Another perspective on Default Reasoning *

Daniel Lehmann
Institute of Computer Science,
Hebrew University, Jerusalem 91904 (Israel)

Abstract

The lexicographic closure of any given finite set D of normal defaults is defined. A conditional assertion $a \mid\sim b$ is in this lexicographic closure if, given the defaults D and the fact a, one would conclude b. The lexicographic closure is essentially a rational extension of D, and of its rational closure, defined in a previous paper. It provides a logic of normal defaults that is different from the one proposed by R. Reiter and that is rich enough not to require the consideration of non-normal defaults. A large number of examples are provided to show that the lexicographic closure corresponds to the basic intuitions behind Reiter's logic of defaults.

1 Plan of this paper

Section 2 is a general introduction, describing the goal of this paper, in relation with Reiter's Default Logic and the program proposed in [12] by Lehmann and Magidor. Section 3 first discusses at length some general principles of the logic of defaults, with many examples, and, then, puts this paper in perspective relatively to previous work. Section 4 sets the stage for this paper by describing the intuitive meaning of default information and the formal representation used in this paper for defaults. It singles out two different possible interpretations for defaults: a *prototypical* and a *presumptive*

*This work was partially supported by the Jean and Helene Alfassa fund for research in Artificial Intelligence

reading. Section 5 briefly discusses the relation between defaults and material implications. This paper proposes a meaning to any set D of defaults. This meaning is presented in a complex construction, that is described in full only in Section 9. The different aspects of this construction are presented separately at first. In Section 6, the meaning of a set consisting of a single default will be studied. Reiter's proposal does not enable the use of a default the antecedent of which is not known to hold. The new perspective of this paper allows much more sophisticated ways of using default information. In particular the default $(a : b)$ may be used to conclude that, if b is known to be false, then a should be presumed to be false too. Section 7 is a short digression on non-normal defaults. It is shown that such defaults can never be understood if one requires that the closure of a set of defaults be rational. Section 8 studies interacting normal defaults that have the same rank (or strength). We propose that, in the case of contradictory defaults of the same rank, we try to satisfy as many as possible. This proposal is in disagreement with D. Poole's [15], but in agreement with the Maximal Entropy approach of [6]. It is shown that this idea guarantees rationality. In Section 9 a formal description of our complete proposal is given. First, a model-theoretic construction is presented: given a finite set D of normal defaults, a modular model is defined and the lexicographic closure of D is the rational consequence relation defined by this model. Then, an equivalent characterization in terms of maxiconsistent sets is given. Section 10 presents examples and the description of the answer provided by our proposal. One of those shows how and why this proposal disagrees with the Maximal Entropy approach. Section 11 is a concluding discussion.

2 Introduction

In [16], R. Reiter proposed a formal framework for Default Reasoning. Its focal point is the definition of an extension. In [17], R. Reiter and G. Criscuolo found that, in this framework, one must consider non-normal defaults. Non-normal defaults have, since then, been taken as the basic piece of *default* information by the logic programming community. An alternative point of view is propounded here. An answer is provided to the following question: given a set D of normal defaults, what are the normal defaults that should be considered as following from D, or entailed by D? This answer provides a Logic of Defaults that does not suffer from the problems of multiple extensions or the inabilty of Reiter's system to cope satisfactorily with disjunctive information. There is no need to consider non-normal defaults. In [12], M. Magidor and the author proposed, as their first thesis [Thesis 1.1 there], that the set of defaults entailed by any set D be rational. This requirement is met. The rational closure of a set D, defined there, is not the set looked for, since it does not provide for inheritance of generic properties to exceptional subclasses. In their second thesis [Thesis 5.25 in Section 5.9], they proposed to look for some uniform way of constructing a rational superset of the rational closure of a knowledge base. The answer provided here, the lexicographic closure, is almost such a set, and a simple variation meets the condition in full. Independently, in [1], Benferhat&al. proposed a similar lexicographic construction based on an unspecified ordering of single defaults. When one applies their construction to the ordering on single defaults defined in [12], one obtains the lexicographic closure presented in this paper. Its computational complexity has been studied in [2] and [9]: it is in Δ_2^p and is NP-hard and co-NP-hard. The lexicographic closure is a syntactic construction in the sense of [14], i.e., it is sensitive to the presentation of the default information.

3 Nonmonotonic Reasoning

3.1 The Rational Enterprise

We shall briefly summarize [8] and [12] and set up the stage. This section was prepared in collaboration with David Makinson. Some nonmonotonic inference relations are better behaved than others. In particular, there are some simple closure conditions that appear highly desirable: *reflexivity, left-logical-equivalence, right-weakening, and, or,* and *cautious monotonicity.* The family of relations that satisfy those properties is closed under intersection. Therefore, given a set $\mathbf{K}$ of ordered pairs (a, b) of formulas (which we shall write $a \mid\!\sim b$ to remind us that they are meant to be elements of an inference relation), there is a natural and convincing way of defining a distinguished superset of $\mathbf{K}$ that satisfies those conditions: simply put $\mathbf{K}^p$, called the preferential closure of $\mathbf{K}$, to be the intersection of all supersets of $\mathbf{K}$ that satisfy the six conditions above.

However, there are other desirable "closure" (in a broad sense) properties that are much more difficult to deal with. *Rational monotonicity* defines a family of relations that is not closed under intersection. Other desirable conditions appear to be incapable of a purely formal expression, but may be conveyed intuitively and are illustrable by examples. Because of their informal nature, their identification is not cut and dried, but four seem to be of particular interest:

1. the presumption of typicality,

2. the presumption of independence,

3. priority to typicality, and

4. respect for specificity.

There may be other desirable properties.

(1) The *presumption of typicality* begins where rational monotonicity leaves off. Suppose $p \mid\!\sim x \in \mathbf{K}$. By rational monotonicity, the *closure* of $\mathbf{K}$, $\mathbf{K}^+$, will contain either $p \wedge q \mid\!\sim x$ or $p \mid\!\sim \neg q$. But which? No guideline is given. The presumption of typicality (it may as well be called "a weak presumption of monotonicity") tells us that, in the absence of a convincing reason to accept the latter, we should prefer the former.

Example 1 if $\mathbf{K}$ has $p \mid\!\sim x$ as its sole element, there is no apparent reason why the relation $\mathbf{K}^+ \supseteq \mathbf{K}$ that we regard as "generated" by $\mathbf{K}$ should contain $p \mid\!\sim \neg q$. Hence, it should contain $p \wedge q \mid\!\sim x$. *Note*: In this and all examples, $p, q, r \ldots x, y, z$ are understood to be *distinct* atomic formulas, i.e., propositional variables.

(2) The *presumption of independence* is a sharpening of the presumption of typicality, and thus a stronger presumption of monotonicity. For, even if typicality is lost with respect to one consequent, we may still presume typicality with respect to another, "unless there is reason to the contrary".

