Representation Independence of Nonmonotonic Inference Relations

Manfred Jaeger
Max-Planck-Institut für Informatik
Im Stadtwald
66123 Saarbrücken
Germany
email: jaeger@mpi-sb.mpg.de

Abstract

A logical concept of representation independence is developed for nonmonotonic logics, including probabilistic inference systems. The general framework then is applied to several nonmonotonic logics, particularly propositional probabilistic logics. For these logics our investigation leads us to modified inference rules with greater representation independence.

1 INTRODUCTION

Entropy maximization is a rule for probabilistic inference for whose application to problems in artificial intelligence there exist several independent and very strong arguments (Grove, Halpern & Koller 1992),(Paris & Vencovská 1990). Unfortunately, though, there is a major drawback for which the maximum entropy inference rule has often been criticized: the result of the inference depends on how given information is represented.

The probably best known example used to illustrate this point is the "Life on Mars" example, a rendition of which may be given as follows: the belief that the probability for the existence of life on mars exceeds 0.6 may be expressed by the statement

$$\tau_1 :\equiv P(LoM) \geq 0.6,$$

using the vocabulary $\{LoM\}$ of propositional variables. Entropy maximization, applied to this chosen language and constraint τ_1 then yields $P(LoM) = 0.6$. Alternatively, we might choose the language $\{ALoM, PLoM\}$ containing propositional variables for "Animal Life on Mars" and "Plant Life on Mars", and express our belief by

$$\tau_2 :\equiv P(ALoM \vee PLoM) \geq 0.6.$$

Entropy maximization here yields $P(ALoM \vee PLoM) = 0.75$. Thus, the result of the maximum entropy inference rule is dependent on the choice of language, and seemingly equivalent statements yield different results.

Even though the charge of representation dependence against maximum entropy methods has been around for a long time, until recently there has neither been any precise explication of what representation independence actually is, nor a systematic investigation into the properties with regard to this property of other inference rules.

The first rigorous examination of these issues has been presented by Halpern and Koller (1995). They proceed from a definition of when two different structures, or state spaces, each equipped with a set of probability distributions, are alternative representations (namely, when on structure can be "faithfully embedded" in the other), and then call a probabilistic inference procedure representation independent when applied to the two structures it picks out corresponding preferred probability measures.

In this paper the question of representation independence will be tackled from a somewhat different perspective, and in a much wider context: first, rather than looking at specific structures and their embeddings, we here consider the purely logical question and ask: when do two sets of formulas represent the same information, and when is an inference relation defined on formulas of some language representation independent? From this syntactic perspective it is natural to extend the scope of the enquiry: a probabilistic inference rule, like entropy maximization, formally defines a nonmonotonic inference relation $\mid\!\sim$ that can be studied with respect to the same formal properties as have been investigated for nonmonotonic logics (Kraus, Lehmann & Magidor 1990), (Gabbay 1985). Conversely, a concept of representation independence, framed entirely in terms of formulas and entailment relations, may be applied to a large class of nonmonotonic logics, not only probabilistic ones. In this paper we will develop the necessary tools for the investigation of representation independence of nonmonotonic inference relations, and take some steps towards clarifying the degree of representation (in-)dependence of existing nonmonotonic logics.

2 THE LOGICAL BACKBONE

In a similar spirit as Gabbay (1985), Kraus et al. (1990) and Makinson (1994) we will take a very general and abstract view of nonmonotonic logics. The definition we here give of what a nonmonotonic logic is puts into focus two elements that usually are either assumed only implicitly for a nonmonotonic logic, or not deemed necessary at all: the existence of a monotonic "background" entailment relation, and the possible dependence of the nonmonotonic entailment relation on the underlying vocabulary. The first of these elements will be crucial for our definition of when two knowledge bases are alternative representations; the second might be regarded as a borderline case of representation dependence.

The following definition is an adaptation of standard definitions in generalized model theory (e.g. (Ebbinghaus 1985)), tailored for the purpose at hand by the distinction of two entailment relations. Here and elsewhere we denote the powerset of X by $\mathscr{P}(X)$.

Definition 2.1 A *logic* $\mathscr{L}$ consists of

- a class of sets S, called the class of *vocabularies* of $\mathscr{L}$, which is closed under intersections and finite unions,
- for each vocabulary S a set L_S of *expressions* of $\mathscr{L}$, such that $S \subseteq S'$ implies $\mathrm{L}_S \subseteq \mathrm{L}_{S'}$.
- for each vocabulary S a relation $\vdash_S$ on $\mathscr{P}(\mathrm{L}_S)\times\mathrm{L}_S$ (the *classical entailment relation*) that is *monotonic* (i.e. $\Phi \subseteq \Phi' \subseteq \mathrm{L}_S$ and $\Phi \vdash_S \phi$ implies $\Phi' \vdash_S \phi$), and that has the *reduct property* (i.e. $\Phi \vdash_{S'} \phi$ and $S' \subseteq S$ implies $\Phi \vdash_S \phi$).

$\mathscr{L}$ is a *nonmonotonic logic* if in addition there is

- for each vocabulary S a relation $\mid\!\sim_S$ on $\mathscr{P}(\mathrm{L}_S) \times \mathrm{L}_S$ (*the nonmonotonic entailment relation*) with $\vdash_S \subseteq \mid\!\sim_S$.

Note that for the sake of simplicity we have already built the property of *supraclassicality* into the definition of a nonmonotonic logic. The concept of nonmonotonic logics is not restricted in any way by making the existence of a classical entailment relation part of the definition: whenever we have a nonmonotonic logic that lacks a natural concept of classical entailment, we can extend it to fit definition 2.1 by simply defining $\vdash_S := \emptyset$ for all S.

By virtue of the reduct property we can delete the subscript in the classical entailment relation $\vdash_S$, and simply write $\Phi \vdash \phi$, meaning that $\Phi \vdash_S \phi$ for any S with $\Phi \cup \{\phi\} \subseteq \mathrm{L}_S$. Generally, we may not expect the nonmonotonic inference relation of $\mathscr{L}$ to possess the reduct property. One counterexample is supplied by the probabilistic *center of mass* inference rule, see (Paris & Vencovská 1992) (the maximum entropy principle, on the other hand, satisfies the reduct property). Therefore, for the nonmonotonic entailment relation $\mid\!\sim_S$ we have to retain the subscript S, unless the reduct property has been established for $\mid\!\sim$ in the logic under consideration.

3 REPRESENTATIONAL VARIANTS

Before the question can be addressed, what it means for the nonmonotonic inference relation $\mid\!\sim$ to be representation independent, we have to clarify what it means for two knowledge bases Φ, Ψ to represent the same information. This will be formalized entirely in terms of the classical entailment relation $\vdash$. Throughout this section we will therefore be concerned only with the classical part of a logic.

Reconsider our introductory example. Intuitively, both τ_1 and τ_2 represent the same information with respect to the existence of life on mars. It is not immediate, however, how this intuition can be captured by a formal logical property of τ_1 and τ_2. Clearly, we are looking for something weaker than logical equivalence, because for example with $\psi :\equiv P(ALoM) < 0.3 \rightarrow P(PLoM) > 0.3$ we have $\tau_2 \vdash \psi$, but $\tau_1 \not\vdash \psi$.

This, of course, is not surprising: by saying that τ_1 and τ_2 provide the same information with respect to the existence of life on mars, we do not mean to imply that τ_1 and τ_2 provide the same information regarding any other statement that can be formulated in either the vocabulary of τ_1, or the vocabulary of τ_2. Only with regard to statements that can be represented in either vocabulary will we expect the same inferences from τ_1 and τ_2. This "common ground" of the two languages can be defined by generalizing what is known as an *interpretation* in model theory (see (Hodges 1993)).

Definition 3.1 Let $\mathscr{L}$ be a logic, S, S' vocabularies of $\mathscr{L}$. An *abstract interpretation* in $\mathscr{L}$ of L_S into $\mathrm{L}_{S'}$ with *admissibility conditions* $\Omega(f) \subset \mathrm{L}_{S'}$ is a mapping

$$f : \mathrm{L}_S \to \mathrm{L}_{S'}$$

such for all $\Phi \cup \{\psi\} \subseteq \mathrm{L}_S$:

$$\Phi \vdash_S \psi \quad \Rightarrow \quad f(\Phi) \cup \Omega(f) \vdash_{S'} f(\psi), \tag{1}$$

where $f(\Phi) := \{f(\phi) \mid \phi \in \Phi\}$. The class of abstract interpretations in $\mathscr{L}$ is denoted $\mathit{Int}(\mathscr{L})$.

Example 3.2 Let $\mathscr{L}^{\mathrm{prop}}$ be propositional logic with languages $\mathrm{L}_S^{\mathrm{prop}}$. Let S and S' be two sets of propositional variables. The extension of any function

$$f : \; S \to \mathrm{L}_{S'}^{\mathrm{prop}}.$$

to $\mathrm{L}_S^{\mathrm{prop}}$ via the conditions $f(\neg\phi) \equiv \neg f(\phi)$; $f(\phi \vee \psi) \equiv f(\phi) \vee f(\psi)$ is an abstract interpretation of $\mathrm{L}_S^{\mathrm{prop}}$ into $\mathrm{L}_{S'}^{\mathrm{prop}}$ with $\Omega(f) = \emptyset$. We call it a *propositional interpretation* and denote the class of propositional interpretations by $\mathit{PI}(\mathscr{L}^{\mathrm{prop}})$.

As in this example, throughout this paper we will only encounter abstract interpretations for which we may let $\Omega(f) = \emptyset$. The admissibility conditions become relevant, for example, when we move to the standard concept of interpretations in first-order logic. Here, an atomic formula $h(x) = y \in \mathrm{L}_S$ is mapped to a formula $\phi(x, y) \in \mathrm{L}_{S'}$. The admissibility condition $\Omega(f)$ then would have to contain the condition that ϕ is functional, i.e. include the axiom $\forall x \exists^{=1} y \phi(x, y)$.

