Artificial Intelligence 319 (2023) 103917

Contents lists available at ScienceDirect

Artificial Intelligence

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

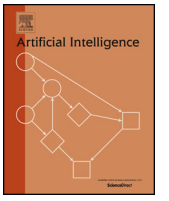

Situated conditional reasoning ☆

Giovanni Casini [a,b,*], Thomas Meyer [b,a,*], Ivan Varzinczak [c,b,a,*]

[a] CNR–ISTI, Italy
[b] CAIR, University of Cape Town, South Africa
[c] LIASD, Université Paris 8, France

ARTICLE INFO

Article history:
Received 13 September 2021
Received in revised form 12 March 2023
Accepted 27 March 2023
Available online 30 March 2023

Keywords:
Conditional reasoning
Non-monotonic reasoning
Counterfactual reasoning
Defeasible reasoning
Belief revision

ABSTRACT

Conditionals are useful for modelling many forms of everyday human reasoning but are not always sufficiently expressive to represent the information we want to reason about. In this paper, we make a case for a form of *situated* conditional. By 'situated', we mean that there is a context, based on an agent's beliefs and expectations, that works as background information in evaluating a conditional, and we allow such a context to vary. These conditionals are able to distinguish, for example, between *expectations* and *counterfactuals*. Formally, they are shown to generalise the conditional setting in the style of Kraus, Lehmann, and Magidor. We show that situated conditionals can be described in terms of a set of rationality postulates. We then propose an intuitive semantics for these conditionals and present a representation result which shows that our semantic construction corresponds exactly to the description in terms of postulates. With the semantics in place, we define a form of entailment for situated conditional knowledge bases, which we refer to as *minimal closure*. Finally, we proceed to show that it is possible to reduce the computation of minimal closure to a series of propositional entailment and satisfiability checks. While this is also the case for rational closure, it is somewhat surprising that the result carries over to minimal closure.

1. Introduction

Conditionals are at the heart of human everyday reasoning and play an important role in the logical formalisation of reasoning. They can usually be interpreted in many ways: as necessity [2,3], as presumption [4–6], normative [7,8], causal [9,10], probabilistic [11–13], counterfactual [14,15], and many others. Two very common interpretations that are also strongly interconnected are conditionals representing *expectations* ('If it is a bird, then presumably it flies') and conditionals representing *counterfactuals* ('If Napoleon had won at Waterloo, the whole of Europe would be speaking French'). Although they are connected by virtue of being conditionals, the types of reasoning they aim to model differ somewhat. For instance, in the first example above, the premise of the conditional is consistent with what is believed, while in the second example, the premise is inconsistent with an agent's beliefs. That this point is problematic can be made concrete with an extended version of the (admittedly overused) penguin example.

☆ This article is an extended and elaborated version of a paper presented at the 35th AAAI Conference on Artificial Intelligence (AAAI 2021) [1].

* Corresponding authors.
E-mail address: tmeyer@cs.uct.ac.za (T. Meyer).

https://doi.org/10.1016/j.artint.2023.103917
0004-3702/

Example 1.1. Suppose we know that birds usually fly, that penguins are birds that usually do not fly, that dodos were birds that usually did not fly, and that dodos do not exist anymore. As outlined in more detail in Example 3.1 later on, the standard preferential semantic approach to representing conditionals [5] is limited in that it allows for two forms of representation of an agent's beliefs. On the one, it would be impossible to distinguish between atypical (exceptional) entities, such as penguins, and non-existing entities, such as dodos (they are equally exceptional). On the other, it would be possible to draw this type of distinction, but at the expense of being unable to reason coherently about counterfactuals—the agent would be forced to conclude anything from the (nowadays absurd) existence of dodos.

By 'reasoning coherently' about counterfactuals, we mean being able to derive new information in a non-trivial way, formalising reasoning patterns that can be recognised as 'rational'. As mentioned in the above example and formally exemplified in Example 3.1 later on, using the classical semantic solutions to reason simultaneously about what is actually plausible and what is counterfactually plausible has some strong limitations. For example, the formalism we use could force us to reason about counterfactuals in a trivial way, relying on the *ex falso quodlibet* principle. In this work, we introduce a logic of *situated* conditionals to overcome precisely this problem. By 'situated', we mean that there is a context, based on an agent's beliefs and expectations, which works as background information in evaluating a conditional. This is particularly important when we consider counterfactual conditionals. The central insight is that adding an explicit notion of situation to standard conditionals allows for a refined semantics of this enriched language in which the problems described in Example 1.1 can be dealt with adequately. It also allows us to reason coherently with counterfactual conditionals such as 'Had Mauritius not been colonised, the dodo would not fly'.[1] That is, the premise of a counterfactual can be inconsistent with the agent's beliefs without lapsing into the triviality of the *ex falso quodlibet* principle. Moreover, it is possible to reason coherently with situated conditionals without knowing whether their premises are plausible or counterfactual. In the case of penguins and dodos, for example, it allows us to state that penguins usually do not fly in a situation where penguins exist and that dodos usually do not fly in a situation where dodos also exist while being unaware of whether or not penguins and dodos actually exist. At the same time, it remains possible to make classical statements specifying what necessarily holds (e.g., stating that penguins and dodos are birds as a necessary fact).

Counterfactual and defeasible reasoning have been important topics of research in knowledge representation and reasoning since the beginning of the AI endeavour [16,17]. Still, they have usually been formalised as the same form of conditional reasoning. While research on defeasible reasoning has always been quite active, the importance of counterfactuals in AI has become more apparent only recently, especially considering the role that counterfactuals play in causal reasoning [18] and in eXplainable AI (XAI) [19–21]. The increasing attention in the area of XAI to this topic means that a detailed formal analysis of counterfactuals and their associated reasoning systems is not just a timely research endeavour but a necessary one. In this context, we put forward a framework for managing both counterfactual reasoning and reasoning based on expectations—one that avoids some of the limitations associated with previous conditional approaches based on a preferential semantics.

The remainder of the paper is organised as follows. Section 2 outlines the formal background on propositional logic and on the preferential approach to conditionals on which our work is based. Section 3 is the heart of the paper. It describes the language of situated conditionals, furnishes it with an appropriate and intuitive semantics, and motivates the corresponding logic through examples, formal postulates, and a formal representation result. With the basics of the logic in place, Section 4 defines a form of entailment for it that is based on the well-known notion of *rational closure* [5]. As such, it plays a role similar to the one that rational closure plays for reasoning with conditionals—it is a basic form of entailment on which other forms of logical consequence can be constructed. Section 5 shows that, from a computational perspective, the version of entailment we propose in the previous section is reducible to classical propositional reasoning. Section 6 reviews related work, while Section 7 concludes and considers future avenues to explore. Long proofs are in the appendix.

2. Formal background

In this paper, we assume a finite set of propositional *atoms* $\mathcal{P}$ and use $p, q, \ldots$ as metavariables to denote its elements. Sentences of the underlying propositional language are denoted by $\alpha, \beta, \ldots$, and are built up from the atomic propositions and the standard Boolean connectives in the usual way. The set of all propositional sentences is denoted by $\mathcal{L}$.

A *valuation* (alias *world*) is a function from $\mathcal{P}$ into $\{0, 1\}$. The set of all valuations is denoted $\mathcal{U}$, and we use $u, v, \ldots$ to denote its elements. Whenever it eases presentation, we represent valuations as sequences of atoms (e.g., p) and barred atoms (e.g., $\overline{p}$), with the usual understanding. As an example, if $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$, with the atoms standing for, respectively, 'being a bird', 'being a flying creature', and 'being a penguin', then the valuation $\mathsf{b}\bar{\mathsf{f}}\mathsf{p}$ conveys the idea that b is true, f is false, and p is true.

With $v \Vdash \alpha$ we denote the fact that the valuation v *satisfies* the sentence α. Given $\alpha \in \mathcal{L}$, with $[\![\alpha]\!] \stackrel{\text{def}}{=} \{v \in \mathcal{U} \mid v \Vdash \alpha\}$ we denote its *models*. For $X \subseteq \mathcal{L}$, $[\![X]\!] \stackrel{\text{def}}{=} \bigcap_{\alpha \in X} [\![\alpha]\!]$. We say $X \subseteq \mathcal{L}$ (classically) *entails* $\alpha \in \mathcal{L}$, denoted $X \models \alpha$, if $[\![X]\!] \subseteq [\![\alpha]\!]$. Given a set of valuations V, $\mathsf{sent}(V)$ indicates a sentence characterising the set V. That is, $\mathsf{sent}(V)$ is a propositional sentence satisfied by all, and only, the valuations in V.

[1] The extinction of dodos in the 17th century is considered to be a consequence of the colonisation of Mauritius.

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

Fig. 1. A ranked interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$ with both $\mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$ and $\mathcal{U}^{\infty}_{\mathscr{R}}$ nonempty.

2.1. KLM-style rational defeasible consequence

A *defeasible consequence relation* $\mathrel{|\!\sim}$ is a binary relation on $\mathcal{L}$. Intuitively, the fact that $(\alpha, \beta) \in \mathrel{|\!\sim}$, which is usually represented as the statement $\alpha \mathrel{|\!\sim} \beta$, captures the idea that "β is a defeasible consequence of α", or, in other words, that "if α, then usually (alias normally, or typically) β". The relation $\mathrel{|\!\sim}$ is said to be *preferential* [4] if it satisfies the well-known KLM preferential postulates below:

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

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

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

If on top of the preferential postulates the relation $\mathrel{|\!\sim}$ also satisfies the following rational monotonicity postulate, then $\mathrel{|\!\sim}$ is said to be *rational*:

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

The merits of these postulates have been addressed extensively in the literature [4,22], and we shall not repeat them here.

A suitable semantics for rational consequence relations is provided by ordered structures called *ranked interpretations* (alternative semantics have also been proposed in the literature, in particular, Spohn's ordinal conditional functions [23]).

Definition 2.1 *(Ranked Interpretation).* A **ranked interpretation** $\mathscr{R}$ is a total function from $\mathcal{U}$ to $\mathbb{N} \cup \{\infty\}$, satisfying the following **convexity property**: for every $u \in \mathcal{U}$ and every $i \in \mathbb{N}$, if $\mathscr{R}(u) = i$, then, for every j s.t. $0 \leq j < i$, there is a $u' \in \mathcal{U}$ for which $\mathscr{R}(u') = j$.

For a given ranked interpretation $\mathscr{R}$ and valuation v, we denote with $\mathscr{R}(v)$ the *rank of* v. The number $\mathscr{R}(v)$ indicates the degree of *atypicality* of v. So the valuations judged most typical are those with rank 0, while those with an infinite rank are deemed so atypical as to be seen as implausible. We can therefore partition the set $\mathcal{U}$ w.r.t. $\mathscr{R}$ into the set of *plausible* valuations $\mathcal{U}^{\mathtt{f}}_{\mathscr{R}} \stackrel{\text{def}}{=} \{u \in \mathcal{U} \mid \mathscr{R}(u) \in \mathbb{N}\}$, and *implausible* valuations $\mathcal{U}^{\infty}_{\mathscr{R}} \stackrel{\text{def}}{=} \mathcal{U} \setminus \mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$. It goes without saying that $\mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$ or $\mathcal{U}^{\infty}_{\mathscr{R}}$ (but not both) can be empty. (Throughout the paper, we shall use the symbol $\mathtt{f}$ to refer to finiteness.) With $[\![i]\!]_{\mathscr{R}}$, for $i \in \mathbb{N} \cup \{\infty\}$, we indicate all the valuations with rank i in $\mathscr{R}$ (we omit the subscript whenever it is clear from the context).

Assuming $\mathcal{P} = \{\mathsf{b}, \mathsf{f}, \mathsf{p}\}$, with the intuitions as above, Fig. 1 shows an example of a ranked interpretation.

Let $\mathscr{R}$ be a ranked interpretation and let $\alpha \in \mathcal{L}$. Then $[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}} \stackrel{\text{def}}{=} \mathcal{U}^{\mathtt{f}}_{\mathscr{R}} \cap [\![\alpha]\!]$, and $\min[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}} \stackrel{\text{def}}{=} \{u \in [\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}} \mid \mathscr{R}(u) \leq \mathscr{R}(v)$, for all $v \in [\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}}\}$. A defeasible consequence relation $\alpha \mathrel{|\!\sim} \beta$ can be given an intuitive semantics in terms of ranked interpretations as follows: $\alpha \mathrel{|\!\sim} \beta$ is *satisfied in* $\mathscr{R}$ (denoted $\mathscr{R} \Vdash \alpha \mathrel{|\!\sim} \beta$) if $\min[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}} \subseteq [\![\beta]\!]$, with $\mathscr{R}$ referred to as a *ranked model* of $\alpha \mathrel{|\!\sim} \beta$. In the example in Fig. 1, we have $\mathscr{R} \Vdash \mathsf{b} \mathrel{|\!\sim} \mathsf{f}$, $\mathscr{R} \Vdash \neg(\mathsf{p} \rightarrow \mathsf{b}) \mathrel{|\!\sim} \bot$, $\mathscr{R} \Vdash \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}$, $\mathscr{R} \nVdash \mathsf{f} \mathrel{|\!\sim} \mathsf{b}$, and $\mathscr{R} \Vdash \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \mathsf{b}$. It is easily verified that $\mathscr{R} \Vdash \neg\alpha \mathrel{|\!\sim} \bot$ if and only if $\mathcal{U}^{\mathtt{f}}_{\mathscr{R}} \subseteq [\![\alpha]\!]$. Hence we frequently abbreviate $\neg\alpha \mathrel{|\!\sim} \bot$ as α. Two defeasible statements $\alpha \mathrel{|\!\sim} \beta$ and $\gamma \mathrel{|\!\sim} \delta$ are said to be *rank equivalent* if they have the same ranked models, i.e., if for every ranked interpretation $\mathscr{R}$, $\mathscr{R} \Vdash \alpha \mathrel{|\!\sim} \beta$ if and only if $\mathscr{R} \Vdash \gamma \mathrel{|\!\sim} \delta$.

The correspondence between rational consequence relations and ranked interpretations is formalised by the following representation result.

Theorem 2.1 *(Lehmann & Magidor [5]; Gärdenfors & Makinson [24]). A defeasible consequence relation $\mathrel{|\!\sim}$ is rational iff there is a ranked interpretation $\mathscr{R}$ such that, for every pair of formulae α and β, $\alpha \mathrel{|\!\sim} \beta$ iff $\mathscr{R} \Vdash \alpha \mathrel{|\!\sim} \beta$.*

2.2. Rational closure

It is possible to represent knowledge as a set of defeasible statements and to use such a set to infer other defeasible statements from it. This is the stance adopted by Lehmann and Magidor [5]. A *conditional knowledge base* $\mathcal{C}$ is a finite set of

defeasible statements of the form $\alpha \mathrel{|\!\sim} \beta$, with $\alpha, \beta \in \mathcal{L}$. As before, in knowledge bases, we shall also abbreviate $\neg\alpha \mathrel{|\!\sim} \bot$ with α. As an example, let $\mathcal{C} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \rightarrow \mathsf{b}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}\}$.

Given a conditional knowledge base $\mathcal{C}$, a *ranked model* of $\mathcal{C}$ is a ranked interpretation satisfying all statements in $\mathcal{C}$. As it turns out, the ranked interpretation in Fig. 1 is a ranked model of the above $\mathcal{C}$. It is not hard to see that, in every ranked model of $\mathcal{C}$, the valuations $\overline{\mathsf{b}}\mathsf{f}\mathsf{p}$ and $\overline{\mathsf{b}}\overline{\mathsf{f}}\mathsf{p}$ are deemed implausible—note, however, that they are still possible from a logical point of view, which is the reason why they feature in all ranked interpretations.

A conditional knowledge base $\mathcal{C}$ is *consistent* if it has a ranked model $\mathscr{R}$ s.t. $[\![0]\!]_{\mathscr{R}} \neq \emptyset$. That is, $\mathcal{C}$ is consistent if it has a ranked model $\mathscr{R}$ that does not satisfy $\top \mathrel{|\!\sim} \bot$. Two conditional knowledge bases are *rank equivalent* if they have exactly the same ranked models.

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

To be more precise, we can take the defeasible conditionals in $\mathcal{C}$ as the core elements of a defeasible consequence relation $\mathrel{|\!\sim}^{\mathcal{C}}$. By closing the latter under the preferential postulates (in the sense of exhaustively applying them as rules), we get a *preferential extension* of $\mathrel{|\!\sim}^{\mathcal{C}}$. 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 $\mathrel{|\!\sim}^{\mathcal{C}}$, which we denote by $\mathrel{|\!\sim}^{\mathcal{C}}_{PC}$. It turns out that the preferential closure of $\mathrel{|\!\sim}^{\mathcal{C}}$ contains exactly the conditionals entailed by $\mathcal{C}$. (Hence, the notions of closure of and entailment from a conditional knowledge base are two sides of the same coin.)

The same process and definitions from above 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 $\mathrel{|\!\sim}^{\mathcal{C}}$. Nevertheless, as pointed out by Lehmann and Magidor [5, Section 4.2], the intersection of all such rational extensions does not generally yield 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 $\mathcal{C} \models_{\mathscr{R}} \alpha \mathrel{|\!\sim} \beta$ if every ranked model of $\mathcal{C}$ also satisfies $\alpha \mathrel{|\!\sim} \beta$, is *monotonic* (in that it is defined as a standard Tarskian entailment relation). Therefore rank entailment 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 [25–31], 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* [25, 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).

Multiple equivalent characterisations of RC have been proposed [5,32,26,33,34], and here we rely on the one by Giordano and others [29]. Assume an ordering $\preceq_{\mathcal{C}}$ on all ranked models of a knowledge base $\mathcal{C}$, which is defined as follows: $\mathscr{R}_1 \preceq_{\mathcal{C}} \mathscr{R}_2$, if, for every $v \in \mathcal{U}$, $\mathscr{R}_1(v) \leq \mathscr{R}_2(v)$. Intuitively, ranked models lower down in the ordering correspond to descriptions of the world in which the typicality of each situation (valuation) is maximised. It is easy to see that $\preceq_{\mathcal{C}}$ is a weak partial order. Giordano et al. [29] showed that there is a unique $\preceq_{\mathcal{C}}$-minimal element. The rational closure of $\mathcal{C}$ is defined in terms of this minimum ranked model of $\mathcal{C}$.

Definition 2.2 *(Rational closure).* Let $\mathcal{C}$ be a conditional knowledge base, and let $\mathscr{R}^{\mathcal{C}}_{RC}$ be the minimum element of $\preceq_{\mathcal{C}}$ on ranked models of $\mathcal{C}$. The **rational closure** of $\mathcal{C}$ is the defeasible consequence relation $\mathrel{|\!\sim}^{\mathcal{C}}_{RC} \stackrel{\text{def}}{=} \{\alpha \mathrel{|\!\sim} \beta \mid \mathscr{R}^{\mathcal{C}}_{RC} \Vdash \alpha \mathrel{|\!\sim} \beta\}$.

As an example, Fig. 1 shows the minimum ranked model of $\mathcal{C} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \rightarrow \mathsf{b}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}\}$ w.r.t. $\preceq_{\mathcal{C}}$. Hence we have that $\neg\mathsf{f} \mathrel{|\!\sim} \neg\mathsf{b}$ is in the rational closure of $\mathcal{C}$ (but note it is not in the preferential closure of $\mathcal{C}$).

Observe that there are two levels of typicality at work for rational closure, namely *within* ranked models of $\mathcal{C}$, where valuations lower down are viewed as more typical, but also *between* ranked models of $\mathcal{C}$, where ranked models lower down in the ordering are viewed as more typical. The most typical ranked model $\mathscr{R}^{\mathcal{C}}_{RC}$ is the one in which valuations are as typical as $\mathcal{C}$ 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 can be and have been constructed [25,35,36,30,31].

3. Situated conditionals

We now turn to the heart of the paper, the introduction of a logic-based formalism for the specification of and reasoning with situated conditionals. For a more detailed motivation, let us consider a more technical version of the penguin-dodo example introduced in Section 1.

Example 3.1. We know that birds usually fly ($\mathsf{b} \mathrel{|\!\sim} \mathsf{f}$), and that penguins are birds ($\mathsf{p} \rightarrow \mathsf{b}$) that usually do not fly ($\mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}$). Also, we know that dodos were birds ($\mathsf{d} \rightarrow \mathsf{b}$) that usually did not fly ($\mathsf{d} \mathrel{|\!\sim} \neg\mathsf{f}$), and that dodos do not exist anymore. Using the standard ranked semantics (Definition 2.1), we have two ways of modelling the information above.

∞	$\mathcal{U} \setminus ([\![0]\!] \cup [\![1]\!] \cup [\![2]\!])$
2	$\mathsf{p}\bar{\mathsf{d}}\mathsf{bf}$ $\bar{\mathsf{p}}\mathsf{dbf}$ pdbf
1	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{bf}$ $\mathsf{pdb}\bar{\mathsf{f}}$ $\bar{\mathsf{p}}\mathsf{db}\bar{\mathsf{f}}$ $\mathsf{pdb}\bar{\mathsf{f}}$
0	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{bf}$ $\bar{\mathsf{p}}\bar{\mathsf{d}}\bar{\mathsf{b}}\mathsf{f}$ $\bar{\mathsf{p}}\bar{\mathsf{d}}\bar{\mathsf{b}}\bar{\mathsf{f}}$

∞	$\mathcal{U} \setminus ([\![0]\!] \cup [\![1]\!] \cup [\![2]\!])$
2	$\mathsf{p}\bar{\mathsf{d}}\mathsf{bf}$
1	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{bf}$ $\mathsf{pdb}\bar{\mathsf{f}}$
0	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{bf}$ $\bar{\mathsf{p}}\bar{\mathsf{d}}\bar{\mathsf{b}}\mathsf{f}$ $\bar{\mathsf{p}}\bar{\mathsf{d}}\bar{\mathsf{b}}\bar{\mathsf{f}}$

Fig. 2. Left: minimal ranked model of the KB in Example 3.1 satisfying $\top \mathrel{|\!\sim} \neg\mathsf{d}$. Right: minimal ranked model of the KB expanded with $\mathsf{d} \mathrel{|\!\sim} \bot$.

