

A Logic of Extended Probability

Nic Wilson

School of Computing and Mathematical Sciences
Oxford Brookes University
Gipsy Lane, Headington, Oxford, OX3 0BP, U.K.
pnwilson@brookes.ac.uk

Abstract

This paper shows how the logic of gambles corresponding to Peter Walley's system of Imprecise Probability can be extended to allow gambles involving infinitesimal values and infinite values. This logic can then be used for reasoning with infinitesimal probabilities alongside conventional reasoning with linear constraints on probabilities. The proof theory is shown to be sound and complete for finite input sets.

Keywords. Imprecise Probability, Probabilistic Logic, infinitesimals, order of magnitude probabilistic reasoning.

1 Introduction

Very unlikely propositions can sometimes cause problems for theories of probability. The proposition may, for example, relate to an event that has not occurred before in our experience. If we assign a precise probability to such a proposition, then the value can be very arbitrary, and thus can lead to very arbitrary inferences, and arbitrary decisions. Even if a probability interval $[\delta, \epsilon]$ is assigned for such a proposition, related problems arise; the values δ and ϵ will tend to be also very arbitrary; if we're cautious, giving a wide interval, the inferences will often be disappointingly weak; if on the other hand, we give a narrow interval, this suffers from similar problems as assigning a precise probability; an interval of intermediate width can suffer from either (or even both) problems. Similarly, utility values of very large magnitudes can be problematic.

An alternative is to consider such values of probability as essentially infinitesimal, and such values of utility as of infinite magnitude. An earlier paper [11] constructed a theory of probability which allowed infinitesimal values of probability as well as the usual values. As shown, in [9, 10, 11], this theory of probability can be used in a simple way in the interpretation

and semantics of a number of theories of uncertain reasoning, for example, non-monotonic consequence relations [1, 4]; NCFs [7], possibility functions [3] and counterfactual probabilities [2].

In this paper, this work is developed into a logic that can reason with both conventional and infinitesimal probability; it is a unified framework that can be used for automated deduction with gambles (or linear constraints on probabilities), and also order of magnitude statements about probabilities. It thus unifies a probabilistic logic and a form of possibilistic reasoning.

Section 2 reiterates from [12] the expression of Imprecise Probability [8] as a logic. Section 3 describes the simple theory of infinitesimals given in [11]. Section 4 brings the two together, constructing a logic that can be used for reasoning with linear constraints on probabilities which can take both real and infinitesimal values; the proof of the main result is sketched in section 5. The proof theory is illustrated with an example in section 6. The logic can be also used for making a parametrised family of conventional probabilistic inferences as shown in section 7. Section 8 summarises the main results and discusses some avenues for further exploration.

2 The Logic of Gambles

This section gives the construction of the logic of gambles and basic properties of the logic, taken from [12] (the results are derived from results in [8]).

Let Ω be a finite set of possibilities, exactly one of which must be true. A gamble on Ω is a function from Ω to \mathbb{R} . If you were to accept gamble X and it was ω that turned out to be true then you would gain $X(\omega)$ utiles (so you would lose if $X(\omega) < 0$). An agent's beliefs are elicited by asking them to tell us a set \mathcal{G} of gambles they find acceptable, i.e., gambles they would be happy to accept.

For $\lambda \in \mathbb{R}$ we write λ for the constant gamble defined

by $\lambda(\omega) = \lambda$ for all $\omega \in \Omega$. If $\lambda > 0$ we should certainly consider gamble λ acceptable since, whatever happens, we gain. If $\lambda < 0$ then we should certainly not accept gamble λ since, whatever happens, we lose. Addition and subtraction of gambles is defined pointwise, so that for gambles X and Y , for each $\omega \in \Omega$, $(X - Y)(\omega) = X(\omega) - Y(\omega)$. For $A \subseteq \Omega$ we define gamble A to be the indicator function of A , so that $A(\omega) = 1$ if $\omega \in A$ and $A(\omega) = 0$ otherwise. Any such gamble A would seem acceptable since you couldn't lose. On the other hand, the gamble $A - 1$ (given by $(A - 1)(\omega) = 0$ if $\omega \in A$; $(A - 1)(\omega) = -1$ otherwise) would only be acceptable to you if you were certain that A were true. If $\Omega = \{\omega_1, \omega_2\}$ and X is the gamble $\{\omega_1\} - 0.5$ (so that $X(\omega_1) = 0.5$ and $X(\omega_2) = -0.5$) and you considered ω_1 more likely than ω_2 then it seems that you should consider the gamble X acceptable.

