First-Order Conditional Logic Revisited*

Nir Friedman
Institute of Computer Science
Hebrew University
Jerusalem, 91904 Israel
nir@cs.huji.ac.il

Joseph Y. Halpern
Dept. of Computer Science
Cornell University
Ithaca, NY 14853
halpern@cs.cornell.edu

Daphne Koller
Dept. of Computer Science
Stanford University
Stanford, CA 94305-9010
koller@cs.stanford.edu

August 31, 1998

Abstract

Conditional logics play an important role in recent attempts to formulate theories of default reasoning. This paper investigates first-order conditional logic. We show that, as for first-order probabilistic logic, it is important not to confound *statistical* conditionals over the domain (such as "most birds fly"), and *subjective* conditionals over possible worlds (such as "I believe that Tweety is unlikely to fly"). We then address the issue of ascribing semantics to first-order conditional logic. As in the propositional case, there are many possible semantics. To study the problem in a coherent way, we use *plausibility structures*. These provide us with a general framework in which many of the standard approaches can be embedded. We show that while these standard approaches are all the same at the propositional level, they are significantly different in the context of a first-order language. Furthermore, we show that plausibilities provide the most natural extension of conditional logic to the first-order case: We provide a sound and complete axiomatization that contains only the KLM properties and standard axioms of first-order modal logic. We show that most of the other approaches have additional properties, which result in an inappropriate treatment of an infinitary version of the *lottery paradox*.

*A preliminary version of this paper appears in *Proc. National Conference on Artificial Intelligence (AAAI '96)*, 1996, pp. 1305–1312. Some of this work was done while all three authors were at the IBM Almaden Research Center, supported by the Air Force Office of Scientific Research (AFSC) under Contract F49620-91-C-0080; some was done while Daphne Koller was at U.C. Berkeley, supported by a University of California President's Postdoctoral Fellowship; and some was done while Nir Friedman was at Stanford University. This work was also partially supported by NSF grants IRI-95-03109 and IRI-96-25901.

1 Introduction

In recent years, conditional logic has come to play a major role as an underlying foundation for default reasoning. Two proposals that have received a lot of attention [Geffner 1992; Goldszmidt, Morris, and Pearl 1993] are based on conditional logic. Unfortunately, while it has long been recognized that first-order expressive power is necessary for a default reasoning system, most of the work on conditional logic has been restricted to the propositional case. In this paper, we investigate the syntax and semantics of *first-order conditional logic*, with the ultimate goal of providing a first-order default reasoning system.

Many seemingly different approaches have been proposed for giving semantics to conditional logic, including preferential structures [Lewis 1973; Boutilier 1994; Kraus, Lehmann, and Magidor 1990], ϵ-semantics [Adams 1975; Pearl 1989], *possibility theory* [Dubois and Prade 1991], and *κ-rankings* [Spohn 1988; Goldszmidt and Pearl 1992]. In preferential structures, for example, a model consists of a set of possible worlds, ordered by a preference ordering $\prec$. If $w \prec w'$, then the world w is strictly more preferred/more normal than w'. The formula $\mathit{Bird} \rightarrow \mathit{Fly}$ holds if in the most preferred worlds in which *Bird* holds, *Fly* also holds. (See Section 2 for more details about this and the other approaches.)

The extension of these approaches to the first-order case seems deceptively easy. After all, we can simply have a preference ordering on first-order, rather than propositional, worlds. However, there is a subtlety here. As in the case of first-order *probabilistic* logic [Bacchus 1990; Halpern 1990], there are two distinct ways to define conditionals in the first-order case. In the probabilistic case, the first corresponds to (objective) statistical statements, such as "90% of birds fly". The second corresponds to subjective degree of belief statements, such as "the probability that Tweety (a particular bird) flies is 0.9". The first is captured by putting a probability distribution over the domain (so that the probability of the set of flying birds is 0.9 that of the set of birds), while the second is captured by putting a probability on the set of possible worlds (so that the probability of the set of worlds where Tweety flies is 0.9 that of the set of worlds where Tweety is a bird). The same phenomenon occurs in the case of first-order conditional logic. Here, we can have a measure (e.g., a preference order) over the domain, or a measure over the set of possible worlds. The first would allow us to capture qualitative statistical statements such as "most birds fly", while the second would allow us to capture subjective beliefs such as "I believe that the bird Tweety is likely to fly". It is important to have a language that allows us to distinguish between these two very different statements. Having distinguished between these two types of conditionals, we can ascribe semantics to each of them using any one of the standard approaches.

There have been previous attempts to formalize first-order conditional logic; some are the natural extension of some propositional formalism [Delgrande 1987; Brafman 1997], while others use alternative approaches [Lehmann and Magidor 1990; Schlechta 1995]. (See Sections 5, 8, and 9. for a more detailed discussion of the alternative approaches.) How do we make sense of this plethora of alternatives? Rather than investigating them separately, we use a single common framework that generalizes almost all of them. This framework uses a notion of uncertainty called a *plausibility measure*, introduced by Friedman and Halpern [1995]. A plausibility measure associates with set of worlds its *plausibility*, which is just an element in a partially ordered space. Probability measures are a subclass of plausibility measures, in which the plausibilities lie in $[0,1]$, with the standard ordering. Friedman and Halpern [1998] show that the different standard approaches to conditional logic can all be mapped to plausibility measures, if we interpret $Bird \rightarrow Fly$ as "the set of worlds where $Bird \wedge Fly$ holds has greater plausibility than that of the set of worlds where $Bird \wedge \neg Fly$ holds".

The existence of a single unifying framework has already proved to be very useful in the case of propositional conditional logic. In particular, it allowed Friedman and Halpern [1998] to explain the intriguing "coincidence" that all of the different approaches to conditional logic result in an identical reasoning system, characterized by the *KLM postulates* [Kraus, Lehmann, and Magidor 1990]. In this paper, we show that plausibility spaces can also be used to clarify the semantics of first-order conditional logic. However, we show that, unlike the propositional case, the different approaches lead to different properties in the first-order case. Of course, these are properties that require quantifiers and therefore cannot be expressed in a propositional language. We show that, in some sense, plausibilities provide the most natural extension of conditional logic to the first-order case. We provide sound and complete axiomatizations for both the subjective and statistical variants of first-order conditional logic that contain only the KLM properties and the standard axioms of first-order modal logic.[1] Essentially the same axiomatizations are shown to be sound and complete for the first-order version of ϵ-semantics, but the other approaches are shown to satisfy additional properties.

One might think that it is not so bad for a conditional logic to satisfy additional properties. After all, there are some properties—such as indifference to irrelevant information—that we would *like* to be able to get. Unfortunately, the additional properties that we get from using these approaches are not the ones we want. The properties we get are related to the treatment of *exceptional individuals*. This issue is perhaps best illustrated by the *lottery paradox* [Kyburg 1961].[2] Suppose we believe about a lottery that any particular individual typically does not win the lottery. Thus we get

$$\forall x(true \rightarrow \neg Winner(x)). \tag{1}$$

However, we believe that typically someone does win the lottery, that is

$$true \rightarrow \exists x\, Winner(x). \tag{2}$$

Let *Lottery* be the conjunction of (1) and (2).

Unfortunately, in many of the standard approaches, such as Delgrande's [1987] version of first-order preferential structures, from (1) we can conclude

$$true \rightarrow \forall x(\neg Winner(x)). \tag{3}$$

[1]By way of contrast, there is no (recursively enumerable) axiomatization of of either statistical or subjective first-order probabilistic logic; the validity problem for these logics is highly undecidable (Π^2_1 complete) [Abadi and Halpern 1994].

[2]We are referring to Kyburg's original version of the lottery paradox [Kyburg 1961], and not to the finitary version discussed by Poole [1991]. As Poole showed, any logic of defaults that satisfies certain minimal properties—properties which are satisfied by all the logics we consider—is bound to suffer from his version of the lottery paradox.

Intuitively, from (1) it follows that in the most preferred worlds, each individual d does not win the lottery. Therefore, in the most preferred worlds, no individual wins. This is exactly what (3) says. Since (2) says that in the most preferred worlds, some individual wins, it follows that there are no most preferred worlds, i.e., we have $true \rightarrow false$. While this may be consistent (as it is in Delgrande's logic), it implies that all defaults hold, which is surely not what we want. Of all the approaches, only ϵ-semantics and plausibility structures, both of which are fully axiomatized by the first-order extension of the KLM axioms, do not suffer from this problem.

It may seem that this problem is perhaps not so serious. After all, how often do we reason about lotteries? But, in fact, this problem arises in many situations which are clearly of the type with which we would like to deal. Assume, for example, that we express the default "birds typically fly" as Delgrande does, using the statement

$$\forall x(Bird(x) \rightarrow Fly(x)). \tag{4}$$

Suppose we also believe that Tweety is a bird that does not fly. There are a number of ways we can capture beliefs in conditional logic. The most standard [Friedman and Halpern 1997] is to identify belief in φ with φ typically being true, that is, with $true \rightarrow \varphi$. Using this approach, our knowledge base would contain the statement $true \rightarrow Bird(Tweety) \wedge \neg Fly(Tweety)$we could similarly conclude $true \rightarrow false$. Again, this is surely not what we want.

Our framework allows us to deal with these problems. Using plausibilities, *Lottery* does not not imply $true \rightarrow false$, since (3) does not follow from (1). That is, the lottery paradox simply does not exist if we use plausibilities. The flying bird example is somewhat more subtle. If we take Tweety to be a *nonrigid designator* (so that it might denote different individuals in different worlds), the two statements are consistent, and the problem disappears. If, however, Tweety is a rigid designator, the pair is inconsistent, as we would expect.[3]

This inconsistency suggests that we might not always want to use (4) to represent "birds typically fly". After all, the former is a statement about a property believed to hold of each individual bird, while the latter is a statement about the class of birds. As argued in [Bacchus, Grove, Halpern, and Koller 1996], defaults often arise from statistical facts about the domain. That is, the default "birds typically fly" is often a consequence of the empirical observation that "almost all birds fly". By defining a logic which allows us to express statistical conditional statements, we provide the user an alternative way of representing such defaults. We would, of course, like such statements to impact our beliefs about individual birds. In [Bacchus, Grove, Halpern, and Koller 1996], the same issue was addressed in the probabilistic context, by presenting an approach for going from statistical knowledge bases to subjective degrees of belief. We leave the problem of providing a similar mechanism for conditional logic to future work.

The rest of this paper is organized as follows. In Section 2, we review the various approaches to conditional logic in the propositional case; we also review the definition of plausibility measures from [Friedman and Halpern 1998] and show how they provide a common

[3] To see this, note that if Tweety is a rigid designator, then $Bird(Tweety) \rightarrow Fly(Tweety)$ is a consequence of (4). See the discussion in Section 3 for more details on this point.

framework for these different approaches. In the next three sections, we focus on first-order subjective conditional logic. In Section 3, we describe the syntax for the language and ascribe semantics to formulas using plausibility. In Section 4, we provide a sound and complete axiomatization for first-order subjective conditional assertions. In Section 5, we discuss the generalization of the other propositional approaches to the first-order subjective case, by investigating their behavior with respect to the lottery paradox. We also provide a brief comparison to some of the other approaches suggested in the literature. In Sections 6, 7, and 8, we go through the same exercise for first-order statistical conditional logic, describing the syntax and semantics, providing a complete axiomatization, and comparing to other approaches. We conclude in Section 9 with discussion and some directions for further work.

2 Propositional conditional logic

The syntax of propositional conditional logic is simple. We start with a set Φ of propositions and close off under the usual propositional connectives ($\neg$, $\vee$, $\wedge$, and $\Rightarrow$, denoting, negation, disjunction, conjunction and material implication, respectively) and the conditional connective $\rightarrow$. That is, if φ and ψ are formulas in the language, so are $\neg\varphi, \varphi \vee \psi, \varphi \wedge \psi, \varphi \Rightarrow \psi$, and $\varphi \rightarrow \psi$.

Many semantics have been proposed in the literature for conditionals. Most of them involve structures of the form (W, X, π), where W is a set of possible worlds, $\pi(w)$ is a truth assignment to primitive propositions, and X is some "measure" on W such as a preference ordering [Lewis 1973; Kraus, Lehmann, and Magidor 1990].[4] We now describe some of the proposals in the literature, and then show how they can be generalized. Given a structure (W, X, π), let $[\![\varphi]\!] \subseteq W$ be the set of worlds satisfying φ.

- A *possibility measure* [Dubois and Prade 1990] Poss is a function Poss : $2^W \mapsto [0, 1]$ such that $\mathrm{Poss}(W) = 1$, $\mathrm{Poss}(\emptyset) = 0$, and $\mathrm{Poss}(A) = \sup_{w \in A}(\mathrm{Poss}(\{w\})$. A *possibility structure* is a tuple (W, Poss, π), where Poss is a possibility measure on W. It satisfies a conditional $\varphi \rightarrow \psi$ if either $\mathrm{Poss}([\![\varphi]\!]) = 0$ or $\mathrm{Poss}([\![\varphi \wedge \psi]\!]) > \mathrm{Poss}([\![\varphi \wedge \neg\psi]\!])$ [Dubois and Prade 1991]. That is, either φ is impossible, in which case the conditional holds vacuously, or $\varphi \wedge \psi$ is more possible than $\varphi \wedge \neg\psi$.
- A *κ-ranking* (or *ordinal ranking*) on W (as defined by [Goldszmidt and Pearl 1992], based on ideas that go back to [Spohn 1988]) is a function $\kappa : 2^W \rightarrow \mathbb{N}^*$, where $\mathbb{N}^* = \mathbb{N} \cup \{\infty\}$, 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. A *κ-structure* is a tuple (W, κ, π), where κ is an ordinal ranking on W. It satisfies a conditional $\varphi \rightarrow \psi$ if either $\kappa([\![\varphi]\!]) = \infty$ or $\kappa([\![\varphi \wedge \psi]\!]) < \kappa([\![\varphi \wedge \neg\psi]\!])$.

