose a semantics with an accompanying representation result for a weaker form of rationality enforcing the rule of Disjunctive Rationality. After this we investigate a notion of closure of (or entailment from) a conditional knowledge base under Disjunctive Rationality. Our analysis is in terms of a set of postulates, all reasonable at first glance, that one can expect a suitable notion of closure to satisfy. Following that we propose a specific construction for the Disjunctive Rational Closure of a conditional knowledge base and assess its suitability in the light of the postulates put forward in the previous section. We conclude with some remarks on future directions of investigation.

Formal Preliminaries

Let \mathcal{P} be a finite set of propositional *atoms*. We use p, q, \dots as meta-variables for atoms. Propositional sentences are denoted by α, β, \dots , and are recursively defined as usual: $\alpha ::= \top \mid \perp \mid \mathcal{P} \mid \neg\alpha \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid \alpha \rightarrow \alpha \mid \alpha \leftrightarrow \alpha$. We use \mathcal{L} to denote the set of all propositional sentences.

With $\mathcal{U} \stackrel{\text{def}}{=} \{0, 1\}^{\mathcal{P}}$, we denote the set of all propositional *valuations*, where 1 represents truth and 0 falsity. We use v, u, \dots , possibly with primes, to denote valuations. We shall sometimes represent valuations as sequences of atoms (e.g., p) and barred atoms (e.g., \bar{p}), where the presence of an 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 $\mathcal{P} = \{b, f, p\}$, 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 is represented as $b\bar{f}p$.

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

KLM-Style Rational Defeasible Consequence

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

A *defeasible consequence relation* \sim is a binary relation on \mathcal{L} . We say \sim is *preferential* (Kraus, Lehmann, and Magidor 1990) if it satisfies the rules:

(Ref) $\alpha \sim \alpha$	(LLE) $\frac{\models \alpha \leftrightarrow \beta, \alpha \sim \gamma}{\beta \sim \gamma}$
(And) $\frac{\alpha \sim \beta, \alpha \sim \gamma}{\alpha \sim \beta \wedge \gamma}$	(Or) $\frac{\alpha \sim \gamma, \beta \sim \gamma}{\alpha \vee \beta \sim \gamma}$
(RW) $\frac{\alpha \sim \beta, \models \beta \rightarrow \gamma}{\alpha \sim \gamma}$	(CM) $\frac{\alpha \sim \beta, \alpha \sim \gamma}{\alpha \wedge \beta \sim \gamma}$

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

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

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

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

Given \mathcal{R} , we call $\mathcal{R}(v)$ the *rank* of v w.r.t. \mathcal{R} . The intuition is that valuations with a lower rank are deemed more

normal (or typical) than those with a higher rank, while those with an infinite rank are regarded as so atypical as to be ‘forbidden’. Given a ranked interpretation \mathcal{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 ∞).

Figure 1 depicts an example of a ranked interpretation for $\mathcal{P} = \{b, f, p\}$. (Plausible valuations are associated with the colour blue, whereas the implausible ones with red.)




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

Given \mathcal{R} and $\alpha \in \mathcal{L}$, with $\llbracket \alpha \rrbracket^{\mathcal{R}}$ we denote the set of plausible valuations satisfying α in \mathcal{R} . With $\mathcal{R}(\alpha) \stackrel{\text{def}}{=} \min\{\mathcal{R}(v) \mid v \in \llbracket \alpha \rrbracket^{\mathcal{R}}\}$ we denote the *rank* of α in \mathcal{R} . By convention, if $\llbracket \alpha \rrbracket^{\mathcal{R}} = \emptyset$, we let $\mathcal{R}(\alpha) = \infty$. Defeasible consequence of the form $\alpha \sim \beta$ is then given a semantics in terms of ranked interpretations in the following way: We say $\alpha \sim \beta$ is *satisfied in* \mathcal{R} (denoted $\mathcal{R} \models \alpha \sim \beta$) if $\mathcal{R}(\alpha) < \mathcal{R}(\alpha \wedge \neg \beta)$. (Here we adopt Jaeger’s (1996) convention that $\infty < \infty$ always holds.) Equivalently, $\alpha \sim \beta$ iff β holds in all the most normal α -valuations. If $\mathcal{R} \models \alpha \sim \beta$, we say \mathcal{R} is a *ranked model* of $\alpha \sim \beta$. In the example in Figure 1, we have $\mathcal{R} \models b \sim f$, $\mathcal{R} \models \neg(p \rightarrow b) \sim \perp$, $\mathcal{R} \models p \sim \neg f$, $\mathcal{R} \not\models f \sim b$, and $\mathcal{R} \models p \wedge \neg b \sim b$.