Gambles can represent a varied set of probability statements. Another way of viewing gambles is as linear constraints on an unknown chance distribution, and indeed any such linear inequality can be represented as a gamble. For example, the constraint $\Pr(A) \geq 0.6$ would be represented by the gamble $A - 0.6$; the constraint $\Pr(A) \leq \Pr(B) + 0.1$ would be represented by the gamble $B + 0.1 - A$; the constraint $\Pr(A|B) \leq 0.9$ (which is taken to satisfied if $\Pr(B) = 0$) by the gamble $0.9B - (A \cap B)$. In general a linear constraint $\lambda_1 \Pr(A_1) + \dots + \lambda_k \Pr(A_k) \geq \mu$ is represented by the gamble $X = \lambda_1 A_1 + \dots + \lambda_k A_k - \mu$. The constraint $\lambda_1 \Pr(A_1) + \dots + \lambda_k \Pr(A_k) \leq \mu$ is equivalent to the constraint $(-\lambda_1) \Pr(A_1) + \dots + (-\lambda_k) \Pr(A_k) \geq -\mu$ and so is represented by the gamble $-X$. Hence $\lambda_1 \Pr(A_1) + \dots + \lambda_k \Pr(A_k) = \mu$ is represented by the pair of gambles $\{X, -X\}$. The language does, however, have some limitations: independence relationships or constraints such as $\Pr(A|B) \leq \Pr(A)$ cannot be represented simply in terms of gambles, since they are non-linear.

The Language

Let \mathcal{L} be the set of gambles on Ω . The constant gamble -1 is written as \perp .

Proof Theory

We are going to define an inference relation \vdash , where for $\mathcal{C} \subseteq \mathcal{L}$ and $X \in \mathcal{L}$, $\mathcal{C} \vdash X$ is intended to mean 'Given that I'm prepared to accept any gamble in \mathcal{C} , then I should also be prepared to accept X '. The axioms and inference rules used to define \vdash will be justified by the semantics (with probability distributions as models) given below. However, they can be justified directly [8].

Axiom Schema: X , for any X with $\min_{\omega \in \Omega} X(\omega) \geq 0$.

For any such X we can't lose so it should be acceptable.

Inference Rule (Schema) 1: For any $\lambda \in \mathbb{R}$ such that $\lambda \geq 0$ the inference rule *From X deduce λX* .

This relates to the situation where stakes are multiplied by a factor λ .

Inference Rule 2: *From X and Y deduce $X + Y$* .

This relates to the combination of two gambles.

Then, in the usual way, we say that $\mathcal{C} \subseteq \mathcal{L}$ proves $X \in \mathcal{L}$, abbreviated to $\mathcal{C} \vdash X$, if there is a finite sequence X_1, \dots, X_n of gambles in \mathcal{L} with $X_n = X$ and each X_i is either an element of \mathcal{C} , an axiom or is produced by an inference rule from earlier elements in the sequence.

For $\mathcal{C} \subseteq \mathcal{L}$ define $\text{TH}(\mathcal{C})$ to be the set $\{X \in \mathcal{L} : \mathcal{C} \vdash X\}$. The relation \vdash is reflexive i.e., $\text{TH}(\mathcal{C}) \supseteq \mathcal{C}$; it is monotonic i.e., $\mathcal{C} \subseteq \Theta \Rightarrow \text{TH}(\mathcal{C}) \subseteq \text{TH}(\Theta)$; it is transitive, i.e., $\text{TH}(\text{TH}(\mathcal{C})) \subseteq \text{TH}(\mathcal{C})$ (and therefore $\text{TH}(\text{TH}(\mathcal{C})) = \text{TH}(\mathcal{C})$) and compact i.e., if $\mathcal{C} \vdash X$ then there exists finite $\mathcal{C}' \subseteq \mathcal{C}$, with $\mathcal{C}' \vdash X$.

Semantics

The set of models \mathcal{M} of \mathcal{L} is defined to be the set of probability distributions on Ω . For probability distribution P and gamble X , $P(X)$ is defined to be the expected value of X , i.e., $\sum_{\omega \in \Omega} P(\omega)X(\omega)$. We say that probability distribution P satisfies gamble X (written $P \models X$) iff $P(X) \geq 0$, that is, if and only if the expected utility is non-negative, i.e., iff we would expect (in the long run) not to lose money from X if P represented objective chances. Note that P satisfies the constraint $\lambda_1 P(a_1) + \dots + \lambda_k P(a_k) \geq \mu$ if and only if $P \models X$ where $X = \lambda_1 a_1 + \dots + \lambda_k a_k - \mu$ is the gamble representing the constraint.

We extend \models in the usual way: P satisfies set of gambles \mathcal{C} (written $P \models \mathcal{C}$) if and only if it satisfies every gamble X in \mathcal{C} , and the semantic entailment relation \models is defined by $\mathcal{C} \models X$ iff for all models P , $[P \models \mathcal{C} \Rightarrow P \models X]$.

