Plausibility Measures and Default Reasoning

Nir Friedman
University of California, Berkeley
and
Joseph Y. Halpern
Cornell University

We introduce a new approach to modeling uncertainty based on *plausibility measures*. This approach is easily seen to generalize other approaches to modeling uncertainty, such as probability measures, belief functions, and possibility measures. We focus on one application of plausibility measures in this paper: default reasoning. 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. While this was viewed as a surprise, we show here that it is almost inevitable. In the framework of plausibility measures, we can give a necessary condition for the KLM axioms to be sound, and an additional condition necessary and sufficient to ensure that the KLM axioms are complete. This additional condition is so weak that it is almost always met whenever the axioms are sound. In particular, it is easily seen to hold for all the proposals made in the literature.

Categories and Subject Descriptors: F.4.1 [**Mathematical Logic and Formal Languages**]: Mathematical Logic; I.2.3 [**Artificial Intelligence**]: Deduction and Theorem Proving—*Nonmonotonic reasoning and belief revision*; I.2.4 [**Artificial Intelligence**]: Knowledge Representation Formalisms and Methods—*Representation Languages*

General Terms: Theory

Additional Key Words and Phrases: Conditional Logic, Default Reasoning, ϵ-Semantics, κ-rankings, Plausibility Measures, Possibility Measures, Preferential Orderings Nonmonotonic Inference.

1. INTRODUCTION

We must reason and act in an uncertain world. There may be uncertainty about the state of the world, uncertainty about the effects of our actions, and uncertainty

Some of this work was done while both authors were at the IBM Almaden Research Center. The first author was also at Stanford while much of the work was done. IBM and Stanford's support are gratefully acknowledged. The work was also supported in part by the Air Force Office of Scientific Research (AFSC), under Contract F49620-91-C-0080 and grant F94620-96-1-0323 and by NSF under grants IRI-95-03109 and IRI-96-25901. The first author was also supported in part by an IBM Graduate Fellowship and by the Rockwell Science Center. A preliminary version of this paper appears in *Proceedings, Thirteenth National Conference on Artificial Intelligence*, 1996, pp. 1297–1304.
Name: Nir Friedman
Address: Computer Science Division, 387 Soda Hall, University of California, Berkeley, CA 94720. email: `nir@cs.berkeley.edu`; url: `http://www.cs.berkeley.edu/~nir`
Name: Joseph Y. Halpern
Address: Computer Science Department, Cornell University, Ithaca, NY 14853.
email: `halpern@cs.cornell.edu`; url: `http://www.cs.cornell.edu/home/halpern`

about other agents' actions. The standard approach to modeling uncertainty is probability theory. In recent years, researchers, motivated by varying concerns including a dissatisfaction with some of the axioms of probability and a desire to represent information more qualitatively, have introduced a number of generalizations and alternatives to probability, such as Dempster-Shafer belief functions [Shafer 1976] and possibility theory [Dubois and Prade 1990]. Our aim here is to introduce what is perhaps the most general approach possible, which uses what we call *plausibility measures*. A plausibility measure associates with a set a *plausibility*, which is just an element in a partially ordered space. The only property that we impose is that the plausibility of a set be at least as large as the plausibility of any of its subsets. Every systematic approach for dealing with uncertainty that we are aware of can be viewed as a plausibility measure. Given how little structure we have imposed on plausibility measures, this is perhaps not surprising. Nevertheless, as we hope to demonstrate in this and other papers, plausibility measures provide us with a natural setting in which to examine various approaches to reasoning about uncertainty.

The focus of this paper is (propositional) default reasoning. There have been many approaches to default reasoning proposed in the literature (see [Ginsberg 1987; Gabbay et al. 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 by Kraus, Lehmann, and Magidor [1990]), and many of the recent approaches to default reasoning, including *preferential structures* [Kraus et al. 1990; Shoham 1987], *ϵ-semantics* [Adams 1975; Geffner 1992b; Pearl 1989], *possibilistic structures* [Dubois and Prade 1991], and *κ-rankings* [Goldszmidt and Pearl 1992; Spohn 1988], 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." As we shall show in this paper, plausibility measures help us understand why this should be so.

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 condition. 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 suggests that it will be difficult to find an interesting approach that gives semantics to defaults with respect to a collection $\mathcal{P}$ of structures and goes beyond beyond KLM. This result thus provides added insight into and justification for approaches such as those of [Bacchus et al. 1993; Geffner 1992a; Goldszmidt and Pearl 1992; Goldszmidt et al. 1993; Lehmann and 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 (rather than requiring that d be true in all structures in $\mathcal{P}$ that satisfy Δ).

This paper is organized as follows. In Section 2, we introduce plausibility measures and show how they generalize various other proposals for capturing uncertainty. 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 prove our main results: 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 Sections 6, 7, and 8, we expand on three independent topics related to our results. In Section 6, we show that qualitative plausibility measures are more expressive than previous semantics considered in the literature. In Section 7, we consider related work, focusing on the relationship between our approach to plausibility and *epistemic entrenchment* [Gärdenfors and Makinson 1988; Gärdenfors and Makinson 1989; Grove 1988]. In Section 8, 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]. We conclude in Section 9 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).[1]

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

[1] Frequently it is also assumed that μ satisfies *countable additivity*, i.e., if A_i, $i > 0$, are pairwise disjoint, then $\mu(\bigcup_i A_i) = \sum_i \mu(A_i)$.

we make is

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

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 $\mathcal{F}$.[2] 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 D from $\leq_D$, $<_D$, $\top_D$, and $\perp_D$ whenever it is clear from context. We also frequently omit the $\mathcal{F}$ when describing a plausibility space when its role is not that significant, and just denote a plausibility space as a pair (W, Pl) rather than $(W, \mathcal{F}, \mathrm{Pl})$.

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 and Klir 1992] is a function $f : 2^W \mapsto [0,1]$, that satisfies A1 and some continuity constraints. A *possibility measure* [Dubois and Prade 1990] Poss is a fuzzy measure such that $\mathrm{Poss}(W) = 1$, $\mathrm{Poss}(\emptyset) = 0$, and $\mathrm{Poss}(A) = \sup_{w \in A}(\mathrm{Poss}(\{w\}))$.

—An *ordinal ranking* (or *κ-ranking*) κ on W (as defined by [Goldszmidt and Pearl 1992], based on ideas that go back to [Spohn 1988]) is a function mapping subsets of W to $\mathbb{N}^* = \mathbb{N} \cup \{\infty\}$ 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 $x \leq_{\mathbb{N}^*} y$ if and only if $y \leq x$ under the usual ordering on the ordinals.

—A *preference ordering* on W is a strict partial order $\prec$ over W [Kraus et al. 1990; Shoham 1987]. Intuitively, $w \prec w'$ holds if w is *preferred* to w'.[3] 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 on worlds as well as the truth values of defaults.

—A *parameterized probability distribution* (PPD) on W is a sequence $\{\mathrm{Pr}_i : i \geq 0\}$ of probability measures over W. Such structures provide semantics for defaults in *ϵ-semantics* [Pearl 1989; Goldszmidt et al. 1993]. In Section 4 we show how to map PPDs on W to plausibility measures on W in a way that preserves the truth-values of conditionals.

[2] In dealing with *conditional* plausibility, the domain D plays a more significant role [Friedman and Halpern 1995].

[3] We follow the standard notation for preference here [Lewis 1973; Kraus et al. 1990], which uses the (perhaps confusing) convention of placing the more likely (or less abnormal) world on the left of the $\prec$ operator.

Plausibility structures are motivated by much the same concerns as two other recent symbolic generalizations of probability by Darwiche [1992] and Weydert [1994b]. 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 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 and Halpern 1995] for a more detailed comparison to [Darwiche 1992; Weydert 1994b].

Given the simplicity and generality of plausibility measures, we were not surprised to discover that Weber [1991] recently defined a notion of *uncertainty measures*, which is a slight generalization of plausibility measure (in that domains more general than algebras of sets are allowed), and that Greco [1987] defined a notion of L-fuzzy measures which is somewhat more restricted than plausibility measures in that the range D is a complete lattice. We expect that others have used variants of this notion as well, although we have not found any further references in the literature. To the best of our knowledge, we are the first to carry out a systematic investigation of the connection between plausibility measures and default reasoning.

3. APPROACHES TO DEFAULT REASONING: A REVIEW

Defaults are statements of the form "if ϕ then typically/likely/by default ψ", which we denote $\phi{\rightarrow}\psi$. For example, the default "birds typically fly" is represented $\mathit{Bird}{\rightarrow}\mathit{Fly}$. Formally, we assume that there is a "base" language $\mathcal{L}$, defined over a set Φ of propositions, that includes the usual propositional connectives, $\wedge, \vee, \neg, \Rightarrow$ and has a consequence relation $\vdash_{\mathcal{L}}$. The language of defaults $\mathcal{L}_{\rm def}$ contains statements of the form $\phi{\rightarrow}\psi$, where $\phi, \psi \in \mathcal{L}$.

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 be 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 [Kraus et al. 1990], consists of the following axiom and rules of inference.

LLE. If $\vdash_{\mathcal{L}} \phi \Leftrightarrow \phi'$, then from $\phi{\rightarrow}\psi$ infer $\phi'{\rightarrow}\psi$ (left logical equivalence)

RW. If $\vdash_{\mathcal{L}} \psi \Rightarrow \psi'$, then from $\phi{\rightarrow}\psi$ infer $\phi{\rightarrow}\psi'$ (right weakening)

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

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

OR. From $\phi_1{\rightarrow}\psi$ and $\phi_2{\rightarrow}\psi$ infer $\phi_1 \vee \phi_2{\rightarrow}\psi$

CM. From $\phi{\rightarrow}\psi_1$ and $\phi{\rightarrow}\psi_2$ infer $\phi \wedge \psi_2{\rightarrow}\psi_1$ (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 $\phi_2{\rightarrow}\psi$ from $\phi_1{\rightarrow}\psi$. RW describes a similar property of the consequent: If ψ (logically) entails ψ', then we can deduce $\phi{\rightarrow}\psi'$ from $\phi{\rightarrow}\psi$.

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, then 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_2 holds when ϕ holds (as would be expected, given the default) should not cause us to retract the default conclusion ψ_1.

This system of rules is called system **P** in [Kraus et al. 1990]. The notation $\Delta \vdash_{\mathbf{P}} \phi{\rightarrow}\psi$ denotes that $\phi{\rightarrow}\psi$ 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 consistent with $\vdash_{\mathcal{L}}$ to formulas in $\mathcal{L}$, 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 $M = (W, X, \pi)$ and a formula $\phi \in \mathcal{L}$, we take $[\![\phi]\!]_M \subseteq W$ to be the set of worlds satisfying ϕ, i.e., $[\![\phi]\!]_M = \{w \in W : \pi(w)(\phi) = \mathbf{true}\}$. We omit the subscript M when it plays no role or is clear from the context.

The first semantic proposal was provided by Kraus, Lehmann and Magidor [1990], using ideas that go back to [Hansson 1969; Lewis 1973; Shoham 1987]. Recall that a preference ordering on W is strict partial order (i.e., an irreflexive and transitive relation) $\prec$ over W. A *preferential structure* is a tuple $(W, \prec, \pi)$, where $\prec$ is a strict partial order on W.[4] The intuition [Shoham 1987] is that a preferential structure satisfies a conditional $\phi{\rightarrow}\psi$ if all the most preferred worlds (i.e., the minimal worlds according to $\prec$) in $[\![\phi]\!]$ satisfy ψ. However, there may be no minimal worlds in $[\![\phi]\!]$. This can happen if $[\![\phi]\!]$ 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 $[\![\phi]\!]$ whenever $[\![\phi]\!]$ is not empty; this is the assumption actually made in [Kraus et al. 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 [Bossu and Siegel 1985; Boutilier 1994; Lewis 1973]. Roughly speaking, $\phi{\rightarrow}\psi$ is true if, from a certain point on, whenever ϕ is true, so is ψ. More formally,

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

---

[4] We note that the formal definition of preferential structures in [Kraus et al. 1990; Lehmann and 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.

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

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

THEOREM 3.1. [Lehmann and Magidor 1992; Boutilier 1994] $\Delta \models_{\rm p} \phi\rightarrow\psi$ *if and only if* $\Delta \vdash_{\mathbf{P}} \phi\rightarrow\psi$.

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 restricting 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* $\phi\rightarrow\psi$, denoted $\Delta \models_{\rm r} \phi\rightarrow\psi$, if every rational structure that satisfies Δ also satisfies $\phi\rightarrow\psi$.[5]

THEOREM 3.2. [Lehmann and Magidor 1992] $\Delta \models_{\rm r} \phi\rightarrow\psi$ *if and only if* $\Delta \vdash_{\mathbf{P}} \phi\rightarrow\psi$.

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

This is perhaps somewhat surprising, since it is is known that rational structures do satisfy the following additional property, known as *rational monotonicity* [Kraus et al. 1990; Lehmann and Magidor 1992]:

RM If $\phi\rightarrow\psi_1$ and $\phi\not\rightarrow\neg\psi_2$ then $\phi \wedge \psi_2\rightarrow\psi_1$.

Note that RM is almost the same as CM, except that $\phi\rightarrow\psi_2$ is replaced by the weaker $\phi\not\rightarrow\neg\psi_2$.

How can the existence of this additional property be consistent with the fact both that rational and preferential structures are characterized by system **P**? The key point is that although RM is an additional property satisfied by rational structures, it is not one that is expressible in the language of defaults (because we do not allow negated defaults). As we shall see in Section 8, once we move to a richer language, rational structures are distinguishable form arbitrary preferential structures.

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

[5]Rational entailment should not be confused with the notion of *rational closure*, also defined by Lehmann and Magidor [1992].

The formal definition needs a bit of machinery.[6] 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 $\phi \rightarrow \psi$ if the conditional probability ψ given ϕ goes to 1 in the limit. Formally, $\phi \rightarrow \psi$ is satisfied if $\lim_{i\rightarrow\infty} \Pr_i([\![\psi]\!] | [\![\phi]\!]) = 1$ (where $\Pr_i([\![\psi]\!] | [\![\phi]\!])$ is taken to be 1 if $\Pr_i([\![\phi]\!]) = 0$). Δ *ϵ-entails* $\phi \rightarrow \psi$, denoted $\Delta \models_\epsilon \phi \rightarrow \psi$, if every PPD structure that satisfies all the defaults in Δ also satisfies $\phi \rightarrow \psi$. Surprisingly, Geffner shows that ϵ-entailment is equivalent to preferential entailment.

THEOREM 3.3. [Geffner 1992b] *$\Delta \models_\epsilon \phi \rightarrow \psi$ if and only if $\Delta \vdash_{\mathbf{P}} \phi \rightarrow \phi$.*[7]

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

These two approaches are again characterized by the KLM properties.

THEOREM 3.4. [Geffner 1992b; Dubois and Prade 1991] *The following are equivalent:*

(a) $\Delta \models_{Poss} \phi \rightarrow \psi$
(b) $\Delta \models_\kappa \phi \rightarrow \psi$
(c) $\Delta \vdash_{\mathbf{P}} \phi \rightarrow \psi$.

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* (for $\mathcal{L}$) is a tuple $PL = (W, \mathcal{F}, \mathrm{Pl}, \pi)$, where $(W, \mathcal{F}, \mathrm{Pl})$ is a plausibility space and π maps each possible world to a truth assignment to the formulas in $\mathcal{L}$ that is consistent with $\vdash_{\mathcal{L}}$ in the obvious sense. Since we will be interested in events that correspond to formulas, we require that $[\![\phi]\!] \in \mathcal{F}$ for all formula $\phi \in \mathcal{L}$. For ease of exposition, when describing a plausibility structure for $\mathcal{L}$, we assume that $\mathcal{F} = \{[\![\phi]\!] : \phi \in \mathcal{L}\}$. Just as with plausibility spaces, we typically omit the algebra $\mathcal{F}$ from the description of a plausibility structure. We define $PL \models_{PL} \phi \rightarrow \psi$ if either $\mathrm{Pl}([\![\phi]\!]) = \perp$ or $\mathrm{Pl}([\![\phi \wedge \psi]\!]) > \mathrm{Pl}([\![\phi \wedge \neg\psi]\!])$.

[6] We adopt the presentation used in [Goldszmidt et al. 1993].
[7] Geffner's result is stated in terms of the original formulation of ϵ-entailment, as described in [Pearl 1989]. However, results of [Goldszmidt et al. 1993] show that the formulation we describe here is equivalent to the original one.

Notice that if Pl is a probability function Pr, then $\phi \rightarrow \psi$ holds exactly if either $\Pr([\![\phi]\!]) = 0$ or $\Pr([\![\psi]\!] | [\![\phi]\!]) > 1/2$. 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 in Section 3), 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.

In fact, Lemma 4.1 shows that there is a general procedure for mapping any approach that satisfies the KLM postulates to plausibility measures. Before describing this general construction, we briefly sketch its instantiation in the case of PPDs and preferential structures.

We start with PPDs. Let $PP = (W, \{\Pr_i\}, \pi)$ be a PPD structure. Let Pl_{PP} be a plausibility measure on W such that

$$\mathrm{Pl}_{PP}(A) \leq \mathrm{Pl}_{PP}(B) \text{ if and only if } \lim_{i \to \infty} \Pr_i(B | A \cup B) = 1. \qquad (1)$$

It is easy to check that such a plausibility measure exists and that $(W, \mathrm{Pl}_{PP}, \pi)$ satisfies the same defaults as PP. We note that this construction, as well as others in the remainder of the paper, specifies only the relative order of plausibilities of events, and does not describe the domain of plausibility values. It is easy to check that as long as the ordering constraints are consistent with reflexivity, transitivity, and A1, we can always construct a matching plausibility domain.[8] From here on, we treat such ordering constraints as though they define a plausibility measure.

We stress that this embedding, which is sufficient for the purpose of this work, is not the only one possible. To see this, suppose that A and B are disjoint sets such that $\Pr_i(A) = \Pr_i(B)$ for all i. One might argue that the plausibility of A and B should be equal. Yet our definition would make $\mathrm{Pl}(A)$ and $\mathrm{Pl}(B)$ incomparable since $\Pr_i(B|A \cup B) = 0.5$ for all i.

The construction for mapping preferential structures into plausibility structures is slightly more complex. Suppose we are given a preferential structure $(W, \prec, \pi)$. 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 every set of elements in D has a least upper bound in D). It is not hard to show that D is well-defined (i.e., there is a unique, up to renaming, smallest set) and that taking $\mathrm{Pl}_\prec(A)$ to be the least upper bound of $\{d_w : w \in A\}$ gives us the following property:

$\mathrm{Pl}_\prec(A) \leq \mathrm{Pl}_\prec(B)$ 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 $w'' \prec w'$. (2)

Again, it is easy to check that $(W, \mathrm{Pl}_\prec, \pi)$ satisfies the same defaults as $(W, \prec, \pi)$.

[8] For example, we can take the domain of the plausibility measure to consist of sets of logically equivalent formulas, partially ordered so as to satisfy the constraints.

We now present our general construction.

LEMMA 4.1. *Let W be a set of possible worlds and let π be a function that maps each world in W to a truth assignment to $\mathcal{L}$. Let $T \subseteq \mathcal{L}_{\mathrm{def}}$ be a set of defaults that is closed under the rules of system* **P** *that satisfies the following condition:*

> $(*)$ *if $\phi \rightarrow \psi \in T$, $[\![\phi]\!] = [\![\phi']\!]$, and $[\![\psi]\!] = [\![\psi']\!]$, then $\phi' \rightarrow \psi' \in T$, for all formulas $\phi, \phi', \psi, \psi' \in \mathcal{L}$.*

There is a plausibility structure $PL_T = (W, \mathrm{Pl}_T, \pi)$ such that $\mathrm{Pl}_T([\![\phi]\!]) \leq \mathrm{Pl}_T([\![\psi]\!])$ if and only if $\phi \vee \psi \rightarrow \psi \in T$. Moreover, $PL_T \models \phi \rightarrow \psi$ if and only if $\phi \rightarrow \psi \in T$.

