Plausibility Measures and Default Reasoning

Nir Friedman
Stanford University
Gates Building 1A
Stanford, CA 94305-9010
nir@cs.stanford.edu

Joseph Y. Halpern
IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120–6099
halpern@almaden.ibm.com

Abstract

In recent years, a number of different semantics for defaults have been proposed, such as preferential structures, ϵ-semantics, possibilistic structures, and κ-rankings, that have been shown to be characterized by the same set of axioms, known as the KLM properties (for Kraus, Lehmann, and Magidor). While this was viewed as a surprise, we show here that it is almost inevitable. We do this by giving yet another semantics for defaults that uses *plausibility measures*, a new approach to modeling uncertainty that generalize other approaches, such as probability measures, belief functions, and possibility measures. We show that all the earlier approaches to default reasoning can be embedded in the framework of plausibility. We then provide a necessary and sufficient condition on plausibilities for the KLM properties to be sound, and an additional condition necessary and sufficient for the KLM properties to be complete. These conditions are easily seen to hold for all the earlier approaches, thus explaining why they are characterized by the KLM properties.

1 Introduction

There have been many approaches to default reasoning proposed in the literature (see (Ginsberg 1987; Gabbay, Hogger, & Robinson 1993) for overviews). The recent literature has been guided by a collection of axioms that have been called the *KLM properties* (since they were discussed in (Kraus, Lehmann, & Magidor 1990)), and many of the recent approaches to default reasoning, including preferential structures (Kraus, Lehmann, & Magidor 1990; Shoham 1987), ϵ-semantics (Adams 1975; Geffner 1992b; Pearl 1989), possibilistic structures (Dubois & Prade 1991), and κ-rankings (Goldszmidt & Pearl 1992; Spohn 1987), have been shown to be characterized by these properties. This has been viewed as somewhat surprising, since these approaches seem to capture quite different intuitions. As Pearl (1989) said of the equivalence between ϵ-semantics and preferential structures, "It is remarkable that two totally different interpretations of defaults yield identical sets of conclusions and identical sets of reasoning machinery."

The goal of this paper is to explain why all these approaches are characterized by the KLM properties. Our key tool is the use of yet another approach for giving semantics to defaults, that makes use of *plausibility measures* (Friedman & Halpern 1995a). A plausibility measure associates with a set a *plausibility*, which is just an element in a partially ordered space. The only property that we require is that the plausibility of a set is at least as great as the plausibility of any of its subsets. Probability distributions, Dempster-Shafer belief functions (Shafer 1976), and possibility measures (Dubois & Prade 1990) are all easily seen to be special cases of plausibility measures. Of more interest to us here is that all the approaches that have been used to give semantics to defaults that can be characterized by the KLM properties can be embedded into the plausibility framework.

In fact, we show much more. All of these approaches can be understood as giving semantics to defaults by considering a class $\mathcal{P}$ of structures (preferential structures, possibilistic structures, etc.). A default d is then said to follow from a knowledge base Δ of defaults if all structures in $\mathcal{P}$ that satisfy Δ also satisfy d. We define a notion of *qualitative plausibility measure*, and show that the KLM properties are sound in a plausibility structure if and only if it is qualitative. Moreover, as long as a class $\mathcal{P}$ of plausibility structures satisfies a minimal richness condition, we show that the KLM properties will completely characterize default reasoning in $\mathcal{P}$. We then show that when we map preferential structures (or possibilistic structures or any of the other structures considered in the literature on defaults) into plausibility structures, we get a class of qualitative structures that is easily seen to satisfy the richness conditioning. This explains why the KLM axioms characterize default reasoning in all these frameworks. Far from being surprising that the KLM axioms are complete in all these cases, it is almost inevitable.

The KLM properties have been viewed as the "conservative core" of default reasoning (Pearl 1989), and much recent effort has been devoted to finding principled methods of going beyond KLM. Our result shows that any approach that gives semantics to defaults with respect to a collection $\mathcal{P}$ of structures will almost certainly not go beyond KLM. This result thus provides added insight into and justification for approaches such as those of (Bacchus *et al.* 1993; Geffner 1992a; Goldszmidt & Pearl 1992; Goldszmidt, Morris, & Pearl 1993; Lehmann & Magidor 1992; Pearl 1990) that, roughly speaking, say d follows from Δ if d is true in a particular structure $P_\Delta \in \mathcal{P}$ that satisfies Δ

(not necessarily all structures in $\mathcal{P}$ that satisfy Δ).

This paper is organized as follows. In Section 2, we review the relevant material from (Friedman & Halpern 1995a) on plausibility measures. In Section 3, we review the KLM properties and various approaches to default reasoning that are characterized by these properties. In Section 4, we show how the various notions of default reasoning considered in the literature can all be viewed as instances of plausibility measures. In Section 5, we define qualitative plausibility structures, show that the KLM properties are sound in a structure if and only if it is qualitative, and provide a weak richness condition that is necessary and sufficient for them to be complete. In Section 6, we discuss how plausibility measures can be used to give semantics to a full logic of conditionals, and compare this with the more traditional approach (Lewis 1973). In the full paper (Friedman & Halpern 1995b) we also consider the relationship between our approach to plausibility and *epistemic entrenchment* (Gärdenfors & Makinson 1988). We conclude in Section 7 with a discussion of other potential applications of plausibility measures.

2 Plausibility Measures

A probability space is a tuple $(W, \mathcal{F}, \mu)$, where W is a set of worlds, $\mathcal{F}$ is an algebra of *measurable* subsets of W (that is, a set of subsets closed under union and complementation to which we assign probability) and μ is a *probability measure*, that is, a function mapping each set in $\mathcal{F}$ to a number in $[0, 1]$ satisfying the well-known Kolmogorov axioms ($\mu(\emptyset) = 0$, $\mu(W) = 1$, and $\mu(A \cup B) = \mu(A) + \mu(B)$ if A and B are disjoint).

A plausibility space is a direct generalization of a probability space. We simply replace the probability measure μ by a *plausibility measure* Pl, which, rather than mapping sets in $\mathcal{F}$ to numbers in $[0, 1]$, maps them to elements in some arbitrary partially ordered set. We read Pl(A) as "the plausibility of set A". If Pl$(A) \leq$ Pl(B), then B is at least as plausible as A. Formally, a *plausibility space* is a tuple $S = (W, \mathcal{F}, \text{Pl})$, where W is a set of worlds, $\mathcal{F}$ is an algebra of subsets of W, and Pl maps the sets in $\mathcal{F}$ to some set D, partially ordered by a relation $\leq_D$ (so that $\leq_D$ is reflexive, transitive, and anti-symmetric). We assume that D is *pointed*: that is, it contains two special elements $\top_D$ and $\perp_D$ such that $\perp_D \leq_D d \leq_D \top_D$ for all $d \in D$; we further assume that Pl$(W) = \top_D$ and Pl$(\emptyset) = \perp_D$. The only other assumption we make is

A1. If $A \subseteq B$, then Pl$(A) \leq_D$ Pl(B).