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

Rational Closure

One can also view defeasible consequence as formalising some form of (defeasible) conditional and bring it down to the level of statements. Such was the stance adopted by Lehmann and Magidor (1992). A *conditional knowledge base* \mathcal{KB} is a finite set of statements of the form $\alpha \sim \beta$, with $\alpha, \beta \in \mathcal{L}$. In knowledge bases, we shall abbreviate $\neg \alpha \sim \perp$ with α . As an example, let $\mathcal{KB} = \{b \sim f, p \rightarrow b, p \sim \neg f\}$. Given a conditional knowledge base \mathcal{KB} , a *ranked model* of \mathcal{KB} is a ranked interpretation satisfying all statements in \mathcal{KB} . The ranked interpretation in Figure 1 is a ranked model of the above \mathcal{KB} . 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} , \sim is expected to adhere to the KLM rules.

To be more precise, we can take the defeasible conditionals in \mathcal{KB} as the core elements of a defeasible consequence

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

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

Assume a partial order $\preceq_{\mathcal{KB}}$ on all ranked models of a knowledge base \mathcal{KB} , defined as follows: $\mathcal{R}_1 \preceq_{\mathcal{KB}} \mathcal{R}_2$, if, for every $v \in \mathcal{U}$, $\mathcal{R}_1(v) \leq \mathcal{R}_2(v)$. Giordano et al. (2015) showed that there is a unique $\preceq_{\mathcal{KB}}$ -minimal element. The rational closure of \mathcal{KB} is defined in terms of this minimum ranked model of \mathcal{KB} .

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

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

Rational closure is commonly viewed as the *basic* (although not the only acceptable) form of non-monotonic entailment, on which other, more venturous forms of entailment can be and have been constructed (Booth et al. 2019; Casini et al. 2014; Casini, Meyer, and Varzinczak 2019; Kern-Isberner 2001; Lehmann 1995).

Disjunctive Rationality and Intervals

One may argue there are cases where Rational Monotonicity is too strong a rule to enforce and for which a weaker defeasible consequence relation would suffice (Giordano et al.

2010; Makinson 1994). Nevertheless, doing away with rationality, i.e., sticking to the preferential rules only, is not particularly appropriate in a defeasible-reasoning context. Indeed, as widely known in the literature, preferential systems induce entailment relations that are monotonic (Lehmann and Magidor 1992). In that respect, here we are interested in defeasible consequence relations (or conditionals) that do not necessarily satisfy RM while still encapsulating some form of rationality. A case in point is that of the Disjunctive Rationality (DR) rule (Kraus, Lehmann, and Magidor 1990):

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

Intuitively, DR says that if one may draw a conclusion from a disjunction of premises, then one should be able to draw this conclusion from at least one of these premises taken alone. Kraus, Lehmann, and Magidor (1990) 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 (Lehmann and Magidor 1992). Therefore, DR is a weaker form of rationality, as its name suggests. Given that, Disjunctive Rationality is indeed a suitable candidate for the type of investigation we have in mind.