PROOF. See Appendix A.1. □

THEOREM 4.2. *(a) Let $PP = \{\mathrm{Pr}_i\}$ be a PPD on W. There is a plausibility measure Pl_{PP} on W such that $(W, \{\mathrm{Pr}_i\}, \pi) \models_\epsilon \phi \rightarrow \psi$ if and only if $(W, \mathrm{Pl}_{PP}, \pi) \models_{PL} \phi \rightarrow \psi$.*

(b) Let $\prec$ be a preference ordering on W. There is a plausibility measure $\mathrm{Pl}_\prec$ on W such that $(W, \prec, \pi) \models_p \phi \rightarrow \psi$ if and only if $(W, \mathrm{Pl}_\prec, \pi) \models_{PL} \phi \rightarrow \psi$.

PROOF. We start with part (a). Set $T_{PP} = \{\phi \rightarrow \psi : (W, PP, \pi) \models_\epsilon \phi \rightarrow \psi\}$. Theorem 3.3 implies that T_{PP} is closed under the KLM rules. Moreover, since T is constructed from a PPD over W using π it satisfies the requirement of Lemma 4.1. It follows that we can construct a plausibility structure that satisfies the requirements of the theorem. It is also easy to verify that this construction agrees with the construction described earlier, in that it satisfies constraint (1). To see this, let $A, B \in \mathcal{F}$. Then, according to our assumptions, there are formulas ϕ and ψ such that $A = [\![\phi]\!]$ and $B = [\![\psi]\!]$. By definition, $\mathrm{Pl}_{PP}(A) \leq \mathrm{Pl}_{PP}(B)$ if and only if $(W, PP, \pi) \models_\epsilon (\phi \vee \psi) \rightarrow \psi$. From definition of $\models_\epsilon$ we immediately get (1).

The proof of part (b) is identical, using Theorem 3.1. Again, an analogous argument easily shows that $\mathrm{Pl}_\prec$ satisfies (2). □

Thus, each of the semantic approaches to default reasoning that were considered in Section 3 can be mapped into plausibility structures in a way that preserves the semantics of defaults. We remark that these mapping are not unique. For example, Freund [1996] gives an alternative mapping from preferential structures to plausibility measures.

Other semantics for defaults can also be mapped into plausibility measures using the general technique of Lemma 4.1. In most cases, we can also establish a direct relationship between these semantics and plausibility measures. For example, the *coherent filters* approach of [Ben-David and Ben-Eliyahu 1994; Schlechta 1995] can be mapped to plausibility, as shown by Schlechta [1996], and Weydert's *full ranking measures* [1994a] are easily seen to be a special case of plausibility measures.

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 $\mathcal{S}$ is a class of plausibility structures, then a knowledge base Δ entails $\phi \rightarrow \psi$ with respect to $\mathcal{S}$, denoted $\Delta \models_{\mathcal{S}} \phi \rightarrow \psi$, if every plausibility structure $PL \in \mathcal{S}$ that satisfies all the defaults in Δ also satisfies $\phi \rightarrow \psi$.* ∎

The classes of structures we are interested in include $\mathcal{S}^{PL}$, the class of all plausibility structures, and $\mathcal{S}^{Poss}$, $\mathcal{S}^{\kappa}$, $\mathcal{S}^{p}$, $\mathcal{S}^{r}$, and $\mathcal{S}^{\epsilon}$, 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, rational and PPD structures, the mapping is the one described in Theorem 4.2.) 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 $\mathcal{S}^{PL}$. In particular, they do not hold in probability structures. For a counterexample, consider a plausibility structure $PL = (W, \mathrm{Pl}, \pi)$, where Pl is actually a probability measure Pr such that $\Pr([\![q \wedge r]\!]) = 0.2$ and $\Pr([\![q \wedge \neg r]\!]) = \Pr([\![\neg q \wedge r]\!]) = 0.4$. Thus, $\Pr([\![q]\!]) = .6$ and $\Pr([\![q]\!] | [\![r]\!]) = 1/3$. Recall that if $\Pr([\![\phi]\!]) > 0$, then $PL \models_{PL} \phi {\rightarrow} \psi$ if and only if $\Pr([\![\psi]\!] | [\![\phi]\!]) > .5$. Thus, $PL \models_{PL} (true{\rightarrow}q) \wedge (true{\rightarrow}r)$, but $PL \not\models_{PL} true{\rightarrow}(q \wedge r)$ and $PL \not\models_{PL} r{\rightarrow}q$. 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 $\Delta \vdash_{\mathbf{P}'} \phi{\rightarrow}\psi$, then $\Delta \models_{\mathcal{S}^{PL}} \phi{\rightarrow}\psi$.*

PROOF. See Appendix A.2. □

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 $[\![\phi]\!] = A$, $[\![\psi_1]\!] = B_1$, 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$'$.*

PROOF. See Appendix A.2. □

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 also satisfies the AND rule. Somewhat surprisingly, a plausibility measure Pl that satisfies A2 also satisfies CM. Moreover, Pl satisfies the non-vacuous case of the OR rule. That is, if $\mathrm{Pl}([\![\phi_1]\!]) > \perp$, then from $\phi_1 \rightarrow \psi$ and $\phi_2 \rightarrow \psi$ we can conclude $(\phi_1 \vee \phi_2) \rightarrow \psi$.[9] 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 in addition to A1. We say $PL = (W, \mathrm{Pl}, \pi)$ is a *qualitative plausibility structure* if (W, Pl) is a qualitative plausibility space. Let $\mathcal{S}^{QPL}$ consist of all qualitative plausibility structures.

THEOREM 5.4. *$\mathcal{S} \subseteq \mathcal{S}^{QPL}$ if and only if for all Δ, ϕ, and ψ, if $\Delta \vdash_{\mathbf{P}} \phi \rightarrow \psi$ then $\Delta \models_{\mathcal{S}} \phi \rightarrow \psi$.*

PROOF. See Appendix A.2. □

Thus, the KLM axioms are sound for qualitative plausibility structures. We remark that Theorem 5.4 provides not only a sufficient but a necessary condition for a set of plausibility structures to satisfy the KLM properties: If the KLM axioms are sound with respect to $\mathcal{S}$, then all $PL \in \mathcal{S}$ must be 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 $\mathcal{S}^{Poss}$, $\mathcal{S}^{\kappa}$, $\mathcal{S}^{\epsilon}$, $\mathcal{S}^{p}$, and $\mathcal{S}^{r}$ is a subset of $\mathcal{S}^{QPL}$.*

PROOF. See Appendix A.2. □

It follows from Theorems 5.4 and 5.5 that the KLM properties hold in all the approaches to default reasoning 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 hold for all these approaches.

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

DEFINITION 5.6. *We say that $\mathcal{S}$ is* rich *if for every collection $\phi_1, \ldots, \phi_n$, $n > 1$, of mutually exclusive formulas, there is a plausibility structure $PL = (W, \mathrm{Pl}, \pi) \in \mathcal{S}$ such that:*

$$\mathrm{Pl}([\![\phi_1]\!]) > \mathrm{Pl}([\![\phi_2]\!]) > \cdots > \mathrm{Pl}([\![\phi_n]\!]) = \perp. \blacksquare$$

The richness requirement is quite mild. It says that we do not have *a priori* constraints on the relative plausibilities of a collection of disjoint sets. Theorem 5.7

[9]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.

shows that every collection of plausibility measures that we have considered thus far can be easily shown to satisfy this richness condition. More importantly, Theorem 5.8 shows that richness is a necessary and sufficient condition to ensure that the KLM properties are complete.

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

PROOF. Let $\phi_1, \ldots, \phi_n$, $n > 1$, be mutually exclusive formulas and let $W = \{w_1, \ldots, w_{n-1}\}$. Since $\phi_1, \ldots, \phi_n$ are mutually exclusive, we can construct a mapping π that maps each world in W to a truth assignment such that $[\![\phi_i]\!] = w_i$ for all $1 \leq i < n$, and $[\![\phi_n]\!] = \emptyset$. Recall that we need to find a plausibility measure Pl such that $\mathrm{Pl}([\![\phi_1]\!]) > \mathrm{Pl}([\![\phi_2]\!]) > \cdots > \mathrm{Pl}([\![\phi_n]\!]) = \perp$. It is easy to find a plausibility measure Pl satisfying this property such that (W, Pl, π) is in $\mathcal{S}^{Poss}$, $\mathcal{S}^{\kappa}$, $\mathcal{S}^{p}$, $\mathcal{S}^{r}$, $\mathcal{S}^{\epsilon}$, or $\mathcal{S}^{QPL}$. For example, to get a structure in $\mathcal{S}^{Poss}$, we define $\mathrm{Poss}(w_i) = 1 - \frac{i}{n}$. To get a structure in $\mathcal{S}^{p}$, we define Pl to correspond to the preference ordering $w_1 \prec w_2 \prec \ldots \prec w_{n-1}$. $\square$

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

PROOF. See Appendix A.2. $\square$

Note that Theorem 5.8 deals with what is usually considered to be *weak* completeness. The *strong* notion of completeness would require us to remove the restriction that Δ is finite from the statement of the theorem. It is possible to find a stronger notion of richness that corresponds to strong completeness, but the details are somewhat cumbersome, so we do not provide them here. Note that if $\models_{\mathcal{S}}$ is compact, then weak completeness implies strong completeness.

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

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

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 $\mathcal{S}$ of qualitative structures satisfying the richness constraint, and defining $\Delta \models_{\mathcal{S}} \phi \rightarrow \psi$ to hold if $\phi \rightarrow \psi$ is true in every structure in $\mathcal{S}$ that satisfies Δ, cannot lead to new properties for entailment.

Thus, to go beyond KLM, we need to either consider interesting non-rich classes of structures or to define a notion of entailment from Δ that does not consider all the structures of a given class. We are not aware of any work that takes the first approach, although 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 the language includes primitive propositions p

and q, and we consider all structures where p is independent of q in the sense that if any of $true \rightarrow q$, $p \rightarrow q$, and $\neg p \rightarrow q$ holds, then the others also do. This means that discovering either p or $\neg p$ does not affect whether or not q is believed.[10] 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 and Pearl 1992; Goldszmidt et al. 1993; Lehmann and Magidor 1992; Pearl 1990] has taken the second approach. Roughly speaking, this approach can be viewed as taking the basic idea of preferential semantics—placing a preference ordering on worlds—one step further: We try to get from a knowledge base a set of preferred *structures* (where the structures themselves put a preference ordering on worlds)—for example, in [Goldszmidt et al. 1993], the PPD of maximum entropy is considered—and carry out all reasoning in these preferred structures. We believe that plausibility measures will provide insight into techniques for choosing such preferred structures. For example, we might want to prefer structures where things are "as independent as possible". We believe that it should be possible to capture this notion in a reasonable way using plausibility; we defer this to future work. (See [Friedman and Halpern 1995] for discussion on independence in the context of plausibility.)

6. EXPRESSIVENESS OF QUALITATIVE PLAUSIBILITY MEASURES

In the previous section we showed that all approaches to default reasoning are instances of qualitative plausibility structures. We now show that each of the classes considered in Theorem 5.5 is a strict subset of $\mathcal{S}^{QPL}$. This is clearly true in a trivial sense. For example, if we consider a qualitative plausibility measure whose range is $[1,2]$, it cannot be either a possibility measure or a κ-ranking. To get around this problem, we define two plausibility spaces (W, Pl) and (W, Pl') (resp. , two plausibility structures (W, Pl, π) and (W, Pl', π)) to be *order-equivalent* if for $A, B \subseteq W$, we have $\mathrm{Pl}(A) \leq \mathrm{Pl}(B)$ if and only if $\mathrm{Pl}'(A) \leq \mathrm{Pl}'(B)$.

We claim that for each of the classes of plausibility structures considered in Theorem 5.5, there is a qualitative plausibility structure that is not order-equivalent to any element of that class. This is almost immediate in the case of $\mathcal{S}^{Poss}$ and $\mathcal{S}^{\kappa}$. Since both require $\leq$ to be a total order, a qualitative plausibility structure (W, Pl, π) such that Pl does not place a total order on the plausibility of subsets cannot be order-equivalent to an element of $\mathcal{S}^{Poss}$ or $\mathcal{S}^{\kappa}$. We say a plausibility structures (W, Pl, π) is *totally ordered* if Pl places a total order on subsets. As the following proposition show, there are even totally-ordered qualitative plausibility structures that are not order-equivalent to any possibility structure or ordinal ranking structure.

PROPOSITION 6.1. *There is a totally-ordered qualitative plausibility structure that is not order-equivalent to any structure in $\mathcal{S}^{Poss}$, $\mathcal{S}^{\kappa}$, $\mathcal{S}^{\mathrm{P}}$, or $\mathcal{S}^{\epsilon}$.*

[10] We remark that if we define independence appropriately in plausibility structures, this property does indeed hold; see [Friedman and Halpern 1995].

Proof. Define a plausibility measure Pl on $\{a,b,c\}$ such that $\mathrm{Pl}(\{a\}) = \mathrm{Pl}(\{b\}) = \mathrm{Pl}(\{c\}) = \mathrm{Pl}(\{b,c\}) = \mathrm{Pl}(\{a,c\}) = 1/2$ and $\mathrm{Pl}(\{a,b\}) = \mathrm{Pl}(\{a,b,c\}) = 1$. It is straightforward to check that Pl is qualitative and totally ordered. Moreover, we have $\mathrm{Pl}(\{c\}) < \mathrm{Pl}(\{a,b\})$, although neither $\mathrm{Pl}(\{c\}) < \mathrm{Pl}(\{a\})$ nor $\mathrm{Pl}(\{c\}) < \mathrm{Pl}(\{b\})$ hold. It is easy to see that there can be no possibility measure, κ-ranking, preference ordering, or PPD on $\{a,b,c\}$ such that the corresponding plausibility space is order-equivalent to $(\{a,b,c\},\mathrm{Pl})$. For example, if Poss is a possibility measure on $\{a,b,c\}$ such that $\mathrm{Poss}(\{c\}) < \mathrm{Poss}(\{a,b\}$, then we must have either $\mathrm{Poss}(\{c\}) < \mathrm{Poss}(\{a\}$ or $\mathrm{Poss}(\{c\}) < \mathrm{Poss}(\{b\})$. A similar observation holds for κ-rankings. This plausibility space also cannot be equivalent to one that arises from the construction of Lemma 4.1, since the construction never gives disjoint sets the same plausibility. Since $\mathrm{Pl}(\{a\}) = \mathrm{Pl}(\{b\})$, the result follows. □

If all that we are interested in is default reasoning, then all that matters is the relative plausibility of disjoint sets. We say that two plausibility spaces (W,Pl) and (W,Pl') (resp. two plausibility structures (W,Pl,π) and (W,Pl',π)) are *default-equivalent* if for all disjoint subsets A and B of W, we have $\mathrm{Pl}(A) < \mathrm{Pl}(B)$ if and only if $\mathrm{Pl}'(A) < \mathrm{Pl}'(B)$. Clearly, if structures (W,Pl,π) and (W,Pl',π) are default-equivalent, then they satisfy the same defaults.

We can strengthen Proposition 6.1 so that it applies to default-equivalence in the case of possibility measures, κ-rankings, and preferential orders.

Proposition 6.2. *There is a totally-ordered qualitative plausibility structure that is not default-equivalent to any structure in $\mathcal{S}^{Poss}$, $\mathcal{S}^{\kappa}$, or $\mathcal{S}^{p}$.*

Proof. The plausibility space described in the proof of Proposition 6.1 also provides a counterexample for default-equivalence in the case of $\mathcal{S}^{Poss}$, $\mathcal{S}^{\kappa}$, and $\mathcal{S}^{p}$. □

Notice that Proposition 6.2 does not apply to $\mathcal{S}^{\epsilon}$. Consider the PPD $(\Pr_1, \Pr_2, \ldots)$ such that $\Pr_n(a) = 1/n$, $\Pr_{2n-1}(b) = 1 - 1/n$, $\Pr_{2n-1}(c) = 0$, $\Pr_{2n}(b) = 0$, $\Pr_{2n}(c) = 1 - 1/n$ for all $n \geq 1$. It is easy to check that the plausibility space arising from this PPD is default-isomorphic to the one in Proposition 6.1.

It is not hard to construct a trivial plausibility structure that is not default-isomorphic to any structure in $\mathcal{S}^{\epsilon}$: Consider the trivial plausibility measure on $\{a\}$ such that $\mathrm{Pl}(\{a\}) = \bot$. This cannot be default-isomorphic to a structure in $\mathcal{S}^{\epsilon}$, since if Pl' is a plausibility measure in such a structure, we must have $\mathrm{Pl}'(\{a\}) = \top > \mathrm{Pl}'(\emptyset)$. But this is essentially all that can go wrong. We say that a plausibility space (W,Pl) (resp. plausibility structure (W,Pl,π)) is *normal* (following Lewis [1973]) if $\mathrm{Pl}(W) > \bot$. It is easy to see that all structures in $\mathcal{S}^{\epsilon}$, $\mathcal{S}^{\kappa}$, and $\mathcal{S}^{Poss}$ are normal.

Theorem 6.3. *If $PL \in \mathcal{S}^{QPL}$ is a normal plausibility structure for a countable language $\mathcal{L}$, then there is a structure $PL' \in \mathcal{S}^{\epsilon}$ that is default-equivalent to PL.*

Proof. See Appendix A.3. □

Corollary 6.4. *If (W,Pl,π) is a normal, qualitative plausibility structure for a countable language $\mathcal{L}$, then there exists a structure $(W,\mathrm{Pl}',\pi) \in \mathcal{S}^{\epsilon}$ such that $(W,\mathrm{Pl},\pi) \models \phi{\rightarrow}\psi$ if and only if $(W,\mathrm{Pl}',\pi) \models \phi{\rightarrow}\psi$ for all $\phi,\psi \in \mathcal{L}$.*

Thus, with respect to conditional statements in a countable language, $\mathcal{S}^\epsilon$ is as expressive (in a strong sense) as $\mathcal{S}^{QPL}$. However, for uncountable languages, there is a difference between $\mathcal{S}^{QPL}$ and $\mathcal{S}^\epsilon$: Probability distributions can assign positive weight only to a countable number of pairwise disjoint events, while qualitative plausibility measures do not suffer from such constraints.

7. EPISTEMIC ENTRENCHMENT

There has been much work related to defaults and plausibility. It can roughly be divided into three categories. The first consists of various approaches to dealing with uncertainty such as the ones mentioned in Section 2. For a more detailed comparison to such approaches see [Friedman and Halpern 1995]. The second category consists of semantics for defaults that are discussed at length in Section 4. The final category includes semantics for defaults that are linguistic in nature. The most well known approach of this kind is *epistemic entrenchment* [Gärdenfors and Makinson 1988; Grove 1988]. This has been proposed as a semantics for *belief revision* [Gärdenfors 1988]. Recently, Gärdenfors and Makinson [1989] proposed using a similar notion of expectation ordering as a semantics for default reasoning. We briefly review their approach here.