$\langle \infty, \infty \rangle$	$[\![\mathsf{p} \wedge \neg\mathsf{b}]\!] \cup [\![\mathsf{d} \wedge \neg\mathsf{b}]\!]$
$\langle \infty, 1 \rangle$	$\bar{\mathsf{p}}\mathsf{dbf}$ pdbf
$\langle \infty, 0 \rangle$	$\bar{\mathsf{p}}\mathsf{db}\bar{\mathsf{f}}$ $\mathsf{pdb}\bar{\mathsf{f}}$
$\langle \mathtt{f}, 2 \rangle$	$\mathsf{p}\bar{\mathsf{d}}\mathsf{bf}$
$\langle \mathtt{f}, 1 \rangle$	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{bf}$ $\mathsf{p}\bar{\mathsf{d}}\mathsf{b}\bar{\mathsf{f}}$
$\langle \mathtt{f}, 0 \rangle$	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{bf}$ $\bar{\mathsf{p}}\bar{\mathsf{d}}\bar{\mathsf{b}}\mathsf{f}$ $\bar{\mathsf{p}}\bar{\mathsf{d}}\bar{\mathsf{b}}\bar{\mathsf{f}}$

Fig. 3. Epistemic interpretation for $\mathcal{P} = \{\mathsf{b}, \mathsf{d}, \mathsf{f}, \mathsf{p}\}$.

The first option is to formalise what an agent believes by referring to the valuations with rank 0 in a ranked interpretation. That is, the agent believes α is true if and only if $\top \mathrel{|\!\sim} \alpha$ holds. In such a case, $\top \mathrel{|\!\sim} \neg\mathsf{d}$ means that the agent believes that dodos do not exist. The minimal model for this conditional knowledge base is shown in Fig. 2 (left). The main limitation of this representation is that all exceptional entities have the same status as dodos since they cannot be satisfied at rank 0. Hence, one of the consequences of the agent's beliefs is the statement $\top \mathrel{|\!\sim} \neg\mathsf{p}$, just as we have $\top \mathrel{|\!\sim} \neg\mathsf{d}$, and, as a result, we are not able to distinguish between the status of the dodos (they do not exist anymore) and the status of the penguins (they do exist and are simply exceptional birds).

The second option is to represent what an agent believes in terms of all valuations with finite ranks. That is, an agent believes α to hold if and only if $\neg\alpha \mathrel{|\!\sim} \bot$ holds. If dodos do not exist, we add the statement $\mathsf{d} \mathrel{|\!\sim} \bot$. The minimal model for this case is depicted in Fig. 2 (right). Here we can distinguish between what is considered false (dodos exist) and what is exceptional (penguins), but we are unable to reason coherently about counterfactuals since from $\mathsf{d} \mathrel{|\!\sim} \bot$ we can conclude anything about dodos (via $\models \bot \rightarrow \alpha$ and RW, for any $\alpha \in \mathcal{L}$).

A *situated conditional* (SC for short) is a statement of the form $\alpha \mathrel{|\!\sim}_\gamma \beta$, with $\alpha, \beta, \gamma \in \mathcal{L}$, which is read as 'given the situation γ, β usually holds on condition that α holds'. Formally, a situated conditional $\mathrel{|\!\sim}$ is a ternary relation on $\mathcal{L}$. We shall write $\alpha \mathrel{|\!\sim}_\gamma \beta$ as an abbreviation for $\langle \alpha, \beta, \gamma \rangle \in \mathrel{|\!\sim}$. To provide a suitable semantics for SCs, we define a refined version of the ranked interpretations of Section 2 that we refer to as *epistemic interpretations*. Following the provision of the semantics, we illustrate it with a representative example in Example 3.2.

A ranked interpretation can differentiate between plausible valuations (those in $\mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$) but not between implausible ones (those in $\mathcal{U}^{\infty}_{\mathscr{R}}$). In contrast, an epistemic interpretation can also tell implausible valuations apart. We thus distinguish between two classes of valuations: plausible valuations with a *finite rank*, and implausible valuations with an *infinite rank*. Within implausible valuations, we further distinguish between those considered as *possible* and those that would be *impossible*. This is formalised by assigning to each valuation u a tuple of the form $\langle \mathtt{f}, i \rangle$, where $i \in \mathbb{N}$, or $\langle \infty, i \rangle$, where $i \in \mathbb{N} \cup \{\infty\}$. The $\mathtt{f}$ in $\langle \mathtt{f}, i \rangle$ is meant to indicate that u has a *finite rank*, while the ∞ in $\langle \infty, i \rangle$ is intended to denote that u has an *infinite rank*, where finite ranks are viewed as more typical than infinite ranks. Implausible valuations that are considered possible have an infinite rank $\langle \infty, i \rangle$, where $i \in \mathbb{N}$, while those considered impossible have the infinite rank $\langle \infty, \infty \rangle$, where $\langle \infty, \infty \rangle$ is taken to be less expected than any of the other infinite ranks.

To capture this formally, let $Rk \stackrel{\text{def}}{=} \{\langle \mathtt{f}, i \rangle \mid i \in \mathbb{N}\} \cup \{\langle \infty, i \rangle \mid i \in \mathbb{N} \cup \{\infty\}\}$ denote henceforth the set of all possible *ranks*. We define the total ordering $\preceq$ over Rk as follows: $\langle x_1, y_1 \rangle \preceq \langle x_2, y_2 \rangle$ if $x_1 = x_2$ and $y_1 \leq y_2$, or $x_1 = \mathtt{f}$ and $x_2 = \infty$, where $i < \infty$ for all $i \in \mathbb{N}$.

Definition 3.1 *(Epistemic Interpretation).* An **epistemic interpretation** $\mathscr{E}$ is a total function from $\mathcal{U}$ to Rk for which the following **convexity** property holds: (*i*) for every $u \in \mathcal{U}$ and every $i \in \mathbb{N}$, if $\mathscr{E}(u) = \langle \mathtt{f}, i \rangle$, then, for all j s.t. $0 \leq j < i$, there is a $u_j \in \mathcal{U}$ s.t. $\mathscr{E}(u_j) = \langle \mathtt{f}, j \rangle$, and (*ii*) for every $u \in \mathcal{U}$ and every $i \in \mathbb{N}$, if $\mathscr{E}(u) = \langle \infty, i \rangle$, then, for all j s.t. $0 \leq j < i$, there is a $u_j \in \mathcal{U}$ s.t. $\mathscr{E}(u_j) = \langle \infty, j \rangle$.

Observe that the version of convexity satisfied by epistemic interpretations is a straightforward extension of the convexity of ranked interpretations (Definition 2.1). Fig. 3 depicts an epistemic interpretation in our running example.

Casini et al. [37] have a similar definition of epistemic interpretations, but they do not allow for the rank $\langle \infty, \infty \rangle$.

We let $\mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \stackrel{\text{def}}{=} \{u \in \mathcal{U} \mid \mathscr{E}(u) = \langle \mathtt{f}, i \rangle$, for some $i \in \mathbb{N}\}$ and $\mathcal{U}^{\infty}_{\mathscr{E}} \stackrel{\text{def}}{=} \{u \in \mathcal{U} \mid \mathscr{E}(u) = \langle \infty, i \rangle$, for some $i \in \mathbb{N}\}$. Note that $\mathcal{U}^{\infty}_{\mathscr{E}}$ does *not* contain valuations with rank $\langle \infty, \infty \rangle$. We let $\min[\![\alpha]\!]_{\mathscr{E}} \stackrel{\text{def}}{=} \{u \in [\![\alpha]\!] \mid \mathscr{E}(u) \preceq \mathscr{E}(v)$, for all $v \in [\![\alpha]\!]\}$, $\min[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{E}} \stackrel{\text{def}}{=} \{u \in [\![\alpha]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \mid \mathscr{E}(u) \preceq \mathscr{E}(v)$, for all $v \in [\![\alpha]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{E}}\}$, and $\min[\![\alpha]\!]^{\infty}_{\mathscr{E}} \stackrel{\text{def}}{=} \{u \in [\![\alpha]\!] \cap \mathcal{U}^{\infty}_{\mathscr{E}} \mid \mathscr{E}(u) \preceq \mathscr{E}(v)$, for all $v \in [\![\alpha]\!] \cap \mathcal{U}^{\infty}_{\mathscr{E}}\}$.

Observe that epistemic interpretations are allowed to have no plausible valuations ($\mathcal{U}^{\mathtt{f}}_{\mathscr{E}} = \emptyset$), as well as no implausible valuations that are possible ($\mathcal{U}^{\infty}_{\mathscr{E}} = \emptyset$). This means it is possible that $\mathscr{E}(u) = \langle \infty, \infty \rangle$ for all $u \in \mathcal{U}$, in which case $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$,

for all α, β, γ (cf. Definition 3.2 below). Epistemic interpretations also allow for cases where all valuations are possible (that is, either plausible or implausible but possible). This corresponds to the case where an epistemic interpretation does not have any valuation with rank $\langle \infty, \infty \rangle$.

Armed with the notion of epistemic interpretation, we can provide an intuitive semantics to situated conditionals.

Definition 3.2 *(Satisfaction and generation of situated conditionals).* Let $\mathscr{E}$ be an epistemic interpretation and let $\alpha, \beta, \gamma \in \mathcal{L}$. We say $\mathscr{E}$ **satisfies** $\alpha \mathrel{|\!\sim_\gamma} \beta$, denoted as $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\gamma} \beta$ and often abbreviated as $\alpha \mathrel{|\!\sim^{\mathscr{E}}_\gamma} \beta$, if

$$\begin{cases} \min[\![\alpha \wedge \gamma]\!]^{\mathtt{f}}_{\mathscr{E}} \subseteq [\![\beta]\!], & \text{if } [\![\gamma]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \neq \emptyset; \\ \min[\![\alpha \wedge \gamma]\!]^{\infty}_{\mathscr{E}} \subseteq [\![\beta]\!], & \text{otherwise.} \end{cases}$$

We say that $\mathscr{E}$ **generates** the situated conditional $\mathrel{|\!\sim}$ if, for every $\alpha, \beta, \gamma \in \mathcal{L}$, $\langle \alpha, \beta, \gamma \rangle \in \mathrel{|\!\sim}$ iff $\alpha \mathrel{|\!\sim^{\mathscr{E}}_\gamma} \beta$.

Intuitively, the satisfaction of situated conditionals works as follows. If the situation γ is compatible with the plausible part of $\mathscr{E}$ (the valuations in $\mathcal{U}^{\mathtt{f}}_{\mathscr{E}}$), then $\alpha \mathrel{|\!\sim_\gamma} \beta$ holds if the most typical plausible models of $\alpha \wedge \gamma$ are also models of β. On the other hand, if the situation γ is not compatible with the plausible part of $\mathscr{E}$, i.e., all models of γ have an infinite rank, then $\alpha \mathrel{|\!\sim_\gamma} \beta$ holds if the most typical implausible (but possible) models of $\alpha \wedge \gamma$ are also models of β.

An immediate corollary of Definition 3.2 is that the rational conditionals defined in terms of ranked interpretations can be simulated with SCs by setting the situation to $\top$.

Definition 3.3 *(Extracted ranked interpretation).* For an epistemic interpretation $\mathscr{E}$, we define the **ranked interpretation** $\mathscr{R}^{\mathscr{E}}$ **extracted from** $\mathscr{E}$ as follows: for $u \in \mathcal{U}^{\mathtt{f}}_{\mathscr{E}}$, $\mathscr{R}^{\mathscr{E}}(u) = i$, where $\mathscr{E}(u) = \langle \mathtt{f}, i \rangle$, and $\mathscr{R}^{\mathscr{E}}(u) = \infty$ for $u \in \mathcal{U} \setminus \mathcal{U}^{\mathtt{f}}_{\mathscr{E}}$.

Corollary 3.1. *Let $\mathscr{E}$ be an epistemic interpretation. Then $\mathscr{R}^{\mathscr{E}} \Vdash \alpha \mathrel{|\!\sim} \beta$ iff $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\top} \beta$.*

Proof. Assume $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\top} \beta$. Then, by definition, we have $\min[\![\alpha \wedge \top]\!]^{\mathtt{f}}_{\mathscr{E}} \subseteq [\![\beta]\!]$ if $\mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \neq \emptyset$, and $\min[\![\alpha \wedge \top]\!]^{\infty}_{\mathscr{E}} \subseteq [\![\beta]\!]$ otherwise. If the former is the case, then, by the construction of $\mathscr{R}^{\mathscr{E}}$, we have $\min[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}^{\mathscr{E}}} \subseteq [\![\beta]\!]$, and therefore $\mathscr{R}^{\mathscr{E}} \Vdash \alpha \mathrel{|\!\sim} \beta$. If, instead, the latter holds, then $[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{E}} = \emptyset$, from which it follows that $[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}^{\mathscr{E}}} = \emptyset$, and therefore $\mathscr{R}^{\mathscr{E}} \Vdash \alpha \mathrel{|\!\sim} \beta$. For the other direction, assume $\mathscr{R}^{\mathscr{E}} \Vdash \alpha \mathrel{|\!\sim} \beta$. If $[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}^{\mathscr{E}}} = \emptyset$, then, from the construction of $\mathscr{R}^{\mathscr{E}}$, we have $[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{E}} = \emptyset$, from which we get $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\top} \beta$. If $[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}^{\mathscr{E}}} \neq \emptyset$, then, since $\min[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{R}^{\mathscr{E}}} \subseteq [\![\beta]\!]$, we must have $\min[\![\alpha]\!]^{\mathtt{f}}_{\mathscr{E}} \subseteq [\![\beta]\!]$, too. From the latter it follows that $\min[\![\alpha \wedge \top]\!]^{\mathtt{f}}_{\mathscr{E}} \subseteq [\![\beta]\!]$, and therefore $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\top} \beta$. □

The principal advantage of situated conditionals and their associated enriched semantics in terms of epistemic interpretations is that they allow us to represent different degrees of epistemic involvement, with the finite ranks (the plausible valuations) representing the expectations of an agent. So $\top \mathrel{|\!\sim_\top} \alpha$ being satisfied in $\mathscr{E}$ indicates that α is expected. What an agent believes to be true corresponds to what is true in all the valuations with finite ranks. That is, the agent believes α to be true if and only if $\mathscr{E} \Vdash \neg\alpha \mathrel{|\!\sim_\top} \bot$, and we will abbreviate $\neg\alpha \mathrel{|\!\sim_\top} \bot$ with α, extending to epistemic interpretations the convention introduced above for ranked interpretations (see Section 2.2).

Another advantage of our framework is that it also allows for reasoning counterfactual: we can express that dodos would not fly if they existed in a coherent way. We can talk about dodos in a counterfactual situation or context, for example, assuming that Mauritius had never been colonised ($\neg\mathsf{mc}$): the conditional $\mathsf{d} \mathrel{|\!\sim_{\neg\mathsf{mc}}} \neg\mathsf{f}$ is read as 'In the situation of Mauritius not having been colonised, the dodo would not fly'. Importantly, we can reason coherently with a situated conditional, even when not knowing whether its premises are plausible or counterfactual. To do so, it is sufficient to introduce statements of the form $\alpha \mathrel{|\!\sim_\alpha} \beta$. If α is plausible, this conditional is evaluated in the context of the finite ranks, exactly as if $\alpha \mathrel{|\!\sim_\top} \beta$ were being evaluated. On the other hand, if $\alpha \mathrel{|\!\sim_\top} \bot$ holds, $\alpha \mathrel{|\!\sim_\alpha} \beta$ will be evaluated referring to the infinite ranks. So, in the case of penguins and dodos, $\mathsf{p} \mathrel{|\!\sim_\mathsf{p}} \neg\mathsf{f}$ and $\mathsf{d} \mathrel{|\!\sim_\mathsf{d}} \neg\mathsf{f}$ express the information that penguins usually do not fly in the situation of penguins existing, and that dodos usually do not fly in the situation of dodos existing, regardless of whether the agent is aware of penguins or dodos existing or not. In contrast, a statement such as $\mathsf{d} \mathrel{|\!\sim_\top} \neg\mathsf{f}$ cannot be used to reason counterfactually about dodos, once we are aware that they do not exist (that is, $\mathsf{d} \mathrel{|\!\sim_\top} \bot$): given the latter, once we consider all the valuations satisfying $\top$ (that is, all the valuations), we have to evaluate every defeasible conditional $\mathsf{d} \mathrel{|\!\sim} \alpha$ (for any α) looking at the valuations with finite ranks. In all such valuations, the sentence d is not satisfied; hence any SC $\mathsf{d} \mathrel{|\!\sim_\top} \alpha$, for any α, would be satisfied. Also, note that it is still possible to impose that something necessarily holds, both in plausible and counterfactual situations. The conditional $\alpha \mathrel{|\!\sim_\alpha} \bot$ holds only in epistemic interpretations in which all valuations satisfying α have $\langle \infty, \infty \rangle$ as their rank. The following example illustrates these claims more concretely.

Example 3.2. Consider the following rephrasing of the statements in Example 3.1. 'Birds usually fly' becomes $\mathsf{b} \mathrel{|\!\sim_\top} \mathsf{f}$. Defeasible information about penguins and dodos is modelled using $\mathsf{p} \mathrel{|\!\sim_\mathsf{p}} \neg\mathsf{f}$ and $\mathsf{d} \mathrel{|\!\sim_\mathsf{d}} \neg\mathsf{f}$. Given that dodos don't exist anymore, the statement $\mathsf{d} \mathrel{|\!\sim_\top} \bot$ leaves open the existence of dodos in the infinite ranks, which allow for coherent reasoning under the assumption that dodos exist (the situation d). Moreover, information such as dodos and penguins necessarily being birds

can be modelled by the conditionals $\mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim}_{\mathsf{p}\wedge\neg\mathsf{b}} \bot$ and $\mathsf{d} \wedge \neg\mathsf{b} \mathrel{|\!\sim}_{\mathsf{d}\wedge\neg\mathsf{b}} \bot$, relegating the valuations in $[\![\mathsf{p} \wedge \neg\mathsf{b}]\!] \cup [\![\mathsf{d} \wedge \neg\mathsf{b}]\!]$ to the rank $\langle\infty, \infty\rangle$. Fig. 3 (below Definition 3.1) shows a model of these statements. (We shall address how certain models of given conditionals are excluded from the picture in Section 4, where we define a suitable form of entailment from a set of situated conditionals.)

Next, we consider the class of situated conditionals from the perspective of a list of *situated* rationality postulates in the KLM style. We start with the following ones:

$$\text{(Ref)}\ \ \alpha \mathrel{|\!\sim}_\gamma \alpha \qquad\qquad \text{(LLE)}\ \ \frac{\models \alpha \leftrightarrow \beta,\ \alpha \mathrel{|\!\sim}_\gamma \delta}{\beta \mathrel{|\!\sim}_\gamma \delta}$$

$$\text{(And)}\ \ \frac{\alpha \mathrel{|\!\sim}_\gamma \beta,\ \alpha \mathrel{|\!\sim}_\gamma \delta}{\alpha \mathrel{|\!\sim}_\gamma \beta \wedge \delta} \qquad \text{(Or)}\ \ \frac{\alpha \mathrel{|\!\sim}_\gamma \delta,\ \beta \mathrel{|\!\sim}_\gamma \delta}{\alpha \vee \beta \mathrel{|\!\sim}_\gamma \delta}$$

$$\text{(RW)}\ \ \frac{\alpha \mathrel{|\!\sim}_\gamma \beta,\ \models \beta \rightarrow \delta}{\alpha \mathrel{|\!\sim}_\gamma \delta} \qquad \text{(RM)}\ \ \frac{\alpha \mathrel{|\!\sim}_\gamma \beta,\ \alpha \not\mathrel{|\!\sim}_\gamma \neg\delta}{\alpha \wedge \delta \mathrel{|\!\sim}_\gamma \beta}$$

Observe that they correspond exactly to the original KLM postulates, except that the notion of situation has been added. As for α and β, the γ occurring in the postulates should be viewed as a meta-variable ranging over $\mathcal{L}$.

Definition 3.4 *(Basic situated conditional).* An SC $\mathrel{|\!\sim}$ is a **basic situated conditional** (BSC, for short) if it satisfies the situated rationality postulates.

An immediate corollary of this definition is that for a BSC with the situation γ fixed, $\mathrel{|\!\sim}_\gamma$ is a rational conditional. We then get the following result.

Theorem 3.1. *Every epistemic interpretation generates a BSC (see Definition 3.2). Nevertheless, the converse does not hold, i.e., some BSCs cannot be generated by any epistemic state.*

The reason why the converse of Theorem 3.1 does not hold is that the structure of a BSC is completely independent of the situation γ referred to in the situated KLM postulates. As a very simple instance of this problem, observe that BSCs are not even syntax-independent w.r.t. the situation. That is, we may have $\alpha \mathrel{|\!\sim}_\gamma \beta$ but $\alpha \not\mathrel{|\!\sim}_\delta \beta$, where $\gamma \equiv \delta$. To put it another way, a BSC is simply a rational defeasible consequence relation with the situation playing no role in determining the BSC's structure. To remedy this, we require BSCs to satisfy the following additional postulates:

$$\text{(Inc)}\ \ \frac{\alpha \mathrel{|\!\sim}_\gamma \beta}{\alpha \wedge \gamma \mathrel{|\!\sim}_\top \beta} \qquad \text{(Vac)}\ \ \frac{\top \not\mathrel{|\!\sim}_\top \neg\gamma,\ \alpha \wedge \gamma \mathrel{|\!\sim}_\top \beta}{\alpha \mathrel{|\!\sim}_\gamma \beta}$$

$$\text{(Ext)}\ \ \frac{\gamma \equiv \delta}{\alpha \mathrel{|\!\sim}_\gamma \beta \text{ iff } \alpha \mathrel{|\!\sim}_\delta \beta} \qquad \text{(SupExp)}\ \ \frac{\alpha \mathrel{|\!\sim}_{\gamma\wedge\delta} \beta}{\alpha \wedge \gamma \mathrel{|\!\sim}_\delta \beta}$$

$$\text{(SubExp)}\ \ \frac{\delta \mathrel{|\!\sim}_\top \bot,\ \alpha \wedge \gamma \mathrel{|\!\sim}_\delta \beta}{\alpha \mathrel{|\!\sim}_{\gamma\wedge\delta} \beta}$$

We shall refer to these as the *situated AGM postulates* for reasons to be outlined below.

Definition 3.5 *(Full situated conditional).* A BSC is a **full SC** (FSC) if it satisfies the situated AGM postulates.

One way to interpret the addition of a situation to conditionals, from a technical perspective, is to think of it as similar to *belief revision*. That is, $\alpha \mathrel{|\!\sim}_\gamma \beta$ can be thought of as stating that if a revision with γ has taken place, then β will hold on condition that α holds. With this view of situated conditionals, the situated AGM postulates above are seen as versions of the AGM postulates for belief revision [38]. The names of these postulates were chosen with the names of their AGM analogues in mind. The situated AGM postulates can be motivated intuitively as follows.