The relation \models is reflexive, transitive, monotonic, but not compact.

The proof theory is sound with respect to the model theory:

Theorem 1 (Soundness) *For gamble X and set of gambles \mathcal{C} , $\mathcal{C} \vdash X \Rightarrow \mathcal{C} \models X$.*

Completeness holds for finite sets of premises:

Theorem 2 (Finite Completeness) *For gamble X and finite set of gambles \mathcal{G} , if $\mathcal{G} \models X$ then $\mathcal{G} \vdash X$.*

We come close to completeness even for infinite sets of premises:

Theorem 3 (Almost-Completeness) *For gamble X and set of gambles \mathcal{G} , if $\mathcal{G} \models X$ then $\mathcal{G} \vdash X + \epsilon$ for all $\epsilon > 0$.*

Corollary (Consistency Completeness) *For $\mathcal{G} \subseteq \mathcal{L}$, $\mathcal{G} \vdash \perp \iff \mathcal{G} \models \perp$.*

3 A Simple Theory of Infinitesimals

This section gives the construction of the simple theory of infinitesimals defined in [11]. The extended reals \mathbb{R}^* are formed by adding an ‘infinitesimal’ element ϵ , to the real numbers. Probability is then defined using the usual axioms.

3.1 The Extended Reals \mathbb{R}^*

Let the extended reals \mathbb{R}^* be $\mathbb{R}(\epsilon)$, the field of rational functions in (dummy variable) ϵ over the field \mathbb{R} (see [5], page 122). Each element of \mathbb{R}^* can be written as a pair p/q where p and q are polynomial functions in ϵ , and p/q represents the same element of \mathbb{R}^* as r/s if and only if ps is the same polynomial as qr . \mathbb{R}^* clearly contains a copy of \mathbb{R} : for $x \in \mathbb{R}$, the ratio of constant polynomials $x/1$ is in \mathbb{R}^* , and we’ll denote this element of \mathbb{R}^* also by x . In particular the element 0 of \mathbb{R}^* is the function which has constant value 0.

Every non-zero element r of \mathbb{R}^* can be uniquely expressed as $\bar{r}\epsilon^{\hat{r}}r'$, where $\bar{r} \in \mathbb{R} \setminus \{0\}$, \hat{r} is in \mathbb{Z} , the set of integers, and $r' \in \mathbb{R}^*$ is such that $r'(0) = 1$. Define $\hat{0} = \infty$. The function $r \mapsto \hat{r}$ gives the order of magnitude (in terms of powers of ϵ) of element r of \mathbb{R}^* . Also for $r \in \mathbb{R}^*$ let \underline{r} be $\lim_{\epsilon \rightarrow 0} r(\epsilon)$, so that

$$\underline{r} = \begin{cases} 0 & \text{if } \hat{r} > 0; \\ \infty & \text{if } \hat{r} < 0 \text{ and } \bar{r} > 0; \\ -\infty & \text{if } \hat{r} < 0 \text{ and } \bar{r} < 0; \\ \bar{r} & \text{if } \hat{r} = 0. \end{cases}$$

If $r \in \mathbb{R}$ then $\underline{r} = r$. For $r \in \mathbb{R}^*$, \underline{r} will be referred to as ‘the real part of r ’. Extended reals r and s will be said to be approximately equal if $\underline{r - s} = 0$.

3.2 The Ordering on \mathbb{R}^*

If $r = p/q \in \mathbb{R}^*$ where p, q are polynomials in ϵ , then, for $x \in \mathbb{R}$, $r(x)$ (the value of r when x is substituted for ϵ) is a real number, if $q(x) \neq 0$. For $r, s \in \mathbb{R}^*$, define relation $>$ by $r > s$ if and only if there exists strictly positive real number y such that $r(x) > s(x)$ for all real x with $0 < x < y$.

ϵ is intended to be a very small positive number, so we define $r > s$ iff r is bigger than s for small enough ϵ . Relations $<$, \geq and \leq are defined from relation $>$ in the usual way, e.g., $r \leq s$ if and only if $s > r$ or $s = r$. This ordering makes \mathbb{R}^* an ordered field (see [5], page 261).

Though they are defined as functions, elements of \mathbb{R}^* should be thought of as numbers; ϵ is a positive number smaller than any strictly positive real number, ϵ^2 is an even smaller positive number, $\epsilon - \epsilon^2$ is between the two, though much closer to ϵ , and so on.