Thus, a set must be at least as plausible as any of its subsets.

Some brief remarks on the definition: We have deliberately suppressed the domain D from the tuple S, since the choice of D is not significant in this paper. All that matters is the ordering induced by $\leq_D$ on the subsets in $\mathcal{F}$. The algebra $\mathcal{F}$ also does not play a significant role in this paper; for our purposes, it suffices to take $\mathcal{F} = 2^W$. We have chosen to allow the generality of having an algebra of measurable sets to make it clear that plausibility spaces generalize probability spaces. For ease of exposition, we omit the $\mathcal{F}$ from here on in, always taking it to be 2^W, and just denote a plausibility space as a pair (W, Pl). As usual, we define the ordering $<_D$ by taking $d_1 <_D d_2$ if $d_1 \leq_D d_2$ and $d_1 \neq d_2$. We omit the subscript S from $\leq_D$, $<_D$, $\top_D$ and $\perp_D$ whenever it is clear from context.

Clearly plausibility spaces generalize probability spaces. We now briefly discuss a few other notions of uncertainty that they generalize:

- A *belief function Bel* on W is a function *Bel* : $2^W \rightarrow [0, 1]$ satisfying certain axioms (Shafer 1976). These axioms certainly imply property A1, so a belief function is a plausibility measure.
- A *fuzzy measure* (or a *Sugeno measure*) f on W (Wang & Klir 1992) is a function $f : 2^W \mapsto [0, 1]$, that satisfies A1 and some continuity constraints. A *possibility measure* (Dubois & Prade 1990) Poss is a fuzzy measure such that Poss$(W) = 1$, Poss$(\emptyset) = 0$, and Poss$(A) = \sup_{w \in A}(\text{Poss}(\{w\})$.
- An *ordinal ranking* (or *κ-ranking*) κ on W (as defined by (Goldszmidt & Pearl 1992), based on ideas that go back to (Spohn 1987)) is a function mapping subsets of W to $I\!N^* = I\!N \cup \{\infty\}$ such that $\kappa(W) = 0$, $\kappa(\emptyset) = \infty$, and $\kappa(A) = \min_{w \in A}(\kappa(\{w\}))$. Intuitively, an ordinal ranking assigns a degree of surprise to each subset of worlds in W, where 0 means unsurprising and higher numbers denote greater surprise. It is easy to see that if κ is a ranking on W, then (W, κ) is a plausibility space, where $x \leq_{I\!N^*} y$ if and only if $y \leq x$ under the usual ordering on the ordinals.
- A *preference ordering* on W is a partial order $\prec$ over W (Kraus, Lehmann, & Magidor 1990; Shoham 1987). Intuitively, $w \prec w'$ holds if w is *preferred* to w'.[1] Preference orders have been used to provide semantics for *default* (i.e., conditional) statements. In Section 4 we show how to map preference orders on W to plausibility measures on W in a way that preserves the ordering of events of the form $\{w\}$ as well as the truth values of defaults.
- A *parametrized probability distribution* (PPD) on W is a sequence $\{\text{Pr}_i : i \geq 0\}$ of probability measures over W. Such structures provide semantics for defaults in *ϵ-semantics* (Pearl 1989; Goldszmidt, Morris, & Pearl 1993). In Section 4 we show how to map PPDs into plausibility structures in a way that preserves the truth-values of conditionals.

Plausibility structures are motivated by much the same concerns as two other recent symbolic generalizations of probability by Darwiche (1992) and Weydert (1994). Their approaches have a great deal more structure though. They start with a domain D and several algebraic operations that have properties similar to the usual arithmetic operations (e.g., addition and multiplication) over $[0, 1]$. The result

[1]We follow the standard notation for preference here (Lewis 1973; Kraus, Lehmann, & Magidor 1990), which uses the (perhaps confusing) convention of placing the more likely world on the left of the $\prec$ operator.

is an algebraic structure over the domain D that satisfies various properties. Their structures are also general enough to capture all of the examples above except preferential orderings. These orderings cannot be captured precisely because of the additional structure. Moreover, as we shall see, by starting with very little structure and adding just what we need, we can sometimes bring to light issues that may be obscured in richer frameworks. We refer the interested reader to (Friedman & Halpern 1995a) for a more detailed comparison to (Darwiche 1992; Weydert 1994).

3 Approaches to Default Reasoning: A Review

Defaults are statements of the form "if φ then typically/likely/by default ψ", which we denote $\varphi{\rightarrow}\psi$. For example, the default "birds typically fly" is represented *Bird*$\rightarrow$*Fly*. There has been a great deal of discussion in the literature as to what the appropriate semantics of defaults should be, and what new defaults should by entailed by a knowledge base of defaults. For the most part, we do not get into these issues here. While there has been little consensus on what the "right" semantics for defaults should be, there has been some consensus on a reasonable "core" of inference rules for default reasoning. This core, known as the KLM properties, was suggested by (Kraus, Lehmann, & Magidor 1990), and consists of the following axiom and rules of inference (where we use $\Rightarrow$ to denote material implication):

LLE. From $\varphi \Leftrightarrow \varphi'$ and $\varphi{\rightarrow}\psi$ infer $\varphi'{\rightarrow}\psi$ (left logical equivalence)

RW. From $\psi \Rightarrow \psi'$ and $\varphi{\rightarrow}\psi$ infer $\varphi{\rightarrow}\psi'$ (right weakening)

REF. $\varphi{\rightarrow}\varphi$ (reflexivity)

AND. From $\varphi{\rightarrow}\psi_1$ and $\varphi{\rightarrow}\psi_2$ infer $\varphi{\rightarrow}\psi_1 \wedge \psi_2$

OR. From $\varphi_1{\rightarrow}\psi$ and $\varphi_2{\rightarrow}\psi$ infer $\varphi_1 \vee \varphi_2{\rightarrow}\psi$

CM. From $\varphi{\rightarrow}\psi_1$ and $\varphi{\rightarrow}\psi_2$ infer $\varphi \wedge \psi_1{\rightarrow}\psi_2$ (cautious monotonicity)

LLE states that the syntactic form of the antecedent is irrelevant. Thus, if φ_1 and φ_2 are equivalent, we can deduce $\varphi_2{\rightarrow}\psi$ from $\varphi_1{\rightarrow}\psi$. RW describes a similar property of the consequent: If ψ (logically) entails ψ', then we can deduce $\varphi{\rightarrow}\psi'$ from $\varphi{\rightarrow}\psi$. This allows us to can combine default and logical reasoning. REF states that φ is always a default conclusion of φ. AND states that we can combine two default conclusions: If we can conclude by default both ψ_1 and ψ_2 from φ, we can also conclude $\psi_1 \wedge \psi_2$ from φ. OR states that we are allowed to reason by cases: If the same default conclusion follows from each of two antecedents, then it also follows from their disjunction. CM states that if ψ_1 and ψ_2 are two default conclusions of φ, then discovering that ψ_1 holds when φ holds (as would be expected, given the default) should not cause us to retract the default conclusion ψ_2.