Example 2 Suppose $\mathbf{K} = \{p \mid\!\sim x, p \mid\!\sim \neg q\}$. Presumption of typicality cannot be used to support $p \wedge q \mid\!\sim x$, since $\mathbf{K}^+$ is known to contain $p \mid\!\sim \neg q$. Presumption of independence tells us we should expect x to be independent of q, and therefore unaffected by the truth of q. Therefore, it tells us, we should accept $p \wedge q \mid\!\sim x$.

Example 3 Suppose $\mathbf{K} = \{p \mathrel{|\!\sim} x, p \wedge q \mathrel{|\!\sim} \neg x, p \mathrel{|\!\sim} y\}$. Notice that $p \mathrel{|\!\sim} \neg q$ is in $\mathbf{K}^p$, the preferential closure of $\mathbf{K}$, and, therefore, the presumption of typicality cannot convince us to accept $p \wedge q \mathrel{|\!\sim} y$. But, we should presume that x is independent from y, as there is no reason to think otherwise, and put $p \wedge q \mathrel{|\!\sim} y$ in the desired consequence relation $\mathbf{K}^+ \supseteq \mathbf{K}$.

Remark: The two conditions above may be interpreted as related to and strengthening the condition of rational monotonicity. The difference between rational monotonicity, on one hand, and the presumptions of typicality and independence is subtle, and may be easily overlooked. Rational monotonicity is a constraint on the product $\mathbf{K}^+ \supseteq \mathbf{K}$, whereas presumptions of typicality and independence are best understood as rough and partial guides to the construction of a desirable $\mathbf{K}^+$.

(3) *Priority to typicality* tells us that, in a situation of clash between two inferences, one of them based on the presumption of typicality, the other one based on the presumption of independence, then we should prefer the former. Two examples are provided now.

Example 4 Suppose $\mathbf{K} = \{p \mathrel{|\!\sim} x, p \wedge q \mathrel{|\!\sim} \neg x\}$. The presumption of typicality offers $p \wedge q \wedge r \mathrel{|\!\sim} \neg x$, since there is no compelling reason to accept $p \wedge q \mathrel{|\!\sim} \neg r$. The presumption of independence offers both $p \wedge q \wedge r \mathrel{|\!\sim} \neg x$ and $p \wedge q \wedge r \mathrel{|\!\sim} x$. It clearly would not be justified to draw both conclusions. Priority to typicality, tells us to prefer the former.

Example 5 Suppose $\mathbf{K} = \{p \mathrel{|\!\sim} x, \mathbf{true} \mathrel{|\!\sim} q, q \mathrel{|\!\sim} \neg x\}$. The presumption of independence, acting on the last assertion of $\mathbf{K}$, offers $q \wedge p \mathrel{|\!\sim} \neg x$. This is in conflict with $p \wedge q \mathrel{|\!\sim} x$ that is offerred by presumption of typicality, from the first assertion. Priority to typicality says we should prefer the latter conclusion.

(4) *Respect for specificity* tells us that, in case of clash between two presumptions, one of them based on an assertion with a more specific antecedent than the other, we should prefer the conclusion based on the more specific antecedent. This principle is generally accepted and has been discussed in the literature. It is somewhat difficult to formalize: what does “based on” mean? It is closely related to the priority to typicality principle described just above, but the exact relationship between those two principles still needs clarification. In examples 4 and 5, the priority given to typicality achieves

precisely the respect for specificity we are looking for. In example 4, we prefer to use $p \wedge q \mathrel{|\!\sim} \neg x$ to $p \mathrel{|\!\sim} x$ also because $p \wedge q$ is strictly more specific than p, i.e., $p \wedge q \models p$. In example 5, we prefer to use $p \mathrel{|\!\sim} x$ to $q \mathrel{|\!\sim} \neg x$ also because p is defeasibly more specific than q, since, from $\mathbf{true} \mathrel{|\!\sim} q$, we shall conclude $p \mathrel{|\!\sim} q$ by presumption of typicality, or, preferably, presumption of independence. Another, more technical, reason to view p as more specific than q is that the rank (the definition found in [12] is explained at the end of 9.1) of p is strictly greater than the rank of q.

Of course, along with the above principles, one should also not forget *avoidance of junk*: the desired $\mathbf{K}^+ \supseteq \mathbf{K}$ should avoid gratuitous additions (otherwise, e.g. the total relation would do). In other words, $\mathbf{K}^+$ should be, in some sense, "least" among the supersets of $\mathbf{K}$ satisfying the desired conditions. "Least" should certainly imply minimal in the set-theoretic sense, i.e., no strict subset is acceptable, but cannot mean "included in any acceptable superset", since our family is not closed under intersection.

In [12], a construction is given, that, given any (finite) set $\mathbf{K}$ of pairs $a \mathrel{|\!\sim} b$ provides a rational extension $\overline{\mathbf{K}}$ such that $\mathbf{K} \subseteq \mathbf{K}^p \subseteq \overline{\mathbf{K}} = \overline{\mathbf{K}}^p = \overline{\overline{\mathbf{K}}}$ that behaves well so far as the presumption of typicality and the respect for specificity are concerned. However it does not pay much heed to the presumption of independence. For example it does not legitimize the conclusion $p \wedge q \mathrel{|\!\sim} y$ in example 3 above. The purpose of the present paper is to propose a different construction that performs better in this last respect, whilst not losing satisfaction of the other formal and informal properties.

3.2 Related Work

R. Reiter's [16] was certainly one of the most influential papers in the field of knowledge representation. It proposed a way of dealing with *default* information. In short, it proposed to represent such information as normal defaults and to define the meaning of a set of normal defaults as the set of extensions it provides to any set of sentences. In a follow-up paper [17], R. Reiter and G. Criscuolo remarked that, in many instances, the simple-minded formalization of situations involving more than one (normal) default was not adequate: the extension semantics enforced some unexpected and undesirable consequences. They proposed to cure this problem by considering an extended class of defaults: semi-normal defaults.

In this paper, a different perspective on default reasoning is proposed.

Normal defaults are considered and sets of normal defaults are given a meaning that is different from the one proposed in [16]. With this meaning, the interactions between defaults are as expected and the consideration of non-normal defaults is superflous. This perspective is in line with the first thesis of [12], that requires a set of defaults to define a rational consequence relation. It is also almost in line with the second thesis of the same paper, that requires a set of defaults to define a consequence relation that extends the rational closure of the set of defaults, and a straightforward variation will be shown to extend rational closure. This goal of implementing Reiter's program, but with different techniques, is similar to David Poole's [15]. The present paper also shares some technical insights with Poole's. It may be considered also as a close relative of the maximal entropy approach of [6, 7], but the semantics proposed here is different from the one obtained from maximal entropy considerations. This paper is a descendant of [11]. The main ideas of the lexicographic construction proposed in this paper have been, independently, proposed in [1]. There, the initial ordering of single defaults was left for the user to chose. A specific ordering of single defaults is used here.

