

UC Berkeley

Faculty Publications

Title

Preferential Structures for Comparative Probabilistic Reasoning

Permalink

<https://escholarship.org/uc/item/40c139d1>

Authors

Harrison-Trainor, Matthew
Holliday, Wesley Halcrow
Icard, Thomas Frederick, III

Publication Date

2016-11-01

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed

Preferential Structures for Comparative Probabilistic Reasoning*

Matthew Harrison-Trainer

Group in Logic and the Methodology of Science
University of California, Berkeley

Wesley H. Holliday

Department of Philosophy and
Group in Logic and the Methodology of Science
University of California, Berkeley

Thomas F. Icard, III

Department of Philosophy and
Symbolic Systems Program
Stanford University

Abstract

Qualitative and quantitative approaches to reasoning about uncertainty can lead to different logical systems for formalizing such reasoning, even when the language for expressing uncertainty is the same. In the case of reasoning about relative likelihood, with statements of the form $\varphi \succsim \psi$ expressing that φ is at least as likely as ψ , a standard qualitative approach using preordered preferential structures yields a dramatically different logical system than a quantitative approach using probability measures. In fact, the standard preferential approach validates principles of reasoning that are incorrect from a probabilistic point of view. However, in this paper we show that a natural modification of the preferential approach yields exactly the same logical system as a probabilistic approach—not using single probability measures, but rather sets of probability measures. Thus, the same preferential structures used in the study of non-monotonic logics and belief revision may be used in the study of comparative probabilistic reasoning based on imprecise probabilities.

Introduction

Central to being an intelligent agent is the ability to reason adequately under, and even about, uncertainty. Such reasoning requires representing this uncertainty in some way. In this paper, we study *qualitative* forms of reasoning about likelihood, a subject that can be traced back to the foundations of mathematical probability theory (de Finetti 1937). Such reasoning may concern absolute or relative likelihood, both of which have been studied in some detail within AI. The former involves statements of the form $L\varphi$ expressing that φ is likely (Halpern and Rabin 1987; Halpern and McAllester 1989). The latter—our topic in this paper—involves statements of the form $\varphi \succsim \psi$ expressing that φ is at least as likely as ψ .¹ This representation language has been considered under several different interpretations:

- **Probabilistic** interpretation (Seeger 1971; Gärdenfors 1975; van der Hoek 1996; Alon and Heifetz 2014; Delgrande and Renne 2015), building on the theory of qualitative probability (de Finetti 1937; Narens 2007);

*Postprint of AAAI 2017 paper, corrected to include a distinguished subset W_{\succsim} in Definitions 2-3 and 5 (resp. W_+ before Theorem 3) to match the semantics of Holliday and Icard 2013 (resp. van der Hoek 1996) when $R(w) = R(v)$ for all $w, v \in W$.

¹A different reading of $\varphi \succsim \psi$ as *belief in φ is at least as strong as belief in ψ* has been used to motivate another mathematical interpretation (Ghosh and de Jongh 2013).

- **Possibilistic** interpretation (del Cerro and Herzig 1991; Dubois, Fargier, and Prade 1998; Dubois, Fargier, and Perny 2003; Touazi, Cayrol, and Dubois 2015);
- **Preferential** interpretation (Halpern 1996; Halpern 1997).

Each of these interpretations captures a distinct intuition about relative likelihood reasoning, and they may be appropriate for different settings (Halpern 2003). A natural question is: to what extent are these interpretations compatible with one another? Our focus is on the relationship between the probabilistic and preferential interpretations. Existing approaches based on the preferential interpretation (Halpern 1997) depart rather dramatically from those based on probabilistic interpretations—in fact, relating closely to the possibilistic interpretation. We demonstrate how a natural modification of the approach—based on qualitative preferential structures—yields a calculus for reasoning about relative likelihood that is, from a logical point view, indistinguishable from standard forms of probabilistic reasoning.

Specifically, we prove soundness and completeness for the comparative logic of imprecise probabilities (Alon and Heifetz 2014) with respect to a class of preferential models that can be independently motivated. This shows how the two approaches can be seen as fundamentally compatible. The methods we use to prove these results further lead to a complexity result: the logic itself we prove to be NP-complete. This justifies the intuitive claim that this style of comparative probabilistic reasoning is relatively elementary, being no more complex than ordinary Boolean reasoning.

Preferential Structures

Preferential structures have been applied widely across the field of representation and reasoning, including in the study of non-monotonic logics (Shoham 1987; Kraus, Lehmann, and Magidor 1990), belief revision (Katsuno and Mendelzon 1991; Friedman and Halpern 1994a), and conditional belief in games (Board 2004; Baltag and Smets 2008). Restricted classes of preferential structures are also used for conditional logics (Lewis 1973; Friedman and Halpern 1994b). Here we are interested in the application of these structures to the study of likelihood comparisons, whereby an ordering on a set of states is “lifted” to an ordering on the set of events over those states, i.e., on the powerset.

Language

The language \mathcal{L} of relative likelihood, generated from a set Prop of propositional variables, is given by the grammar

$$\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \psi) \mid (\varphi \succ \psi),$$

where $p \in \text{Prop}$. We define \vee , \rightarrow , \leftrightarrow , \perp , and \top as usual. The *modal depth* of a formula of \mathcal{L} is defined by:

1. $md(p) = 0$ for $p \in \text{Prop}$;
2. $md(\neg\varphi) = md(\varphi)$;
3. $md(\varphi \wedge \psi) = \max\{md(\varphi), md(\psi)\}$;
4. $md(\varphi \succ \psi) = \max\{md(\varphi), md(\psi)\} + 1$.

The *length* $|\varphi|$ of φ is the number of symbols in φ .

Preferential Semantics

All of the models for \mathcal{L} that we will consider will add some extra structure to the following common base models.

Definition 1. A *state space model* is a tuple $M = \langle W, V \rangle$ where W is a nonempty set and $V: \text{Prop} \rightarrow \wp(W)$.

Preferential models then add a preorder on the space W .

Definition 2. A *preferential model* is an $\mathcal{M} = \langle W, \succeq, V \rangle$ where $\langle W, V \rangle$ is a state space model and \succeq is a preorder (reflexive and transitive binary relation) on a nonempty subset W_\succeq of W .

In a multi-agent setting, one could consider a family $\{\succeq_{w,a}\}_{w \in W, a \in \text{Agents}}$ of preorders on W , but in this paper we can make our essential points with a single preorder.