Let $\mathcal{L}$ be some logical language that includes the usual propositional connectives with a compact consequence relation $\vdash_{\mathcal{L}}$ that satisfies the axioms of propositional logic. An *expectation ordering* $\trianglelefteq$ on $\mathcal{L}$ is a relation over formulas in L that satisfies the following requirements:

E1. $\trianglelefteq$ is transitive,

E2. if $\vdash_{\mathcal{L}} \phi \Rightarrow \psi$ then $\phi \trianglelefteq \psi$,

E3. for any ϕ and ψ, either $\phi \trianglelefteq \phi \land \psi$ or $\psi \trianglelefteq \phi \land \psi$.

Intuitively, $\phi \trianglelefteq \psi$ if ψ is as at least as expected as ϕ, so the agent would not retract his belief in ψ before retracting his belief in ϕ. We do not go here into the motivation for E1–E3. It is not hard to verify that E1–E3 imply that $\trianglelefteq$ is a total preorder on $\mathcal{L}$.

An *expectation structure* is a pair $E = (\mathcal{L}, \trianglelefteq)$, where $\trianglelefteq$ is an expectation ordering on $\mathcal{L}$. Intuitively, E satisfies $\phi \!\rightarrow\! \psi$ if ψ is the consequence of formulas that are expected given ϕ. This definition hinges on the choice of formulas that are expected given ϕ. Gärdenfors and Makinson take these to be the formulas that are more expected than $\neg\phi$. Formally, an expectation structure $E = (\mathcal{L}, \trianglelefteq)$ satisfies a default $\phi \!\rightarrow\! \psi$ if $\{\phi\} \cup \{\xi : \neg\phi \lhd \xi\} \vdash_{\mathcal{L}} \psi$, where $\phi \lhd \psi$ holds if $\phi \trianglelefteq \psi$ and not $\psi \trianglelefteq \phi$. The following result is almost immediate from the definitions:

THEOREM 7.1. [Gärdenfors and Makinson 1989] *Let $E = (\mathcal{L}, \trianglelefteq)$ be an expectation structure. $E \models \phi \!\rightarrow\! \psi$ if and only if $\vdash_{\mathcal{L}} \phi \Rightarrow \psi$ or $(\phi \Rightarrow \neg\psi) \lhd (\phi \Rightarrow \psi)$.*

While this definition seems quite different than the one described in Section 4, the two are in fact closely related. Notice that $\phi \Rightarrow \neg\psi$ is equivalent to $\neg(\phi \land \psi)$, while $\phi \Rightarrow \psi$ is equivalent to $\neg(\phi \land \neg\psi)$. Thus, the second clause in the theorem above, which says $(\phi \Rightarrow \neg\psi) \lhd (\phi \Rightarrow \psi)$, can be viewed as saying $\neg(\phi \land \psi)$ is less expected than $\neg(\phi \land \neg\psi)$; this is clearly much in the spirit of the second clause in the definition of defaults in plausibility, which says that $(\phi \land \psi)$ must be more plausible than $(\phi \wedge \neg\psi)$. If we identify p being more plausible than q with $\neg p$ being less expected than $\neg q$, they are equivalent. The first clause in the theorem, $\vdash_L \phi \Rightarrow \psi$, corresponds to the vacuous case that ϕ has plausibility $\perp$ in our definition. However, as we now show, Gärdenfors and Makinson treat the vacuous case in a somewhat nonstandard manner (which can still be captured using plausibility).

To make the relationship between expectation orderings and plausibility precise, let $E = (\mathcal{L}, \trianglelefteq)$ be an expectation structure. We say that a set $V \subseteq \mathcal{L}$ is *consistent* if for all $\phi_1, \ldots, \phi_n \in V$, we have $\not\vdash_{\mathcal{L}} \neg(\phi_1 \wedge \ldots \wedge \phi_n)$. V is a *maximal consistent* set if it is consistent and for each $\phi \in \mathcal{L}$, either $\phi \in V$ or $\neg\phi \in V$. We now construct a plausibility structure $PL_E = (W_E, \mathrm{Pl}_E, \pi_E)$. We define $W_E = \{w_V : V$ is a maximally consistent subset of $\mathcal{L}\}$ and $\pi_E(w_V)(p) = \mathbf{true}$ if $p \in V$. Finally, we need to define Pl_E. The obvious choice is to define $\mathrm{Pl}_E([\![\phi]\!]) \leq \mathrm{Pl}_E([\![\psi]\!])$ if and only if $(\neg\psi) \trianglelefteq (\neg\phi)$. It is easy to see that this implies that $(\phi \Rightarrow \neg\phi) \triangleleft (\phi \Rightarrow \psi)$ if and only if $\mathrm{Pl}_E([\![\phi \wedge \psi]\!]) > \mathrm{Pl}_E([\![\phi \wedge \neg\psi]\!])$. Thus, E and PL_E agree on non-vacuous defaults. However, suppose that $\mathrm{Pl}_E([\![\phi]\!]) = \perp$ for some consistent formula ϕ. It follows that $PL_E \models \phi{\rightarrow}\psi$ for any ψ. On the other hand, it is not hard to show that $E \models \phi{\rightarrow}\psi$ if and only if $\vdash_L \phi \Rightarrow \psi$.[11] We can easily modify the definition of Pl_E to avoid this problem: We define $\mathrm{Pl}_E([\![\phi]\!]) \leq \mathrm{Pl}_E([\![\psi]\!])$ either if $(\neg\psi) \trianglelefteq (\neg\phi)$ and it is not the case that $true \trianglelefteq \neg\psi$ or if $[\![\phi]\!] \subseteq [\![\psi]\!]$. With this modified definition, we get the desired result.

PROPOSITION 7.2. *If E is an expectation structure, then PL_E is a plausibility structure. Furthermore, $E \models \phi{\rightarrow}\psi$ if and only if $PL_E \models \phi{\rightarrow}\psi$.*

PROOF. See Appendix A.4. □

We now examine default entailment with respect to expectation orderings. Let $\mathcal{S}^E$ be the set of plausibility structures that correspond to expectation structures. It is not hard to prove that

THEOREM 7.3. *$\mathcal{S}^E$ is a subset of $\mathcal{S}^{QPL}$.*

PROOF. It is straightforward to verify that if E satisfies E1–E3, then PL_E satisfies A2 and A3. □

It immediately follows that the KLM properties are sound for default entailment with expectation orderings, i.e., with respect to $\mathcal{S}^E$. The KLM properties, however, are not complete with respect to $\mathcal{S}^E$. For example, if p and q are arbitrary primitive propositions, then $p{\rightarrow}false$ entails $q{\rightarrow}false$. This example is a consequence of a property Gärdenfors and Makinson call *consistency preservation*:

—$E \models \phi{\rightarrow}false$ if and only if $\vdash_{\mathcal{L}} \phi \Rightarrow false$.

This property states that ϕ is totally *implausible*—that is, has plausibility $\perp$—if and only if it is inconsistent. This implies that no $\mathrm{Pl}_E \in \mathcal{S}^E$ satisfies $p{\rightarrow}false$, and hence $p{\rightarrow}false$ entails, vacuously, all other defaults. Thus, one cannot specify that

[11] Proof sketch: The "if" direction follows from Theorem 7.1. For the "only if" direction, $\mathrm{Pl}_E([\![\phi]\!]) = \perp$, so we must have $true \trianglelefteq \neg\phi$. Since $\vdash_L \neg\phi \Rightarrow (\phi \Rightarrow \neg\psi)$, it follows from E1 and E2 that $true \trianglelefteq (\phi \Rightarrow \neg\psi)$. Moreover, since $\vdash_L (\phi \Rightarrow \psi) \Rightarrow true$ we have $(\phi \Rightarrow \psi) \trianglelefteq true$. Thus, we cannot have $(\phi \Rightarrow \neg\psi) \triangleleft (\phi \Rightarrow \psi)$, for otherwise, by E2, we would have $true \triangleleft true$, a contradiction.

events such as $white \wedge black$ are impossible in the database Δ; these constraints must be somehow embedded in $\vdash_{\mathcal{L}}$.

We note that expectation orderings are similar to plausibility measures in that they order events. However, there are several differences. First, expectation orderings use formulas to denote events. (We remark that there are similar formulations of probability theory [Jeffreys 1961] that are based on a linguistic description of events.) Secondly, as shown by our construction of PL_E, expectation orderings order events according to the implausibility of their complements. (This type of ordering is usually called the *dual order* [Dubois and Prade 1990; Friedman and Halpern 1995; Shafer 1976].) Thirdly, the treatment of the vacuous case is slightly different. This difference leads to additional properties of default entailment.

8. CONDITIONAL LOGIC

Up to now, we 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 disjunctions of defaults. It is easy to extend all the approaches we 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 closing 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., $(p \rightarrow q) \rightarrow r$).

We note that the connections between default reasoning and conditional logics are well-known; see [Boutilier 1994; Crocco and Lamarre 1992; Kraus et al. 1990; Katsuno and Satoh 1991]. We gloss over the subtle philosophical differences between the two here.

The semantics of conditional logic is similar to the semantics of defaults. As with defaults, we evaluate conditional statements such as $\phi \rightarrow \psi$ by comparing the plausibility of those worlds that satisfy $\phi \wedge \psi$ to the plausibility of those worlds that satisfy $\phi \wedge \neg\psi$. Unlike default reasoning, conditional logic allows us to combine defaults with propositional statements. Thus, $p \wedge (q \rightarrow r)$ is a formula of conditional logic, and is satisfied if both p and $q \rightarrow r$ are satisfied. The truth of a formula such as $p \wedge (q \rightarrow r)$ depends on the world; $p \wedge (q \rightarrow r)$ might be true in w_1 and false in w_2 if p is true at w_1 and not at w_2.

Conditional logic also allows us to consider nested conditionals. For example, to evaluate $(p \rightarrow q) \rightarrow r$, we need to consider the plausibility of the worlds that satisfy r and $p \rightarrow q$ and compare them to the plausibility of worlds that satisfy $\neg r$ and $p \rightarrow q$. In the structures we considered in the preceding sections, a statement such as $p \rightarrow q$ is determined by the global plausibility measure. Thus, the set of worlds that satisfy $p \rightarrow q$ is either the empty set or W (i.e., all possible worlds). It is not hard to show that, as a result of this, we can *denest* nested conditional statements. That is, every formula is equivalent to one without nested conditionals. (See [Friedman and Halpern 1994] for a proof of this well-known observation.) The usual definition of conditional logic [Lewis 1973] provides a nontrivial semantics for nested conditionals by associating with each world a different preferential order over worlds. We can give a similar definition based on plausibility measures.

A *(generalized) plausibility structure* is a tuple $(W, \mathcal{P}, \pi)$ where W and π are, as usual, a set of worlds and a mapping from worlds to truth assignments, and $\mathcal{P}$ maps each world w to a plausibility space $(W_w, \mathcal{F}_w, \mathrm{Pl}_w)$ where $W_w \subseteq W$. Intuitively, $(W_w, \mathcal{F}_w, \mathrm{Pl}_w)$ describes the agent's plausibility when she is in world w. We can view the plausibility structures we defined in previous sections to be a special case of generalized plausibility structures where $\mathcal{P}(w)$ is the same for all worlds w. For the remainder of this section we focus on generalized plausibility structures, but continue to refer to them as plausibility structures.

Given a plausibility structure $PL = (W, \mathcal{P}, \pi)$, 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; for $\rightarrow$, we use the definition already given:

—$(PL, w) \models p$ if $\pi(w) \models p$ for a primitive proposition p

—$(PL, w) \models \neg\phi$ if $(PL, w) \not\models \phi$

—$(PL, w) \models \phi \wedge \psi$ if $(PL, w) \models \phi$ and $(PL, w) \models \psi$

—$(PL, w) \models \phi \rightarrow \psi$ if either $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) = \perp$ or $\mathrm{Pl}_w([\![\phi \wedge \psi]\!]_{(PL,w)}) > \mathrm{Pl}_w([\![\phi \wedge \neg\psi]\!]_{(PL,w)})$, where we define $[\![\phi]\!]_{(PL,w)} = \{w \in W_w : (\mathrm{Pl}, w) \models \phi\}$.

We can similarly define generalized structures that use preferential orderings, κ-rankings, ϵ-semantics, or possibility measures instead of plausibility measures. As before, all of these structures can be embedded in qualitative plausibility structures. We denote by $\mathcal{S}_c^{QPL}$ the class of all qualitative generalized plausibility structures, and similarly denote the subclasses that correspond to various semantics (e.g., $\mathcal{S}_c^p$ is the class that consists of plausibility structures based on preference orderings).

We saw that with default reasoning, we could not distinguish between plausibility, possibility, preference orderings, ϵ-semantics, and κ-rankings. What happens when we move to the richer language of conditional logic? As we shall see, this richer language allows us to make some finer distinctions.

We start by examining preferential structures. There is a complete axiomatization for conditional logic with respect to preferential structures due to Burgess [1981] called System **C**, consisting of the following axioms and inference rules:

C0. All substitution instances of propositional tautologies

C1. $\phi \rightarrow \phi$

C2. $((\phi \rightarrow \psi_1) \wedge (\phi \rightarrow \psi_2)) \Rightarrow (\phi \rightarrow (\psi_1 \wedge \psi_2))$

C3. $((\phi_1 \rightarrow \psi) \wedge (\phi_2 \rightarrow \psi)) \Rightarrow ((\phi_1 \vee \phi_2) \rightarrow \psi)$

C4. $((\phi_1 \rightarrow \phi_2) \wedge (\phi_1 \rightarrow \psi)) \Rightarrow ((\phi_1 \wedge \phi_2) \rightarrow \psi)$

R1. From ϕ and $\phi \Rightarrow \psi$ infer ψ

RC1. From $\phi \Leftrightarrow \phi'$ infer $(\phi \rightarrow \psi) \Rightarrow (\phi' \rightarrow \psi)$

RC2. From $\psi \Rightarrow \psi'$ infer $(\phi \rightarrow \psi) \Rightarrow (\phi \rightarrow \psi')$

System **C** can be viewed as a generalization of system **P**. The richer language lets us replace a rule like AND by the axiom C2. Similarly, C1, C3, C4, RC1, and RC2 are the analogues of REF, OR, CM, LLE, and RW, respectively. We need C0 and R1 to deal with propositional reasoning.

THEOREM 8.1. [Burgess 1981] *System* **C** *is a sound and complete axiomatization of* $\mathcal{L}^C$ *with respect to* $\mathcal{S}_c^p$.

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

THEOREM 8.2. *System* **C** *is a sound and complete axiomatization of* $\mathcal{L}^C$ *with respect to* $\mathcal{S}_c^{QPL}$.

The proof of Theorem 8.1 given by Burgess is quite complicated. We can get a simpler direct proof of Theorem 8.2, without going through Theorem 8.1, by using standard techniques of modal logic. We provide the details in Appendix A.5.

Theorems 8.1 and 8.2 show that, even in the richer framework of conditional logic, we cannot distinguish between preferential orders and plausibility, at least not axiomatically. What about the other approaches we have been considering?

Not surprisingly, conditional logic does allow us to distinguish rational structures from arbitrary preferential structures, because now we can express RM within the language, using the following axiom:

C5. $\phi \rightarrow \psi \wedge \neg(\phi \rightarrow \neg\xi) \Rightarrow \phi \wedge \xi \rightarrow \psi$

Does C5 (together with system **C**) characterize $\mathcal{S}_c^r$? Almost, but not quite. We say that a plausibility measure Pl is *rational* if it satisfies the following two properties:

A4. For all pairwise disjoint sets A, B and C, if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$, then $\mathrm{Pl}(A) < \mathrm{Pl}(C)$ or $\mathrm{Pl}(C) < \mathrm{Pl}(B)$.

A5. For all pairwise disjoint sets A, B and C, if $\mathrm{Pl}(A) < \mathrm{Pl}(B \cup C)$, then $\mathrm{Pl}(A) < \mathrm{Pl}(B)$ or $\mathrm{Pl}(A) < \mathrm{Pl}(C)$.

A4 says that ordering of disjoint sets is modular. (Recall that an ordering is modular if there are no three elements x, y, z such that $x > y$ and z incomparable to both x and y.) A5 says that the plausibility of $B \cup C$ is essentially the maximum of the plausibility of B and C. Thus, $B \cup C$ cannot be more plausible than A, unless one of the components is more plausible than A.

It is not hard to show that C5 is valid in, and only in, systems where the plausibility ordering is rational.

PROPOSITION 8.3. *Let* $\mathcal{S} \subseteq \mathcal{S}_c^{QPL}$. *C5 is valid in* $\mathcal{S}$ *if and only if all structures in* $\mathcal{S}$ *are rational.*

PROOF. See Appendix A.5. □

It is easy to verify that rational preference orderings, κ-rankings, and possibility measures are all classes of rational structures. It immediately follows that C5 is valid in each of $\mathcal{S}_c^r$, $\mathcal{S}_c^{Poss}$, and $\mathcal{S}_c^\kappa$. On the other hand, C5 it is not valid in $\mathcal{S}_c^p$, since we can easily construct preferential structures that violate A4 and A5.

As we said above, condition A4 states that the ordering is "almost" modular, in the sense that, when restricted to pairwise disjoint sets, it is modular. It is not surprising to see modularity arise in this context. It is well known that a modular ordering induces a total order. More precisely, if $<$ is a modular order on some set and we define $x \leq y$ as $y \not< x$, then $\leq$ is a *total* order, that is, a partial order such that either $x \leq y$ or $y \leq x$ for all x and y. All approaches that satisfy rational monotonicity that have been proposed in the literature involve structures where there is a total or modular order on worlds (e.g., rational preference orderings, κ-rankings, and possibility measures).

We say that a plausibility measure Pl is a *ranking* if it satisfies the following two properties:

A4′. $\leq_D$ is a total order; that is, either $\mathrm{Pl}(A) \leq_D \mathrm{Pl}(B)$ or $\mathrm{Pl}(B) \leq_D \mathrm{Pl}(A)$ for all sets A, B.

A5′. $\mathrm{Pl}(A \cup B) = \max(\mathrm{Pl}(A), \mathrm{Pl}(B))$ for all sets A, B.

It is easy to see that A4′ implies A4, and that in the presence of A4′, A5′ implies A5 (A4′ is required to ensure that the two plausibilities values have a maximum). Thus, any ranking is a rational measure. The opposite, however, is not true. It is easy to verify that A4 and A5 do not imply A4′ and A5′. Thus, there is a discrepancy between the properties that are necessary to satisfy C5 and those studied in the literature. However, we now show that if we care only about defaults, then there is no difference between rational structures and *ranked (plausibility) structures*, where the plausibility measure is a ranking.

THEOREM 8.4. *If* (W, Pl) *be a rational qualitative plausibility space, then there is a default-equivalent plausibility space* (W, Pl') *such that* Pl' *is a ranking.*

PROOF. See Appendix A.5. □

COROLLARY 8.5. *If* $PL = (W, \mathcal{P}, \pi)$ *is a rational plausibility structure, then there is a ranked plausibility structure* $PL' = (W, \mathcal{P}', \pi)$ *such that* $(PL, w) \models \phi$ *if and only if* $(PL', w) \models \phi$ *for all worlds* w *and formulas* $\phi \in \mathcal{L}^C$.

Conditional logic allows us to capture another property that we encountered earlier. Recall that measures based on PPDs, possibility measures, and κ-rankings are all *normal*, that is, that $\mathrm{Pl}(W) > \perp$. This property corresponds to the axiom

C6. $\neg(\mathit{true} \rightarrow \mathit{false})$.

It is not hard to show that C6 is valid in each of $\mathcal{S}_c^{Poss}, \mathcal{S}_c^{\kappa}$, and $\mathcal{S}^{\epsilon}$.

Using C5 and C6 we can characterize $\mathcal{S}_c^{\epsilon}$, $\mathcal{S}_c^{\kappa}$, and $\mathcal{S}_c^{Poss}$.

THEOREM 8.6. *(a)* **C**+{*C6*} *is a sound and complete axiomatization of* $\mathcal{S}_c^{\epsilon}$.