4 What is default information?

Default information is information about the way things usually are. The paradigmatical example of such information, that has been used by most researchers in the field, is *birds fly*. Syntactically, a default is a pair of propositions that will be written as $(a : b)$, where a and b are formulas (of a propositional calculus for this paper). Remember that only normal defaults are considered, so that $(a : b)$ is our notation for Reiter's $\frac{a:b}{b}$. The default $(\mathbf{true} : b)$ will written as $(: b)$. Given a set D of defaults representing some background information about the way things typically behave and a formula a representing our knowledge of the situation at hand, we shall ask what formulas should be accepted as presumably true. The meaning of a set of defaults D will therefore be understood as the set of pairs (conditional assertions) $c \mathrel{|\!\sim} d$ it *entails*, i.e., for which d should be presumed to be true if c is our knowledge about the specific situation, i.e. represents the conjunction of the facts we know to be true. It is probably reasonable to expect that the conditional assertion $c \mathrel{|\!\sim} d$ be entailed by a set D containing the default $(c : d)$, but this will be discussed in the sequel. Notice, that we may, as well, consider that a set of normal defaults entails a set a defaults, confusing "snake" ($\mid\!\sim$) and colon (:).

The meaning of defaults is a delicate affair and it will be discussed in depth now. In [17], a *prototypical* reading is proposed: *birds fly* being understood as *typical birds fly*. But, there is another possible reading: *birds are presumed to fly unless there is evidence to the contrary*. This second reading will be called the *presumptive* reading. The conclusions of this paper may be summarized in three sentences. The two readings above are *almost* equivalent when isolated defaults are concerned, they are *not* when sets of defaults are concerned. The rational closure construction of [12] is the correct formalization of the prototypical reading. The presumptive reading is the one intended by Default Logic and its formalization is the topic of this paper. The distinction between the two readings will be explained with an example. This example is formally equivalent to the *Swedes* example described informally in [12, page 4]. In this example, as in all other examples of this paper, the formulas appearing in the defaults will be represented by meaningless letters and not, as customary in the field, by meaningful sentences. The remark that logic, the study of deductive processes, may be concerned only with the form of the propositions, and not with their meaning, dates back to Aristotle. The use of semantically loaded formulas and the import of the reader's knowledge of the world may only hamper the study of the formal properties of nonmonotonic deduction (that should perhaps be called induction). When a given example is formally isomorphic to some well known folklore example (or at least to some possible formalization of it), it will be pointed out.

Example 6 (Swedes) Let p and q be different propositional variables. Let D be the set of two defaults: $\{(: p), (: q)\}$. Accepting D means that we believe, by default, that p is true, and also that q is true. Following the prototypical reading, then, *typically p is true* and *typically q is true*. Following the presumptive reading *p is presumed to be true unless there is evidence to the contrary* and *q is presumed to be true unless there is evidence to the contrary*. Suppose now that we have the information that $\neg p \vee \neg q$ is true, i.e., at least one of p or q is false.

Using the prototypical reading, we shall conclude that the situation at hand is *not* typical. In such a case none of our two defaults is applicable: *typically p is true, but this is not a typical situation*, and therefore we shall not conclude, even by default (i.e. defeasibly) that $p \vee q$ holds true.

Using the presumptive reading, on the contrary, we shall conclude that $p \vee q$ should be presumed to be true unless there is evidence to the contrary, and since there is no evidence of this sort, it should be presumed to be true. We should therefore presume that exactly one of p and q holds.

5 Default vs. Material Implication

A very natural feeling is that the meaning of any single default $(a : b)$ should be closely related to the meaning of the material implication $a \rightarrow b$. This last formula will be called the *material counterpart* of the default $(a : b)$. Similarly the meaning of a set of defaults D should be related to the meaning of the set of its material counterparts.

It turns out that, both in the rational closure construction of [12] and in the construction proposed in this paper, the meaning of a default $(a : b)$ (that is an element of the set D of defaults accepted by a reasoner) in the presence of knowledge c, either its material counterpart $a \rightarrow b$ or void (i.e. equivalent to a tautology: **true**). Both constructions may therefore be described by pinpointing, given specific information c, which of the defaults of D are meaningful. If D_c is this set and M_c the set of material counterparts of D_c, then d should be presumed true iff d is a logical consequence of $c \cup M_c$, i.e., $M_c, c \models d$. This semantics fits well into the *implicit content* framework proposed in [18].

6 Single Defaults

Let D be the singleton set $\{(a : b)\}$, where a and b are arbitrary formulas. We propose the following meaning to D:

- if the information at hand c is consistent with $a \rightarrow b$, i.e., $c \not\models a \wedge \neg b$ then the default is meaningful and d is presumed iff $c, a \rightarrow b \models d$

- otherwise, the default is meaningless and d is presumed iff $c \models d$.

An equivalent, more model-theoretic description, of the consequence relation determined by D is the following: the rational consequence relation that is defined by the modular model in which the propositional models are ranked in two levels: on the bottom level (the more normal one) all models that satisfy the material implication $a \rightarrow b$, on the top level all other models.

This is the most natural understanding of the default information *if a is true then b is presumably true*, and completely in line with D. Poole's [15] treatment of defaults. Notice, though, that it does not always agree with Reiter's treatment and only almost agrees with rational closure. If the information at hand c logically implies a, then the perspective proposed here agrees with Reiter's: d is presumed to be true iff d is an element of the unique extension of $(D, \{c\})$.

To see the difference with Reiter's treatment, suppose a and b are different propositional variables and consider c to be $\neg b$. The perspective defended here will support the claim that $\neg a$ should be presumed to be true, i.e., a should be presumed to be false. For Reiter, on the contrary, there is a unique extension: $\mathcal{C}n(\neg b)$ ($\mathcal{C}n$ is the logical consequence operator) and therefore we should not presume anything about a. Similarly if c is **true**, the present perspective will support $a \rightarrow b$, whereas Reiter will not.

The comparison with rational closure is more subtle. Our perspective agrees with rational closure except when $a \models \neg b$. This is quite an out of the ordinary situation: a is logically equivalent to something of the form $\neg b \wedge e$, and the default is of the form *if b is false and e is true, then assume b is true*. Such a default will probably never be used in practice, but its consideration is nevertheless enlightening. In such a situation, the present perspective claims that the meaning of the default $(\neg b \wedge e : b)$ is that all models that satisfy $\neg b \wedge e$ are on the top level. In other terms, if $c \not\models \neg b \wedge e$, the default is meaningful and means $\neg b \wedge e \rightarrow b$, which is logically equivalent to $e \rightarrow b$, but if $c \models \neg b \wedge e$, then the default is meaningless. The treatment of this last case is well in line with the presumptive reading: if b is known to be false, then b should not be assumed to be true. If we look at the way rational closure deals with this case, we see that it agrees with the present perspective in the first case (i.e., if $c \not\models \neg b \wedge e$) but disagrees with it in the second case. Rational closure accepts any conclusion from the information that e is true and b is false. This is in line with the prototypical reading of the default: *if e is true and b is false, then typically b is true* may only mean that it is inconsistent for $e \wedge \neg b$ to be true and therefore one should conclude anything when this happens.