[4]We could also consider a more general definition, in which one associates a different "measure" with each world, as done by Lewis [1973]. It is straightforward to extend our definitions to handle this. Since this issue is orthogonal to the main point of the paper, we do not discuss it further here.

- A *preference ordering* on W is a partial order $\prec$ over W [Kraus, Lehmann, and Magidor 1990; Shoham 1987]. Intuitively, $w \prec w'$ holds if w is *preferred* to w'. A *preferential structure* is a tuple $(W, \prec, \pi)$, where $\prec$ is a partial order on W. The intuition [Shoham 1987] is that a preferential structure satisfies a conditional $\varphi \rightarrow \psi$ if all the most preferred worlds (i.e., the minimal worlds according to $\prec$) in $[\![\varphi]\!]$ satisfy ψ. However, there may be no minimal worlds in $[\![\varphi]\!]$. This can happen if $[\![\varphi]\!]$ contains an infinite descending sequence $\ldots \prec w_2 \prec w_1$. What do we do in these structures? There are a number of options: the first is to assume that, there are no infinite descending sequences, i.e., that $\prec$ is *well-founded*; this is essentially the assumption made by Kraus, Lehmann, and Magidor [1990].[5] A yet more general definition—one that works even if $\prec$ is not well-founded—is given in [Lewis 1973; Boutilier 1994]. Roughly speaking, $\varphi \rightarrow \psi$ is true if, from a certain point on, whenever φ is true, so is ψ. More formally,

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

  It is easy to verify that this definition is equivalent to the earlier one if $\prec$ is well founded.

- A *parameterized probability distribution* (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 over W. Intuitively, it satisfies a conditional $\varphi \rightarrow \psi$ if the conditional probability ψ given φ goes to 1 in the limit. Formally, $\varphi \rightarrow \psi$ is satisfied if $\lim_{i \rightarrow \infty} \Pr_i([\![\psi]\!] | [\![\psi]\!]) = 1$ (where $\Pr_i([\![\psi]\!] | [\![\varphi]\!])$ is taken to be 1 if $\Pr_i([\![\varphi]\!]) = 0$). PPD structures were introduced in [Goldszmidt, Morris, and Pearl 1993] as a reformulation of Pearl's *ϵ-semantics* [Pearl 1989].

These variants are quite different from each other. As Friedman and Halpern [1998] show, we can provide a uniform framework for all of them using the notion of plausibility measures.

A *plausibility measure* Pl on W is a function that maps subsets of W to elements in some arbitrary partially ordered set. We read $\mathrm{Pl}(A)$ as "the plausibility of set A". If $\mathrm{Pl}(A) \leq \mathrm{Pl}(B)$, 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 (that is, a set of subsets closed under union and complementation), and Pl maps the sets in $\mathcal{F}$ to some set D, partially ordered by a relation $\leq$ (so that $\leq$ is reflexive, transitive, and anti-symmetric). To simplify notation, we typically omit the algebra $\mathcal{F}$ from the description of the plausibility space. As usual, we define the ordering $<$ by taking $d_1 < d_2$ if $d_1 \leq d_2$ and $d_1 \neq d_2$. We assume that D is *pointed*: that is, it contains two special elements $\top$ and $\bot$ such that

[5] Actually, they make a weaker assumption, called *smoothness*, that for each formula φ, there are minimal worlds in $[\![\varphi]\!]$, i.e., that $\prec$ is well-founded on the sets of interest. All the results we prove for well-founded preferential structures hold for smooth ones as well.

$\perp \leq d \leq \top$ for all $d \in D$; we further assume that $\mathrm{Pl}(W) = \top$ and $\mathrm{Pl}(\emptyset) = \perp$. Since we want a set to be at least as plausible as any of its subsets, we require:

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

Clearly, plausibility spaces generalize probability spaces. Other approaches to dealing with uncertainty, such as possibility measures, κ-rankings, and *belief functions* [Shafer 1976], are also easily seen to be plausibility measures.

We can give semantics to conditionals using plausibility in much the same way as it is done using possibility. A *plausibility structure* is a tuple $PL = (W, \mathrm{Pl}, \pi)$, where Pl is a plausibility measure on W. We then define:

- $PL \models \varphi \rightarrow \psi$ if either $\mathrm{Pl}([\![\varphi]\!]) = \perp$ or $\mathrm{Pl}([\![\varphi \wedge \psi]\!]) > \mathrm{Pl}([\![\varphi \wedge \neg\psi]\!])$.

Intuitively, $\varphi \rightarrow \psi$ holds vacuously if φ is impossible; otherwise, it holds if $\varphi \wedge \psi$ is more plausible than $\varphi \wedge \neg\psi$. It is easy to see that this semantics for conditionals generalizes the semantics of conditionals in possibility structures and κ-structures. We are implicitly assuming here that $[\![\varphi]\!]$ is in $\mathcal{F}$ (i.e., in the domain of Pl) for each formula φ.

As shown in [Friedman and Halpern 1998, Theorem 4.2], it also generalizes the semantics of conditionals in preferential structures and PPD structures. More precisely, a mapping is given from preferential structures (resp., PPD structures) to plausibility structures such that the semantics of defaults are preserved. For future reference, we sketch these constructions here.

For PPDs, it is quite straightforward. Given a PPD $PP = (\mathrm{Pr}_1, \mathrm{Pr}_2, \ldots)$ on a space W, we can define a plausibility measure Pl_{PP} such that $\mathrm{Pl}_{PP}(A) \leq \mathrm{Pl}_{PP}(B)$ iff $\lim_{i \to \infty} \mathrm{Pr}_i(B | A \cup B) = 1$. It can then be shown that $((W, PP, \pi), w) \models \varphi$ iff $((W, \mathrm{Pl}_{PP}, \pi), w) \models \varphi$ for all $w \in W$ and interpretations π.

The mapping of 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:

$$\mathrm{Pl}_\prec(A) \leq \mathrm{Pl}_\prec(B) \text{ if and only if for all } w \in A - B \text{, there is a world } w' \in B \text{ such that } w' \prec w \text{ and there is no } w'' \in A - B \text{ such that } w'' \prec w'. \tag{5}$$

It is then easy to check that $((W, \prec, \pi), w) \models \varphi$ if and only if $((W, \mathrm{Pl}_\prec, \pi), w) \models \varphi$, for all $w \in W$ and interpretations π.

These results show that our semantics for conditionals in plausibility structures generalizes the various approaches examined in the literature. Does it capture our intuitions about conditionals? In the AI literature, there has been discussion of the right properties of default statements (which are essentially conditionals). While there has been little consensus on what the "right" properties for defaults should be, there has been some consensus on a reasonable "core" of inference rules for default reasoning. This core, is known as the KLM properties [Kraus, Lehmann, and Magidor 1990]. We briefly list these properties here:

LLE. If $\vdash \varphi \Leftrightarrow \varphi'$[6], then from $\varphi \rightarrow \psi$ infer $\varphi' \rightarrow \psi$ (Left Logical Equivalence)

RW. If $\vdash \psi \Rightarrow \psi'$, then from $\varphi \rightarrow \psi$ infer $\varphi \rightarrow \psi'$ (Right Weakening)

REF. $\varphi \rightarrow \varphi$ (Reflexivity)

AND. From $\varphi \rightarrow \psi_1$ and $\varphi \rightarrow \psi_2$ infer $\varphi \rightarrow \psi_1 \wedge \psi_2$ (And)

OR. From $\varphi_1 \rightarrow \psi$ and $\varphi_2 \rightarrow \psi$ infer $\varphi_1 \vee \varphi_2 \rightarrow \psi$ (Or)

CM. From $\varphi \rightarrow \psi_1$ and $\varphi \rightarrow \psi_2$ infer $\varphi \wedge \psi_2 \rightarrow \psi_1$ (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 $\varphi_2 \rightarrow \psi$ from $\varphi_1 \rightarrow \psi$. RW describes a similar property of the consequent: If ψ (logically) entails ψ', then we can deduce $\varphi \rightarrow \psi'$ from $\varphi \rightarrow \psi$. This allows us to can combine default and logical reasoning. REF states that φ is always a default conclusion of φ. AND states that we can combine two default conclusions: If we can conclude by default both ψ_1 and ψ_2 from φ, 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.

Do conditionals in plausibility structures satisfy the KLM properties? They always satisfy REF, LLE, and RW, but they do not in general satisfy AND, OR, and CM. To satisfy the KLM properties we must limit our attention to plausibility structures that satisfy the following two conditions:

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 $\mathrm{Pl}(A) > \mathrm{Pl}(B \cup C)$.

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

A plausibility space (W, Pl) is *qualitative* if it satisfies A2 and A3. A plausibility structure (W, Pl, π) is qualitative if (W, Pl) is a qualitative plausibility space. Friedman and Halpern [1998] show that, in a very general sense, qualitative plausibility structures capture default reasoning. More precisely, the KLM properties are sound with respect to a class

[6] where $\vdash$ denotes provability in propositional logic

of plausibility structures if and only if the class consists of qualitative plausibility structures. Furthermore, a very weak condition is necessary and sufficient in order for the KLM properties to be a complete axiomatization of conditional logic. As a consequence, once we consider a class of structures where the KLM axioms are sound, it is almost inevitable that they will also be complete with respect to that class. This explains the somewhat surprising fact that KLM properties characterize default entailment not just in preferential structures, but also in ϵ-semantics, possibility measures, and κ-rankings. Each one of these approaches corresponds, in a precise sense, to a class of qualitative plausibility structures. These results show that plausibility structures provide a unifying framework for the characterization of default entailment in these different logics.

3 First-order subjective conditional logic

We now want to generalize conditional logic to the first-order case. As mentioned above, there are two distinct notions of conditionals in first-order logic, one involving statistical conditionals and one involving subjective conditionals. For each of these, we use a different syntax, analogous to the syntax used in [Halpern 1990] for the probabilistic case. In the next three sections, we focus on the subjective case; in Sections 6, 7, and 8, we consider the statistical case.

The syntax for subjective conditional logic is fairly straightforward. Let Φ be a first-order vocabulary, consisting of predicate and function symbols. (As usual, constant symbols are viewed as 0-ary function symbols.) Starting with atomic formulas of first-order logic, we form more complicated formulas by closing off under truth-functional connectives (i.e., $\wedge, \vee, \neg$, and $\Rightarrow$), first-order quantification, and the modal operator $\rightarrow$. Thus, a typical formula is $\forall x(P(x) \rightarrow \exists y(Q(x,y) \rightarrow R(y)))$. Let $\mathcal{L}^{subj}(\Phi)$ be the resulting language (the "subj" stands for "subjective", since the conditionals are viewed as expressing subjective degrees of belief). We typically omit the Φ if it is clear from context or irrelevant.

We can ascribe semantics to subjective conditionals using any one of the approaches described in the previous section. However, since we can embed all of the approaches within the class of plausibility structures, we use these as the basic semantics. As in the propositional case, we can then analyze the behavior of the other approaches simply by restricting attention to the appropriate subclass of plausibility structures.

To give semantics to $\mathcal{L}^{subj}(\Phi)$, we use *(first-order) subjective plausibility structures* over Φ. These are tuples of the form $PL = (Dom, W, \mathrm{Pl}, \pi)$, where Dom is a domain, (W, Pl) is a plausibility space and $\pi(w)$ is an interpretation assigning to each predicate symbol and function symbol in Φ a predicate or function of the right arity over Dom. As usual, a *valuation* maps each variable to an element of Dom. We define the set of worlds that satisfy φ given the valuation v to be $[\![\varphi]\!]_{(PL,v)} = \{w : (PL, w, v) \models \varphi\}$. (We omit the subscript whenever it is clear from context.) For subjective conditionals, we have

- $(PL, w, v) \models \varphi \rightarrow \psi$ if $\mathrm{Pl}([\![\varphi]\!]_{(PL,v)}) = \perp$ or $\mathrm{Pl}([\![\varphi \wedge \psi]\!]_{(PL,v)}) > \mathrm{Pl}([\![\varphi \wedge \neg\psi]\!]_{(PL,v)})$.

The semantics of atomic formulas and quantifiers is the same as in first-order logic. As an example, for the atomic formula $P(x, \mathbf{c})$, we have

- $(PL, w, v) \models P(x, \mathbf{c})$ if $(v(x), \pi(w)(\mathbf{c})) \in \pi(w)(P)$.

Note that $\pi(w)(\mathbf{c})$ is the interpretation of the constant $\mathbf{c}$ in the world w. There may be a different interpretation of $\mathbf{c}$ in each world; that is, we may have $\pi(w)(\mathbf{c}) \neq \pi(w')(\mathbf{c})$ if $w \neq w'$. Thus, $\mathbf{c}$ is *nonrigid*. We return to this issue below. Similarly, $\pi(w)(P)$ is the interpretation of P in w.

To give the semantics of quantification, it is useful to define a family of equivalence relations $\sim_X$ on valuations, where X is a set of variables. We write $v \sim_X v'$ if v and v' agree on the values they give to all variables except possibly those in X. If X is the singleton $\{x\}$, we write $\sim_x$ instead of $\sim_{\{x\}}$.