(b) **C** +{*C5, C6*} *is a sound and complete axiomatization of* $\mathcal{S}_c^{\kappa}$ *and* $\mathcal{S}_c^{Poss}$.

PROOF. See Appendix A.5. □

9. CONCLUSIONS

We feel that this paper makes three major contributions: the introduction of plausibility measures, the unification of all earlier results regarding the KLM properties into one framework, and a general result showing the inevitability of these properties.

Do we really need plausibility measures? That depends on the language we are interested in. If all we are interested in is propositional default reasoning and the KLM properties, then, as is well known (and our results emphasize), many different approaches turn out to be equivalent in expressive power. If we move to the richer language of propositional conditional logic, then, as the results of

Section 8 show, we start to see some differences (that are captured by axioms C5 and C6, which correspond to rationality and normality, respectively), although plausibility structures and preferential structures continue to be characterized by the same axioms. As we show in a companion paper [Friedman et al. 1996], once we move to first-order conditional logic, more significant differences start to appear. The extra expressive power of plausibility structures makes them more appropriate than preferential structures for providing semantics for first-order default reasoning. This difference is due to the fact 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; see Lemma A.9.) In finite structures, preferential orders and plausibility measures are equi-expressive. The differences that we observe between them in first-order conditional logic are due to the fact that in first-order reasoning, infinite structures play a more important role.

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 and Halpern 1995]. 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 and Halpern 1996], we use plausibilistic independence to define a plausibilistic analogue of Markov chains. In future work we plan to explore further the properties and applications of plausibility structures.

ACKNOWLEDGMENTS

The authors are grateful to Ronen Brafman, Adnan Darwiche, Moises Goldszmidt, Adam Grove, Daphne Koller, Daniel Lehmann, Karl Schlechta, and Zohar Yakhini for useful discussions relating to this work.

APPENDIX

A. DETAILED PROOFS

A.1 Proofs for Section 4

LEMMA 4.1. *Let W be a set of possible worlds and let π be a function that maps each world in W to a truth assignment to $\mathcal{L}$. Let $T \subseteq \mathcal{L}_{\mathrm{def}}$ be a set of defaults that is closed under the rules of system* **P** *that satisfies the following condition:*

($*$) *if $\phi \rightarrow \psi \in T$, $[\![\phi]\!] = [\![\phi']\!]$, and $[\![\psi]\!] = [\![\psi']\!]$, then $\phi' \rightarrow \psi' \in T$, for all formulas $\phi, \phi', \psi, \psi' \in \mathcal{L}$.*

There is a plausibility structure $PL_T = (W, \mathrm{Pl}_T, \pi)$ such that $\mathrm{Pl}_T([\![\phi]\!]) \leq \mathrm{Pl}_T([\![\psi]\!])$ if and only if $\phi \vee \psi \rightarrow \psi \in T$. Moreover, $PL_T \models \phi \rightarrow \psi$ if and only if $\phi \rightarrow \psi \in T$.

PROOF. We start by noting that the condition imposed on T ensures that Pl_T is well defined, i.e., sets that are described by different formulas are compared in a consistent manner. We now examine whether there exists a plausibility measure (W, Pl_T) satisfying ($*$). It suffices to show that the order relation on Pl_T is reflexive, transitive, and satisfies A1.

Reflexivity. Applying REF and LLE, we get that $(\phi \vee \phi){\rightarrow}\phi \in T$ for all ϕ.

Transitivity. This is a direct consequence of the following lemma of Kraus, Lehmann and Magidor.

LEMMA A.1. [Kraus et al. 1990, Lemma 5.5] *Let T be a set of defaults closed under applications of the rules of system* **P**. *Then if both $\phi_1 \vee \phi_2{\rightarrow}\phi_2$ and $\phi_2 \vee \phi_3{\rightarrow}\phi_3$ are in T, then so is $\phi_1 \vee \phi_3{\rightarrow}\phi_3$.*

A1. Assume $[\![\phi]\!] \subseteq [\![\psi]\!]$. It follows that $[\![\psi]\!] = [\![\psi \vee \phi]\!]$. Since T is closed under REF we get that $(\phi \vee \psi){\rightarrow}(\phi \vee \psi) \in T$. Using $(*)$, we get that $(\phi \vee \psi){\rightarrow}\psi \in T$. Thus, $\mathrm{Pl}_T([\![\phi]\!]) \leq \mathrm{Pl}_T([\![\psi]\!])$.

Finally, we need to show that $(W, \mathrm{Pl}_T, \pi) \models \phi{\rightarrow}\psi$ if and only if $\phi{\rightarrow}\psi \in T$. We start by observing that LLE, REF, and AND imply that $\phi{\rightarrow}\psi \in T$ if and only if $(\phi \wedge \neg\psi) \vee (\phi \wedge \psi){\rightarrow}\phi \wedge \psi \in T$. We conclude that

$$(**)\ \phi{\rightarrow}\psi \in T \text{ if and only if } \mathrm{Pl}_T([\![\phi \wedge \neg\psi]\!]) \leq \mathrm{Pl}_T([\![\phi \wedge \psi]\!]).$$

Thus, it suffices to show that $\mathrm{Pl}_T([\![\phi \wedge \neg\psi]\!]) \leq \mathrm{Pl}_T([\![\phi \wedge \psi]\!])$ if and only if either $\mathrm{Pl}_T([\![\phi \wedge \neg\psi]\!]) < \mathrm{Pl}_T([\![\phi \wedge \psi]\!])$ or $\mathrm{Pl}_T([\![\phi]\!]) = \bot$. The "if" direction is trivial. For the "only if" direction, suppose by way of contradiction that we have $\mathrm{Pl}_T([\![\phi]\!]) > \bot$, $\mathrm{Pl}_T([\![\phi \wedge \neg\psi]\!]) \leq \mathrm{Pl}_T([\![\phi \wedge \psi]\!])$, and $\mathrm{Pl}_T([\![\phi \wedge \psi]\!]) \leq \mathrm{Pl}_T([\![\phi \wedge \neg\psi]\!])$. From $(**)$, we have that $\phi{\rightarrow}\psi \in T$ and $\phi{\rightarrow}\neg\psi \in T$. By AND and $(*)$, we have that $\phi{\rightarrow}\mathit{false} \in T$. But then $\mathrm{Pl}_T([\![\phi]\!]) = \bot$, which contradicts our assumptions. □

A.2 Proofs for Section 5

THEOREM 5.2. *If $\Delta \vdash_{\mathbf{P}'} \phi{\rightarrow}\psi$, then $\Delta \models_{\mathcal{S}^{PL}} \phi{\rightarrow}\psi$.*

PROOF. We need to show that LLE, RW, and REF are sound in $\mathcal{S}^{PL}$. Let $PL = (W, \mathrm{Pl}, \pi)$ be a plausibility structure. We proceed as follows.

LLE. Assume that $\vdash_{\mathcal{L}} \phi \Leftrightarrow \phi'$. Then, by definition, $[\![\phi]\!] = [\![\phi']\!]$. The soundness of LLE immediately follows.

RW. Assume that $\vdash_{\mathcal{L}} \psi \Rightarrow \psi'$ and that $PL \models \phi{\rightarrow}\psi$. We want to show that $PL \models \phi{\rightarrow}\psi'$. If $\mathrm{Pl}([\![\phi]\!]) = \bot$, this is immediate. On the other hand, if $\mathrm{Pl}([\![\phi]\!]) > \bot$, then $\mathrm{Pl}([\![\phi \wedge \neg\psi]\!]) < \mathrm{Pl}([\![\phi \wedge \psi]\!])$. Since $\vdash_{\mathcal{L}} \psi \Rightarrow \psi'$ we have that $[\![\psi]\!] \subseteq [\![\psi']\!]$. It follows that $[\![\phi \wedge \neg\psi']\!] \subseteq [\![\phi \wedge \neg\psi]\!]$ and $[\![\phi \wedge \psi]\!] \subseteq [\![\phi \wedge \psi']\!]$. Using A1, we conclude that $\mathrm{Pl}([\![\phi \wedge \psi']\!] > \mathrm{Pl}([\![\phi \wedge \neg\psi']\!]$, so $PL \models \phi{\rightarrow}\psi'$.

REF. By definition, $[\![\phi \wedge \neg\phi]\!] = \emptyset$ and $\mathrm{Pl}(\emptyset) = \bot$. Thus, if $\mathrm{Pl}([\![\phi]\!]) > \bot$, then $PL \models \phi{\rightarrow}\phi$. On the other hand, if $\mathrm{Pl}([\![\phi]\!]) = \bot$, then $PL \models \phi{\rightarrow}\phi$ vacuously.

□

PROPOSITION 5.3. *A plausibility measure satisfies A2 if and only if it satisfies A2′.*

PROOF. Let (W, Pl) be a plausibility space. Assume that Pl satisfies A2. Let A, B_1, and B_2 be sets such that $\mathrm{Pl}(A \cap B_1) > \mathrm{Pl}(A \cap \overline{B_1})$ and $\mathrm{Pl}(A \cap B_2) > \mathrm{Pl}(A \cap \overline{B_2})$. Set $C = A \cap B_1 \cap B_2$, $D = A \cap B_1 \cap \overline{B_2}$, and $E = A \cap \overline{B_1}$. It is easy to verify that C, D, and E are pairwise disjoint. Since $C \cup D = A \cap B_1$, we have that $\mathrm{Pl}(C \cup D) > \mathrm{Pl}(E)$. Moreover, since $C \cup E \supseteq A \cap B_2$, $D \subseteq A \cap \overline{B_2}$, and $\mathrm{Pl}(A \cap B_2) > \mathrm{Pl}(A \cap \overline{B_2})$, we can apply A1 and conclude that $\mathrm{Pl}(C \cup E) > \mathrm{Pl}(D)$. Applying A2, we conclude that $\mathrm{Pl}(C) > \mathrm{Pl}(D \cup E)$, i.e., $\mathrm{Pl}(A \cap B_1 \cap B_2) > \mathrm{Pl}(A \cap \overline{(B_1 \cap B_2)})$. This gives us A2′.

Now assume that Pl satisfies A2′. Let C, D, and E be pairwise disjoint sets such that $\mathrm{Pl}(C \cup D) > \mathrm{Pl}(E)$ and $\mathrm{Pl}(C \cup E) > \mathrm{Pl}(D)$. Let $A = C \cup D \cup E$, $B_1 = C \cup D$, and $B_2 = C \cup E$. Then we have $\mathrm{Pl}(A \cap B_1) = \mathrm{Pl}(C \cup D) > \mathrm{Pl}(E) = \mathrm{Pl}(A \cap \overline{B_1})$ and $\mathrm{Pl}(A \cap B_2) = \mathrm{Pl}(C \cup E) > \mathrm{Pl}(D) = \mathrm{Pl}(A \cap \overline{B_2})$. From A2′, we have that $\mathrm{Pl}(A \cap B_1 \cap B_2) > \mathrm{Pl}(A \cap \overline{(B_1 \cap B_2)})$, i.e., $\mathrm{Pl}(C) > \mathrm{Pl}(D \cup E)$. This gives us A2. □

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

Proof. To prove the "if" direction it suffices to show that each rule in system **P** is sound in qualitative structures. Let $PL = (W, \mathrm{Pl}, \pi)$ be a qualitative plausibility structure. The soundness of LLE, RW, and REF is proved in Theorem 5.2. To deal with the remaining cases, we proceed as follows.

AND. Assume that $PL \models \phi {\rightarrow} \psi_1$ and $PL \models \phi {\rightarrow} \psi_2$. If $\mathrm{Pl}([\![\phi]\!]) = \bot$, then $PL \models \phi {\rightarrow} \psi_1 \wedge \psi_2$ vacuously. Assume that $\mathrm{Pl}([\![\phi]\!]) > \bot$. Let $A = [\![\phi]\!]$, $B_1 = [\![\psi_1]\!]$, and $B_2 = [\![\psi_2]\!]$. Since $PL \models \phi {\rightarrow} \psi_1$ and $PL \models \phi {\rightarrow} \psi_2$, we have that $\mathrm{Pl}(A \cap B_1) > \mathrm{Pl}(A \cap \overline{B_1})$ and $\mathrm{Pl}(A \cap B_2) > \mathrm{Pl}(A \cap \overline{B_2})$. Proposition 5.3 states that Pl satisfies A2′, and thus we get that $\mathrm{Pl}(A \cap B_1 \cap B_2) > \mathrm{Pl}(A \cap \overline{(B_1 \cap B_2)})$, so $PL \models \phi {\rightarrow} \psi_1 \wedge \psi_2$.

CM. Again assume that $PL \models \phi {\rightarrow} \psi_1$ and $PL \models \phi {\rightarrow} \psi_2$. If $\mathrm{Pl}([\![\phi \wedge \psi_1]\!]) = \bot$, then $PL \models \phi \wedge \psi_1 {\rightarrow} \psi_2$ vacuously. Assume that $\mathrm{Pl}([\![\phi \wedge \psi_1]\!]) > \bot$. Let A, B_1, and B_2 be defined as in the treatment of AND above. Again, we have $\mathrm{Pl}((A \cap B_1 \cap B_2) > \mathrm{Pl}(A \cap \overline{(B_1 \cap B_2)})$. Since $A \cap B_1 \cap \overline{B_2} \subseteq A \cap \overline{(B_1 \cap B_2)}$, we conclude that $\mathrm{Pl}(A \cap B_1 \cap B_2) > \mathrm{Pl}(A \cap B_1 \cap \overline{B_2})$. Thus, $PL \models \phi \wedge \psi_1 {\rightarrow} \psi_2$.

OR. Assume that $PL \models \phi_1 {\rightarrow} \psi$ and $PL \models \phi_2 {\rightarrow} \psi$. If $\mathrm{Pl}([\![\phi_1]\!]) = \mathrm{Pl}([\![\phi_2]\!]) = \bot$, then applying A3 we get that $\mathrm{Pl}([\![\phi_1 \vee \phi_2]\!]) = \bot$ and thus $PL \models (\phi_1 \vee \phi_2) {\rightarrow} \psi$ vacuously. So assume that $\mathrm{Pl}([\![\phi_1]\!]) > \bot$. (Identical argument works if $\mathrm{Pl}([\![\phi_2]\!]) > \bot$.) Set $A = [\![(\phi_1 \vee \phi_2) \wedge \psi]\!]$, $B = [\![\phi_1 \wedge \neg\psi]\!]$, and $C = [\![\phi_2 \wedge \neg\phi_1 \wedge \neg\psi]\!]$. To prove that $PL \models (\phi_1 \vee \phi_2) {\rightarrow} \psi$, we must show that $\mathrm{Pl}(A) > \mathrm{Pl}(B \cup C)$. Since $PL \models \phi_1 {\rightarrow} \psi$, we have that $\mathrm{Pl}(A) \geq \mathrm{Pl}([\![\phi_1 \wedge \psi]\!]) > \mathrm{Pl}(B)$. If $\mathrm{Pl}([\![\phi_2]\!]) = \bot$, then $\mathrm{Pl}(C) = \bot$ and we conclude that $\mathrm{Pl}(A) > \mathrm{Pl}(C)$. On the other hand, if $\mathrm{Pl}([\![\phi_2]\!]) > \bot$, then since $PL \models \phi_2 {\rightarrow} \psi$, we have $\mathrm{Pl}(A) \geq \mathrm{Pl}([\![\phi_2 \wedge \psi]\!]) > \mathrm{Pl}([\![\phi_2 \wedge \neg\psi]\!]) \geq \mathrm{Pl}(C)$. From A2, $\mathrm{Pl}(A) > \mathrm{Pl}(B)$, and $\mathrm{Pl}(A) > \mathrm{Pl}(C)$, we get that $\mathrm{Pl}(A) > \mathrm{Pl}(B \cup C)$. Thus, $PL \models (\phi_1 \vee \phi_2) {\rightarrow} \psi$.

To prove the "only if" direction we have to show that if there is some $PL = (W, \mathrm{Pl}, \pi)$ in $\mathcal{S}$ that is not qualitative, then the KLM properties are not sound with respect to P. Assume that Pl does not satisfy A2. Since we have assumed that $\mathcal{F} = \{[\![\phi]\!] : \phi \in \mathcal{L}\}$, there are formulas ϕ, ψ_1, and ψ_2, such that $[\![\phi]\!]$, $[\![\psi_1]\!]$, and $[\![\psi_2]\!]$ are pairwise disjoint, $\mathrm{Pl}([\![\phi \vee \psi_2]\!]) > \mathrm{Pl}([\![\psi_1]\!])$, $\mathrm{Pl}([\![\phi \vee \psi_1]\!]) > \mathrm{Pl}([\![\psi_2]\!])$, and yet $\mathrm{Pl}([\![\phi]\!]) \not> \mathrm{Pl}([\![\psi_1 \wedge \psi_2]\!])$. Thus, $PL \models (\phi \vee \psi_1 \vee \psi_2) {\rightarrow} \neg\psi_1$ and $PL \models (\phi \vee \psi_1 \vee \psi_2) {\rightarrow} \neg\psi_2$. However, $PL \not\models (\phi \vee \psi_1 \vee \psi_2) {\rightarrow} (\neg\psi_1 \wedge \neg\psi_2)$. This shows that the AND rule is not sound in $\mathcal{S}$. Now assume that there is some $PL = (W, \mathrm{Pl}, \pi)$ that does not satisfy A3. Thus, there are formulas ϕ_1 and ϕ_2 such that $\mathrm{Pl}([\![\phi_1]\!]) = \mathrm{Pl}([\![\phi_2]\!]) = \bot$ and $\mathrm{Pl}([\![\phi_1 \vee \phi_2]\!]) > \bot$. We conclude that $PL \models \phi_1 {\rightarrow} \mathit{false}$ and $PL \models \phi_2 {\rightarrow} \mathit{false}$, but $PL \not\models (\phi_1 \vee \phi_2) {\rightarrow} \mathit{false}$. This shows that the OR rule is not sound in $\mathcal{S}$. □

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

PROOF. It is straightforward to verify that A2 and A3 hold for each structure in $\mathcal{S}^{Poss}$, $\mathcal{S}^{\kappa}$, $\mathcal{S}^{\epsilon}$, $\mathcal{S}^{p}$, and $\mathcal{S}^{r}$.

We start with $\mathcal{S}^{Poss}$. Let (W, Poss, π) be a possibility structure. To prove A2, assume that $A, B, C \subseteq W$ are pairwise disjoint sets such that $\mathrm{Poss}(A \cup B) > \mathrm{Poss}(C)$ and $\mathrm{Poss}(A \cup C) > \mathrm{Poss}(B)$. Since $\mathrm{Poss}(A \cup B) = \max(\mathrm{Poss}(A), \mathrm{Poss}(B))$, we have that $\max(\mathrm{Poss}(A), \mathrm{Poss}(B)) > \mathrm{Poss}(C)$ and that $\max(\mathrm{Poss}(A), \mathrm{Poss}(C)) > \mathrm{Poss}(B)$. It easily follows that $\mathrm{Poss}(A) > \max(\mathrm{Poss}(B), \mathrm{Poss}(C)) = \mathrm{Poss}(B \cup C)$, as required by A2. To prove A3, suppose that $\mathrm{Poss}(A) = \mathrm{Poss}(B) = 0$. Since $\mathrm{Poss}(A \cup B) = \max(\mathrm{Poss}(A), \mathrm{Poss}(B))$, we have that $\mathrm{Poss}(A \cup B) = 0$, as required by A3.

The proof for $\mathcal{S}^{\kappa}$ is identical (replacing max and 0 with min and ∞, respectively).