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

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

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/upper rank, i.e., it is a strict partial order that additionally satisfies the *interval condition*: if $u \prec v$ and $u' \prec v'$, then either $u \prec v'$ or $u' \prec v$. Furthermore, every interval order over any subset of \mathcal{U} can be defined from an interval-based interpretation in this way. See the work of Fishburn (1985) for more details, and also that of Rott (2014), who more recently explored interval orders in the context of belief contraction.

Figure 2 illustrates an example of an interval-based interpretation for $\mathcal{P} = \{b, f, p\}$. Whenever the intervals associated to valuations u and v overlap, the intuition is that both valuations are incomparable in \mathcal{I} ; otherwise the leftmost interval is seen as more preferred than the rightmost one.




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

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 (iii) of the previous definition). With $\mathcal{L}(\alpha) \stackrel{\text{def}}{=} \min\{\mathcal{L}(v) \mid v \in \llbracket \alpha \rrbracket^{\mathcal{I}}\}$ and $\mathcal{U}(\alpha) \stackrel{\text{def}}{=} \min\{\mathcal{U}(v) \mid v \in \llbracket \alpha \rrbracket^{\mathcal{I}}\}$ we denote, respectively, the *lower* and the *upper rank* of α in \mathcal{I} . By convention, if $\llbracket \alpha \rrbracket^{\mathcal{I}} = \emptyset$, we let $\mathcal{L}(\alpha) = \mathcal{U}(\alpha) = \infty$. We say $\alpha \sim \beta$ is satisfied in \mathcal{I} (denoted $\mathcal{I} \Vdash \alpha \sim \beta$) if $\mathcal{U}(\alpha) < \mathcal{L}(\alpha \wedge \neg\beta)$. (Recall the convention that $\infty < \infty$.) As an example, in the interval-based interpretation of Figure 2, we have $\mathcal{I} \Vdash b \sim f$, $\mathcal{I} \Vdash p \sim \neg f$, and $\mathcal{I} \not\Vdash \neg f \sim \neg p$ (contrary to the ranked interpretation \mathcal{R} in Figure 1, which endorses the latter).

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

Theorem 1 *A defeasible consequence relation \sim is disjunctive if and only if there is \mathcal{I} such that $\sim = \sim_{\mathcal{I}}$.*

Towards Disjunctive Rational Closure

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

Proposition 1 *Given \mathcal{KB} , (i) the disjunctive closure of \mathcal{KB} coincides with its preferential closure $\sim_{PC}^{\mathcal{KB}}$. (ii) There exists \mathcal{KB} such that $\sim_{PC}^{\mathcal{KB}}$ does not satisfy DR.*

This result suggests the quest for a suitable definition of entailment under DR 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 $\sim_*^{\mathcal{KB}}$ the special defeasible consequence relation that we are looking for. Next, we consider some desirable properties for the mapping from \mathcal{KB} to $\sim_*^{\mathcal{KB}}$, 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} .

Basic Postulates

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

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

D-Rationality $\sim_*^{\mathcal{KB}}$ is a disjunctive consequence relation.

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

Definition 4 *For $\alpha, \beta, \gamma, \delta \in \mathcal{L}$, $\alpha \sim \beta$ is **equivalent** to $\gamma \sim \delta$ if $\models (\alpha \leftrightarrow \gamma) \wedge (\beta \leftrightarrow \delta)$. \mathcal{KB} and \mathcal{KB}' are **equivalent** ($\mathcal{KB} \equiv \mathcal{KB}'$), if there is a bijection $f : \mathcal{KB} \rightarrow \mathcal{KB}'$ s.t. each $\alpha \sim \beta \in \mathcal{KB}$ is equivalent to $f(\alpha \sim \beta)$.*

We can then express a weak form of syntax independence:

Equivalence If $\mathcal{KB}_1 \equiv \mathcal{KB}_2$, then $\sim_*^{\mathcal{KB}_1} = \sim_*^{\mathcal{KB}_2}$.