The new perspective does not always support each member of the rational closure, but the reader may check that the solution it supports is always rational (in the technical sense of [12]). A proof of this, in a more general setting, will be given in Section 9. How come our proposal is different from rational closure, that seemed to be the only reasonable one? Let $\mathbf{K}$ be the conditional knowledge base containing the single assertion $e \wedge \neg b \mid\!\sim b$. The rational relation proposed here in place of the rational closure does not contain $e \wedge \neg b \mid\!\sim b$. It is not an extension of $\mathbf{K}$ and therefore does not satisfy Thesis 5.25: "The set of assertions entailed by any set of assertions $\mathbf{K}$ is a rational superset of the rational closure of $\mathbf{K}$" of [12]. This departure from Thesis 5.25 is not central to our proposal and a slight variant of it would satisfy Thesis 5.25 by treating differently only those useless defaults discussed just above. This variant does not seem to fully fit the presumptive reading of defaults, though. If we denote by $\mathbf{K}^l$ the (lexicographic) construction proposed in this paper, the variant we have in mind may be defined in the following way: accept $a \mid\!\sim b$ iff either a has a rank for $\mathbf{K}$ and $a \mid\!\sim b \in \mathbf{K}^l$, or a has no rank. This variant gives a superset (sometimes strict) of $\mathbf{K}^l$, that is also a superset (sometimes strict) of the rational closure $\overline{\mathbf{K}}$.

7 Seminormal defaults

This paper will show that, if one accepts a semantics that is different from Reiter's, the reasons that compelled him to introduce non-normal defaults disappear, and one may restrict oneself to normal defaults. The reader may well ask whether one would not like to consider, anyway, a more general form of defaults: the semi-normal defaults. A semi-normal default $\frac{a : e \wedge b}{b}$ means that *if a is known to be true and there is no evidence that $e \wedge b$ is false, then b should be presumed to be true.* Let a be a tautology, i.e., **true** and e and b be different propositional variables (q and p respectively). Suppose we accept the semi-normal default $\frac{: q \wedge p}{p}$. There is general agreement about the following points:

- if the information at hand c is a tautology, i.e., we have no specific information, we should presume that p is true, since there is no evidence that $q \wedge p$ does not hold,
- if c is $\neg q$, i.e., we know for sure that q does not hold, we should not use the default information and therefore we should *not* presume p,

- but, if we have no specific information, we should *not* presume that q holds (why should we?).

The three points above provide a counter-example to the rule of Rational Monotonicity of [12]: we accept $\mathbf{true} \mathrel{|\!\sim} p$, but neither $\mathbf{true} \mathrel{|\!\sim} \neg\neg q$, nor $\neg q \mathrel{|\!\sim} p$. Even the simplest isolated non-normal default cannot be given a rational interpretation. This remark is very important in view of the fact that the efforts to harness Logic Programming to Nonmonotonic Reasoning take as their basic component rules of the form

$$a \leftarrow b\,, \neg c$$

meaning *conclude a if b has been concluded and c cannot be concluded.* This is essentially equivalent to considering the semi-normal default

$$\frac{b : a \wedge \neg c}{a},$$

or to considering the not even semi-normal default

$$\frac{b : \neg c}{a}$$

and will lead to a consequence relation that is not rational. All we have shown here is that non-normal defaults or the logic programming approach to nonmonotonic reasoning are incompatible with the property of rational monotonicity, which is central to this and previous papers.

8 Competing but equal defaults

After dealing, in Section 6, with single defaults, we shall treat now the more interesting case of a set of interacting normal defaults. In general, given a set of defaults D, this set defines a ranking of the defaults, as explained in [12]. This ranking will be described in full in Section 9.1. The ranking of a default $(a : b)$ relative to D depends only on its antecedent a and, as we shall see in Section 10, defaults of higher ranking (they correspond to exceptions) should be considered stronger than those of lower ranking. In this Section, we shall deal with the case all defaults have the same rank, i.e., all defaults are equal in strength and none of them correspond to an exception. This happens only when all elements of D have rank zero, as will be clear in Section 9.1. It is clear that, when considering such defaults, we should always assume that as many defaults as possible are satisfied (i.e. not violated). We should therefore always prefer violating a smaller set of defaults to violating a larger one. One may hesitate about the meaning to be given to "smaller": set inclusion or smaller size. The main conclusion of our considerations will be that sets of defaults should (for rationality's sake) be compared by their size, not by set inclusion.

We choose an example isomorphic to the musicians example of [4, Section 4.4], but we shall first ask about it questions that are different from those asked traditionally.

Example 7 (Musicians) Let p, q and r be different propositional variables. Let D contain the three following defaults: $\{(: p), (: q), (: r)\}$. In other words, p, q and r are assumed to hold by default. If we learn that $c \stackrel{\mathrm{def}}{=} \neg p \wedge \neg q \vee \neg r$ holds, i.e., that either both p and q are false, contrary to expectation, or, also contrary to expectation, r does not hold, what should we assume? Should one of the two possibilities $(\neg p \wedge \neg q)$ and $\neg r$ be assumed more likely than the other one?

In [15], D. Poole claims we should not. He claims there are two different maximal subsets of the material counterpart of D consistent with c (two bases for c): $\{p, q\}$ and $\{r\}$ and he proposes that we presume true only those formulas that are both in $\mathcal{C}n(r, c)$ and in $\mathcal{C}n(p, q, c)$. In particular, we should not presume p to be true. But, we should presume $p \leftrightarrow q$ to hold. Also, if we learn that $c \wedge \neg p$ holds, we should presume true only those formulas that are both in $\mathcal{C}n(r, c, \neg p) = \mathcal{C}n(r, c)$ and in $\mathcal{C}n(q, c, \neg p)$. In particular, we should not presume the truth of $p \leftrightarrow q$. D. Poole's proposal, therefore, does not satisfy the Rational Monotonicity principle.

Guided by Thesis 1.1 of [12], that requires Rational Monotonicity, a slight modification of Poole's ideas will be put forward now. This modification is also supported by the Maximal Entropy approach of [6]. The two bases above should not be considered equivalently plausible. The larger one, which contains two defaults should be considered more plausible than the one containing only a single default. In other terms, situations that violate two defaults should be considered less plausible than those that violate only one default. Here is a model-theoretic description. We shall consider the (propositional) models of our language, and rank them by the number of defaults of D they violate. A model m violates a default $a \mathrel{|\!\sim} b$ iff it does not satisfy the material implication $a \rightarrow b$, i.e., iff $m \models a \wedge \neg b$. The most normal models are those that violate no default of D: they constitute the bottom level (zero) of our modular model. Slightly less normal are those models that violate one single default: they constitute level one of our model. In general, level i is constituted by all models that violate exactly i members of D. The nonmonotonic consequence relation defined by this model is the one defined by D. It is rational, since the model described is ranked and consequence relations defined by modular models are rational (Lemma 3.9 of [12]). Coming back to the Musicians example: if we learn that $c \stackrel{\rm def}{=} \neg p \wedge \neg q \vee \neg r$ holds, $\neg r$ should be presumed true. We should therefore presume p to be true.