With abstract interpretations at our disposal, we can define what it means for two knowledge bases Φ, Ψ of a logic $\mathscr{L}$ to represent the same information with respect to a common ground – which is any language L_{S^*} that can be interpreted in both the language of Φ and the language of Ψ.

Definition 3.3 Let L_{S^*}, L_{S_1}, L_{S_2} be languages in some logic $\mathscr{L}$. Let $f : \; \mathrm{L}_{S^*} \to \mathrm{L}_{S_1}$ and $g : \; \mathrm{L}_{S^*} \to \mathrm{L}_{S_2}$ be abstract interpretations with admissibility conditions $\Omega(f)$ and $\Omega(g)$. Let $\Phi \subseteq \mathrm{L}_{S_1}$, $\Psi \subseteq \mathrm{L}_{S_2}$. Φ and Ψ are called *representational variants with respect to f and g*, written $\Phi \overset{f,g}{\longleftrightarrow} \Psi$, iff for all $\alpha \in \mathrm{L}_{S^*}$:

$$\Phi \cup \Omega(f) \vdash_{S_1} f(\alpha) \;\text{ iff }\; \Psi \cup \Omega(g) \vdash_{S_2} g(\alpha).$$

The concept of representational varianthood being central to our development of representation independence, we shall here digress for a short while from the straightforward presentation of our main topic, and for the remainder of this section take a closer look at representational variants.

First, we turn to the question of how given knowledge bases Φ and Ψ can be proved to be representational variants, and to the even more elementary question of when a given mapping $f : \mathrm{L}_S \to \mathrm{L}_{S'}$ is an abstract interpretation.

For this purpose we here narrow down the class of logics to be considered to such logics in which the concepts introduced in definitions 3.1 and 3.3 as purely syntactical relations correspond to relations between semantical structures, so that questions about representational varianthood can be answered by constructing and comparing semantical structures.

Still adopting standard notions from generalized model theory, we consider logics that have a *model theoretic presentation*, and the *Boole property.* For a logic to have a model theoretic presentation means that for every vocabulary S there is defined a class of *S-structures* M, and a *satisfaction relation* $\models_S$ between S-structures and sets of S-expressions, such that $\Phi \vdash_S \psi$ iff for all S-structures M: $M \models_S \Phi$ implies $M \models_S \psi$. In propositional logic, for example, an S-structure is simply a truth assignment to the propositional variables, in first-order logic we have the usual model-theoretic S-structures. For $\Phi \subseteq \mathrm{L}_S$ we denote by $\mathrm{Mod}(\Phi, S)$ the class of S-structures M with $M \models \Phi$. Two classes $\mathcal{M}$ and $\mathcal{N}$ of S-structures are called *elementarily equivalent*, written $\mathcal{M} \sim \mathcal{N}$, when for all $\phi \in \mathrm{L}_S$: $M \models \phi$ for all $M \in \mathcal{M}$ iff $N \models \phi$ for all $N \in \mathcal{N}$. A logic with a model theoretic presentation has the Boole-property iff for every $\phi, \psi \in \mathrm{L}_S$ there exist expressions $\neg\phi, \phi \vee \psi \in \mathrm{L}_S$ such that for every S-structure M: $M \models \neg\phi$ iff $M \not\models \phi$, and $M \models \phi \vee \psi$ iff $M \models \phi$ or $M \models \psi$.

Logics with a model-theoretic presentation and the Boole property permit us to apply the usual rules of classical logic for the manipulation of boolean connectives with respect to the classical entailment relation. Particularly, in any such logic the deduction theorem is valid: $\Phi \cup \{\phi\} \vdash \psi$ iff $\Phi \vdash \phi \to \psi$ ($\phi \to \psi$ being, of course, and abbreviation for $\neg\phi \vee \psi$).

We now obtain the following connection between S-structures and interpretations:

Lemma 3.4 Let $\mathscr{L}$ be a logic with a model theoretic presentation, $f : \mathrm{L}_S \to \mathrm{L}_{S'}$, $\Omega(f) \subset \mathrm{L}_{S'}$. If for every S'-structure M with $M \models \Omega(f)$ there exists an S-structure N such that for all $\phi \in \mathrm{L}_S$:

$$M \models f(\phi) \;\text{ iff }\; N \models \phi \tag{2}$$

then f is an abstract interpretation with admissibility conditions $\Omega(f)$.

When $\mathscr{L}$ has the Boole-property, and f is compatible with boolean connectives, i.e. $f(\neg\phi) \equiv \neg f(\phi)$, $f(\phi \vee \psi) \equiv f(\phi) \vee f(\psi)$ ($\phi, \psi \in \mathrm{L}_S$), then the converse also holds: if f is an abstract interpretation with admissibility conditions $\Omega(f)$ then for each $M \models \Omega(f)$ there exists an S-structure N with (2).

Proof: For the first part of the lemma let f and $\Omega(f)$ be given. Let $\Phi \cup \{\psi\} \subseteq \mathrm{L}_S$ with $\Phi \vdash \psi$. If $M \not\models f(\Phi)$ for all S'-structures M with $M \models \Omega(f)$ then (1) is trivially satisfied. Now suppose that $M \models f(\Phi) \cup \Omega(f)$, and let N be an S-structure satisfying (2). Then $N \models \Phi \cup \{\psi\}$, hence $M \models f(\psi)$, and therefore $f(\Phi) \cup \Omega(f) \vdash f(\psi)$.

For the second part assume that f is an abstract interpretation with admissibility conditions $\Omega(f)$ that is compatible with boolean connectives. Let $M \models \Omega(f)$ be given, and define $\Phi := \{\phi \in \mathrm{L}_S \mid M \models f(\phi)\}$. Assume that there does not exist an S-structure N with $N \models \Phi$. Then $\Phi \vdash \phi \wedge \neg\phi$ for arbitrary $\phi \in \mathrm{L}_S$, and by (1) $f(\Phi) \cup \Omega(f) \vdash f(\phi) \wedge \neg f(\phi)$, contradicting the assumption that $M \models f(\Phi) \cup \Omega(f)$. □

When for an abstract interpretation $f : \mathrm{L}_S \to \mathrm{L}_{S'}$ and any S'-structure M an S-structure N with (2) exists, then f is said to *define associated structures*. The mapping that assigns an associated structure N to S'-structures M is denoted by $\bar{f}$. Generally, this mapping will not be uniquely determined by condition (2), and $\bar{f}(M)$ can be any (usually canonical) selection of one of the associated S-structures for M.

Example 3.5 Let $S = \{A_1, \ldots, A_k\}$, $S' = \{B_1, \ldots, B_l\}$ be propositional vocabularies, $f : \mathrm{L}_S \to \mathrm{L}_{S'}$; $f(A_i) :\equiv \phi_{A_i}$ a propositional interpretation.

Let $M : S' \to \{\mathit{true}, \mathit{false}\}$ be a truth assignment. Then the unique truth assignment $\bar{f}(M) : S \to \{\mathit{true}, \mathit{false}\}$ for which (2) holds is defined by

$$\bar{f}(M)(A_i) := M(\phi_{A_i}) \quad (A_i \in S).$$

A criterion for representational varianthood now is given by the following lemma.

Lemma 3.6 Let $\mathscr{L}$ be a logic with a model theoretic presentation. Let $f : \mathrm{L}_{S^*} \to \mathrm{L}_{S_1}$, $g : \mathrm{L}_{S^*} \to \mathrm{L}_{S_2}$ be abstract interpretations that define associated structures $\bar{f}(\cdot), \bar{g}(\cdot)$. Let $\Phi \subseteq \mathrm{L}_{S_1}$, $\Psi \subseteq \mathrm{L}_{S_2}$. Then $\Phi \overset{f,g}{\longleftrightarrow} \Psi$ iff

$$\bar{f}(\mathrm{Mod}(\Phi \cup \Omega(f), S_1)) \sim \bar{g}(\mathrm{Mod}(\Psi \cup \Omega(g), S_2)).$$

The proof is immediate from the definitions.

Example 3.7 Let S_1 be the set of propositional variables *LoM* (Life on Mars), *LoE* (Life on Earth), *LiSS* (Life in Solar System), *IntL* (Intelligent Life), and *EalL* (Earth-like Life). Let S_2 consist of *ALoM* (Animal Life on Mars), *PLoM* (Plant Life on Mars), *LiSS* (Life in Solar System), and *HL* (Human Life).

Let

$$\begin{array}{ll} \phi \equiv (\mathit{LoM} \to \mathit{LiSS} \wedge \neg \mathit{IntL}) \wedge (\mathit{LoE} \to \mathit{EalL}) & \in \mathrm{L}_{S_1} \\ \psi \equiv (\mathit{ALoM} \vee \mathit{PLoM}) \to (\mathit{LiSS} \wedge \neg \mathit{HL}) & \in \mathrm{L}_{S_2} \end{array}$$

Here we think of the propositional variables rather as relation symbols, designating the set of all objects that are "life on mars", "intelligent life", and so on, rather than as propositions "there exists life on mars", etc.

The only propositional variable that the two vocabularies have in common is *LiSS*. Still we can extract the same information about the relation of life on mars, life in the solar system, and human life from the two formulas. To make this precise, consider the propositional interpretations f, g of $S^* := \{A, B, C\}$ into S_1 and S_2, respectively defined by

$$\begin{array}{lll} f: & A \mapsto \mathit{LoM}, & B \mapsto (\mathit{IntL} \wedge \mathit{EalL}), \\ & C \mapsto \mathit{LiSS}; & \\ g: & A \mapsto (\mathit{ALoM} \vee \mathit{PLoM}), & B \mapsto \mathit{HL}, \\ & C \mapsto \mathit{LiSS}. & \end{array}$$

We may now use lemma 3.6 to show that $\phi \overset{f,g}{\longleftrightarrow} \psi$. For this example, and for propositional logic in general, this is particularly simple, because here for any classes $\mathcal{M}$, $\mathcal{N}$ of structures $\mathcal{M} \sim \mathcal{N}$ is equivalent to $\mathcal{M} = \mathcal{N}$. In the current example it is readily verified that both $\bar{f}(\mathrm{Mod}(\phi, S_1))$ and $\bar{g}(\mathrm{Mod}(\psi, S_2))$ are equal to the set $\{(t,f,t), (f,f,t), (f,t,t), (f,t,f), (f,f,f)\}$ of truth assignments to the variables (A, B, C), which corresponds to the set of models of the S^*-formula $A \to C \wedge \neg B$.