Next, consider $\mathcal{S}^{\epsilon}$. Let $(W, \{\mathrm{Pr}_i\}, \pi)$ be a PPD structure and let $(W, \mathrm{Pl}_{PP}, \pi)$ be the corresponding structure in $\mathcal{S}^{\epsilon}$. To prove A2, assume that $A, B, C \subseteq W$ are pairwise disjoint sets such that $\mathrm{Pl}_{PP}(A \cup B) > \mathrm{Pl}_{PP}(C)$ and $\mathrm{Pl}_{PP}(A \cup C) > \mathrm{Pl}_{PP}(B)$. We want to show that $\mathrm{Pl}_{PP}(A) > \mathrm{Pl}_{PP}(B \cup C)$. According to the construction of Theorem 4.2, we need to show that $\lim_{i\to\infty} \mathrm{Pr}_i(A \mid A \cup B \cup C) = 1$ and that $\lim_{i\to\infty} \mathrm{Pr}_i(B \cup C \mid A \cup B \cup C) \neq 1$ (the limit can be undefined in this case). Since $\mathrm{Pl}_{PP}(A \cup B) > \mathrm{Pl}_{PP}(C)$ and $\mathrm{Pl}_{PP}(A \cup C) > \mathrm{Pl}_{PP}(B)$, we have that

$$\lim_{i\to\infty} \mathrm{Pr}_i(A \cup B \mid A \cup B \cup C) = 1,\ \lim_{i\to\infty} \mathrm{Pr}_i(A \cup C \mid A \cup B \cup C) = 1 \quad (3)$$
$$\lim_{i\to\infty} \mathrm{Pr}_i(B \mid A \cup B \cup C) \neq 1,\ \lim_{i\to\infty} \mathrm{Pr}_i(C \mid A \cup B \cup C) \neq 1 \quad (4)$$

To prove that $\lim_{i\to\infty} \mathrm{Pr}_i(A \mid A \cup B \cup C) = 1$, fix $\epsilon > 0$. From (3), we have that there is an n_ϵ such that for all $i > n_\epsilon$, $\mathrm{Pr}_i(A \cup B \mid A \cup B \cup C) > 1 - \frac{\epsilon}{2}$ and $\mathrm{Pr}_i(A \cup C \mid A \cup B \cup C) > 1 - \frac{\epsilon}{2}$. Let $i > n_\epsilon$. There are two cases. If $\mathrm{Pr}_i(A \cup B \cup C) = 0$, then $\mathrm{Pr}_i(A \mid A \cup B \cup C) = 1$ by definition. If $\mathrm{Pr}_i(A \mid A \cup B \cup C) > 0$, we use the disjointness of A, B, and C to get

$$\mathrm{Pr}_i(A \mid A \cup B \cup C) + \mathrm{Pr}_i(B \mid A \cup B \cup C) \ > \ 1 - \frac{\epsilon}{2}$$
$$\mathrm{Pr}_i(A \mid A \cup B \cup C) + \mathrm{Pr}_i(C \mid A \cup B \cup C) \ > \ 1 - \frac{\epsilon}{2}$$

This implies that $\mathrm{Pr}_i(A \mid A \cup B \cup C) + \mathrm{Pr}_i(A \cup B \cup C \mid A \cup B \cup C) > 2 - \epsilon$. Since $\mathrm{Pr}_i(A \cup B \cup C \mid A \cup B \cup C) = 1$, we get that $\mathrm{Pr}_i(A \mid A \cup B \cup C) > 1 - \epsilon$. We conclude that $\mathrm{Pr}_i(A \mid A \cup B \cup C) > 1 - \epsilon$ for all $i > n_\epsilon$, and thus $\lim_{i\to\infty} \mathrm{Pr}_i(A \mid A \cup B \cup C) = 1$.

To prove that $\lim_{i\to\infty} \mathrm{Pr}_i(B \cup C \mid A \cup B \cup C) \neq 1$, it suffices to find a subsequence on which $\mathrm{Pr}_i(B \cup C \mid A \cup B \cup C) \to 0$. Let $i_1, i_2, \ldots, i_j, \ldots$ be the sequence of indexes such that $\mathrm{Pr}_{i_j}(A \cup B \cup C) > 0$. This sequence must be infinite, for otherwise, since $\mathrm{Pr}_i(B \mid A \cup B \cup C) = 1$ whenever $\mathrm{Pr}_i(A \cup B \cup C) = 0$, we would have that $\lim_{i\to\infty} \mathrm{Pr}_i(B \mid A \cup B \cup C) = 1$, contradicting (4). From $\lim_{i\to\infty} \mathrm{Pr}_i(A \mid A \cup B \cup C) = 1$, we have that $\lim_{j\to\infty} \mathrm{Pr}_{i_j}(A \mid A \cup B \cup C) = 1$. Moreover, since $\mathrm{Pr}_{i_j}(B \cup C \mid A \cup B \cup C) = 1 - \mathrm{Pr}_{i_j}(A \mid A \cup B \cup C)$, we get that $\lim_{j\to\infty} \mathrm{Pr}_{i_j}(B \cup C \mid A \cup B \cup C) = 0$. We conclude that $\lim_{i\to\infty} \mathrm{Pr}_i(B \cup C \mid A \cup B \cup C) \neq 1$.

To prove A3, assume that $A, B \subseteq W$ are such that $\mathrm{Pl}_{PP}(A) = \mathrm{Pl}_{PP}(B) = \perp$. By the construction of Theorem 4.2, we have that $\mathrm{Pl}_{PP}(A) \leq \perp = \mathrm{Pl}_{PP}(\emptyset)$ if $\lim_{i\to\infty} \mathrm{Pr}_i(\emptyset \mid A \cup \emptyset) = 1$. This implies that there is an index n_A such that $\mathrm{Pr}_i(A) = 0$ for all $i > n_A$. Similarly, there is an n_B such that $\mathrm{Pr}_i(B) = 0$ for all $i > n_B$. Hence, $\Pr_i(A \cup B) = 0$ for all $i > \max(n_A, n_B)$. We conclude that $\mathrm{Pl}_{PP}(A \cup B) = \bot$.

Finally, we consider $\mathcal{S}^p$ and $\mathcal{S}^r$. Let $(W, \prec, \pi)$ be a preference structure and let $(W, \mathrm{Pl}_\prec, \pi)$ be the corresponding structure in $\mathcal{S}^p$. To prove A2, assume that $A, B, C \subseteq W$ are pairwise disjoint sets such that $\mathrm{Pl}_\prec(A \cup B) > \mathrm{Pl}_\prec(C)$ and $\mathrm{Pl}_\prec(A \cup C) > \mathrm{Pl}_\prec(B)$. We want to show that $\mathrm{Pl}_\prec(A) > \mathrm{Pl}_\prec(B \cup C)$. It is easy to verify that the construction of Theorem 4.2 is such that disjoint sets cannot have equal plausibilities. Thus, it suffices to show that $\mathrm{Pl}_\prec(A) \geq \mathrm{Pl}_\prec(B \cup C)$. That is, for all $w \in B \cup C$ there is a world $w' \in A$ such that (a) $w' \prec w$ and (b) there is no $w'' \in B \cup C$ such that $w'' \prec w'$. Let $w_{BC} \in B \cup C$. Without loss of generality, we can assume that $w \in B$. Since $\mathrm{Pl}_\prec(A \cup C) > \mathrm{Pl}_\prec(B)$, there is a world $w_{AC} \in A \cup C$ such that $w_{AC} \prec w_{BC}$ and for all $w_B \in B$, $w_B \not\prec w_{AC}$. There are three cases: (1) If $w_{AC} \in C$, then since $\mathrm{Pl}_\prec(A \cup B) > \mathrm{Pl}_\prec(C)$, there is a world $w_{AB} \in A \cup B$ such that $w_{AB} \prec w_{AC}$ and for all $w_C \in C$, $w_C \not\prec w_{AB}$. Since $w_B \not\prec w_{AC}$ for all $w_B \in B$, we get that $w_{AB} \in A$. For requirement (a), by the transitivity of $\prec$ we have that $w_{AB} \prec w_{BC}$. For requirement (b), suppose that $w'' \in B \cup C$. If $w'' \in C$, then we have that $w'' \not\prec w_{AB}$. On the other hand, if $w'' \in B$, then we have that $w'' \not\prec w_{AC}$, and by transitivity $w'' \not\prec w_{AB}$. (2) If $w_{AC} \in A$ and there is a world $w_C \in C$ such that $w_C \prec w_{AC}$, then since $\mathrm{Pl}_\prec(A \cup B) > \mathrm{Pl}_\prec(C)$, there is a world w_{AB} such that $w_{AB} \prec w_C$ and for all worlds $w'' \in C$, $w'' \not\prec w_{AB}$. Again, it follows that $w_{AC} \in A$ and satisfies (a) and (b). Finally, (3) if $w_{AC} \in A$ and for all $w_C \in C$, $w_C \not\prec w_{AC}$, then it is easy to check that w_{AC} satisfies (a) and (b).

To prove A3, we note that the construction of Theorem 4.2 is such that $\mathrm{Pl}_\prec(A) = \bot$ if and only if $A = \emptyset$. A3 immediately follows. $\square$

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

PROOF. For the "if" direction, assume that $\mathcal{S}$ is not rich. We need to show that system $\mathbf{P}$ is not complete for $\models_{\mathcal{S}}$. It is sufficient to construct Δ, ϕ, and ψ such that $\Delta \models_{\mathcal{S}} \phi \rightarrow \psi$ but $\Delta \not\vdash_{\mathbf{P}} \phi \rightarrow \psi$.

We start with a lemma, whose straightforward proof is left to the reader.

LEMMA A.2. *Let $\phi_1, \ldots, \phi_n$ be a collection of mutually exclusive formulas. Let Δ consist of the default $\phi_n \rightarrow$ false and the defaults $\phi_i \vee \phi_j \rightarrow \phi_i$ for all $1 \leq i < j \leq n$. Then $(W, \mathrm{Pl}, \pi) \models \Delta$ if and only if there is some j with $1 \leq j \leq n$ such that*

$$\mathrm{Pl}([\![\phi_1]\!]) > \mathrm{Pl}([\![\phi_2]\!]) > \cdots > \mathrm{Pl}([\![\phi_j]\!]) = \cdots = \mathrm{Pl}([\![\phi_n]\!]) = \bot.$$

Since $\mathcal{S}$ is not rich, there is a collection $\phi_1, \ldots, \phi_n$ that is a counterexample to the definition of richness. Let Δ be the set of defaults defined in Lemma A.2. We claim that if $(W, \mathrm{Pl}, \pi) \in \mathcal{S}$ satisfies all the defaults in Δ, then $\mathrm{Pl}([\![\phi_{n-1}]\!]) = \bot$. To see this, assume that $\mathrm{Pl}([\![\phi_{n-1}]\!]) > \bot$. Then according to Lemma A.2, $\mathrm{Pl}([\![\phi_1]\!]) > \cdots > \mathrm{Pl}([\![\phi_{n-1}]\!]) > \mathrm{Pl}([\![\phi_n]\!]) = \bot$, but this contradicts the assumption that the sequence $\phi_1, \ldots, \phi_n$ is a counterexample to richness. Since $\mathrm{Pl}([\![\phi_{n-1}]\!]) = \bot$ in every structure that satisfies Δ, we conclude that $\Delta \models_{\mathcal{S}} \phi_{n-1} \rightarrow$*false*.

We now show that $\Delta \not\vdash_{\mathbf{P}} \phi_{n-1} \rightarrow$*false*. The easiest way of proving this is by using Theorem 3.1. All we need to show is that there is a preferential structure that satisfies Δ but does not satisfies $\phi_{n-1} \rightarrow$*false*. Let $W = \{w_1, \ldots, w_{n-1}\}$, let $\prec$ be such that $w_i \prec w_j$ for all $i < j$, and let π be such that $\pi(w_i)(\phi_j) = \mathbf{true}$ if and only if $i = j$. It is straightforward to verify that $(W, \prec, \pi)$ satisfies Δ. However, it easy to see that $(W, \prec, \pi) \not\models \phi_{n-1} {\rightarrow} \mathit{false}$. This concludes the proof of "if" direction.

For the "only if" direction, assume that there is some Δ and $\phi {\rightarrow} \psi$ such that $\Delta \models_{\mathcal{S}} \phi {\rightarrow} \psi$ but $\Delta \not\vdash_{\mathbf{P}} \phi {\rightarrow} \psi$. Using Theorem 3.1 we get that $\Delta \not\models_{\mathrm{p}} \phi {\rightarrow} \psi$. Thus, there is some preferential structure $P = (W, \prec, \pi)$ that satisfies the defaults in Δ but not $\phi {\rightarrow} \psi$. In fact, as the following lemma shows, we can assume that P is a *linear* structure.

LEMMA A.3. [Friedman and Halpern 1994] *Let Δ be a finite set of defaults. If there is a preferential structure that satisfies Δ and does not satisfy $\phi {\rightarrow} \psi$, then there is a preferential structure $P = (W, \prec, \pi)$ such that $W = \{w_1, \ldots, w_n\}$, $w_i \prec w_j$ for all $i < j$, $P \models \Delta$ and $P \not\models \phi {\rightarrow} \psi$.*

We now use P to construct a sequence of formulas that will be a counterexample to the richness of $\mathcal{S}$. Let $p_1, \ldots, p_m$ be the propositions that appear in Δ and $\phi {\rightarrow} \psi$. Since Δ is finite, there is a finite number of such propositions. We note that whether a default $\phi {\rightarrow} \psi$ is satisfied in P depends only on the minimal world satisfying ϕ. If $\pi(w_i)$ and $\pi(w_j)$ for some $i < j$ agree on the truth of $p_1, \ldots, p_m$, then w_j cannot be a minimal world for any formula defined using only $p_1 \ldots p_m$. Thus, we can assume, without loss of generality, that for all $w_i \neq w_j$, there is some p_k that is assigned a different truth value by each of the two worlds. We now construct formulas that characterize the truth assignment to $p_1, \ldots, p_m$ in each world in W. Let

$$\phi_i = \bigwedge_{\{j : \pi(w_i)(p_j) = \mathbf{true}\}} p_j \wedge \bigwedge_{\{j : \pi(w_i)(p_j) = \mathbf{false}\}} \neg p_j$$

for $i = 1, \ldots, n$, and let $\phi_{n+1} = \neg(\phi_1 \vee \ldots \vee \phi_n)$. It is easy to verify that these formulas are mutually exclusive.

We now claim that if PL is a plausibility structure where $\mathrm{Pl}([\![\phi_1]\!]) > \cdots > \mathrm{Pl}([\![\phi_{n+1}]\!]) = \bot$, then PL satisfies the defaults in Δ but not $\phi {\rightarrow} \psi$. This will suffice to prove that $\mathcal{S}$ is not rich, since if $\mathcal{S}$ contains such a structure we get a contradiction to the assumption that $\Delta \models_{\mathcal{S}} \phi {\rightarrow} \psi$.

Let PL be a plausibility structure where $\mathrm{Pl}([\![\phi_1]\!]) > \cdots > \mathrm{Pl}([\![\phi_{n+1}]\!]) = \bot$. We want to show that $PL \models \xi {\rightarrow} \xi'$ if and only if $P \models \xi {\rightarrow} \xi'$, for all formulas ξ and ξ' defined using only $p_1, \ldots, p_m$.

Let ξ, ξ' be formulas defined over $p_1, \ldots, p_m$. Assume that $P \models \xi {\rightarrow} \xi'$. There are two cases: either (a) ξ is not satisfied in W, or (b) the minimal world satisfying ξ also satisfies ξ'. In case (a), it is easy to see that $\vdash_{\mathcal{L}} \xi \Rightarrow \phi_{n+1}$. From A1, we have that $\mathrm{Pl}(\xi) = \bot$, and thus $PL \models \xi {\rightarrow} \xi'$ vacuously. In case (b) assume that w_i is the minimal world satisfying ξ. Since $P \models \xi {\rightarrow} \xi'$ we have that $\pi(w_i) \models \xi \wedge \xi'$. A simple argument shows that $\vdash_{\mathcal{L}} \phi_i \Rightarrow \xi \wedge \xi'$. Thus, using A1, we get that $\mathrm{Pl}([\![\xi \wedge \xi']\!]) \geq \mathrm{Pl}([\![\phi_i]\!])$. Since w_i is the minimal world satisfying ξ and it also satisfies ξ', we have that $\xi \wedge \neg\xi'$ is not satisfied by $w_1, \ldots, w_i$. Since $\phi_1, \ldots, \phi_{n+1}$ are exhaustive, we have that $\vdash_{\mathcal{L}} (\xi \wedge \neg\xi') \Rightarrow (\phi_{i+1} \vee \ldots \vee \phi_{n+1})$. Thus, $\mathrm{Pl}([\![\xi \wedge \neg\xi']\!]) \leq \mathrm{Pl}([\![\phi_{i+1} \vee \ldots \vee \phi_{n+1}]\!])$. By repeated applications of A2 and the fact that $\mathrm{Pl}([\![\phi_i]\!]) > \mathrm{Pl}([\![\phi_j]\!])$ for all $j > i$, we get that $\mathrm{Pl}([\![\phi_i]\!]) > \mathrm{Pl}[\![\xi \wedge \neg\xi']\!])$. We conclude that $\mathrm{Pl}([\![\xi \wedge \xi']\!]) > \mathrm{Pl}([\![\xi \wedge \neg\xi']\!])$ and thus $PL \models \xi {\rightarrow} \xi'$.

Now assume that $P \not\models \xi {\rightarrow} \xi'$. Thus, there is a minimal world w_i that satisfies ξ; moreover, w_i does not satisfy ξ'. This implies that $P \models \xi {\rightarrow} \neg\xi'$ and $\mathrm{Pl}([\![\xi]\!]) > \bot$.

Applying the proof in the previous paragraph, we have that $\mathrm{Pl}([\![\xi \wedge \neg\xi']\!]) > \mathrm{Pl}([\![\xi \wedge \xi']\!])$, and thus $PL \not\models \xi \rightarrow \xi'$. $\square$

A.3 Proofs for Section 6

We now prove Theorem 6.3. We start with two preliminary lemmas.

LEMMA A.4. *Let $(W, \mathcal{F}, \mathrm{Pl})$ be a normal qualitative plausibility space such that $\mathcal{F}$ is finite, let $A^*, B^* \in \mathcal{F}$ be disjoint sets such that $\mathrm{Pl}(A^*) \not< \mathrm{Pl}(B^*)$, and let $x \geq 2$. Then there is a probability measure $\Pr$ over W such that $\Pr(B^*|A^* \cup B^*) \leq \frac{1}{2}$; moreover if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$ then $\Pr(B|A \cup B) \geq 1 - \frac{1}{x}$, for all disjoint sets $A, B \in \mathcal{F}$.*

PROOF. Let $A_1, \ldots, A_n$ be the *atoms* of $\mathcal{F}$, i.e., each $A_i \neq \emptyset$ is in $\mathcal{F}$ and there is no nonempty $B \in \mathcal{F}$ such that $B \subset A_i$. Since $\mathcal{F}$ is finite, every set in $\mathcal{F}$ is a disjoint union of atoms.