- $(PL, w, v) \models \forall x \varphi$ if $(\mathrm{Pl}, w, v') \models \varphi$ for all valuations v' such that $v' \sim_x v$.

Because terms are not rigid designators, we cannot substitute terms for universally quantified variables. (A similar phenomenon holds in other modal logics where terms are not rigid [Garson 1977].) For example, let $N\varphi$ be an abbreviation for $\neg\varphi \rightarrow \mathit{false}$. Notice that $(PL, w) \models N\varphi$ if $\mathrm{Pl}([\![\neg\varphi]\!]) = \perp$; i.e., $N\varphi$ asserts that the plausibility of $\neg\varphi$ is the same as that of the empty set, so that φ is true "almost everywhere".[7] Suppose $\mathbf{c}$ is a constant that does not appear in the formula φ. It is not hard to see that $\forall x(\neg N\varphi(x)) \Rightarrow (\neg N\varphi(\mathbf{c}))$ is not valid in our framework; that is, we cannot substitute constants for universally quantified variables. To see this, let $\varphi(x)$ be the formula $P(x)$, where P is a unary predicate. Consider the plausibility structure $PL = (\{d_1, d_2\}, \{w_1, w_2\}, \mathrm{Pl}, \pi)$, where π is such that $\mathbf{c}$ is d_1 in world w_1 and d_2 in world w_2, the extension of P in w_1 is $\{d_1\}$ and the extension of P in w_2 is $\{d_2\}$, and Pl is such that $\mathrm{Pl}(\{w_1\}) = \mathrm{Pl}(\{w_2\} \neq \perp$. It is easy to see that $(PL, w_1) \models \forall x(\neg NP(x)) \wedge NP(c)$.

We could substitute $\mathbf{c}$ for x in $\forall x \varphi(x)$ if $\mathbf{c}$ were rigid. We can get the effect of rigidity by assuming that $\exists x(N(x = \mathbf{c}))$ holds. Thus, we do not lose expressive power by not assuming rigidity.

As in first-order logic, a *sentence* is a formula with no free variables. It is easy to check that, just as in first-order logic, the truth of a sentence is independent of the valuation. Thus, if φ is a sentence, we often write $(PL, w) \models \varphi$ rather than $(PL, w, v) \models \varphi$.

4 Axiomatizing first-order subjective conditional logic

We now want to show that plausibility structures provide an appropriate semantics for a first-order logic of defaults. As in the propositional case, this is true only if we restrict attention to qualitative plausibility structures, i.e., those satisfying conditions A2 and A3

[7] N stands for "necessary".

above. Let $\mathcal{P}_{subj}^{QPL}$ be the class of all subjective qualitative plausibility structures. We provide a sound and complete axiom system for $\mathcal{P}_{subj}^{QPL}$, and show that it is the natural extension of the KLM properties to the first-order case.

The system $\mathbf{C}^{subj}$ consists of all generalizations of the following axioms (where φ is a *generalization* of ψ if φ is of the form $\forall x_1 \ldots \forall x_n \psi$) and rules. In the axioms x and y denote variables, while t denotes an arbitrary term. $\mathbf{C}^{subj}$ consists of three parts. The first set of axioms (C0–C5 together with the rules MP, R1, and R2) is simply the standard axiomatization of propositional conditional logic [Hughes and Cresswell 1968]; the second set (axioms F1–F5) consists of the standard axioms of first-order logic [Enderton 1972]. the final set (F6–F7) contains standard axioms relating the two [Hughes and Cresswell 1968]. These axioms describe the interaction between N and equality, and hold because we are essentially treating variables as rigid designators.

C0. All instances of propositional tautologies

C1. $\varphi \rightarrow \varphi$

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

C3. $((\varphi_1 \rightarrow \psi) \wedge (\varphi_2 \rightarrow \psi)) \Rightarrow ((\varphi_1 \vee \varphi_2) \rightarrow \psi)$

C4. $((\varphi_1 \rightarrow \varphi_2) \wedge (\varphi_1 \rightarrow \psi)) \Rightarrow ((\varphi_1 \wedge \varphi_2) \rightarrow \psi)$

C5. $[(\varphi \rightarrow \psi) \Rightarrow N(\varphi \rightarrow \psi)] \wedge [\neg(\varphi \rightarrow \psi) \Rightarrow N\neg(\varphi \rightarrow \psi)]$

F1. $\forall x \varphi \Rightarrow \varphi[x/t]$, where t is *substitutable* for x in the sense discussed below and $\varphi[x/t]$ is the result of substituting t for all free occurrences of x in φ (see [Enderton 1972] for a formal definition)

F2. $\forall x(\varphi \Rightarrow \psi) \Rightarrow (\forall x \varphi \Rightarrow \forall x \psi)$

F3. $\varphi \Rightarrow \forall x \varphi$ if x does not occur free in φ

F4. $x = x$

F5. $x = y \Rightarrow (\varphi \Rightarrow \varphi')$, where φ is a quantifier-free and $\rightarrow$-free formula and φ' is obtained from φ by replacing zero or more occurrences of x in φ by y

F6. $x = y \Rightarrow N(x = y)$

F7. $x \neq y \Rightarrow N(x \neq y)$

MP. From φ and $\varphi \Rightarrow \psi$ infer ψ

R1. From $\varphi_1 \Leftrightarrow \varphi_2$ infer $\varphi_1 \rightarrow \psi \Leftrightarrow \varphi_2 \rightarrow \psi$

R2. From $\psi_1 \Rightarrow \psi_2$ infer $\varphi \rightarrow \psi_1 \Rightarrow \varphi \rightarrow \psi_2$.

It remains to explain the notion of "substitutable" in F1. Clearly we cannot substitute a term t for x with free variables that might be captured by some quantifiers in φ; for example, while $\forall x \exists y(x \neq y)$ is true as long as the domain has at least two elements, if we substitute y for x, we get $\exists y(y \neq y)$, which is surely false. In the case of first-order logic, it suffices to define "substitutable" so as to make sure this does not happen (see [Enderton 1972] for details). However, in modal logics such as this one, we have to be a little more careful. As we observed in Section 3, we cannot substitute terms for universally quantified variables in a modal context, since terms are not in general rigid. Thus, we require that if φ is a formula that has occurrences of $\rightarrow$, then the only terms that are substitutable for x in φ are other variables.

We claim that $\mathbf{C}^{subj}$ is the weakest "natural" first-order extension of the KLM properties. The bulk of the propositional fragment of this axiom system (axioms C1–C4, R1, and R2) corresponds precisely to the KLM properties. For example, C1 is just REF, C2 is AND, R1 is LLE, and so on. The remaining axiom (C5) captures the fact that the plausibility function Pl is independent of the world. We could consider a more general semantics where the plausibility measure used depends on the world (see [Friedman and Halpern 1998, Section 8]); in this case, we would drop C5. This property does not appear in [Kraus, Lehmann, and Magidor 1990] since they do not allow nesting of conditionals. As discussed above, the remaining axioms are standard properties of first-order modal logic.

The system $\mathbf{C}^{subj}$ characterizes first-order default reasoning in this framework:

Theorem 4.1:: $\mathbf{C}^{subj}$ *is a sound and complete axiomatization of* $\mathcal{L}^{subj}$ *with respect to* $\mathcal{P}^{QPL}_{subj}$.

Proof: The proof combines ideas from the standard Henkin-style completeness proof for first-order logic [Enderton 1972] with the proof of completeness for propositional conditional logic given in [Friedman and Halpern 1998]. The details can be found in the appendix. ∎

5 Other approaches to first-order subjective conditional logic

In the previous section we showed that $\mathbf{C}^{subj}$ is sound and complete with respect to $\mathcal{P}^{QPL}_{subj}$. What happens if we use one of the approaches described in Section 2 to give semantics to conditionals? As noted above, we can associate with each of these approach a subset of qualitative plausibility structures. Let $\mathcal{P}^{p,w}_{subj}, \mathcal{P}^{p}_{subj}, \mathcal{P}^{\kappa}_{subj}, \mathcal{P}^{poss}_{subj}$, and $\mathcal{P}^{\epsilon}_{subj}$ be the subsets of $\mathcal{P}^{QPL}_{subj}$ that correspond to well-founded preference orderings, preference orderings, κ-rankings, possibility measures, and PPDs, respectively. From Theorem 4.1, we immediately get

Theorem 5.1:: $\mathbf{C}^{subj}$ *is sound in* $\mathcal{P}^{p,w}_{subj}$, $\mathcal{P}^{p,s}_{subj}$, $\mathcal{P}^{p}_{subj}$, $\mathcal{P}^{\kappa}_{subj}$, $\mathcal{P}^{poss}_{subj}$, *and* $\mathcal{P}^{\epsilon}_{subj}$.

Is $\mathbf{C}^{subj}$ complete with respect to these approaches? Even at the propositional level, it is well known that because κ rankings and possibility measures induce plausibility measures that are total (rather than partial) orders, they satisfy the following additional property:

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

In addition, the plausibility measures induced by κ rankings, possibility measures, and ϵ semantics are easily seen to have the property that $\top > \bot$. This leads to the following axiom:

C7. $\neg(\mathit{true}\rightarrow \mathit{false})$.

In the propositional setting, these additional axioms and the basic propositional conditional system (i.e., C0–C5, MP, LLE, and RW) lead to sound and complete axiomatization of the corresponding (propositional) structures. (See [Friedman and Halpern 1998, Section 8].)

Does the same phenomenon occur in the first-order case? For ϵ-semantics, it does.

Theorem 5.2:: $\mathbf{C}^{subj}$+C7 *is a sound and complete axiomatization of* $\mathcal{L}^{subj}$ *with respect to* $\mathcal{P}^{\epsilon}_{subj}$.

Proof: We combine ideas from the proof of Theorem 4.1 with results from [Friedman and Halpern 1998] showing how a plausibility structure satisfying C7 can be viewed as a PPD structure. The details are in the appendix. ∎

Although ϵ-semantics has essentially the same expressive power in the first-order case as plausibility measures, this is not the case for the other approaches that are characterized by the KLM properties in the propositional case. These approaches all satisfy properties beyond $\mathbf{C}^{subj}$, C6, and C7. And these additional properties are ones that we would argue are undesirable, since they cause the lottery paradox. Recall that *Lottery*, the formula that represents the lottery paradox, is the conjunction of two formulas: (1) $\forall x(\mathit{true}\rightarrow\neg \mathit{Winner}(x))$ states that every individual is unlikely to win the lottery, while (2) $\mathit{true}\rightarrow\exists x\, \mathit{Winner}(x)$ states that is is likely that some individual does win the lottery. We start by showing that *Lottery* is consistent in $\mathcal{P}^{QPL}_{subj}$.

Example 5.3:: We define a first-order subjective plausibility structure $PL_{lot} = (Dom_{lot}, W_{lot}, \mathrm{Pl}_{lot}, \pi_{lot})$ as follows: Dom_{lot} is a countable domain consisting of the individuals $1, 2, 3, \ldots$; W_{lot} consists of a countable number of worlds $w_1, w_2, w_3, \ldots$; Pl_{lot} gives the empty set plausibility 0, each non-empty finite set plausibility 1/2, and each infinite set plausibility 1; finally, the denotation of *Winner* in world w_i according to π_{lot} is the singleton set $\{d_i\}$ (that is, in world w_i the lottery winner is individual d_i). It is easy to check that $[\![\neg \mathit{Winner}(d_i)]\!] = W - \{w_i\}$, so $\mathrm{Pl}_{lot}([\![\neg \mathit{Winner}(d_i)]\!]) = 1 > 1/2 = \mathrm{Pl}([\![\mathit{Winner}(d_i)]\!])$; hence, PL_{lot} satisfies (1). On the other hand, $[\![\exists x\, \mathit{Winner}(x)]\!] = W$, so $\mathrm{Pl}_{lot}([\![\exists x\, \mathit{Winner}(x)]\!]) > \mathrm{Pl}_{lot}([\![\neg\exists x\, \mathit{Winner}(x)]\!])$; hence PL_{lot} satisfies (2). It is also easy to verify that Pl_{lot} is a qualitative measure, i.e., satisfies A2 and A3. A similar construction allows us to capture a situation where birds typically fly but we know that Tweety does not fly. ∎

What happens to the lottery paradox in the other approaches? First consider well-founded preferential structures, i.e., $\mathcal{P}^{p,w}_{subj}$. In these structures, $\varphi\rightarrow\psi$ holds if ψ holds in all the preferred worlds that satisfy φ. Thus, (1) implies that for any domain element d, d

is not a winner in the most preferred worlds. On the other hand, (2) implies that in the most preferred worlds, some domain element wins. Together both imply that there are no preferred worlds. When, in general, does an argument of this type go through? As we now show, it is a consequence of the following generalization of A2.

A2*. If $\{A_i : i \in I\}$ are pairwise disjoint sets, $A = \cup_{i \in I} A_i$, $0 \in I$, and for all $i \in I - \{0\}$, $\mathrm{Pl}(A - A_i) > \mathrm{Pl}(A_i)$, then $\mathrm{Pl}(A_0) > \mathrm{Pl}(A - A_0)$.