The idea of the preferential semantics for \mathcal{L} is to lift the order \succeq on W to an order \succeq^\uparrow on $\wp(W)$. A natural way to do so was studied by (Halpern 1996; Halpern 1997), following an earlier proposal by (Lewis 1973) in a different setting: let $A \succeq^l B$ iff for all $b \in B$ there is $a \in A$ such that $a \succeq b$. This is equivalent to there being a function $f: B \rightarrow A$ that is *inflationary* with respect to \succeq , i.e., $f(x) \succeq x$ for all $x \in B$.

Definition 3 (Inflationary function semantics). Given a preferential model $\mathcal{M} = \langle W, \succeq, V \rangle$, state $w \in W$, and formula $\varphi \in \mathcal{L}$, we define the satisfaction relation $\mathcal{M}, w \models_l \varphi$ by:

1. $\mathcal{M}, w \models_l p$ iff $w \in V(p)$;
2. $\mathcal{M}, w \models_l \neg\varphi$ iff $\mathcal{M}, w \not\models_l \varphi$;
3. $\mathcal{M}, w \models_l \varphi \wedge \psi$ iff $\mathcal{M}, w \models_l \varphi$ and $\mathcal{M}, w \models_l \psi$;
4. $\mathcal{M}, w \models_l \varphi \succ \psi$ iff $\llbracket \varphi \rrbracket^{\mathcal{M}} \succeq^l \llbracket \psi \rrbracket^{\mathcal{M}}$,

where $\llbracket \alpha \rrbracket^{\mathcal{M}} = \{w \in W_\succeq \mid \mathcal{M}, w \models_l \alpha\}$.

For this semantics and all those to follow, a formula is *valid* over a class of models according to the semantics iff it is satisfied by every state in every model in the class.

The motivation for studying the logic of the particular lifting operation above was in part to establish connections with default reasoning and with possibilistic models. In particular, the axiom (L4) below is the central principle in the axiomatization of this class of models.

Definition 4. The set of theorems of IL (the logic of *inflationary lifting*) is the smallest set of formulas that contains all tautologies of propositional logic, is closed under *modus ponens* (if $\varphi \in \text{IL}$ and $\varphi \rightarrow \psi \in \text{IL}$, then $\psi \in \text{IL}$) and *necessitation* (if $\varphi \in \text{IL}$, then $\varphi \succ \top \in \text{IL}$), and contains all instances of the following axiom schemas:

- (L1) $\varphi \succ \varphi$
- (L2) $(\perp \succ (\psi \wedge \neg\varphi)) \rightarrow (\varphi \succ \psi)$
- (L3) $((\varphi \succ \psi) \wedge (\psi \succ \chi)) \rightarrow (\varphi \succ \chi)$
- (L4) $((\varphi \succ \psi) \wedge (\varphi \succ \chi)) \rightarrow (\varphi \succ (\psi \vee \chi))$
- (I1) $(\varphi \succ \psi) \rightarrow ((\varphi \succ \psi) \succ \top)$
- (I2) $\neg(\varphi \succ \psi) \rightarrow (\neg(\varphi \succ \psi) \succ \top)$.

Theorem 1 (Halpern 2003). The logic IL is sound and complete with respect to preferential models according to the inflationary lifting semantics.

While the principle (L4) may be sensible in some contexts, it is patently incompatible with probabilistic reasoning, where taking the union of many improbable events can lead to probable events. A natural question is whether there is an alternative lifting operation that might accord better with probabilistic reasoning. Such a lifting was proposed in (Holliday and Icard 2013): instead of requiring an inflationary function B to A , let us say that $A \succeq^i B$ iff there is an inflationary *injection* from B to A . Call this the *inflationary injection semantics* for the language \mathcal{L} .

Definition 5 (Inflationary injection semantics). Given a preferential model $\mathcal{M} = \langle W, \succeq, V \rangle$, state $w \in W$, and formula $\varphi \in \mathcal{L}$, we define $\mathcal{M}, w \models \varphi$ in the same way as Definition 3, except with clause 4 replaced by:

4. $\mathcal{M}, w \models \varphi \succ \psi$ iff $\llbracket \varphi \rrbracket^{\mathcal{M}} \succeq^i \llbracket \psi \rrbracket^{\mathcal{M}}$,

where $\llbracket \alpha \rrbracket^{\mathcal{M}} = \{w \in W_\succeq \mid \mathcal{M}, w \models \alpha\}$.

The intuition behind this alternative is clear. If we are trying to determine whether event A is at least as likely as event B , we might try to find, for each possible B state, an A state that is as least as likely as it. However, we cannot match multiple B states with the same A state, since together these B states may become more likely than the A state. The inflationary injection semantics generalizes this intuition to require each B state to correspond with its own unique A state that weakly dominates it in likelihood.

We would now like to verify that this intuition does indeed lead to a system for reasoning about relative likelihood that validates intuitive principles. In particular, we would like a completeness theorem along the lines of Theorem 1. In that direction we first introduce a previously studied logic for reasoning about probabilistic semantics for language \mathcal{L} .

Probabilistic Semantics

Probabilistic models replace the preorder \succeq by a set \mathcal{P} of probability measures. A singleton set \mathcal{P} represents *sharp probabilities*, while multiple disagreeing measures in \mathcal{P} can be used to represent *imprecise probabilities*.

Definition 6. A *multi-measure model* is a tuple $\mathfrak{M} = \langle W, \mathcal{P}, V \rangle$ where $\langle W, V \rangle$ is a state space model and \mathcal{P} is a set of finitely additive measures $\mu: \wp(W) \rightarrow [0, 1]$.

A *single-measure model* is a multi-measure model in which $|\mathcal{P}| = 1$. If $\mathcal{P} = \{\mu\}$, we write $\mathfrak{M} = \langle W, \mu, V \rangle$.

Likelihood comparisons $\varphi \succ \psi$ are now decided by a unanimity rule over \mathcal{P} as follows.

Definition 7 (Multi-measure semantics). Given a multi-measure model $\mathfrak{M} = \langle W, \mathcal{P}, V \rangle$, state $w \in W$, and formula $\varphi \in \mathcal{L}$, we define $\mathfrak{M}, w \models \varphi$ with the usual clauses for propositional variables and Boolean connectives, plus:

$$\mathfrak{M}, w \models \varphi \succsim \psi \text{ iff } \forall \mu \in \mathcal{P}: \mu(\llbracket \varphi \rrbracket^{\mathfrak{M}}) \geq \mu(\llbracket \psi \rrbracket^{\mathfrak{M}})$$

where $\llbracket \varphi \rrbracket^{\mathfrak{M}} = \{w \in W \mid \mathfrak{M}, w \models \varphi\}$.