The one axiom of the real numbers that \mathbb{R}^* lacks is the completeness axiom, that any set bounded above has a least upper bound. However, this is the case for any ordered field extension of \mathbb{R} containing an infinitesimal. Consider any ordered field extension R of the reals \mathbb{R} . Let \mathcal{E} be the set of positive infinitesimals in R , i.e., all $r \in R$ such that $r > 0$ and for all strictly positive reals x , $r < x$. Assume $\mathcal{E} \neq \emptyset$ (i.e., that R contains an infinitesimal). \mathcal{E} is bounded above by any strictly positive real. Suppose it has a least upper bound a . a must be in \mathcal{E} because if for some strictly positive $x \in \mathbb{R}$, $x \leq a$ then $x/2 < a$ and $x/2$ is an upper bound for \mathcal{E} , which contradicts a being the least upper bound. However $a \in \mathcal{E}$ implies that $2a \in \mathcal{E}$, contradicting a being an upper bound for \mathcal{E} .

3.3 Extended Probability

To define extended probability and utility, the usual definitions suffice, except using \mathbb{R}^* instead of \mathbb{R} .

An extended utility function (or extended gamble) on Ω is a function from Ω to \mathbb{R}^* . An extended probability measure P over Ω is defined to be a function from 2^Ω to \mathbb{R}^* , satisfying $P(\Omega) = 1$, and for $A, B \subseteq \Omega$ such that $A \cap B = \emptyset$, $P(A \cup B) = P(A) + P(B)$.

Extended probability theory has many of the same properties as probability theory, since it shares all the axioms, apart from the completeness axiom of the real numbers.

An interpretation of extended probability is given in [11]. Briefly, a proposition E with very small but unknown probability is assigned the value ϵ . The statement ‘the probability of A is ϵ ’ is then interpreted as that A and E are (precisely) equiprobable. The meaning of statements of the form ‘The probability of A is r ’, for various extended reals r , is constructed using a sequence of thought experiments using independent, mutually exclusive, or conditional propositions; for example, ‘the probability of A is ϵ^2 ’ is interpreted as meaning that A is equiprobable with $E \wedge F$ where F is some proposition independent of E and equiprobable with E .

4 A Logic of Extended Gambles

We can extend the language, proof theory and semantics of the logic of gambles to allow infinitesimals and infinities just by replacing \mathbb{R} by \mathbb{R}^* in the various definitions.

Hence the language \mathcal{L}^* is defined to be the set of extended gambles, i.e., the set of functions from Ω to \mathbb{R}^* . As before, \perp is an abbreviation for the constant extended gamble -1 .

Proof theory

The proof theory is generated by the following axiom schema and inference rules:

Axiom Schema: X , for any X with $\min_{\omega \in \Omega} X(\omega) \geq 0$.

Inference Rule (Schema) 1: For any $\lambda \in \mathbb{R}^*$ such that $\lambda \geq 0$ the inference rule *From X deduce λX* .

Inference Rule 2: *From X and Y deduce $X + Y$* .

So, in the usual way, we say that $\mathcal{M} \subseteq \mathcal{L}^*$ proves $X \in \mathcal{L}^*$, abbreviated to $\mathcal{M} \vdash^* X$, if there is a finite sequence X_1, \dots, X_n of extended gambles in \mathcal{L}^* with $X_n = X$ and each X_i is either an element of \mathcal{M} , an axiom, or is produced by an inference rule from earlier elements in the sequence.

Like \vdash , the relation \vdash^* is reflexive, monotonic, transitive and compact.

Semantics

The set of models \mathcal{M}^* of \mathcal{L}^* is defined to be the set of extended probability distributions on Ω , i.e., functions P from Ω to \mathbb{R}^* such that for all $\omega \in \Omega$, $P(\omega) \geq 0$ and $\sum_{\omega \in \Omega} P(\omega) = 1$. For extended probability distribution P and extended gamble X , $P(X)$ is defined, as before, to be the expected value of X , i.e., $\sum_{\omega \in \Omega} P(\omega)X(\omega)$; we say that probability distribution P satisfies extended gamble X (written $P \models^* X$) iff $P(X) \geq 0$. The relation \models^* is extended in the usual way: P satisfies set of extended gambles \mathcal{M} , written $P \models^* \mathcal{M}$, if and only if it satisfies every extended gamble X in \mathcal{M} , and the semantic entailment relation \models^* is defined by $\mathcal{M} \models^* X$ iff for all models P , $[P \models^* \mathcal{M} \Rightarrow P \models^* X]$.

Set of extended gambles \mathcal{M} is said to be consistent if it has a model, i.e., if there exists $P \in \mathcal{M}^*$ such that $P \models^* \mathcal{M}$. It can be seen that \mathcal{M} is inconsistent if and only if $\mathcal{M} \models^* \perp$.