Recall that A2 states that if A_0, A_1, and A_2 are disjoint, $\mathrm{Pl}(A_0 \cup A_1) > \mathrm{Pl}(A_2)$, and $\mathrm{Pl}(A_0 \cup A_2) > \mathrm{Pl}(A_1)$, then $\mathrm{Pl}(A_0) > \mathrm{Pl}(A_1 \cup A_2)$. It is easy to check that for any finite number of sets, a similar property follows from A1 and A2 by induction. A2* asserts that a condition of this type holds even for an infinite collection of sets. This is not implied by A1 and A2. To see this, consider the plausibility model PL_{lot} from Example 5.3. Take A_0 to be empty and take A_i, $i > 1$, to be the singleton consisting of the world w_i. Then $\mathrm{Pl}_{lot}(A - A_i) = 1 > 1/2 = \mathrm{Pl}_{lot}(A_i)$, but $\mathrm{Pl}_{lot}(A_0) = 0 < 1 = \mathrm{Pl}(\cup_{i>0} A_i)$. Hence, A2* does not hold for plausibility structures in general. It does, however, hold for certain subclasses:

Proposition 5.4:: *A2* holds in every plausibility structure in* $\mathcal{P}^{p,w}_{subj}$ *and* $\mathcal{P}^{\kappa}_{subj}$.

Proof: See the appendix. ▮

A consequence of A2* is the following axiom, called ∀3 by Delgrande:

∀3. $\forall x(\varphi \rightarrow \psi) \Rightarrow (\varphi \rightarrow \forall x \psi)$ if x does not occur free in φ.

This axiom can be viewed as an infinitary version of axiom C2 (which is essentially KLM's And Rule), for (abusing notation somewhat) in a domain D, ∀3 essentially says:

$$\wedge_{d \in D}(\varphi \rightarrow \psi[x/d]) \Rightarrow (\varphi \rightarrow \wedge_{d \in D}\, \psi[x/d]).$$

Proposition 5.5:: ∀3 *is valid in all plausibility structures satisfying A2*.*

Proof: See the appendix. ▮

Since A2* holds in $\mathcal{P}^{p,w}_{subj}$ and $\mathcal{P}^{\kappa}_{subj}$, it follows that ∀3 does as well. Moreover, it is easy to see that the axiom ∀3 leads to the lottery paradox: From $\forall x(\mathit{true} \rightarrow \neg \mathit{Winner}(x))$, ∀3 allows us to conclude $\mathit{true} \rightarrow \forall x(\neg \mathit{Winner}(x))$.

A2* does not hold in $\mathcal{P}^{poss}_{subj}$ and $\mathcal{P}^{p}_{subj}$. In fact, the infinite lottery is consistent in these classes, although a somewhat unnatural model is required to express it, as the following example shows.

Example 5.6:: Consider the possibility structure $(Dom_{lot}, W_{lot}, \mathrm{Poss}, \pi_{lot})$, where all the components besides Poss are just as in the plausibility structure PL_{lot} from Example 5.3 and $\mathrm{Poss}(w_i) = i/(i+1)$. This means that if $i > j$, then it is more possible that individual i

wins than individual j. Moreover, this possibility approaches 1 as i increases. It is not hard to show that this possibility structure satisfies formulas (1) and (2).

A preferential structure in the same spirit also captures the lottery paradox. Consider the preferential structure $(Dom_{lot}, W_{lot}, \prec, \pi_{lot})$, where all the components besides $\prec$ are just as in the plausibility structure PL_{lot}, and we have $\ldots w_3 \prec w_2 \prec w_2 \prec w_1$. Thus, again we have that if $i > j$, then it is more likely that individual i wins than individual j. (More precisely, the world where individual i wins is preferred to that where individual j wins.) It is easy to verify that this preferential structure (which is obviously not well-founded) also satisfies *Lottery*. ∎

Although *Lottery* is satisfiable in possibility structures and preferential structures, a slight variant of it is not. Consider a *crooked lottery*, where there is one individual who is more likely to win than the rest, but is still unlikely to win. To formalize this in the language, we add the following formula that we call *Crooked*:

$$\exists y \forall x (x \neq y \Rightarrow ((\mathit{Winner}(x) \vee \mathit{Winner}(y)) \rightarrow \mathit{Winner}(y)))$$

This formula states that there is an individual who is more likely to win than the rest. To see this, recall that $(\varphi \vee \psi) \rightarrow \psi$ implies that either $\mathrm{Pl}([\![\varphi \vee \psi]\!]) = \perp$ (which cannot happen here because of the first clause of *Crooked*) or $\mathrm{Pl}([\![\varphi]\!]) < \mathrm{Pl}([\![\psi]\!])$. We take the crooked lottery to be formalized by the formula *Lottery* $\wedge$ *Crooked*.

It is easy to model the crooked lottery using plausibility. Consider the structure $PL'_{lot} = (Dom_{lot}, W_{lot}, \mathrm{Pl}'_{lot}, \pi_{lot})$, which is identical to PL_{lot} except for the plausibility measure Pl'_{lot}. We define $\mathrm{Pl}'_{lot}(w_1) = 3/4$; $\mathrm{Pl}'_{lot}(w_i) = 1/2$ for $i > 1$; $\mathrm{Pl}'_{lot}(A)$ of a finite set A is 3/4 if $w_1 \in A$, and 1/2 if $w_1 \notin A$; and $\mathrm{Pl}_{lot}(A) = 1$ for infinite A. It is easy to verify that PL'_{lot} satisfies *Crooked*, taking d_1 to be the special individual who is most likely to win (since $\mathrm{Pl}([\![\mathit{Winner}(d_1)]\!]) = 3/4 > 1/2 = \mathrm{Pl}([\![\mathit{Winner}(d_i)]\!])$ for $i > 1$). It is also easy to verify that $\mathrm{Pl}'_{lot} \models \mathit{Lottery}$.

On the other hand, the crooked lottery cannot be captured in $\mathcal{P}^{poss}_{subj}$ and $\mathcal{P}^{p}_{subj}$. To show this, we take a slight detour.

Consider the following two properties:

A2†. If $\{A_i : i \in I\}$ are pairwise disjoint sets, $A = \cup_{i \in I} A_i$, $0 \in I$, and for all $i \in I - \{0\}$, $\mathrm{Pl}(A_0) > \mathrm{Pl}(A_i)$, then $\mathrm{Pl}(A_0) \not< \mathrm{Pl}(A - A_0)$.

A3* If $\{A_i : i \in I\}$ are sets such that $\mathrm{Pl}(A_i) = \perp$, then $\mathrm{Pl}(\cup_i A_i) \perp$.

It is easy to see that A2† is implied by A2*. Suppose that Pl satisfies A2* and the preconditions of A2†. By A1 we have that $\mathrm{Pl}(A_0) > \mathrm{Pl}(A_i)$ implies that $\mathrm{Pl}(A - A_i) > \mathrm{Pl}(A_i)$. Thus by A2* we have that $\mathrm{Pl}(A_0) > \mathrm{Pl}(A - A_0)$, and therefore $\mathrm{Pl}(A_0) \not< \mathrm{Pl}(A - A_0)$. Moreover, A2† can hold in structures that do not satisfy A2*.

Proposition 5.7:: *A2† holds in every plausibility structure in $\mathcal{P}^{p}_{subj}$ and $\mathcal{P}^{poss}_{subj}$.*

Proof: See the appendix. ∎

A3* is an infinitary version of A3. It is easy to verify that it holds in all the approaches we consider, except plausibility measures and ϵ-semantics.

Proposition 5.8:: *A3* holds in every plausibility structure in* $\mathcal{P}^{p}_{subj}$, $\mathcal{P}^{p,w}_{subj}$, $\mathcal{P}^{\kappa}_{subj}$ *and* $\mathcal{P}^{poss}_{subj}$.

Proof: The proof is straightforward and left as an exercise to the reader. ∎

A3* has elegant axiomatic consequences.

Proposition 5.9:: *The axiom*

$$\forall x N\varphi \Rightarrow N(\forall x\varphi)$$

is sound in structures satisfying A3. Moreover, the axiom*

$$\forall x(\varphi \rightarrow \psi) \Rightarrow ((\exists x\varphi) \Rightarrow \psi), \quad \textit{if } x \textit{ does not appear free in } \psi$$

is sound in structures satisfying A2 and A3*.*[8]

Finally, we show that when A2† and A3* hold, the crooked lottery is (almost) inconsistent.

Proposition 5.10:: *The formula Lottery* $\wedge$ *Crooked* $\Rightarrow$ (*true*$\rightarrow$*false*) *is valid in structures satisfying A2† and A3*.*

Proof: See the appendix. ∎Notice that, since A2† and A3* are valid in $\mathcal{P}^{poss}_{subj}$, it immediately follows that *Lottery* $\wedge$ *Crooked* is unsatisfiable in $\mathcal{P}^{poss}_{subj}$.

To summarize, the discussion in this section shows that, once we move to first-order logic, kappa-rankings, possibility structures and preferential structures satisfy extra properties over and above those characterized by $\mathbf{C}^{subj}$ (and C6 and C7). We identified these properties both in terms of the constraints on the plausibility measures allowed by these semantics (e.g., conditions A2*, A2†, and A3*), and in terms of corresponding properties in the language (e.g., axioms and the variants of the lottery example). Our analysis leaves open the question of complete axiomatization of first-order conditional logic with respect to these classes of structures.

6 First-order statistical conditional logic

In the next three section, we analyze the statistical version of first-order conditional logic in much the same way we did the subjective version.

[8] The latter axiom can be viewed as an infinitary version of the OR Rule (C3), just as ∀3 can be viewed as an infinitary version of the AND Rule (C2).

The syntax for statistical conditionals is fairly straightforward. Let Φ be a first-order vocabulary, consisting of predicate and function symbols. (As usual, constant symbols are viewed as 0-ary function symbols.) Starting with atomic formulas of first-order logic, we form more complicated formulas by closing off under truth-functional connectives (i.e., $\wedge, \vee, \neg$, and $\Rightarrow$), first-order quantification, and the family of modal operators $\varphi \rightsquigarrow_X \psi$, where X is a set of distinct variables.[9] We denote the resulting language $\mathcal{L}^{stat}(\Phi)$. (We typically omit the Φ if it is clear from context.) The intuitive reading of $\varphi \rightsquigarrow_X \psi$ is that almost all of the X's that satisfy φ also satisfy ψ. Thus, the $\rightsquigarrow_X$ modality binds the variables X in φ and ψ, just as $\forall x$ binds the occurrences of x in $\forall x \varphi$. A typical formula in this language is $\exists y (P(x,y) \rightsquigarrow_x Q(x,y))$, which can be read "there is some y such that most x's satisfying $P(x,y)$ also satisfy $Q(x,y)$". Note that we allow arbitrary nesting of first-order and modal operators. For simplicity, we assume that all variables used in formulas come from the set $\{x_1, x_2, x_3, \ldots\}$.

To give semantics to $\mathcal{L}^{stat}(\Phi)$, we use *(first-order) statistical plausibility structures* (over Φ), which generalize the semantics of statistical probabilistic structures [Bacchus 1990; Halpern 1990] and statistical preferential structures [Brafman 1997]. Statistical plausibility structures over Φ are tuples of the form $PL = (Dom, \pi, \mathrm{Pl})$, where Dom is a domain, π is an interpretation assigning each predicate symbol and function symbol in Φ a predicate or function of the right arity over Dom, and Pl is a plausibility measure on Dom^∞ (a countable product of copies of Dom) that satisfies one restriction, described below. Note that we can identify Dom^∞ with the set of all valuations by associating a valuation v with an infinite sequence $(d_1, d_2, \ldots)$ of elements in Dom, where $v(x_i) = d_i$. Thus, we can view Pl as defining a plausibility measure on the space of valuations.

We require that Pl treats all variables uniformly, in the following sense:

- REN. If h is a *finite permutation* of the natural numbers (formally, $h : I\!N \rightarrow I\!N$ is a bijection such that $h(n) = n$ for all but finitely many elements $n \in I\!N$), then $\mathrm{Pl}(A^h) = \mathrm{Pl}(A)$ for all $A \subseteq Dom^\infty$, where $A^h = \{(d_{h(1)}, d_{h(2)}, d_{h(3)}, \ldots) : (d_1, d_2, d_3, \ldots) \in A\}$.

REN assures us, for example, that if $A \subseteq Dom$, then $\mathrm{Pl}(A \times Dom^\infty)$, the plausibility of a valuation giving x_1 a value in A, is the same as the plausibility of a valuation giving x_2 a value in A (i.e., $\mathrm{Pl}(Dom \times A \times Dom^\infty)$) and, in fact, the same as the plausibility of a valuation given x_k a value in A, for all k. As the name suggests, REN guarantees that we can rename variables, so that $\varphi \rightsquigarrow_X \psi$ will be equivalent to $\varphi[x/y] \rightsquigarrow_{X[x/y]} \psi[x/y]$ if y does not occur φ or ψ, where $X[x/y]$ is the result of replacing x in X by y (if $x \in X$; otherwise $X[x/y] = X$).

We may want to put a number of other restrictions on Pl, to make it act like a product measure, as Brafman [1997] does. While we believe such requirements may be quite reasonable, we do not make them here, to simplify the presentation. We discuss this issue further in Section 8.

[9] This syntax is borrowed from Brafman [1997], which in turn is based on that of [Bacchus 1990; Halpern 1990], except that in the earlier papers, the subscript X was taken to be a *sequence* of variables, rather than a set. Since the order of the variables is irrelevant, taking it to be a set seems more natural.

Given a statistical structure PL and a valuation v, we can associate with every formula φ a truth value in a straightforward way. For an atomic formula such as $P(x, \mathbf{c})$, we have

- $(PL, v) \models P(x, \mathbf{c})$ if $(v(x), \pi(\mathbf{c})) \in \pi(P)$.