The next two lemmas will be useful for our later results. The first reduces formulas to a convenient form.

Lemma 1. In each of the semantics (multi-measure, inflationary function, inflationary injection) above, every formula $\varphi \in \mathcal{L}$ is equivalent to a formula ψ of modal depth ≤ 1 . Moreover, we may take ψ to be a disjunction in which the length of each disjunct is $O(|\varphi|)$.

Proof. Suppose $md(\varphi) > 1$, which implies there is an inequality $\alpha \succsim \beta$ inside the scope of another inequality, where $md(\alpha) = md(\beta) = 0$. Define $\varphi' := \varphi_1 \vee \varphi_2$ by

$$\varphi' := (\varphi[\top/\alpha \succsim \beta] \wedge \alpha \succsim \beta) \vee (\varphi[\perp/\alpha \succsim \beta] \wedge \neg(\alpha \succsim \beta)).$$

For any model \mathbb{M} with domain W , either $\llbracket \alpha \succsim \beta \rrbracket^{\mathbb{M}} = W$ or $\llbracket \alpha \succsim \beta \rrbracket^{\mathbb{M}} = \emptyset$, which implies that $\llbracket \varphi \rrbracket^{\mathbb{M}} = \llbracket \varphi' \rrbracket^{\mathbb{M}}$. The number of inequalities in the scope of other inequalities in φ'_i is less than that in φ , and the length of φ'_i is greater than that of φ by a constant amount. Thus, repeating this process on φ'_1 and φ'_2 , we eventually obtain the desired ψ . \square

The next lemma bounds the number of measures we need.

Lemma 2. If there is a multi-measure model $\langle W, \mathcal{P}, V \rangle$ and $w \in W$ that satisfies φ , then there is a $\mathcal{P}_0 \subseteq \mathcal{P}$ whose size is $O(|\varphi|)$ such that w satisfies φ in $\langle W, \mathcal{P}_0, V \rangle$.

Proof. By Lemma 1 and propositional logic, φ is equivalent to a disjunction of conjunctions of the form

$$\neg(\varphi_1 \succsim \psi_1) \wedge \cdots \wedge \neg(\varphi_n \succsim \psi_n) \wedge \xi$$

where ξ is a conjunction of modal depth 0 formulas and inequalities $\varphi \succsim \psi$ between such formulas. It suffices to satisfy one of the disjuncts, and for this it suffices to have one measure μ_i for each of the n negated inequalities. Let $\mathcal{P}_0 = \{\mu_1, \dots, \mu_n\}$. Since the satisfaction of the positive inequalities in ξ is preserved under discarding measures from \mathcal{P} , we have that φ is still satisfied by w in $\langle W, \mathcal{P}_0, V \rangle$. \square

Logical Systems SP and IP

For logics of probability models, we first define the logic of single-measure models. To do so, it will be helpful to introduce some abbreviations from (Seegerberg 1971; Gärdenfors 1975). Given formulas $\varphi_1, \dots, \varphi_n, \psi_1, \dots, \psi_n$ and $0 \leq k \leq n$, define C_k to be the disjunction of all conjunctions

$$f_1 \varphi_1 \wedge \cdots \wedge f_n \varphi_n \wedge g_1 \psi_1 \wedge \cdots \wedge g_n \psi_n$$

where exactly k of the f 's and k of the g 's are the empty string, and the rest are \neg . Thus, C_k is true at a state w iff exactly k of the φ 's and k of the ψ 's are true at w . Then let

$$(\varphi_1, \dots, \varphi_n) \equiv (\psi_1, \dots, \psi_n) := C_0 \vee \cdots \vee C_n,$$

which is true at a state w iff the number of φ 's true at w is exactly the same as the number of ψ 's true at w .

Definition 8. The set of theorems of SP (the logic of *sharp probability*) is the smallest set of formulas that contains all tautologies of propositional logic, is closed under *modus ponens* and *necessitation* as in Definition 4, and contains all instances of **(I1)**, **(I2)**, and the following for all $n \in \mathbb{N}$:²

$$\mathbf{(A0)} (\varphi \succsim \psi) \vee (\psi \succsim \varphi)$$

$$\mathbf{(A1)} \varphi \succsim \perp \quad \mathbf{(A2)} \varphi \succsim \varphi \quad \mathbf{(A3)} \neg(\perp \succsim \top)$$

$$\mathbf{(A4)} [(\varphi_1 \succsim \psi_1) \wedge \cdots \wedge (\varphi_n \succsim \psi_n) \wedge$$

$$(\varphi_1, \dots, \varphi_n, \varphi') \equiv (\psi_1, \dots, \psi_n, \psi') \succsim \top] \rightarrow (\psi' \succsim \varphi').$$

It is easy to see that each axiom of SP is valid over single-measure models. In the case of **(A4)**, if the set of states in \mathfrak{M} that make exactly the same number of formulas from $(\varphi_1, \dots, \varphi_n, \varphi')$ true as from $(\psi_1, \dots, \psi_n, \psi')$ has measure 1, and $\mu(\llbracket \varphi_i \rrbracket^{\mathfrak{M}}) \geq \mu(\llbracket \psi_i \rrbracket^{\mathfrak{M}})$ for $1 \leq i \leq n$, then it is impossible to have $\mu(\llbracket \varphi' \rrbracket^{\mathfrak{M}}) > \mu(\llbracket \psi' \rrbracket^{\mathfrak{M}})$. The soundness of SP is thus straightforward, but the completeness of SP is far less so. The completeness proof in (Seegerberg 1971; Gärdenfors 1975) uses a famous representation theorem of (Kraft, Pratt, and Seidenberg 1959; Scott 1964).

Theorem 2 (Seegerberg 1971, Gärdenfors 1975). SP is sound and complete with respect to the class of single-measure models.

As shown by van der Hoek (van der Hoek 1996), SP can be given a simple semantics using what we will call *distinguished state models* $M = \langle W, W_+, V \rangle$ where (W, V) is a state space model and $\emptyset \neq W_+ \subseteq W$. Then we define

$$M, w \models \varphi \succsim \psi \text{ iff } |\llbracket \varphi \rrbracket^M| \geq |\llbracket \psi \rrbracket^M|,$$

where $\llbracket \varphi \rrbracket^M = \{w \in W_+ \mid M, w \models \varphi\}$ and $|\cdot|$ is cardinality.

Theorem 3 (van der Hoek 1996). SP is sound and complete with respect to the class of finite distinguished state models with the cardinality semantics above.