Together, Inc and Vac require that when the situation (or revision with) γ is compatible with what is currently plausible, then a conditional w.r.t. the situation γ (a 'revison by' γ) is the same as a conditional where the situation is $\top$ (where there is no 'revision' at all), but with γ added to the premise of the conditional. Ext ensures that BSCs are syntax-independent of the situation. Finally, SupExp and SubExp together require that if the situation δ is implausible (that is, the 'revision' with δ is incompatible with what is plausible), then a conditional w.r.t. the situation $\gamma \wedge \delta$ (a 'revision by' $\gamma \wedge \delta$) is the same as a conditional where the situation (or 'revision') is δ, but with γ added to the premise of the conditional.

It turns out that FSCs are characterised by epistemic interpretations, resulting in the following representation result.

Theorem 3.2. *Every epistemic interpretation generates an FSC. Every FSC can be generated by an epistemic interpretation.*

The AGM-savvy reader may have noticed that the following two obvious analogues of the suite of situated AGM postulates are missing from our list above.

(Succ) $\alpha \mathrel{|\!\sim_\gamma} \gamma$ (Cons) $\top \mathrel{|\!\sim_\gamma} \bot$ iff $\gamma \equiv \bot$

Succ requires situations to matter: a 'revision' by γ will always be successful. Cons states that we obtain an inconsistency only when the situation is inconsistent.

It turns out that Succ holds for epistemic interpretations: it follows from the combination of the situated KLM and AGM postulates. On the other hand, just one direction of Cons holds.

Corollary 3.2. *Every FSC satisfies Succ, but there are FSCs for which Cons does not hold. However, the right-to-left direction of Cons holds: If $\gamma \equiv \bot$, then $\top \mathrel{|\!\sim_\gamma} \bot$.*

Proof. To prove that Succ holds, it suffices, by Theorem 3.1, to show that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\gamma} \gamma$ for all epistemic interpretations $\mathscr{E}$ and all α, γ. To see that this holds, observe that $[\![\alpha \wedge \gamma]\!]_{\mathscr{E}} \subseteq [\![\gamma]\!]_{\mathscr{E}}$.

To prove that Cons does not hold, it suffices, by Theorem 3.1, to show that there is an epistemic interpretation $\mathscr{E}$ such that $\mathscr{E} \Vdash \top \mathrel{|\!\sim_\gamma} \bot$ but $\gamma \not\equiv \bot$. To construct such an $\mathscr{E}$, let $\mathcal{U}^{\mathsf{f}}_{\mathscr{E}} = \mathcal{U}^{\infty}_{\mathscr{E}} = \emptyset$ (and so $\mathscr{E}(u) = \langle \infty, \infty \rangle$ for all $u \in \mathcal{U}$). It is easy to see that by picking any γ s.t. $\gamma \not\equiv \bot$ the result follows.

To prove that if $\gamma \equiv \bot$ then $\top \mathrel{|\!\sim_\gamma} \bot$, note that, by Definition 3.2 and Theorem 3.2, $\top \mathrel{|\!\sim_\gamma} \bot$ iff $\min[\![\top \wedge \gamma]\!]^{\infty}_{\mathscr{E}} \subseteq [\![\bot]\!]$ whenever $\gamma \equiv \bot$, which holds since $\min[\![\top \wedge \gamma]\!]^{\infty}_{\mathscr{E}} = [\![\bot]\!] = \emptyset$. □

We conclude this section by considering the following two properties.

(Incons) $\alpha \mathrel{|\!\sim_\bot} \beta$ (Cond) If $\gamma \not|\!\sim_\top \bot$, then $\alpha \wedge \gamma \mathrel{|\!\sim_\top} \beta$ iff $\alpha \mathrel{|\!\sim_\gamma} \beta$

Incons requires that all conditionals hold when the situation is inconsistent, while Cond requires that conditionals w.r.t. the situation γ be equivalent to the same conditional with γ added to the premise and with a tautologous situation (i.e., the situation is $\top$), provided that γ is not inconsistent w.r.t. the tautologous situation.

Proposition 3.1. *Every FSC satisfies Incons and Cond.*

Proof. To prove that Incons holds, it suffices, by Theorem 3.1, to show that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\bot} \beta$ for all epistemic interpretations $\mathscr{E}$, and all α, β. To see that this holds, observe that $[\![\alpha \wedge \bot]\!]_{\mathscr{E}} = \emptyset$.

To prove that Cond holds, it suffices, by Theorem 3.1, to show that if $\mathscr{E} \nVdash \gamma \mathrel{|\!\sim_\top} \bot$, then $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim_\top} \beta$ iff $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\gamma} \beta$ for all epistemic interpretations $\mathscr{E}$, and all α, β, γ. So, suppose that $\mathscr{E} \nVdash \gamma \mathrel{|\!\sim_\top} \bot$. By Definition 3.2, this means that $\mathcal{U}^{\mathsf{f}}_{\mathscr{E}} \cap [\![\gamma]\!] \neq \emptyset$ and also that $\mathcal{U}^{\mathsf{f}}_{\mathscr{E}} \cap [\![\top]\!] \neq \emptyset$. From this, by Definition 3.2, we need to show that $[\![\alpha \wedge \gamma \wedge \top]\!]^{\mathsf{f}}_{\mathscr{E}} \subseteq [\![\beta]\!]$ iff $[\![\alpha \wedge \gamma]\!]^{\mathsf{f}}_{\mathscr{E}} \subseteq [\![\beta]\!]$ for the result to hold, which follows immediately. □

4. Reasoning with situated conditionals

The previous section provides a framework for characterising the class of full situated conditionals in terms of epistemic interpretations. In this section, we move to an investigation of how we can reason within this framework. More precisely, the question of interest is the following: given a finite set of situated conditionals, or a *situated conditional knowledge base* (SCKB) $\mathcal{KB}$, which situated conditionals can be said to be *entailed* from it? That is, for example, given an SCKB consisting of the conditionals 'birds typically fly' ($\mathsf{b} \mathrel{|\!\sim_\top} \mathsf{f}$), 'penguins are birds' ($\mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim_{\mathsf{p}\wedge\neg\mathsf{b}}} \bot$), 'emperor penguins are penguins' ($\mathsf{ep} \wedge \neg\mathsf{p} \mathrel{|\!\sim_{\mathsf{ep}\wedge\neg\mathsf{p}}} \bot$), and 'penguins typically do not fly' ($\mathsf{p} \mathrel{|\!\sim_{\mathsf{p}}} \neg\mathsf{f}$), should or should we not derive that emperor penguins typically do not fly ($\mathsf{ep} \mathrel{|\!\sim_{\mathsf{ep}}} \neg\mathsf{f}$)? In a non-monotonic framework, it is generally not appropriate to consider entailment relations that are Tarskian in nature, i.e., that determine the consequences of a knowledge base by looking at what holds in *all* the models of the knowledge base. This is because such entailment relations are, by definition, monotonic: let $\mathcal{KB}$ and $\mathcal{KB}'$ be two SCKB's s.t. $\mathcal{KB} \subseteq \mathcal{KB}'$, and let $\alpha \mathrel{|\!\sim_\gamma} \beta$ be a consequence of $\mathcal{KB}$. That is, it is satisfied by every model of $\mathcal{KB}$. Since every model of $\mathcal{KB}'$ is also a model of $\mathcal{KB}$, $\alpha \mathrel{|\!\sim_\gamma} \beta$ is satisfied by every model of $\mathcal{KB}'$ too. That is, $\alpha \mathrel{|\!\sim_\gamma} \beta$ is also a consequence of $\mathcal{KB}'$. Because of monotonicity, if we reason with a Tarskian approach in a conditional setting, we are relegated to very weak inferences. For example, if we only know that 'birds typically fly' ($\mathsf{b} \mathrel{|\!\sim_\top} \mathsf{f}$) and 'robins are birds' ($\mathsf{r} \wedge \neg\mathsf{b} \mathrel{|\!\sim_{\mathsf{r}\wedge\neg\mathsf{b}}} \bot$), we could not even draw a simple conclusion based on property inheritance such as 'robins typically fly' ($\mathsf{r} \mathrel{|\!\sim_{\mathsf{r}}} \mathsf{f}$), since the initial information can also be satisfied by interpretations in which robins are atypical non-flying birds. Hence in a monotonic framework, we cannot, for example, reason under the principle of *presumption of typicality* [25], assuming that everything behaves according to our expectations unless we are explicitly informed that this is not the case.

Since Tarskian inference relations tend to be too weak, inferentially speaking [5], in the framework of non-monotonic reasoning, more suitable entailment relations can be defined by choosing a single model of the knowledge base that satisfies some desirable postulates. It is generally accepted that there is not a unique entailment relation for defeasible reasoning, with different forms of entailment being dependent on the kind of reasoning one wants to model [25,31]. In the framework

of preferential semantics, rational closure, recalled in Section 2, is generally recognised as a core form of entailment, with other suitable forms of entailment being refinements of rational closure.

We now define a form of entailment for situated conditionals, which we call *minimal closure* (MC). It is based on a semantic construction that reformulates in the framework of situated conditionals the semantic construction characterising rational closure for defeasible conditionals: we adapt the notion of a minimal model [29], recalled in Section 2, for our framework, and show that for any SCKB the minimal model is unique.

In the rest of the section, we proceed as follows: we first define the notion of consistency in the present setting. Then, we connect the notions of satisfaction in epistemic states for situated and classic conditionals, respectively. Such a connection, expressed in particular by the content of Corollary 4.1 below, allows us to use known results, regarding entailment relations for defeasible conditionals in ranked interpretations, for the definition of an entailment relation for situated conditionals and epistemic interpretations. In particular, using known results about minimal ranked models [29], we can define our minimal closure on top of the well-known rational closure. Moving from that, in Section 5, we define a decision procedure for entailment that is based on a series of decision steps for classical propositional logic.

Starting from rational closure, which is defined for defeasible conditionals and ranked interpretations, we define minimal closure, defined for situated conditionals and epistemic interpretations, using the connections between the former framework and the latter one. First, we can extend the notion of consistency for defeasible conditionals to situated conditionals. We have seen (Section 2.2) that a set $\mathcal{C}$ of defeasible conditionals is *consistent* if and only if it has a ranked model $\mathscr{R}$ s.t. $[\![0]\!]_{\mathscr{R}} \neq \emptyset$. Such a condition indicates that the agent has a consistent set of expectations since such a model does not satisfy the conditional $\top \mathrel{|\!\sim} \bot$, which captures absurdity in the conditional framework. This condition can easily be translated into our framework.

Definition 4.1 *(SCKB consistency).* An SCKB is **consistent** if it has an epistemic model $\mathscr{E}$ s.t. $[\![\langle \mathtt{f}, 0\rangle]\!]_{\mathscr{E}} \neq \emptyset$.

In other words, an SCKB is consistent if it has an epistemic model $\mathscr{E}$ that does not satisfy $\top \mathrel{|\!\sim_\top} \bot$. $[\![\langle \mathtt{f}, 0\rangle]\!]_{\mathscr{E}}$ is a notation for epistemic interpretations that mirrors the notation $[\![0]\!]_{\mathscr{R}}$ for ranked interpretations, that is, $[\![\langle x, y\rangle]\!]_{\mathscr{E}}$ represents the set of worlds that have rank $\langle x, y\rangle$ in $\mathscr{E}$. On the other hand, given Corollary 3.1, $\top \mathrel{|\!\sim_\top} \bot$ is a situated conditional that has the same meaning as the defeasible conditional $\top \mathrel{|\!\sim} \bot$, that is, an agent believing $\top \mathrel{|\!\sim_\top} \bot$ believes to be in an inconsistent situation since it expects $\bot$ to hold.

Given Corollary 3.1, we can intuitively introduce a notion of satisfaction of defeasible conditionals also for epistemic interpretations:

$$\mathscr{E} \Vdash \alpha \mathrel{|\!\sim} \beta \quad \text{iff} \quad \mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\top} \beta$$

Note that an epistemic interpretation $\mathscr{E}$ satisfies exactly the same defeasible conditionals of its extracted ranked interpretation $\mathscr{R}^{\mathscr{E}}$ (see Definition 3.3). That is, the ranks specified in the interval $\mathcal{U}^{\infty}_{\mathscr{E}} \cup [\![\langle \infty, \infty\rangle]\!]$ are totally irrelevant w.r.t. the satisfaction of the defeasible conditionals of the form $\alpha \mathrel{|\!\sim} \beta$. We can also intuitively define the converse operation w.r.t. the extraction of a ranked interpretation from an epistemic one (Definition 3.3): we can *extract* an epistemic interpretation from a given ranked interpretation. Such an extraction is simply a direct translation of the ranks of the ranked interpretations into the formalism of the epistemic interpretations, simply associating the value $\langle \infty, \infty\rangle$ to all the worlds that have the rank ∞ in the ranked interpretation.

Definition 4.2 *(Extracted epistemic interpretation).* For a ranked interpretation $\mathscr{R}$, we define the **epistemic interpretation** $\mathscr{E}^{\mathscr{R}}$ **extracted from** $\mathscr{R}$ as follows: for $u \in \mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$, $\mathscr{E}^{\mathscr{R}}(u) = \langle \mathtt{f}, i\rangle$, where $\mathscr{R}(u) = i$, and $\mathscr{E}^{\mathscr{R}}(u) = \langle \infty, \infty\rangle$, for $u \in \mathcal{U} \setminus \mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$.

It is easy to see that $\mathscr{R}$ and $\mathscr{E}^{\mathscr{R}}$ are equivalent w.r.t. the satisfaction of defeasible conditionals.

The following corollary of Proposition 3.1, which is simply a semantic reformulation of the postulate Cond, will be central in connecting the satisfaction of situated conditionals to that of defeasible ones.

Corollary 4.1. *For every epistemic interpretation $\mathscr{E}$, if $\mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \cap [\![\gamma]\!] \neq \emptyset$, then $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_\gamma} \beta$ iff $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim} \beta$.*

Proof. Since it is just a semantic reformulation of the postulate Cond, it follows directly from the proof that Cond holds (Proposition 3.1). □

Given Corollary 4.1, we define a simple transformation of situated conditional knowledge bases.

Definition 4.3. Let $\mathcal{KB}$ be an SCKB; with $\mathcal{KB}^{\wedge}$ we denote its *conjunctive classical form*, defined as follows: $\mathcal{KB}^{\wedge} \stackrel{\text{def}}{=} \{\alpha \wedge \gamma \mathrel{|\!\sim} \beta \mid \alpha \mathrel{|\!\sim_\gamma} \beta \in \mathcal{KB}\}$.

We can use the conjunctive classical form to define two models for an SCKB: the *classical epistemic model* and the *minimal epistemic model*. The former will allow us to prove that checking the logical consistency of an SCKB can be reduced to a

consistency check in propositional logic (see Corollary 4.2 below). The latter is the epistemic model, which characterises the MC of an SCKB. Given an SCKB $\mathcal{KB}$, both its classical epistemic model and its minimal epistemic model are defined starting from the minimal ranked model of its conjunctive classical form $\mathcal{KB}^\wedge$ (see Definitions 4.4 and 4.6 below).

Definition 4.4 *(Classical epistemic model).* Let $\mathcal{KB}$ be an SCKB, $\mathcal{KB}^\wedge$ its conjunctive classical form, and $\mathscr{R}$ the minimal ranked model of $\mathcal{KB}^\wedge$. The **classical epistemic model** of $\mathcal{KB}$ is the epistemic interpretation $\mathscr{E}^{\mathscr{R}}$ extracted from $\mathscr{R}$ (see Definition 4.2).

Since $\mathscr{R}$ is a ranked model of $\mathcal{KB}^\wedge$, so is $\mathscr{E}^{\mathscr{R}}$. We need to check whether $\mathscr{E}^{\mathscr{R}}$ is also a model of $\mathcal{KB}$.

Proposition 4.1. *Let $\mathcal{KB}$ be an SCKB, and let $\mathscr{E}^{\mathscr{R}}$ be defined as in Definition 4.4. Then, we have that $\mathscr{E}^{\mathscr{R}}$ is a model of $\mathcal{KB}$.*

Proof. Let $\alpha \mathrel{|\!\sim_\gamma} \beta \in \mathcal{KB}$. Since $\mathscr{E}^{\mathscr{R}}$ is an epistemic model of $\mathcal{KB}^\wedge$ and we have Corollary 4.1, if $[\![\gamma]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{E}^{\mathscr{R}}} \neq \emptyset$, then we conclude $\mathscr{E}^{\mathscr{R}} \Vdash \alpha \mathrel{|\!\sim_\gamma} \beta$. Otherwise, suppose $[\![\gamma]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{E}^{\mathscr{R}}} = \emptyset$. Since $\mathscr{E}^{\mathscr{R}}$ is an extracted epistemic interpretation (Definition 4.2), its only infinite rank is $\langle\infty,\infty\rangle$, and we have $[\![\gamma]\!] \subseteq [\![\langle\infty,\infty\rangle]\!]$, which implies $[\![\alpha \wedge \gamma]\!] \subseteq [\![\langle\infty,\infty\rangle]\!]$, which in turn implies $\mathscr{E}^{\mathscr{R}} \Vdash \alpha \mathrel{|\!\sim_\gamma} \beta$. □

From Proposition 4.1 and Corollary 4.1, we can prove the following result.

Proposition 4.2. *Let $\mathcal{KB}$ be an SCKB. $\mathcal{KB}$ has an epistemic model with $[\![\langle \mathtt{f}, 0\rangle]\!] \neq \emptyset$ iff $\mathcal{KB}^\wedge$ has a ranked model with $[\![0]\!] \neq \emptyset$.*

Proof. Proposition 4.1 and Definitions 4.2 and 4.4 show that if $\mathcal{KB}^\wedge$ has a ranked model with $[\![0]\!] \neq \emptyset$, then $\mathcal{KB}$ has an epistemic model with $[\![\langle \mathtt{f}, 0\rangle]\!] \neq \emptyset$. For the opposite direction, assume that $\mathcal{KB}$ has an epistemic model $\mathscr{E}$ s.t. $[\![\langle \mathtt{f}, 0\rangle]\!]_{\mathscr{E}} \neq \emptyset$. From $\mathscr{E}$, we define an epistemic model $\mathscr{E}_{rk}$ in the following way:

$$\mathscr{E}_{rk}(u) = \begin{cases} \mathscr{E}(u), & \text{if } \mathscr{E}(u) = \langle \mathtt{f}, i\rangle \text{ for some } i; \\ \langle\infty,\infty\rangle, & \text{otherwise.} \end{cases}$$

Clearly $[\![\langle \mathtt{f}, 0\rangle]\!]_{\mathscr{E}_{rk}} \neq \emptyset$. It is easy to check that $\mathscr{E}_{rk}$ is an epistemic model of $\mathcal{KB}$. Moreover, thanks to Corollary 4.1, we can prove that it is also an epistemic model of $\mathcal{KB}^\wedge$: for every $\alpha \mathrel{|\!\sim_\gamma} \beta \in \mathcal{KB}$, if $[\![\gamma]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{E}_{rk}} \neq \emptyset$, then $\mathscr{E}_{rk} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim} \beta$ by Corollary 4.1; if $[\![\gamma]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{E}_{rk}} = \emptyset$, then $[\![\alpha \wedge \gamma]\!] \subseteq [\![\langle\infty,\infty\rangle]\!]$, and we can conclude $\mathscr{E}_{rk} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim} \beta$.

Let $\mathscr{R}$ be the ranked model corresponding to $\mathscr{E}_{rk}$, that is,

$$\mathscr{R}(u) = \begin{cases} i, & \text{if } \mathscr{E}_{rk}(u) = \langle \mathtt{f}, i\rangle \text{ for some } i; \\ \infty, & \text{otherwise.} \end{cases}$$

We have that $[\![\langle \mathtt{f}, 0\rangle]\!]_{\mathscr{E}_{rk}} \neq \emptyset$ implies $[\![0]\!] \neq \emptyset$. Since for every pair of valuations u, v in $\mathcal{U}$, u is preferred to v in $\mathscr{E}_{rk}$ iff u is preferred to v in $\mathscr{R}$, it is easy to see that if $\mathscr{E}_{rk}$ is an epistemic model of $\mathcal{KB}^\wedge$, then $\mathscr{R}$ is a ranked model of $\mathcal{KB}^\wedge$. □

Proposition 4.2 tells us that the consistency of an SCKB $\mathcal{KB}$ corresponds to the consistency of the conditional knowledge base $\mathcal{KB}^\wedge$, the conjunctive form of $\mathcal{KB}$. By linking the satisfaction of an SCKB $\mathcal{KB}$ to the satisfaction of its conjunctive form $\mathcal{KB}^\wedge$, we can define a simple method for checking the consistency of an SCKB, based on the *materialisation* $\overline{\mathcal{KB}^\wedge}$ of $\mathcal{KB}^\wedge$. The materialisation $\overline{\mathcal{C}}$ of a set of defeasible conditionals $\mathcal{C}$ is the set of material implications corresponding to the conditionals in $\mathcal{C}$, defined in the following way:

$$\overline{\mathcal{C}} \stackrel{\text{def}}{=} \{\alpha \rightarrow \beta \mid \alpha \mathrel{|\!\sim} \beta \in \mathcal{C}\}$$

Corollary 4.2. *An SCKB $\mathcal{KB}$ is consistent iff $\overline{\mathcal{KB}^\wedge} \not\models \bot$.*

This corollary is an immediate consequence of Proposition 4.2 and the well-known property that a finite set of defeasible conditionals is consistent if and only if its materialisation is a consistent propositional knowledge base [5, Lemma 5.21].

Example 4.1. Consider an SCKB $\mathcal{KB} = \{\alpha \mathrel{|\!\sim_\alpha} \bot, \top \mathrel{|\!\sim_\top} \alpha \wedge \beta\}$. The meaning of $\alpha \mathrel{|\!\sim_\alpha} \bot$ is that α is necessarily false, while $\top \mathrel{|\!\sim_\top} \alpha \wedge \beta$ indicates that the agent presumes that $\alpha \wedge \beta$ holds (see Example 3.2). Clearly, this is a simple inconsistent knowledge base since it is not rational to consider α as presumably true and necessarily false at the same time. In fact, its only epistemic model is the epistemic model in which all the worlds have rank $\langle\infty,\infty\rangle$.