Like \models , the relation \models^* is reflexive, transitive and monotonic, and not compact. A counter-example to compactness is when $\mathcal{M} = \{(A - 1 + \varepsilon^n) : n \in \mathbb{N}\}$,

where A is some non-empty proper subset of Ω . If $P \models^* \mathcal{M}$, then for all n , $P(A - 1 + \varepsilon^n) \geq 0$, so for all n , $P(A) \geq 1 - \varepsilon^n$. This implies that $P(A) = 1$, and so we have $\mathcal{M} \models^* A - 1$. However, if \mathcal{M}' is any finite subset of \mathcal{M} , then it is not the case that $\mathcal{M}' \models^* A - 1$ (contradicting compactness)—since we can choose P such that $P(A) = 1 - \varepsilon^n$ for some n large enough so that $P \models^* \mathcal{M}'$, but P isn't a model of $A - 1$, as $P(A - 1) = -\varepsilon^n < 0$.

We still have Soundness:

Theorem 4 (Soundness of \vdash^* w.r.t. \models^*) *For extended gamble X and set of extended gambles \mathcal{M} , $\mathcal{M} \vdash^* X \Rightarrow \mathcal{M} \models^* X$.*

The proof is straight-forward: we just check the soundness of the axioms and inference rules; if Z is one of the axiom schema, then we certainly have $P \models^* Z$ for any model P ; if $\lambda \geq 0$ and $P \models^* X$ and $P \models^* Y$ then $P \models^* \lambda X$ and $X + Y$. The result then follows using induction on the length of a proof in the logic.

As before, we cannot hope to have full completeness of \vdash^* with respect to \models^* because \models^* is not compact. However we do have completeness for finite sets \mathcal{M} . The following is the main result of the paper:

Theorem 5 (Finite Completeness of \vdash^* w.r.t. \models^*) *For extended gamble X and finite set of extended gambles \mathcal{M} , if $\mathcal{M} \models^* X$ then $\mathcal{M} \vdash^* X$.*

The proof is sketched in the next section.

Unfortunately we do not have *Almost-Completeness* (c.f. section 2) as the following example shows.

Counter-Example to Almost-Completeness

Let \mathbb{R}_+^* be the set of strictly positive elements of \mathbb{R}^* , i.e., $\{r \in \mathbb{R}^* : r > 0\}$. We say that $\delta \in \mathbb{R}_+^*$ is a positive infinitesimal in \mathbb{R}^* if $\delta < x$ for all $x \in \mathbb{R}$ with $x > 0$. Let \mathcal{E} be the set of positive infinitesimals in \mathbb{R}^* .

Let A be some non-empty proper subset of Ω . Define the set of extended gambles \mathcal{M}_1 to be $\{A - \delta : \delta \in \mathcal{E}\}$ and \mathcal{M}_2 to be $\{r - A : r \in \mathbb{R}_+^* - \mathcal{E}\}$.

Now $\mathcal{M} = \mathcal{M}_1 \cup \mathcal{M}_2$ is inconsistent, i.e., it has no model: for suppose $P \models^* \mathcal{M}$, and let $s = P(A)$; then $P \models^* \mathcal{M}_1$ which implies $s \geq \delta$ for all $\delta \in \mathcal{E}$; also $P \models^* \mathcal{M}_2$, so $s \leq r$ for all $r \in \mathbb{R}_+^* - \mathcal{E}$; but there is no value $s \in \mathbb{R}^*$ satisfying these constraints (the second constraint implies $s \in \mathcal{E}$; but then $s < 2s \in \mathcal{E}$ which contradicts the first constraint).

However, it can be shown that it is not the case that $\mathcal{M} \vdash^* \perp$.

This example thus shows that it not the case that

, $\vdash^* \perp \iff$, $\models^* \perp$, so it is a counter-example to Consistency Completeness.

If Almost-Completeness held, then , $\models^* \perp$ would imply that , $\vdash^* \perp + \frac{1}{2}$; however, , $\vdash^* \perp + \frac{1}{2}$ implies (using inference rule 1) , $\vdash^* \perp$ which, as observed above, is not the case. Therefore Almost Completeness does not hold.

5 Sketch of Proof of Finite Completeness (Theorem 5)

A subset of \mathcal{L}^* is said to be a *convex cone* (over \mathbb{R}^*) if it is closed under multiplication by non-negative scalars (in \mathbb{R}^*) and closed under addition, i.e., iff it is closed under the two inference rule schemas of the proof theory. For , $\subseteq \mathcal{L}^*$, define $C(,)$ to be the unique smallest convex cone containing , . Every element of $C(,)$ can then be expressed as a linear combination of the elements in , with non-negative co-efficients. A convex cone Δ is said to be *finite* if $\Delta = C(,)$ for some finite , .