The two formulas ϕ and ψ in this example differ in one aspect that we will later find to be quite essential: while ψ actually is the formula $g(A \to C \wedge \neg B)$, it is not the case that $\phi \equiv f(A \to C \wedge \neg B)$. Moreover, ϕ is not even equivalent to any formula $f(\alpha)$ with $\alpha \in \mathrm{L}_{S^*}$, because ϕ makes a nontrivial statement about the propositional variable *LoE*, which does not appear in the range of f.

The following example is useful to caution us against expecting more from the relation $\Phi \overset{f,g}{\longleftrightarrow} \Psi$ than there actually is in it. For this example and further use below we introduce the notation $\mathit{id}(\mathrm{L}_S)$ to denote the "identity interpretation" of L_S, i.e. the abstract interpretation $\mathit{id} : \mathrm{L}_S \to \mathrm{L}_S$ with $\mathit{id}(\alpha) \equiv \alpha$ for all $\alpha \in \mathrm{L}_S$, and $\Omega(\mathit{id}) = \emptyset$.

Example 3.8 Let $S^* = S_1 = \{A\}$, $S_2 = \{B, C\}$ be propositional vocabularies, $f = id(\mathrm{L}_{S^*})$, $g : \mathrm{L}_{S^*}^{\mathrm{prop}} \to \mathrm{L}_{S_2}^{\mathrm{prop}}$; $g(A) :\equiv B \vee C$ be propositional interpretations. Let $\phi :\equiv A$ and $\psi :\equiv B \wedge C$. Then $\phi \overset{f,g}{\longleftrightarrow} \psi$. It is not true, however, that also $\neg\phi \overset{f,g}{\longleftrightarrow} \neg\psi$, because $\neg\phi \vdash f(\neg A)$, whereas $\neg\psi \not\vdash \neg(B \vee C) \equiv g(\neg A)$.

By a similar example it can also be shown that $\phi \overset{f,g}{\longleftrightarrow} \psi$ and $\phi' \overset{f,g}{\longleftrightarrow} \psi'$ does not imply $\phi \wedge \phi' \overset{f,g}{\longleftrightarrow} \psi \wedge \psi'$.

Interestingly, however, in any logic with a model-theoretic presentation and the Boole property $\phi \overset{f,g}{\longleftrightarrow} \psi$ and $\phi' \overset{f,g}{\longleftrightarrow} \psi'$ does imply $\phi \vee \phi' \overset{f,g}{\longleftrightarrow} \psi \vee \psi'$, because for $\alpha \in \mathrm{L}_{S^*}$ we then have $\phi \vee \phi' \vdash f(\alpha)$ iff $\phi \vdash f(\alpha)$ and $\phi' \vdash f(\alpha)$ iff $\psi \vdash g(\alpha)$ and $\psi' \vdash g(\alpha)$ iff $\psi \vee \psi' \vdash g(\alpha)$.

4 REPRESENTATION INDEPENDENCE

Having defined what it means that Φ and Ψ represent the same information (with respect to a commonly interpretable language L_{S^*}), we can address our main issue, and define representation (in)dependence of $\mid\!\sim$.

Now that we will be dealing with the nonmonotonic entailment relation of a logic, somewhat greater care than in the preceding section has to be taken with respect to specifying what vocabulary is being assumed as defining the background language in which a knowledge base is to be evaluated. Basically, this just means that we have to more conscientiously use the properly indexed entailment operator $\mid\!\sim_S$. In addition, it will be convenient to also add the specification of a vocabulary S with $\Phi \subseteq \mathrm{L}_S$ as an external attribute to a knowledge base. We therefore define the class of *knowledge bases* of $\mathscr{L}$ as the set of pairs

$$KB(\mathscr{L}) = \{(\Phi, S) \mid \Phi \subseteq \mathrm{L}_S;\ S \text{ vocabulary}\}.$$

Whenever the specification of the vocabulary S can be dispensed with, either because the logic under consideration is known to have the reduct property for $\mid\!\sim$, or because the underlying vocabulary is clear or irrelevant in the given context, we will continue to simply write Φ for (Φ, S).

The general thrust of the definition of representation independence should be quite obvious: the nonmonotonic logic $\mathscr{L}$ is representation independent (r.i.) if, whenever (Φ, S_1) and (Ψ, S_2) represent the same information about L_{S^*}, then $\mathscr{L}$ also permits the same nonmonotonic inferences from (Φ, S_1) and (Ψ, S_2) about L_{S^*}, i.e. if $\Phi \overset{f,g}{\longleftrightarrow} \Psi$, then for all $\alpha \in \mathrm{L}_{S^*}$:

$$\Phi \cup \Omega(f) \mid\!\sim_{S_1} f(\alpha) \ \text{ iff } \ \Psi \cup \Omega(g) \mid\!\sim_{S_2} g(\alpha). \qquad (*)$$

This is essentially the definition of r.i. that we will give below. However, a few complications are inevitable: we can not hope to obtain $(*)$ for any interesting nonmonotonic logic for arbitrary Φ, Ψ, f, and g. Usually, we will have to look a little closer at how Φ, Ψ, f, and g can interact in a specific logic, and try to obtain only qualified statements of the form $(*)$.

There are basically two ways to impose restrictions on the statements of this form: we may consider only a special set of abstract interpretations, not arbitrary ones, and we may restrict the set of admissible knowledge bases Φ and Ψ, where these sets may also depend on the interpretations considered.

The first of these limitations is really inevitable: whenever we consider a specific logic $\mathscr{L}$, we will be concerned with certain classes of natural interpretations, e.g. in the context of propositional logic the propositional interpretations of example 3.2, and not with any abstract interpretation in the sense of definition 3.1. Restrictions of the second type, on the other hand, usually mean a real compromising of the strength of the results obtained.

Definition 4.1 Let $\mathscr{L}$ be a nonmonotonic logic. Let

$$\boldsymbol{F} \subseteq Int(\mathscr{L}) \times Int(\mathscr{L}) \times KB(\mathscr{L}) \times KB(\mathscr{L}).$$

$\mathscr{L}$ is *representation independent with respect to* $\boldsymbol{F}$, iff $(f, g, (\Phi, S_1), (\Psi, S_2)) \in \boldsymbol{F}$ and $\Phi \overset{f,g}{\longleftrightarrow} \Psi$ implies $\forall \alpha \in \mathrm{L}_{S^*}$:

$$\Phi \cup \Omega(f) \mid\!\sim_{S_1} f(\alpha) \ \text{ iff } \ \Psi \cup \Omega(g) \mid\!\sim_{S_2} g(\alpha), \qquad (3)$$

where L_{S^*} is the common domain of f and g.

Note that in this definition the condition $(f, g, (\Phi, S_1), (\Psi, S_2)) \in \boldsymbol{F}$ in itself does not guarantee that f and g are defined on the same language L_{S^*}. However, this is implicit in the additional condition $\Phi \overset{f,g}{\longleftrightarrow} \Psi$.

The following examples show how the general concept of representation independence of definition 4.1 subsumes special properties previously considered in the literature.

Example 4.2 The reduct property for the nonmonotonic entailment relations $\mid\!\sim_S$ may be represented in the framework of definition 4.1 as follows. Let $\boldsymbol{F}_{RP}$ contain all tuples of the form $(f, f, (\Phi, S_1), (\Phi, S_2))$ where $f = id(\mathrm{L}_S)$ for some L_S with $\Phi \subseteq \mathrm{L}_S$. For any such tuple $\Phi \overset{f,f}{\longleftrightarrow} \Phi$ then holds, and nonmonotonic entailment in $\mathscr{L}$ has the reduct property iff $\mathscr{L}$ is r.i. with respect to F_{RP}.

Example 4.3 One of the simplest changes of representation one can devise is by renaming the nonlogical symbols in an expression. Since definition 2.1 does not require that the expressions in L_S will actually be strings constructed from elements of S, we may not have a feasible concept of renaming for every logic in the sense of definition 2.1. Therefore, assume that $\mathscr{L}$ is a logic in which each bijection $f : S \to S'$ of vocabularies (respecting such attributes as sort and arity of symbols, if these exist) induces a bijection $f : \mathrm{L}_S \to \mathrm{L}_{S'}$ of the corresponding languages such that f is an abstract interpretation of L_S in $\mathrm{L}_{S'}$, and f^{-1} is an abstract interpretation of $\mathrm{L}_{S'}$ in L_S with $\Omega(f) = \Omega(f^{-1}) = \emptyset$ (this just means that the classical part or $\mathscr{L}$ has the *renaming property* (Ebbinghaus 1985)).

Now let $\boldsymbol{F}_R$ contain all tuples of the form $(f, g, (\Phi, S), (g(\Phi), S'))$ where $f = id(\mathrm{L}_S)$, and $g : S \to S'$ is a renaming. By the renaming property of $\vdash$ any such tuple already satisfies $\Phi \overset{f,g}{\longleftrightarrow} g(\Phi)$, and nonmonotonic entailment in $\mathscr{L}$ is invariant under renaming iff $\mathscr{L}$ is r.i. with respect to $\boldsymbol{F}_R$.

Example 4.4 Another basic form of changing representation is by choosing a logically equivalent knowledge base (in the sense of the classical entailment relation) in the same language. When such a change does not affect the set of formulas entailed by $\mid\!\sim$ a nonmonotonic logic $\mathscr{L}$ is said to have the property of *left logical equivalence* (cf. (Makinson 1994)).