We can actually check the inconsistency of $\mathcal{KB}$ easily: according to Proposition 4.2, $\mathcal{KB}$ has an epistemic model with $[\![\langle \mathtt{f}, 0\rangle]\!] \neq \emptyset$ iff $\mathcal{KB}^\wedge$ has a ranked model with $[\![0]\!] \neq \emptyset$, and, in turn, $\mathcal{KB}^\wedge$ has a ranked model with $[\![0]\!] \neq \emptyset$ iff $\overline{\mathcal{KB}^\wedge} \not\models \bot$ (Corollary 4.2). $\mathcal{KB}^\wedge = \{\alpha \wedge \alpha \mathrel{|\!\sim} \bot, \top \wedge \top \mathrel{|\!\sim} \alpha \wedge \beta\}$, and $\overline{\mathcal{KB}^\wedge} = \{(\alpha \wedge \alpha) \rightarrow \bot, (\top \wedge \top) \rightarrow \alpha \wedge \beta\}$, which is logically equivalent to the set $\{\neg\alpha, \alpha \wedge \beta\}$, which is clearly inconsistent (that is, $\overline{\mathcal{KB}^\wedge} \models \bot$).

Hence, the classical epistemic model allows reducing SCKB consistency checking to a simple propositional satisfiability checking. This is because it is a direct translation of a ranked interpretation into an equivalent epistemic interpretation. At the same time, since classical epistemic models do not cater for an immediate definition of appropriate forms of entailment (at least in a non-monotonic setting), we now move to the definition of the *minimal epistemic model*, referring to the minimality order introduced for ranked interpretations in Section 2. We need to adapt, in an intuitive way, the notion of minimality defined for ranked interpretations to the present framework. In Section 3, we defined a total ordering $\preceq$ over the tuples $\langle x, y\rangle$ representing the ranks in epistemic interpretations. Let the ordering $\preceq_{\mathcal{KB}}$ on all the epistemic models of an SCKB $\mathcal{KB}$ be defined as follows: $\mathscr{E}_1 \preceq_{\mathcal{KB}} \mathscr{E}_2$, if, for every $v \in \mathcal{U}$, $\mathscr{E}_1(v) \preceq \mathscr{E}_2(v)$. We use $\mathscr{E}_1 \prec_{\mathcal{KB}} \mathscr{E}_2$ to denote its strict counterpart ($\mathscr{E}_1 \preceq_{\mathcal{KB}} \mathscr{E}_2$ and $\mathscr{E}_2 \not\preceq_{\mathcal{KB}} \mathscr{E}_1$).

Definition 4.5 *(Minimal epistemic model).* Let $\mathcal{KB}$ be a consistent SCKB, and $\mathcal{E}_{\mathcal{KB}}$ be the set of its epistemic models. $\mathscr{E} \in \mathcal{E}_{\mathcal{KB}}$ is a **minimal epistemic model of** $\mathcal{KB}$ if there is no $\mathscr{E}' \in \mathcal{E}_{\mathcal{KB}}$ s.t. $\mathscr{E}' \prec_{\mathcal{KB}} \mathscr{E}$.

We first define the construction of a model, given a consistent SCKB $\mathcal{KB}$. Then we prove that it is actually the unique minimal epistemic model of $\mathcal{KB}$ w.r.t. the ordering $\prec_{\mathcal{KB}}$.

Definition 4.6 *(Construction of a minimal epistemic model).* Let $\mathcal{KB}$ be a consistent SCKB, $\mathcal{KB}^{\wedge}$ its conjunctive classical form, and let $\mathscr{R}$ be the minimal ranked model of $\mathcal{KB}^{\wedge}$. We pick out in a set $\mathcal{KB}_{\infty}$ the conditionals in $\mathcal{KB}$ associated with a situation that has infinite rank in $\mathscr{R}$, that is,

- $\mathcal{KB}_{\infty} \stackrel{\text{def}}{=} \{\alpha \mathrel{|\!\sim_{\gamma}} \beta \in \mathcal{KB} \mid \mathscr{R}(\gamma) = \infty\}$.

And from $\mathcal{KB}_{\infty}$ we define the set $\mathcal{KB}^{\wedge}_{\infty\downarrow}$:

- $\mathcal{KB}^{\wedge}_{\infty\downarrow} \stackrel{\text{def}}{=} \{\alpha \wedge \gamma \mathrel{|\!\sim} \beta \mid \alpha \mathrel{|\!\sim_{\gamma}} \beta \in \mathcal{KB}_{\infty}\} \cup \{\mathsf{sent}(\mathcal{U}^{\mathtt{f}}_{\mathscr{R}}) \mathrel{|\!\sim} \bot\}$.

We construct the interpretation $\mathscr{E}_{\mathcal{KB}}$ in the following way:

1. For every $u \in \mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$, if $\mathscr{R}(u) = i$, then $\mathscr{E}_{\mathcal{KB}}(u) = \langle \mathtt{f}, i\rangle$;
2. Let $\mathscr{R}'$ be the minimal ranked model of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$. For every $u \in \mathcal{U}^{\infty}_{\mathscr{R}}$, if $\mathscr{R}'(u) = i$, with $i \in \mathbb{N} \cup \{\infty\}$, then $\mathscr{E}_{\mathcal{KB}}(u) = \langle \infty, i\rangle$.

Definition 4.6 proceeds as follows. First, we want to partition the conditionals that can be considered plausible (that is, the associated situation can be satisfied by valuations with a finite rank) from those that must be considered implausible (that is, the associated situation can be satisfied only by valuations with infinite ranks). This is the set $\mathcal{KB}_{\infty}$. According to Definition 3.2, given an epistemic interpretation, a conditional $\alpha \mathrel{|\!\sim_{\gamma}} \beta$ is evaluated w.r.t. plausible valuations if and only if γ is satisfied by some plausible valuation. It is evaluated w.r.t. implausible valuations otherwise. Given an SCKB $\mathcal{KB}$, γ is not satisfied by any plausible valuation in any model of $\mathcal{KB}$ if and only if $\gamma \mathrel{|\!\sim_{\top}} \bot$ is satisfied by every model of $\mathcal{KB}$, which, by Corollaries 4.1 and 3.1, justifies the use of the minimal ranked model $\mathscr{R}$ of the conjunctive form $\mathcal{KB}^{\wedge}$ for identifying $\mathcal{KB}_{\infty}$. We then specify the minimal configuration satisfying $\mathcal{KB}$, considering first the finite ranks, and then the infinite ones. Corollary 4.1 tells us that, w.r.t. the plausible situations (i.e., finite ranks), the minimal configuration is associated with the conjunctive classical form. Hence, we refer again to the minimal ranked model $\mathscr{R}$ of $\mathcal{KB}^{\wedge}$ to decide the configuration of the plausible valuations (Point 1 in Definition 4.6). We move to configure the infinite ranks, which need to have the minimal configuration satisfying $\mathcal{KB}_{\infty}$, the counterfactual conditionals in our knowledge base. In order to decide such a configuration, we consider $\mathcal{KB}^{\wedge}_{\infty\downarrow}$: all the conditionals in $\mathcal{KB}_{\infty}$, plus the conditional negating the formula characterising all the plausible valuations ($\mathsf{sent}(\mathcal{U}^{\mathtt{f}}_{\mathscr{R}}) \mathrel{|\!\sim} \bot$). The idea behind the use of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$ is as follows. We want to construct a minimal ranking of the counterfactual situations. In order to do that, we need to consider matters from the perspective of being in the counterfactual situations. To do that, we introduce $\mathsf{sent}(\mathcal{U}^{\mathtt{f}}_{\mathscr{R}}) \mathrel{|\!\sim} \bot$. In this way, all the situations that are plausible w.r.t. $\mathcal{KB}$ are now considered impossible, and the plausible situations w.r.t. $\mathcal{KB}^{\wedge}_{\infty\downarrow}$ are the ones that were counterfactual w.r.t. $\mathcal{KB}$. The rank of the implausible valuations in $\mathscr{E}_{\mathcal{KB}}$ is then determined by the rank of the same valuations in the minimal model of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$: $\mathscr{R}'$ defines the minimal configuration satisfying the conditionals in $\mathcal{KB}^{\wedge}_{\infty\downarrow}$, and, at Point 2 in Definition 4.6, we put such a configuration 'on top' of the finite ranks to define $\mathscr{E}_{\mathcal{KB}}$. There is the possibility that the conditional knowledge base $\mathcal{KB}^{\wedge}_{\infty\downarrow}$ is not consistent (see Section 2.2). In such a case, Definition 4.6 still holds: the only model of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$ is the one associating to every valuation the rank ∞, and consequently for every $u \in \mathcal{U}^{\infty}_{\mathscr{R}}$, $\mathscr{E}_{\mathcal{KB}}(u) = \langle \infty, \infty\rangle$.

We need to prove that $\mathscr{E}_{\mathcal{KB}}$ is an epistemic model of $\mathcal{KB}$, and that, moreover, it is the unique minimal epistemic model of $\mathcal{KB}$.

Let $\mathscr{E}$ be an epistemic interpretation. We build an interpretation $\mathscr{E}^{\infty}_{\downarrow}$, the *counterfactual shifting* of $\mathscr{E}$, in the following way. For every valuation u,

$$\mathscr{E}^{\infty}_{\downarrow}(u) \stackrel{\text{def}}{=} \begin{cases} \langle \mathtt{f}, i\rangle, & \text{if } \mathscr{E}(u) = \langle \infty, i\rangle, \text{ with } i < \infty; \\ \langle \infty, \infty\rangle, & \text{otherwise.} \end{cases}$$

∞	$\mathcal{U} \setminus (\llbracket 0 \rrbracket \cup \llbracket 1 \rrbracket \cup \llbracket 2 \rrbracket)$
2	$\mathsf{p}\bar{\mathsf{d}}\mathsf{bf}$
1	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{b}\bar{\mathsf{f}}$, $\mathsf{p}\bar{\mathsf{d}}\mathsf{b}\bar{\mathsf{f}}$,
0	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{bf}$, $\bar{\mathsf{p}}\mathsf{d}\bar{\mathsf{b}}\mathsf{f}$, $\bar{\mathsf{p}}\bar{\mathsf{d}}\bar{\mathsf{b}}\bar{\mathsf{f}}$

Fig. 4. Minimal ranked model of $\mathcal{KB}^{\wedge}$ in Example 4.2.

$\langle \infty, \infty \rangle$	$\llbracket \mathsf{p} \wedge \neg \mathsf{b} \rrbracket \cup \llbracket \mathsf{d} \wedge \neg \mathsf{b} \rrbracket$
$\langle \infty, 1 \rangle$	$\bar{\mathsf{p}}\mathsf{dbf}$, pdbf
$\langle \infty, 0 \rangle$	$\bar{\mathsf{p}}\mathsf{dbf}$, $\mathsf{pdb}\bar{\mathsf{f}}$
$\langle \mathtt{f}, 2 \rangle$	$\mathsf{p}\bar{\mathsf{d}}\mathsf{bf}$
$\langle \mathtt{f}, 1 \rangle$	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{b}\bar{\mathsf{f}}$, $\mathsf{p}\bar{\mathsf{d}}\mathsf{b}\bar{\mathsf{f}}$
$\langle \mathtt{f}, 0 \rangle$	$\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{bf}$, $\bar{\mathsf{p}}\bar{\mathsf{d}}\mathsf{b}\bar{\mathsf{f}}$, $\bar{\mathsf{p}}\bar{\mathsf{d}}\bar{\mathsf{b}}\bar{\mathsf{f}}$

Fig. 5. Minimal epistemic model of the knowledge base in Example 4.2.

Intuitively, $\mathscr{E}^{\infty}_{\downarrow}$ simply shifts the infinite ranks in $\mathscr{E}$ to the finite ranks. For $\mathscr{E}^{\infty}_{\downarrow}$, we can prove a lemma corresponding to Corollary 4.1.

Lemma 4.1. *For every epistemic interpretation $\mathscr{E}$, if $\mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \cap \llbracket \gamma \rrbracket = \emptyset$, then $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim_{\gamma}} \beta$ iff $\mathscr{E}^{\infty}_{\downarrow} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim} \beta$.*

Using Corollary 4.1 and Lemma 4.1, it becomes easy to prove that $\mathscr{E}_{\mathcal{KB}}$ is indeed an epistemic model of $\mathcal{KB}$.

Proposition 4.3. *Let $\mathcal{KB}$ be a consistent SCKB, and let $\mathscr{E}_{\mathcal{KB}}$ be the epistemic interpretation built as in Definition 4.6. Then, $\mathscr{E}_{\mathcal{KB}}$ is an epistemic model of $\mathcal{KB}$.*

Proof. Let $\mathcal{KB}_{\infty}$ be defined as in Definition 4.6. We distinguish two possible cases.

- $\alpha \mathrel{|\!\sim_{\gamma}} \beta \in \mathcal{KB} \setminus \mathcal{KB}_{\infty}$, that is, $\mathscr{E}_{\mathcal{KB}}(\gamma) = \langle \mathtt{f}, i \rangle$, for some i. By the construction of $\mathscr{E}_{\mathcal{KB}}$ (Definition 4.6), $\mathscr{E}_{\mathcal{KB}}$ is an epistemic model of $\mathcal{KB}^{\wedge}$, that is, it is an epistemic model of $\alpha \wedge \gamma \mathrel{|\!\sim} \beta$. From Corollary 4.1, it follows that $\mathscr{E}_{\mathcal{KB}} \Vdash \alpha \mathrel{|\!\sim_{\gamma}} \beta$.
- $\alpha \mathrel{|\!\sim_{\gamma}} \beta \in \mathcal{KB}_{\infty}$, that is, $\mathscr{E}_{\mathcal{KB}}(\gamma) = \langle \infty, i \rangle$, for some i. By the construction of $\mathscr{E}_{\mathcal{KB}}$ (Definition 4.6), $\mathscr{E}_{\mathcal{KB}}$ is an epistemic model of $\mathcal{KB}^{\wedge}$, that is, it is an epistemic model of $\alpha \wedge \gamma \mathrel{|\!\sim} \beta$. Let $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}$ be the counterfactual shifting of $\mathscr{E}_{\mathcal{KB}}$. From Lemma 4.1, we know that, since $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim} \beta$, $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow} \Vdash \alpha \mathrel{|\!\sim_{\gamma}} \beta$ holds. Since $\llbracket \alpha \wedge \gamma \rrbracket_{\mathscr{E}_{\mathcal{KB}}} = \llbracket \alpha \wedge \gamma \rrbracket^{\infty}_{\mathscr{E}_{\mathcal{KB}}} = \llbracket \alpha \wedge \gamma \rrbracket_{\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}}$, for every $u \in \mathcal{U}$, we have $u \in \llbracket \alpha \wedge \gamma \rrbracket_{\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}}$ iff $u \in \llbracket \alpha \wedge \gamma \rrbracket_{\mathscr{E}_{\mathcal{KB}}}$, that is, $\mathscr{E}_{\mathcal{KB}} \Vdash \alpha \mathrel{|\!\sim_{\gamma}} \beta$.

Therefore, for every $\alpha \mathrel{|\!\sim_{\gamma}} \beta \in \mathcal{KB}$, we have $\mathscr{E}_{\mathcal{KB}} \Vdash \alpha \mathrel{|\!\sim_{\gamma}} \beta$, and the result follows. □

We proceed by showing that $\mathscr{E}_{\mathcal{KB}}$ above is actually the only minimal epistemic model of $\mathcal{KB}$.

Proposition 4.4. *Let $\mathcal{KB}$ be a consistent SCKB, and let $\mathscr{E}_{\mathcal{KB}}$ be the epistemic interpretation built as in Definition 4.6. Then $\mathscr{E}_{\mathcal{KB}}$ is the only minimal epistemic model of $\mathcal{KB}$.*

Example 4.2. Assume the SCKB $\mathcal{KB} = \{\mathsf{b} \mathrel{|\!\sim_{\top}} \mathsf{f}, \mathsf{p} \mathrel{|\!\sim_{\mathsf{p}}} \neg\mathsf{f}, \mathsf{d} \mathrel{|\!\sim_{\mathsf{d}}} \neg\mathsf{f}, \mathsf{d} \mathrel{|\!\sim_{\top}} \bot, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim_{\mathsf{p}\wedge\neg\mathsf{b}}} \bot, \mathsf{d} \wedge \neg\mathsf{b} \mathrel{|\!\sim_{\mathsf{d}\wedge\neg\mathsf{b}}} \bot\}$ from Example 3.2. Then we have $\mathcal{KB}^{\wedge} = \{\mathsf{b} \wedge \top \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \wedge \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{d} \wedge \mathsf{d} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{d} \wedge \top \mathrel{|\!\sim} \bot, \mathsf{p} \wedge \neg\mathsf{b} \wedge \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot, \mathsf{d} \wedge \neg\mathsf{b} \wedge \mathsf{d} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$, which is rank equivalent to $\{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{d} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{d} \mathrel{|\!\sim} \bot, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot, \mathsf{d} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$. Fig. 4 depicts the minimal ranked model of $\mathcal{KB}^{\wedge}$. Following Definition 4.6, we have $\mathcal{KB}_{\infty} = \{\mathsf{d} \mathrel{|\!\sim_{\mathsf{d}}} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim_{\mathsf{p}\wedge\neg\mathsf{b}}} \bot, \mathsf{d} \wedge \neg\mathsf{b} \mathrel{|\!\sim_{\mathsf{d}\wedge\neg\mathsf{b}}} \bot\}$. From $\mathcal{KB}_{\infty}$, we get $\mathcal{KB}^{\wedge}_{\infty\downarrow} = \{\mathsf{d} \wedge \mathsf{d} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \wedge \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot, \mathsf{d} \wedge \neg\mathsf{b} \wedge \mathsf{d} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot, (\mathsf{p} \rightarrow \mathsf{b}) \wedge \neg\mathsf{d} \mathrel{|\!\sim} \bot\}$, which is rank equivalent to $\{\mathsf{d} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot, \mathsf{d} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot, (\mathsf{p} \rightarrow \mathsf{b}) \wedge \neg\mathsf{d} \mathrel{|\!\sim} \bot\}$ (note that $(\mathsf{p} \rightarrow \mathsf{b}) \wedge \neg\mathsf{d} \mathrel{|\!\sim} \bot \in \mathcal{KB}^{\wedge}_{\infty\downarrow}$ since $(\mathsf{p} \rightarrow \mathsf{b}) \wedge \neg\mathsf{d} \mathrel{|\!\sim} \bot$ corresponds to the conditional $\mathsf{sent}(\mathcal{U}^{\mathtt{f}}_{\mathscr{E}}) \mathrel{|\!\sim} \bot$, as indicated in Definition 4.6). Following Steps 1 and 2 in Definition 4.6, we construct the minimal epistemic model of the original knowledge base, which is shown in Fig. 5.

The minimal closure of $\mathcal{KB}$ is defined in terms of the minimum epistemic model of $\mathcal{KB}$ constructed in this way.

Definition 4.7 *(Minimal entailment and closure).* $\alpha \mathrel{|\!\sim_{\gamma}} \beta$ is **minimally entailed** by an SCKB $\mathcal{KB}$, denoted as $\mathcal{KB} \models_m \alpha \mathrel{|\!\sim_{\gamma}} \beta$, if $\mathscr{E}_{\mathcal{KB}} \Vdash \alpha \mathrel{|\!\sim_{\gamma}} \beta$, where $\mathscr{E}_{\mathcal{KB}}$ is the minimal model of $\mathcal{KB}$. The corresponding closure operation

$$\mathcal{C}_m(\mathcal{KB}) \stackrel{\text{def}}{=} \{\alpha \mathrel{|\!\sim_{\gamma}} \beta \mid \mathcal{KB} \models_m \alpha \mathrel{|\!\sim_{\gamma}} \beta\}$$

is the **minimal closure** of $\mathcal{KB}$.

∞	$\bar{\mathsf{b}}\bar{\mathsf{f}}\mathsf{p}$, $\bar{\mathsf{b}}\mathsf{f}\mathsf{p}$
2	$\mathsf{b}\mathsf{f}\mathsf{p}$
1	$\mathsf{b}\bar{\mathsf{f}}\bar{\mathsf{p}}$, $\mathsf{b}\bar{\mathsf{f}}\mathsf{p}$
0	$\bar{\mathsf{b}}\bar{\mathsf{f}}\bar{\mathsf{p}}$, $\bar{\mathsf{b}}\mathsf{f}\bar{\mathsf{p}}$, $\mathsf{b}\mathsf{f}\bar{\mathsf{p}}$

Fig. 6. Minimal ranked model of the knowledge base $\mathcal{C} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$.

Example 4.3. We proceed from Example 4.2. Looking at the model in Fig. 5, we are able to check what is minimally entailed. For every $\alpha \mathrel{|\!\sim}_\gamma \beta \in \mathcal{KB}$, $\mathcal{KB} \models_m \alpha \mathrel{|\!\sim}_\gamma \beta$. In particular, while $\mathcal{KB} \models_m \mathsf{d} \mathrel{|\!\sim}_\top \bot$, we do not have $\mathcal{KB} \models_m \mathsf{d} \mathrel{|\!\sim}_\mathsf{d} \bot$, that is, it is possible to reason counterfactually about dodos. From the point of view of the actual situation (that is, in the situation $\top$), we can conclude anything about dodos, since they do not exist. Indeed, we have both $\mathcal{KB} \models_m \mathsf{d} \mathrel{|\!\sim}_\top \neg\mathsf{f}$ and $\mathcal{KB} \models_m \mathsf{d} \mathrel{|\!\sim}_\top \mathsf{f}$. Nevertheless, we are able to reason coherently about dodos once we assume a point of view in which they would exist. To witness, we have $\mathcal{KB} \models_m \mathsf{d} \mathrel{|\!\sim}_\mathsf{d} \neg\mathsf{f}$, but $\mathcal{KB} \not\models_m \mathsf{d} \mathrel{|\!\sim}_\mathsf{d} \mathsf{f}$.

Definition 4.6 shows that the minimal epistemic model can be defined using the minimal ranked models for two sets of defeasible conditionals, $\mathcal{KB}^\wedge$ and $\mathcal{KB}^\wedge_{\infty\downarrow}$. If a valuation is associated with a finite rank i in the minimal ranked model of $\mathcal{KB}^\wedge$, then we associate to it the corresponding rank $\langle \mathtt{f}, i\rangle$ in the minimal epistemic model. All the other valuations, those that have rank ∞ in the minimal model of $\mathcal{KB}^\wedge$, will have a rank determined by the minimal ranked model of $\mathcal{KB}^\wedge_{\infty\downarrow}$: for each one of such valuations, if its rank in the minimal model of $\mathcal{KB}^\wedge_{\infty\downarrow}$ is i ($i \in \mathbb{N} \cup \{\infty\}$), it will have the rank $\langle \infty, i\rangle$ in the minimal epistemic model.