Weaker notions of equivalence between knowledge bases are possible (see, e.g., (Beierle, Eichhorn, and Kern-Isberner 2017)), leading to stronger forms of syntax independence.

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

Infra-Rationality $\sim_*^{\mathcal{KB}} \subseteq \sim_{RC}^{\mathcal{KB}}$.

Minimality Postulates

Echoing a fundamental principle of reasoning in general and of non-monotonic reasoning in particular is a property requiring $\sim_*^{\mathcal{KB}}$ 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 $\sim_*^{\mathcal{KB}}$ to be a set-theoretically *minimal* disjunctive consequence relation that extends \mathcal{KB} .

Example 1 *Suppose the only knowledge we have is a single conditional saying “birds normally fly”, i.e., $\mathcal{KB} = \{b \sim f\}$. Assuming just two variables, we have a unique \subseteq -minimal disjunctive consequence relation extending \mathcal{KB} , which is given by the interval-based interpretation \mathcal{I} in Figure 3. Indeed, the conditional $b \sim f$ is saying precisely that $bf \prec bf$, 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 $\neg f \sim_{\mathcal{I}} \neg b$ here. Note the rational closure in this example does endorse this latter conclusion, thus providing evidence that the rational closure arguably gives some unwarranted conclusions.*

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




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

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




Figure 4: Left and center: Interval-based models of the two \subseteq -minimal extensions of $\sim^{\mathcal{KB}}$, for $\mathcal{KB} = \{m \sim s, \neg m \sim s\}$. Right: Interval-based model of the union of the two \subseteq -minimal extensions of $\sim^{\mathcal{KB}}$.

In the light of Example 2, a question that arises is what to do when one has more than a single \subseteq -minimal extension of $\sim^{\mathcal{KB}}$. Proposition 1 already tells us we cannot, in general, take their intersection. However we might still expect the following postulates as reasonable.

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

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

(Note, given Proposition 1, the postulate above follows from Inclusion and D-Rationality.)

Representation Independence

What should the answer be in Example 2? Intuitively, faced with the choice of which of the pairs $ms \prec \bar{m}\bar{s}$ or $\bar{m}s \prec m\bar{s}$ to include, and in the absence of any reason to prefer either one, it seems the right thing to do is to include both, and thereby let the interval-based interpretation depicted in Figure 4 (right) 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* (Marquis and Schwind 2014). A symbol translation (on \mathcal{P}) is a function $\sigma : \mathcal{P} \rightarrow \mathcal{L}$. It 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. Similarly, given \mathcal{KB} and $\sigma(\cdot)$, we denote by $\sigma(\mathcal{KB})$ the knowledge base obtained by replacing each conditional $\alpha \sim \beta$ in \mathcal{KB} by $\sigma(\alpha) \sim \sigma(\beta)$. A specific family of symbol translations are the *negation-swapping* ones, i.e., when $\sigma(p) \in \{p, \neg p\}$, for all $p \in \mathcal{P}$. We propose the following postulate:

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

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

We note that Weydert (2003) and Jaeger (1996) also consider representation independence in the context of conditional inference, but in slightly different frameworks.

Cumulativity Postulates

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

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

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

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

Theorem 2 There is no method $*$ satisfying all of Inclusion, D-Rationality, Equivalence, Vacuity, Cautious Monotonicity and Negated Representation Independence.

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

$$(m \leftrightarrow s) \sim^{\mathcal{KB}}_* s \text{ or } (\neg m \leftrightarrow s) \sim^{\mathcal{KB}}_* s \quad (1)$$

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

$$(m \leftrightarrow s) \sim^{\mathcal{KB}}_* s \text{ iff } (\neg m \leftrightarrow s) \sim^{\mathcal{KB}}_* s \quad (2)$$