Note that now we write $\pi(\mathbf{c})$ rather than $\pi(w)(\mathbf{c})$. We no longer have different worlds as we did in the subjective case. Thus, the issue of rigid vs. nonrigid designators does not arise in the statistical case.

We again treat quantification just as we do in first-order logic, so

- $(PL, v) \models \forall x\varphi$ iff $(PL, v') \models \varphi$ for all $v' \sim_x v$.

The interesting case, of course, comes in giving semantics to formulas of the form $\varphi \rightsquigarrow_X \psi$. In this case we have

- $(PL, v) \models \varphi \rightsquigarrow_X \psi$ if either $\mathrm{Pl}(v' : (PL, v') \models \varphi,\ v' \sim_X v\}) = \perp$ or $\mathrm{Pl}(\{v' : (PL, v') \models \varphi \wedge \psi,\ v' \sim_X v\}) > \mathrm{Pl}(\{v' : (PL, v') \models \varphi \wedge \neg\psi,\ v' \sim_X v\})$.

Again, we implicitly assume here that for each valuation v, vector X of variables, and formula φ, the set of valuations $\{v' : (PL, v') \models \varphi,\ v' \sim_X v\}$ is in $\mathcal{F}$, the domain of Pl.

As before, if φ is a sentence, then the truth of φ is independent of v; thus, we write $PL \models \varphi$ rather than $PL, v \models \varphi$.

7 Axiomatizing first-order statistical conditional logic

We can axiomatize first-order statistical conditional logic in much the same way as we did first-order subjective conditional logic. Again, we restrict attention to structures where the plausibility measures on Dom^∞ are qualitative; let $\mathcal{P}_{stat}^{QPL}$ be the class of all such structures.

Let $\mathbf{C}^{stat}$ consists of all generalizations of the following axioms, together with the inference rule MP. In the axioms, we write $\forall X\sigma$, where $X = \{x_1, \ldots, x_m\}$, as an abbreviation for $\forall x_1 \ldots \forall x_m \sigma$.

C0′. All instances of valid formulas of first-order logic with equality

C1′. $\varphi \rightsquigarrow_X \varphi$

C2′. $((\varphi \rightsquigarrow_X \psi_1) \wedge (\varphi \rightsquigarrow_X \psi_2)) \Rightarrow (\varphi \rightsquigarrow_X \psi_1 \wedge \psi_2)$

C3′. $((\varphi_1 \rightsquigarrow_X \psi) \wedge (\varphi_2 \rightsquigarrow_X \psi)) \Rightarrow ((\varphi_1 \vee \varphi_2) \rightsquigarrow_X \psi)$

C4′. $(\varphi_1 \rightsquigarrow_X \varphi_2 \wedge \varphi_1 \rightsquigarrow_X \psi) \Rightarrow \varphi_1 \wedge \varphi_2 \rightsquigarrow_X \psi$

R1′. $\forall X(\varphi_1 \Leftrightarrow \varphi_2) \Rightarrow ((\varphi_1 \rightsquigarrow_X \psi \Rightarrow (\varphi_2 \rightsquigarrow_X \psi)$

R2′. $\forall X(\psi_1 \Rightarrow \psi_2) \Rightarrow ((\varphi \rightsquigarrow_X \psi_2 \Rightarrow (\varphi \rightsquigarrow_X \psi_2)$

U. $\forall X\psi \Rightarrow (\varphi \rightsquigarrow_X \psi)$

Ren. $\varphi \rightsquigarrow_X \psi \Rightarrow \varphi[x/y] \rightsquigarrow_{X[x/y]} \psi[x/y]$, if y does not occur in φ or ψ.

As the notation suggests, C1′–C4′, R1′, and R2′ are the obvious analogues C1–C4, R1, and R2, except that R1′ and R2′ are now axioms rather than inference rules. C0′ subsumes C0 and F1–F5 in $\mathbf{C}^{subj}$. Note that we no longer need a special notion of substitutivity; there is only one world, and there are no concerns regarding the substitution of nonrigid terms into a modal context. For similar reasons, there is no analogue of F6 and F7 here. Ren and U are analogues of similar axioms for statistical probabilistic structures [Bacchus 1990; Halpern 1990]; here we need to require REN to ensure that Ren holds.

Theorem 7.1:: $\mathbf{C}^{stat}$ *is a sound and complete axiomatization of* $\mathcal{L}^{stat}$ *with respect to* $\mathcal{P}^{QPL}_{stat}$.

Proof: The basic idea similar to that of the proof of Theorem 4.1; indeed, the proof is even simpler. See the appendix for details. ∎

8 Other approaches to first-order statistical conditional logic

We have already remarked that we can construct "statistical" first-order analogues of all the approaches considered in the propositional case. We omit the formal definitions here. Let $\mathcal{P}^{p,w}_{stat}, \mathcal{P}^{p}_{stat}, \mathcal{P}^{\kappa}_{stat}, \mathcal{P}^{poss}_{stat}$, and $\mathcal{P}^{\epsilon}_{stat}$ be the subsets of $\mathcal{P}^{QPL}_{stat}$ that correspond to well-founded preference orderings, preference orderings, κ-rankings, possibility measures, and PPDs, respectively. The results are similar to those in Section 5, so we just sketch them here.

With κ-rankings and possibility measures, we need to require the obvious analogues of C6 and C7, namely

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

C7′ $\neg(\mathit{true} \rightsquigarrow_X \mathit{false})$

As we would expect, ϵ-semantics satisfies C7′ but not necessarily C6′.[10]

We have the following analogue of Theorem 5.2.

Theorem 8.1:: $\mathbf{C}^{stat}$+C7′ *is a sound and complete axiomatization of* $\mathcal{L}^{stat}$ *with respect to* $\mathcal{P}^{\epsilon}_{stat}$.

Proof: Follows from the proof of Theorem 7.1 using the same techniques as those used to prove Theorem 5.2 from Theorem 4.1. We omit further details here. ∎

Statistical plausibility structures based on well-founded preferential structures and κ-rankings also satisfy the following analogue of ∀3:

[10] Brafman [1997] discusses *pointed* PPDs, in which all the relevant limits are guaranteed to exist; for pointed PPDs, C6′ holds as well.

$\forall 3'$. $\forall y(\varphi \rightsquigarrow_X \psi) \Rightarrow (\varphi \rightsquigarrow_X \forall y \psi)$ if y does not occur free in φ or in X.

Interestingly, Brafman [1997] shows that $\mathbf{C}^{stat}$ together with C6′, C7′, and $\forall 3'$ is complete with respect to totally-ordered well-founded preferential structures.[11] These are essentially identical to the structures generated by κ-rankings. Thus, we have the following result.

Theorem 8.2:: [Brafman 1997] $\mathbf{C}^{stat} + \{$C6′, C7′, $\forall 3'\}$ *is a sound and complete axiomatization of* $\mathcal{L}^{stat}$ *with respect to* $\mathcal{P}^{\kappa}_{stat}$.

In light of Brafman's result, it seems likely that $\mathbf{C}^{stat}+\{$C6′, $\forall 3'\}$ is a sound and complete axiomatization of $\mathcal{L}^{stat}$ with respect to $\mathcal{P}^{p,w}_{stat}$, although we have not checked details.

Just as in the subjective case, $\forall 3'$ is not valid in statistical possibility structures or (non-well-founded) preferential structures, but a variant of the crooked lottery example does give us a valid formula for these structures too that does not follow from $\mathbf{C}^{stat}+$ {C6′, C7′}.

Up to now, we have put minimal structure on the plausibility measure on Dom^∞. In the case of statistical probability structures, the probability measure was assumed to be the product measure induced by a probability measure on *Dom*. We can make an analogous assumption in the case of ϵ-semantics, possibility measures, and κ-rankings. For example, if we start with a possibility measure Poss on *Dom*, we can define Poss^∞ on Dom^∞ by taking $\mathrm{Poss}^\infty(d_1, d_2, \ldots) = \inf_i \mathrm{Poss}(d_i)$, and taking $\mathrm{Poss}^\infty(A) = \sup_{\vec{d} \in A} \mathrm{Poss}^\infty(\vec{d})$ for $A \subseteq Dom^\infty$. A similar construction works for κ-rankings, except inf is replaced by + and sup is replaced by min. We get extra properties if we assume such a product measure construction, although the exact properties depend on the underlying notion of likelihood that we start with. For example, one property we get in all cases is the following:

- If $A, A' \subseteq Dom^n$ and $B \subseteq Dom^m$, $\vec{y} = \langle y_1, \ldots, y_n \rangle$, $\vec{z} = \langle z_1, \ldots, z_m \rangle$, $\vec{y}$ and $\vec{z}$ are disjoint, $\mathrm{Pl}(\{v : (v(y_1), \ldots, v(y_n)) \in A\}) \leq \mathrm{Pl}(\{v : (v(y_1), \ldots, v(y_n)) \in A'\})$ then $\mathrm{Pl}(\{v : (v(y_1), \ldots, v(y_n), v(z_1), \ldots, v(z_m)) \in A \times B\}) \leq \mathrm{Pl}(\{v : (v(y_1), \ldots, v(y_n), v(z_1), \ldots, v(z_m)) \in A' \times B\})$.

This property is captured by the axiom

$$\varphi \rightsquigarrow_X \psi \Rightarrow \varphi \wedge \varphi' \rightsquigarrow_X \psi,$$

where the set of variables free in φ' is disjoint from the set of variables free in $\varphi \wedge \psi$.

Whether or not we assume that Pl is generated as a product measure somehow, once we have $\forall 3'$ as an axiom (or the closely related variant as in the crooked lottery example), we

[11] Actually, there are a number of minor differences between the framework we have presented and that of Brafman. For example, Brafman assumes that there is a separate order defined on Dom^n, for each finite n, rather than one order defined on Dom^∞. The two approaches are essentially equivalent—we could have used either one here. The connection to valuations is perhaps clearer when we consider Dom^∞. He also has the axiom $(\varphi \rightsquigarrow_X \psi) \Rightarrow (\exists x \varphi \Rightarrow \exists x(\varphi \wedge \psi))$ instead of C7′. It is not hard to show that these axioms are equivalent in the presence of all the other axioms.

get the problems in the statistical case similar to those we saw in the subjective case. For example, suppose $\forall 3'$ is valid. Consider the statement

$$\forall y(\mathit{true} \rightsquigarrow_x \neg \mathit{Married}(x, y)).$$

This states that for any individual y, most individuals are not married to y. This seems reasonable since each y is married to at most one individual, which clearly constitutes a small fraction of the population. $\forall 3'$ then gives us

$$\mathit{true} \rightsquigarrow_x \forall y \neg \mathit{Married}(x, y).$$

That is, most people are not married! This certainly does not seem to be a reasonable conclusion.

It is straightforward to construct similar examples for the statistical variants of the other approaches, again, with the exception of plausibility structures and ϵ-semantics. We note that these problems occur for precisely the same reasons they occur in the subjective case. In particular, $\forall 3'$ holds whenever the plausibility measure on Dom^∞ satisfies A2*.

This shows that, just as for the subjective case, we need the greater generality of plausibility measures and ϵ-semantics to correctly model first-order statistical reasoning about conditionals.

We observe that problems similar to the lottery paradox occur in the approach of Lehmann and Magidor [1990], which can be viewed as a hybrid of subjective and statistical conditionals based on on preferential structures. More precisely, rather than putting a preferential ordering on worlds or on valuations, they put an ordering on world-valuation pairs. While this greater flexibility allows them to avoid some problems associated with putting an order solely on worlds or on valuations, the fundamental difficulty still remains.

Finally, we observe that the approach of [Schlechta 1995], which is based on a novel representation of "large" subsets, is in the spirit of our notion of statistical defaults (although his language is somewhat less expressive than ours).

9 Discussion

We have considered a number of different approaches to ascribing semantics to both a subjective and statistical first-order logic of conditionals in a number of ways. Our analysis shows that, once we move to the first-order case, significant differences arise between approaches that were shown to be equivalent in the propositional case. This vindicates the intuition that there are significant differences between these approaches, which the propositional language is simply too weak to capture. The analysis also supports our choice of plausibility structures as the semantics for first-order conditional logic; it shows that, with the exception of ϵ-semantics, all the previous approaches have significant shortcomings, which manifest themselves in lottery-paradox type situations. Plausibility also lets us home in on what properties of an approach give us lead to an infinitary AND rule like $\forall 3$.

What does all this say about default reasoning? As we have argued, statements like "birds typically fly" should perhaps be thought of as statistical statements, and should thus be represented as $Bird(x) \rightsquigarrow_x Fly(x)$. Such a representation gives us a logic of defaults, in which statements such as "birds typically fly" and "birds typically do not fly" are inconsistent, as we would expect.

Of course, what we really want to do with such typicality statements is to draw default conclusions about individuals. Suppose we believe such a typicality statement. What other beliefs should follow? In general, $\forall x(Bird(x) \rightarrow Fly(x))$ does not follow; we should not necessarily believe that *all* birds are likely to fly. We may well know that Tacky the penguin [Lester 1988] does not fly . As long as Tacky is a rigid designator, this is simply inconsistent with believing that all birds are likely to fly. In the absence of information about any particular bird, $\forall x(Bird(x) \rightarrow Fly(x))$ may well be a reasonable belief to hold. Moreover, no matter what we know about exceptional birds, it seems reasonable to believe $true \rightsquigarrow_x (Bird(x) \rightarrow Fly(x))$: almost all birds are likely to fly (assuming we have a logic that allows the obvious combination of statistical and subjective plausibility).