We provided, just above, a model-theoretic description of our proposal. An equivalent description in terms of "bases", in the spirit of [15], is provided now. The same bases were also considered in [1]. Let $E = \{e_i\}$ be a finite set of formulas. Let c be a formula.

Definition 1 *A subset F of E is said to be a maxbase for c iff c is consistent with F and there is no subset F' of E, $\mid F' \mid > \mid F \mid$ that is consistent with c.*

Theorem 1 *Let D be a set of defaults such that all elements of D have rank zero (with respect to D). Let E be the set of material implications corresponding to the defaults of D. The consequence relation defined by the model-theoretic description above is characterized by:*

$$a \mathrel{|\!\sim} b \text{ iff for every maxbase } F \ of \ E \text{ for } a \quad F, a \models b. \tag{1}$$

.

Theorem 1 implies that the relation defined by Equation (1) is rational.
Proof: First, some remarks. Let n be the size of the set D.

1. If a is satisfied by some model of level i ($0 \leq i \leq n$), then, there is a maxbase for a, and all maxbases for a are of size larger or equal to $n - i$.

2. If F is a maxbase of size k ($0 \leq k \leq n$) for a, then, there is a model of level $n - k$ that satisfies a, and no model of lower level satisfies a.

3. There is no maxbase for a iff a is a logical contradiction.

Suppose $a \mathrel{|\!\sim} b$. If no model satisfies a, a is a logical contradiction and $X, a \models b$ for any X. Suppose, then, that a is satisfied by some model of level i, but by no model of lower level. Any model of level i that satisfies a, satisfies b, by hypothesis. Let F be a maxbase for a. By remark (2), $n - \mid F \mid = i$. Any model that satisfies F is obviously of level less or equal to $n - \mid F \mid = i$. Any model that satisfies F and a is therefore of level i and satisfies b, by hypothesis. We conclude that $F, a \models b$.

Suppose, now, that for any maxbase F for a we have $F, a \models b$. If there is no maxbase for a, then, by remark (3), a is a logical contradiction and $a \mathrel{|\!\sim} b$. Suppose, then, the maxbases for a are of size k. There is, by remark (2), a model of a of level $n - k$, and there is no model of level less than $n - k$ that satisfies a. We must show that any model of a of level $n - k$ satisfies b. Let m be such a model. Since m violates $n - k$ defaults, it satisfies a set M of k defaults. But M is consistent with a, since $m \models a$. The size of M is the size of the maxbases for a, therefore M is a maxbase for a and, since $M, a \models b$, we conclude that $m \models b$. ∎

Example 8 (Musicians, continued) We shall now describe our solution to the questions traditionally asked about the musicians' example and generally used to demonstrate that counterfactuals do not satisfy Rational Monotonicity. Suppose our specific information is $c \stackrel{\text{def}}{=} p \wedge \neg r \vee \neg p \wedge r$. There are two maxbases: $\{p, q\}$ and $\{q, r\}$. We shall therefore presume that q holds and we shall *not* presume that $d \stackrel{\text{def}}{=} q \wedge r \vee \neg q \wedge \neg r$ holds. This is the common wisdom and the present perspective subscribes to it.

Suppose now that our specific information is $c \wedge \neg d$, or, equivalently, $p \wedge q \wedge \neg r \vee \neg p \wedge \neg q \wedge r$. The common wisdom, defended in [4], would like to convince us that we should not presume q to be true. The position defended here, presumes that q is true (and also p and $\neg r$) because this situation violates only one default ($: r$) whereas the other possible situation: $\neg p \wedge \neg q \wedge r$ violates two defaults.

The reader may suspect that our policy gives results that are extremely sensitive to the way the defaults are presented. Indeed, the way defaults are presented is important, and our perspective on defaults does not enjoy the nice global properties of rational closure described in [12, Section 5.5] that make it invariant under the addition or deletion of entailed defaults. Two examples of this phenomenon will be described now. The first one shows that the addition to D of a default entailed by D may add new conclusions. The second one shows that the addition to D of a default entailed by D may force us to withdraw previous conclusions. The examples presented are very simple and natural and should convince the reader that any presumptive reading of defaults leads to a high sensitivity to the presentation of the default information. This sensitivity is, probably, a drawback of the lexicographic closure. The following examples should convince the reader that this problem is inevitably brought about by a presumptive understanding of defaults. If we had decided to consider multisets of defaults, instead of sets, thus allowing certain (stronger) defaults to appear a number of times in D, our construction would had been sensitive to the number of times each default appears in D.

Example 9 (Adding entailed defaults may add conclusions) Let D be the singleton $\{(: p \wedge q)\}$. The default (identifying defaults and conditional assertions) $(: p)$ is obviously entailed by D. But the default $(\neg p \vee \neg q : p)$ is not entailed by D, the antecedent being inconsistent with the only default of D. Nevertheless $(\neg p \vee \neg q : p)$ is entailed by the set $\{(: p \wedge q), (: p)\}$, since its antecedent is consistent with the second default. The behavior of the corresponding Poole system is the same.

Example 10 (Adding entailed defaults may delete conclusions) Let D be the set $\{(: p), (: q)\}$. Both defaults $(: p \leftrightarrow q)$ and $(\neg p : q)$ are entailed by D. But $(\neg p : q)$ is *not* entailed by the set $\{(: p), (: q), (: p \leftrightarrow q)\}$, since the antecedent is consistent with both the last defaults separately but not together, and there are therefore two maxbases. In this case also, the behavior of the corresponding Poole systems is the same.

9 Lexicographic closure

9.1 Introduction and definition

In the previous Section we discussed the treatment of conflicting defaults that had the same precedence. We shall now treat arbitrary conflicting defaults, and define the construction we propose in full generality. We must take into account the fact that defaults may have different weight, or precedence. Fortunately, the correct definition of the relative precedence of defaults has been obtained in a previous work. Given a finite set of defaults D, the precedence of a default is given by its rank (higher rank means higher precedence), i.e., by the rank of its antecedent as defined in [12, Section 2.6]. The definition presented here is equivalent to the original definition, by Corollary 5.22 there.

We shall now remind the reader of this definition. Let D be a finite set of defaults, and $\tilde{D}$ the set of its material counterparts. Let a be a formula. We shall put $E_0 = D$. If a does not have rank less than i, but is consistent with $\tilde{E}_i$, it has rank i. The set E_{i+1} is the subset of E_i that contains all defaults $(a : b)$ of D for which a does not have rank less or equal to i. We shall put $D_i = E_i - E_{i+1}$, and let D_∞ be the set of all elements of D that have no rank, i.e., have infinite rank. Elements of D_∞ have precedence over all other defaults. Notice that, since D is finite, all D_i's are empty after a certain point, except possibly D_∞. There is a k such that for any i, $k \leq i < \infty$, $D_i = \emptyset$. The smallest such number k will be called the *order* of the set D. The set D may be partitioned into $D_\infty \oplus D_{k-1} \oplus D_{k-2} \oplus \ldots \oplus D_0$.