We can describe Pl using a set of defaults. Let $p_1, \ldots, p_n$ be a collection of distinct propositions. For each set $A \in \mathcal{F}$ we define $\phi_A = \bigvee_{A_i \subseteq A} p_i$ if A is nonempty, and define $\phi_A = \mathit{false}$ if A is empty. Let

$$\begin{aligned}\Delta \;=\; & \{(\phi_A \vee \phi_B) \rightarrow \phi_B : A, B \in \mathcal{F}, A \cap B = \emptyset, \mathrm{Pl}(A) < \mathrm{Pl}(B)\} \cup \\ & \{(p_i \wedge p_j) \rightarrow \mathit{false} : i \neq j\} \cup \\ & \{\neg(p_1 \vee \ldots \vee p_n) \rightarrow \mathit{false}\}.\end{aligned}$$

Let π be a truth assignment to W such that $\pi(w)(p_i) = \mathbf{true}$ if and only if $w \in A_i$. Then it is easy to check that $PL = (W, \mathcal{F}, \mathrm{Pl}, \pi)$ satisfies Δ and $PL \not\models (\phi_{A^*} \vee \phi_{B^*}) \rightarrow \phi_{B^*}$. Since Pl is a qualitative plausibility measure by assumption, we have that $PL \in \mathcal{S}^{QPL}$. Using Theorem 5.8, we get that $\Delta \not\vdash_{\mathbf{P}} (\phi_{A^*} \vee \phi_{B^*}) \rightarrow \phi_{B^*}$. By Theorem 3.1, there is a preferential structure $P = (W_P, \prec, \pi_P)$ such that P satisfies Δ but $P \not\models (\phi_{A^*} \vee \phi_{B^*}) \rightarrow \phi_{B^*}$. Applying Lemma A.3 we can assume that P is linear, i.e., $W_P = \{v_1, \ldots, v_m\}$ for some $m \leq n$, and $\prec$ is such that $v_i \prec v_j$ if $i < j$. Moreover, we have that W_P not empty: Since PL is normal, we have that $PL \not\models \mathit{true} \rightarrow \mathit{false}$. Hence, $P \not\models \mathit{true} \rightarrow \mathit{false}$, which implies that there are some minimal worlds that satisfy *true*.

We now construct the required probability measure over W. We start by noting that the defaults in Δ imply that for each world v_i, $\pi_P(v_i)$ assigns **true** to exactly one proposition. Thus, without loss of generality, we can assume that $\pi_P(v_i)(p_j) =$**true** if and only if $i = j$. We define Pr by assigning a probability to each atom. The probability of all other sets is induced from this assignment: $\Pr(A) = \sum_{A_i \subseteq A} \Pr(A_i)$, for all $A \in \mathcal{F}$. Let A_i be an atom. We define $\Pr(A_i) = \alpha \cdot (\frac{1}{x})^i$, if $i \leq m$; otherwise we define $\Pr(A_i) = 0$. The constant α is a normalization constant that ensures that the probability of atoms sum to 1. It is easy to verify that $\alpha = (x-1)\frac{x^m}{x^m - 1}$.

We now show that Pr satisfies the requirements of the lemma. Assume that $A, B \in \mathcal{F}$ are disjoint and $\mathrm{Pl}(A) < \mathrm{Pl}(B)$. We want to show that $\Pr(B|A \cup B) \geq 1 - \frac{1}{x}$. By definition, $(\phi_A \vee \phi_B) \rightarrow \phi_B \in \Delta$, and thus $P \models (\phi_A \vee \phi_B) \rightarrow \phi_B$. There are two cases: either (a) there is no world in W_P that satisfies $(\phi_A \vee \phi_B)$ or (b) the minimal world that satisfies $(\phi_A \vee \phi_B)$ also satisfies ϕ_B. In case (a), it immediately follows that if $A_i \subseteq A \cup B$, then $i > m$. We conclude that $\Pr(A) = \Pr(B) = 0$. Thus, $\Pr(B|A \cup B) = 1$. (Recall that if $\Pr(A \cup B) = 0$ then, by convention,

$\Pr(B|A \cup B) = 1$.) In case (b), assume that v_i is the minimal world satisfying $\phi_A \vee \phi_B$. Since $P \models (\phi_A \vee \phi_B) \rightarrow \phi_B$, it must be the case that v_i satisfies ϕ_B. This implies that $\Pr(B) \geq \Pr(A_i)$. Since v_i is the minimal world that satisfies $\phi_A \vee \phi_B$, and since v_i does not satisfy ϕ_B (since A and B are disjoint) we conclude that $\Pr(A) \leq \sum_{j>i} \Pr(A_j)$. Simple calculation show that $\sum_{j>i} \Pr(A_j) \leq \Pr(A_i) \cdot \frac{1}{x-1}$. Thus, $(x-1)\Pr(A) \leq \Pr(A_i) \leq \Pr(B)$. This implies that $(x-1)(\Pr(A)+\Pr(B)) \leq x\Pr(B)$. Since A and B are disjoint, we have that $\Pr(A \cup B) = \Pr(A) + \Pr(B)$. We get that $(x-1)\Pr(A \cup B) \leq x\Pr(B)$, so $\Pr(B|A \cup B) \geq \frac{x-1}{x}$.

Finally we have to show that $\Pr(B^*|A^* \cup B^*) \leq \frac{1}{2}$. Since $P \not\models (\phi_{A^*} \vee \phi_{B^*}) \rightarrow \phi_{B^*}$, the minimal world, v_i, satisfying $\phi_{A^*} \vee \phi_{B^*}$ does not satisfy ϕ_{B^*}. Since A^* and B^* are disjoint, this implies that v_i satisfies ϕ_{A^*}. The argument above shows that $\Pr(A^*|A^* \cup B^*) \geq 1 - \frac{1}{x}$ and $\Pr(A^* \cup B^*) > 0$. Therefore, $\Pr(B^*|A^* \cup B^*) \leq \frac{1}{x} \leq \frac{1}{2}$. □

LEMMA A.5. *Let $(W, \mathcal{F}, \mathrm{Pl})$ be a normal qualitative plausibility space such that $\mathcal{F}$ is finite. Then there is a PPD $PP = \{\Pr_n : n \geq 1\}$ such that, for all disjoint $A, B \in \mathcal{F}$,*

—*if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$, then $\Pr_n(B|A \cup B) \geq 1 - \frac{1}{n+1}$ for all n,*

—*if $\mathrm{Pl}(A) \not< \mathrm{Pl}(B)$, then for all n, there is an m such that $n \leq m < n + |\mathcal{F}|^2$ and $\Pr_m(B|A \cup B) \leq \frac{1}{2}$.*

PROOF. Let $A_0, \ldots, A_{|\mathcal{F}|-1}$ be some enumeration of the members of $\mathcal{F}$. We define the PPD $PP = \{\Pr_n : n \geq 1\}$ as follows. For each $n \geq 1$ we define i_n, j_n to be the unique integers such that $n = k \cdot |\mathcal{F}|^2 + i_n \cdot |\mathcal{F}| + j_n$ for some positive integer k. According to Lemma A.4 there exists a probability distribution $\Pr_n$ on W such that if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$, then $\Pr_n(B|A \cup B) \geq 1 - \frac{1}{n+1}$, for all disjoint sets $A, B \in F$. Moreover, if $\mathrm{Pl}(A_{i_n}) \not< \mathrm{Pl}(A_{j_n})$, we can ensure that $\Pr_n(A_{j_n}|A_{i_n} \cup A_{j_n}) \leq \frac{1}{2}$. It is easy to verify that this PPD satisfies the requirements of lemma. □

THEOREM 6.3. *If $PL \in \mathcal{S}^{QPL}$ is a normal plausibility structure for a countable language $\mathcal{L}$, then there is a structure $PL' \in \mathcal{S}^{\epsilon}$ that is default-equivalent to PL.*

PROOF. Let $PL = (W, \mathcal{F}, \mathrm{Pl}, \pi)$ be a normal qualitative plausibility for a countable language $\mathcal{L}$. Since $\mathcal{L}$ is countable, so is $\mathcal{F} = \{[\![\phi]\!] : \phi \in \mathcal{L}\}$. Let $A_1, A_2, \ldots$ be an enumeration of the sets in $\mathcal{F}$. Without loss of generality, we can assume that $A_1 = W$. Since Pl is normal, we must have $\mathrm{Pl}(A_1) > \perp$. For each k, let $\mathcal{F}_k$ be the minimal algebra that contains $A_1, \ldots, A_k$. Clearly, $\mathcal{F}_k$ is a finite algebra, and Pl restricted to $\mathcal{F}_k$ is normal. Thus, there is a PPD $PP^k = \{\Pr_1^k, \Pr_2^k, \ldots\}$ that satisfies the conditions of Lemma A.5.

We now use elements of these sequences to construct the desired PPD. We define PP to be the sequence that consists of a segment of PP^1, followed by a segment of PP^2, and so on, such that the length of the segment from PP^k is $|F_k|^2$:

$$\{\Pr_1^1, \ldots, \Pr_{1+|\mathcal{F}_1|^2-1}^1, \Pr_2^2, \ldots, \Pr_{2+|\mathcal{F}_2|^2-1}^2, \ldots, \Pr_k^k, \ldots, \Pr_{k+|\mathcal{F}_k|^2-1}^k, \ldots\}.$$

To show that the plausibility measure that corresponds to (W, PP, π) is default-equivalent to PL we need to show that for all disjoint $A_i, A_j \in \mathcal{F}$, $\mathrm{Pl}(A_i) < \mathrm{Pl}(A_j)$ if and only if $\lim_{n\to\infty} \Pr_n(A_j|A_i \cup A_j) = 1$.

Assume that $\mathrm{Pl}(A_i) < \mathrm{Pl}(A_j)$. Let $\epsilon > 0$. Set $m_\epsilon = \max(i, j, \lceil \frac{1}{\epsilon} \rceil)$, and $N_\epsilon = \sum_{k=1}^{m_\epsilon} |\mathcal{F}_k|^2$. Let $n > N_\epsilon$. Then our construction is such that $\Pr_n = \Pr_l^k$ for some $l \geq k \geq m_\epsilon$. Since $k \geq max(i, j)$, we have that $A_i, A_j \in \mathcal{F}_k$. Moreover, according to Lemma A.5, $\Pr_l^k(A_j | A_i \cup A_j) \geq 1 - \frac{1}{l+1} \geq 1 - \frac{1}{m_\epsilon + 1} \geq 1 - \epsilon$. Thus, we conclude that $\lim_{n\to\infty} \Pr_n(A_j | A_i \cup A_j) = 1$.

Assume that $\mathrm{Pl}(A_i) \not< \mathrm{Pl}(A_j)$. Let $k > \max(i, j)$. According to Lemma A.5, there is an m such that $k \leq m < k + |\mathcal{F}_k|^2$ such that $\Pr_m^k(A_j | A_i \cup A_j) \leq \frac{1}{2}$. Moreover, by the our construction, there is an n such that $\Pr_n = \Pr_m^k$. Thus, for infinitely many n, $\Pr_n(A_j | A_i \cup A_j) \leq \frac{1}{2}$. We conclude that $\lim_{n\to\infty} \Pr_n(A_j | A_i \cup A_j) \neq 1$. □

A.4 Proofs for Section 7

PROPOSITION 7.2. *If E is an expectation structure, then PL_E is a plausibility structure. Furthermore, $E \models \phi \rightarrow \psi$ if and only if $PL_E \models \phi \rightarrow \psi$.*

PROOF. Suppose $E = (\mathcal{L}, \trianglelefteq)$ is an expectation structure, and $PL_E = (W_E, \mathrm{Pl}_E, \pi_E)$. To show that PL_E is a plausibility structure, we need to check that the ordering defined by Pl_E is reflexive, transitive, and satisfies A1. Note that it is easy to show, using standard arguments, that $[\![\phi]\!] \subseteq [\![\psi]\!]$ if and only if $\vdash_{\mathcal{L}} \phi \Rightarrow \psi$. The proof that PL_E is plausibility measure follows in a straightforward manner and is left as an exercise to the reader.

We now show that $E \models \phi \rightarrow \psi$ if and only if $PL_E \models \phi \rightarrow \psi$. Assume that $E \models \phi \rightarrow \psi$. According to Theorem 7.1 there are two possible cases: either (a) $\vdash_{\mathcal{L}} \phi \Rightarrow \psi$ or (b) $(\phi \Rightarrow \neg\psi) \triangleleft (\phi \Rightarrow \psi)$. In case (a), we apply REF and RW (which are valid in all plausibility structures) to get that $PL_E \models \phi \rightarrow \psi$. Now consider case (b). It is clear that $\mathit{true} \not\trianglelefteq (\phi \Rightarrow \neg\psi)$, for if $\mathit{true} \trianglelefteq (\phi \Rightarrow \neg\psi)$, then $\mathit{true} \triangleleft (\phi \Rightarrow \psi)$ by E1. We also have $(\phi \Rightarrow \psi) \trianglelefteq \mathit{true}$ by using E2, so by E1 again, we get that $\mathit{true} \triangleleft \mathit{true}$, which is a contradiction. Since $\mathit{true} \not\trianglelefteq (\phi \Rightarrow \neg\psi)$ and $(\phi \Rightarrow \neg\psi) \triangleleft (\phi \Rightarrow \psi)$, by definition, we have that $\mathrm{Pl}_E([\![\phi \wedge \neg\psi]\!]) \leq \mathrm{Pl}_E([\![\phi \wedge \psi]\!])$. Moreover, since $\mathit{true} \not\trianglelefteq (\phi \Rightarrow \neg\psi)$, we have that $\phi \wedge \psi$ is consistent. Thus, $\not\vdash_{\mathcal{L}} \phi \wedge \psi \Rightarrow \phi \wedge \neg\psi$. Furthermore, since $(\phi \Rightarrow \neg\psi) \triangleleft (\phi \Rightarrow \psi)$, we also have that $(\phi \Rightarrow \psi) \not\trianglelefteq (\phi \Rightarrow \neg\psi)$. This implies, by definition, that $\mathrm{Pl}_E([\![\phi \wedge \neg\psi]\!]) \not\geq \mathrm{Pl}_E([\![\phi \wedge \psi]\!])$. Hence, $\mathrm{Pl}_E([\![\phi \wedge \neg\psi]\!]) < \mathrm{Pl}_E([\![\phi \wedge \psi]\!])$, and so $PL_E \models \phi \rightarrow \psi$.

Now assume that $PL_E \models \phi \rightarrow \psi$. If $\vdash_{\mathcal{L}} \phi \Rightarrow \psi$, then $E \models \phi \rightarrow \psi$. Assume that $\not\vdash_{\mathcal{L}} \phi \Rightarrow \psi$. This implies that $\phi \wedge \neg\psi$ is consistent in $\mathcal{L}$. We claim that $\mathrm{Pl}_E([\![\phi]\!]) > \bot$. To see this, assume that $\mathrm{Pl}_E([\![\phi]\!] \leq \mathrm{Pl}([\![\mathit{false}]\!])$. Since $\mathit{true} \trianglelefteq \neg\mathit{false}$, the definition of Pl_E implies that $\vdash_{\mathcal{L}} \phi \Rightarrow \mathit{false}$. But this contradict the assumption that ϕ is consistent in $\mathcal{L}$. We conclude that $\mathrm{Pl}_E([\![\phi]\!]) > \bot$, and thus since $PL_e \models \phi \rightarrow \psi$, $\mathrm{Pl}_E([\![\phi \wedge \psi]\!]) > \mathrm{Pl}_E([\![\phi \wedge \neg\psi]\!])$. It is straightforward to show that since $\phi \wedge \neg\psi$ is consistent in $\mathcal{L}$, we get that $\neg(\phi \wedge \psi) \triangleleft \neg(\phi \wedge \neg\psi)$. Rewriting this equation, we get that $(\phi \Rightarrow \neg\psi) \triangleleft (\phi \Rightarrow \psi)$, and thus $E \models \phi \rightarrow \psi$. □

A.5 Proofs for Section 8

We now want to prove Theorem 8.2. For the proof, it is useful to define $N\phi$ as an abbreviation for $\neg\phi \rightarrow \mathit{false}$. (The N operator is called the *outer modality* in [Lewis 1973].) Expanding the definition of $\rightarrow$, we get that $N\phi$ holds at w if and only if $\mathrm{Pl}([\![\neg\phi]\!]) = \bot$. Thus, $N\phi$ holds if $\neg\phi$ is considered completely implausible. Thus, it implies that ϕ is true "almost everywhere". The following lemma collects some properties of N that will be needed in the proof.

LEMMA A.6. *(a)* $\vdash_{\mathbf{C}} N(\phi \wedge \psi) \Rightarrow N\phi$.
(b) $\vdash_{\mathbf{C}} (N\phi \wedge N\psi) \Rightarrow N(\phi \wedge \psi)$.
(c) $\vdash_{\mathbf{C}} (N\phi \wedge N(\phi \Rightarrow \psi)) \Rightarrow N\psi$.
(d) *If* $\vdash_{\mathbf{C}} \phi$ *then* $\vdash_{\mathbf{C}} N\phi$.
(e) $\vdash_{\mathbf{C}} N\phi \Rightarrow (\psi \rightarrow \phi)$.
(f) $\vdash_{\mathbf{C}} (N(\phi \Leftrightarrow \phi') \wedge N(\psi \Leftrightarrow \psi')) \Rightarrow ((\phi \rightarrow \psi) \Leftrightarrow (\phi' \rightarrow \psi'))$.

PROOF. Recall that $N\phi$ is defined as $\neg\phi \rightarrow \mathit{false}$.
We start with part (a).

1. $\vdash_{\mathbf{C}} (\neg(\phi \wedge \psi) \rightarrow \mathit{false}) \Rightarrow (\neg(\phi \wedge \psi) \rightarrow \neg\phi)$	RC2
2. $\vdash_{\mathbf{C}} (\neg(\phi \wedge \psi) \rightarrow \mathit{false}) \Rightarrow ((\neg(\phi \wedge \psi) \wedge \neg\phi) \rightarrow \mathit{false})$	1, C4
3. $\vdash_{\mathbf{C}} (\neg(\phi \wedge \psi) \rightarrow \mathit{false}) \Rightarrow (\neg\phi \rightarrow \mathit{false})$	2, RC1
4. $\vdash_{\mathbf{C}} N(\phi \wedge \psi) \Rightarrow N\phi$	3 rewritten

To prove part (b), we proceed as follows:

1. $\vdash_{\mathbf{C}} ((\neg\phi \rightarrow \mathit{false}) \wedge (\neg\psi \rightarrow \mathit{false})) \Rightarrow ((\neg\phi \vee \neg\psi) \rightarrow \mathit{false})$	C3
2. $\vdash_{\mathbf{C}} ((\neg\phi \rightarrow \mathit{false}) \wedge (\neg\psi \rightarrow \mathit{false})) \Rightarrow (\neg(\phi \wedge \psi) \rightarrow \mathit{false})$	1, RC1
3. $\vdash_{\mathbf{C}} (N\phi \wedge N\psi) \Rightarrow N(\phi \wedge \psi)$	2 rewritten

For part (c), we proceed as follows:

1. $\vdash_{\mathbf{C}} (N\phi \wedge N(\phi \Rightarrow \psi)) \Rightarrow N(\phi \wedge (\phi \Rightarrow \psi))$	(b)
2. $\vdash_{\mathbf{C}} (N\phi \wedge N(\phi \Rightarrow \psi)) \Rightarrow N(\phi \wedge \psi)$	1, RC1
3. $\vdash_{\mathbf{C}} (N\phi \wedge N(\phi \Rightarrow \psi)) \Rightarrow N\psi$	(a), 2

To prove part (d), we assume that $\vdash_{\mathbf{C}} \phi$.