For $T, X \in \mathcal{L}^*$ define TX to be $\sum_{\omega \in \Omega} T(\omega)X(\omega)$. For , $\subseteq \mathcal{L}^*$, define , $^+$ (known as the *dual cone* of ,) to be $\{T \in \mathcal{L}^* : TX \geq 0 \text{ for all } X \in , \}$.

Theorem 6 (Finite \mathbb{R}^* -convex cones are reflexive) *If Δ is a finite convex cone over \mathbb{R}^* then it is reflexive, i.e., $(\Delta^+)^+ = \Delta$.*

This is the analogue of a fundamental property of convex cones over \mathbb{R} , which requires a lengthy proof; the proof of this in [6] was used to prove Theorem 6: each step of the proof was examined (where necessary breaking it down into components) to confirm that the step is valid also for \mathbb{R}^* -convex cones (the key point to check was that the completeness property of \mathbb{R} was not used, only the properties of an ordered field, which are shared by \mathbb{R}^*).

For , $\subseteq \mathcal{L}^*$, define $\text{TH}(,)$ to be the syntactic consequences of , , i.e., $\{X : , \vdash^* X\}$, and define $\text{Th}(,)$ to be the semantic consequences of , , i.e., $\{X : , \models^* X\}$. Theorem 4 (Soundness) is equivalent to the statement that for all , , $\text{TH}(,) \subseteq \text{Th}(,)$. Theorem 5 (Finite Completeness) is equivalent to: for all finite , , $\text{Th}(,) \subseteq \text{TH}(,)$.

Proposition 1

Let $\mathcal{R} = \{R \in \mathcal{L}^* : \text{for all } \omega \in \Omega, R(\omega) \geq 0\}$. Let , be a subset of \mathcal{L}^* . Define $[,]$ to be , $^+ \cap \mathcal{R}$. For each $\omega \in \Omega$ define the indicator gamble I_ω by $I_\omega(\omega') = 0$ if $\omega \neq \omega'$, and $I_\omega(\omega) = 1$. Let $\mathcal{I} = \{I_\omega : \omega \in \Omega\}$.

- (i) $\text{Th}(,) = [,]^+$;
- (ii) $[,] = [\text{TH}(,)]$;

(iii) $\text{TH}(,) = C(, \cup \mathcal{I})$;

(iv) $[C(, \cup \mathcal{I})] = (C(, \cup \mathcal{I}))^+$;

(v) $\text{Th}(,) = ((\text{TH}(,))^+)^+$.

Sketch of proof: (i) follows easily using the observation that the models of , are normalised elements of $[,]$. (ii) is a version of Soundness and can be proved in the same way. (iii) is immediate from the definitions. (iv) holds because if $TX \geq 0$ for all $X \in \mathcal{I}$ then $T \in \mathcal{R}$. (v): applying, in turn, equalities (i), (ii), (iii), (iv) and (iii) gives $\text{Th}(,) = [,]^+ = [\text{TH}(,)]^+ = [C(, \cup \mathcal{I})]^+ = ((C(, \cup \mathcal{I}))^+)^+ = ((\text{TH}(,))^+)^+$.

Proof of Theorem 5: As , is finite, by Proposition 1(iii), $\text{TH}(,)$ is a finite convex cone, so is reflexive by Theorem 6: $\text{TH}(,) = ((\text{TH}(,))^+)^+$. Hence, by Proposition 1(v), $\text{Th}(,) = \text{TH}(,)$, proving Finite Completeness.

6 Example

The following is an example to illustrate the mechanics of the proof theory.

Suppose we are given the following constraints on extended probabilities:

$$P(A|B) \leq \varepsilon^3, \quad P(B|C) \leq \varepsilon, \quad P(A|A \cup \bar{C}) \geq 0.1,$$

and the logical piece of information $A \subseteq B \cap C$,

where \bar{C} is the negation of the proposition C , and a constraint is taken to be satisfied if the probability is undefined (owing to the denominator being zero).

What upper bounds can we give for $P(A|C)$?

Firstly we must represent propositions A, B and C as subsets of a set of possibilities Ω . A natural way of doing this that doesn't assume anything more than is given in the problem is to set $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5\}$ and $A = \{\omega_1\}$, $B = \{\omega_1, \omega_2, \omega_3\}$, $C = \{\omega_1, \omega_2, \omega_4\}$. (Other representations of Ω, A, B and C are of course possible, but the choice doesn't affect derived probabilities. Although the choice of Ω etc. is not part of the logic, it could be automated.)