As we are going to see in the next section, since the construction of the minimal epistemic model relies on the construction of two minimal ranked models, it is possible to decide whether an SC is in the minimal entailment of an SCKB fully relying on a series of propositional decision steps.

5. Computing entailment from situated conditional knowledge bases

In this section, we define a procedure to decide whether a conditional is in the minimal closure of an SCKB. The procedure is described by Algorithm 6, `MinimalClosure`, and it relies on a series of propositional entailment checks. Hence, it can be implemented on top of any propositional reasoner.

We will start by looking at Algorithms `Exceptional` (1), `ComputeRanking` (2), `Rank` (3), and `RationalClosure` (4), which formalise known procedures (see the work of Freund [39] and of Casini and Straccia [40, Section 2]) that together define a decision procedure for rational closure (RC). As indicated in Section 2, on the semantic side, the RC of a knowledge base $\mathcal{C}$ containing defeasible conditionals can be characterised using the *minimal ranked model* $\mathscr{R}^\mathcal{C}_{RC}$ [29], that is, $\alpha \mathrel{|\!\sim} \beta$ is in the RC of a set of defeasible conditionals $\mathcal{C}$ if and only if $\mathscr{R}^\mathcal{C}_{RC} \Vdash \alpha \mathrel{|\!\sim} \beta$ (Definition 2.2).

It has been proved [39,40] that $\alpha \mathrel{|\!\sim} \beta$ is in the RC of $\mathcal{C}$, that is, $\mathscr{R}^\mathcal{C}_{RC} \Vdash \alpha \mathrel{|\!\sim} \beta$, if and only if $\mathtt{RationalClosure}(\mathcal{C}, \alpha \mathrel{|\!\sim} \beta)$ returns `true`. In what follows, we provide an explanation of all the algorithms involved in the process. We shall often refer to Fig. 1 (repeated in Fig. 6 for the reader's convenience), which is the minimal ranked model of the knowledge base $\mathcal{C} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$.

- $\mathtt{Exceptional}(\mathcal{C})$ (Algorithm 1) takes as input a finite set $\mathcal{C}$ of defeasible conditionals and gives back the exceptional elements, that is, the conditionals $\alpha \mathrel{|\!\sim} \beta$ s.t. $\top \mathrel{|\!\sim} \neg\alpha$ holds in the minimal ranked model of $\mathcal{C}$. For example, from Fig. 6, one can check that the conditionals $\mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}$ and $\mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot$ are exceptional, since none of the valuations in layer 0 satisfies p, and in fact $\mathtt{Exceptional}(\mathcal{C}) = \{\mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$. The procedure fully relies on a series of decision steps in classical propositional logic, since it uses the *materialisation* of the KB $\mathcal{C}$ (see Section 4).
- $\mathtt{ComputeRanking}(\mathcal{C})$ (Algorithm 2) ranks each conditional in the KB $\mathcal{C}$ w.r.t. its exceptionality level. $\mathcal{E}_0$ contains all the conditionals, $\mathcal{E}_1$ the exceptional ones w.r.t. $\mathcal{E}_0$, and so on. $\mathcal{E}_\infty$ contains the fixed point of the exceptionality procedure, that is, the conditionals having antecedents that cannot be satisfied in any valuation that is ranked as finite in any ranked model of $\mathcal{C}$. $\mathtt{ComputeRanking}(\mathcal{C})$ returns $\mathcal{E}_0 = \mathcal{C} = \{\mathsf{b} \mathrel{|\!\sim} \mathsf{f}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$, $\mathcal{E}_1 = \{\mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$, $\mathcal{E}_\infty = \{\mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$.
- $\mathtt{Rank}(\mathcal{C}, \alpha)$ (Algorithm 3) decides the rank of a proposition, that is, the lowest rank in the minimal ranked model containing a valuation that satisfies the proposition. For example, the reader can check that $\mathtt{Rank}(\mathcal{C}, \neg\mathsf{p}) = 0$, $\mathtt{Rank}(\mathcal{C}, \mathsf{p}) = 1$, $\mathtt{Rank}(\mathcal{C}, \mathsf{p} \wedge \mathsf{f}) = 2$, $\mathtt{Rank}(\mathcal{C}, \mathsf{p} \wedge \neg\mathsf{b}) = \infty$, values that, for each of the propositions, correspond exactly to the lowest layer in the minimal ranked model in which there is a valuation satisfying the proposition (see Fig. 6).
- $\mathtt{RationalClosure}(\mathcal{C}, \alpha \mathrel{|\!\sim} \beta)$ (Algorithm 4) tells us whether $\alpha \mathrel{|\!\sim} \beta$ is in the RC of $\mathcal{C}$, that is, whether $\mathscr{R}^\mathcal{C}_{RC} \Vdash \alpha \mathrel{|\!\sim} \beta$. For example, $\mathtt{RationalClosure}(\mathcal{C}, \mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f})$ is `true`, since: $\mathtt{Rank}(\mathcal{C}, \mathsf{p}) = 1$, $\mathcal{E}_1 = \{\mathsf{p} \mathrel{|\!\sim} \neg\mathsf{f}, \mathsf{p} \wedge \neg\mathsf{b} \mathrel{|\!\sim} \bot\}$, and $\overline{\mathcal{E}_1} \cup \{\mathsf{p}\} \models \neg\mathsf{f}$.

Note that all the procedures fully rely on a series of decision steps in classical propositional logic.

Algorithm 1: `Exceptional`$(\mathcal{C})$.

input : a set of defeasible conditionals $\mathcal{C}$
output: $\mathcal{E} \subseteq \mathcal{C}$ s.t. $\mathcal{E}$ is exceptional w.r.t. $\mathcal{C}$

$\mathcal{E} := \emptyset$
$\overline{\mathcal{C}} := \{\alpha \rightarrow \beta \mid \alpha \mathrel{|\!\sim} \beta \in \mathcal{C}\}$
foreach $\alpha \mathrel{|\!\sim} \beta \in \mathcal{C}$ **do**
  if $\overline{\mathcal{C}} \models \neg\alpha$ **then**
    $\mathcal{E} := \mathcal{E} \cup \{\alpha \mathrel{|\!\sim} \beta\}$
  end
end
return $\mathcal{E}$

Algorithm 2: `ComputeRanking`$(\mathcal{C})$.

input : a set of defeasible conditionals $\mathcal{C}$
output: an exceptionality ranking $r_{\mathcal{C}}$

$i := 0$
$\mathcal{E}_0 := \mathcal{C}$
$\mathcal{E}_1 :=$ `Exceptional`$(\mathcal{E}_0)$
while $\mathcal{E}_{i+1} \neq \mathcal{E}_i$ **do**
  $i := i + 1$
  $\mathcal{E}_{i+1} :=$ `Exceptional`$(\mathcal{E}_i)$
end
$\mathcal{E}_\infty := \mathcal{E}_i$
$r_{\mathcal{C}} := (\mathcal{E}_0, \ldots, \mathcal{E}_{i-1}, \mathcal{E}_\infty)$
return $r_{\mathcal{C}}$

Algorithm 3: `Rank`$(\mathcal{C}, \alpha)$.

input : a set of defeasible conditionals $\mathcal{C}$, a proposition α
output: the rank $rk_{\mathcal{C}}(\alpha)$ of α

$r_{\mathcal{C}} = (\mathcal{E}_0, \ldots, \mathcal{E}_n, \mathcal{E}_\infty) :=$ `ComputeRanking`$(\mathcal{C})$
foreach $0 \leq i \leq n$ **do**
  $\overline{\mathcal{E}_i} := \{\alpha \rightarrow \beta \mid \alpha \mathrel{|\!\sim} \beta \in \mathcal{E}_i\}$
end
$\overline{\mathcal{E}_\infty} := \{\alpha \rightarrow \beta \mid \alpha \mathrel{|\!\sim} \beta \in \mathcal{E}_\infty\}$
$i := 0$
while $\overline{\mathcal{E}_i} \models \neg\alpha$ *and* $i \leq n$ **do**
  $i := i + 1$
end
if $i \leq n$ **then**
  $rk_{\mathcal{C}}(\alpha) := i$
end
else
  if $\overline{\mathcal{E}_\infty} \not\models \neg\alpha$ **then**
    $rk_{\mathcal{C}}(\alpha) := i + 1$
  end
  else
    $rk_{\mathcal{C}}(\alpha) := \infty$
  end
end
return $rk_{\mathcal{C}}(\alpha)$

Algorithm 4: `RationalClosure`$(\mathcal{C}, \alpha \mathrel{|\!\sim} \beta)$.

input : a set of defeasible conditionals $\mathcal{C}$, a query $\alpha \mathrel{|\!\sim} \beta$
output: `true`, if $\mathcal{C} \models_{RC} \alpha \mathrel{|\!\sim} \beta$, `false` otherwise

$r_{\mathcal{KB}} = (\mathcal{E}_0, \ldots, \mathcal{E}_n, \mathcal{E}_\infty) :=$ `ComputeRanking`$(\mathcal{C})$
$r :=$ `Rank`$(\mathcal{C}, \alpha)$
return $\overline{\mathcal{E}_r} \cup \{\alpha\} \models \beta$

Algorithms `Partition` (5) and `MinimalClosure` (6) are novel. They define a procedure to decide minimal entailment $\models_m$, given an SCKB, and they are built on top of `ComputeRanking`, `Rank`, and `RationalClosure`. Let us go through them:

Algorithm 5: `Partition`($\mathcal{KB}$).

input : an SCKB $\mathcal{KB}$
output: the conjunctive forms $\mathcal{KB}^{\wedge}$ and $\mathcal{KB}^{\wedge}_{\infty\downarrow}$

$\mathcal{KB}^{\wedge} := \{\alpha \wedge \gamma \mathrel{|\sim} \beta \mid \alpha \mathrel{|\sim}_{\gamma} \beta \in \mathcal{KB}\}$
$r_{\mathcal{KB}^{\wedge}} = (\mathcal{E}_0, \ldots, \mathcal{E}_n, \mathcal{E}_{\infty}) := \texttt{ComputeRanking}(\mathcal{KB}^{\wedge})$
$\mathcal{KB}_{\infty} := \emptyset$
foreach $\alpha \mathrel{|\sim}_{\gamma} \beta \in \mathcal{KB}$ **do**
  if $\texttt{Rank}(\mathcal{KB}^{\wedge}, \gamma) = \infty$ **then**
    $\mathcal{KB}_{\infty} := \mathcal{KB}_{\infty} \cup \{\alpha \mathrel{|\sim}_{\gamma} \beta\}$
  end
end
$\mu := \bigwedge\{\neg\alpha \mid \alpha \mathrel{|\sim} \beta \in \mathcal{E}_{\infty}\}$
$\mathcal{KB}^{\wedge}_{\infty\downarrow} := \{\alpha \wedge \gamma \mathrel{|\sim} \beta \mid \alpha \mathrel{|\sim}_{\gamma} \beta \in \mathcal{KB}_{\infty}\} \cup \{\mu \mathrel{|\sim} \bot\}$
return $\mathcal{KB}^{\wedge}, \mathcal{KB}^{\wedge}_{\infty\downarrow}$

Algorithm 6: `MinimalClosure`($\mathcal{KB}, \alpha \mathrel{|\sim}_{\gamma} \beta$).

input : an SCKB $\mathcal{KB}$, a query $\alpha \mathrel{|\sim}_{\gamma} \beta$
output: `true`, if $\mathcal{KB} \models_m \alpha \mathrel{|\sim}_{\gamma} \beta$, `false` otherwise

$\mathcal{KB}^{\wedge}, \mathcal{KB}^{\wedge}_{\infty\downarrow} := \texttt{Partition}(\mathcal{KB})$ $\overline{\mathcal{KB}^{\wedge}} = \{(\alpha \wedge \gamma) \rightarrow \beta \mid \alpha \wedge \gamma \mathrel{|\sim} \beta \in \mathcal{KB}^{\wedge}\}$
if $\overline{\mathcal{KB}^{\wedge}} \models \bot$ **then**
  return `true`
end
else
  if $\texttt{Rank}(\mathcal{KB}^{\wedge}, \gamma) < \infty$ **then**
    return $\texttt{RationalClosure}(\mathcal{KB}^{\wedge}, \alpha \wedge \gamma \mathrel{|\sim} \beta)$
  end
  else
    return $\texttt{RationalClosure}(\mathcal{KB}^{\wedge}_{\infty\downarrow}, \alpha \wedge \gamma \mathrel{|\sim} \beta)$
  end
end

- `Partition`($\mathcal{KB}$) (Algorithm 5) takes as input an SCKB $\mathcal{KB}$ and identifies the set $\mathcal{KB}_{\infty}$ and the set of defeasible conditionals $\mathcal{KB}^{\wedge}_{\infty\downarrow}$, in a way that, as we shall prove, corresponds to Definition 4.6. That is, $\mathcal{KB}_{\infty}$ is the set of conditionals of which the situations are ranked as infinite w.r.t. $\mathcal{KB}^{\wedge}$.
- `MinimalClosure`($\mathcal{KB}, \alpha \mathrel{|\sim}_{\gamma} \beta$) (Algorithm 6) tells us whether $\alpha \mathrel{|\sim}_{\gamma} \beta$ is in the minimal closure of $\mathcal{KB}$. First, the algorithm checks if $\mathcal{KB}$ is a consistent SCKB (see Definition 4.1): by Corollary 4.2, it is sufficient to check whether $\overline{\mathcal{KB}^{\wedge}} \models \bot$. Then, in case it is consistent, it checks the rank of the situation γ. If the situation's rank is finite, then it checks whether the conjunctive form $\alpha \wedge \gamma \mathrel{|\sim} \beta$ is in the RC of $\mathcal{KB}^{\wedge}$. Otherwise, it checks whether the conjunctive form $\alpha \wedge \gamma \mathrel{|\sim} \beta$ is in the RC of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$.

We need to prove that Algorithm 6 is complete and correct w.r.t. minimal entailment $\models_m$. Before the main theorem, we need to prove the following lemma.

Lemma 5.1. *Let $\mathcal{KB}$ be a consistent SCKB, let $\mathcal{KB}^{\wedge}$ be its conjunctive classical form, and let $\mathscr{R}$ be the minimal ranked model of $\mathcal{KB}^{\wedge}$. Moreover, let μ be defined as in Algorithm 5, and let $\mathsf{sent}(\mathcal{U}^{\mathfrak{f}}_{\mathscr{R}})$ be as in Definition 4.6. Then we have that μ is logically equivalent to $\mathsf{sent}(\mathcal{U}^{\mathfrak{f}}_{\mathscr{R}})$.*

Proof. First, we prove that $\mathsf{sent}(\mathcal{U}^{\mathfrak{f}}_{\mathscr{R}}) \models \mu$. Let $\alpha \mathrel{|\sim} \beta \in \mathcal{E}_{\infty}$. This implies that $rk_{\mathcal{KB}^{\wedge}}(\alpha) = \infty$, that is, all the valuations satisfying α have rank ∞. That is, $\mathcal{U}^{\mathfrak{f}}_{\mathscr{R}} \subseteq [\![\neg\alpha]\!]$ for every α s.t. $\alpha \mathrel{|\sim} \beta \in \mathcal{E}_{\infty}$. That implies

$$\mathcal{U}^{\mathfrak{f}}_{\mathscr{R}} \subseteq \bigcap\{[\![\neg\alpha]\!]_{\mathscr{R}} \mid \alpha \mathrel{|\sim} \beta \in \mathcal{E}_{\infty}\},$$

and, consequently, $\mathsf{sent}(\mathcal{U}^{\mathfrak{f}}_{\mathscr{R}}) \models \mu$.

Now we prove that $\mu \models \mathsf{sent}(\mathcal{U}^{\mathfrak{f}}_{\mathscr{R}})$. Assume this is not the case. That is, there is a valuation $w \in \mathcal{U}^{\infty}_{\mathscr{R}}$ s.t. $w \Vdash \mu$. Let n be the highest finite rank in $\mathscr{R}$, and consider the ranked model $\mathscr{R}'$ obtained from $\mathscr{R}$ just by re-assigning the valuation w from the rank ∞ to the rank $n+1$ (note that if the valuation w is the only valuation in $\mathcal{U}^{\infty}_{\mathscr{R}}$ and, consequently, $\mathcal{U}^{\infty}_{\mathscr{R}'} = \emptyset$, then $\mathscr{R}'$ is still a ranked interpretation since it is compatible with Definition 2.1). $\mathscr{R}'$ is preferred to $\mathscr{R}$, and it is easy to see that $\mathscr{R}'$ is a ranked model of $\mathcal{KB}$: for every $\alpha \mathrel{|\sim} \beta \in \mathcal{E}_i$, for some $i < \infty$, there is a valuation in a lower rank satisfying $\alpha \wedge \beta$, while for every $\alpha \mathrel{|\sim} \beta \in \mathcal{E}_{\infty}$, $w \Vdash \neg\alpha$, and consequently w is irrelevant w.r.t. the satisfaction of $\alpha \mathrel{|\sim} \beta$ by $\mathscr{R}'$, since it is not in $\min[\![\alpha]\!]^{\mathfrak{f}}_{\mathscr{R}'}$. Hence, we have that $\mathscr{R}' \prec_{\mathcal{KB}} \mathscr{R}$, against the hypothesis that $\mathscr{R}$ is the minimal element in $\prec_{\mathcal{KB}}$, which leads to a contradiction. Therefore, $\mu \models \mathsf{sent}(\mathcal{U}^{\mathfrak{f}}_{\mathscr{R}})$. □

Now we can state the main result of the present section.

Theorem 5.1. *Let $\mathcal{KB}$ be an SCKB.* `MinimalClosure`$(\mathcal{KB}, \alpha \mathrel{|\!\sim_\gamma} \beta)$ *returns* `true` *iff* $\mathcal{KB} \models_m \alpha \mathrel{|\!\sim_\gamma} \beta$.

Example 5.1. Let us model a more practically-oriented scenario. The agent knows that the Kitchen has been cleaned ($\neg\mathsf{ck} \mathrel{|\!\sim_\top} \bot$), and has a series of (defeasible) expectations: the pan is clean (cl) and positioned in Cupboard1 (cb1) ($\top \mathrel{|\!\sim_\top} \mathsf{cl}$ and $\top \mathrel{|\!\sim_\top} \mathsf{cb1}$), but in case the pan is in Cupboard2 (cb2), the agent will need a stool (st) to reach the pan ($\mathsf{cb2} \mathrel{|\!\sim_\top} \mathsf{st}$). We can also model the agent's expectations about counterfactual situations, that is, situations that are not compatible with the information the agent has about the actual situation: if the kitchen has not been cleaned, the pan will presumably be in the sink ($\top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \mathsf{si}$), and it will be dirty ($\top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \neg\mathsf{cl}$). Also, we have some constraints that *must necessarily* hold, simply stating that the pan must be in exactly one place: $\neg\mathsf{cb1} \wedge \neg\mathsf{cb2} \wedge \neg\mathsf{si} \mathrel{|\!\sim_{\neg\mathsf{cb1}\wedge\neg\mathsf{cb2}\wedge\neg\mathsf{si}}} \bot$, $\mathsf{cb1} \wedge \mathsf{cb2} \mathrel{|\!\sim_{\mathsf{cb1}\wedge\mathsf{cb2}}} \bot$, $\mathsf{cb1} \wedge \mathsf{si} \mathrel{|\!\sim_{\mathsf{cb1}\wedge\mathsf{si}}} \bot$, $\mathsf{cb2} \wedge \mathsf{si} \mathrel{|\!\sim_{\mathsf{cb2}\wedge\mathsf{si}}} \bot$. Note that the conditionals $\alpha \mathrel{|\!\sim_\alpha} \bot$ impose that the valuations satisfying α can be placed only in rank $\langle\infty, \infty\rangle$, that is, $\neg\alpha$ cannot be falsified, even in the counterfactual situations (see Example 3.2).