Unfortunately, we do not have a general approach that will let us go from believing that birds typically fly to believing that almost all birds are likely to fly. Nor do we have an approach that allows us to conclude that Tweety is likely to fly given that birds typically fly and Tweety is a bird (and that we know nothing else about Tweety). These issues were addressed in the first-order setting by both Lehmann and Magidor [1990] and Delgrande [1988]. The key feature of their approaches, as well as other propositional approaches rests upon getting a suitable notion of irrelevance. While we also do not have a general solution to the problem of irrelevance, we believe that plausibility structures give us the tools to study it in an abstract setting. We suspect that many of the intuitions behind probabilistic approaches that allow us to cope with irrelevance [Bacchus, Grove, Halpern, and Koller 1996; Koller and Halpern 1996] can also be brought to bear here. We hope to return to this issue in future work.

A Proofs

Theorem 4.1: *$\mathbf{C}^{subj}$ is a sound and complete axiomatization of $\mathcal{L}^{subj}$ with respect to $\mathcal{P}^{QPL}_{subj}$.*

Proof: A formula φ is said to be *consistent* with $\mathbf{C}^{subj}$ if $\mathbf{C}^{subj} \not\vdash \neg\varphi$. A finite set of formulas $\{\sigma_1, \ldots, \sigma_k\}$ is consistent with $\mathbf{C}^{subj}$ if their conjunction $\sigma_1 \wedge \ldots \wedge \sigma_k$ is consistent with $\mathbf{C}^{subj}$. An infinite set Σ of formulas is consistent with $\mathbf{C}^{subj}$ if every finite subset of Σ is consistent with $\mathbf{C}^{subj}$. Finally, a set Σ is said to be a *maximal consistent set of sentences* if (1) it consists only of sentences (recall that a sentence is a formula with no free variables), (2) it is consistent and (3) no strict superset of Σ consisting only of sentences is consistent. In the discussion below, all maximal consistent sets are maximal consistent sets of sentences; however, the other consistent sets we construct may include formulas that are not sentences.

Our goal is to show that a formula φ is consistent with $\mathbf{C}^{subj}$ iff it is satisfiable in a first-order plausibility structure. As usual, this clearly suffices to prove completeness. We can also assume without loss of generality that φ is a sentence, for using standard arguments of first order logic (see [Enderton 1972, p. 109]) we can show that if $y_1, \ldots, y_m$ are the free variables in φ, then φ is provable iff its universal closure $\forall y_1 \ldots \forall y_m \varphi$ is provable, and hence φ is consistent iff $\exists y_1 \ldots \exists y_m \varphi$ is consistent.

Let $\mathcal{C}$ be a countable set of constant symbols not in Φ, let Φ' consist of the symbols in Φ that actually appear in φ, and let $\Phi^+ = \Phi' \cup \mathcal{C}$.[12] As usual in Henkin-style completeness proofs, we construct a structure satisfying φ using maximal consistent subsets of $\mathbf{C}^{subj}$ (in the language $\mathcal{L}^{subj}(\Phi^+)$).

A maximal consistent subset A of $\mathbf{C}^{subj}$ is said to be $\mathcal{C}$-*good* if (1) $\neg\forall x\psi \in A$ implies $\neg\psi[x/\mathbf{c}] \in A$ for some $\mathbf{c} \in \mathcal{C}$ and (2) $\forall x\psi \in A$ implies $\psi[x/\mathbf{c}] \in A$ for all $\mathbf{c} \in \mathcal{C}$. Note that property (2) holds automatically for maximal consistent sets in first-order logic, but does not hold in general in our logic, because of our restriction on F1. Intuitively, A is $\mathcal{C}$-good if the constants in $\mathcal{C}$ are rigid designators such that every domain element is the interpretation of some constant in $\mathcal{C}$.

The proof now proceeds according to the following steps:

1. We show that there is a $\mathcal{C}$-good maximal consistent set C^* that includes φ. This follows closely the standard Henkin-style completeness proof for first-order logic [Enderton 1972].

2. We construct a structure PL by using the formulas in C^*. This step uses techniques from [Enderton 1972] for defining the domain, and from [Friedman and Halpern 1998] for defining the set of possible worlds and the plausibility measure over them

3. We show that $\mathrm{Pl} \models \varphi$. Again, this argument is in the spirit of the standard Henkin-style completeness for first-order logic.

For the first step, we proceed as follows. Let $\sigma_0, \sigma_1, \ldots$ be an enumeration of the formulas in the language $\mathcal{L}^{subj}(\Phi^+)$. We inductively construct a sequence $A_0, A_1, \ldots$ of finite sets of formulas such that A_n is consistent. Let $A_0 = \{\varphi\}$. Let A_{k+1} consist of A_k together with the formula $\neg\forall x\sigma_{k+1} \Rightarrow \neg\sigma_{k+1}[x/\mathbf{c}]$, where $\mathbf{c}$ is a constant in $\mathcal{C}$ that does not appear in any of the formulas in A_k or in σ_{k+1}. (This is possible since A_k is finite.) Intuitively, $\neg\forall x\sigma_{k+1} \Rightarrow \neg\sigma_{k+1}[x/\mathbf{c}]$ says that $\mathbf{c}$ provides a witness to the fact that $\neg\sigma_{k+1}$ does not hold for all x, if such a witness is necessary.

We claim that $A^* = \cup_k A_k$ is consistent. This follows from the following somewhat more general lemma.

Lemma A.1:: *If B_0 is a finite consistent set of formulas, and for $k \geq 0$, $B_{k+1} = B_k \cup \{\neg\forall x\sigma \Rightarrow \sigma[x/\mathbf{c}]\}$ for some formula σ and constant $\mathbf{c}$ that does not appear in B_k or σ, then $\cup_k B_k$ is consistent.*

[12] This proof would go through without change if we took Φ^+ to be $\Phi \cup \mathcal{C}$. However, for the proof of Theorem 5.2, it is useful to restriction attention to a language that is guaranteed to be countable.

Proof: From the definition of consistency, it clearly suffices to prove that B_k is consistent for all $k \geq 0$. We do this by induction on k. By assumption B_0 is consistent. Suppose B_k is consistent but B_{k+1} is not. Suppose $B_{k+1} = B_k \cup \{\neg\forall x\sigma \Rightarrow \neg\sigma[x/\mathbf{c}]\}$. Identifying B_k with the conjunction of formulas in B_k, it then follows that

$$\mathbf{C}^{subj} \vdash B_k \Rightarrow (\neg\forall x\sigma \wedge \sigma[x/\mathbf{c}]).$$

Since $\mathbf{c}$ does not appear in any of the formulas in B_k or σ, a standard argument from first-order logic (see [Enderton 1972, p. 116]) can be used to show that

$$\mathbf{C}^{subj} \vdash B_k \Rightarrow (\neg\forall x\sigma \wedge \forall x\sigma),$$

contradicting the consistency of B_k. ∎

Let B^* consist of all the formulas in A^* together with all the formulas of the form $\forall x\sigma \Rightarrow \sigma[x/\mathbf{c}]$, where σ is a formula in $\mathcal{L}^{subj}$ and $\mathbf{c} \in \mathcal{C}$. We claim that B^* is consistent. For suppose not. Then there exists a finite set of formulas in A^*, say A', and a finite set of formulas $B' \subseteq \mathcal{L}^{subj}$, and a finite set of constants $\mathcal{C}' \subseteq \mathcal{C}$ such that

$$\mathbf{C}^{subj} \vdash A' \Rightarrow ((\wedge_{\sigma\in B'}\forall x\sigma) \wedge (\vee_{\sigma\in B',\mathbf{c}\in\mathcal{C}'}\neg\sigma[x/\mathbf{c}])).$$

Suppose $\mathcal{C}' = \{\mathbf{c}_1, \ldots, \mathbf{c}_k\}$. Let $y_1, \ldots, y_k$ be fresh variables, that do not appear in the formulas in A' or B'. Let A'' and B'' be the result of replacing all occurrences of $\mathbf{c}_1, \ldots, \mathbf{c}_k$ in the formulas in A' and B', respectively, by the variables $y_1, \ldots, y_k$. Again, using standard techniques [Enderton 1972, p. 116], we have

$$\mathbf{C}^{subj} \vdash A'' \Rightarrow ((\wedge_{\sigma\in B''}\forall x\sigma) \wedge (\vee_{\sigma\in B'',y\in\{y_1,\ldots,y_k\}}\neg\sigma[x/y])).$$

It follows from F1 that $(\wedge_{\sigma\in B''}\forall x\sigma) \wedge (\vee_{\sigma\in B'',y\in\{y_1,\ldots,y_k\}}\neg\sigma[x/y])$ is inconsistent. Thus, A'' must be inconsistent. But it follows from Lemma A.1 that A'' is consistent. This gives us the desired contradiction.

Let $B^\dagger$ consist of all the sentences in B^*. Using standard techniques, we can extend $B^\dagger$ to a maximal consistent set of sentences: We construct a sequence $C_0, C_1, \ldots$ of consistent sets of sentences by taking $C_0 = B^\dagger$ and $C_{k+1} = C_k \cup \{\sigma_k\}$ if $C_k \cup \{\sigma_k\}$ is consistent and σ_k is a sentence, and $C_{k+1} = C_k$ otherwise. Let $C^* = \cup_{k=1}^{\infty} C_k$. C^* is then easily seen to be a maximal consistent set of sentences. Moreover, our construction guarantees that C^* is $\mathcal{C}$-good and contains φ. This completes Step 1 of the proof.

We now proceed to the second step of the proof, where we construct a first-order subjective plausibility structure based on C^*. First, however, we need two more definitions in order to allow us to characterize the domain and the set of possible worlds in our desired plausibility structure.

- We define an equivalence relation $\sim$ on $\mathcal{C}$ by defining $\mathbf{c} \sim \mathbf{c}'$ if $\mathbf{c} = \mathbf{c}' \in C^*$. Let $[\mathbf{c}] = \{\mathbf{c}' : \mathbf{c} \sim \mathbf{c}\}$. As we shall see, these equivalence classes will be the domain elements in our structure.

- If A is a set of formulas, define $A/N = \{\psi : N\psi \in A\}$. The worlds in our structure will be all $\mathcal{C}$-good maximal consistent sets A of sentences such that $A/N = C^*/N$.

We want to ensure that global properties, such as equality of domain elements and conditional statements, are true in all worlds in the structure. As we shall see, our construction is such that formulas in C^*/N are true in all these worlds. Thus, we need the following lemma.

Lemma A.2:: *If ψ is of the form $\psi' \rightarrow \psi''$ or $\mathbf{c} = \mathbf{c}'$, for $\mathbf{c}, \mathbf{c}' \in \mathcal{C}$ and A is a $\mathcal{C}$-good maximal consistent set, $\psi \in A$ iff $\psi \in A/N$.*

Proof: Suppose A is a $\mathcal{C}$-good maximal consistent set of formulas and ψ is of the form $\psi' \rightarrow \psi''$. If $\psi \in A$, then $N\psi \in A$ by C5, so $\psi \in A/N$. Conversely, if $\psi \in A/N$, then $N\psi \in A$. Suppose, by way of contradiction, that $\psi \notin A$. Since A is a maximal consistent set, we must have $\neg\psi \in A$. By C5, it follows that $N\neg\psi \in A$. That is, both $\psi \rightarrow \mathit{false}$ and $\neg\psi \rightarrow \mathit{false}$ are in A. By C3, it follows that $\mathit{true} \rightarrow \mathit{false} \in A$. By RW, we have that both $\mathit{true} \rightarrow \psi'$ and $\mathit{true} \rightarrow \psi''$ are in A. By C4, we have that $\psi = \psi' \rightarrow \psi'' \in A$, contradicting our assumption.

Now suppose ψ is of the form $\mathbf{c} = \mathbf{c}'$. By F6, we have that $\forall x, y(x = y \Rightarrow N(x = y)) \in A$. Since A is $\mathcal{C}$-good, it follows that $\mathbf{c} = \mathbf{c}' \Rightarrow N(\mathbf{c} = \mathbf{c}') \in A$. Similarly, by F7, we have that $\mathbf{c} \neq \mathbf{c}' \Rightarrow N(\mathbf{c} \neq \mathbf{c}') \in A$. We can now show that $\psi \in A$ iff $\psi \in A/N$ just as we did in the case of $\psi' \rightarrow \psi''$, replacing the use of C5 by these consequences of F6 and F7. ∎

We construct a first-order subjective plausibility structure $PL = (\mathit{Dom}, W, \mathrm{Pl}, \pi)$ as follows:

- $\mathit{Dom} = \{[\mathbf{c}] : \mathbf{c} \in \mathcal{C}\}$
- $W = \{w : w$ is a $\mathcal{C}$-good maximal consistent set of sentences and $w/N = C^*PBox$ (i.e., $N\psi \in w$ iff $N\psi \in C^*$) $\}$
- π is defined so that $\pi(w)(\mathbf{c}) = [\mathbf{c}]$ for the constants $\mathbf{c} \in \mathcal{C}$ and $\pi(w)$ for the symbols in Φ' is determined by the atomic sentences in w in the obvious way (see [Enderton 1972, p. 131]). For example, if P is a binary predicate, then $([\mathbf{c}], [\mathbf{c}']) \in \pi(w)(P)$ iff $P(\mathbf{c}, \mathbf{c}')$ is one of the formulas in w.
- Pl is defined so that for all formulas ψ and ψ', we have $\mathrm{Pl}([\psi]) \leq \mathrm{Pl}([\psi'])$ iff $(\psi \vee \psi') \rightarrow \psi' \in C^*$, where $[\sigma] = \{w : \sigma \in w\}$.