This system of rules is called system **P** in (Kraus, Lehmann, & Magidor 1990). The notation $\Delta \vdash_{\mathbf{P}} \varphi{\rightarrow}\psi$ denotes that $\varphi{\rightarrow}\psi$ can be deduced from Δ using these inference rules.

There are a number of well-known semantics for defaults that are characterized by these rules. We sketch a few of them here, referring the reader to the original references for further details and motivation. All of these semantics involve structures of the form (W, X, π), where W is a set of possible worlds, $\pi(w)$ is a truth assignment to primitive propositions, and X is some "measure" on W such as a preference ordering, a κ-ranking, or a possibility measure. We define a little notation that will simplify the discussion below. Given a structure (W, X, π), we take $[\![\varphi]\!] \subseteq W$ to be the set of of worlds satisfying φ, i.e., $[\![\varphi]\!] = \{w \in W : \pi(w)(\varphi) = \mathbf{true}\}$.

The first semantic proposal was provided by Kraus, Lehmann and Magidor (1990), using ideas that go back to (Lewis 1973; Shoham 1987). Recall that a preference ordering on W is partial order (i.e., irreflexive and transitive relation) $\prec$ over W. A *preferential structure* is a tuple $(W, \prec, \pi)$, where $\prec$ is a partial order on W.[2] The intuition (Shoham 1987) is that a preferential structure satisfies a conditional $\varphi{\rightarrow}\psi$ if all the the most preferred worlds (i.e., the minimal worlds according to $\prec$) in $[\![\varphi]\!]$ satisfy ψ. However, there may be no minimal worlds in $[\![\varphi]\!]$. This can happen if $[\![\varphi]\!]$ contains an infinite descending sequence $\ldots \prec w_2 \prec w_1$. What do we do in these structures? There are a number of options: the first is to assume that, for each formula φ, there are minimal worlds in $[\![\varphi]\!]$; this is the assumption actually made in (Kraus, Lehmann, & Magidor 1990), where it is called the *smoothness* assumption. A yet more general definition—one that works even if $\prec$ is not smooth—is given in (Lewis 1973; Boutilier 1994). Roughly speaking, $\varphi{\rightarrow}\psi$ is true if, from a certain point on, whenever φ is true, so is ψ. More formally,

> $(W, \prec, \pi)$ satisfies $\varphi{\rightarrow}\psi$, if for every world $w_1 \in [\![\varphi]\!]$, there is a world w_2 such that (a) $w_2 \preceq w_1$ (so that w_2 is at least as normal as w_1), (b) $w_2 \in [\![\varphi \wedge \psi]\!]$, and (c) for all worlds $w_3 \prec w_2$, we have $w_3 \in [\![\varphi \Rightarrow \psi]\!]$ (so any world more normal than w_2 that satisfies φ also satisfies ψ).

It is easy to verify that this definition is equivalent to the earlier one if $\prec$ is smooth. A knowledge-base Δ *preferentially entails* $\varphi{\rightarrow}\psi$, denoted $\Delta \models_{\mathrm{p}} \varphi{\rightarrow}\psi$, if every preferential structure that satisfies (all the defaults in) Δ also satisfies $\varphi{\rightarrow}\psi$.

Lehmann and Magidor show that preferential entailment is characterized by system **P**.

[2] We note that the formal definition of preferential structures in (Kraus, Lehmann, & Magidor 1990; Lehmann & Magidor 1992) is slightly more complex. Kraus, Lehmann, and Magidor distinguish between *states* and *worlds*. In their definition, a preferential structure is an ordering over states together with a labeling function that maps states to worlds. They take worlds to be truth assignments to primitive propositions. Our worlds thus correspond to states in their terminology, since we allow two worlds $w \neq w'$ such that $\pi(w) = \pi(w')$. Despite these minor differences, all the results that we prove for our version of preferential structures hold (with almost no change in proof) for theirs.

Theorem 3.1: (Lehmann & Magidor 1992; Boutilier 1994) $\Delta \models_{\rm p} \varphi \rightarrow \psi$ *if and only if* $\Delta \vdash_{\mathbf{P}} \varphi \rightarrow \psi$.

Thus, reasoning with preferential structures corresponds in a precise sense to reasoning with the core properties of default reasoning.

As we mentioned earlier, we usually want to add additional inferences to those sanctioned by the core. Lehmann and Magidor (1992) hoped to do so by limiting attention to a special class of preferential structures. A preferential structure $(W, \prec, \pi)$ is *rational* if $\prec$ is a *modular* order, so that for all worlds $u, v, w \in W$, if $w \prec v$, then either $u \prec v$ or $w \prec u$. It is not hard to show that modularity implies that possible worlds are clustered into equivalence classes, each class consisting of worlds that are incomparable to one another, with these classes being totally ordered. Thus, rational structures form a "well-behaved" subset of preferential structures. Unfortunately, Lehmann and Magidor showed that restricting to rational structures gives no additional properties (at least, as far as the limited language of defaults is concerned). We say that a knowledge base Δ *rationally entails* $\varphi \rightarrow \psi$, denoted $\Delta \models_{\rm r} \varphi \rightarrow \psi$, if every rational structure that satisfies Δ also satisfies $\varphi \rightarrow \psi$.[3]

Theorem 3.2: (Lehmann & Magidor 1992) $\Delta \models_{\rm r} \varphi \rightarrow \psi$ *if and only if* $\Delta \vdash_{\mathbf{P}} \varphi \rightarrow \psi$.

Thus, we do not gain any new patterns of default inference when we restrict our attention to rational structures.

Pearl (1989) considers a semantics for defaults grounded in probability, using an approach due to Adams (1975). In this approach, a default $\varphi \rightarrow \psi$ denotes that $\Pr(\psi \mid \varphi)$ is extremely high, i.e., almost 1. Roughly speaking, a collection Δ of defaults implies a default $\varphi \rightarrow \psi$ if we can ensure that $\Pr(\varphi | \psi)$ is arbitrarily close to 1, by taking the probabilities of the defaults in Δ to be sufficiently high.