Let $\mathcal{KB} = \{\neg\mathsf{cb1} \wedge \neg\mathsf{cb2} \wedge \neg\mathsf{si} \mathrel{|\!\sim_{\neg\mathsf{cb1}\wedge\neg\mathsf{cb2}\wedge\neg\mathsf{si}}} \bot, \mathsf{cb1} \wedge \mathsf{cb2} \mathrel{|\!\sim_{\mathsf{cb1}\wedge\mathsf{cb2}}} \bot, \mathsf{cb1} \wedge \mathsf{si} \mathrel{|\!\sim_{\mathsf{cb1}\wedge\mathsf{si}}} \bot, \mathsf{cb2} \wedge \mathsf{si} \mathrel{|\!\sim_{\mathsf{cb2}\wedge\mathsf{si}}} \bot, \neg\mathsf{ck} \mathrel{|\!\sim_\top} \bot, \top \mathrel{|\!\sim_\top} \mathsf{cl}, \top \mathrel{|\!\sim_\top} \mathsf{cb1}, \mathsf{cb2} \mathrel{|\!\sim_\top} \mathsf{st}, \top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \mathsf{si}, \top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \neg\mathsf{cl}\}$ be an SCKB formalising the scenario in Example 5.1. We apply Algorithm 5, `Partition`, to $\mathcal{KB}$:

- The algorithm creates the conjunctive form $\mathcal{KB}^\wedge = \{\neg\mathsf{cb1} \wedge \neg\mathsf{cb2} \wedge \neg\mathsf{si} \mathrel{|\!\sim} \bot, \mathsf{cb1} \wedge \mathsf{cb2} \mathrel{|\!\sim} \bot, \mathsf{cb1} \wedge \mathsf{si} \mathrel{|\!\sim} \bot, \mathsf{cb2} \wedge \mathsf{si} \mathrel{|\!\sim} \bot, \neg\mathsf{ck} \mathrel{|\!\sim} \bot, \top \mathrel{|\!\sim} \mathsf{cl}, \top \mathrel{|\!\sim} \mathsf{cb1}, \mathsf{cb2} \mathrel{|\!\sim} \mathsf{st}, \neg\mathsf{ck} \mathrel{|\!\sim} \mathsf{si}, \neg\mathsf{ck} \mathrel{|\!\sim} \neg\mathsf{cl}\}$ (we have simplified the formulas in the conditionals w.r.t. the definition of $\mathcal{KB}^\wedge$ in Section 4, for example substituting formulas $\alpha \wedge \alpha$ or $\alpha \wedge \top$ with α).
- Calling algorithm `ComputeRanking`, we rank $\mathcal{KB}^\wedge$ in $\mathcal{E}_0 = \{\top \mathrel{|\!\sim} \mathsf{cl}, \top \mathrel{|\!\sim} \mathsf{cb1}\} \cup \mathcal{E}_1$, $\mathcal{E}_1 = \{\mathsf{cb2} \mathrel{|\!\sim} \mathsf{st}\} \cup \mathcal{E}_\infty$, $\mathcal{E}_\infty = \{\neg\mathsf{cb1} \wedge \neg\mathsf{cb2} \wedge \neg\mathsf{si} \mathrel{|\!\sim} \bot, \mathsf{cb1} \wedge \mathsf{cb2} \mathrel{|\!\sim} \bot, \mathsf{cb1} \wedge \mathsf{si} \mathrel{|\!\sim} \bot, \mathsf{cb2} \wedge \mathsf{si} \mathrel{|\!\sim} \bot, \neg\mathsf{ck} \mathrel{|\!\sim} \bot, \neg\mathsf{ck} \mathrel{|\!\sim} \mathsf{si}, \neg\mathsf{ck} \mathrel{|\!\sim} \neg\mathsf{cl}\}$.
- We then call the procedure `Rank`$(\mathcal{KB}^\wedge, \gamma)$ for every formula γ appearing in some conditional $\alpha \mathrel{|\!\sim} \beta$ in $\mathcal{KB}$. It turns out that `Rank`$(\mathcal{KB}^\wedge, \gamma) = \infty$ for $\gamma \in \{\neg\mathsf{cb1} \wedge \neg\mathsf{cb2} \wedge \neg\mathsf{si}, \mathsf{cb1} \wedge \mathsf{cb2}, \mathsf{cb1} \wedge \mathsf{si}, \mathsf{cb2} \wedge \mathsf{si}, \neg\mathsf{ck}\}$. Consequently, $\mathcal{KB}_\infty = \{\neg\mathsf{cb1} \wedge \neg\mathsf{cb2} \wedge \neg\mathsf{si} \mathrel{|\!\sim_{\neg\mathsf{cb1}\wedge\neg\mathsf{cb2}\wedge\neg\mathsf{si}}} \bot, \mathsf{cb1} \wedge \mathsf{cb2} \mathrel{|\!\sim_{\mathsf{cb1}\wedge\mathsf{cb2}}} \bot, \mathsf{cb1} \wedge \mathsf{si} \mathrel{|\!\sim_{\mathsf{cb1}\wedge\mathsf{si}}} \bot, \mathsf{cb2} \wedge \mathsf{si} \mathrel{|\!\sim_{\mathsf{cb2}\wedge\mathsf{si}}} \bot, \top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \mathsf{si}, \top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \neg\mathsf{cl}\}$.
- Eventually, the algorithm constructs the set $\mathcal{KB}^\wedge_{\infty\downarrow}$: first, from $\mathcal{E}_\infty$, we define μ as $\bigwedge\{\mathsf{cb1} \vee \mathsf{cb2} \vee \mathsf{si}, \neg\mathsf{cb1} \vee \neg\mathsf{cb2}, \neg\mathsf{cb1} \vee \neg\mathsf{si}, \neg\mathsf{cb2} \vee \neg\mathsf{si}, \mathsf{ck}\}$; then we set $\mathcal{KB}^\wedge_{\infty\downarrow}$ as $\{\neg\mathsf{cb1} \wedge \neg\mathsf{cb2} \wedge \neg\mathsf{si} \mathrel{|\!\sim} \bot, \mathsf{cb1} \wedge \mathsf{cb2} \mathrel{|\!\sim} \bot, \mathsf{cb1} \wedge \mathsf{si} \mathrel{|\!\sim} \bot, \mathsf{cb2} \wedge \mathsf{si} \mathrel{|\!\sim} \bot, \neg\mathsf{ck} \mathrel{|\!\sim} \mathsf{si}, \neg\mathsf{ck} \mathrel{|\!\sim} \neg\mathsf{cl}, \mu \mathrel{|\!\sim} \bot\}$.

Once we have $\mathcal{KB}^\wedge$ and $\mathcal{KB}^\wedge_{\infty\downarrow}$, we can give queries to Algorithm 6 (`MinimalClosure`). For example, we can check whether the agent should expect the pan to be in the sink ($\top \mathrel{|\!\sim_\top} \mathsf{si}$).

- Given $\mathcal{KB}^\wedge$, we define its materialisation $\overline{\mathcal{KB}^\wedge}$, which contains the implications $(\alpha \wedge \gamma) \rightarrow \beta$ corresponding to the conditionals $\alpha \wedge \gamma \mathrel{|\!\sim} \beta$ in $\mathcal{KB}^\wedge$. Using $\overline{\mathcal{KB}^\wedge}$, the algorithm checks whether the knowledge base $\mathcal{KB}$ is inconsistent by checking whether $\overline{\mathcal{KB}^\wedge} \models \bot$ (the reader can check that it is not the case.)
- We then have to check the rank of the situation $\top$ in $\top \mathrel{|\!\sim_\top} \mathsf{si}$, which, being $\top$, must be 0. Hence, semantically, since $\top$ cannot be an exceptional proposition, $\top \mathrel{|\!\sim_\top} \mathsf{si}$ is a conditional whose satisfaction needs to be checked w.r.t. the valuations in the finite ranks of the minimal epistemic model of $\mathcal{KB}$, in particular, w.r.t. the valuations in the rank $\langle \mathtt{f}, 0\rangle$. This corresponds to checking in Algorithm `MinimalClosure` whether $\top \mathrel{|\!\sim} \mathsf{si}$ is in the rational closure of $\mathcal{KB}^\wedge$. That is, whether `RationalClosure`$(\mathcal{KB}^\wedge, \top \mathrel{|\!\sim} \mathsf{si})$ returns `true`.
  In the procedure `RationalClosure`$(\mathcal{KB}^\wedge, \top \mathrel{|\!\sim} \mathsf{si})$, the rank 0 is associated to $\top$, and $\mathcal{E}_0 = \mathcal{KB}^\wedge$. Consequently, $\top \mathrel{|\!\sim} \mathsf{si}$ is in the rational closure of $\mathcal{KB}^\wedge$ if and only if $\overline{\mathcal{E}_0} \models \mathsf{si}$, which is not the case. Actually, we have that $\top \mathrel{|\!\sim_\top} \neg\mathsf{si}$ is in the minimal closure of $\mathcal{KB}$, since, due to the presence of $\top \rightarrow \mathsf{cb1}$ and $(\mathsf{cb1} \wedge \mathsf{si}) \rightarrow \bot$ in $\overline{\mathcal{E}_0}$, we have $\overline{\mathcal{E}_0} \models \neg\mathsf{si}$.

We now consider a counterfactual situation, checking whether the agent believes that, in case the kitchen has not been cleaned, the pan is not in Cupboard2 ($\top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \neg\mathsf{cb2}$).

- As for the previous query, the algorithm starts by checking whether $\mathcal{KB}$ is consistent.
- We then have to check the rank of the situation $\neg\mathsf{ck}$ in $\top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \neg\mathsf{cb2}$. Since in $\mathcal{KB}$ we have the conditional $\neg\mathsf{ck} \mathrel{|\!\sim_\top} \bot$, that is, the agent knows that the kitchen has been cleaned, the immediate conclusion is that `Rank`$(\mathcal{KB}^\wedge) = \infty$.
- Hence, semantically, $\top \mathrel{|\!\sim_{\neg\mathsf{ck}}} \neg\mathsf{cb2}$ is a conditional that needs to be checked w.r.t. the valuations in the infinite ranks of the minimal epistemic model of $\mathcal{KB}$. This corresponds to checking whether $\neg\mathsf{ck} \mathrel{|\!\sim} \neg\mathsf{cb2}$ follows from $\mathcal{KB}^\wedge_{\infty\downarrow}$, that is, whether `RationalClosure`$(\mathcal{KB}^\wedge_{\infty\downarrow}, \neg\mathsf{ck} \mathrel{|\!\sim} \neg\mathsf{cb2})$ returns `true`. `RationalClosure`$(\mathcal{KB}^\wedge_{\infty\downarrow}, \neg\mathsf{ck} \mathrel{|\!\sim} \neg\mathsf{cb2})$ associates the rank 0 to $\neg\mathsf{ck}$, and $\mathcal{E}_0 = \mathcal{KB}^\wedge_{\infty\downarrow}$. Consequently, $\neg\mathsf{ck} \mathrel{|\!\sim} \neg\mathsf{cb2}$ is in the rational closure of $\mathcal{KB}^\wedge_{\infty\downarrow}$ if and only if $\overline{\mathcal{KB}^\wedge_{\infty\downarrow}} \cup \{\neg\mathsf{ck}\} \models \neg\mathsf{cb2}$, which is the case, since $\overline{\mathcal{KB}^\wedge_{\infty\downarrow}}$ contains $\neg\mathsf{ck} \rightarrow \mathsf{si}$ and $\mathsf{cb2} \wedge \mathsf{si} \rightarrow \bot$.

5.1. Computational complexity of minimal entailment

We now turn our attention to the computational complexity of deciding minimal entailment. We have seen that the entire procedure can be reduced to a sequence of classical propositional entailment tests, with propositional entailment known to be co-NP-complete. Therefore, we have to check, given an SCKB as input, how many classical entailment tests are required in the worst case. We examine each algorithm in turn.

- Given a set of defeasible conditionals $\mathcal{C}$, Algorithm `Exceptional` performs $|\mathcal{C}|$ propositional entailment tests.
- Given a set of defeasible conditionals $\mathcal{C}$, Algorithm `ComputeRanking` runs the algorithm `Exceptional` at most $|\mathcal{C}|$ times in the case where each conditional from $\mathcal{C}$ has a distinct antecedent, and each rank contains exactly one conditional. In such a case, we have that the first iteration of the algorithm `Exceptional` performs $|\mathcal{C}|$ entailment checks, the second one $|\mathcal{C}| - 1$ entailment checks, the third one $|\mathcal{C}| - 2$ entailment checks, and so on. That is, the i-th iteration of `Exceptional` performs $|\mathcal{C}| - i + 1$ propositional entailment checks. So there are fewer than $|\mathcal{C}|^2$ entailment checks and hence Algorithm `ComputeRanking` performs a polynomial number of propositional entailment checks. Note that, given a conditional knowledge base $\mathcal{C}$, we need to run `ComputeRanking` only once.
- Given a set of defeasible conditionals $\mathcal{C}$ and a formula α, Algorithm `Rank` calls `ComputeRanking` (which performs at most $|\mathcal{C}|^2$ entailment checks), and then performs at most a number of entailment checks that corresponds to the number of ranks, which is $|\mathcal{C}|$ at most. Hence Algorithm `Rank` performs a polynomial number of propositional entailment checks.
- Given a set of defeasible conditionals $\mathcal{C}$ and a conditional $\alpha \mathrel{|\!\sim} \beta$, Algorithm `RationalClosure` calls Algorithm `ComputeRanking` once and Algorithm `Rank` once, plus it makes a final entailment check. Hence, the algorithm performs a polynomial number of propositional entailment checks.
- Given an SCKB $\mathcal{KB}$, Algorithm `Partition` runs Algorithm `ComputeRanking` once and Algorithm `Rank` at most $|\mathcal{KB}|$ times. Since $|\mathcal{KB}^{\wedge}| = |\mathcal{KB}|$, running `ComputeRanking` consists of $|\mathcal{KB}|^2$ entailment checks at most. The same holds for each run of `Rank`. Hence running `Partition` consists of at most $|\mathcal{KB}|^2 \cdot (|\mathcal{KB}| + 1) = |\mathcal{KB}|^3 + |\mathcal{KB}|^2$ entailment checks.
- Given an SCKB $\mathcal{KB}$ and a situated conditional $\alpha \mathrel{|\!\sim_\gamma} \beta$, Algorithm `MinimalClosure` runs Algorithm `Partition` once, followed by one entailment check (line 2), one call to Algorithm `Rank` and one call to algorithm `RationalClosure` (with either $\mathcal{KB}^{\wedge}$ or $\mathcal{KB}^{\wedge}_{\infty,\downarrow}$ as argument):
  - `Partition` performs at most $|\mathcal{KB}|^3 + |\mathcal{KB}|^2$ entailment checks.
  - `Rank` performs at most $|\mathcal{KB}|^2$ entailment checks.
  - `RationalClosure` performs at most $|\mathcal{KB}|^2$ entailment checks.

  Hence Algorithm `MinimalClosure` performs a polynomial number of propositional entailments checks.

In summary then, deciding minimal entailment using Algorithm `MinimalClosure` involves a polynomial number of propositional entailment checks, and is therefore in $\mathrm{P}^{\mathrm{coNP}} = \Delta_2^{\mathrm{P}}$. Whether this decision problem is Δ_2^{P}-complete is currently an open question.

6. Related work

With regard to the distinction between a plausible and an implausible state of affairs, a similar distinction has been used by Booth et al. [41], where some pieces of information are considered *credible* while others are not, and a new piece of information is accepted only if it is credible. Nonetheless, in case it is not credible, its plausibility increases every time such a piece of information is iterated. The distinction between plausible and implausible valuations links such an approach with our proposal, but the reasoning problems they model are different. Booth et al. deal with the credibility of a new piece of information, also considering whether the agent is repeatedly exposed to such a piece of information. Here we deal with the distinction between expectations and counterfactuals: given an SC $\alpha \mathrel{|\!\sim_\gamma} \beta$, we could say that if γ is *credible*, then the defeasible conditional $\alpha \mathrel{|\!\sim} \beta$ is evaluated w.r.t. one ranked interpretation (represented by the finite ranks), while it is evaluated w.r.t. another ranked interpretation (represented by the infinite ranks) otherwise.

The connection between our conditional system and belief change is already made quite clear by the *situated AGM postulates* (See Section 3), but it still needs to be properly investigated. Such an investigation should proceed not only from the point of view of the possible definition of interesting revision operators corresponding to our situated conditionals (for example, via some modified version of the Ramsey test) but also from the point of view of the definition of appropriate revision operators modelling the dynamics of SCKB's, in line with what has been done for conditional knowledge bases [42,43]. With respect to this latter problem, the work of Booth et al. [41] offers an interesting perspective on modelling the dynamics of a semantics with plausible and implausible state of affairs.

The literature on the notion of context, which is akin to our use of situation, is vast, and several formalisations and applications of it have been studied across many areas within AI [44–48]. The role of context in conditional-like statements has been explored recently, in particular in defeasible reasoning over description logic ontologies and within semantic frameworks that are closely related to ours. Britz and Varzinczak [49,50], for example, have put forward a notion of defeasible class inclusion parameterised by atomic roles. Their semantics allows for multiple preference relations on objects,

$\langle \infty, \infty \rangle$	$[\![\mathsf{p} \wedge \neg \mathsf{b}]\!] \cup [\![\mathsf{d} \wedge \neg \mathsf{b}]\!]$
$\langle \infty_2, 0 \rangle$	$[\![\mathsf{u}]\!] \setminus ([\![\mathsf{p} \wedge \neg \mathsf{b}]\!] \cup [\![\mathsf{d} \wedge \neg \mathsf{b}]\!])$
$\langle \infty_1, 1 \rangle$	$\overline{\mathsf{p}}\mathsf{dbf}\overline{\mathsf{u}}$, $\mathsf{pdbf}\overline{\mathsf{u}}$
$\langle \infty_1, 0 \rangle$	$\overline{\mathsf{p}}\mathsf{dbf}\overline{\mathsf{u}}$, $\mathsf{pdb}\overline{\mathsf{f}}\overline{\mathsf{u}}$
$\langle \mathtt{f}, 2 \rangle$	$\mathsf{p}\overline{\mathsf{d}}\mathsf{bf}\overline{\mathsf{u}}$
$\langle \mathtt{f}, 1 \rangle$	$\overline{\mathsf{p}}\overline{\mathsf{d}}\mathsf{bf}\overline{\mathsf{u}}$, $\mathsf{p}\overline{\mathsf{d}}\mathsf{b}\overline{\mathsf{f}}\overline{\mathsf{u}}$
$\langle \mathtt{f}, 0 \rangle$	$\overline{\mathsf{p}}\overline{\mathsf{d}}\mathsf{bf}\overline{\mathsf{u}}$, $\overline{\mathsf{p}}\mathsf{d}\overline{\mathsf{b}}\mathsf{f}\overline{\mathsf{u}}$, $\overline{\mathsf{p}}\mathsf{d}\overline{\mathsf{b}}\overline{\mathsf{f}}\overline{\mathsf{u}}$

Fig. 7. An example of an interpretation extending the epistemic interpretation with further infinite levels for more complex counterfactual reasoning.

which is more general than our single-preference approach, and allows for objects to be compared in more than one way. This makes normality (or typicality) context-dependent and gives more flexibility from a modelling perspective. Giordano and Gliozzi [51] consider reasoning about multiple aspects in defeasible description logics where the notion of aspect (or context) is linked to concept names (alias, atoms) also in a multi-preference semantics.

When compared with our framework, neither of the above-mentioned approaches allows for reasoning about objects that are 'forbidden' by the background knowledge. In that respect, our proposal is complementary to theirs, and a contextual form of class inclusion along the lines of the ternary $\mid\!\sim$ here studied, with potential applications going beyond that of defeasible reasoning in ontologies, is worth exploring as future work.

7. Concluding remarks

In this paper, we have made the case for the provision of a simple situated form of conditional. We have shown, using a number of representative examples, that it is sufficiently general to be used in several application domains. The proposed situated conditionals have an intuitive semantics which is based on a semantic construction that has proved to be quite useful in the area of belief change, and is more general and also more fine-grained than the standard preferential semantics. We also showed that the proposed conditionals can be described in terms of a set of postulates. We provided a representation result, showing that the postulates capture exactly the constructions obtained from the proposed semantics. An analysis in terms of the postulates shows that these situated conditionals are suitable for knowledge representation and reasoning, in particular when reasoning about information that is incompatible with background knowledge.

With the basic semantic structures in place, we then proceeded to define a form of entailment for situated conditional knowledge bases that is based on the widely-accepted notion of rational closure for KLM-style reasoning. Moreover, we showed that, like rational closure, entailment for situated conditional knowledge bases is reducible to classical propositional reasoning.

Note that the semantics we have proposed in the present work can easily be refined further. Our framework allows only for the distinction between the *plausible* situations (the valuations with a finite rank that, with different degrees of expectation, define the agent's beliefs), and the *implausible* ones (the valuations with an infinite rank that are not compatible with the agent's beliefs but are still conceivable with different levels of expectation). All other valuations have the *inconceivable* rank $\langle \infty, \infty \rangle$.

There is a fairly straightforward way of refining the framework by allowing for different ranks of the kind $\langle \infty_1, i \rangle$, $\langle \infty_2, i \rangle$, etc., ($i \in \mathbb{N}$). To illustrate the point, assume we add *unicorns* (u) to our vocabulary, and consider the interpretation in Fig. 7. In such a model, we would be able to represent the fact that we believe that unicorns would not exist even if we move to situations in which dodos exist ($\mathsf{u} \mid\!\sim_{\mathsf{d}} \bot$), represented by the ranks $\langle \infty_1, 0 \rangle$ and $\langle \infty_1, 1 \rangle$. We would also be able to move to a further level of implausibility (∞_2) in which the existence of unicorns is considered, making it possible to reason coherently about them. For now, we shall leave such a refinement of our semantic framework for future work.

The work described in this paper assumes classical propositional logic as the underlying logical formalism, but it is worthwhile to consider extending this to other, more expressive logics. In this regard, an extension to Description Logics is perhaps an obvious starting point, particularly since rational closure has already been reformulated for this case [29,52,40, 50]. A different kind of extension of the work presented here is one in which other forms of entailment are investigated. For this, the obvious initial candidate is lexicographic closure [25] and its variants [36,31,53]. More generally, we intend to investigate an extension to the class of entailment relations studied by Casini et al. [31].

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.

Data availability

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

Acknowledgements

The work of Giovanni Casini was partially supported by TAILOR (Foundations of Trustworthy AI Integrating Reasoning, Learning and Optimization), a project funded by EU Horizon 2020 research and innovation programme under GA No 952215.

Ivan Varzinczak has been partially supported by: (*i*) the ANR Chaire IA BE4musIA: BElief change FOR better MUlti-Source Information Analysis; (*ii*) the CNRS/Royal Society project *Non-Classical Reasoning for Enhanced Ontology-based Semantic Technologies*, and (*iii*) the project AGGREEY ANR-22-CE23-0005 of the French National Research Agency (ANR).

Thanks to the anonymous reviewers for their helpful suggestions.

Appendix A. Proofs

Theorem 3.1. *Every epistemic interpretation generates a BSC (see Definition 3.2). Nevertheless, the converse does not hold, i.e., some BSCs cannot be generated by any epistemic state.*

Proof. Consider any epistemic interpretation $\mathscr{E}$ and pick any $\gamma \in \mathcal{L}$. We consider three disjoint and covering cases.

Case 1: If $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma]\!] \neq \emptyset$, then define $\mathscr{R}$ from $\mathscr{E}$ as follows: (*i*) for all $u \in \mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma]\!]$, $\mathscr{R}(u) \stackrel{\text{def}}{=} i$, where $\mathscr{E}(u) = \langle \mathfrak{f}, i\rangle$; (*ii*) for all $u \in \mathcal{U} \setminus \mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma]\!]$, $\mathscr{R}(u) \stackrel{\text{def}}{=} \infty$. It follows from Definition 3.2 and the definition of satisfaction of $\mathrel{|\!\sim}$-statements in ranked interpretations that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$ iff $\mathscr{R} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim} \beta$. From Theorem 2.1, it follows that the $\mathrel{|\!\sim}$ generated by $\mathscr{R}$ satisfies the original KLM postulates. Hence, it follows that $\mathrel{|\!\sim}_\gamma$ satisfies the situated rationality postulates.

Case 2: If $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma]\!] = \emptyset$ but $\mathcal{U}^{\infty}_{\mathscr{E}} \cap [\![\gamma]\!] \neq \emptyset$, then define $\mathscr{R}$ from $\mathscr{E}$ as follows: (*i*) for all $u \in \mathcal{U}^{\infty}_{\mathscr{E}} \cap [\![\gamma]\!]$, $\mathscr{R}(u) \stackrel{\text{def}}{=} i$, where $\mathscr{E}(u) = \langle \infty, i\rangle$; (*ii*) for all $u \in \mathcal{U} \setminus (\mathcal{U}^{\infty}_{\mathscr{E}} \cap [\![\gamma]\!])$, $\mathscr{R}(u) \stackrel{\text{def}}{=} \infty$. It follows from Definition 3.2 and the definition of satisfaction for $\mathrel{|\!\sim}$-statements in ranked interpretations that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$ iff $\mathscr{R} \Vdash \alpha \mathrel{|\!\sim} \beta$. From Theorem 2.1, it follows that the $\mathrel{|\!\sim}$ generated by $\mathscr{R}$ satisfies the original KLM postulates. For this specific γ it then follows that $\mathrel{|\!\sim}_\gamma$ satisfies the situated rationality postulates.