In their completeness proof for propositional conditional logic, Friedman and Halpern [1998] define a plausibility measure Pl in a similar way and show that it is well defined (i.e., if $[\psi] = [\psi']$ and $[\sigma] = [\sigma']$, then $(\psi \vee \sigma) \rightarrow \sigma \in C^*$ iff $(\psi' \vee \sigma') \rightarrow \sigma' \in C^*$) and that it satisfies A1, A2, and A3, so that Pl is a qualitative plausibility measure. (See the proof of Theorem 8.2 in [Friedman and Halpern 1998].) The same proof applies here without change, so we do not repeat it.

Finally, we move to the third and last step of the proof: showing that PL satisfy φ. To do so we prove the standard *truth lemma*, namely

(∗) $(PL, w) \models \psi$ iff $\psi \in w$, for all $w \in W$.

We prove (*) by a straightforward induction on the depth of nesting of $\rightarrow$ in ψ, with a subinduction on structure.

If ψ is an atomic formula of the form $\mathbf{c} = \mathbf{c}'$, this follows from Lemma A.2 and the definition of π. For other atomic formulas, this is immediate from the definition of π.

If ψ is a conjunction or a negation, it is immediate from the induction hypothesis.

If ψ has the form $\psi' \rightarrow \psi''$, then we have

$$
\begin{array}{lll}
 & (PL, w) \models \psi' \rightarrow \psi'' & \\
\text{iff} & \mathrm{Pl}([\![\psi' \wedge \psi'']\!]) > \mathrm{Pl}([\![\psi' \wedge \neg\psi'']\!]) \text{ or } \mathrm{Pl}([\![\psi']\!]) = \bot & \\
\text{iff} & \mathrm{Pl}([\psi' \wedge \psi'']) > \mathrm{Pl}([\psi' \wedge \neg\psi'']) \text{ or } \mathrm{Pl}([\psi']) = \bot & \text{[induction hypothesis]} \\
\text{iff} & ((\psi' \wedge \psi'') \vee (\psi' \wedge \neg\psi'')) \rightarrow (\psi' \wedge \psi'') \in C^* \text{ or } \psi' \rightarrow \mathit{false} \in C^* & \text{[by definition of Pl]} \\
\text{iff} & \psi' \rightarrow \psi'' \in C^* \text{ or } \psi' \rightarrow \mathit{false} \in C^* & \text{[by LLE,C1,C2]} \\
\text{iff} & \psi' \rightarrow \psi'' \in C^* & \text{[by RW]} \\
\text{iff} & N(\psi' \rightarrow \psi'') \in C^* & \text{[by Lemma A.2]} \\
\text{iff} & N(\psi' \rightarrow \psi'') \in w & \text{[since } w \in W\text{]} \\
\text{iff} & \psi' \rightarrow \psi'' \in w & \text{[by Lemma A.2]}
\end{array}
$$

Finally, if ψ has the form $\forall x \psi'$, then we have

$$
\begin{array}{lll}
 & (PL, w) \models \forall x \psi & \\
\text{iff} & (PL, w) \models \psi[x/\mathbf{c}] \text{ for all } \mathbf{c} \in \mathcal{C} & \\
\text{iff} & \psi[x/\mathbf{c}] \in w \text{ for all } \mathbf{c} \in \mathcal{C} & \text{[induction hypothesis]} \\
\text{iff} & \forall x \psi \in w & \text{[since } w \text{ is } \mathcal{C}\text{-good]}
\end{array}
$$

This completes the proof of (∗). It now follows that $(PL, C^*) \models \varphi$. We can easily get from this a structure over the vocabulary Φ that satisfies the formula φ: we simply define π in an arbitrary way for the symbols in $\Phi - \Phi'$, and ignore the interpretation of the symbols in $\mathcal{C}$. This completes the proof of the theorem. ∎

Theorem 5.2: $\mathbf{C}^{subj}$+C7 *is a sound and complete axiomatization of* $\mathcal{L}^{subj}$ *with respect to* $\mathcal{P}^{\epsilon}_{subj}$.

Proof: As in the proof of Theorem 4.1, it suffices to show that if φ is consistent with $\mathbf{C}^{subj}$+C7, then it is satisfiable in a structure in $\mathcal{P}^{\epsilon}_{subj}$. The first steps in the proof mimic those of the proof of Theorem 4.1. We define $\Phi^+ = \Phi' \cup \mathcal{C}$ as in the proof of Theorem 4.1. We then construct a plausibility structure $PL = (D, W, \mathrm{Pl}, \pi)$ satisfying φ by considering maximal ($\mathbf{C}^{subj}$+C7)–consistent sets of sentences. It is easy to see that the structure PL is what is called in [Friedman and Halpern 1998] (following [Lewis 1973]) *normal*: we must have $\mathrm{Pl}(W) > \bot$ (otherwise C7 would not be valid in PL). By Theorem 6.3 in [Friedman and Halpern 1998], it follows that there is a PPD PP on W that satisfies the same defaults. More precisely, if $PL_{pp} = (D, W, \mathrm{Pl}_{PP}, \pi)$, where Pl_{PP} is the plausibility measure corresponding to PP, as described in Section 2, then we have $(PL, w) \models \psi \rightarrow \psi'$ iff $(PL_{PP}, w) \models \psi \rightarrow \psi'$.[13] A straightforward induction on the structure of formulas now shows that PL and PL_{PP} agree on all sentences in $\mathcal{L}^{subj}(\Phi^+)$. This gives us the desired PPD structure satisfying φ, and completes the proof. ∎

Proposition 5.4: *A2* holds in every plausibility structure in* $\mathcal{P}^{p,w}_{subj}$ *and* $\mathcal{P}^{\kappa}_{subj}$.

Proof: We start with $\mathcal{P}^{\kappa}_{subj}$. Suppose $PL = (D, W, \mathrm{Pl}_\kappa, \pi) \in \mathcal{P}^{\kappa}_{subj}$, where κ is the ranking to which Pl_κ corresponds. Since lower ranks correspond to greater plausibility, we have $\kappa(A) \leq \kappa(B)$ iff $\mathrm{Pl}_\kappa(A) \geq \mathrm{Pl}_\kappa(B)$. Let $\{A_i : i \in I\}$ be a collection of pairwise disjoint sets such that $\kappa(A - A_i) < \kappa(A_i)$ for all $i \in I - \{0\}$. We claim that (1) $\kappa(A) = \kappa(A_0)$, (2) $\kappa(A) < \kappa(A_i)$ for $i \in I - \{0\}$, and (3) $\kappa(A_0) < \kappa(\cup_{i \in I - \{0\}} A_i)$. (2) follows immediately from the assumption that $\kappa(A - A_i) < \kappa(A_i)$, since $\kappa(A) \leq \kappa(A - A_i)$. (1) follows from (2) and the observations that (a) $\kappa(A) = \min_{i \in I} \kappa(A_i)$ and (b) the range of κ is the natural numbers. Finally, (3) follows from (1) and (2) and the observation that $\kappa(\cup_{i \in I - \{0\}} A_i) = \min_{i \in I - \{0\}} \kappa(A_i)$. From (1), (2), and (3), it is immediate that $\kappa(A_0) < \kappa(A - A_0)$.

The argument in the case of $\mathcal{P}^{p,w}_{subj}$ is similar in spirit. Suppose $PL = (D, W, \mathrm{Pl}_\prec, \pi) \in \mathcal{P}^{p,w}_{subj}$, where $\mathrm{Pl}_\prec$ is constructed from a partial order $\prec$ on W as described in Section 2. Again, let $\{A_i : i \in I\}$ be a collection of pairwise disjoint sets such that $\mathrm{Pl}_\prec(A - A_i) < \mathrm{Pl}_\prec(A_i)$ for all $i \in I - \{0\}$. Recall that $\mathrm{Pl}_\prec(A) \leq \mathrm{Pl}_\prec(B)$ 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 $w'' \prec w'$. Thus, to show that $\mathrm{Pl}_\prec(A_0) \geq \mathrm{Pl}_\prec(A - A_0)$, we must show that if $w \in A - A_0$, then there exists some $w' \in A_0$ such that $w' \prec w$ and there is no $w'' \in A - A_0$ such that $w'' \prec w'$. Suppose not. Then we construct an infinite decreasing sequence $\ldots w_k \prec w_{k-1} \prec \cdots \prec w_0$, contradicting the assumption that $\prec$ is well founded. We proceed as follows. Let $w_0 = w$. Suppose inductively we have constructed $w_0, \ldots, w_k$. If $w_k \in A_0$, by assumption, there is some $w_{k+1} \in A - A_0$ such that $w_{k+1} \prec w_k$. If $w_k \in A - A_0$, then $w_k \in A_i$ for some $i \neq 0$. Since $\mathrm{Pl}_\prec(A - A_i) > \mathrm{Pl}_\prec(A_i)$, it follows from the construction of $\mathrm{Pl}_\prec$ that there must be some $w_{k+1} \in A - A_i$ such that $w_{k+1} \prec w_k$. This completes the inductive proof, and gives us the desired contradiction. It is easy to see that because A_0 and $A - A_0$ are disjoint, the fact that $\mathrm{Pl}_\prec(A_0) \geq \mathrm{Pl}_\prec(A - A_0)$ implies that $\mathrm{Pl}_\prec(A_0) > \mathrm{Pl}_\prec(A - A_0)$, as desired. ∎

Proposition 5.5: $\forall 3$ *is valid in all plausibility structures satisfying A2*.*

Proof: Suppose $PL = (D, W, \mathrm{Pl}, \pi)$ satisfies A2* and $(PL, w, v) \models \forall x(\varphi \rightarrow \psi)$, where x does not appear free in φ. It follows that

$$[\![\varphi \wedge \psi]\!]_{PL,w,v'} > [\![\varphi \wedge \neg\psi]\!]_{PL,w,v'}, \text{ for all valuations } v' \text{ such that } v' \sim_x v. \qquad (6)$$

[13]Theorem 6.3 in [Friedman and Halpern 1998] applies only to countable languages, so it is important that we use Φ^+ here, rather than $\Phi \cup \mathcal{C}$, which may not be countable.

Let $A = [\![\varphi]\!]_{PL,w,v}$. (Note that we have $A = [\![\varphi]\!]_{PL,w,v'}$ for all v' such that $v' \sim_x v$, since x is not free in φ.) For each $d \in D$, let $A_d = [\![\varphi \wedge \neg\psi]\!]_{PL,w,v_d}$, where $v_d \sim_x v$ and $v_d(x) = d$. Let $A' = [\![\varphi \wedge \forall x\psi]\!]$. Note that $A = A' \cup (\cup_{d\in D} A_d)$ and that A' is disjoint from A_d, for each $d \in D$. The sets A_d are not necessarily disjoint. Thus, let B_d be such that $B_d \subseteq A_d$, the sets B_d are pairwise disjoint, and $\cup_{d\in D} B_d = \cup_{d\in D} A_d$. (We can always find such sets B_d. If D is countable, say $D = \{1, 2, 3, \ldots\}$ without loss of generality, then we can take $B_1 = A_1$ and $B_{k+1} = A_{k+1} - (B_1 \cup \ldots B_k)$. If D is uncountable, we must first well-order D; then a similar inductive construction works.) Thus, we have $A = A' \cup (\cup_{d\in D} B_d)$, and all the sets on the right-hand side are pairwise disjoint.

From (6), it follows that $\mathrm{Pl}(A - A_d) > \mathrm{Pl}(A_d)$, for all $d \in D$, so clearly $\mathrm{Pl}(A - B_d) > \mathrm{Pl}(B_d)$, since $B_d \subseteq A_d$. From A2*, it follows that $\mathrm{Pl}(A') > \mathrm{Pl}(A - A')$. Thus, $(PL, w, v) \models \varphi \rightarrow \forall x\psi$, as desired. ∎

Proposition 5.7: *A2† holds in every plausibility structure in $\mathcal{P}^{p}_{subj}$ and $\mathcal{P}^{poss}_{subj}$.*

Proof: We start with $\mathcal{P}^{poss}_{subj}$. Suppose that $POSS = (D, W, \mathrm{Poss}, \pi)$ is a possibility structure. Let $\{A_i : i \in I\}$ be a collection of pairwise disjoint sets such that $\mathrm{Poss}(A_0) > \mathrm{Poss}(A_i)$ for all $i \in I - \{0\}$. This implies that $\mathrm{Poss}(A_0) \geq \sup_{i\in I-\{0\}} \mathrm{Poss}(A_i) = \mathrm{Poss}(A - A_0)$. We immediately get that $\mathrm{Poss}(A_0) \not< \mathrm{Poss}(A - A_0)$.

The argument in the case of $\mathcal{P}^{p}_{subj}$ is similar in spirit, although somewhat more involved. Suppose $PL = (D, W, \mathrm{Pl}_\prec, \pi) \in \mathcal{P}^{p}_{subj}$, where $\mathrm{Pl}_\prec$ is constructed from a partial order $\prec$ on W as described in Section 2. Again, let $\{A_i : i \in I\}$ be a collection of pairwise disjoint sets such that $\mathrm{Pl}_\prec(A_0) > \mathrm{Pl}_\prec(A_i)$ for all $i \in I - \{0\}$.