We remind the reader that the rational closure of the set D, defined in [10] and studied in depth in [12] (see Theorem 5.17 and Lemma 2.24) is the set of defaults $\overline{D}$ that consists of all defaults $(a : b)$ such that the rank of a is strictly less than the rank of $a \wedge \neg b$ (equivalently, the rank of $a \wedge b$ is strictly less than the rank of $a \wedge \neg b$), or such that a has no rank. We shall now define another closure for D, the lexicographic closure. We define the lexicographic closure by way of a modular model in which every model is ranked by the set of defaults it violates. A similar presentation may be used to define rational closure, it will also be described.

9.2 The model-theoretic description

As usual, we shall suppose a finite set D of defaults is given. We shall describe the consequence relation defined by D, the lexicographic closure of D, D^l as the consequence relation defined by a certain modular model, $\mathcal{M}_D$. To define this model, we need to order the propositional models by some modular ordering. We shall order the propositional models by ordering the sets of defaults (of D) that they violate: each model m violates a set $D_m \subseteq D$ of defaults. How should we order the subsets of D? Intuitively, we are looking for a "degree of seriousness". We prefer to violate a "lighter" set of defaults than a more serious one, i.e., a propositional model that violates a lighter set of defaults is more normal than a model that violates a more serious set. There are two criteria that must be taken into account when deciding which of two sets is more serious:

- the size of the set: the smaller the set, the less serious it is. We have seen in Section 8, that "smaller" should be taken here to mean "of smaller size", and
- the seriousness of the elements of the set, it is less serious to violate a less specific default than a more specific default, i.e. a default of lower rank than a default of higher rank.

The reader should notice here that our definition is in no way circular. The lexicographic closure is defined in terms of a specific modular model that is, in turn, defined in terms of the ranks of the formulas involved. These ranks have been defined above, by a straightforward inductive definition. The ranks of the formulas have a close relationship with the ordering of the modular model that defines the rational closure of D, but this is a different modular model. In fact, the model we are describing now is a refinement of the model that defines rational closure (a level may split into a number of sublevels). The next question, now, is how should we compose those two criteria? The principle of rationality will trace the way for us. We want a modular ordering on the subsets of D. Each one of the criteria above gives a modular ordering. Is there a general way to combine two modular orderings and obtain a modular ordering? Yes, a lexicographic (i.e. consider one criterion as the principal criterion, the other as secondary) composition of modular orderings is a modular ordering. Which of the two criteria above should be considered as the major criterion? Clearly the second one: specificity. We should prefer violating two defaults of low specificity to violating one of high specificity.

Example 11 Let $D = \{(: p), (: q), (: x), (y : \neg x), (y : r)\}$. Suppose our assumptions are $y \wedge (\neg p \wedge \neg q \vee \neg r)$. You may imagine that p, q and x are generic properties (of birds, say), and that y is a class of birds that are exceptional with respect to x. The property r is a generic property of y birds. Suppose we have a bird that is part of the y class, and is known to be, either exceptional with respect to two generic properties of birds, or exceptional with respect to one generic property of the sub-class y. Presumption of typicality (see Section 3.1), from the last default of D proposes the conclusion r (and therefore $\neg p$ and $\neg q$). Presumption of independence proposes the conclusions p and q (and therefore $\neg r$). Priority to typicality convinces us to accept the former and reject the latter.

Therefore, to decide which of two sets of defaults is more serious, we shall partition those sets into subsets of defaults of equal ranks and compare (by size) rank by rank, starting with the higher ranks. As soon (in terms of ranks) as a decision can be made, we decide and stop.

Definition 2 *Let D be a set of defaults and k its order. To every subset $X \subseteq D$ may be associated a $k+1$-tuple of natural numbers: $\langle n_0, \ldots, n_k \rangle$, where $n_0 = \mid D_\infty \cap X \mid$, $n_1 = \mid D_{k-1} \cap X \mid$, and in general, for $i = 1, \ldots, k$, $n_i = \mid D_{k-i} \cap X \mid$. In other terms, n_0 is the number of defaults of X that have no rank and, for $0 < i \leq k$, n_i is the number of defaults of X that have rank $k-i$. We shall order the subsets of D by the natural lexicographic ordering on their associated tuples. This is a strict modular partial ordering: it will be denoted by $\prec$, (the* seriousness ordering*).*

The seriousness ordering on sets of defaults is used to order the propositional models: $m \prec m'$ iff $V(m) \prec V(m')$, where $V(m) \subseteq D$ is the set of defaults violated by m. This modular ordering on models defines a modular preferential model, that, in turn defines a consequence relation, D^l, the lexicographic closure of D. The reader may check that all examples treated in this paper conform to the definition above.

Let us, now, before we give an alternative description of lexicographic closure, briefly digress to see that rational closure may be defined by a specific seriousness ordering, different from the one defined in Definition 2.

Definition 3 *Let X_1 and X_2 be subsets of a set D of defaults and k its order. Let n^1_i and n^2_i, for $i = 1, \ldots, k$, be the size of the partitions of X_1 and X_2 respectively. Let m^j be the smallest i such that $n^j_i \neq 0$, for $j = 1, 2$. We shall write $X_1 \ll X_2$ iff $m^1 > m^2$.*

Clearly $X_1 \ll X_2$ implies $X_1 \prec X_2$, i.e., $\ll$ is coarser than $\prec$.

Theorem 2 *Suppose a has a finite rank. The conditional assertion $a \mid\!\sim b$ is a member of the rational closure of D iff it is satisfied by the modular model* *in which each propositional model is ranked by the $\ll$ ordering on the set of defaults it violates.*

Proof: Suppose the rank of a is strictly less than that of $a \wedge \neg b$, and that m is a propositional model that satisfies a and is minimal among those for the $\ll$ ordering. If a has rank l, there is a model that satisfies a and violates no default of D of rank greater or equal to l. We conclude that m violates no such default and therefore satisfies no formula of rank strictly greater than l. The model m does not satisfy $a \wedge \neg b$, and therefore satisfies b.

Suppose, now, that all propositional models that satisfy a and are minimal in the $\ll$ ordering for that property, also satisfy b. Let the rank of a be k. Since there is a propositional model that satisfies a and violates no default of rank greater or equal to k, all models that satisfy a and violate no default of rank greater or equal to k satisfy b. We conclude that the rank of $a \wedge \neg b$ is greater than k. ∎

We may now show that the lexicographic closure is a superset of the rational closure, at least for defaults of finite rank, thus almost complying with Thesis 5.25 of [12].

Theorem 3 *If a has a finite rank and $a \mid\!\sim b$ is an element of the rational closure of D, then it is an element of its lexicographic closure.*

Proof: Suppose a has a finite rank and $a \mid\!\sim b$ is an element of the rational closure of D. By Theorem 2, $a \mid\!\sim b$ is satisfied in the modular model defined by $\ll$. It is therefore satisfied in any modular model defined by a finer relation. We noticed, just following Definition 3, that $\prec$ is such a finer relation. ∎

A characterization of the lexicographic closure in terms of bases will be described now.