The logic SP is not sound, however, with respect to all multi-measure models, as axiom **(A0)** is clearly invalid; for one measure $\mu \in \mathcal{P}$ we may have $\mu(\llbracket \varphi \rrbracket^{\mathfrak{M}}) > \mu(\llbracket \psi \rrbracket^{\mathfrak{M}})$ while for another measure $\mu' \in \mathcal{P}$ we may have $\mu'(\llbracket \psi \rrbracket^{\mathfrak{M}}) > \mu'(\llbracket \varphi \rrbracket^{\mathfrak{M}})$, in which case neither $\varphi \succsim \psi$ nor $\psi \succsim \varphi$ is true in \mathfrak{M} according to Definition 7. Thus, for multi-measure models we must drop the axiom **(A0)**; but if we only drop **(A0)** from SP, then the resulting logic will be sound but not *complete* with respect to multi-measure models (Harrison-Trainor, Holliday, and Icard 2016b). We must not only drop **(A0)** but also strengthen **(A4)** as follows.

Definition 9. The set of theorems of IP (the logic of *imprecise probability*) is defined in the same way as SP except without axiom **(A0)** and with **(A4)** replaced by:

$$\mathbf{(A4')} [(\varphi_1 \succsim \psi_1) \wedge \cdots \wedge (\varphi_n \succsim \psi_n) \wedge$$

$$(\varphi_1, \dots, \varphi_n, \underbrace{\varphi', \dots, \varphi'}_{k \text{ times}}) \equiv (\psi_1, \dots, \psi_n, \underbrace{\psi', \dots, \psi'}_{k \text{ times}}) \succsim \top]$$

$$\rightarrow (\psi' \succsim \varphi').$$

Again soundness is not difficult, while completeness relies on a representation theorem in (Alon and Lehrer 2014).

Theorem 4 (Alon and Heifetz 2014). IP is sound and complete with respect to the class of all multi-measure models.

²We use the labeling of axioms in (Alon and Heifetz 2014).

Main Result

Our main result is to prove that the logic one obtains from finite preferential models with the inflationary injection semantics is exactly the logic IP of imprecise probability that one obtains from multi-measure models.

Theorem 5. IP is sound and complete with respect to the class of finite preferential models with the inflationary injection semantics. Moreover, it is complete with respect to the class of finite preferential models in which \succeq is a *total* preorder on W_{\succeq} (i.e., for all $w, v \in W_{\succeq}$, $w \succeq v$ or $v \succeq w$).

We will first prove soundness in the next section and then completeness in the following section.

Soundness of IP

In proving the soundness of IP, the main task is show that the axiom (A4') is valid according to the inflationary injection semantics. To do so, we will use a result about inflationary injections from (Harrison-Trainor, Holliday, and Icard 2016a). This result involves the following relation between sequences of sets, which matches the relation expressed by the second line in the display of axiom (A4') above.

Definition 10. Let S be a finite set. For any two sequences $\langle E_1, \dots, E_k \rangle$ and $\langle F_1, \dots, F_k \rangle$ of events from $\mathcal{P}(S)$,

$$\langle E_1, \dots, E_k \rangle =_0 \langle F_1, \dots, F_k \rangle$$

if and only if for all $s \in S$, the cardinality of $\{i \mid s \in E_i\}$ is equal to the cardinality of $\{i \mid s \in F_i\}$.

If $\langle E_1, \dots, E_k \rangle =_0 \langle F_1, \dots, F_k \rangle$, then we say that the two sequences are *balanced*; every state appears the same number of times on the left side as on the right side.

The following lemma is easily seen to follow from Lemma 3.6 and the proof of Lemma 3.7 in (Harrison-Trainor, Holliday, and Icard 2016a).

Lemma 3 (Harrison-Trainor, Holliday, and Icard 2016). Let $\langle W, \succeq \rangle$ be a finite preorder. If

$$\langle E_1, \dots, E_n, \underbrace{A, \dots, A}_{r \text{ times}} \rangle =_0 \langle F_1, \dots, F_n, \underbrace{B, \dots, B}_{r \text{ times}} \rangle$$

are balanced sequences of subsets of W , and there is an inflationary injection $f_i: F_i \rightarrow E_i$ for all i , then there is an inflationary injection $g: A \rightarrow B$.

We will now prove our soundness theorem, using Lemma 3 to verify axiom (A4').³

Theorem 6. IP is sound with respect to finite preferential models according to the inflationary injection semantics.

Proof. Obviously all tautologies are valid, and the rules of modus ponens and necessitation preserve validity. It only remains to show that (A1)–(A3), (A4'), and (I1)–(I2) hold at every state w in each preferential model $\mathcal{M} = \langle W, \succeq, V \rangle$.

(A1) holds via the trivial injection $\emptyset \rightarrow \llbracket \varphi \rrbracket^{\mathcal{M}}$. (A2) holds via the identity injection $\llbracket \varphi \rrbracket^{\mathcal{M}} \rightarrow \llbracket \varphi \rrbracket^{\mathcal{M}}$. (A3) holds because there is no injection $W \rightarrow \emptyset$.

³Lemma 3 and Theorem 6 hold more generally for models in which there is no infinite sequence x_0, x_1, x_2, \dots of distinct states with $x_{n+1} \succeq x_n$ (Harrison-Trainor, Holliday, and Icard 2016a).

For (A4'), given

$$\begin{aligned} \mathcal{M}, w \models (\varphi_1 \succsim \psi_1) \wedge \dots \wedge (\varphi_n \succsim \psi_n) \wedge \\ (\varphi_1, \dots, \varphi_n, \underbrace{\varphi', \dots, \varphi'}_{k \text{ times}}) \equiv (\psi_1, \dots, \psi_n, \underbrace{\psi', \dots, \psi'}_{k \text{ times}}) \succsim \top, \end{aligned}$$

we have to show that $\mathcal{M}, w \models \psi' \succsim \varphi'$.