1. $\vdash_{\mathbf{C}} \phi$	assumption
2. $\vdash_{\mathbf{C}} \mathit{true} \Leftrightarrow \phi$	1, C0
3. $\vdash_{\mathbf{C}} \mathit{false} \rightarrow \mathit{false}$	C1
4. $\vdash_{\mathbf{C}} \neg\phi \rightarrow \mathit{false}$	3, RC1

To prove part (e), we proceed as follows:

1. $\vdash_{\mathbf{C}} (\psi \wedge \phi) \rightarrow \phi$	C1, RC2
2. $\vdash_{\mathbf{C}} (\neg\phi \rightarrow \mathit{false}) \Rightarrow (\neg\phi \rightarrow \phi) \wedge (\neg\phi \rightarrow \psi)$	RC2
3. $((\neg\phi \rightarrow \phi) \wedge (\neg\phi \rightarrow \psi)) \Rightarrow ((\psi \wedge \neg\phi) \rightarrow \phi$	C4
4. $\vdash_{\mathbf{C}} (((\psi \wedge \phi) \rightarrow \phi) \wedge ((\psi \wedge \neg\phi) \rightarrow \phi))) \Rightarrow (\psi \rightarrow \phi)$	C3, RC1
5. $\vdash_{\mathbf{C}} N\phi \Rightarrow (\psi \rightarrow \phi)$	1, 2, 3, 4

Finally, we prove part (f).

1. $\vdash_{\mathbf{C}} N(\phi \Leftrightarrow \phi') \Rightarrow (N(\phi' \Rightarrow \phi) \wedge N(\phi \Rightarrow \phi'))$	(a)
2. $\vdash_{\mathbf{C}} N(\phi' \Rightarrow \phi) \Rightarrow ((\phi' \wedge \neg\phi) \rightarrow \psi)$	definition of N, RC2

3. $\vdash_{\mathbf{C}} N(\phi \Rightarrow \phi') \Rightarrow (\phi \rightarrow (\phi \Rightarrow \phi'))$ (e)
4. $\vdash_{\mathbf{C}} N(\phi \Rightarrow \phi') \Rightarrow (\phi \rightarrow \phi')$ 3, C1, C2, RC2
5. $\vdash_{\mathbf{C}} (\phi \rightarrow \phi') \land (\phi \rightarrow \psi) \Rightarrow ((\phi' \land \phi) \rightarrow \psi)$ C4
6. $\vdash_{\mathbf{C}} N(\phi \Leftrightarrow \phi') \land (\phi \rightarrow \psi) \Rightarrow (\phi' \rightarrow \psi)$ 1, 2, 3, 4, 5, C3, RC1
7. $\vdash_{\mathbf{C}} N(\psi \Leftrightarrow \psi') \Rightarrow (\phi' \rightarrow (\psi \Rightarrow \psi'))$ (e), RC2
8. $\vdash_{\mathbf{C}} N(\phi \Leftrightarrow \phi') \land N(\psi \Leftrightarrow \psi') \land (\phi \rightarrow \psi) \Rightarrow (\phi' \rightarrow \psi')$ 6, 7, C2, RC2

□

THEOREM 8.2. *System* **C** *is a sound and complete axiomatization of* $\mathcal{L}^C$ *with respect to* $\mathcal{S}_c^{QPL}$.

PROOF. It is easy to verify that System **C** is sound in $\mathcal{S}_c^{QPL}$. To prove completeness, we have to show that if $\models_{\mathcal{S}_c^{QPL}} \phi$, then $\vdash_{\mathbf{C}} \phi$. This is equivalent to showing that if $\not\vdash_{\mathbf{C}} \phi$ (i.e., $\neg\phi$ is consistent) then $\not\models_{\mathcal{S}_c^{QPL}} \phi$ (i.e., $\neg\phi$ is satisfiable).

We construct a *canonical* qualitative plausibility structure PL such that for all $\xi \in \mathcal{L}^C$ we have that if ξ is consistent, then $(PL, w) \models \xi$ for some world w, using standard techniques. Recall that a set of formulas $V \subseteq \mathcal{L}^C$ is a *maximal* $\vdash_{\mathbf{C}}$-*consistent* set if it is consistent with respect to $\vdash_{\mathbf{C}}$ and for each $\phi \in \mathcal{L}^C$, either $\phi \in V$ or $\neg\phi \in V$. Let V be a set of formulas. We define $V/N = \{\phi : N\phi \in V\}$.

Define $PL = (W, \mathcal{P}, \pi)$ as follows:

—$W = \{w_V : V \subseteq \mathcal{L}^C$ is a maximal $\vdash_{\mathbf{C}}$-consistent set of formulas $\}$,
—$\mathcal{P}(w_V) = (W_w, \mathcal{F}, \mathrm{Pl}_{w_V})$ where
 —$W_{w_V} = \{w_U : V/N \subseteq U\}$,
 —$\mathcal{F}_{w_V} = \{[\phi]_{w_V} : \phi \in \mathcal{L}^C\}$, where $[\phi]_{w_V} = \{w_U \in W_{w_V} : \phi \in U\}$,
 —Pl_{w_V} is such that $\mathrm{Pl}_{w_V}([\phi]_{w_V}) \le \mathrm{Pl}_{w_V}([\psi]_{w_V})$ if and only if $(\phi \lor \psi) \rightarrow \psi \in V$,
—$\pi(w_V)(p) = \mathbf{true}$ if and only if $p \in V$.

We want to show that PL is a qualitative plausibility structure and that $(PL, w_V) \models \phi$ if and only if $\phi \in V$ for all formulas ϕ and worlds w_V.

We first need to establish that Pl_{w_V} is well-defined, for all $w_V \in W$. Let V be a maximal $\vdash_{\mathbf{C}}$-consistent set. We need to show that if $[\phi]_{w_V} = [\phi']_{w_V}$ and $[\psi]_{w_V} = [\psi']_{w_V}$, then $(\phi \lor \psi) \rightarrow \psi \in V$ if and only $(\phi' \lor \psi') \rightarrow \psi' \in V$.

We claim that $[\phi]_{w_V} = [\phi']_{w_V}$ if and only if $V/N \cup \{\neg(\phi \Leftrightarrow \phi')\}$ is inconsistent. The "if" direction is obvious. For the "only if" direction, assume that $V/N \cup \{\neg(\phi \Leftrightarrow \phi')\}$ is consistent. This implies that there is a maximal consistent set U such that $\Delta \cup \{\neg(\phi \Leftrightarrow \phi')\} \subseteq U$. Clearly, w_U is in the symmetric difference between $[\phi]_{w_V}$ and $[\phi']_{w_V}$. Thus, $[\phi]_{w_V} \neq [\phi']_{w_V}$.

Now assume that $[\phi]_{w_V} = [\phi']_{w_V}$ and $[\psi]_{w_V} = [\psi']_{w_V}$. Therefore, as we have just shown, $V/N \cup \{\neg(\phi \Leftrightarrow \phi')\}$ and $V/N \cup \{\neg(\psi \Leftrightarrow \psi)\}$ are both inconsistent. Thus, there exists a formula δ which is the conjunction of a finite number of formulas in V/N such that $\vdash_{\mathbf{C}} \delta \Rightarrow (\phi \Leftrightarrow \phi') \land (\psi \Leftrightarrow \psi)$. Using part (b) of Lemma A.6, and the fact that if $\phi \in V/N$, then $N\phi \in V$, we get that $N\delta \in V$. From parts (c) and (d) of Lemma A.6, we get that $\vdash_{\mathbf{C}} N\delta \Rightarrow N(\phi \Leftrightarrow \phi') \land N(\psi \Leftrightarrow \psi')$. Finally, using part (f) of Lemma A.6 we get that

$$\vdash_{\mathbf{C}} N\delta \Rightarrow ((\phi \rightarrow \psi) \Leftrightarrow (\phi' \rightarrow \psi')).$$

Thus, $\phi \rightarrow \psi \in V$ if and only if $\phi' \rightarrow \psi' \in V$. This suffices to prove that Pl_{w_V} is well-defined.

To see that PL is a qualitative plausibility structure, first note that the definition of Pl_{w_V} mirrors the construction in Lemma 4.1. It is easy to prove that the set of defaults $\{\phi \rightarrow \psi \in V\}$ is closed under the KLM rules. Thus, we can immediately use the proof of Lemma 4.1 to show that $(W_{w_V}, \mathrm{Pl}_{w_V})$ is a plausibility space. The proof that PL satisfies A2 and A3 is identical to the proof of Theorem 5.4. Thus, $PL \in \mathcal{S}_c^{QPL}$.

We next show that $(PL, w_V) \models \phi$ if and only if $\phi \in V$. This is done by induction on the structure of ϕ. The only case of interest is if ϕ is of the form $\phi' \rightarrow \psi$. Here again the proof is identical to the proof of Lemma 4.1. The only difference is the use of axioms in system **C** instead of the corresponding rules in system **P**.

Let ξ be a $\vdash_{\mathbf{C}}$-consistent formula. Using standard arguments, it is easy to show that there is some maximal $\vdash_{\mathbf{C}}$-consistent set V_ξ such that $\xi \in V_\xi$. Thus, $(PL, w_{V_\xi}) \models \xi$, so ξ is satisfiable in $\mathcal{S}_c^{QPL}$. $\square$

PROPOSITION 8.3. *Let $\mathcal{P} \subseteq \mathcal{S}_c^{QPL}$. C5 is valid in $\mathcal{S}$ if and only if all structures in $\mathcal{S}$ are rational.*

PROOF. To prove the "if" direction it suffices to show that C5 is sound in rational qualitative structures. Let $PL = (W, \mathcal{P}, \pi)$ be a rational qualitative plausibility structure. Assume that $(PL, w) \models (\phi \rightarrow \psi)$ and $(PL, w) \models \neg(\phi \wedge \xi \rightarrow \psi)$. We need to prove that $(PL, w) \models \phi \rightarrow \neg\xi$. If $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) = \perp$, then $(PL, w) \models \phi \rightarrow \neg\xi$ vacuously, and we are done. Now assume that $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) > \perp$. Let $A = [\![\phi]\!]_{(PL,w)}$, $B = [\![\psi]\!]_{(PL,w)}$, and $C = [\![\xi]\!]_{(PL,w)}$. Since $(PL, w) \models \phi \rightarrow \psi$, we have that $\mathrm{Pl}_w(A \cap B) > \mathrm{Pl}_w(A \cap \overline{B})$. Since Pl_w satisfies A5, we have that either $\mathrm{Pl}_w(A \cap B \cap C) > \mathrm{Pl}_w(A \cap \overline{B})$ or $\mathrm{Pl}_w(A \cap B \cap \overline{C}) > \mathrm{Pl}_w(A \cap \overline{B})$. However, since $(PL, w) \models \neg(\phi \wedge \xi \rightarrow \psi)$, we have that $\mathrm{Pl}_w(A \cap B \cap C) \not> \mathrm{Pl}_w(A \cap \overline{B} \cap C)$. This implies that

$$\mathrm{Pl}_w(A \cap B \cap C) \not> \mathrm{Pl}_w(A \cap \overline{B}). \tag{5}$$

Thus, we conclude that

$$\mathrm{Pl}_w(A \cap B \cap \overline{C}) > \mathrm{Pl}_w(A \cap \overline{B}). \tag{6}$$

Applying A4 with (6) as the antecedent, we get that either $\mathrm{Pl}_w(A \cap B \cap C) > \mathrm{Pl}_w(A \cap \overline{B})$ or $\mathrm{Pl}_w(A \cap B \cap \overline{C}) > \mathrm{Pl}_w(A \cap B \cap C)$. Since the former contradicts (5), we conclude that $\mathrm{Pl}_w(A \cap B \cap \overline{C}) > \mathrm{Pl}_w(A \cap B \cap C)$. Using A1, A2, and (6), we get that $\mathrm{Pl}_w(A \cap \overline{C}) \geq \mathrm{Pl}_w(A \cap B \cap \overline{C}) > \mathrm{Pl}_w((A \cap B \cap C) \cup (A \cap \overline{B})) \geq \mathrm{Pl}_w(A \cap C)$. Thus, $(PL, w) \models \phi \rightarrow \neg\xi$.

To prove the "only if" direction, we have to show that if there is some $PL = (W, \mathcal{P}, \pi)$ in $\mathcal{S}$ that is not rational, then C5 is violated.

Suppose that PL does not satisfy A4. Since we have assumed that $\mathcal{F}_w = \{[\![\phi]\!]_{(PL,w)} : \phi \in \mathcal{L}\}$, there is a world w, and formulas ϕ, ψ, and ξ such that $[\![\phi]\!]_{(PL,w)}$, $[\![\psi]\!]_{(PL,w)}$, and $[\![\xi]\!]_{(PL,w)}$ are pairwise disjoint, $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) > \mathrm{Pl}_w([\![\psi]\!]_{(PL,w)})$, and yet $\mathrm{Pl}_w([\![\xi]\!]_{(PL,w)}) \not> \mathrm{Pl}_w([\![\psi]\!]_{(PL,w)})$ and $\mathrm{Pl}_w([\![\xi]\!]_{(PL,w)}) \not< \mathrm{Pl}_w([\![\phi]\!]_{(PL,w)})$. Since $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) > \mathrm{Pl}_w([\![\psi]\!]_{(PL,w)})$, we have that $\mathrm{Pl}_w([\![\phi \vee \xi]\!]_{(PL,w)}) > \mathrm{Pl}_w([\![\psi]\!]_{(PL,w)})$. Thus, $(PL, w) \models (\phi \vee \psi \vee \xi) \rightarrow \neg\psi$. Moreover, since $\mathrm{Pl}_w([\![\xi]\!]_{(PL,w)}) \not< \mathrm{Pl}_w([\![\phi]\!]_{(PL,w)})$, we have that $\mathrm{Pl}_w([\![\xi]\!]_{(PL,w)}) > \perp$. Since we also have that $\mathrm{Pl}_w([\![\xi]\!]_{(PL,w)}) \not>

$\mathrm{Pl}_w([\![\psi]\!]_{(PL,w)})$, we conclude that $(PL, w) \models \neg((\xi\vee\psi)\rightarrow\neg\psi)$. Finally, since $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) \not> \mathrm{Pl}_w([\![\xi]\!]_{(PL,w)})$, we have that $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) \not> \mathrm{Pl}_w([\![\xi\vee\psi]\!]_{(PL,w)})$. Thus, $(PL, w) \models \neg((\phi\vee\psi\vee\xi)\rightarrow\neg(\psi\vee\xi))$. Define α as $\phi\vee\psi\vee\xi$, β as $\psi\vee\xi$ and γ as $\neg\psi$. We have just proved that $(PL, w) \models (\alpha\rightarrow\gamma)\wedge\neg(\alpha\rightarrow\neg\beta)\wedge\neg(\alpha\wedge\beta\rightarrow\gamma)$, which contradicts C5.

Now suppose that PL does not satisfy A5. Then there is a world w, and formulas ϕ, ψ, and ξ, such that $[\![\phi]\!]_{(PL,w)}$, $[\![\psi]\!]_{(PL,w)}$, and $[\![\xi]\!]_{(PL,w)}$ are pairwise disjoint, $\mathrm{Pl}_w([\![\phi\vee\psi]\!]_{(PL,w)}) > \mathrm{Pl}_w([\![\xi]\!]_{(PL,w)})$, and yet $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) \not> \mathrm{Pl}_w([\![\xi]\!]_{(PL,w)})$ and $\mathrm{Pl}_w([\![\psi]\!]_{(PL,w)}) \not> \mathrm{Pl}_w([\![\xi]\!]_{(PL,w)})$. Since $\mathrm{Pl}_w([\![\phi\vee\psi]\!]_{(PL,w)}) > \mathrm{Pl}_w([\![\xi]\!]_{(PL,w)}) \geq \perp$, either $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) > \perp$ or $\mathrm{Pl}_w([\![\psi]\!]_{(PL,w)}) > \perp$, for otherwise, using A3 we would get that $\mathrm{Pl}_w([\![\phi\vee\psi]\!]_{(PL,w)}) = \perp$. Without loss of generality, we assume that $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) > \perp$. Since $\mathrm{Pl}_w([\![\phi\vee\psi]\!]_{(PL,w)}) > \mathrm{Pl}_w([\![\xi]\!]_{(PL,w)})$, we have that $(PL, w) \models (\phi\vee\psi\vee\xi)\rightarrow\neg\xi$. Since $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) > \perp$ and $\mathrm{Pl}_w([\![\phi]\!]_{(PL,w)}) \not> \mathrm{Pl}_w([\![\xi]\!]_{(PL,w)})$, we have that $(PL, w) \models \neg((\phi\vee\xi)\rightarrow\neg\xi)$. Finally, since $\mathrm{Pl}_w([\![\psi]\!]_{(PL,w)}) \not> \mathrm{Pl}_w([\![\xi]\!]_{(PL,w)})$, we have that $\mathrm{Pl}_w([\![\psi]\!]_{(PL,w)}) \not> \mathrm{Pl}_w([\![\phi\vee\xi]\!]_{(PL,w)})$, and thus $(PL, w) \models \neg((\phi\vee\psi\vee\xi)\rightarrow\psi))$. Define α as $\phi\vee\psi\vee\xi$, β as $\neg\psi$ and γ as $\neg\xi$. We have proved that $(PL, w) \models (\alpha\rightarrow\gamma)\wedge\neg(\alpha\rightarrow\neg\beta)\neg(\alpha\wedge\beta\rightarrow\gamma)$, which contradicts C5. $\square$

We next want to prove Theorem 8.4. We first need a lemma.

LEMMA A.7. *Let (W, Pl) be a qualitative plausibility space, and let $<^*$ be a binary relation on subsets of W such that $A <^* B$ if there is some set $C \subseteq W$ such that $A \cap C = \emptyset$ and $\mathrm{Pl}(A) < \mathrm{Pl}(C) \leq \mathrm{Pl}(B)$. Then*

(a) $<^*$ is a strict order; that is irreflexive, transitive and anti-symmetric,

(b) if $A \cap B = \emptyset$, then $A <^* B$ if and only if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$,

(c) if Pl is rational, then $<^*$ is modular, and

(d) if Pl is rational and $A <^* (A \cup B)$, then $B \not<^* (A \cup B)$.

PROOF. We start with part (a). The definition of $<^*$ implies that $A <^* B$ only if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$. Thus, we get that $<^*$ is irreflexive and anti-symmetric. We now show that $<^*$ is transitive. Assume that $A <^* B$ and $B <^* C$. Since $A <^* B$, there is a set D such that $A \cap D = \emptyset$, and $\mathrm{Pl}(A) < \mathrm{Pl}(D) \leq \mathrm{Pl}(B)$. Moreover, since $B <^* C$, we have that $\mathrm{Pl}(B) < \mathrm{Pl}(C)$. Thus, we get that $\mathrm{Pl}(D) < \mathrm{Pl}(C)$, and hence $A <^* C$.

For part (b), let A and B be disjoint sets. If $\mathrm{Pl}(A) < \mathrm{Pl}(B)$, then $A <^* B$ using the set $C = B$. On the other hand, if $A <^* B$, then $\mathrm{Pl}(A) < \mathrm{Pl}(B)$.

For part (c), assume that Pl is rational, and that $A <^* B$. Let $C \subseteq W$. We have to show that either $A <^* C$ or $C <^* B$. Since $A <^* B$, there is a set D such that $A \cap D = \emptyset$ and $\mathrm{Pl}(A) < \mathrm{Pl}(D) \leq \mathrm{Pl}(B)$. Since D is the disjoint union of $(D \cap C)$ and $D - C$, we can apply A5 and get that either $\mathrm{Pl}(A) < \mathrm{Pl}(D \cap C)$ or $\mathrm{Pl}(A) < \mathrm{Pl}(D - C)$. If $\mathrm{Pl}(A) < \mathrm{Pl}(D \cap C)$, then since $\mathrm{Pl}(D \cap C) \leq \mathrm{Pl}(C)$, we get that $A <^* C$ and we are done. If $\mathrm{Pl}(A) < \mathrm{Pl}(D - C)$, since A, $D - C$ and $C - A$ are pairwise disjoint, we get from A4 that either $\mathrm{Pl}(A) < \mathrm{Pl}(C - A)$ or $\mathrm{Pl}(C-A) < \mathrm{Pl}(D-C)$. If $\mathrm{Pl}(A) < \mathrm{Pl}(C-A)$, we again get that $A <^* C$ and we are done. If $\mathrm{Pl}(C-A) < \mathrm{Pl}(D-C)$, using A2, we get that $\mathrm{Pl}((C-A)\cup A) < \mathrm{Pl}(D-C)$. We conclude that $\mathrm{Pl}(C) < \mathrm{Pl}(D - C)$. Since $\mathrm{Pl}(D - C) \leq \mathrm{Pl}(D) \leq \mathrm{Pl}(B)$, we get that $C <^* B$.