9.3 Bases

Definition 4 *Let a be a formula, and B a subset of D. We shall say that B is a* basis *for a iff a is consistent with $\tilde{B}$, the material counterpart of B, and, B is maximal* with respect to the seriousness ordering *for this property.*

The following lemma may help explain the structure of bases, but is not used in the sequel.

Lemma 1 *If a has rank i (a has no rank is understood as a having an infinite rank) and B is a basis for a, then, for any $j \geq i$, $D_j \subseteq B$.*

In other terms, any basis for a is full, for all indexes larger or equal to the rank of a.
Proof: Since a has rank i, for any $j \geq i$, a is consistent with E_j and therefore with $D_j \cup B \cap E_{j+1} \subseteq E_j$. ∎

Theorem 4 *The default $(a : b)$ is in D^l, the lexicographic closure D iff, for any basis B for a, $\tilde{B}, a \models b$.*

Proof: The proof is a generalization of that of Theorem 1. Let k be the order of D. Let $d_i = \mid D_i \mid$ for $i = 0, \ldots, k-1, \infty$.

1. If a is satisfied by some model of seriousness level $(i_0, \ldots, i_k)$, then, there is a basis for a of level $(d_\infty - i_0, d_{k-1} - i_1, \ldots, d_0 - i_k)$, and all bases for a have this level.

2. If B is a basis of level $(l_0, \ldots, l_k)$ for a, then, there is a model of level $(d_\infty - l_0, d_{k-1} - l_1, \ldots, d_0 - l_k)$ that satisfy a, and no model of strictly smaller seriousness satisfy a.

3. There is no basis for a iff a is a logical contradiction.

Suppose, first, that $a \mathrel{|\!\sim} b \in D^l$. If no model satisfies a, a is a logical contradiction and $X, a \models b$ for any X. Suppose, then, that a is satisfied by some model of level $(i_0, \ldots, i_k)$, but by no model of lower level. Any model of level $(i_0, \ldots, i_k)$ that satisfies a, satisfies b, by hypothesis. Let B be a basis for a, of seriousness $(b_\infty, \ldots, b_0)$. By remark (2), $d_j - b_{k-j} = i_{k-j}$, for $j = 0, \ldots, k-1$, and $d_\infty - b_0 = i_0$. Any model that satisfies B is obviously of seriousness level less or equal to $(d_\infty - b_0, d_{k-1} - b_1, \ldots, d_0 - b_k)$, i.e., of level less or equal $i_0, \ldots, i_k$. Any model that satisfies B and a is therefore of level $i_0, \ldots, i_k$, and satisfies b, by hypothesis. We conclude that $B, a \models b$.

Suppose, now, that for any basis B for a we have $B, a \models b$. If there is no basis for a, then, by remark (3), a is a logical contradiction and $a \mathrel{|\!\sim} b \in D^l$. Suppose, then, the bases for a are of seriousness $(b_0, \ldots, b_k)$. There is, by remark (2), a model of a of level $l = (d_\infty - b_0, \ldots, d_0 - b_k)$, and there is no model of level less than l that satisfies a. We must show that any model of a of level l satisfies b. Let m be such a model. Since m violates $d_j - b_{k-j}$

defaults of rank j, it satisfies a set M containing b_{k-j} defaults of rank j. But M is consistent with a, since $m \models a$. The seriousness of M is $(b_0, \ldots, b_k)$, the seriousness of the bases for a. Therefore M is a basis for a and, since $M, a \models b$, we conclude that $m \models b$.

∎

We shall now describe the lexicographic closure of a number of sets of defaults, some of them well-known from the literature.

10 Examples

In this Section, motivating examples will be described, indicating, for each of them the conclusions endorsed by the lexicographic closure. My goal is to, gradually, convince the reader, that each one of the decisions taken in the process of defining the lexicographic closure was reasonable. My goal is *not* to convince the reader that lexicographic closure provides the *intuitively* correct answer once the propositional variables have been interpreted in some manner that is well-known in the folklore of the field, because I believe that, in most cases, we intuitively treat interpreted formulas in a meaning-dependent manner, not in the formal, meaning-independent way that is the hallmark of logic. In other terms, once the variables are interpreted, there is no way of knowing whether the intuitive conclusions come from formal logical considerations or from world knowledge the reasoner has about the situations or objects the interpreted variables refer to. A first example exemplifies why one needs to give precedence to defaults describing exceptional cases over those that describe more normal cases.

Example 12 (Exceptions) Let $D = \{(: p), (: q), (\neg p : \neg q)\}$. Here p and q hold by default, and when p does not hold, then, by default q does not hold either. Suppose we know that p does not hold. Then, obviously we cannot use the first default. But, we could use the second one to conclude q or the third one to conclude $\neg q$. We obviously want to presume $\neg q$, and we need to say that the third default has precedence over the second one. Fortunately, it is not difficult to justify why the third default has precedence over the other ones. It (i.e., its antecedent) has rank one whereas the two other defaults have rank zero. This comes from the fact that the antecedents of the first two defaults (**true**), describe some unexceptional situation, while $\neg p$ describes an unexpected, exceptional situation since p is, by default, presumed to be true.

The technical description of the lexicographic closure follows. The order of D is two. The first two defaults of D have rank zero, the last one has rank one. The most normal models (those that have level zero), are those that satisfy p and q (and therefore $\neg p \rightarrow \neg q$). On level one, we find those models that satisfy $\neg p \rightarrow \neg q$ and violate exactly one of p or q: the models that satisfy p and $\neg q$. The third level contains those models that violate both p and q, but satisfy $\neg p \rightarrow \neg q$, i.e. the models satisfying $\neg p$ and $\neg q$. The fourth level contains all models violating $\neg p \rightarrow \neg q$, but satisfying p and q; it is empty. The fifth level contains all models violating $\neg p \rightarrow \neg q$, and exactly one of p and q: it contains one model. The sixth level is empty.

So, we must give precedence to defaults of higher rank over defaults of lower rank. Notice that rank is really all the difference between the default $(a : b)$ and the default $(: a \rightarrow b)$. The second one has always rank zero, while the first one may have a much higher rank (if a is presumed to be false) and is therefore more powerful. The reader may easily check that the first two defaults of D have rank zero, whereas the third one has rank one. Let us now treat the similar but more classical penguin example.

Example 13 (Penguins) Let $D = \{(p : q), (r : p), (r : \neg q)\}$. The default $(: \neg r)$ is entailed by D and the second and third defaults have rank one, whereas the first default has rank zero. The defaults $(r : \neg q)$ and $(p \wedge r : \neg q)$ are entailed by D whereas $(r : q)$ is not.

Technically, the rank of D is two. At level zero: all models satisfying $p \rightarrow q$ and $\neg r$. At level one: all models satisfying p and $\neq q$, and those satisfying r, $\neg p$ and $\neg q$. Level two is empty. Level three: all models satisfying r, q and p. Level four: all models satisfying r, q and $\neg p$. Level five: empty.