Case 3: If $[\![\gamma]\!] \subseteq \mathcal{U} \setminus (\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cup \mathcal{U}^{\infty}_{\mathscr{E}})$, then $\mathscr{R}(u) \stackrel{\text{def}}{=} \infty$ for all $u \in [\![\gamma]\!]$. Again, it follows from Definition 3.2 and the definition of satisfaction for $\mathrel{|\!\sim}$ in ranked interpretations that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$ iff $\mathscr{R} \Vdash \alpha \mathrel{|\!\sim} \beta$. From Theorem 2.1, it follows that the $\mathrel{|\!\sim}$ generated by $\mathscr{R}$ satisfies the original KLM postulates. For this specific γ, it then follows that $\mathrel{|\!\sim}_\gamma$ satisfies the situated rationality postulates.

Putting the three cases above together, it then follows immediately that the situated conditional $\mathrel{|\!\sim}_\gamma$ obtained from $\mathscr{E}$ satisfies the situated rationality postulates.

Now, in order to show that the converse does not hold, consider the language generated from (and only) $\{\mathsf{p}, \mathsf{q}\}$. Note first that there is a ranked interpretation $\mathscr{R}$ such that $\mathscr{R} \Vdash \alpha \mathrel{|\!\sim} \beta$ iff $\mathsf{p} \wedge \mathsf{q} \wedge \alpha \models \beta$. From Theorem 2.1, it follows that $\mathrel{|\!\sim}$, defined in this way, is a rational conditional, and therefore satisfies the situated KLM postulates. Similarly, there is a ranked interpretation $\mathscr{R}'$ such that $\mathscr{R}' \Vdash \alpha \mathrel{|\!\sim} \beta$ iff $\mathsf{p} \wedge \mathsf{q} \wedge \alpha \models \beta$. From Theorem 2.1, it follows that $\mathrel{|\!\sim}$, defined in this way, is a rational conditional, and therefore satisfies the situated KLM postulates. Now, define a situated conditional by letting $\alpha \mathrel{|\!\sim}_{\mathsf{p}} \beta$ iff $\mathsf{p} \wedge \mathsf{q} \wedge \alpha \models \beta$, and $\alpha \mathrel{|\!\sim}_\gamma \beta$ iff $\alpha \models \beta$, for every γ other than p. It then follows immediately that this situated conditional is a BSC. However, it is easy to see that it cannot be generated by an epistemic interpretation. To see why, observe that $\mathsf{p} \mathrel{|\!\sim}_{\mathsf{p}} \mathsf{q}$, but that $\mathsf{p} \not\mathrel{|\!\sim}_{\mathsf{p}\vee\mathsf{p}} \mathsf{q}$. □

Theorem 3.2. *Every epistemic interpretation generates an FSC. Every FSC can be generated by an epistemic interpretation.*

Proof. Let $\mathscr{E}$ be an epistemic interpretation and let $\gamma \in \mathcal{L}$. Suppose $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma]\!] \neq \emptyset$. Then if $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$, it follows by Definition 3.2 that $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim}_\top \beta$. On the other hand, if $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma]\!] = \emptyset$, then $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim}_\top \beta$. This means that the situated conditional $\mathrel{|\!\sim}$ obtained from $\mathscr{E}$ as follows satisfies Inc: $\alpha \mathrel{|\!\sim}_\gamma \beta$ iff $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$.

Suppose $\mathscr{E} \nVdash \top \mathrel{|\!\sim}_\top \neg\gamma$. This means that $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma]\!] \neq \emptyset$. Then if $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim}_\top \beta$, it follows by Definition 3.2 that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$. This means that the situated conditional $\mathrel{|\!\sim}$ obtained from $\mathscr{E}$ as follows satisfies Vac: $\alpha \mathrel{|\!\sim}_\gamma \beta$ iff $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$.

That the situated conditional obtained from $\mathscr{E}$ as follows satisfies Ext follows immediately from Definition 3.2: $\alpha \mathrel{|\!\sim}_\gamma \beta$ iff $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_\gamma \beta$.

For SupExp we consider two cases. For Case 1, if $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma \wedge \delta]\!] \neq \emptyset$, then the result follows easily. For Case 2, suppose $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\gamma \wedge \delta]\!] = \emptyset$. If $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\delta]\!] = \emptyset$, then the result follows easily. Otherwise the result follows from the fact that $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \cap [\![\alpha \wedge \gamma \wedge \delta]\!] = \emptyset$.

For SubExp, suppose that $\mathscr{E} \Vdash \delta \mathrel{|\!\sim}_\top \bot$. This means $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim}_\delta \beta$ implies that $\mathcal{U}^{\infty}_{\mathscr{E}} \cap [\![\alpha \wedge \gamma \wedge \delta]\!] \subseteq [\![\beta]\!]$, from which it follows that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma\wedge\delta} \beta$.

For the converse, consider any FSC $\mathrel{|\!\sim}$. We construct an epistemic interpretation $\mathscr{E}$ as follows. First, consider $\mathrel{|\!\sim}_\top$. Since it satisfies the situated KLM postulates, there is a ranked interpretation $\mathscr{R}$ such that $\mathscr{R} \Vdash \alpha \mathrel{|\!\sim} \beta$ iff $\alpha \mathrel{|\!\sim}_\top \beta$. We set $\mathcal{U}^{\mathfrak{f}}_{\mathscr{E}} \stackrel{\text{def}}{=} \mathcal{U}^{\mathfrak{f}}_{\mathscr{R}}$, and for all $u \in \mathcal{U}^{\mathfrak{f}}_{\mathscr{E}}$, we let $\mathscr{E}(u) \stackrel{\text{def}}{=} \langle \mathfrak{f}, \mathscr{R}(u)\rangle$. Next, let $\mathcal{U}' \stackrel{\text{def}}{=} \mathcal{U} \setminus \mathcal{U}^{\mathfrak{f}}_{\mathscr{E}}$. Let $k^{\mathfrak{f}}$ be a formula such that $[\![k^{\mathfrak{f}}]\!] = \mathcal{U}^{\mathfrak{f}}_{\mathscr{E}}$. Similarly, let k^{∞} be a formula such that $[\![k^{\infty}]\!] = \mathcal{U}'$. Now, consider $\mathrel{|\!\sim}_{k^\infty}$. Since it satisfies the situated KLM postulates, there is a ranked interpretation $\mathscr{R}'$ such that $\mathscr{R}' \Vdash \alpha \mathrel{|\!\sim} \beta$ iff $\alpha \mathrel{|\!\sim}_{k^\infty} \beta$. We let $\mathcal{U}^{\infty}_{\mathscr{E}} \stackrel{\text{def}}{=} \{u \in \mathcal{U}' \mid \mathscr{R}'(u) \neq \infty\}$, and for all $u \in \mathcal{U}'$, we let $\mathscr{E}(u) \stackrel{\text{def}}{=} \langle \infty, \mathscr{R}'(u)\rangle$. Observe that for some $u \in \mathcal{U}'$ it may be the case that $\mathscr{E}(u) = \langle \infty, \infty\rangle$, which means that for such a u,

$u \notin \mathcal{U}^{\infty}_{\mathscr{E}}$. It is easily verified that $\mathscr{E}$ is indeed an epistemic interpretation. Next we show that $\alpha \mathrel{|\!\sim}_{\gamma} \beta$ iff $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$. We do so by considering two cases.

Case 1: $\mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \cap [\![\gamma]\!] \neq \emptyset$. Note first that it follows easily from the construction of $\mathscr{E}$ that $\alpha \mathrel{|\!\sim}_{\top} \beta$ iff $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\top} \beta$. Suppose $\alpha \mathrel{|\!\sim}_{\gamma} \beta$. By Inc, $\alpha \wedge \gamma \mathrel{|\!\sim}_{\top} \beta$ and therefore $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim}_{\top} \beta$, and $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$, by definition. Conversely, suppose $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$. Then by definition, $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim}_{\top} \beta$, and therefore $\alpha \wedge \gamma \mathrel{|\!\sim}_{\top} \beta$. Since $\top \not\mathrel{|\!\sim}_{\top} \neg\gamma$, it then follows from Vac that $\alpha \mathrel{|\!\sim}_{\gamma} \beta$.

Case 2: $\mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \cap [\![\gamma]\!] \neq \emptyset$. By the construction of $\mathscr{E}$, it follows that $\alpha \mathrel{|\!\sim}_{k^{\infty}} \beta$ iff $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{k^{\infty}} \beta$. Suppose $\alpha \mathrel{|\!\sim}_{\gamma} \beta$. Note that $\gamma \equiv k^{\infty}$. By Ext, $\alpha \mathrel{|\!\sim}_{\gamma \wedge k^{\infty}} \beta$ and so, by SupExp, $\alpha \wedge \gamma \mathrel{|\!\sim}_{k^{\infty}} \beta$. It then follows that $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim}_{k^{\infty}} \beta$ and, by Definition 3.2, that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$. Conversely, suppose that $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$. Then $\mathscr{E} \Vdash \alpha \wedge \gamma \mathrel{|\!\sim}_{k^{\infty}} \beta$, by Definition 3.2, and, therefore, using Ext, we have $\alpha \wedge \gamma \mathrel{|\!\sim}_{\gamma \wedge k^{\infty}} \beta$. Note that $\mathscr{E} \Vdash k^{\mathtt{f}} \mathrel{|\!\sim}_{\top} \bot$ and therefore $k^{\mathtt{f}} \mathrel{|\!\sim}_{\top} \bot$. By SubExp it then follows that $\alpha \mathrel{|\!\sim}_{\gamma \wedge k^{\infty}} \beta$, and by Ext that $\alpha \mathrel{|\!\sim}_{\gamma} \beta$ holds. □

Lemma 4.1. *For every epistemic interpretation $\mathscr{E}$, if $\mathcal{U}^{\mathtt{f}}_{\mathscr{E}} \cap [\![\gamma]\!] = \emptyset$, then $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$ iff $\mathscr{E}^{\infty}_{\downarrow} \Vdash \alpha \wedge \gamma \mathrel{|\sim} \beta$.*

Proof. In case $\alpha \wedge \gamma$ is not logically consistent, the lemma holds since $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$ and $\mathscr{E}^{\infty}_{\downarrow} \Vdash \alpha \wedge \gamma \mathrel{|\sim} \beta$ for any β. Hence we assume that $\alpha \wedge \gamma$ is logically consistent.

Let $\mathscr{E} \Vdash \neg\gamma$, that is, there are no valuations in the finite ranks satisfying γ. Then the satisfaction of the conditionals with situation γ must be checked, referring to the valuations that are ranked as infinite. $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$ implies two possible situations: either there are some valuations among the ones in $[\![\gamma]\!]$ that are ranked as infinite and satisfy $\alpha \wedge \gamma$, and among them, all the minimal ones satisfy also β; or all the valuations satisfying $\alpha \wedge \gamma$ have rank $\langle \infty, \infty \rangle$. γ has finite rank in $\mathscr{E}^{\infty}_{\downarrow}$, or the rank $\langle \infty, \infty \rangle$. In the latter case, we have $\mathscr{E}^{\infty}_{\downarrow} \Vdash \alpha \wedge \gamma \mathrel{|\sim} \beta$. In the former case, the rank of γ in $\mathscr{E}$ is $\langle \infty, i \rangle$, with $i < \infty$, that is, the rank of $\gamma \wedge \alpha$ in $\mathscr{E}^{\infty}_{\downarrow}$ is $\langle \mathtt{f}, j \rangle$, for some j s.t. $i \leq j < \infty$, or $\langle \infty, \infty \rangle$. In the latter case, again, it is straightforward to conclude $\mathscr{E}^{\infty}_{\downarrow} \Vdash \alpha \wedge \gamma \mathrel{|\sim} \beta$. In the former case, we have $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$, and the construction of $\mathscr{E}^{\infty}_{\downarrow}$ imposes that the minimal valuations in $[\![\alpha \wedge \gamma]\!]$ satisfy also β, that is, $\mathscr{E}^{\infty}_{\downarrow} \Vdash \alpha \wedge \gamma \mathrel{|\sim} \beta$.

The proof is analogous in the opposite direction. Let $\mathscr{E} \Vdash \neg\gamma$ and $\mathscr{E}{\downarrow}^{\infty} \Vdash \alpha \wedge \gamma \mathrel{|\sim} \beta$. Either the minimal valuations in $\mathscr{E}{\downarrow}^{\infty}$ satisfying $\alpha \wedge \gamma$ are in rank $\langle \mathtt{f}, i \rangle$, for some $i < \infty$, and they all satisfy β (Case 1), or they are in $\langle \infty, \infty \rangle$ (Case 2). Since $\mathscr{E} \Vdash \neg\gamma$, in $\mathscr{E}$ all the valuations satisfying γ are in $\mathcal{U}^{\infty}_{\mathscr{E}} \cup \langle \infty, \infty \rangle$, and consequently the satisfaction of the SC $\alpha \mathrel{|\!\sim}_{\gamma} \beta$ needs to be judged considering the valuations in $\mathcal{U}^{\infty}_{\mathscr{E}} \cup \langle \infty, \infty \rangle$. If we are in Case 1, since a valuation w has rank $\langle \mathtt{f}, i \rangle$ in $\mathscr{E}{\downarrow}^{\infty}$ iff w has rank $\langle \infty, i \rangle$ in $\mathscr{E}$ (see the definition of *counterfactual shifting* in Section 4), we have that the minimal valuations in $\mathscr{E}$ satisfying $\alpha \wedge \gamma$ have rank $\langle \infty, i \rangle$, for some $i < \infty$, and they all satisfy β. If we are in Case 2, since $\mathscr{E} \Vdash \neg\gamma$, a valuation w that satisfies γ can have rank $\langle \infty, \infty \rangle$ in $\mathscr{E}{\downarrow}^{\infty}$ only if w has rank $\langle \infty, \infty \rangle$ in $\mathscr{E}$ (again, see the definition of *counterfactual shifting* in Section 4 and consider that no valuation with a finite rank satisfies γ in $\mathscr{E}$); hence we have that the minimal valuations in $\mathscr{E}$ satisfying $\alpha \wedge \gamma$ are in rank $\langle \infty, \infty \rangle$. In both cases, we have $\mathscr{E} \Vdash \alpha \mathrel{|\!\sim}_{\gamma} \beta$. □

Proposition 4.4. *Let $\mathcal{KB}$ be a consistent SCKB, and let $\mathscr{E}_{\mathcal{KB}}$ be the epistemic interpretation built as in Definition 4.6. Then $\mathscr{E}_{\mathcal{KB}}$ is the only minimal epistemic model of $\mathcal{KB}$.*

Proof. We divide the proof into two parts. First, we prove that $\mathscr{E}_{\mathcal{KB}}$ is a minimal epistemic model, then that it is also the *only* minimal epistemic model.

Regarding minimality, we proceed by contradiction. We know, by Proposition 4.3, that $\mathscr{E}_{\mathcal{KB}}$ is an epistemic model of $\mathcal{KB}$. Assume it is *not* minimal, that is, assume there is an epistemic model $\mathscr{E}'$ of $\mathcal{KB}$ s.t., for every $u \in \mathcal{U}$, $\mathscr{E}'(u) \leq \mathscr{E}_{\mathcal{KB}}(u)$, and there is a $w \in \mathcal{U}$ s.t. $\mathscr{E}'(w) < \mathscr{E}_{\mathcal{KB}}(w)$. Regarding the ranking of w, we have two possibilities:

Case 1. $\mathscr{E}_{\mathcal{KB}}(w) = \langle \mathtt{f}, i \rangle$, for some i, and $\mathscr{E}'(w) = \langle \mathtt{f}, j \rangle$, for some $j < i$. Let $\mathcal{KB}^{\mathtt{f}}_{\mathscr{E}'} = \{\alpha \mathrel{|\!\sim}_{\gamma} \beta \in \mathcal{KB} \mid \mathscr{E}' \nVdash \neg\gamma\}$. By Corollary 4.1, $\mathscr{E}' \Vdash \alpha \wedge \gamma \mathrel{|\sim} \beta$, for every $\alpha \mathrel{|\!\sim}_{\gamma} \beta \in \mathcal{KB}^{\mathtt{f}}_{\mathscr{E}'}$. Consider the ranked interpretation $\mathscr{R}'$ defined as:

$$\mathscr{R}'(u) = \begin{cases} i, & \text{if } \mathscr{E}'(u) = \langle \mathtt{f}, i \rangle, \text{ for some } i; \\ \infty, & \text{otherwise.} \end{cases}$$

$\mathscr{R}'$ above is clearly a ranked model of every $\alpha \wedge \gamma \mathrel{|\sim} \beta$ s.t. $\alpha \mathrel{|\!\sim}_{\gamma} \beta \in \mathcal{KB}^{\mathtt{f}}_{\mathscr{E}'}$. Since $\mathscr{R}'$ has only one infinite rank, ∞, $\mathscr{R}'$ is also a ranked model of every $\alpha \wedge \gamma \mathrel{|\sim} \beta$ s.t. $\alpha \mathrel{|\!\sim}_{\gamma} \beta \in \mathcal{KB} \setminus \mathcal{KB}^{\mathtt{f}}_{\mathscr{E}'}$, since the minimal valuations satisfying their premises are in $[\![\langle \infty, \infty \rangle]\!]$, and consequently they are trivially satisfied. Hence, $\mathscr{R}'$ is a ranked model of $\mathcal{KB}^{\wedge}$.
By Definition 4.6, $\mathscr{E}_{\mathcal{KB}}$ has been built using the minimal ranked model $\mathscr{R}$ of $\mathcal{KB}^{\wedge}$. But now we have also a ranked model $\mathscr{R}'$ of $\mathcal{KB}^{\wedge}$ s.t. $\mathscr{R}'$ is a ranked model of $\mathcal{KB}^{\wedge}$ and, moreover, $\mathscr{R}'$ is preferred to $\mathscr{R}$: by Definition 4.6, for every $u \in \mathcal{U}$,

$$\mathscr{R}(u) = \begin{cases} i, & \text{if } \mathscr{E}_{\mathcal{KB}}(u) = \langle \mathtt{f}, i \rangle, \text{ for some } i; \\ \infty, & \text{otherwise.} \end{cases}$$

Since we have assumed that $\mathscr{E}'(u) \leq \mathscr{E}_{\mathcal{KB}}(u)$ for every u, and there is a w s.t. $\mathscr{E}_{\mathcal{KB}}(w) = \langle \mathtt{f}, i \rangle$, for some i, and $\mathscr{E}'(w) = \langle \mathtt{f}, j \rangle$, for some $j < i$, we have that $\mathscr{R}'(u) \leq \mathscr{R}(u)$ for every u, and there is a w s.t. $\mathscr{R}(w) = i$, for some

$i \in \mathbb{N}$, and $\mathscr{R}'(w) = j$, for some $j < i$. That is, for every $u \in \mathcal{U}$ $\mathscr{R}'(u) \leq \mathscr{R}(u)$, and $\mathscr{R}'(w) < \mathscr{R}(w)$. Hence $\mathscr{R}$ is not the minimal model of $\mathcal{KB}^{\wedge}$, and this leads to a contradiction.

Case 2. $\mathscr{E}_{\mathcal{KB}}$ and $\mathscr{E}'$ are identical w.r.t. the finite ranks, and $\mathscr{E}_{\mathcal{KB}}(w) = \langle \infty, i \rangle$, for some i. We have two subcases: $\mathscr{E}'(w) = \langle \infty, j \rangle$, for some $j < i$, or $\mathscr{E}'(w) = \langle \mathtt{f}, j \rangle$, for some j. The latter subcase leads to a contradiction: it can be proved analogously to Case 1. It remains to prove the first subcase.

The proof is still close to the one for Case 1 above; we simply have to refer to the counterfactual shiftings of $\mathscr{E}_{\mathcal{KB}}$ and $\mathscr{E}'$, $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}$ and $\mathscr{E}'^{\infty}_{\downarrow}$ (see page 11). Since $\mathscr{E}_{\mathcal{KB}}$ and $\mathscr{E}'$ are epistemic models of $\mathcal{KB}^{\wedge}$, $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}$ and $\mathscr{E}'^{\infty}_{\downarrow}$ are epistemic models of $\mathcal{KB}^{\wedge}_{\infty}$, and $\mathscr{E}'^{\infty}_{\downarrow}$ is preferred to $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}$. From $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}$ and $\mathscr{E}'^{\infty}_{\downarrow}$, we can extract two ranked interpretations, $\mathscr{R}^{\infty}_{\mathcal{KB}}$ and $\mathscr{R}'^{\infty}$ (see Definition 3.3), that are both epistemic models of $\mathcal{KB}^{\wedge}_{\infty}$. In the construction of $\mathscr{E}_{\mathcal{KB}}$, following Definition 4.6, we have used for the infinite ranks the ranked interpretation $\mathscr{R}^{\infty}_{\mathcal{KB}}$, which, also by Definition 4.6, must be the minimal ranked model of $\mathcal{KB}^{\wedge}_{\infty}$. But in the present case, $\mathscr{R}^{\infty}_{\mathcal{KB}}$ cannot be the minimal ranked model of $\mathcal{KB}^{\wedge}_{\infty}$, since $\mathscr{R}'^{\infty}$ is a ranked model of $\mathcal{KB}^{\wedge}_{\infty}$ that is preferred to $\mathscr{R}^{\infty}_{\mathcal{KB}}$. This leads to a contradiction.

To conclude this part, in all the possible cases, if $\mathscr{E}_{\mathcal{KB}}$ is not a minimal epistemic model of $\mathcal{KB}$, then we end up with a contradiction. Hence $\mathscr{E}_{\mathcal{KB}}$ must be a minimal epistemic model of $\mathcal{KB}$.

The final step consists in proving that $\mathscr{E}_{\mathcal{KB}}$ is the *only* minimal epistemic model of $\mathcal{KB}$. The procedure is again by contradiction, assuming that $\mathscr{E}_{\mathcal{KB}}$ is not the only minimal epistemic model of $\mathcal{KB}$. Hence, let $\mathscr{E}'$ be another minimal epistemic model of $\mathcal{KB}$. The structure of the proof actually mirrors the one for the previous part, about the minimality of $\mathscr{E}_{\mathcal{KB}}$. Again, we can distinguish two main cases.