We have inflationary injections $f_i: \llbracket \psi_i \rrbracket^{\mathcal{M}} \rightarrow \llbracket \varphi_i \rrbracket^{\mathcal{M}}$. There is also an inflationary injection

$$W \rightarrow$$

$$\llbracket (\varphi_1, \dots, \varphi_n, \underbrace{\varphi', \dots, \varphi'}_{k \text{ times}}) \rrbracket^{\mathcal{M}} \equiv (\psi_1, \dots, \psi_n, \underbrace{\psi', \dots, \psi'}_{k \text{ times}}) \rrbracket^{\mathcal{M}}.$$

Since W is finite, it follows that

$$W =$$

$$\llbracket (\varphi_1, \dots, \varphi_n, \underbrace{\varphi', \dots, \varphi'}_{k \text{ times}}) \rrbracket^{\mathcal{M}} \equiv (\psi_1, \dots, \psi_n, \underbrace{\psi', \dots, \psi'}_{k \text{ times}}) \rrbracket^{\mathcal{M}},$$

which implies

$$\begin{aligned} \langle \llbracket \varphi_1 \rrbracket^{\mathcal{M}}, \dots, \llbracket \varphi_n \rrbracket^{\mathcal{M}}, \underbrace{\llbracket \varphi' \rrbracket^{\mathcal{M}}, \dots, \llbracket \varphi' \rrbracket^{\mathcal{M}}}_{k \text{ times}} \rangle \\ =_0 \langle \llbracket \psi_1 \rrbracket^{\mathcal{M}}, \dots, \llbracket \psi_n \rrbracket^{\mathcal{M}}, \underbrace{\llbracket \psi' \rrbracket^{\mathcal{M}}, \dots, \llbracket \psi' \rrbracket^{\mathcal{M}}}_{k \text{ times}} \rangle. \end{aligned}$$

Then by Lemma 3, there is an inflationary injection $\llbracket \varphi' \rrbracket^{\mathcal{M}} \rightarrow \llbracket \psi' \rrbracket^{\mathcal{M}}$, and so $\mathcal{M}, w \models \psi' \succsim \varphi'$.

For (I1), suppose that $\mathcal{M}, w' \models \varphi \succsim \psi$, say via an inflationary injection $f: \llbracket \psi \rrbracket^{\mathcal{M}} \rightarrow \llbracket \varphi \rrbracket^{\mathcal{M}}$. Then in fact, for all $w, \mathcal{M}, w \models \varphi \succsim \psi$ via f . Hence $\llbracket \varphi \succsim \psi \rrbracket^{\mathcal{M}} = W$, and so $\mathcal{M}, w \models (\varphi \succsim \psi) \succsim \top$. For (I2), by a similar argument, if $\mathcal{M}, w \models \neg(\varphi \succsim \psi)$, then for all $w', \mathcal{M}, w' \models \neg(\varphi \succsim \psi)$, so that $\llbracket \neg(\varphi \succsim \psi) \rrbracket^{\mathcal{M}} = W$, and $\mathcal{M}, w \models \neg(\varphi \succsim \psi) \succsim \top$. \square

Completeness of IP

Our method of proving completeness is to show that any finite multi-measure model can be transformed into a finite preferential model that satisfies the same formulas according to the inflationary injection semantics. We will then invoke the completeness of IP with respect to finite multi-measure models, as proved by (Alon and Heifetz 2014), to establish the completeness of IP with respect to finite preferential models with the inflationary injection semantics.

We begin with a restatement of Lemma 3.2.2 in (van der Hoek 1996).

Lemma 4. For any finite single-measure model $\mathfrak{M} = \langle W, \mu, V \rangle$, there is a distinguished state model $M^\# = \langle W^\#, W^\#_+, V^\# \rangle$ with $W \subseteq W^\#$ such that for all $w \in W$ and $\varphi \in \mathcal{L}$:

$$\mathfrak{M}, w \models \varphi \text{ iff } M^\#, w \models \varphi.$$

The transformation first rationalizes the probability measure, then normalizes the measure so that it takes on only integer values (but does not necessarily assign measure one to the whole state space), and finally duplicates each point in the state space that has nonzero measure with a number of points equal to its measure, which become the points in $W^\#_+$.

Using Lemma 4, we can transform multi-measure models into preferential models as follows.

Lemma 5. For any multi-measure model $\mathfrak{M} = \langle W, \mathcal{P}, V \rangle$ with W and \mathcal{P} finite and $w \in W$, there is a finite preferential model \mathcal{M} and state v such that for all $\varphi \in \mathcal{L}$:

$$\mathfrak{M}, w \models \varphi \text{ iff } \mathcal{M}, v \models \varphi.$$

Moreover, we may take \succeq in \mathcal{M} to be total on its field W_{\succeq} .

Proof. Where $\mathcal{P} = \{\mu_i\}_{i \in I}$, define for each $i \in I$ a single-measure model $\mathfrak{M}_i = \langle W_i, \mu_i, V_i \rangle$ where $W_i = W \times \{i\}$ and $V_i(p) = V(p) \times \{i\}$. By Lemma 4, for each \mathfrak{M}_i there is a distinguished state model $M_i^\# = \langle W_i^\#, W_{i+}^\#, V_i^\# \rangle$ with $W_i \subseteq W_i^\#$ such that for all $v \in W_i$ and $\varphi \in \mathcal{L}$:

$$\mathfrak{M}_i, v \models_\mu \varphi \text{ iff } M_i^\#, v \models \varphi. \quad (1)$$

Without loss of generality, assume that the domains $W_i^\#$ are pairwise disjoint. Now define the preferential model $\mathcal{M} = \langle W', \succeq, V' \rangle$ as follows:

$$(a) W' = \bigcup_{i \in I} W_i^\#;$$

$$(b) V'(p) = \bigcup_{i \in I} V_i^\#(p);$$

$$(c) w \succeq v \text{ iff for some } i \in I, \text{ we have } w, v \in W_{i+}^\#.$$

First, note that for any formula φ of modal depth 0:

$$\llbracket \varphi \rrbracket^{\mathcal{M}} = \bigcup_{i \in I} \llbracket \varphi \rrbracket^{M_i^\#}. \quad (2)$$

Next, we claim that for any $i \in I$ and formula φ :

$$\mathfrak{M}_i, w \models \varphi \text{ iff } \mathcal{M}, \langle w, i \rangle \models \varphi. \quad (3)$$

By Lemma 1, it suffices to consider formulas of modal depth ≤ 1 . The proof of (3) is by induction with obvious atomic and Boolean cases. For the modal case, consider $\varphi \succsim \psi$ where φ and ψ are of modal depth 0. If $\mathfrak{M}_i, w \models \varphi \succsim \psi$, then for all $i \in I$, $\mathfrak{M}_i, \langle w, i \rangle \models \varphi \succsim \psi$ and hence $M_i^\#, \langle w, i \rangle \models \varphi \succsim \psi$ by (1). Since $M_i^\#, \langle w, i \rangle \models \varphi \succsim \psi$, there is an injection $f_i: \llbracket \psi \rrbracket^{M_i^\#} \rightarrow \llbracket \varphi \rrbracket^{M_i^\#}$. Then $g = \bigcup_{i \in I} f_i$ is an injection; g is inflationary by (c) above; and $g: \llbracket \psi \rrbracket^{\mathcal{M}} \rightarrow \llbracket \varphi \rrbracket^{\mathcal{M}}$ by (2). Thus, $\mathcal{M}, \langle w, i \rangle \models \varphi \succsim \psi$.