The following notation is used for gambles: $(9, 8, 7, 6, 5)$ is taken to be the gamble X with $X(\omega_1) = 9$, $X(\omega_2) = 8$ etc. The above inequalities must be converted into gambles. The constraint $P(A|B) \leq \varepsilon^3$ is equivalent to $\varepsilon^3 P(B) - P(A \cap B) \geq 0$, and hence can be expressed as the gamble $X_1 = \varepsilon^3 B - A \cap B$, which may be written more explicitly as $(\varepsilon^3 - 1, \varepsilon^3, \varepsilon^3, 0, 0)$. Thus P satisfies the first constraint if and only if $P \models^* X_1$.

Similarly, the second constraint can be represented by the gamble $X_2 = (\varepsilon - 1, \varepsilon - 1, 0, \varepsilon, 0)$, and

the third constraint can be represented by $X_3 = (0.9, 0, -0.1, 0, -0.1)$. Let $\mathcal{G} = \{X_1, X_2, X_3\}$.

From X_3 can be derived $Y_1 = (0.9, 0, -0.1, 0, 0)$: this is because $Y_1 - X_3$ is non-negative, so is one of the axioms in the axiom schema, and we can then derive using inference rule 2, $Y_1 = X_3 + (Y_1 - X_3)$. The first inference rule then enables us to derive $10Y_1$ and $\frac{1}{\varepsilon^3}X_1$, so we can derive (using inference rule 2) $Y_2 = 10Y_1 + \frac{1}{\varepsilon^3}X_1$, which equals $(10 - \frac{1}{\varepsilon^3}, 1, 0, 0, 0)$. We can add this to X_2 to derive $Y_3 = (\varepsilon + 9 - \frac{1}{\varepsilon^3}, \varepsilon, 0, \varepsilon, 0)$, and we're essentially there (since we've made $Y_3(\omega_3) = Y_3(\omega_5) = 0$ and $Y_3(\omega_2) = Y_3(\omega_4)$).

For any model P of \mathcal{G} , by the soundness of the proof theory, $P \models^* Y_3$, i.e., $(9 - \frac{1}{\varepsilon^3})P(\omega_1) + \varepsilon(P(\omega_1) + P(\omega_2) + P(\omega_4)) \geq 0$, so $(9 - \frac{1}{\varepsilon^3})P(A \cap C) + \varepsilon P(C) \geq 0$, which can be rearranged to give

$$P(A|C) \leq \frac{\varepsilon^4}{1 - 9\varepsilon^3} \quad (\text{or } P(C) = 0).$$

This upper bound is just slightly more than ε^4 , and less than e.g., $1.001\varepsilon^4$.

In fact, this is best upper bound we can put on $P(A|C)$ because it can be shown that there is a model P of \mathcal{G} , with $P(A|C) = \frac{\varepsilon^4}{1 - 9\varepsilon^3}$.

7 Substituting Real Values for ε

Since the interpretation of extended probability given in [11] involves considering ε as a small positive unknown real number, one might wonder, if we take an inference $\mathcal{G} \vdash^* X$, whether we can replace all occurrences of ε by a real number x and obtain a sound (conventional) probabilistic inference (as defined in section 2).

Let $X[\varepsilon/x]$ be the gamble obtained when extended gamble X has each occurrence of ε replaced by a positive real number x . Similarly, if \mathcal{G} is a set of extended gambles, let $\mathcal{G}[\varepsilon/x] = \{Y[\varepsilon/x] : Y \in \mathcal{G}\}$. In fact $X[\varepsilon/x]$ may not be well-defined, since some component may have a zero denominator; for example if X is the constant extended gamble $\varepsilon/(1 - \varepsilon)$ then $X[\varepsilon/1]$ is not well-defined. However, it is always the case that for all small enough x , $X[\varepsilon/x]$ is well-defined (because Ω is finite). Hence, if \mathcal{G} is finite, for all small enough x , $\mathcal{G}[\varepsilon/x]$ is a set of well-defined gambles.

Lemma 1 *Let \mathcal{G} be a finite set of extended gambles, and suppose that $\mathcal{G} \vdash^* X$ for some extended gamble X . Then there exists strictly positive $y \in \mathbb{R}$ such that for all $0 < x < y$, $\mathcal{G}[\varepsilon/x] \vdash X[\varepsilon/x]$.*

This result also holds for infinite \mathcal{G} , as long as there exists some strictly positive $y' \in \mathbb{R}$ such that for all $Y \in \mathcal{G}$, $Y[\varepsilon/x]$ is well-defined for all $0 < x < y'$; this

condition is required to ensure that for small enough x , $\mathcal{G}[\varepsilon/x]$ is a well-defined set of gambles.

Sketch of proof: By definition of \vdash^* there exists a sequence of extended gambles X_1, X_2, \dots, X_n with $X_n = X$, where each X_i is either an element of \mathcal{G} , an axiom, or is produced from earlier elements in the sequence using one of the two inference rules. It can be shown inductively that, if x is small enough, $X_1[\varepsilon/x], X_2[\varepsilon/x], \dots, X_n[\varepsilon/x]$ is a proof of $X[\varepsilon/x]$ from $\mathcal{G}[\varepsilon/x]$ in the logic of gambles.