The formal definition needs a bit of machinery.[4] Recall that a PPD on W is a sequence $\{\Pr_i : i \geq 0\}$ of probability measures over W. A *PPD structure* is a tuple $(W, \{\Pr_i : i \geq 0\}, \pi)$, where $\{\Pr_i\}$ is PPD on W. Intuitively, it satisfies a conditional $\varphi \rightarrow \psi$ if the conditional probability ψ given φ goes to 1 in the limit. Formally, $\varphi \rightarrow \psi$ is satisfied if $\lim_{i \rightarrow \infty} \Pr_i([\![\psi]\!] | [\![\psi]\!]) = 1$ (Goldszmidt, Morris, & Pearl 1993) (where $\Pr_i([\![\psi]\!] | [\![\varphi]\!])$ is taken to be 1 if $\Pr_i([\![\varphi]\!]) = 0$). Δ *ϵ-entails* $\varphi \rightarrow \psi$, denoted $\Delta \models_{\epsilon} \varphi \rightarrow \psi$, if every PPD structure that satisfies all the defaults in Δ also satisfies $\varphi \rightarrow \psi$. Surprisingly, Geffner shows that ϵ-entailment is equivalent to preferential entailment.

Theorem 3.3: (Geffner 1992b) $\Delta \models_{\epsilon} \varphi \rightarrow \psi$ *if and only if* $\Delta \vdash_{\mathbf{P}} \varphi \rightarrow \varphi$.[5]

Possibility measures and ordinal rankings provide two more semantics for defaults. A *possibility structure* is a tuple $PS = (W, \text{Poss}, \pi)$ such that Poss is a possibility measure on W. We say $PS \models_{Poss} \varphi \rightarrow \psi$ if either $\text{Poss}([\![\varphi]\!]) = 0$ or $\text{Poss}([\![\varphi \wedge \psi]\!]) > \text{Poss}([\![\varphi \wedge \neg\psi]\!])$. Intuitively, $\varphi \rightarrow \psi$ holds vacuously if φ is impossible; otherwise, it holds if $\varphi \wedge \psi$ is more "possible" than $\varphi \wedge \neg\psi$. For example, *Bird*$\rightarrow$*Fly* is satisfied when *Bird* $\wedge$ *Fly* is more possible than *Bird* $\wedge \neg$*Fly*. Similarly, an *ordinal ranking structure* is a tuple $R = (W, \kappa, \pi)$ if κ is a an ordinal ranking on W. We say that $R \models_{\kappa} \varphi \rightarrow \psi$ if either $\kappa([\![\varphi]\!]) = \infty$ or $\kappa([\![\varphi \wedge \psi]\!]) < \kappa([\![\varphi \wedge \neg\psi]\!])$. We say that Δ *possibilistically entails* $\varphi \rightarrow \psi$, denoted $\Delta \models_{Poss} \varphi \rightarrow \psi$ (resp., Δ *κ-entails* $\varphi \rightarrow \psi$, denoted $\Delta \models_{\kappa} \varphi \rightarrow \psi$) if all possibility structures (resp., all ordinal ranking structures) that satisfy Δ also satisfy $\varphi \rightarrow \psi$.

These two approaches are again characterized by the KLM properties.

Theorem 3.4: (Geffner 1992b; Dubois & Prade 1991) $\Delta \models_{Poss} \varphi \rightarrow \psi$ *if and only if* $\Delta \models_{\kappa} \varphi \rightarrow \psi$ *if and only if* $\Delta \vdash_{\mathbf{P}} \varphi \rightarrow \psi$.

Why do we always seem to end up with the KLM properties? As we are about to show, thinking in terms of plausibility measures provides the key to understanding this issue.

4 Default Reasoning Using Plausibility

We can give semantics to defaults using plausibility measures much as we did using possibility measures. A *plausibility structure* is a tuple $PL = (W, \text{Pl}, \pi)$, where (W, Pl) is a plausibility space and π maps each possible world to a truth assignment. We define $PL \models_{PL} \varphi \rightarrow \psi$ if either $\text{Pl}([\![\varphi]\!]) = \perp$ or $\text{Pl}([\![\varphi \wedge \psi]\!]) > \text{Pl}([\![\varphi \wedge \neg\psi]\!])$.

Notice that if Pl is a probability function Pr, then $\varphi \rightarrow \psi$ holds exactly if either $\Pr([\![\varphi]\!]) = 0$ or $\Pr([\![\psi]\!] | [\![\varphi]\!]) > 1/2$. How does this semantics for defaults compare to others that have been given in the literature? It is immediate from the definitions that the semantics we give to defaults in possibility structures is the same as that given to them if we view these possibility structures as plausibility structures (using the obvious mapping described above, and similarly for ordinal ranking structures. What about preferential structures and PPD structures? Can we map them into plausibility structures while still preserving the semantics of defaults? As we now show, we can.

Theorem 4.1: **(a)** *Let* $\prec$ *be a preference ordering on* W. *There is a plausibility measure* $\text{Pl}_{\prec}$ *on* W *such that* $(W, \prec, \pi) \models_{p} \varphi \rightarrow \psi$ *iff* $(W, \text{Pl}_{\prec}, \pi) \models_{PL} \varphi \rightarrow \psi$.

(b) *Let* $PP = \{\Pr_i\}$ *be a PPD on* W. *There is a plausibility measure* Pl_{PP} *on* W *such that* $(W, \{\Pr_i\}, \pi) \models_{\epsilon} \varphi \rightarrow \psi$ *iff* $(W, \text{Pl}_{PP}, \pi) \models_{PL} \varphi \rightarrow \psi$.

Proof: We first sketch the proof of part (a). Let $\prec$ be a preference order on W. We define a plausibility measure $\text{Pl}_{\prec}$ on W as follows. Let D_0 be the domain of plausibility values consisting of one element d_w for every element $w \in W$. We use $\prec$ to determine the order of these elements: $d_v < d_w$ if $w \prec v$. (Recall that $w \prec w'$ denotes that w is preferred to w'.) We then take D to be the smallest set containing D_0 closed under least upper bounds (so that

[3] Rational entailment should not be confused with the notion of *rational closure*, also defined by Lehmann and Magidor.

[4] We adopt the presentation used in (Goldszmidt, Morris, & Pearl 1993).

[5] Geffner's result is stated in terms of the original formulation of ϵ-entailment, as described in (Pearl 1989). However, results of (Goldszmidt, Morris, & Pearl 1993) show that the formulation we describe here is equivalent to the original one.

every set of elements in D has a least upper bound in D). In the full paper, we show that this construction results in the following ordering over subsets of W:

> $\mathrm{Pl}_{\prec}(A) \leq \mathrm{Pl}_{\prec}(B)$ if and only if for all $w \in A - B$, there is a world $w' \in B$ such that $w' \prec w$ and there is no $w'' \in A - B$ such that $w'' \prec w'$.

It is not hard to show that $\mathrm{Pl}_{\prec}$ satisfies the requirements of the theorem.

We next sketch the proof of part (b). Let $PP = \{\mathrm{Pr}_1, \mathrm{Pr}_2, \ldots, \}$ be a PPD on W. We define Pl_{PP} so that $\mathrm{Pl}_{PP}(A) \leq \mathrm{Pl}_{PP}(B)$ iff $\lim_{i\to\infty} \mathrm{Pr}_i(B|A \cup B) = 1$. It is easy to see that this definition uniquely determines the effect of Pl_{PP}. It is also easy to show that Pl_{PP} satisfies the requirements of the theorem. ∎

Thus, each of the semantic approaches to default reasoning that were described above can be mapped into plausibility structures in a way that preserves the semantics of defaults.