Putting (1) and (2) together gives us both $(m \leftrightarrow s) \sim^{\mathcal{KB}}_* s$ and $(\neg m \leftrightarrow s) \sim^{\mathcal{KB}}_* s$. Now, let $\mathcal{KB}' = \mathcal{KB} \cup \{(m \leftrightarrow s) \sim$

s}. By Cautious Monotonicity, $\vdash_*^{\mathcal{KB}} \subseteq \vdash_*^{\mathcal{KB}'}$. In particular, $(\neg m \leftrightarrow s) \vdash_*^{\mathcal{KB}'}$ s. It can be checked that the preferential closure of \mathcal{KB}' is itself a disjunctive consequence relation. In fact, it corresponds to the interval-based interpretation on the left of Figure 4. Hence, by Vacuity, this particular interval-based interpretation corresponds also to $\vdash_*^{\mathcal{KB}'}$. But, we can see from the figure $(\neg m \leftrightarrow s) \not\vdash_*^{\mathcal{KB}'}$ s—contradiction. ■

Theorem 2 is both surprising and disappointing, since all the properties mentioned seem to be intuitive and desirable. What can we do in the face of this result? Our strategy will be to construct a method that can satisfy as many of these properties as possible. We now provide our candidate for such a method - the disjunctive rational closure.

A Construction for Disjunctive Rational Closure

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

To specify $\mathcal{I}_{DC}^{\mathcal{KB}}$, we will construct the pair $\langle \mathcal{L}_{DC}^{\mathcal{KB}}, \mathcal{U}_{DC}^{\mathcal{KB}} \rangle$ 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 $\mathcal{R}_{RC}^{\mathcal{KB}}$ of \mathcal{KB} as a point of departure. Starting with the lower ranks, we simply set, for all $v \in \mathcal{U}$: $\mathcal{L}_{DC}^{\mathcal{KB}}(v) \stackrel{\text{def}}{=} \mathcal{R}_{RC}^{\mathcal{KB}}(v)$. That is, the lower ranks are given by the rational closure.

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

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

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

Case 1: v does not verify any of the conditionals in \mathcal{KB} in $\mathcal{R}_{RC}^{\mathcal{KB}}$. In this case, we set: $\mathcal{U}_{DC}^{\mathcal{KB}}(v) \stackrel{\text{def}}{=} \max\{\mathcal{R}_{RC}^{\mathcal{KB}}(u) \mid \mathcal{R}_{RC}^{\mathcal{KB}}(u) \neq \infty\}$.

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

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

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

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

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

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

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

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

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

Let us revisit the examples we have seen throughout the paper, to see what DRC gives.

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

Example 5 Returning to Example 2, with $\mathcal{KB} = \{m \sim s, \neg m \sim s\}$, the rational closure yields $\mathcal{R}_{RC}^{\mathcal{KB}}(ms) = \mathcal{R}_{RC}^{\mathcal{KB}}(ms) = 0$ and $\mathcal{R}_{RC}^{\mathcal{KB}}(m\bar{s}) = \mathcal{R}_{RC}^{\mathcal{KB}}(m\bar{s}) = 1$, which gives us the lower ranks. The valuation ms verifies only $m \sim s$, and so $\mathcal{U}_{DC}^{\mathcal{KB}}(ms) = t_{RC}^{\mathcal{KB}}(m, s) = \mathcal{R}_{RC}^{\mathcal{KB}}(m \wedge \neg s) - 1 = 1 - 1 = 0$. Similarly, the valuation $\bar{m}s$ verifies only $\neg m \sim s$ and so, by analogous reasoning, $\mathcal{U}_{DC}^{\mathcal{KB}}(\bar{m}s) = t_{RC}^{\mathcal{KB}}(\neg m, s) = 0$. So both valuations are assigned the interval $(0, 0)$ by $\mathcal{I}_{DC}^{\mathcal{KB}}$. The other two valuations, which verify neither conditional in \mathcal{KB} , are assigned $(1, 1)$. In this case, $\mathcal{I}_{DC}^{\mathcal{KB}}$ returns just the rational closure of \mathcal{KB} , as pictured in Figure 4 (right).