Conversely, if $\mathfrak{M}_i, w \not\models \varphi \succsim \psi$, then there is an $i \in I$ with $\mathfrak{M}_i, \langle w, i \rangle \not\models \varphi \succsim \psi$ and hence $M_i^\#, \langle w, i \rangle \not\models \varphi \succsim \psi$ by (1). Thus, there is no injection $f_i: \llbracket \psi \rrbracket^{M_i^\#} \rightarrow \llbracket \varphi \rrbracket^{M_i^\#}$. Now suppose for reductio that there is an inflationary injection $g: \llbracket \psi \rrbracket^{\mathcal{M}} \rightarrow \llbracket \varphi \rrbracket^{\mathcal{M}}$. Given (c), that g is an inflationary injection implies that $g \upharpoonright W_{i+}^\#$ is an injection from $\llbracket \psi \rrbracket^{M_i^\#}$ to $\llbracket \varphi \rrbracket^{M_i^\#}$, contradicting the above. Thus, there is no such inflationary injection g , so $\mathcal{M}, \langle w, i \rangle \not\models \varphi \succsim \psi$.

Finally, we will show how to transform \mathcal{M} into a model \mathcal{M}^t where \succeq^t is total on W_{\succeq^t} . Just as the domain of \mathcal{M} is the union of disjoint domains $W_i^\#$ for $i \in I$, the domain of \mathcal{M}^t will be the union of disjoint domains W_i^t . Suppose the index set I with which we began is $I = \{0, \dots, m\}$. We define the domains W_i^t inductively as follows, for some $d \notin W'$:

$$(d_1) W_0^t = W_0^\# \times \{d\};$$

$$(d_2) W_{n+1}^t = W_{n+1}^\# \times \left(\bigcup_{i \leq n} W_i^t \cup \{d\} \right);$$

$$(d_3) W^t = \bigcup_{i \in I} W_i^t.$$

We multiply $W_0^\#$ by $\{d\}$ so that every state in W^t will be a pair, the first coordinate of which is a state from W' . Then W_1^t is the result of making $|W_0^t| + 1$ -many copies of each state in $W_1^\#$; W_2^t is the result of making $|W_0^t \cup W_1^t| + 1$ -many copies of each state in $W_2^\#$; and so on.

We define the valuation V^t by

$$(e) V^t(p) = \{ \langle w, v \rangle \in W^t \mid w \in V'(p) \},$$

so each copy of a state in \mathcal{M}^t has the same propositional valuation that the original state had in \mathcal{M} .

Where $W_{n+}^t = \{ \langle w, v \rangle \in W_n^t \mid w \in W_{n+}^\# \}$, define

$$(f) x \succeq^t y \text{ iff for some } k \leq \ell, x \in W_{k+}^t \text{ and } y \in W_{\ell+}^t.$$

Then clearly \succeq^t is a total preorder on its field.

We claim that for any $\langle w, v \rangle \in W^t$ and $\varphi \in \mathcal{L}$:

$$\mathcal{M}^t, \langle w, v \rangle \models \varphi \text{ iff } \mathcal{M}, w \models \varphi. \quad (4)$$

The proof is by induction on φ with the atomic case given by the definition of V^t . The Boolean cases are also routine.

For the modal case, as before we may consider $\varphi \succsim \psi$ where φ and ψ are of modal depth 0. If $\mathcal{M}, w \models \varphi \succsim \psi$, then in \mathcal{M} there is a \succeq -inflationary injection $f: \llbracket \psi \rrbracket^{\mathcal{M}} \rightarrow \llbracket \varphi \rrbracket^{\mathcal{M}}$. By the inductive hypothesis, we have

$$\llbracket \psi \rrbracket^{\mathcal{M}^t} = \{ \langle x, z \rangle \in W^t \mid x \in \llbracket \psi \rrbracket^{\mathcal{M}} \} \quad (5)$$

$$\llbracket \varphi \rrbracket^{\mathcal{M}^t} = \{ \langle y, z \rangle \in W^t \mid y \in \llbracket \varphi \rrbracket^{\mathcal{M}} \}. \quad (6)$$

Since f is \succeq -inflationary, $f(x) = y$ implies that x and y belong to the same $W_{i+}^\#$ by (c) above. Thus, by (d₁)-(d₃), we have that x and y get copied by the same elements in W^t : $\langle x, z \rangle \in W^t$ iff $\langle y, z \rangle \in W^t$. Then given (5)-(6), it follows that the function f^t sending each $\langle x, z \rangle \in \llbracket \psi \rrbracket^{\mathcal{M}^t}$ to $\langle f(x), z \rangle \in \llbracket \varphi \rrbracket^{\mathcal{M}^t}$ is a \succeq^t -inflationary injection.

Now suppose $\mathcal{M}, w \not\models \varphi \succsim \psi$, so there is no \succeq -inflationary injection from $\llbracket \psi \rrbracket^{\mathcal{M}}$ to $\llbracket \varphi \rrbracket^{\mathcal{M}}$. By construction of \mathcal{M} , it follows that for some $i \in I$, there is no injection from $\llbracket \psi \rrbracket^{\mathcal{M}} \cap W_i^\#$ to $\llbracket \varphi \rrbracket^{\mathcal{M}} \cap W_i^\#$. Thus,

$$|\llbracket \psi \rrbracket^{\mathcal{M}} \cap W_i^\#| > |\llbracket \varphi \rrbracket^{\mathcal{M}} \cap W_i^\#|.$$

It follows by (d₁)-(d₃) and (5)-(6) that

$$|\llbracket \psi \rrbracket^{\mathcal{M}^t} \cap W_i^t| > |\llbracket \varphi \rrbracket^{\mathcal{M}^t} \cap W_i^t| + \left| \bigcup_{j < i} W_j^t \right|. \quad (7)$$

Now by the definition of \succeq^t in (f), any \succeq^t -inflationary injection from $\llbracket \psi \rrbracket^{\mathcal{M}^t} \cap W_i^t$ to $\llbracket \varphi \rrbracket^{\mathcal{M}^t}$ must be an injection from $\llbracket \psi \rrbracket^{\mathcal{M}^t} \cap W_i^t$ to $\llbracket \varphi \rrbracket^{\mathcal{M}^t} \cap \bigcup_{j < i} W_j^t$, but such an injection cannot exist by (7). Thus, there is no \succeq^t -inflationary injection from $\llbracket \psi \rrbracket^{\mathcal{M}^t}$ to $\llbracket \varphi \rrbracket^{\mathcal{M}^t}$, so $\mathcal{M}^t, \langle w, v \rangle \not\models \varphi \succsim \psi$. \square

Using Lemma 5, we now exploit the known completeness result for multi-measure models to obtain the following.