This property, too, can be very easily captured as a form of representation independence in the sense of definition 4.1. Let $\boldsymbol{F}_{LLE}$ contain all tuples of the form $(f, f, (\Phi, S_1), (\Psi, S_2))$ where $f = id(\mathrm{L}_S)$ for some S with $\Phi, \Psi \subseteq \mathrm{L}_S$. The condition $\Phi \overset{f,f}{\longleftrightarrow} \Psi$ then means that Φ and Ψ are logically equivalent, and $\mathscr{L}$ is r.i. with respect to $\boldsymbol{F}_{LLE}$ iff $\mathscr{L}$ satisfies left logical equivalence.

Note that $\boldsymbol{F}_{RP} \subseteq \boldsymbol{F}_{LLE}$, so that left logical equivalence as defined here implies the reduct property.

Definition 4.1, once again, is formulated on a purely syntactical level. Pursuing our endeavor to supplement these syntactical concepts with corresponding semantical notions, we shall now describe a semantical approach to representation independence for a certain class of nonmonotonic logics.

As previously we had to narrow down the class of logics to be considered to those in which classical entailment is defined by a model-theoretic presentation, we now have to focus on such logics in which nonmonotonic entailment, too, has a semantical background. A large and natural such class is comprised of those logics in which nonmonotonic entailment has a *preferential model semantics* (Shoham 1987). We here generalize the notion of preferential models in one aspect, for which we need the following preparatory definition.

Definition 4.5 Let $\mathscr{L}$ be a logic with a model-theoretic presentation, S a vocabulary, and $\mathcal{S}$ the class of S-structures. We say that $\mathcal{S}$ is *locally ordered* iff on every subset $\mathcal{S}' \subseteq \mathcal{S}$ a unique partial order is defined. $\mathcal{S}$ is said to be *globally ordered* iff $\mathcal{S}$ is locally ordered, and for each $\mathcal{S}' \subseteq \mathcal{S}$ the order on $\mathcal{S}'$ is given by the restriction to $\mathcal{S}'$ of the order on $\mathcal{S}$. We denote by $\lfloor \mathcal{S}' \rfloor$ the set of elements of $\mathcal{S}'$ that are minimal in the ordering of $\mathcal{S}'$.

Definition 4.6 Let $\mathscr{L}$ be a nonmonotonic logic with a model-theoretic presentation. $\mathscr{L}$ has a *local preferential model semantics* iff for each vocabulary S the class $\mathcal{S}$ is locally ordered, and for each $\Phi \cup \{\phi\} \subseteq \mathrm{L}_S$: $\Phi \mid\!\sim_S \phi$ iff $M \models \phi$ for all $M \in \lfloor Mod(\Phi, S) \rfloor$. $\mathscr{L}$ has a *global preferential model semantics* iff $\mathscr{L}$ has a local preferential model semantics defined by a global order on each $\mathcal{S}$.

Hence, the usual notion of preferential model semantics corresponds to a global preferential model semantics as defined here. The generalization to local preferential model semantics is necessary in order to accommodate certain probabilistic logics (cf. section 5.3), and may also deserve some attention in other contexts. We now obtain the following criterion for representation independence.

Lemma 4.7 Let $\mathscr{L}$ be a logic with a local preferential model semantics, let $\boldsymbol{F}$ be as in definition 4.1 such that all f and g that appear in a tuple of $\boldsymbol{F}$ define associated structures. Then $\mathscr{L}$ is r.i. with respect to $\boldsymbol{F}$ iff for every tuple $(f, g, (\Phi, S_1), (\Psi, S_2)) \in \boldsymbol{F}$ with $\Phi \overset{f,g}{\longleftrightarrow} \Psi$

$$\bar{f}(\lfloor Mod(\Phi \cup \Omega(f), S_1) \rfloor) \sim \bar{g}(\lfloor Mod(\Psi \cup \Omega(g), S_2) \rfloor).$$

The proof is immediate from the definitions.

5 CASE STUDIES

In this section the general framework developed so far is applied to a selection of specific nonmonotonic logics.

5.1 CLOSED WORLD ASSUMPTION AND NEGATION AS FAILURE

The closed world assumption is a very general form of nonmonotonic reasoning that is not restricted to any specific logical framework. Any classical logic $\mathscr{L}$ that has a model theoretic presentation and the Boole property (actually, only any feasible concept of negation is required) can be extended to a nonmonotonic logic $\mathscr{L}^{\text{cwa}}$ by the definition $\Phi \mathrel{|\!\sim} \phi$ iff $\Phi \not\vdash \neg\phi$ (or $\Phi \vdash \neg\phi \wedge \phi$). It is obvious that $\mathscr{L}^{\text{cwa}}$ is r.i. in a very strong sense, namely with respect to $\boldsymbol{F} = \mathit{Neg} \times \mathit{Neg} \times \mathit{KB}(\mathscr{L}^{\text{cwa}}) \times \mathit{KB}(\mathscr{L}^{\text{cwa}})$, where $\mathit{Neg} \subseteq \mathit{Int}(\mathscr{L}^{\text{cwa}})$ is the class of abstract interpretations f with $f(\neg\alpha) \equiv \neg f(\alpha)$ for all α.

In logic programming an approximation of the closed world assumption may be implemented by the negation as failure rule (Clark 1978). Here we lose completely the representation independence of the closed world assumption; in fact, negation as failure does not even satisfy left logical equivalence: from the two logic programs

$$P: \begin{array}{l} Rf(x) \leftarrow Rf(x) \\ Qa \end{array} \quad \text{and} \quad P': \begin{array}{l} Rf(x) \leftarrow Rx \\ Qa \end{array}$$

the same literals are provable (namely, just Qa), so that in the sense of logic programming they are logically equivalent. However, for any literal $L \equiv Rf^n(a)$ $(n \geq 1)$, an attempted proof of L from P will not terminate, while it will fail from P'. Hence we get $P \mathrel{|\!\not\sim} \neg Rf^n(a)$, but $P' \mathrel{|\!\sim} \neg Rf^n(a)$ for all $n \geq 1$.

5.2 RATIONAL CLOSURE

Lehmann's (1989) rational closure and Pearl's (1990) system Z are two systems for evaluating *conditional knowledge bases* that for finitary knowledge bases have been found to be equivalent (Pearl 1990). Pearl (1990) also has shown that this logic can be defined by a preferential model semantics using *κ-rankings* going back to Spohn (1990). We here use this last approach for defining the (finitary) rational closure logic $\mathscr{L}^{\text{rc}}$.

The class of vocabularies of $\mathscr{L}^{\text{rc}}$ consists of all finite propositional vocabularies. An S-expression is a *conditional* $\phi \rightsquigarrow \psi$ with $\phi, \psi \in \mathrm{L}_S^{\text{prop}}$. An S-structure is a *ranking function*

$$\kappa : W_S \to \mathbf{N} \cup \{\infty\}$$

where $W_S = \{w_1, \ldots, w_n\}$ is the set of the $n = 2^{|S|}$ truth assignments for S. A ranking function κ satisfies a conditional $\phi \rightsquigarrow \psi$ iff

$$\mathit{min}\{\kappa(w) \mid w \models \phi \wedge \psi;\ w \in W_S\} < \mathit{min}\{\kappa(w) \mid w \models \phi \wedge \neg\psi;\ w \in W_S\},$$

using the conventions $\mathit{min}\,\emptyset = \infty$ and $\infty < \infty$. The class $\mathcal{S}$ of ranking functions for S is globally ordered by: $\kappa < \kappa'$ iff $\forall w \in W_S: \ \kappa(w) \leq \kappa'(w)$ and $\exists w \in W_S: \ \kappa(w) < \kappa'(w)$. Nonmonotonic entailment in $\mathscr{L}^{\text{rc}}$ then is defined by the preferential model semantics with respect to this ordering.

The definition of $\mathscr{L}^{\text{rc}}$ in one important aspect is of a somewhat different nature than those of many other nonmonotonic logics: in $\mathscr{L}^{\text{rc}}$ the nonmonotonic entailment relation of the logic, i.e. the entailment of conditionals $\phi \rightsquigarrow \psi$ from a set Φ of given conditionals, is not directly an image of the nonmonotonic commonsense reasoning that the logic wants to describe. Such reasoning, rather, is represented by the object-language operator $\rightsquigarrow$, and a commonsense inference of ψ from a given fact ϕ and a set Φ of default rules is modeled in the formal logic by the nonmonotonic entailment of $\phi \rightsquigarrow \psi$ from Φ.

$\mathscr{L}^{\text{rc}}$ satisfies the reduct property for nonmonotonic entailment, so that most of the time a reference to the underlying vocabulary here is unnecessary.

Lehmann & Magidor (1992) have observed that rational closure is invariant under renaming. We here extend this result for $\mathscr{L}^{\text{rc}}$ and show that this logic is r.i. in a more general sense.

We consider abstract interpretations $f^* : \mathrm{L}_S^{\text{rc}} \to \mathrm{L}_{S'}^{\text{rc}}$ obtained by extending propositional interpretations $f \in \mathit{PI}(\mathscr{L}^{\text{prop}})$ to L_S^{rc} via

$$f^*(\phi \rightsquigarrow \psi) :\equiv f(\phi) \rightsquigarrow f(\psi). \tag{4}$$

The following lemma tells us that f^* really is an abstract interpretation for $\mathscr{L}^{\text{rc}}$.

Lemma 5.1 Let $f \in \mathit{PI}(\mathscr{L}^{\text{prop}})$, $f : \mathrm{L}_S^{\text{prop}} \to \mathrm{L}_{S'}^{\text{prop}}$. Then $f^* : \mathrm{L}_S^{\text{rc}} \to \mathrm{L}_{S'}^{\text{rc}}$ as defined by (4) is an abstract interpretation with $\Omega(f^*) = \emptyset$.