5 Default Entailment in Plausibility Structures

In this section we characterize default entailment in plausibility structures. To do so, it is useful to have a somewhat more general definition of entailment in plausibility structures.

Definition 5.1: If $\mathcal{P}$ is a class of plausibility structures, then a knowledge base Δ entails $\varphi{\rightarrow}\psi$ with respect to $\mathcal{P}$, denoted $\Delta \models_{\mathcal{P}} \varphi{\rightarrow}\psi$, if every $PL \in \mathcal{P}$ that satisfies Δ also satisfies $\varphi{\rightarrow}\psi$. ∎

The classes of structures we are interested in include $\mathcal{P}^{PL}$, the class of all plausibility structures, and $\mathcal{P}^{Poss}$, $\mathcal{P}^{\kappa}$, $\mathcal{P}^{p}$, $\mathcal{P}^{r}$, and $\mathcal{P}^{\epsilon}$, the classes that arise from mapping possibility structures, ordinal ranking structures, preferential structures, rational structures, and PPDs, respectively, into plausibility structures. (In the case of possibility structures and ordinal ranking structures, the mapping is the obvious one discussed in Section 2; in the case of preferential and rational structures and PPDs, the mapping is the one described in Theorem 4.1.) Recall that all these mappings preserve the semantics of defaults.

It is easy to check that our semantics for defaults does *not* guarantee that the axioms of system **P** hold in all structures in $\mathcal{P}^{PL}$. In particular, they do not hold in probability structures. It is easy to construct a plausibility structure PL where Pl is actually a probability measure Pr such that $\mathrm{Pr}([\![p \wedge q]\!]) > 0$, $\mathrm{Pr}([\![q]\!]|[\![p]\!]) > .5$, $\mathrm{Pr}([\![r]\!]|[\![p]\!]) > .5$, but $\mathrm{Pr}([\![q \wedge r]\!]|[\![p]\!]) < .5$ and $\mathrm{Pr}([\![r]\!]|[\![p \wedge q]\!]) < .5$. Recall that if $\mathrm{Pr}(\varphi) > 0$ then $\varphi{\rightarrow}\psi$ holds if and only if $\mathrm{Pr}(\psi|\varphi) > .5$. Thus, $PL \models_{PL} (\mathit{true}{\rightarrow}p) \wedge (\mathit{true}{\rightarrow}q)$, but $PL \not\models_{PL} \mathit{true}{\rightarrow}p \wedge q$ and $PL \not\models_{PL} (\mathit{true} \wedge p{\rightarrow}q)$. This gives us a violation of both AND and CM. We can similarly construct a counterexample to OR. On the other hand, as the following result shows, plausibility structures do satisfy the other three axioms of system **P**. Let system **P**$'$ be the system consisting of LLE, RW, and REF.

Theorem 5.2: *If $\Delta \vdash_{\mathbf{P}'} \varphi{\rightarrow}\psi$, then $\Delta \models_{\mathcal{P}^{PL}} \varphi{\rightarrow}\psi$.*

What extra conditions do we have to place on plausibility structures to ensure that AND, OR, and CM are satisfied? We focus first on the AND rule. We want an axiom that cuts out probability functions, but leaves more qualitative notions. Working at a semantic level, taking $[\![\varphi]\!] = A$, $[\![\psi_1]\!] = B_1$, and $[\![\psi_2]\!] = B_2$, and using $\overline{X}$ to denote the complement of X, the AND rule translates to:

A2$'$. For all sets A, B_1, and B_2, if $\mathrm{Pl}(A \cap B_1) > \mathrm{Pl}(A \cap \overline{B_1})$ and $\mathrm{Pl}(A \cap B_2) > \mathrm{Pl}(A \cap \overline{B_2})$, then $\mathrm{Pl}(A \cap B_1 \cap B_2) > \mathrm{Pl}(A \cap \overline{(B_1 \cap B_2)})$.

It turns out that in the presence of A1, the following somewhat simpler axiom is equivalent to A2$'$:

A2. If A, B, and C are pairwise disjoint sets, $\mathrm{Pl}(A \cup B) > \mathrm{Pl}(C)$, and $\mathrm{Pl}(A \cup C) > \mathrm{Pl}(B)$, then $\mathrm{Pl}(A) > \mathrm{Pl}(B \cup C)$.

Proposition 5.3: *A plausibility measure satisfies A2 if and only if it satisfies A2$'$.*

A2 can be viewed as a generalization of a natural requirement of qualitative plausibility: if A, B, and C are pairwise disjoint, $\mathrm{Pl}(A) > \mathrm{Pl}(B)$, and $\mathrm{Pl}(A) > \mathrm{Pl}(C)$, then $\mathrm{Pl}(A) > \mathrm{Pl}(B \cup C)$. Moreover, since A2 is equivalent to A2$'$, and A2$'$ is a direct translation of the AND rule into conditions on plausibility measures, any plausibility structure whose plausibility measure satisfies A2 satisfies the AND rule. Somewhat surprisingly, a plausibility measure Pl that satisfies A2 satisfies CM. Moreover, Pl satisfies the non-vacuous case of the OR rule. That is, if $\mathrm{Pl}([\![\varphi_1]\!]) > \perp$, then from $\varphi_1{\rightarrow}\psi$ and $\varphi_2{\rightarrow}\psi$ we can conclude $(\varphi_1 \vee \varphi_2){\rightarrow}\psi$.[6] To handle the vacuous case of OR we need an additional axiom:

A3. If $\mathrm{Pl}(A) = \mathrm{Pl}(B) = \perp$, then $\mathrm{Pl}(A \cup B) = \perp$.

Thus, A2 and A3 capture the essence of the KLM properties. To make this precise, define a plausibility space (W, Pl) to be *qualitative* if it satisfies A2 and A3. We say $PL = (W, \mathrm{Pl}, \pi)$ is a *qualitative plausibility structure* if (W, Pl) is a qualitative plausibility space. Let $\mathcal{P}^{QPL}$ consist of all qualitative plausibility structures.