Theorem 7. IP is complete with respect to finite preferential models (with \succeq total on W_{\succeq}) according to the inflationary injection semantics.

Proof. If φ is not a theorem of IP, then by Theorem 2 of (Alon and Heifetz 2014), there is a finite multi-measure model that satisfies $\neg\varphi$. We may assume it has only finitely many measures by Lemma 2. Thus, by Lemma 5, there is a finite preferential model (with \succeq total on W_{\succeq}) satisfying $\neg\varphi$ according to the inflationary injection semantics. \square

Complexity of IP

We now consider the problem of deciding whether a formula is consistent in IP (a formula φ being consistent if $\neg\varphi \notin \text{IP}$). We will show that this problem is NP-complete. Our strategy will be to show that if a formula φ is satisfiable, then it is satisfiable in a multi-measure model of polynomially bounded size, in which the measure of each state is a rational number of polynomially bounded size. A similar argument can be given using preferential models, but for the sake of space we will use a known result for measure models.

In particular, we will use a result from (Fagin, Halpern, and Megiddo 1990), where it was shown that the satisfiability problem, for single-measure models, of a logic allowing for rational comparisons of probability is NP-complete. The language they considered is essentially an extension of the formulas of our language of modal depth ≤ 1 to allow for rational comparisons. Thus, the following is immediate from Theorem 2.6 of (Fagin, Halpern, and Megiddo 1990).

Theorem 8. Suppose φ is a formula of modal depth ≤ 1 that is satisfied in some finite single-measure model. Then φ is satisfied in a single-measure model $\langle W, \mu, V \rangle$ with $|W| \leq |\varphi|$ and where the probability assigned to each state is a rational number with size $O(|\varphi| \log |\varphi|)$.

By the size of a rational number, we mean a bound on the binary representations of its numerator and denominator.

From Theorem 8, we obtain an analogous result for multi-measure models using the strategy of Lemma 5.

Theorem 9. Suppose θ is a formula of any modal depth that is satisfied in some finite multi-measure model. Then θ is satisfied in a multi-measure model $\langle W, \mathcal{P}, V \rangle$ with $|W| \leq O(|\theta|^2)$, $|\mathcal{P}| \leq O(|\theta|)$, and where the probabilities assigned to each state are rational numbers with size $O(|\theta| \log |\theta|)$.

Proof. Given a formula θ , applying Lemma 1 and putting each disjunct in disjunctive normal form, as in Lemma 2 we can rewrite θ as a disjunction of formulas

$$\neg(\varphi_1 \succsim \psi_1) \wedge \cdots \wedge \neg(\varphi_n \succsim \psi_n) \wedge \xi$$

of modal depth ≤ 1 , where ξ is a conjunction of modal depth 0 formulas and inequalities $\varphi \succsim \psi$ between modal depth 0 formulas. By Lemma 1, the length of each disjunct is $O(|\theta|)$. Since θ is satisfiable iff one of the disjuncts is, it suffices to show that if a disjunct φ is satisfiable then it is satisfiable in a model of bounded size as in the statement of the theorem.

If the disjunct

$$\varphi := \neg(\varphi_1 \succsim \psi_1) \wedge \cdots \wedge \neg(\varphi_n \succsim \psi_n) \wedge \xi$$

is satisfiable in a finite multi-measure model $\mathfrak{M} = \langle W, \mathcal{P}, V \rangle$ at a state w , then for each $i \leq n$, there is a measure μ_i such that $\mu_i(\llbracket \varphi_i \rrbracket^{\mathfrak{M}}) < \mu_i(\llbracket \psi_i \rrbracket^{\mathfrak{M}})$. Then $\langle W, \mu_i, V \rangle, w \models \neg(\varphi_i \succsim \psi_i) \wedge \xi$. So each formula $\neg(\varphi_i \succsim \psi_i) \wedge \xi$ is satisfiable in a finite single-measure model. Thus, by Theorem 8, each $\neg(\varphi_i \succsim \psi_i) \wedge \xi$ is satisfiable at w_i in a finite single-measure model $\mathfrak{M}_i^* = \langle W_i^*, \mu_i^*, V_i^* \rangle$ with at most $O(|\varphi|)$ states and where the probability assigned to each state is a rational number with size $O(|\varphi| \log |\varphi|)$.

Let W^* be the disjoint union of the W_i^* . We can extend each μ_i^* to all of W^* by having μ_i^* assign measure zero to $W^* - W_i^*$. Let \mathcal{P}^* be the set of the μ_i^* , and let V^* be the valuation on W^* such that $V^*(p)$ is the union of the $V_i^*(p)$'s. Let $\mathfrak{M}^* = \langle W^*, \mathcal{P}^*, V^* \rangle$. Then for any $w = w_i$,

$$\mathfrak{M}^*, w \models \neg(\varphi_1 \succsim \psi_1) \wedge \cdots \wedge \neg(\varphi_n \succsim \psi_n) \wedge \xi.$$

As $|W^*| \leq O(|\varphi|^2)$ and $|\mathcal{P}^*| \leq O(|\varphi|)$, we are done. \square

From Theorem 9, the complexity result easily follows.

Theorem 10. The problem of deciding whether a formula is consistent in the logic IP is NP-complete.

Proof. The problem is clearly NP-hard, as it generalizes the satisfiability problem for propositional logic. It is not hard to see that it is in NP; a certificate which witnesses that a formula θ is satisfiable is the small model from Theorem 9, which is of polynomially bounded size in $|\theta|$; and we can check in polynomial time that θ is true in this model. \square

Discussion

As we have shown, the preferential approach to reasoning about relative likelihood is quite compatible with a probabilistic approach. Beginning with a preorder on states, we can reason about event comparisons in a way that coheres perfectly with quantitative probabilistic reasoning.