By way of contradiction, suppose that $\mathrm{Pl}_\prec(A_0) < \mathrm{Pl}_\prec(A - A_0)$. Let $w \in A_0$. Since $\mathrm{Pl}_\prec(A_0) < \mathrm{Pl}_\prec(A - A_0)$, there is a world $w' \in A - A_0$ such that $w' \prec w$ and there is no $w'' \in A_0$ such that $w'' \prec w'$. Let A_i be the set that contains w'. (There must be such an index, since $A - A_0$ is the union of such sets.) Since $\mathrm{Pl}_\prec(A_i) < \mathrm{Pl}_\prec(A_0)$, there is a world $w'' \in A_0$ such that $w'' \prec w'$. Thus, we get a contradiction. We conclude that $\mathrm{Pl}_\prec(A_0) \not< \mathrm{Pl}_\prec(A - A_0)$. ∎

Proposition 5.9: *The axiom*

$$\forall x N\varphi \Rightarrow N(\forall x\varphi)$$

is sound in structures satisfying A3. Moreover, the axiom*

$$\forall x(\varphi \rightarrow \psi) \Rightarrow ((\exists x\varphi) \Rightarrow \psi), \quad \textit{if } x \textit{ does not appear free in } \psi$$

is sound in structures satisfying A2 and A3*.*

Proof: For the first part of the proposition, suppose that $PL = (D, W, \mathrm{Pl}, \pi)$ be a plausibility structure satisfying A3*. Assume that there is a world $w \in W$ and a valuation v such that

$$(PL, w, v) \models \forall x N\varphi.$$

This means that if $v' \sim_x v$, then

$$(PL, w, v') \models N\varphi \tag{7}$$

For each $d \in D$, let v_d be the valuation such that $v_d \sim_x v$ and $v_d(x) = d$. Let $A_d = \{w' : (PL, w', v_d) \models \varphi\}$. From (7), we have that $\mathrm{Pl}(W - A_d) = \perp$ for all $d \in D$. Using A3*, we get that $PL(W - (\cap_d A_d)) = \perp$. Thus,

$$(PL, w, v) \models N\forall x\varphi$$

as desired.

For the second part of the proposition, suppose that $PL = (D, W, \mathrm{Pl}, \pi)$ be a plausibility structure satisfying A2* and A3*. Assume that there is a world $w \in W$ and a valuation v such that

$$(PL, w, v) \models \forall x(\varphi \rightarrow \psi),$$

where x does not appear free in φ. This means that if $v' \sim_x v$, then

$$(PL, w, v') \models \varphi \rightarrow \psi. \tag{8}$$

Again, for each $d \in D$, let v_d be the valuation such that $v_d \sim_x v$ and $v_d(x) = d$. Let $A_d = \{w' : (PL, w', v_d) \models \varphi\}$ and let $B = \{w' : (PL, w', v) \models \psi\}$.

If $\mathrm{Pl}(A_d) = \perp$ for all d, then by A3* we get that $\mathrm{Pl}(\cup_d A_d) = \perp$. In this case $(PL, w, v) \models (\forall x\varphi) \rightarrow \psi$ vacuously.

On the other hand, if $\mathrm{Pl}(A_d) \neq \perp$ for all $d \in D$, then we must have $\mathrm{Pl}(\cup_{d \in D} A_d) > \perp$. By (8), for each $d' \in D$, we have that either $\mathrm{Pl}(A_d) = \perp$ or $\mathrm{Pl}(A_d \cap B) > \mathrm{Pl}(A_d - B)$. (Note that since x is not free in ψ, we have that $(PL, w', v) \models \psi$ iff $(PL, w', v_d) \models \psi$.) In either case, we can conclude that

$$\mathrm{Pl}((\cup_d A_d) \cap B) > \mathrm{Pl}(A_{d'} - B).$$

Let A'_d be pairwise disjoint sets such that $A'_d \subseteq A_d$, and $\cup_d A'_d = A_d$. (Such sets must exist; see the proof of Proposition 5.5.) Thus, we have that $\mathrm{Pl}((\cup_d A_d) \cap B) > \mathrm{Pl}(A'_d - B)$. Using A2*, we get that

$$\mathrm{Pl}((\cup_d A_d) \cap B) > \mathrm{Pl}(\cup_d A'_d - B) = \mathrm{Pl}(\cup_d A_d - B).$$

We conclude that

$$(PL, w, v) \models (\forall x\varphi) \rightarrow \psi,$$

as desired. ∎

Proposition 5.10: *The formula Lottery* $\wedge$ *Crooked* $\Rightarrow$ *(true*$\rightarrow$*false) true in structures satisfying A2*† *and A3*.*

Proof: Let $PL = (D, W, \mathrm{Pl}, \pi)$ be a plausibility structure satisfying A2† and A3*. Suppose also that $PL \models$ *Lottery* $\wedge$ *Crooked*. Since

$$PL \models \exists y \forall x(x \neq y \Rightarrow ((Winner(x) \vee Winner(y)) \rightarrow Winner(y))),$$

there must be some domain element $d_0 \in D$ and valuation v such that $v(y) = d_0$ and

$$(PL, v) \models \forall x(x \neq y \Rightarrow ((Winner(x) \vee Winner(y)) \rightarrow Winner(y))).$$

This in turn means that if $v' \sim_x v$, then

$$(PL, w, v') \models (Winner(x) \vee Winner(y)) \rightarrow Winner(y). \tag{9}$$

For each $d \in D$, v_d be such that $v_d \sim_x v$ and $v_d(x) = d$. Let $A_d = \{w' \in W : (PL, w, v_d \models Winner(x)\}$, let $A = \cup_d A_d$, and let $B = W - A$.

It is immediate from (9) that either

(a) $\mathrm{Pl}(A_{d_0}) = \mathrm{Pl}(A_{d_0} - A_d) = \bot$ for all $d \neq d_0$, or

(b) $PL(A_{d_0}) > PL(A_d - A_{d_0})$ for all $d \neg d_0$.

Assume that (a) is true. By A3*, we get that $\mathrm{Pl}(\cup_d A_d) = \bot$. Since $PL \models true \rightarrow \exists x\, Winner(x)$, we get that either $\mathrm{Pl}(W) = \bot$ or $\mathrm{Pl}(\cup_d A_d) > \mathrm{Pl}(W - (\cup_d A_d))$. Since the latter inequality is inconsistent, we conclude that $\mathrm{Pl}(W) = \bot$ and, thus, $PL \models true \rightarrow false$, as desired.

Now assume that (b) is true. From A2†, it follows that

$$\mathrm{Pl}(A_{d_0}) \not< \mathrm{Pl}(A - A_{d_0}) \tag{10}$$

Since $PL \models \forall x(true \rightarrow \neg Winner(x))$, we have that

$$(PL, v_{d_0}) \models true \rightarrow \neg Winner(x)$$

Thus,

$$\mathrm{Pl}(A_{d_0}) < \mathrm{Pl}(W - A_{d_0}) = \mathrm{Pl}(B \cup (A - A_{d_0})). \tag{11}$$

Finally, since $PL \models true \rightarrow \exists x\, Winner(x)$, we have that

$$\mathrm{Pl}(B) < \mathrm{Pl}(A) = \mathrm{Pl}(A_{d_0} \cup (A - A_{d_0})). \tag{12}$$

Using A2, (10), and (11), we get that $\mathrm{Pl}(A - A_{d_0}) > \mathrm{Pl}(A_{d_0} \cup B) \geq \mathrm{Pl}(A_{d_0})$. This, however, contradicts (10), showing that (b) is impossible. ∎

Theorem 7.1: $\mathbf{C}^{stat}$ *is a sound and complete axiomatization of* $\mathcal{L}^{stat}$ *with respect to* $\mathcal{P}^{QPL}_{stat}$.

Proof: We proceed much as in the proof of Theorem 4.1, using the same notation. Again we can assume without loss of generality that φ is a sentence. Using standard arguments of first-order logic, we can show that there is a $\mathcal{C}$-good maximal consistent set of sentences C^* that includes φ. (The second property of $\mathcal{C}$-goodness, that $\forall x \psi \in C^*$ implies $\psi[x/\mathbf{c}] \in C^*$, now follows immediately from the axioms, since we can substitute constants in arbitrary contexts, and the proof of the first property is just the standard first-order proof, again because we get to use all the standard first-order axioms with no change.)

We construct a first-order statistical plausibility structure $PL = (Dom, \mathrm{Pl}, \pi)$ by again taking $Dom = \{[\mathbf{c}] : \mathbf{c} \in \mathcal{C}\}$ and defining π so that $\pi(\mathbf{c}) = [\mathbf{c}]$ for the constants $\mathbf{c} \in \mathcal{C}$ and π for the symbols in Φ' is determined by the atomic sentences in C^* in the obvious way. The definition of Pl is also similar in spirit to that in the proof of Theorem 4.1. We take the domain of Pl to be all the definable subsets of Dom^∞. More precisely, given a formula φ, let φ_v be the formula that results by replacing all free occurrences of x_i in φ by some constant in the equivalence class $v(x_i)$. (It doesn't matter which one.) Let $A_\varphi = \{v : \varphi_v \in C^*\}$. The definable subsets are precisely those of the form A_φ for some formula φ. Note that the definable subsets do form an algebra, since $A_\varphi \cup A_\psi = A_{\varphi\vee\psi}$ and $\overline{A_\varphi} = A_{\neg\varphi}$. We define Pl on this algebra by taking $\mathrm{Pl}(A_\varphi) \leq \mathrm{Pl}(A_\psi)$ iff $\varphi \vee \psi \rightsquigarrow_X \psi \in C^*$, where X consists of all the variables free in $\varphi \vee \psi$. We leave it to the reader to check that for every sentence $\psi \in \mathcal{L}^{stat}$, we have $PL \models \psi$ iff $\psi \in C^*$. ∎

Acknowledgments

We would like to thank Ronen Brafman, Ron Fagin, and Adam Grove for their comments on an earlier version of the paper.

References

Abadi, M. and J. Halpern (1994). Decidability and expressiveness for first-order logics of probability. *Information and Computation 112*(1), 1–36.

Adams, E. (1975). *The Logic of Conditionals*. Dordrecht, Netherlands: D. Reidel.

Bacchus, F. (1990). *Representing and Reasoning with Probabilistic Knowledge*. Cambridge, Mass.: MIT Press.

Bacchus, F., A. J. Grove, J. Y. Halpern, and D. Koller (1996). From statistical knowledge bases to degrees of belief. *Artificial Intelligence 87*(1–2), 75–143.

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

Brafman, R. I. (1997). A first-order conditional logic with qualitative statistical semantics. *Journal of Logic and Computation 7*(6), 777–803.

Delgrande, J. P. (1987). A first-order conditional logic for prototypical properties. *Artificial Intelligence 33*, 105–130.

Delgrande, J. P. (1988). An approach to default reasoning based on a first-order conditional logic: revised report. *Artificial Intelligence 36*, 63–90.

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

Dubois, D. and H. Prade (1991). Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proc. Twelfth International Joint Conference on Artificial Intelligence (IJCAI '91)*, pp. 419–424.

Enderton, H. B. (1972). *A Mathematical Introduction to Logic*. New York: Academic Press.

Friedman, N. and J. Y. Halpern (1995). Plausibility measures: a user's manual. In *Proc. Eleventh Conference on Uncertainty in Artificial Intelligence (UAI '95)*, pp. 175–184.

Friedman, N. and J. Y. Halpern (1997). Modeling belief in dynamic systems. part I: foundations. *Artificial Intelligence 95*(2), 257–316.

Friedman, N. and J. Y. Halpern (1998). Plausibility measures and default reasoning. *Journal of the ACM*. Accepted for publication. Also available at `http://www.cs.berkeley.edu/~nir`.

Garson, J. W. (1977). Quantification in modal logic. In D. Gabbay and F. Guenthner (Eds.), *Handbook of Philosophical Logic, Vol. II*, pp. 249–307. Dordrecht, Netherlands: Reidel.

Geffner, H. (1992). *Default reasoning: causal and conditional theories*. MIT Press.

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

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

Halpern, J. Y. (1990). An analysis of first-order logics of probability. *Artificial Intelligence 46*, 311–350.

Hughes, G. E. and M. J. Cresswell (1968). *An Introduction to Modal Logic*. London: Methuen.

Koller, D. and J. Halpern (1996). Irrelevance and conditioning in first-order probabilistic logic. In *Proceedings, Thirteenth National Conference on Artificial Intelligence (AAAI '96)*, pp. 569–576.

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

Kyburg, Jr., H. E. (1961). *Probability and the Logic of Rational Belief*. Middletown, Connecticut: Wesleyan University Press.

Lehmann, D. and M. Magidor (1990). Preferential logics: the predicate calculus case. In *Theoretical Aspects of Reasoning about Knowledge: Proc. Third Conference*, pp. 57–72. San Francisco, Calif.: Morgan Kaufmann.

Lester, H. (1988). *Tacky the Penguin*. Houghton-Mifflin.

Lewis, D. K. (1973). *Counterfactuals*. Cambridge, Mass.: Harvard University Press.

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)*, pp. 505–516. Reprinted in *Readings in Uncertain Reasoning*, G. Shafer and J. Pearl (eds.), Morgan Kaufmann, San Francisco, Calif., 1990, pp. 699–710.

Poole, D. (1991). The effect of knowledge on belief: conditioning, specificity and the lottery paradox in default reasoning. *Artificial Intelligence 49*(1–3), 282–307.

Schlechta, K. (1995). Defaults as generalized quantifiers. *Journal of Logic and Computation 5*(4), 473–494.

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

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