In the present proposal, this precedence of defaults of higher rank is in a sense (or in two ways) absolute: one should not trade the violation of a default of rank $n + 1$ for the violation of any number of defaults of rank less than or equal to n. On this point, the present proposal is in disagreement with the Principle of Maximal Entropy as proposed in [7]. The following example will exhibit this disagreement, it is not meant to support one construction against the other.

Example 14 (Winged Penguins) Let $D = \{(b : w), (b : f), (p : b), (p : \neg f)\}$. The default $(p \wedge (f \vee \neg w) : b)$ is entailed by D, since the only basis for

$p \wedge (f \vee \neg w)$ is the set $\{(p : b), (p : \neg f)\}$, containing all defaults of rank one. In fact, even $(p \wedge (f \vee \neg w) : b \wedge \neg f \wedge \neg w)$ is entailed by D. But the Principle of Maximal Entropy of [7] will consider as equivalent

- to violate two defaults of rank zero: $((b : w)$ and $(b : f)$, and
- to violate one default of rank one: $(p : b)$,

and therefore will not accept $(p \wedge (f \vee \neg w) : b)$. Notice that the set D is Minimal Core in the sense of [7].

Example 15 (Exceptions again) Let

$$D = \{(: r), (: p), (: q), (\neg p : \neg q), (\neg p : \neg r)\}.$$

Suppose our specific information is $\neg p \wedge q$, which means the situation is doubly exceptional: p is presumed true but is in fact false, and when p is false q is presumed false, but it is true. In other words the rank of $\neg p \wedge q$ is two. Should we presume r to be true or false? It is clear we should presume it false, since the default $(\neg p : \neg r)$ talks about a situation closer to the one at hand than the default $(: r)$, and should have precedence over it. But notice that our information shows that the situation described by $\neg p \wedge q$ is exceptional with respect to the one described by $\neg p$.

11 Discussion

The lexicographic closure D^l of a finite set D of defaults is defined by a modular model in which all propositional models appear, at some level. An assertion of the form $a \mid\sim \textbf{false}$ will appear in D^l only if a is a logical contradiction. In other terms, lexicographic closure is, in the terminology of [13], *consistency preserving*. This is indeed one of the hallmarks of Default Reasoning *à la Reiter* as attested by the discussion in [16], and in particular Corollary 2.2 there.

Since there has been a lot of discussion in the literature, in particular in [5] and [3], of the principle of Transitivity, it is probably worth a short discussion.

The question of Transitivity is: should we accept $(a : c)$ on the basis of the two defaults $(a : b)$ and $(b : c)$? The answer proposed here, and which follows from our construction, is that, if we have both $(a : b)$ and $(b : c)$, and if a and b are both of the same rank, then we should also accept $(a : c)$. Note that, if we accept $(a : b)$, then b has rank lower or equal to that of a: indeed if the rank of b was larger than that of a, the rational closure of our set of defaults would include $a \vee b \mathrel{|\!\sim} \neg b$ and therefore $a \mathrel{|\!\sim} \neg b$, and, by Theorem 3, we would accept $(a : \neg b)$ in the lexicographic closure. If the rank of a is strictly greater than that of b, then we are not guaranteed that$(b : c)$ will be part of all bases for a.

Given a finite set D of defaults and a default $(a : b)$, how difficult is it to decide whether $(a : b)$ is entailed by D? This decision seems to require the computation of the ranks of the defaults of D, but this is relatively easy: a quadratic number of satisfiabilty problems. It seems that it also requires the consideration of a possibly large number of subsets of D and seems therefore inherently exponential. The rational closure construction of [12] provides a quick and dirty approximation to this construction in the following sense: if a default belongs to the rational closure it is entailed (up to a slightly different treatment of formulas that have no rank). The case in which all defaults have a Horn structure, needs further study. One may perhaps avoid the exponential blow-up in this case.

12 Acknowledgements

Comments, corrections and suggestions by Michael Freund, Moisés Goldszmidt and David Makinson are gratefully acknowledged. Three anonymous referees (not necessarily disjoint from the previous set) also helped a lot.

References

[1] Salem Benferhat, Claudette Cayrol, Didier Dubois, Jérôme Lang, and Henri Prade. Inconsistency management and prioritized syntax-based entailment. In Ruzena Bajcsy, editor, *Proceedings of the 13th I.J.C.A.I.*, pages 640–645, Chambéry, Savoie, France, August 1993. Morgan Kaufmann.

[2] Claudette Cayrol and Marie-Christine Lagasquie-Schiex. Comparaison de relations d'inférence non monotone: Étude de complexité. Techni-

cal Report 93.23.R, Institut de Recherche en Informatique de Toulouse, September 1993.

[3] Michael Freund, Daniel Lehmann, and Paul H. Morris. Rationality, transitivity, and contraposition. *Artificial Intelligence*, 52(2):191–203, December 1991. Research Note.

[4] Matthew L. Ginsberg. Counterfactuals. *Artificial Intelligence*, 30:35–79, 1986.

[5] Matthew L. Ginsberg. *Readings in Nonmonotonic Reasoning*, chapter 1, pages 1–23. Morgan Kaufmann, Los Altos, CA, 1987.

[6] Moisés Goldszmidt, Paul H. Morris, and Judea Pearl. A maximum entropy approach to nonmonotonic reasoning. In *Proceedings of AAAI-90*, Boston, August 1990.

[7] Moisés Goldszmidt, Paul H. Morris, and Judea Pearl. A maximum entropy approach to nonmonotonic reasoning. *IEEE Transactions of Pattern Analysis and Machine Intelligence*, 1993. To Appear.

[8] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1–2):167–207, July 1990.

[9] Marie-Christine Lagasquie-Schiex, December 1993. Private Communication.

[10] Daniel Lehmann. What does a conditional knowledge base entail? In Ron Brachman and Hector Levesque, editors, *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, Canada, May 1989. Morgan Kaufmann.

[11] Daniel Lehmann. Another perspective on default reasoning. Technical Report TR 92-12, The Leibniz Center for Computer Science, Institute of Computer Science, Hebrew University, Jerusalem, July 1992. Presented at the Bar-Ilan Symposium on Foundations of AI, June 1993.

[12] Daniel Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, May 1992.

[13] David Makinson. General patterns in nonmonotonic reasoning. In D. M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume Vol. 3, Nonmonotonic and Uncertain Reasoning, pages 35–110. Oxford University Press, 1994.

[14] Bernhard Nebel. Belief revision and default reasoning: Syntax-based approaches. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning*, pages 417–428. Morgan Kaufmann, 1991.

[15] David Poole. A logical framework for default reasoning. *Artificial Intelligence*, 36:27–47, 1988.

[16] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[17] Raymond Reiter and Giovanni Criscuolo. Some representational issues in default reasoning. *International Journal of Computers and Mathematics with Applications*, 9:15–27, 1983.

[18] Robert C. Stalnaker. What is a nonmonotonic consequence relation? In *Fourth International Workshop on Nonmonotonic Reasoning*, Plymouth, Vermont, May 1992.