Theorem 5.4: *If $\mathcal{P} \subseteq \mathcal{P}^{QPL}$, then for all Δ, φ and ψ, if $\Delta \vdash_{\mathbf{P}} \varphi{\rightarrow}\psi$, then $\Delta \models_{\mathcal{P}} \varphi{\rightarrow}\psi$.*

Thus, the KLM axioms are sound for qualitative plausibility structures. We remark that Theorem 5.4 provides, in a precise sense, not only a sufficient but a necessary condition for a set of preferential structures to satisfy the KLM properties. As we show in the full paper, if the KLM axioms are sound with respect to $\mathcal{P}$, then even if there is a structure

[6] We remark that if we dropped requirement A1, then we can define properties of plausibilities measures that correspond precisely to CM and OR. The point is that in the presence of A1, A2—which essentially corresponds to AND—implies CM and the non-vacuous case of OR. Despite appearances, A1 does *not* correspond to RW. Semantically, RW says that if A and B are disjoint sets such that $\mathrm{Pl}(A) > \mathrm{Pl}(B)$, and $A \subseteq A'$, $B' \subseteq B$, and A' and B' are disjoint, then $\mathrm{Pl}(A') > \mathrm{Pl}(B')$. While this follows from A1, it is much weaker than A1.

$P = (W, \mathrm{Pl}, \pi) \in \mathcal{P}$ that is not qualitative, P is "essentially qualitative" for all practical purposes. More precisely, we can show that Pl$'$, the restriction of Pl to sets of the form $[\![\varphi]\!]$ is qualitative.

This, of course, leads to the question of which plausibility structures are qualitative. All the ones we have been focusing on are.

Theorem 5.5: *Each of $\mathcal{P}^{Poss}$, $\mathcal{P}^{\kappa}$, $\mathcal{P}^{\epsilon}$, $\mathcal{P}^{p}$, and $\mathcal{P}^{r}$ is a subset of $\mathcal{P}^{QPL}$.*

It follows from Theorems 5.4 and 5.5 that the KLM properties hold in all the approaches to defaults considered in Section 3. While this fact was already known, this result gives us a deeper understanding as to *why* the KLM properties should hold. In a precise sense, it is because A2 and A3 holds for all these approaches. In the full paper we also show that each of the classes considered in Theorem 5.5 is, in a nontrivial sense, a subset of $\mathcal{P}^{QPL}$; this remains true even if we restrict to totally ordered plausibility measures in the case of $\mathcal{P}^{Poss}$ and $\mathcal{P}^{\kappa}$.[7]

We now turn to the problem of completeness. To get soundness we had to ensure that $\mathcal{P}$ did not contain too many structures, in particular, no structures that are not qualitative. To get completeness we have to ensure that $\mathcal{P}$ contains "enough" structures. In particular, if $\Delta \not\vdash_{\mathbf{P}} \varphi \rightarrow \psi$, we want to ensure that there is a plausibility structure $PL \in \mathcal{P}$ such that $PL \models_{PL} \Delta$ and $PL \not\models_{PL} \varphi \rightarrow \psi$. The following weak condition on $\mathcal{P}$ does this.

Definition 5.6: We say that $\mathcal{P}$ is *rich* if for every collection $\varphi_1, \ldots, \varphi_n$ of mutually exclusive formulas, there is a plausibility structure $PL = (W, \mathrm{Pl}, \pi) \in \mathcal{P}$ such that:

$$\perp = \mathrm{Pl}([\![\varphi_1]\!]) < \mathrm{Pl}([\![\varphi_2]\!]) < \cdots < \mathrm{Pl}([\![\varphi_n]\!]). \quad \blacksquare$$

The requirement of richness is quite mild. It says that we do not have *a priori* constraints on the relative plausibilities of a collection of disjoint sets. Certainly every collection of plausibility measures we have considered thus far can be easily shown to satisfy this richness condition.

Theorem 5.7: *Each of $\mathcal{P}^{Poss}$, $\mathcal{P}^{\kappa}$, $\mathcal{P}^{p}$, $\mathcal{P}^{r}$, $\mathcal{P}^{\epsilon}$, and $\mathcal{P}^{QPL}$ is rich.*

More importantly, richness is a necessary and sufficient condition to ensure that the KLM properties are complete.

Theorem 5.8: *A set $\mathcal{P}$ of qualitative plausibility structures is rich if and only if for all Δ and defaults $\varphi \rightarrow \psi$, we have that $\Delta \models_{\mathcal{P}} \varphi \rightarrow \psi$ implies $\Delta \vdash_{\mathbf{P}} \varphi \rightarrow \psi$.*

Putting together Theorems 5.4, 5.5, and 5.8, we get

Corollary 5.9: *For $\mathcal{P} \in \{\mathcal{P}^{Poss}, \mathcal{P}^{\kappa}, \mathcal{P}^{p}, \mathcal{P}^{r}, \mathcal{P}^{\epsilon}, \mathcal{P}^{QPL}\}$, and all Δ, φ, and ψ, we have $\Delta \vdash_{\mathbf{P}} \varphi \rightarrow \psi$ if and only if $\Delta \models_{\mathcal{P}} \varphi \rightarrow \psi$.*

Not only does this result gives us a straightforward and uniform proof that the KLM properties characterize default reasoning in each of the systems considered in Section 3, it gives us a general technique for proving completeness of the KLM properties for other semantics as well. All we have to do is to provide a mapping of the intended semantics into plausibility structures, which is usually straightforward, and then show that the resulting set of structures is qualitative and rich.

Theorem 5.8 also has important implications for attempts to go beyond the KLM properties (as was the goal in introducing rational structures). It says that any semantics for defaults that proceeds by considering a class $\mathcal{P}$ of structures satisfying the richness constraint, and defining $\Delta \models_{\mathcal{P}} \varphi \rightarrow \psi$ to hold if $\varphi \rightarrow \psi$ is true in every structure in $\mathcal{P}$ that satisfies Δ cannot lead to new properties for entailment.

Thus, to go beyond KLM, we either need to consider interesting non-rich classes of structures, or to define a notion of entailment that does not amount to considering what holds in all the structures of a given class. It is possible to construct classes of structures that are arguably interesting and violate the richness constraint. One way is to impose independence constraints. For example, suppose we consider all structures where p is independent of q in the sense that *true*$\rightarrow q$ holds if and only if $p \rightarrow q$ holds if and only if $\neg p \rightarrow q$ holds, so that discovering either p or $\neg p$ does not affect whether or not q is believed.[8] Restricting to such structures clearly gives us extra properties. For example, from *true*$\rightarrow q$ we can infer $p \rightarrow q$, which certainly does not follow from the KLM properties. Such structures do not satisfy the richness constraint, since we cannot have, for example, $\mathrm{Pl}([\![p \wedge q]\!]) > \mathrm{Pl}([\![p \wedge \neg q]\!]) > \mathrm{Pl}([\![\neg p \wedge \neg q]\!]) > \mathrm{Pl}([\![\neg p \wedge q]\!])$. Much of the recent work in default reasoning (Bacchus *et al.* 1993; Geffner 1992a; Goldszmidt & Pearl 1992; Goldszmidt, Morris, & Pearl 1993; Lehmann & Magidor 1992; Pearl 1990) has taken the second approach, of not looking at entailment with respect to a class of structures. Roughly speaking, these approaches can be viewed as taking the basic idea of preferential semantics—placing a preference ordering on worlds—one step further: They try to get from a knowledge base one preferred *structure* (where the structure itself puts a preference ordering on worlds)—for example, in (Goldszmidt, Morris, & Pearl 1993), the PPD of maximum entropy is considered—and carry out all reasoning in that preferred structure. We believe that plausibility measures will provide insight into techniques for choosing such preferred structures, particularly through the use of independence, but the discussion of this issue is beyond the scope of this paper.