Proof: We use lemma 3.4 and show that f^* defines associated structures. Let $\kappa : W_{S'} \to \mathbf{N} \cup \{\infty\}$ be an arbitrary S'-structure. An S-structure $\bar{f}^*(\kappa)$ is defined by letting for $w \in W_S$:

$$\bar{f}^*(\kappa)(w) := \mathit{min}\{\kappa(w') \mid w = \bar{f}(w');\ w' \in W_{S'}\}.$$

For every $\alpha \in \mathrm{L}_S^{\text{prop}}$ we then have

$$\begin{aligned} &\mathit{min}\{\bar{f}^*(\kappa)(w) \mid w \models \alpha;\ w \in W_S\} \\ &\quad = \ \mathit{min}\{\kappa(w') \mid w = \bar{f}(w');\ w \models \alpha;\ w' \in W_{S'}\} \\ &\quad = \ \mathit{min}\{\kappa(w') \mid w' \models f(\alpha);\ w' \in W_{S'}\} \end{aligned}$$

(note that in the special case that α only is true in $w \in W_S$ for which $\{w' \in W_{S'} \mid \bar{f}(w') = w\} = \emptyset$ we receive the value ∞ for the terms in this identity).

It now follows immediately that $\kappa \models f^*(\alpha \rightsquigarrow \beta)$ iff $\bar{f}^*(\kappa) \models \alpha \rightsquigarrow \beta$ for all $\alpha \rightsquigarrow \beta \in \mathrm{L}^{\mathrm{rc}}_S$, and hence that f^* is an abstract interpretation. □

The class of abstract interpretations defined by (4) is denoted $PI(\mathscr{L}^{\mathrm{rc}})$. For simplicity, in the sequel we will no longer distinguish in our notation an interpretation $f \in PI(\mathscr{L}^{\mathrm{rc}})$ from its underlying propositional interpretation $f \in PI(\mathscr{L}^{\mathrm{prop}})$. The distinction is implicit in the argument of the function, which may be either a propositional formula, or a conditional, and in the case of $\bar{f}$, a truth assignment, or a ranking function.

Letting

$$\boldsymbol{F}_R := \{(f,g,\Phi,\Psi) \mid f,g \in PI(\mathscr{L}^{\mathrm{rc}}); \\ \Phi \subseteq f(\mathrm{L}^{\mathrm{rc}}_{S^*}), \Psi \subseteq g(\mathrm{L}^{\mathrm{rc}}_{S^*})\},$$

we now obtain the following result about representation independence of $\mathscr{L}^{\mathrm{rc}}$.

Theorem 5.2 $\mathscr{L}^{\mathrm{rc}}$ is r.i. with respect to $\boldsymbol{F}_R$.

Proof: The theorem is proved using lemma 4.7 by showing that

$$\bar{f}(\lfloor \mathrm{Mod}(\Phi, S_1)\rfloor) = \bar{g}(\lfloor \mathrm{Mod}(\Psi, S_2)\rfloor) \qquad (5)$$

for $(f,g,\Phi,\Psi) \in \boldsymbol{F}_R$ with $\Phi \overset{f,g}{\longleftrightarrow} \Psi$.

For this, let $\kappa_1 \in \lfloor \mathrm{Mod}(\Phi, S_1)\rfloor$. We define a ranking function κ_2 on W_{S_2} by

$$\kappa_2(w) := \bar{f}(\kappa_1)(\bar{g}(w)) \qquad (w \in W_{S_2}).$$

It now has to be verified that $\bar{g}(\kappa_2) = \bar{f}(\kappa_1)$, and that $\kappa_2 \in \lfloor \mathrm{Mod}(\Psi, S_2)\rfloor$.

Let $w \in W_{S^*}$. Then

$$\begin{aligned}\bar{g}(\kappa_2)(w) &= min\{\kappa_2(w') \mid w = \bar{g}(w');\ w' \in W_{S_2}\} \\ &= min\{\bar{f}(\kappa_1)(\bar{g}(w')) \mid w = \bar{g}(w');\ w' \in W_{S_2}\} \\ &= min\{\bar{f}(\kappa_1)(w) \mid w = \bar{g}(w');\ w' \in W_{S_2}\}.\end{aligned}$$

When $\bar{g}^{-1}(w) \neq \emptyset$ we clearly obtain $\bar{f}(\kappa_1)(w)$ on the right hand side of this equation. Else, we get $min\,\emptyset = \infty$. To see that in this case $\bar{f}(\kappa_1)(w) = \infty$ also must hold, let ω be a propositional S^*-formula that is true only under the truth assignment w. From $\bar{g}^{-1}(w) = \emptyset$ it follows that $g(\omega)$ is not satisfiable in $\mathscr{L}^{\mathrm{prop}}$. Hence, in $\mathscr{L}^{\mathrm{rc}}$: $\vdash g(\omega) \rightsquigarrow \mathit{false}$. From $\Phi \overset{f,g}{\longleftrightarrow} \Psi$ it then follows that $\Phi \vdash f(\omega) \rightsquigarrow \mathit{false}$, meaning that

$$\begin{aligned}\infty &= min\{\kappa_1(w') \mid w' \models f(\omega);\ w' \in W_{S_1}\} \\ &= min\{\kappa_1(w') \mid w = \bar{f}(w');\ w' \in W_{S_1}\} \\ &= \bar{f}(\kappa_1)(w).\end{aligned}$$

From $\Psi \subseteq g(\mathrm{L}^{\mathrm{rc}}_{S^*})$ it immediately follows that $\kappa_2 \models \Psi$: for every conditional $g(\alpha) \rightsquigarrow g(\beta) \in \Psi$ ($\alpha, \beta \in \mathrm{L}^{\mathrm{prop}}_{S^*}$) trivially $\Psi \vdash g(\alpha) \rightsquigarrow g(\beta)$, and hence $\Phi \vdash f(\alpha) \rightsquigarrow f(\beta)$. Thus, $\bar{f}(\kappa_1) = \bar{g}(\kappa_2) \models \alpha \rightsquigarrow \beta$, and $\kappa_2 \models g(\alpha) \rightsquigarrow g(\beta)$.

Now assume that there exists $\kappa_2' \in \mathrm{Mod}(\Psi, S_2)$ with $\kappa_2' < \kappa_2$. From the special form of κ_2 (namely, the fact that $w_1, w_2 \in W_{S_2}$ with $\bar{g}(w_1) = \bar{g}(w_2)$ have the same rank in κ_2) it follows that then $\bar{g}(\kappa_2') < \bar{g}(\kappa_2)$. In the same manner as κ_2 was defined we now construct $\kappa_1' \in \mathrm{Mod}(\Phi, S_1)$ with $\bar{f}(\kappa_1') = \bar{g}(\kappa_2')$. By the special form of κ_1', and from $\bar{f}(\kappa_1') < \bar{f}(\kappa_1)$ it now follows that $\kappa_1' < \kappa_1$, contradicting the minimality of κ_1. Thus, $\kappa_2 \in \lfloor \mathrm{Mod}(\Psi, S_2)\rfloor$, which concludes the proof that $\bar{f}(\lfloor \mathrm{Mod}(\Phi, S_1)\rfloor) \subseteq \bar{g}(\lfloor \mathrm{Mod}(\Psi, S_2)\rfloor)$. By symmetry, the converse inclusion holds as well, which proves (5) □

Results similar to theorem 5.2 also hold for default logic (Reiter 1980) and Boutelier's (1994) conditional logic CT4O, among others. In each of these cases representation independence holds with respect to a set $\boldsymbol{F}$ whose definition includes the condition that $(f,g,\Phi,\Psi) \in \boldsymbol{F}$ only when $\Phi \subseteq f(\mathrm{L}_{S^*})$ and $\Psi \subseteq g(\mathrm{L}_{S^*})$. This condition (which we will also encounter in the following section) so far we have only been able to dispense with in case of the simple logic $\mathscr{L}^{\mathrm{cwa}}$. It may be possible in some cases to relax this condition when at the same time the class of interpretations admitted in $\boldsymbol{F}$ is further restricted. One natural restriction one may consider for interpretations f is to demand that for f the converse implication also holds in (1) (this is the syntactic analogue to Halpern and Koller's (1995) condition of "faithfulness" for embeddings of state spaces).

5.3 PROPOSITIONAL PROBABILISTIC LOGIC

We now return to our starting point: probabilistic inference rules, specifically a type of inference rules that have been studied extensively (e.g. (Diaconis & Zabell 1982), (Paris & Vencovská 1992), (Halpern & Koller 1995)), and which might be designated *measure selection inferences.*

In order to fit these rules into our framework, we first have to present them as nonmonotonic logics in the sense of definition 2.1. The classical parts of each of these logics will be the same and similar to several previously defined formalisms (e.g. (Nilsson 1986), (Frisch & Haddaway 1994)). We denote it by $\mathscr{L}^{\mathrm{pp}}$, and define it as follows.

The class of vocabularies of $\mathscr{L}^{\mathrm{pp}}$ is the class of all finite sets S of propositional variables. To define the language $\mathrm{L}_S^{\mathrm{pp}}$, we first inductively define an *S-term* to be either a constant symbol for a rational number, a *probability expression* $P(\phi)$, where $\phi \in \mathrm{L}_S^{\mathrm{prop}}$, or to be of the form $t_1 + t_2$ or $t_1 \cdot t_2$ with S-terms t_1 and t_2. An *atomic S-formula* now is defined as an expression of the form $t_1 \leq t_2$ with S-terms t_1 and t_2. Finally, $\mathrm{L}_S^{\mathrm{pp}}$ consists of the closure of atomic S-formulas under boolean connectives.