In both examples above, DRC returns the right answers.

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

Properties of the Disjunctive Rational Closure

We now turn to the question of which of the postulates from the preceding section are satisfied by DRC. To begin with, we obtain all of the basic postulates proposed there.

Proposition 2 *DRC satisfies Inclusion, D-Rationality, Equivalence and Infra-Rationality.*

We remind the reader that, since Inclusion and D-Rationality hold, DRC also satisfies Preferential Extension. We can also confirm that DRC conforms with our Representation Independence requirement.

Proposition 3 *DRC satisfies Negated Representation Independence.*

DRC essentially inherits this property from rational closure, which can also be shown to satisfy it. Although Jaeger (1996) 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 (1992) that rational closure satisfies both Cautious Monotonicity and Cut, and, in fact, if $\alpha \sim_{RC}^{\mathcal{KB}} \beta$ and $\mathcal{KB}' = \mathcal{KB} \cup \{\alpha \sim \beta\}$, then $\mathcal{R}_{RC}^{\mathcal{KB}} = \mathcal{R}_{RC}^{\mathcal{KB}'}$. We can show the following for DRC.

Proposition 4 *DRC satisfies Cautious Monotonicity, but does not satisfy Cut.*

The reason for the failure of Cut is that by adding a new conditional $\alpha \sim \beta$ to \mathcal{KB} , even when that conditional is already inferred by DRC, we give certain valuations (i.e., those in $\llbracket \alpha \rrbracket$) opportunity to verify one more conditional from the knowledge base in $\mathcal{R}_{RC}^{\mathcal{KB}'}$. This leads, potentially, to a corresponding decrease in their upper ranks $\mathcal{U}_{DC}^{\mathcal{KB}'}$, 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, adding an inferred conditional will never lead to an *increase* in the upper ranks, which means DRC *does* satisfy Cautious Monotonicity.

As we have seen in our impossibility result (Theorem 2), satisfaction of Cautious Monotonicity, plus that of the other properties proved earlier in this section, comes at the cost of

Vacuity, i.e., even if the preferential closure is a disjunctive relation, the output may sanction extra conclusions.

Proposition 5 *DRC does not satisfy Vacuity.*

Concluding Remarks

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

In our quest for a suitable form of entailment ensuring Disjunctive Rationality, we started by putting forward a set of postulates, all reasonable at first glance, characterising its expected behaviour. As it turns out, not all of them can be satisfied simultaneously, which suggests there might be more than one answer to our research question. We have then provided a construction of the disjunctive rational closure of a conditional knowledge base, which infers a set of conditionals intermediate between the preferential closure and the rational closure. Regarding the computational complexity of our construction, space considerations prevent us from providing the details. Nevertheless, we have checked that our construction method runs in time that grows (singly) exponentially with the size of the input, with rational closure of the knowledge base computed offline.

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

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

As for future work, we plan to investigate suitable definitions of a preference relation on the set of interval-based interpretations. We hope our construction can be shown to be the most preferred extension of the knowledge base according to some intuitively defined preference relation, as has been done in the rational case.

Acknowledgments

This work has been partially supported by the CNRS/Royal Society project *Non-Classical Reasoning for Enhanced Ontology-based Semantic Technologies*. Thanks to the anonymous reviewers for helpful suggestions. A preliminary version of this work was presented at the 18th Workshop on Non-Monotonic Reasoning (NMR 2020) (Booth and Varzinczak 2020).