6 A Logic of Defaults

Up to now, we have just focused on whether a set of defaults implies another default. We have not considered a full logic of defaults, with negated defaults, nested defaults, and dis-

[7] Since, for example, the range of a possibility measure is [0,1], there are totally ordered plausibility measures that are not possibility measures, although they may put the same ordering on sets. However, for example, we can have a qualitative plausibility measure on $\{1,2\}$ such that $\mathrm{Pl}(\{1\}) = \mathrm{Pl}(\{2\}) < \mathrm{Pl}(\{1,2\})$. This cannot correspond to a possibility measure, since $\mathrm{Poss}(\{1\}) = \mathrm{Poss}(\{2\})$ would imply that $\mathrm{Poss}(\{1\}) = \mathrm{Poss}(\{1,2\})$.

[8] We remark that if we define independence appropriately in plausibility structures, this property does indeed hold; see (Friedman & Halpern 1995a).

junctions of defaults. It is easy to extend all the approaches we have defined so far to deal with such a logic. *Conditional logic* is a logic that treats $\rightarrow$ as a modal operator. The syntax of the logic is simple: let $\mathcal{L}^C$ be the language defined by starting with primitive propositions, and close off under $\wedge$, $\neg$, and $\rightarrow$. Formulas can describe logical combination of defaults (e.g., $(p \rightarrow q) \vee (p \rightarrow \neg q)$) as well as nested defaults (e.g., $(p \rightarrow q) \rightarrow r$).

The semantics of conditional logic is similar to the semantics of defaults.[9] The usual definition (Lewis 1973) associates with each world a preferential order over worlds. We now give a similar definition based on plausibility measures. Given a preferential structure $PL = (W, \mathrm{Pl}, \pi)$, we define what it means for a formula φ to be true at a world w in *PL*. The definition for the propositional connectives is standard, and for $\rightarrow$, we use the definition already given:

- $(PL, w) \models p$ if $\pi(w) \models p$ for a primitive proposition p
- $(PL, w) \models \neg\varphi$ if $(PL, w) \not\models \varphi$
- $(PL, w) \models \varphi \wedge \psi$ if $(PL, w) \models \varphi$ and $(PL, w) \models \psi$
- $(PL, w) \models \varphi \rightarrow \psi$ if either $\mathrm{Pl}([\![\varphi]\!]_{PL}) = \perp$ or $\mathrm{Pl}([\![\varphi \wedge \psi]\!]_{PL}) > \mathrm{Pl}([\![\varphi \wedge \neg\psi]\!]_{PL})$, where we define $[\![\varphi]\!]_{PL} = \{w \in W : (\mathrm{Pl}, w) \models \varphi\}$.[10]

We now want to axiomatize default reasoning in this framework. Clearly we need axioms and inference rules that generalize those of system **P**. Let $N\varphi$ be an abbreviation for $\neg\varphi \rightarrow$*false*. (This operator is called the *outer modality* in (Lewis 1973).) Expanding the definition of $\rightarrow$, we get that $N\varphi$ holds at w if and only if $\mathrm{Pl}([\![\neg\varphi]\!]) = \perp$. Thus, $N\varphi$ holds if $\neg\varphi$ is considered completely implausible. Thus, it implies that φ is true "almost everywhere". Let system **C** be the system consisting of LLE, RW, and the following axioms and inference rules:

C0. All the propositional tautologies
C1. $\varphi \rightarrow \varphi$
C2. $((\varphi \rightarrow \psi_1) \wedge (\varphi \rightarrow \psi_2)) \Rightarrow (\varphi \rightarrow (\psi_1 \wedge \psi_2))$
C3. $((\varphi_1 \rightarrow \psi) \wedge (\varphi_2 \rightarrow \psi)) \Rightarrow ((\varphi_1 \vee \varphi_2) \rightarrow \psi)$
C4. $((\varphi_1 \rightarrow \varphi_2) \wedge (\varphi_1 \rightarrow \psi)) \Rightarrow ((\varphi_1 \wedge \varphi_2) \rightarrow \psi)$
C5. $[(\varphi \rightarrow \psi) \Rightarrow N(\varphi \rightarrow \psi)] \wedge [\neg(\varphi \rightarrow \psi) \Rightarrow N\neg(\varphi \rightarrow \psi)]$
MP. From φ and $\varphi \Rightarrow \psi$ infer ψ.

It is easy to see that system **C** is very similar to system **P**. The richer language lets us replace a rule like AND by the axiom C2. Similarly, C1 to REF, C3 to OR and C4 to CM. We need C0 and MP to deal with propositional reasoning. Finally, C5 captures the fact that the plausibility function Pl is independent of the world. Thus, if a default is true (false) at some world, it is true (false) at all of them. If we had enriched plausibility structures to allow a different plausibility function Pl_w for each world w (as is done in the general definition of conditional logic (Lewis 1973)) then we would not need this axiom. There is no KLM property analogous to C5 since a formula such as $N(\varphi \rightarrow \psi)$ involves nested $\rightarrow$'s.

It is well known (Lewis 1973; Burgess 1981; Friedman & Halpern 1994) that system **C** captures reasoning in preferential structures.

Theorem 6.1: (Burgess 1981; Friedman & Halpern 1994) *System* **C** *is a sound and complete axiomatization of* $\mathcal{L}^C$ *with respect to* $\mathcal{P}^p$.

Since the axioms of system **C** are clearly valid in all the structures in $\mathcal{P}^{QPL}$ and $\mathcal{P}^p \subseteq \mathcal{P}^{QPL}$, we immediately get

Corollary 6.2: *System* **C** *is a sound and complete axiomatization of* $\mathcal{L}^C$ *with respect to* $\mathcal{P}^{QPL}$.

This result shows that, at least as far as the language $\mathcal{L}^C$ goes, plausibility structures are no more expressive than preferential structures. We return to this issue in Section 7.

The language $\mathcal{L}^C$ allows us to make distinctions that we could not make using just implication between defaults, as in Section 4. For example, consider the following axiom:

C6. $\varphi \rightarrow \psi \wedge \neg(\varphi \wedge \xi \rightarrow \psi) \Rightarrow \varphi \rightarrow \neg\xi$

Axiom C6 corresponds to the rule of *rational monotonicity* discussed in (Kraus, Lehmann, & Magidor 1990; Lehmann & Magidor 1992). It is not hard to show that C6 is valid in systems where the plausibility ordering is modular. In particular, it is valid in each of $\mathcal{P}^r$, $\mathcal{P}^{Poss}$, and $\mathcal{P}^\kappa$, although it is not valid in $\mathcal{P}^p$. In fact, it is well-known that system **C**+C6 is a sound and complete axiomatization of $\mathcal{L}^C$ with respect to $\mathcal{P}^r$ (Burgess 1981).[11]

7 Conclusions

We feel that this paper unifies earlier results regarding the KLM properties, and explains why they arise so frequently. It also points out the advantage of using plausibility measures as a semantics for defaults.