An S-structure is a probability measure μ on the set W_S of truth assignments for S (more precisely: μ is a measure on the algebra $\mathscr{P}(W_S)$). To define the satisfaction relation $\models$ between μ and $\mathrm{L}_S^{\mathrm{pp}}$-formulas, first, for $\phi \in \mathrm{L}_S^{\mathrm{prop}}$ let

$$\mu(\phi) := \mu(\{w \in W_S \mid w(\phi) = \mathit{true}\}).$$

For an atomic S-formula $t_1 \leq t_2$ we then define $\mu \models t_1 \leq t_2$ iff by substituting for each probability expression $P(\phi)$ occurring in t_1 or t_2 the value $\mu(\phi)$ we obtain a valid numerical inequality. The obvious conditions for boolean connectives then complete the definition of the relation $\mu \models \phi$ for arbitrary $\phi \in \mathrm{L}_S^{\mathrm{pp}}$.

The relation $\models$ now provides a model theoretic presentation for an entailment relation $\vdash$ of $\mathscr{L}^{\mathrm{pp}}$.

Next, what Paris & Vencovská (1992) and Halpern & Koller (1995) have called an inference process (or procedure) is turned into a nonmonotonic entailment relation with a preferential model semantics for $\mathscr{L}^{\mathrm{pp}}$. We proceed by first giving a formal definition for a *measure selection function*.

Definition 5.3 For $n \geq 1$ let

$$\Delta^n := \{(x_1, \ldots, x_n) \in \mathbf{R}^n \mid x_i \geq 0,\ \sum x_i = 1\}.$$

For each $n \geq 1$ let $I_n : \mathscr{P}(\Delta^n) \to \mathscr{P}(\Delta^n)$ with $I_n(G) \subseteq G$ for all $G \subseteq \Delta^n$, and $I_n(\pi(G)) = \pi(I_n(G))$ for all permutations π of $\{1, \ldots, n\}$. We then call $I := \cup_n I_n$ a measure selection function

While I is here defined to only work on subsets of Δ^n, it clearly induces a selection rule for measures on any finite algebra $\mathfrak{A}$: the set of probability measures on $\mathfrak{A}$, denoted $\Delta\mathfrak{A}$, after ordering the atoms $a_1, \ldots, a_n$ of $\mathfrak{A}$ in an arbitrary way, can be identified with Δ^n, and a measure selection function I can be used to select a subset $I(G)$ for every $G \subseteq \Delta\mathfrak{A}$. The condition of I_n being compatible with permutations makes $I(G)$ independent of the chosen ordering of the atoms of $\mathfrak{A}$.

Any measure selection function I now induces a preferential model semantics for $\mathscr{L}^{\mathrm{pp}}$. For $\Phi \subseteq \mathrm{L}_S^{\mathrm{pp}}$ we here denote $\mathrm{Mod}(\Phi, S) \subseteq \Delta\mathscr{P}(W_S)$ by $\Delta(\Phi, S)$, and define a partial order on $\Delta(\Phi, S)$ simply by $\mu < \mu'$ iff $\mu \in I(\Delta(\Phi, S))$ and $\mu' \notin I(\Delta(\Phi, S))$. Hence, in this ordering, $I(\Delta(\Phi, S)) = \lfloor \Delta(\Phi, S) \rfloor$. The nonmonotonic logic defined by adding the resulting nonmonotonic entailment relation to $\mathscr{L}^{\mathrm{pp}}$ is denoted by $\mathscr{L}_I^{\mathrm{pp}}$. As already noted in section 2, $\mathscr{L}_I^{\mathrm{pp}}$ may or may not have the reduct property for $\mid\!\sim$, depending on I.

In many cases the measure selection function I is defined by maximizing some (permutation invariant) function $K : \Delta^n \to \mathbf{R} \cup \{\infty\}$. In that case we actually have a global preferential model semantics defined by the preference relation on $\Delta\mathscr{P}(W_S)$: $\mu < \mu'$ iff $K(\mu) > K(\mu')$. The most prominent example, of course, being I_{me}, the maximum entropy selection rule. In some other interesting cases, however, we really only have a local preferential model semantics. The nonmonotonic logic $\mathscr{L}_{I_{\mathrm{cm}}}^{\mathrm{pp}}$ obtained from the center of mass selection rule I_{cm}, for example is not defined by a global preferential model semantics.

The definitions given so far now provide the formal framework for our introductory example from section 1: we get $\tau_1 \mid\!\sim P(\mathit{LoM}) = 0.6$ and $\tau_2 \mid\!\sim P(\mathit{ALoM} \vee \mathit{PLoM}) = 0.75$ in $\mathscr{L}_{I_{\mathrm{me}}}^{\mathrm{pp}}$. To make precise the representation dependence of maximum-entropy inference we first define a suitable subclass of $\mathit{Int}(\mathscr{L}^{\mathrm{pp}})$.

Let S, S' be finite propositional vocabularies, $f : \mathrm{L}_S^{\mathrm{prop}} \to \mathrm{L}_{S'}^{\mathrm{prop}}$ a propositional interpretation. In a straightforward way f then induces an abstract interpretation f^* of $\mathrm{L}_S^{\mathrm{pp}}$ into $\mathrm{L}_{S'}^{\mathrm{pp}}$ by inductively defining: $f^*(r) :\equiv r$ for rational constants r, $f^*(P(\phi)) :\equiv P(f(\phi))$ ($\phi \in \mathrm{L}_S^{\mathrm{prop}}$), and extending these definitions to arbitrary S-terms and $\mathrm{L}_S^{\mathrm{pp}}$-formulas via the canonical definitions for $+, \cdot, \leq, \neg$, and $\vee$.

Though this will be fairly obvious, we may formally appeal to lemma 3.4 in order to show that the mapping f^* really is an abstract interpretation with $\Omega(f^*) = \emptyset$: for $\mu \in \Delta\mathscr{P}(W_{S'})$ let $\bar{f}^*(\mu) \in \Delta(\mathscr{P}(W_S))$ be defined by

$$\bar{f}^*(\mu)(w) := \mu(\{w' \in W_{S'} \mid \bar{f}(w') = w\}) \quad (w \in W_S).$$

To see that (2) holds for μ and $\bar{f}^*(\mu)$ we may use an induction on the structure of $\phi \in \mathrm{L}_S^{\mathrm{pp}}$, founded on the following identity for $\gamma \in \mathrm{L}_S^{\mathrm{prop}}$:

$$\begin{aligned} \mu(f(\gamma)) &= \mu(\{w' \in W_{S'} \mid w'(f(\gamma)) = \mathit{true}\}) \\ &= \mu(\{w' \in W_{S'} \mid \bar{f}(w')(\gamma) = \mathit{true}\}) \\ &= \bar{f}^*(\mu)(\{w \in W_S \mid w(\gamma) = \mathit{true}\}) \\ &= \bar{f}^*(\mu)(\gamma). \end{aligned}$$

$PI(\mathscr{L}^{\mathrm{pp}})$ denotes the class of all abstract interpretations of $\mathscr{L}^{\mathrm{pp}}$ derived from propositional interpretations in the manner here described.

In the sequel, once again, we use f both for the underlying propositional interpretation, and the induced interpretation for $\mathscr{L}^{\mathrm{pp}}$. Note that then, whereas the argument of f may be either $\phi \in \mathrm{L}_S^{\mathrm{prop}}$ or $\phi \in \mathrm{L}_S^{\mathrm{pp}}$, the argument of the associated operator $\bar{f}$ can be either a truth assignment w, or a probability measure μ.

To continue the discussion of our example, let $S^* = \{A\}$, $f : A \mapsto LoM$, $g : A \mapsto ALoM \vee PLoM$. Then τ_1 and τ_2 are respectively $f(\alpha)$ and $g(\alpha)$ for $\alpha \equiv P(A) \geq 0.6$. Furthermore, $\tau_1 \overset{f,g}{\longleftrightarrow} \tau_2$, but for $\beta :\equiv P(A) = 0.6$: $\tau_1 \mid\!\sim f(\beta)$, $\tau_2 \not\mid\!\sim g(\beta)$. Thus, $\mathscr{L}^{\mathrm{pp}}_{I_{\mathrm{me}}}$ is not r.i. with respect to any $\boldsymbol{F}$ with $(f, g, \tau_1, \tau_2) \in \boldsymbol{F}$.

We now show how $\mathscr{L}^{\mathrm{pp}}_{I_{\mathrm{me}}}$ and any other nonmonotonic logic of the type $\mathscr{L}^{\mathrm{pp}}_{I}$ can be modified, so as to be r.i. with respect to

$$\boldsymbol{F}_P := \{(f, g, (\Phi, S_1), (\Psi, S_2)) \mid f, g \in PI(\mathscr{L}^{\mathrm{pp}}),\\ \Phi \subseteq f(\mathrm{L}^{\mathrm{pp}}_{S^*}),\ \Psi \subseteq g(\mathrm{L}^{\mathrm{pp}}_{S^*})\},$$

which clearly contains the tuple (f, g, τ_1, τ_2) from above.