References

- Beierle, C.; Eichhorn, C.; and Kern-Isberner, G. 2017. A transformation system for unique minimal normal forms of conditional knowledge bases. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, 236–245. Springer.
- Bonatti, P. 2019. Rational closure for all description logics. *Artificial Intelligence* 274: 197–223.
- Booth, R.; Casini, G.; Meyer, T.; and Varzinczak, I. 2019. On Rational Entailment for Propositional Typicality Logic. *Artificial Intelligence* 277.
- Booth, R.; and Varzinczak, I. 2020. Towards Conditional Inference under Disjunctive Rationality. In *18th International Workshop on Nonmonotonic Reasoning (NMR)*.
- Britz, K.; Meyer, T.; and Varzinczak, I. 2011. Semantic Foundation for Preferential Description Logics. In *Proc. of the 24th Australasian Joint Conference on Artificial Intelligence*, number 7106 in LNAI, 491–500. Springer.
- Casini, G.; Meyer, T.; Moodley, K.; and Nortjé, R. 2014. Relevant Closure: A New Form of Defeasible Reasoning for Description Logics. In *Proc. of the 14th European Conference on Logics in Artificial Intelligence (JELIA)*, number 8761 in LNCS, 92–106. Springer.
- Casini, G.; Meyer, T.; and Varzinczak, I. 2019. Taking Defeasible Entailment Beyond Rational Closure. In *Proc. of the 16th European Conference on Logics in Artificial Intelligence (JELIA)*, number 11468 in LNCS, 182–197. Springer.
- Fishburn, P. 1985. *Interval Orders and Interval Graphs: A Study of Partially Ordered Sets*. Wiley.
- Freund, M. 1993. Injective Models and Disjunctive Relations. *Journal of Logic and Computation* 3(3): 231–247.
- Gärdenfors, P.; and Makinson, D. 1994. Nonmonotonic inference based on expectations. *Artificial Intelligence* 65(2): 197–245.
- Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2007. Preferential Description Logics. In *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, number 4790 in LNAI, 257–272. Springer.
- Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2010. Preferential vs Rational Description Logics: which one for Reasoning About Typicality? In *Proc. of the European Conference on Artificial Intelligence (ECAI)*, 1069–1070.
- Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2012. A Minimal Model Semantics for Nonmonotonic Reasoning. In *Proc. of the 13th European Conference on Logics in Artificial Intelligence (JELIA)*, number 7519 in LNCS, 228–241. Springer.
- Giordano, L.; Gliozzi, V.; Olivetti, N.; and Pozzato, G. 2015. Semantic characterization of rational closure: From propositional logic to description logics. *Artificial Intelligence* 226: 1–33.
- Goldszmidt, M.; Morris, P.; and Pearl, J. 1993. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(3): 220–232.
- Jaeger, M. 1996. Representation Independence of Nonmonotonic Inference Relations. In *Proc. of the 5th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, 461–472. Morgan Kaufmann Publishers.
- Kern-Isberner, G. 2001. *Conditionals in Nonmonotonic Reasoning and Belief Revision*. Springer, Lecture Notes in Artificial Intelligence.
- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44: 167–207.
- Lehmann, D. 1995. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence* 15(1): 61–82.
- Lehmann, D.; and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55: 1–60.
- Makinson, D. 1994. General patterns in nonmonotonic reasoning. In Gabbay, D.; Hogger, C.; and Robinson, J., eds., *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 3, 35–110. Oxford University Press.
- Marquis, P.; and Schwind, N. 2014. Lost in translation: Language independence in propositional logic – application to belief change. *Artificial Intelligence* 206: 1–24.
- McCarthy, J. 1980. Circumscription, a form of nonmonotonic reasoning. *Artificial Intelligence* 13(1-2): 27–39.
- Pearl, J. 1990. 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)*.
- Reiter, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13(1-2): 81–132.
- Rott, H. 2014. Four Floors for the Theory of Theory Change: The Case of Imperfect Discrimination. In *Proc. of the 14th European Conference on Logics in Artificial Intelligence (JELIA)*, number 8761 in LNCS, 368–382. Springer.
- Shoham, Y. 1988. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press.
- Weydert, E. 2003. System JLZ - rational default reasoning by minimal ranking constructions. *Journal of Applied Logic* 1(3-4): 273–308.