Case 1. $\mathscr{E}_{\mathcal{KB}}$ and $\mathscr{E}'$ differ w.r.t. the ranking of some valuations among the ones ranked as finite. From $\mathscr{E}_{\mathcal{KB}}$ and $\mathscr{E}'$, we can extract, respectively, the ranked models $\mathscr{R}$ and $\mathscr{R}'$, which are both ranked models of $\mathcal{KB}^{\wedge}$. But, by Definition 4.5, $\mathscr{R}$ is the only minimal ranked model of $\mathcal{KB}^{\wedge}$, that is, $\mathscr{R} \prec \mathscr{R}'$, which implies that $\mathscr{E}'$ cannot be a minimal epistemic model of $\mathcal{KB}$.

Case 2. $\mathscr{E}_{\mathcal{KB}}$ and $\mathscr{E}'$ do not differ w.r.t. the ranking of the valuations that are ranked as finite in both of them but differ w.r.t. the ranking of some valuation, w, that is ranked as infinite in one of the two. W.l.o.g., we assume that w is ranked as infinite in $\mathscr{E}_{\mathcal{KB}}$. We have two subcases: $\mathscr{E}'(w) = \langle \infty, j \rangle$, for some j, or $\mathscr{E}'(w) = \langle \mathtt{f}, j \rangle$, for some j. The latter subcase leads to a contradiction: it can be proved analogously to Case 1 using the extracted ranked models. It remains to show the first subcase.

The proof is still analogous to Case 2 above. We refer to the counterfactual shiftings of $\mathscr{E}_{\mathcal{KB}}$ and $\mathscr{E}'$, $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}$ and $\mathscr{E}'^{\infty}_{\downarrow}$. Since $\mathscr{E}_{\mathcal{KB}}$ and $\mathscr{E}'$ are epistemic models of $\mathcal{KB}^{\wedge}$ and they are identical w.r.t. the finite ranks, $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}$ and $\mathscr{E}'^{\infty}_{\downarrow}$ are epistemic models of $\mathcal{KB}^{\wedge}_{\infty}$. From $\mathscr{E}^{\infty}_{\mathcal{KB}\downarrow}$ and $\mathscr{E}'^{\infty}_{\downarrow}$, we can extract two ranked interpretations, $\mathscr{R}^{\infty}_{\mathcal{KB}}$ and $\mathscr{R}'^{\infty}$ (see Definition 3.3), that are both ranked models of $\mathcal{KB}^{\wedge}_{\infty}$. In the construction of $\mathscr{E}_{\mathcal{KB}}$, following Definition 4.5, we have used for the infinite ranks the ranked interpretation $\mathscr{R}^{\infty}_{\mathcal{KB}}$, which, also by Definition 4.5, must be the minimal ranked model of $\mathcal{KB}^{\wedge}_{\infty}$. If $\mathscr{R}^{\infty}_{\mathcal{KB}}$ is the minimal ranked model of $\mathcal{KB}^{\wedge}_{\infty}$, then $\mathscr{R}^{\infty}_{\mathcal{KB}}$ is preferred to $\mathscr{R}'^{\infty}$, and, by construction, $\mathscr{E}_{\mathcal{KB}}$ must be preferred to $\mathscr{E}'$. This leads to a contradiction.

To conclude, if we assume that there is another minimal epistemic model of $\mathcal{KB}$ beyond $\mathscr{E}_{\mathcal{KB}}$, we end up with a contradiction. Hence, $\mathscr{E}_{\mathcal{KB}}$ must be the only minimal epistemic model of $\mathcal{KB}$. □

Theorem 5.1. *Let $\mathcal{KB}$ be an SCKB.* `MinimalClosure`$(\mathcal{KB}, \alpha \mathrel{|\!\sim_{\gamma}} \beta)$ *returns* `true` *iff* $\mathcal{KB} \models_m \alpha \mathrel{|\!\sim_{\gamma}} \beta$.

Proof. We already know that algorithms `Exceptional`, `ComputeRanking`, `Rank` and `RationalClosure` are complete and correct w.r.t. the corresponding semantic notions.

As a first step, we need to prove that algorithm `Partition` returns the correct result, that is, the sets $\mathcal{KB}_{\infty}$ and $\mathcal{KB}^{\wedge}_{\infty\downarrow}$ correspond to the same sets introduced in Definition 4.6.

The correspondence of $\mathcal{KB}_{\infty}$ to the semantic notion introduced in Definition 4.6 is guaranteed by the correctness of algorithm `ComputeRanking` w.r.t. the semantic definition of ranks w.r.t. the rational closure.

To prove the correspondence of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$ to the semantic notion in Definition 4.6, we need to prove also that the defeasible conditionals $\mu \mathrel{|\!\sim} \bot$ and $\mathsf{sent}(\mathcal{U}^{\mathtt{f}}_{\mathscr{R}}) \mathrel{|\!\sim} \bot$ are equivalent, which is an immediate consequence of Lemma 5.1 and the LLE postulate.

Now we can check the correctness of algorithm `MinimalClosure`. We consider the possible cases as presented in the algorithm.

Case 1. $\overline{\mathcal{KB}^{\wedge}} \models \bot$.

By Corollary 4.2, $\overline{\mathcal{KB}^{\wedge}} \models \bot$ iff $\mathcal{KB}$ is inconsistent, and in such a case $\mathcal{KB} \models_m \alpha \mathrel{|\!\sim_{\gamma}} \beta$ for every α, γ, β, and the algorithm behaves correctly.

Case 2. $\overline{\mathcal{KB}^{\wedge}} \not\models \bot$ and `Rank`$(\mathcal{KB}^{\wedge}, \gamma) < \infty$.

We have to prove that in this case, $\alpha \wedge \gamma \mathrel{|\!\sim} \beta$ is in the RC of $\mathcal{KB}^{\wedge}$ iff $\alpha \mathrel{|\!\sim_{\gamma}} \beta$ is in the minimal closure of $\mathcal{KB}$.

Assume $\alpha \wedge \gamma \mid\sim \beta$ is in the RC of $\mathcal{KB}^{\wedge}$, and let $\mathscr{R}$ be the minimal ranked model of $\mathcal{KB}^{\wedge}$. That means that $\min[\![\alpha \wedge \gamma]\!]_{\mathscr{R}} \subseteq [\![\beta]\!]$. Also, since $\texttt{Rank}(\mathcal{KB}^{\wedge}, \gamma) < \infty$, we have that $[\![\gamma]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{R}} \neq \emptyset$. By construction of the minimal epistemic model of $\mathcal{KB}$, $\mathcal{E}_{\mathcal{KB}}$, $\mathcal{U}^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} = \mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$, and the rank of each valuation is the same. Consequently, we have that $[\![\gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} \neq \emptyset$. According to Definition 3.2, we have to check whether $[\![\gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} \subseteq [\![\neg\alpha]\!]$ or $[\![\alpha \wedge \gamma]\!]_{\mathcal{E}_{\mathcal{KB}}} \subseteq [\![\beta]\!]$. From $\mathscr{R} \Vdash \alpha \wedge \gamma \mid\sim \beta$, we single out two possible cases:

- $\texttt{Rank}(\mathcal{KB}^{\wedge}, \alpha) = \infty$. This implies that $[\![\gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} \subseteq [\![\neg\alpha]\!]$.
- Otherwise, in $\mathscr{R}$ we have $\min[\![\alpha \wedge \gamma]\!]^{\mathtt{f}} \subseteq [\![\beta]\!]$. Since $\mathcal{E}_{\mathcal{KB}}$ preserves in $\mathcal{U}^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}}$ the same ranking as in $\mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$, we have $[\![\alpha \wedge \gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} \subseteq [\![\beta]\!]$.

We can conclude that $\mathcal{E}_{\mathcal{KB}} \Vdash \alpha \mid\sim_{\gamma} \beta$.

Now we check the opposite direction: we assume $\mathcal{E}_{\mathcal{KB}} \Vdash \alpha \mid\sim_{\gamma} \beta$. Since $\texttt{Rank}(\mathcal{KB}^{\wedge}, \gamma) < \infty$, by Definition 4.6 we have that $[\![\gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} \neq \emptyset$. The latter, together with $\mathcal{E}_{\mathcal{KB}} \Vdash \alpha \mid\sim_{\gamma} \beta$, implies $[\![\alpha \wedge \gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} \subseteq [\![\beta]\!]$. By Definition 4.6, this condition implies that $[\![\alpha \wedge \gamma]\!]^{\mathtt{f}}_{\mathscr{R}} \subseteq [\![\beta]\!]$, which in turn implies $\mathscr{R} \Vdash \alpha \wedge \gamma \mid\sim \beta$.

Case 3. $\overline{\mathcal{KB}^{\wedge}} \not\models \bot$ and $\texttt{Rank}(\mathcal{KB}^{\wedge}, \gamma) = \infty$.

We have to prove that in this case $\alpha \wedge \gamma \mid\sim \beta$ is in the RC of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$ iff $\alpha \mid\sim_{\gamma} \beta$ is in the minimal closure of $\mathcal{KB}$.

Since $\texttt{Rank}(\mathcal{KB}^{\wedge}, \gamma) = \infty$, we have that $[\![\gamma]\!] \cap \mathcal{U}^{\mathtt{f}}_{\mathscr{R}} = \emptyset$. By construction of the minimal epistemic model of $\mathcal{KB}$, $\mathcal{E}_{\mathcal{KB}}$, $\mathcal{U}^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} = \mathcal{U}^{\mathtt{f}}_{\mathscr{R}}$. Consequently, we have that $[\![\gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} = \emptyset$ and $[\![\gamma]\!]_{\mathcal{E}_{\mathcal{KB}}}$ all have rank $\langle\infty, j\rangle$, for some j.

Assume $\alpha \wedge \gamma \mid\sim \beta$ is in the RC of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$, and let $\mathscr{R}'$ be the minimal ranked model of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$. According to Definition 3.2, we have to check whether $[\![\alpha \wedge \gamma]\!]_{\mathcal{E}_{\mathcal{KB}}} \subseteq [\![\beta]\!]$. Assume this is not the case, that is, $\mathcal{E}_{\mathcal{KB}} \nVdash \alpha \mid\sim_{\gamma} \beta$. Since $[\![\gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} = \emptyset$, all the valuations in $[\![\gamma]\!]_{\mathcal{E}_{\mathcal{KB}}}$ are ranked as infinite, and $\mathcal{E}_{\mathcal{KB}} \nVdash \alpha \mid\sim_{\gamma} \beta$ implies that there is a valuation w in $[\![\alpha \wedge \gamma]\!]_{\mathcal{E}_{\mathcal{KB}}}$ s.t. $w \nVdash \beta$. Let $w \in [\![\langle\infty, i\rangle]\!]$, for some $i < \infty$, and $w \preceq v$, for every $v \in [\![\alpha \wedge \gamma]\!]$. By Definition 4.5, in $\mathscr{R}'$ we have $w \in [\![i]\!]$, for some $i < \infty$, and $w \preceq v$, for every $v \in [\![\alpha \wedge \gamma]\!]$. Hence, we would have $\mathscr{R}' \nVdash \alpha \wedge \gamma \mid\sim \beta$, which is against our hypothesis that $\alpha \wedge \gamma \mid\sim \beta$ is in the RC of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$.

Now we assume $\mathcal{E}_{\mathcal{KB}} \Vdash \alpha \mid\sim_{\gamma} \beta$. Again, since $[\![\gamma]\!]^{\mathtt{f}}_{\mathcal{E}_{\mathcal{KB}}} = \emptyset$, all the valuations in $[\![\gamma]\!]_{\mathcal{E}_{\mathcal{KB}}}$ are ranked as infinite. The latter, together with Definition 4.5, implies that $[\![\gamma]\!]_{\mathcal{E}_{\mathcal{KB}}} = \min[\![\alpha \wedge \gamma]\!]_{\mathscr{R}'}$, and consequently $[\![\gamma]\!]_{\mathcal{E}_{\mathcal{KB}}} \subseteq [\![\beta]\!]$ implies $\min[\![\alpha \wedge \gamma]\!]_{\mathscr{R}'} \subseteq [\![\beta]\!]$. We can conclude $\mathscr{R}' \Vdash \alpha \wedge \gamma \mid\sim \beta$, that is, $\alpha \wedge \gamma \mid\sim \beta$ is in the RC of $\mathcal{KB}^{\wedge}_{\infty\downarrow}$.

We have proved that in all possible cases $\texttt{MinimalClosure}(\mathcal{KB}, \alpha \mid\sim_{\gamma} \beta)$ returns $\texttt{true}$ iff $\mathcal{KB} \models_m \alpha \mid\sim_{\gamma} \beta$. □

References

[1] G. Casini, T. Meyer, I. Varzinczak, Contextual conditional reasoning, in: K. Leyton-Brown, Mausam (Eds.), Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI), AAAI Press, 2021, pp. 6254–6261.

[2] A. Kratzer, Conditional necessity and possibility, in: R. Bäuerle, U. Egli, A. von Stechow (Eds.), Semantics from Different Points of View, in: SSLAN, vol. 6, Springer, 1979, pp. 117–147.

[3] B. Bouchon-Meunier, G. Coletti, C. Marsala, Conditional possibility and necessity, in: B. Bouchon-Meunier, J. Gutiérrez-Ríos, L. Magdalena, R. Yager (Eds.), Technologies for Constructing Intelligent Systems 2: Tools, in: STUDFUZZ, vol. 90, Physica-Verlag, Heidelberg, 2002, pp. 59–71.

[4] S. Kraus, D. Lehmann, M. Magidor, Nonmonotonic reasoning, preferential models and cumulative logics, Artif. Intell. 44 (1990) 167–207.

[5] D. Lehmann, M. Magidor, What does a conditional knowledge base entail?, Artif. Intell. 55 (1992) 1–60.

[6] C. Boutilier, Conditional logics of normality: a modal approach, Artif. Intell. 68 (1) (1994) 87–154.

[7] B. Hansson, An analysis of some deontic logics, Noûs 3 (1969) 373–398.

[8] D. Makinson, L. van der Torre, Input/output logics, J. Philos. Log. 29 (4) (2000) 383–408.

[9] S. Beller, G. Kuhnmünch, What causal conditional reasoning tells us about people's understanding of causality, Think. Reasoning 13 (4) (2007) 426–460.

[10] L. Giordano, C. Schwind, Conditional logic of actions and causation, Artif. Intell. 157 (1) (2004) 239–279.

[11] G. Schurz, Probabilistic semantics for Delgrande's conditional logic and a counterexample to his default logic, Artif. Intell. 102 (1) (1998) 81–95.

[12] P. Snow, Diverse confidence levels in a probabilistic semantics for conditional logics, Artif. Intell. 113 (1) (1999) 269–279.

[13] J. Hawthorne, D. Makinson, The quantitative/qualitative watershed for rules of uncertain inference, Stud. Log. 86 (2) (2007) 247–297.

[14] W. Starr, Counterfactuals, in: E.N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy, summer 2021 edition, Metaphysics Research Lab, Stanford University, 2021.

[15] D. Lewis, Counterfactuals, Blackwell, 1973.

[16] D.G. Bobrow, Editor's preface, Artif. Intell. 13 (1) (1980) 1–4, https://doi.org/10.1016/0004-3702(80)90009-0, special Issue on Non-Monotonic Logic.

[17] M.L. Ginsberg, Counterfactuals, Artif. Intell. 30 (1) (1986) 35–79, https://doi.org/10.1016/0004-3702(86)90067-6.

[18] H. Rott, Difference-making conditionals and the relevant Ramsey test, Rev. Symb. Log. 15 (1) (2022) 133–164, https://doi.org/10.1017/S1755020319000674.

[19] G. Warren, M.T. Keane, R.M.J. Byrne, Features of explainability: how users understand counterfactual and causal explanations for categorical and continuous features in XAI, in: J. Heyninck, T. Meyer, M. Ragni, M. Thimm, G. Kern-Isberner (Eds.), Proceedings of the Workshop on Cognitive Aspects of Knowledge Representation (CAKR), in: CEUR Workshop Proceedings, vol. 3251, CEUR-WS.org, 2022.

[20] R.M.J. Byrne, Counterfactuals in explainable artificial intelligence (XAI): evidence from human reasoning, in: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI), 2019, pp. 6276–6282.

[21] X. Dai, M.T. Keane, L. Shalloo, E. Ruelle, R.M.J. Byrne, Counterfactual explanations for prediction and diagnosis in XAI, in: V. Conitzer, J. Tasioulas, M. Scheutz, R. Calo, M. Mara, A. Zimmermann (Eds.), Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES), ACM, 2022, pp. 215–226.

[22] D.M. Gabbay, Theoretical foundations for non-monotonic reasoning in expert systems, in: K.R. Apt (Ed.), Logics and Models of Concurrent Systems - Conference Proceedings, in: NATO ASI Series, vol. 13, Springer, 1984, pp. 439–457.

[23] W. Spohn, Ordinal conditional functions: a dynamic theory of epistemic states, in: H.L. William (Ed.), Causation in Decision, Belief Change, and Statistics, No. 2 in Proceedings of the Irvine Conference on Probability and Causation, Kluwer, Dordrecht, 1988, pp. 105–134.

[24] P. Gärdenfors, D. Makinson, Nonmonotonic inference based on expectations, Artif. Intell. 65 (2) (1994) 197–245.
[25] D. Lehmann, Another perspective on default reasoning, Ann. Math. Artif. Intell. 15 (1) (1995) 61–82.
[26] 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.
[27] E. Weydert, System JLZ - rational default reasoning by minimal ranking constructions, J. Appl. Log. 1 (3–4) (2003) 273–308.
[28] L. Giordano, V. Gliozzi, N. Olivetti, G. Pozzato, A minimal model semantics for nonmonotonic reasoning, in: L. Fariñas del Cerro, A. Herzig, J. Mengin (Eds.), Proceedings of the 13th European Conference on Logics in Artificial Intelligence (JELIA), in: LNCS, vol. 7519, Springer, 2012, pp. 228–241.
[29] 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.
[30] R. Booth, G. Casini, T. Meyer, I. Varzinczak, On rational entailment for propositional typicality logic, Artif. Intell. 277 (2019) 103–178.
[31] G. Casini, T. Meyer, I. Varzinczak, Taking defeasible entailment beyond rational closure, in: F. Calimeri, N. Leone, M. Manna (Eds.), Proceedings of the 16th European Conference on Logics in Artificial Intelligence (JELIA), in: LNCS, vol. 11468, Springer, 2019, pp. 182–197.
[32] J. Pearl, System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning, in: Proceedings of the 3rd Conference on Theoretical Aspects of Rationality and Knowledge (TARK), 1990, pp. 121–135.
[33] L.C. Hill, J.B. Paris, When maximizing entropy gives the rational closure, J. Log. Comput. 13 (1) (2003) 51–68.
[34] 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) 1–46.
[35] G. Kern-Isberner, Conditionals in Nonmonotonic Reasoning and Belief Revision, LNCS, vol. 2087, Springer, 2001.
[36] G. Casini, T. Meyer, K. Moodley, R. Nortjé, Relevant closure: a new form of defeasible reasoning for description logics, in: E. Fermé, J. Leite (Eds.), Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA), in: LNCS, vol. 8761, Springer, 2014, pp. 92–106.
[37] G. Casini, T. Meyer, I. Varzinczak, Rational defeasible belief change, in: D. Calvanese, E. Erdem (Eds.), Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning (KR), 2020.
[38] C. Alchourrón, P. Gärdenfors, D. Makinson, On the logic of theory change: partial meet contraction and revision functions, J. Symb. Log. 50 (1985) 510–530.
[39] M. Freund, Preferential reasoning in the perspective of Poole default logic, Artif. Intell. 98 (1998) 209–235.
[40] G. Casini, U. Straccia, Rational closure for defeasible description logics, in: T. Janhunen, I. Niemelä (Eds.), Proceedings of the 12th European Conference on Logics in Artificial Intelligence (JELIA), in: LNCS, vol. 6341, Springer-Verlag, 2010, pp. 77–90.
[41] R. Booth, E.L. Fermé, S. Konieczny, R.P. Pérez, Credibility-limited improvement operators, in: T. Schaub, G. Friedrich, B. O'Sullivan (Eds.), Proceedings of the 21st European Conference on Artificial Intelligence (ECAI), vol. 263, IOS Press, 2014, pp. 123–128.
[42] G. Casini, T. Meyer, Belief change in a preferential non-monotonic framework, in: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence (IJCAI), 2017, pp. 929–935.
[43] G. Casini, E. Fermé, T. Meyer, I. Varzinczak, A semantic perspective on belief change in a preferential non-monotonic framework, in: M. Thielscher, F. Toni, F. Wolter (Eds.), Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference (KR), AAAI Press, 2018, pp. 220–229.
[44] A. Bikakis, G. Antoniou, Defeasible contextual reasoning with arguments in ambient intelligence, IEEE Trans. Knowl. Data Eng. 22 (11) (2010) 1492–1506.
[45] C. Ghidini, F. Giunchiglia, Local models semantics, or contextual reasoning=locality+compatibility, Artif. Intell. 127 (2) (2001) 221–259.
[46] L. Serafini, M. Homola, Contextualized knowledge repositories for the semantic web, J. Web Semant. 12 (C) (2012) 64–87.
[47] S. Klarman, V. Gutiérrez-Basulto, Description logics of context, J. Log. Comput. 26 (3) (2013) 817–854.
[48] R.P. Pérez, C. Uzcátegui, Jumping to explanations versus jumping to conclusions, Artif. Intell. 111 (1–2) (1999) 131–169.
[49] K. Britz, I. Varzinczak, Rationality and context in defeasible subsumption, in: F. Ferrarotti, S. Woltran (Eds.), Foundations of Information and Knowledge Systems (FoIKS), Springer, Cham, 2018, pp. 114–132.
[50] K. Britz, I. Varzinczak, Contextual rational closure for defeasible $\mathcal{ALC}$, Ann. Math. Artif. Intell. 87 (1–2) (2019) 83–108.
[51] L. Giordano, V. Gliozzi, Reasoning about multiple aspects in DLs: semantics and closure construction, CoRR, arXiv:1801.07161 [abs].
[52] P. Bonatti, Rational closure for all description logics, Artif. Intell. 274 (2019) 197–223.
[53] G. Casini, U. Straccia, Defeasible inheritance-based description logics, J. Artif. Intell. Res. 48 (2013) 415–473.