Do we really need plausibility measures? If all we are interested in is propositional default reasoning and the KLM properties, then the results of Section 6 show that preferential structures provide us all the expressive power we need. Roughly speaking, this is so because when doing propositional reasoning, we can safely restrict to finite structures. (Technically, this is because we have a *finite model property*: if a formula in $\mathcal{L}^C$ is satisfiable, it is satisfiable in a finite plausibility structure (Friedman & Halpern 1994).) As we show in companion paper (Friedman, Halpern, & Koller 1996), preferential structures and plausibility structures are no longer equally expressive once we move to a first-order logic, precisely because infinite structures now play a more important role. The extra expressive power of plausibility structures makes them more appropriate than

[9]The connections between default reasoning and conditional logics are well-known; see (Boutilier 1994; Kraus, Lehmann, & Magidor 1990; Katsuno & Satoh 1991).

[10]We redefine $[\![\varphi]\!]_{PL}$ since φ can involve conditional statements. Note that if φ does not contain occurrences of $\rightarrow$ then this definition is equivalent to the one we gave earlier. Again, we omit the subscript when it is clear from the context.

[11]To capture $\mathcal{P}^\kappa$ and $\mathcal{P}^{poss}$, we need the additional axiom $\neg(\mathit{true} \rightarrow \mathit{false})$. This axiom together with system **C** also provides a complete axiomatization for $\mathcal{P}^\epsilon$. These results are based on well-known results in conditional logic (Burgess 1981; Friedman & Halpern 1994) and are proved in the full paper.

preferential structures for providing semantics for first-order default reasoning.

Beyond their role in default reasoning, we expect that plausibility measures will prove useful whenever we want to express uncertainty and do not want to (or cannot) do so using probability. For example, we can easily define a plausibilistic analogue of conditioning (Friedman & Halpern 1995a). While this can also be done in many of the other approaches we have considered, we believe that the generality of plausibility structures will allow us to again see what properties of independence we need for various tasks. In particular, in (Friedman & Halpern 1996), we use plausibilistic independence to define a plausibilistic analogue of Markov chains. We plan to further explore the properties and applications of plausibility structures in future work.

Acknowledgements

The authors are grateful to Ronen Brafman, Adnan Darwiche, Moises Goldszmidt, Adam Grove, and Daphne Koller for useful discussions relating to this work. The authors were supported in part by the Air Force Office of Scientific Research (AFSC), under Contract F49620-91-C-0080 and by NSF grant IRI-95-03109. The first author was also supported in part by Rockwell Science Center.

References

Adams, E. 1975. *The Logic of Conditionals*. Reidel.

Bacchus, F.; Grove, A. J.; Halpern, J. Y.; and Koller, D. 1993. Statistical foundations for default reasoning. In *IJCAI '93*, 563–569. Available at http://logos.uwaterloo.ca.

Boutilier, C. 1994. Conditional logics of normality: a modal approach. *Artificial Intelligence* 68:87–154.

Burgess, J. 1981. Quick completeness proofs for some logics of conditionals. *Notre Dame J. of Formal Logic* 22:76–84.

Darwiche, A. 1992. *A Symbolic Generalization of Probability Theory*. Ph.D. Dissertation, Stanford University.

Dubois, D., and Prade, H. 1990. An introduction to possibilistic and fuzzy logics. In Shafer, G., and Pearl, J., eds., *Readings in Uncertain Reasoning*. Morgan Kaufmann.

Dubois, D., and Prade, H. 1991. Possibilistic logic, preferential models, non-monotonicity and related issues. In *IJCAI '91*, 419–424.

Friedman, N., and Halpern, J. Y. 1994. On the complexity of conditional logics. In *KR '94*, 202–213.

Friedman, N., and Halpern, J. Y. 1995a. Plausibility measures: a user's manual. In *UAI '95*, 175–184.

Friedman, N., and Halpern, J. Y. 1995b. Plausibility measures and default reasoning. Technical Report RJ 9959, IBM. Available at `http://robotics.stanford.edu/users/nir`.

Friedman, N., and Halpern, J. Y. 1996. A qualitative Markov assumption and its implications for belief change. Submitted to UAI '96.

Friedman, N.; Halpern, J. Y.; and Koller, D. 1996. Conditional first-order logic revisited. In *AAAI '96*.

Gabbay, D. M.; Hogger, C. J.; and Robinson, J. A., eds. 1993. *Nonmonotonic Reasoning and Uncertain Reasoning*, volume 3 of *Handbook of Logic in Artificial Intelligence and Logic Programming*. Oxford University Press.

Gärdenfors, P., and Makinson, D. 1988. Revisions of knowledge systems using epistemic entrenchment. In *Proc. 2nd Conf. on Theoretical Aspects of Reasoning about Knowledge*. 83–95.

Geffner, H. 1992a. *Default Reasoning*. MIT Press.

Geffner, H. 1992b. High probabilities, model preference and default arguments. *Mind and Machines* 2:51–70.

Ginsberg, M. L., ed. 1987. *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann.

Goldszmidt, M., and Pearl, J. 1992. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *KR '92*, 661–672.

Goldszmidt, M.; Morris, P.; and Pearl, J. 1993. A maximum entropy approach to nonmonotonic reasoning. *IEEE Trans. of Pattern Analysis and Machine Intelligence* 15(3):220–232.

Katsuno, H., and Satoh, K. 1991. A unified view of consequence relation, belief revision and conditional logic. In *IJCAI '91*, 406–412.

Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44:167–207.

Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artificial Intelligence* 55:1–60.

Lewis, D. K. 1973. *Counterfactuals*. Harvard University Press.

Pearl, J. 1989. Probabilistic semantics for nonmonotonic reasoning: a survey. In *KR '89*, 505–516.

Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Theoretical Aspects of Reasoning about Knowledge: Proc. 3rd Conference*, 121–135.

Shafer, G. 1976. *A Mathematical Theory of Evidence*. Princeton University Press.

Shoham, Y. 1987. A semantical approach to nonmonotonic logics. In *Proc. 2nd IEEE Symp. on Logic in Computer Science*, 275–279.

Spohn, W. 1987. Ordinal conditional functions: a dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change and Statistics*, volume 2. Reidel. 105–134.

Wang, Z., and Klir, G. J. 1992. *Fuzzy Measure Theory*. Plenum.

Weydert, E. 1994. General belief measures. In *UAI '94*, 575–582.