To this end, we first take a closer look at the set $\Delta(\Phi, S)$ for $\Phi \subseteq \mathrm{L}^{\mathrm{pp}}_S$. By definition, $\Delta(\Phi, S)$ is a set of probability measures on the algebra $\mathscr{P}(W_S)$. Now consider $\Delta(\tau_2, S_2)$: this is a set of probability measures on $W_{S_2} = \{w_1, \ldots, w_4\}$, where $S_2 = \{ALoM, PLoM\}$. $\Delta(\tau_2, S_2)$ is defined by a single constraint on the subset of truth assignments w_i for which $w_i(ALoM \vee PLoM) = \mathit{true}$. Assume that these are just w_1, w_2, w_3. Then, in order to decide whether $\mu \in \Delta(\tau_2, S_2)$, we only have to consider the restriction of μ to the subalgebra

$$\mathfrak{A}_{\Delta(\tau_2, S_2)} := \{\emptyset, \{w_1, w_2, w_3\}, \{w_4\}, W_{S_2}\}$$

of $\mathscr{P}(W_{S_2})$. Generally, for every finite algebra $\mathfrak{A}$, $G \subseteq \Delta\mathfrak{A}$, and subalgebra $\mathfrak{A}' \subseteq \mathfrak{A}$, we say that G *is defined by constraints on* $\mathfrak{A}'$ iff

$$\forall \mu :\ \mu \in G \ \text{ iff } \ \mu \restriction \mathfrak{A}' \in \{\nu \restriction \mathfrak{A}' \mid \nu \in G\} =: G \restriction \mathfrak{A}'.$$

In (Jaeger 1995) it is shown that when $G \subseteq \Delta\mathfrak{A}$ is defined by constraints on $\mathfrak{A}' \subseteq \mathfrak{A}$ and by constraints on $\mathfrak{A}'' \subseteq \mathfrak{A}$, then G is also defined by constraints on $\mathfrak{A}' \cap \mathfrak{A}''$. Thus, there always exists a unique smallest algebra $\mathfrak{A}_G \subseteq \mathfrak{A}$, such that G is defined by constraints on $\mathfrak{A}_G$. For any measure selection function I we may therefore define a modified measure selection function $\tilde{I}$ by

$$\tilde{I}(G) := \{\nu \mid \nu \restriction \mathfrak{A}_G \in I(G \restriction \mathfrak{A}_G)\}. \qquad (6)$$

Returning to our example, we here obtain

$$\Delta(\tau_1, S_1) \restriction \mathfrak{A}_{\Delta(\tau_1, S_1)} = \Delta(\tau_2, S_2) \restriction \mathfrak{A}_{\Delta(\tau_2, S_2)}\\ = \{(x_1, x_2) \in \Delta^2 \mid x_1 \geq 0.6\}.$$

(assuming suitable corresponding orderings of the atoms of $\mathfrak{A}_{\Delta(\tau_1, S_1)}$ and $\mathfrak{A}_{\Delta(\tau_2, S_2)}$). The maximum entropy element of this set is (0.6,0.4). Using the modified maximum entropy selection rule $\tilde{I}_{\mathrm{me}}$ we then get

$$\tilde{I}_{\mathrm{me}}(\Delta(\tau_1, S_1)) = \{(0.6, 0.4)\}$$
$$\tilde{I}_{\mathrm{me}}(\Delta(\tau_2, S_2)) = \{(x_1, x_2, x_3, x_4) \in \Delta^4 \mid\\ x_1 + x_2 + x_3 = 0.6\},$$

and hence in $\mathscr{L}^{\mathrm{pp}}_{\tilde{I}_{\mathrm{me}}}$: $\tau_2 \mid\!\sim P(ALoM \vee PLoM) = 0.6$, as well as $\tau_1 \mid\!\sim P(LoM) = 0.6$.

A general result on the representation independence of logics $\mathscr{L}^{\mathrm{pp}}_{\tilde{I}}$ with respect to $\boldsymbol{F}_P$ can now be proven essentially by showing that the identity $\Delta(\Phi, S_1) \restriction \mathfrak{A}_{\Delta(\Phi, S_1)} = \Delta(\Psi, S_2) \restriction \mathfrak{A}_{\Delta(\Psi, S_2)}$ generally holds when $\Phi \overset{f,g}{\longleftrightarrow} \Psi$. However in certain cases it may only be true that $\Delta(\Phi, S_1) \restriction \mathfrak{A}_{\Delta(\Phi, S_1)} = \Delta(\Psi, S_2) \restriction \mathfrak{A}_{\Delta(\Psi, S_2)} \times \{0\}$, or $\Delta(\Psi, S_2) \restriction \mathfrak{A}_{\Delta(\Psi, S_2)} = \Delta(\Phi, S_1) \restriction \mathfrak{A}_{\Delta(\Phi, S_1)} \times \{0\}$, i.e. one of the two sets may consist of measures having an additional component, not corresponding to a component of measures from the other set, with constant value 0. In such cases we also write $G[\times 0] = J$ to express that either $G = J$ or $G \times \{0\} = J$. On account of this possibility we have to limit our general theorem to measure selection functions I that satisfy the condition: for all n, $G \subseteq \Delta^n$: $I(G \times \{0\}) = I(G) \times \{0\}$. We call such functions I *dimension independent*. Observe that all common selection functions are dimension independent.

Theorem 5.4 Let I be a dimension independent measure selection function, and $\tilde{I}$ defined by (6). Then $\mathscr{L}^{\mathrm{pp}}_{\tilde{I}}$ is r.i. with respect to $\boldsymbol{F}_P$.

Proof: Let $(f, g, (\Phi, S_1), (\Psi, S_2)) \in \boldsymbol{F}_P$, $\Phi \overset{f,g}{\longleftrightarrow} \Psi$. The theorem is proved using lemma 4.7 by showing that $\bar{f}(\tilde{I}(\Delta(\Phi, S_1))) = \bar{g}(\tilde{I}(\Delta(\Psi, S_2)))$, which is accomplished by proving the identities

$$\bar{f}(\tilde{I}(\Delta(\Phi, S_1))) = \tilde{I}(\bar{f}(\Delta(\Phi, S_1))) \qquad (7)$$

(and the corresponding statement for $\bar{g}, \Psi$) and

$$\bar{f}(\Delta(\Phi, S_1)) = \bar{g}(\Delta(\Psi, S_2)). \qquad (8)$$

Let $W_{S^*} = \{w_1, \ldots, w_n\}$. The (possibly empty) sets $\bar{f}^{-1}(w_i) \subseteq W_{S_1}$ $(i = 1, \ldots, n)$ form a partition of W_{S_1} and therefore are the atoms of a subalgebra $\mathfrak{A}_f$ of $\mathscr{P}(W_{S_1})$. A subalgebra $\mathfrak{A}_{\bar{f}}$ of $\mathscr{P}(W_{S^*})$

is defined by joining all the $w_i \in W_{S^*}$ for which $\bar{f}^{-1}(w_i) = \emptyset$, i.e. the atoms of $\mathfrak{A}_{\bar{f}}$ are the sets $\{w_i\}$ with $\bar{f}^{-1}(w_i) \neq \emptyset$ and $A_{\bar{f}} := \{w_i \mid \bar{f}^{-1}(w_i) = \emptyset\}$ (if this set is nonempty). We can view $\bar{f} : \bar{f}^{-1}(w_i) \mapsto \{w_i\}$ as an embedding of $\mathfrak{A}_f$ into $\mathfrak{A}_{\bar{f}}$, which is an isomorphism iff $A_{\bar{f}} = \emptyset$. Furthermore, taking measures on $\mathfrak{A}_{\bar{f}}$ and $\mathfrak{A}_f$ to be tuples in $\Delta^{k(+1)}$ defined by the corresponding orderings $(w_{i_1}, \ldots, w_{i_k}, A_{\bar{f}})$ and $(\bar{f}^{-1}(w_{i_1}), \ldots, \bar{f}^{-1}(w_{i_k}))$ $(w_{i_j} \in W_{S^*},\ \bar{f}^{-1}(w_{i_j}) \neq \emptyset)$ we obtain

$$\Delta(\Phi, S_1)\restriction\mathfrak{A}_f[\times 0] = \bar{f}(\Delta(\Phi, S_1))\restriction\mathfrak{A}_{\bar{f}}. \tag{9}$$

From the condition that $\Phi \subseteq f(\mathrm{L}^{\mathrm{pp}}_{S^*})$ it follows that $\Delta(\Phi, S_1)$ is defined by constraints on $\mathfrak{A}_f$, because the evaluation of any probability term $P(f(\alpha))$ ($\alpha \in \mathrm{L}^{\mathrm{prop}}_{S^*}$) that may occur in some $\phi \in \Phi$ only depends on the measure of the set $\cup\{\bar{f}^{-1}(w_i) \mid w_i \in W_{S^*}, w_i(\alpha) = \mathit{true}\} \in \mathfrak{A}_f$. Also, $\bar{f}(\Delta(\Phi, S_1))$ is defined by constraints on $\mathfrak{A}_{\bar{f}}$, because $\bar{f}(\mu)(w_i) = 0$ for all $\mu \in \Delta(\Phi, S_1)$ and $w_i \in A_{\bar{f}}$. Therefore, from (9) we obtain

$$\Delta(\Phi, S_1)\restriction\mathfrak{A}_{\Delta(\Phi,S_1)}[\times 0] = \bar{f}(\Delta(\Phi, S_1))\restriction\mathfrak{A}_{\bar{f}(\Delta(\Phi,S_1))}, \tag{10}$$

where corresponding atoms of $\mathfrak{A}_{\bar{f}(\Delta(\Phi,S_1))}$ and $\mathfrak{A}_{\Delta(\Phi,S_1)}$ now are of the form $\{w_{i_1}, \ldots, w_{i_l}\}$ and $\{\bar{f}^{-1}(w_{i_1}), \ldots, \bar{f}^{-1}(w_{i_l})\}$ $(w_{i_j} \in W_{S^*})$. In case that $\mu(\bar{f}^{-1}(w_{i_j})) = 0$ for some $w_{i_j} \in W_{S^*}$ with $\bar{f}^{-1}(w_{i_j}) \neq \emptyset$ and all $\mu \in \Delta(\Phi, S_1)$ there also is a pair of corresponding atoms of the form $\{w_{i_1}, \ldots, w_{i_l}, A_{\bar{f}}\}$ and $\{\bar{f}^{-1}(w_{i_1}), \ldots, \bar{f}^{-1}(w_{i_l})\}$, so that $\mathfrak{A}_{\bar{f}(\Delta(\Phi,S_1))}$ and $\mathfrak{A}_{\Delta(\Phi,S_1)}$ are actually isomorphic.