Finally, we prove part (d). Assume that $A <^* (A \cup B)$. Without loss of generality, we can also assume that $B \cap A = \emptyset$ (if not, replace B by $B - A$). We want to show that $B \not<^* (A \cup B)$. By way of contradiction, assume that $B <^* (A \cup B)$. Since $A <^* B$, there is a set C such that $A \cap C = \emptyset$ and $\mathrm{Pl}(A) < \mathrm{Pl}(C) \leq \mathrm{Pl}(A \cup B)$. Since $C \subseteq (C - B) \cup B$, this implies that $\mathrm{Pl}(A) < \mathrm{Pl}((C - B) \cup B)$. Since A, B, and $C - B$ are pairwise disjoint, we can apply A5 to get that either $\mathrm{Pl}(A) < \mathrm{Pl}(B)$ or $\mathrm{Pl}(A) < \mathrm{Pl}(C - B)$. If $\mathrm{Pl}(A) < \mathrm{Pl}(B)$, since we assumed that $B <^* (A \cup B)$, there is a set D such that $D \cap B = \emptyset$ and $\mathrm{Pl}(B) < \mathrm{Pl}(D) \leq \mathrm{Pl}(A \cup B)$. This implies that $\mathrm{Pl}(B) < \mathrm{Pl}((D - A) \cup B)$. Moreover, since $\mathrm{Pl}(A) < \mathrm{Pl}(B)$, we also have that $\mathrm{Pl}(A) < \mathrm{Pl}((D-A) \cup B)$. Since A, B and $D - A$ are pairwise disjoint, we can apply A2 to get that $\mathrm{Pl}(A \cup B) < \mathrm{Pl}(D - A) \leq \mathrm{Pl}(D)$. But this contradicts the assumption that $\mathrm{Pl}(D) \leq \mathrm{Pl}(A \cup B)$. If $\mathrm{Pl}(A) < \mathrm{Pl}(C - B)$, since A, B, and $C - B$ are pairwise disjoint, we have by A4 that either $\mathrm{Pl}(A) < \mathrm{Pl}(B)$ or $\mathrm{Pl}(B) < \mathrm{Pl}(C - B)$. We have already seen that if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$ we get a contradiction. Now assume that $\mathrm{Pl}(B) < \mathrm{Pl}(C - B)$. Then, since we also have that $\mathrm{Pl}(A) < \mathrm{Pl}(C - B)$, we can apply A2 to get that $\mathrm{Pl}(A \cup B) < \mathrm{Pl}(C - B) \leq \mathrm{Pl}(C)$. But this contradicts our assumption that $\mathrm{Pl}(C) \leq \mathrm{Pl}(A \cup B)$. Thus, we must have $B \not<^* (A \cup B)$. □

THEOREM 8.4. *If (W, Pl) be a rational qualitative plausibility space, then there is a default-equivalent plausibility space (W, Pl') such that Pl' is a ranking.*

PROOF. Let $<^*$ be the relation defined in Lemma A.7. Define a relation $\approx^*$ on sets such that $A \approx^* B$ if neither $A <^* B$ nor $B <^* A$. Since $<^*$ is modular, standard (and straightforward) arguments show that $\approx^*$ is an equivalence relation. We construct a new plausibility measure based on these two ordering. Let $\mathcal{F}/\approx^*$ be the set of equivalence classes of measurable sets. Let $\leq^*$ be the the total order on $\mathcal{F}/\approx^*$ induced by $<^*$. Let Pl' be a plausibility measure on $\mathcal{F}$ whose range is $\mathcal{F}/\approx^*$, defined so that $\mathrm{Pl}'(A) = [A]$ where $[A]$ is the equivalence class containing A. We have to show that Pl' is a plausibility measure. Assume that $A \subseteq B$. Then $\mathrm{Pl}(A) \leq \mathrm{Pl}(B)$ and clearly $B \not<^* A$. Thus, either $A <^* B$ or $A \approx^* B$. We conclude that $\mathrm{Pl}'(A) \leq \mathrm{Pl}'(B)$, as desired. Since $\leq^*$ is a total order, Pl' satisfies A4$'$. Using Lemma A.7(d), we get that $\mathrm{Pl}'(A \cup B) = \max(\mathrm{Pl}'(A), \mathrm{Pl}'(B))$. Thus, Pl' also satisfies A5$'$, and hence is a ranking. Finally, we have to show that (W, Pl') is default-equivalent to (W, Pl). Let A and B be disjoint events. By Lemma A.7, $A <^* B$ if and only if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$. Thus, we conclude that $\mathrm{Pl}'(A) < \mathrm{Pl}'(B)$ if and only if $\mathrm{Pl}(A) < \mathrm{Pl}(B)$. □

We now prove Theorem 8.4. We actually prove some more general results. Consider the following two restrictions on plausibility structures:

R. Pl_w is rational for all worlds w.

N. Pl_w in normal for all worlds w.

THEOREM A.8. *Let $\phi \in \mathcal{L}^C$, let $\mathcal{A}$ be a subset of $\{R, N\}$, and let A be the corresponding subset of $\{C5, C6\}$. Then ϕ is valid in subset of $\mathcal{S}_c^{QPL}$ that satisfies $\mathcal{A}$ if and only if ϕ is provable from system* **C**+*A*.

PROOF. Again, we focus on completeness. We obtain completeness in each case by modifying the proof of Theorem 8.2. We construct a canonical model as in that proof, checking consistency with the extended axiom system. The resulting structure has the property that $(PL, w_V) \models \phi$ if and only if $\phi \in V$. We just need to show that this structure also satisfies the corresponding semantic restrictions.

If C5 is included as an axiom, then C5 is valid in PL. Using Proposition 8.3 we get that PL is rational.

If C6 is included as an axiom and PL is not normal, then there is some world w_V where $\mathrm{Pl}_{w_V}([\![true]\!]_{(PL,w_V)}) = \mathrm{Pl}_{w_V}([\![false]\!]_{(PL,w_V)})$. We get that $(PL, w_V) \models true \rightarrow false$, which contradicts C6. $\Box$

Lemma A.9. *Let $\phi \in \mathcal{L}^C$ and let $\mathcal{A}$ be a subset of $\{R, N\}$. If ϕ is satisfiable in a structure satisfying $\mathcal{A}$, it is satisfiable in a finite structure satisfying $\mathcal{A}$.*

Proof. Let $\phi \in \mathcal{L}^C$ and let $PL = (W, \mathcal{P}, \pi)$ be a structure satisfying $\mathcal{A}$ such that $(PL, w) \models \phi$ for some $w \in W$. We now construct a finite structure PL' that satisfies ϕ. The key idea is that when evaluating ϕ we only examine subformulas of ϕ. Thus, we are interested in distinguishing only between worlds that differ in the evaluation of some subformula of ϕ.

We now make this argument precise. Let $Sub(\phi)$ be the set of subformulas of ϕ. We partition W into equivalence classes: For $w \in W$, define $[w] = \{w' \in W : \forall \psi \in Sub(\phi), (PL, w) \models \psi$ if and only if $(PL, w') \models \psi\}$. Thus, $[w]$ contains all the worlds that are indistinguishable, for the purpose of evaluating ϕ, from w. We now choose, arbitrarily, a *representative* world $\hat{w} \in [w]$ for each equivalence class. These definitions extend to sets of worlds: For $A \subseteq W$, define $[A] = \cup_{w \in A}[w]$, and $\hat{A} = \{\hat{w} : w \in A\}$.

We construct PL' as follows. We set $W' = \hat{W}$. Clearly W' is finite, since there are only a finite number of subformulas of ϕ. Let π' be the restriction of π to W', and let $\mathcal{P}'(\hat{w})(W'_{\hat{w}}, \mathrm{Pl}'_{\hat{w}})$, where $W'_{\hat{w}} = \hat{W}_{\hat{w}}$, and $\mathrm{Pl}'_{\hat{w}}$ is defined so that $\mathrm{Pl}'_{\hat{w}}(A') \leq \mathrm{Pl}_{\hat{w}}(B')$ if $\mathrm{Pl}_{\hat{w}}([A'] \cap W_{\hat{w}}) \leq \mathrm{Pl}_{\hat{w}}([B'] \cap W_{\hat{w}})$.

We need show that PL' is a qualitative plausibility structure. This is easy to prove since $[\cdot]$ preserves subsets, unions, and disjointness of sets: if $A' \subseteq B'$, then $[A'] \subseteq [B']$; $[A' \cup B'] = [A'] \cup [B']$; and $[A' \cap B'] = [A'] \cap [B']$.

Next, we need to show that if $(PL, w) \models \phi$, then $(PL', \hat{w}) \models \phi$. In fact, we prove this property for every subformula ψ of ϕ. The proof is by induction on the structure of of ψ. The only case of interest is for conditional formulas $\psi \rightarrow \xi \in Sub(\phi)$. This case follows easily once we notice that if the induction hypothesis holds for ψ, then $[[\![\psi]\!]_{(PL',\hat{w})}] \cap W_{\hat{w}} = [\![\psi]\!]_{(PL,\hat{w})} \cap W_{\hat{w}}$.

Finally, we have to show that if PL is rational or normal, then so is PL'. Again, this follows easily from the properties of $[\cdot]$. $\Box$

Lemma A.10. *Let $(W, \mathcal{F}, \mathrm{Pl})$ be a rational and normal qualitative plausibility space such that $\mathcal{F}$ is finite. Then there is a κ-ranking κ on W and a possibility measure Poss on W such that for all disjoint $A, B \in \mathcal{F}$, $\mathrm{Pl}(A) > \mathrm{Pl}(B)$ if and only if $\kappa(A) < \kappa(B)$ and $\mathrm{Pl}(A) > \mathrm{Pl}(B)$ if and only if $Poss(A) > Poss(B)$.*

Proof. By Theorem 8.4, there is a ranked plausibility space $(W, \mathcal{F}, \mathrm{Pl}')$ such that Pl' is default-equivalent to Pl. Since $\mathcal{F}$ is finite, the set $\{d : \exists A \in \mathcal{F}, \mathrm{Pl}'(A) = d\}$ is finite. Moreover, since Pl' is a ranking, this set is totally ordered. Let $d_0 > d_1 > \ldots > d_n$ be an ordered enumeration of this set of values.

Let κ be a κ-ranking such that $\kappa(A) = k$ if $\mathrm{Pl}'(A) = d_k > \perp$ and $\kappa(A) = \infty$ if $\mathrm{Pl}'(A) = \perp$. Similarly, let Poss be a possibility measure such that $\mathrm{Poss}(A) = 1 - k/n$

if $\mathrm{Pl}'(A) = d_k > \perp$ and $\mathrm{Poss}(A) = 0$ if $\mathrm{Pl}'(A) = \perp$. It easy to see that since Pl′ is a ranking, we get that $\kappa(A \cup B) = \min(\kappa(A), \kappa(B))$ and that $\mathrm{Poss}(A \cup B) = \max(\mathrm{Poss}(A), \mathrm{Poss}(B))$. It is also easy to see that both measures are equivalent to Pl′ and thus default-equivalent to Pl. □

THEOREM 8.6. *(a)* $\mathbf{C}+\{C6\}$ *is a sound and complete axiomatization of* $\mathcal{S}_c^{\epsilon}$.
(b) $\mathbf{C}+\{C5,\ C6\}$ *is a sound and complete axiomatization of* $\mathcal{S}_c^{\kappa}$ *and* $\mathcal{S}_c^{Poss}$.

PROOF. Part (a) is an immediate corollary of Theorems 6.3 and A.8.

For Part (b), as usual, it is easy to verify soundness. For completeness, it suffices to show that if ϕ is consistent with system $\mathbf{C}+\{$C5,C6$\}$, then it is satisfiable in $\mathcal{S}_c^{\kappa}$ and $\mathcal{S}_c^{Poss}$. Suppose that ϕ is consistent with system $\mathbf{C}+\{$C5,C6$\}$. By Theorem A.8, ϕ is satisfiable in a rational and normal structure in $\mathcal{S}_C^{QPL}$. By Lemma A.9, we can assume that this structure is finite. The result now follows from Lemma A.10. □

REFERENCES

ADAMS, E. 1975. *The Logic of Conditionals*. D. Reidel, Dordrecht, Netherlands.

BACCHUS, F., GROVE, A. J., HALPERN, J. Y., AND KOLLER, D. 1993. Statistical foundations for default reasoning. In *Proc. Thirteenth International Joint Conference on Artificial Intelligence (IJCAI '93)*, pp. 563–569.

BEN-DAVID, S. AND BEN-ELIYAHU, R. 1994. A modal logic for subjective default reasoning. In *Proc. 9th IEEE Symp. on Logic in Computer Science*, pp. 477–486.

BOSSU, G. AND SIEGEL, P. 1985. Saturation, nonmonotonic reasoning and the closed-world assumption. *Artificial Intelligence 25*, 13–63.

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 Journal of Formal Logic 22*, 76–84.

CROCCO, G. AND LAMARRE, P. 1992. On the connection between non-monotonic inference systems and conditional logics. In B. Nebel, C. Rich, and W. Swartout (Eds.), *Proc. Third International Conference on Principles of Knowledge Representation and Reasoning (KR '92)*, pp. 565–571. San Francisco: Morgan Kaufmann.

DARWICHE, A. 1992. *A Symbolic Generalization of Probability Theory*. Ph. D. thesis, Stanford University.

DUBOIS, D. AND PRADE, H. 1990. An introduction to possibilistic and fuzzy logics. In G. Shafer and J. Pearl (Eds.), *Readings in Uncertain Reasoning*, pp. 742–761. San Francisco: Morgan Kaufmann.

DUBOIS, D. AND PRADE, H. 1991. Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proc. Twelfth International Joint Conference on Artificial Intelligence (IJCAI '91)*, pp. 419–424. San Francisco: Morgan Kaufmann.

FREUND, M. 1996. Preferential orders and plausibility measures. Unpublished manuscript.

FRIEDMAN, N. AND HALPERN, J. Y. 1994. On the complexity of conditional logics. In J. Doyle, E. Sandewall, and P. Torasso (Eds.), *Principles of Knowledge Representation and Reasoning: Proc. Fourth International Conference (KR '94)*, pp. 202–213. San Francisco: Morgan Kaufmann.

FRIEDMAN, N. AND HALPERN, J. Y. 1995. Plausibility measures: a user's manual. In P. Besnard and S. Hanks (Eds.), *Proc. Eleventh Conference on Uncertainty in Artificial Intelligence (UAI '95)*, pp. 175–184. San Francisco: Morgan Kaufmann.

FRIEDMAN, N. AND HALPERN, J. Y. 1996. A qualitative Markov assumption and its implications for belief change. In E. Horvits and F. Jenssen (Eds.), *Proc. Twelfth Conference on Uncertainty in Artificial Intelligence (UAI '96)*, pp. 263–273. San Franscisco: Morgan Kaufmann.

Friedman, N., Halpern, J. Y., and Koller, D. 1996. Conditional first-order logic revisited. In *Proc. National Conference on Artificial Intelligence (AAAI '96)*, pp. 1305–1312. Menlo Park, Calif.: AAAI Press.

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, Oxford, U.K.

Gärdenfors, P. 1988. *Knowledge in Flux.* MIT Press, Cambridge, Mass.

Gärdenfors, P. and Makinson, D. 1988. Revisions of knowledge systems using epistemic entrenchment. In M. Vardi (Ed.), *Proc. Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pp. 83–95. San Francisco: Morgan Kaufmann.

Gärdenfors, P. and Makinson, D. 1989. Relations between the theory of change and nonmonotonic logic. In A. Fuhrmann and M. Morreau (Eds.), *Logic of Theory Change. Workshop proceedings*, pp. 185–205. Springer-Verlag.

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

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, San Francisco.

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

Goldszmidt, M. and Pearl, J. 1992. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In B. Nebel, C. Rich, and W. Swartout (Eds.), *Proc. Third International Conference on Principles of Knowledge Representation and Reasoning (KR '92)*, pp. 661–672. San Francisco: Morgan Kaufmann.

Greco, G. H. 1987. Fuzzy integrals and fuzzy measures with their values in complete lattices. *Journal of Mathematical Analysis and Applications 126*, 594–603.

Grove, A. 1988. Two modelings for theory change. *Journal of Philosophical Logic 17*, 157–170.

Hansson, B. 1969. An analysis of some deontic logics. *Nous 3*, 373–398. Reprinted in R. Hilpinen ed. *Deontic Logic: Introductory and systematic readings*, Reidel, Dordrecht, 1971, pp. 121–147.

Jeffreys, H. 1961. *Theory of Probability.* Oxford University Press, Oxford, U.K.

Katsuno, H. and Satoh, K. 1991. A unified view of consequence relation, belief revision and conditional logic. In *Proc. Twelfth International Joint Conference on Artificial Intelligence (IJCAI '91)* (1991), pp. 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, Cambridge, Mass.

Pearl, J. 1989. Probabilistic semantics for nonmonotonic reasoning: a survey. In R. J. Brachman, H. J. Levesque, and R. Reiter (Eds.), *Proc. First International Conference on Principles of Knowledge Representation and Reasoning (KR '89)* (1989), pp. 505–516. Reprinted in *Readings in Uncertain Reasoning*, G. Shafer and J. Pearl (eds.), Morgan Kaufmann, San Francisco, Calif., 1990, pp. 699–710.

Pearl, J. 1990. System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In R. Parikh (Ed.), *Theoretical Aspects of Reasoning about Knowledge: Proc. Third Conference*, pp. 121–135. San Francisco: Morgan Kaufmann.

Schlechta, K. 1995. Defaults as generalized quantifiers. *Journal of Logic and Computation 5*, 473–494.

Schlechta, K. 1996. Filters and partial orders. Submitted for publication.

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

SHOHAM, Y. 1987. A semantical approach to nonmonotonic logics. In *Proc. 2nd IEEE Symp. on Logic in Computer Science* (1987), pp. 275–279. Reprinted in M. L. Ginsberg (Ed.), *Readings in Nonmonotonic Reasoning*, Morgan Kaufman, San Francisco, Calif., 1987, pp. 227–250.

SPOHN, W. 1988. Ordinal conditional functions: a dynamic theory of epistemic states. In W. Harper and B. Skyrms (Eds.), *Causation in Decision, Belief Change and Statistics*, Volume 2, pp. 105–134. Dordrecht, Netherlands: Reidel.

WANG, Z. AND KLIR, G. J. 1992. *Fuzzy Measure Theory*. Plenum Press, New York.

WEBER, S. 1991. Uncertainty measures, decomposability and admissibility. *Fuzzy Sets and Systems 40*, 395–405.

WEYDERT, E. 1994a. Doxastic normality logic: A qualitative probabilistic modal framework for defaults. In A. Fuhrmann and H. Rott (Eds.), *Logic, Action and Information*. New York: de Gruyter.

WEYDERT, E. 1994b. General belief measures. In R. López de Mantarás and D. Poole (Eds.), *Proc. Tenth Conference on Uncertainty in Artificial Intelligence (UAI '94)*, pp. 575–582. San Francisco: Morgan Kaufmann.