These results are of interest not only for the field of representation and reasoning within AI, but also in other domains such as theoretical linguistics. In the area of natural language semantics, for example, what we have called the *inflationary function* semantics was independently proposed as a model of epistemic comparatives, e.g., “at least as likely as” in English (Kratzer 1991). While this is evidently the dominant approach in the field, several authors have argued that the validity of principle (L4) disqualifies it as a model of ordinary speakers’ intuitions about relative likelihood. The specific question has been raised as to whether there is an alternative lifting operation—that is, a different preferential approach—whose logic fits better with ordinary intuitions (Yalcin 2010). Insofar as IP, the logic of sets of measures, matches intuitions, our results answer this question.

More generally, the present work can be seen as part of a broader project to explore possible ways of unifying logical and probabilistic approaches to AI (Russell 2015). In addition to combining logical and probabilistic tools, another important strand of this project is to clarify when well-understood qualitative tools—such as preferential structures—and familiar forms of probabilistic reasoning—such as reasoning about sets of measures—can already be seen as two sides of the same logical coin.

References

- Alon, S., and Heifetz, A. 2014. The logic of Knightian games. *Economic Theory Bulletin* 2(2):161–182.
- Alon, S., and Lehrer, E. 2014. Subjective multi-prior probability: A representation of a partial likelihood relation. *Journal of Economic Theory* 151:476–492.
- Baltag, A., and Smets, S. 2008. A qualitative theory of dynamic belief revision. In Bonanno, G.; van der Hoek, W.; and Wooldridge, M., eds., *Logic and the Foundations of Game and Decision Theory*, Texts in Logic and Games 3. Amsterdam: Amsterdam University Press. 13–60.
- Board, O. 2004. Dynamic interactive epistemology. *Games and Economic Behavior* 49(1):49–80.
- del Cerro, L. F., and Herzog, A. 1991. A modal analysis of possibility theory. *Lecture Notes in Computer Science* 535:11–18.
- Delgrande, J. P., and Renne, B. 2015. The logic of qualitative probability. In Yang, Q., and Wooldridge, M., eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)*, 2904–2910. Palo Alto: AAAI Press.
- Dubois, D.; Fargier, H.; and Perny, P. 2003. Qualitative decision theory with preference relations and comparative uncertainty: An axiomatic approach. *Artificial Intelligence* 138:219–260.
- Dubois, D.; Fargier, H.; and Prade, H. 1998. Comparative uncertainty, belief functions, and accepted beliefs. In Cooper, G., and Moral, S., eds., *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence (UAI '98)*, 113–120. San Francisco: Morgan Kaufmann.
- Fagin, R.; Halpern, J. Y.; and Megiddo, N. 1990. A logic for reasoning about probabilities. *Information and Computation* 87:78–128.
- de Finetti, B. 1937. La prévision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré* 7:1–68.
- Friedman, N., and Halpern, J. Y. 1994a. Conditional logics of belief change. In *Proceedings of AAAI-94*, 915–921. Menlo Park: AAAI Press.
- Friedman, N., and Halpern, J. Y. 1994b. On the complexity of conditional logics. In Doyle, J.; Sandwell, E.; and Torasso, P., eds., *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR '94)*. San Francisco: Morgan Kaufmann. 202–213.
- Gärdenfors, P. 1975. Qualitative probability as an intensional logic. *Journal of Philosophical Logic* 4(2):171–185.
- Ghosh, S., and de Jongh, D. 2013. Comparing strengths of beliefs explicitly. *Logic Journal of the IGPL* 21(3):488–514.
- Halpern, J. Y., and McAllester, D. A. 1989. Likelihood, probability, and knowledge. *Computational Intelligence* 5(2):151–160.
- Halpern, J. Y., and Rabin, M. O. 1987. A logic to reason about likelihood. *Artificial Intelligence* 32(3):379–405.
- Halpern, J. Y. 1996. Defining relative likelihood in partially-ordered preferential structures. In Horvitz, E., and Jensen, F., eds., *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence (UAI '96)*, 299–306. San Francisco: Morgan Kaufmann.
- Halpern, J. Y. 1997. Defining relative likelihood in partially-ordered preferential structures. *Journal of Artificial Intelligence Research* 7:1–24.
- Halpern, J. Y. 2003. *Reasoning about Uncertainty*. Cambridge, Mass.: MIT Press.
- Harrison-Trainor, M.; Holliday, W. H.; and Icard, T. F. 2016a. Inferring Probability Comparisons. *Mathematical Social Sciences*. Forthcoming. Preprint at <http://escholarship.org/uc/item/8br2b074>.
- Harrison-Trainor, M.; Holliday, W. H.; and Icard, T. F. 2016b. A note on cancellation axioms for comparative probability. *Theory and Decision* 80(1):159–166.
- Holliday, W. H., and Icard, T. F. 2013. Measure semantics and qualitative semantics for epistemic modals. In Snider, T., ed., *Proceedings of Semantics and Linguistic Theory (SALT) 23*. LSA and CLC Publications.
- Katsuno, H., and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.
- Kraft, C. H.; Pratt, J. W.; and Seidenberg, A. 1959. Intuitive probability on finite sets. *The Annals of Mathematical Statistics* 30(2):408–419.
- Kratzer, A. 1991. Modality. In von Stechow, A., and Wunderlich, D., eds., *Semantics: An International Handbook of Contemporary Research*. Berlin: de Gruyter. 639–650.
- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence* 44(2):167–207.
- Lewis, D. 1973. *Counterfactuals*. Oxford: Basil Blackwell.
- Narens, L. 2007. *Theories of Probability: An Examination of Logical and Qualitative Foundations*. London: World Scientific Publishing.
- Russell, S. 2015. Unifying logic and probability. *Communications of the ACM* 58(7):88–97.
- Scott, D. 1964. Measurement structures and linear inequalities. *Journal of Mathematical Psychology* 1:233–247.
- Seegerberg, K. 1971. Qualitative probability in a modal setting. In Fenstad, E., ed., *Second Scandinavian Logic Symposium*, 341–352. Amsterdam: North-Holland.
- Shoham, Y. 1987. A semantical approach to nonmonotonic logics. In *Proceedings of the 2nd Symposium on Logic in Computer Science*, 275–279.
- Touazi, F.; Cayrol, C.; and Dubois, D. 2015. Possibilistic reasoning with partially ordered beliefs. *Journal of Applied Logic* 13(4):770–798.
- van der Hoek, W. 1996. Qualitative modalities. *International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems* 4(1):45–59.
- Yalcin, S. 2010. Probability Operators. *Philosophy Compass* 5(11):916–937.