By the dimension independence of I it follows from (9) and (10) that

$$\tilde{I}(\Delta(\Phi, S_1)\restriction\mathfrak{A}_f)[\times 0] = \tilde{I}(\bar{f}(\Delta(\Phi, S_1))\restriction\mathfrak{A}_{\bar{f}}) \tag{11}$$

from which (7) now follows by $\bar{f}(\tilde{I}(\Delta(\Phi, S_1))) = \bar{f}(\tilde{I}(\Delta(\Phi, S_1)\restriction\mathfrak{A}_f))$, $\tilde{I}(\bar{f}(\Delta(\Phi, S_1))) = \tilde{I}(\bar{f}(\Delta(\Phi, S_1))\restriction\mathfrak{A}_{\bar{f}})$, and the fact that for $\mu \in \Delta\mathfrak{A}_f$, $\mu' \in \Delta\mathfrak{A}_{\bar{f}}$ with $\mu[\times 0] = \mu'$ we have $\bar{f}(\mu) = \mu'$.

To prove (8) let $\nu \in \Delta(\Psi, S_2)$. For each $w_i \in W_{S^*}$ let ω_i be a propositional S^*-formula that is true exactly under the truth assignment w_i. For $w_i \in A_{\bar{f}}$ we now have $\vdash P(f(\omega_i)) = 0$, because $f(\omega_i)$ is unsatisfiable. From $\Phi \overset{f,g}{\longleftrightarrow} \Psi$ it then follows that $\Psi \vdash P(g(\omega_i)) = 0$, and hence $\bar{g}(\nu)(\{w_i\}) = 0$. It follows that $\mu(\bar{f}^{-1}(w_i)) := \bar{g}(\nu)(\{w_i\})$ defines a probability measure on $\mathfrak{A}_f$. Taking any extension of μ to $\mathscr{P}(W_{S_1})$ we receive an S_1-structure μ with $\bar{f}(\mu) = \bar{g}(\nu)$.

To see that $\mu \in \Delta(\Phi, S_1)$ let $\phi \in \Phi$, $\phi \equiv f(\alpha)$ for some $\alpha \in \mathrm{L}^{\mathrm{pp}}_{S^*}$. Then $\mu \models \phi$ iff $\bar{f}(\mu) \models \alpha$ iff $\bar{g}(\nu) \models \alpha$ iff $\nu \models g(\alpha)$. By $\Phi \overset{f,g}{\longleftrightarrow} \Psi$ and $\nu \models \Psi$ this last condition clearly is satisfied, and the inclusion from left to right of (8) is proven. The other inclusion follows from the symmetrical argument. □

The theorem is no longer true when the condition $\Phi \subseteq f(\mathrm{L}^{\mathrm{pp}}_{S^*})$, $\Psi \subseteq g(\mathrm{L}^{\mathrm{pp}}_{S^*})$ is dropped from the definition of $\boldsymbol{F}_P$: for $\tau_3 :\equiv P(\mathit{ALoM} \vee \mathit{PLoM}) \geq 0.6 \wedge P(\mathit{ALoM}) < 0.4$ we still have $\tau_1 \overset{f,g}{\longleftrightarrow} \tau_3$. But now $\Delta(\tau_3, S_2)$ is defined by constraints on no smaller algebra than $\mathscr{P}(W_{S_2})$ itself, and $\tilde{I}_{\mathrm{me}}(\Delta(\tau_3, S_2)) = I_{\mathrm{me}}(\Delta(\tau_3, S_2)) = \emptyset$. Thus, we get $\tau_3 \mid\!\sim \mathit{false}$, yet $\tau_1 \not\mid\!\sim \mathit{false}$ in $\mathscr{L}^{\mathrm{pp}}_{\tilde{I}_{\mathrm{me}}}$.

$\boldsymbol{F}_P$ clearly contains the sets $\boldsymbol{F}_{RP}$, $\boldsymbol{F}_R$ and $\boldsymbol{F}_{LLE}$ of examples 4.2-4.4. Thus, for any dimension independent measure selection function I, $\tilde{I}$ possesses the reduct-, renaming-, and left logical equivalence property. For example, $\tilde{I}_{\mathrm{cm}}$ now has the reduct property that I_{cm} is lacking.

The modified maximum entropy selection rule, $\tilde{I}_{\mathrm{me}}$ can be shown to retain the property of I_{me} that Halpern and Koller (1995) have called "enforcing minimal irrelevance": when $\Phi \cup \{\phi\} \subseteq \mathrm{L}^{\mathrm{pp}}_S$, $\Psi \subseteq \mathrm{L}^{\mathrm{pp}}_{S'}$ with $S \cap S' = \emptyset$, then $\Phi \mid\!\sim \phi$ iff $\Phi \cup \Psi \mid\!\sim \phi$ in $\mathscr{L}^{\mathrm{pp}}_{\tilde{I}_{\mathrm{me}}}$. Clearly, $\mathscr{L}^{\mathrm{pp}}_{\tilde{I}_{\mathrm{me}}}$ also is not trivial in the sense of being "essentially entailment" (Halpern & Koller 1995). Finally, the representation independence of $\mathscr{L}^{\mathrm{pp}}_{\tilde{I}_{\mathrm{me}}}$ with respect to $\boldsymbol{F}_P$ implies that the measure selection rule $\tilde{I}_{\mathrm{me}}$ is representation independent (for state spaces consisting of propositional truth assignments and knowledge bases in $\mathit{KB}(\mathscr{L}^{\mathrm{pp}})$) in Halpern and Koller's sense. This result puts Theorem 3.10 in (Halpern & Koller 1995) somewhat into perspective: it shows that this theorem must actually rely on Halpern and Koller's condition that a measure selection function I to qualify as an inference procedure must preserve consistency, i.e. $I(G) \neq \emptyset$ for every $G \neq \emptyset$. This rather severe condition is not satisfied by I_{me}, $\tilde{I}_{\mathrm{me}}$, or most other natural selection rules. Therefore, these rules are outside the scope of theorem 3.10 as stated in (Halpern & Koller 1995). The selection function $\tilde{I}_{\mathrm{me}}$ shows that this theorem can not be extended to selection rules that may violate the preservation of consistency condition.

6 CONCLUSION

We have developed a very general and purely logical framework for analyzing representation independence of nonmonotonic logics. Representation independence does not emerge as a property as clear cut and unequivocal as, say, compactness. Rather, it comes (or fails to come) in many different forms and strengths, in our approach expressed by the parameter $\boldsymbol{F}$. Two specific (classes of) logics have here been studied in some detail: rational closure logic, which has been found to exhibit a degree of representation independence more or less typical for qualitative nonmonotonic logics based on propositional logic, and propositional probabilistic logics, some important instances of which are notorious for their strong representation dependence. We have also seen that these latter logics can be modified so as to make them more representation independent, while at the same time retaining at least some of the desirable properties of the original inference rules.

The specific applications of the general theory given in this paper were rather easy by being restricted to a propositional background and correspondingly simple classes of abstract interpretations. When one moves to logics based on first-order predicate logic things become more complicated, as the concept of an interpretation becomes much more powerful. For this reason one will not be able to obtain results as clean as theorem 5.2 for circumscription, for example.

References

Boutelier, C. (1994), 'Conditional logics of normality: a modal approach', *Artificial Intelligence* **68**, 87–154.

Clark, K. L. (1978), Negation as failure, *in* H. Gallaire & J. Minker, eds, 'Logic and Databases', Plenum Press, New York.

Diaconis, P. & Zabell, S. (1982), 'Updating subjective probability', *Journal of the American Statistical Association* **77**(380), 822–830.

Ebbinghaus, H. D. (1985), Extended logics: The general framework, *in* J. Barwise & S. Feferman, eds, 'Model-Theoretic Logics', Springer-Verlag, pp. 25–76.

Frisch, A. & Haddaway, P. (1994), 'Anytime deduction for probabilistic logic', *Artificial Intelligence* **69**, 93–122.

Gabbay, D. (1985), Theoretical foundations for nonmonotonic reasoning in expert systems, *in* K. Apt, ed., 'Logics and Models of Cuncurrent Systems', Springer-Verlag, Berlin.

Grove, A., Halpern, J. & Koller, D. (1992), Random worlds and maximum entropy, *in* 'Proc. 7th IEEE Symp. on Logic in Computer Science'.

Halpern, J. Y. & Koller, D. (1995), Representation dependence in probabilistic inference, *in* 'Proceedings of IJCAI-95', pp. 1853–1860.

Hodges, W. (1993), *Model Theory*, Cambridge University Press.

Jaeger, M. (1995), Default Reasoning about Probabilities, PhD thesis, Universität des Saarlandes.

Kraus, S., Lehmann, D. & Magidor, M. (1990), 'Nonmonotonic reasoning, preferential models and cumulative logics', *Artificial Intelligence* **44**, 167–207.

Lehmann, D. (1989), What does a conditional knowledge base entail?, *in* 'Proceedings First International Conference on Principles of Knowledge Representation and Reasoning'.

Lehmann, D. & Magidor, M. (1992), 'What does a conditional knowledge base entail?', *Artificial Intelligence* **55**, 1–60.

Makinson, D. (1994), General patterns in nonmonotonic reasoning, *in* 'Handbook of Logic in Artificial Intelligence and Logic Programming', Oxford University Press, pp. 35–110.

Nilsson, N. (1986), 'Probabilistic logic', *Artificial Intelligence* **28**, 71–88.

Paris, J. & Vencovská, A. (1990), 'A note on the inevitability of maximum entropy', *International Journal of Approximate Reasoning* **4**, 183–223.

Paris, J. & Vencovská, A. (1992), 'A method for updating that justifies minimum cross entropy', *International Journal of Approximate Reasoning* **7**, 1–18.

Pearl, J. (1990), System Z: A natural ordering of defaults with tractable applications to default reasoning, *in* 'Proceedings of TARK-90', pp. 121–135.

Reiter, R. (1980), 'A logic for default reasoning', *Artificial Intelligence* **13**, 81–132.

Shoham, Y. (1987), Nonmonotonic logics: Meaning and utility, *in* 'Proceedings of IJCAI-87'.

Spohn, W. (1990), A general non-probabilistic theory of inductive reasoning, *in* R. Shachter, T. Levitt, L. Kanal & J. Lemmer, eds, 'Uncertainty in Artificial Intelligence 4', Elsevier Science Publishers.