In fact, for any particular proof of X from \mathcal{G} , by considering each step, one can generate (often easily) a range of values of x which lead (via the substitution) to sound conventional probabilistic inferences. The logic of extended gambles can thus be used for making a parametrised family of probabilistic inferences of a more usual kind.

This can be illustrated using the example of section 6, where it can be easily checked that each inference in the proof is valid if ε is substituted for any strictly positive value, and so the whole proof is valid for any such substitution; however, the final rearrangement is only valid if $1 - 9\varepsilon^3 > 0$, i.e., if $\varepsilon < 1/\sqrt[3]{9}$.

Lemma 2 *Let $\mathcal{G} \cup \{X\}$ be a finite set of extended gambles. Suppose there exists strictly positive $y \in \mathbb{R}$ such that for all $0 < x < y$, $\mathcal{G}[\varepsilon/x] \vdash X[\varepsilon/x]$. Then $\mathcal{G} \vdash^* X$.*

Note that this result does not hold for infinite \mathcal{G} . Consider, for example, $\mathcal{G} = \{(1 - n\varepsilon) : n = 1, 2, 3, \dots\}$, and any X which is not an axiom (e.g., $X = \perp$). For any $x > 0$, $\mathcal{G}[\varepsilon/x]$ has no model, so $\mathcal{G}[\varepsilon/x] \vdash X[\varepsilon/x]$ (when $X[\varepsilon/x]$ is well-defined). However, \mathcal{G} is satisfied by every model, so it is not the case that $\mathcal{G} \vdash^* X$.

Sketch of proof of Lemma 2: Suppose $P \in \mathcal{M}^*$ is such that $P \not\models^* X$. We must show that $P \not\models^* X$. It can be shown that for all sufficiently small x , $P[\varepsilon/x] \models \mathcal{G}[\varepsilon/x]$; so then, by the hypothesis, for sufficiently small x , $P[\varepsilon/x] \models X[\varepsilon/x]$, i.e., $\sum_{\omega \in \Omega} P[\varepsilon/x](\omega) X[\varepsilon/x](\omega) \geq 0$. Hence $(\sum_{\omega \in \Omega} P(\omega) X(\omega))[\varepsilon/x] \geq 0$ for all sufficiently small x , which implies, using a basic property of the ordering \geq , that $\sum_{\omega \in \Omega} P(\omega) X(\omega) \geq 0$, i.e., $P \models^* X$, as required.

Theorem 7 *Let $\mathcal{G} \cup \{X\}$ be a finite set of extended gambles. The following statements are equivalent:*

(i) $\mathcal{G} \vdash^* X$;

(ii) $\mathcal{G} \models^* X$;

(iii) *there exists strictly positive $y \in \mathbb{R}$ such that for all $0 < x < y$, $\mathcal{G}[\varepsilon/x] \vdash X[\varepsilon/x]$;*

(iv) there exists strictly positive $y \in \mathbb{R}$ such that for all $0 < x < y$, $[\varepsilon/x] \models X[\varepsilon/x]$.

Proof: by Lemma 1, (i) \Rightarrow (iii); by Theorem 1, (iii) \Rightarrow (iv); Lemma 2 gives (iv) \Rightarrow (ii); finally, by Theorem 5, (ii) \Rightarrow (i).

Theorem 7 means that one can give an alternative interpretation of this logic of extended gambles: X follows from finite \mathcal{L} , if and only if for all sufficiently small x , one can infer (in a more conventional probabilistic way) $X[\varepsilon/x]$ from $\mathcal{L}[\varepsilon/x]$.

8 Summary and Discussion

A logic has been constructed which can reason with both conventional and infinitesimal probabilities. Surprisingly, perhaps, for such an expressive logic, it has a very simple proof theory which is complete for finite input sets.

The inferences made by this proof theory are guaranteed to be correct (in the sense of conventional probabilistic reasoning) if we replace the parameter ε by real-valued positive x , so long as x is small enough.

Use of the proof theory, however, may well often not be the most efficient method for deduction; it would be interesting to see to what extent linear programming techniques, for example, could be adapted for this logic.

It is clear that this logic doesn't completely solve the problem of very unlikely propositions; for example, if we want to represent two very small values, we have to precisely state their ratio. A natural extension of this work is therefore to add several incomparable infinitesimals to the reals as briefly discussed in [10]. Also, assuming that a probability of a proposition is arbitrarily small can lead to undesirable inferences: that you should bet against such a proposition at any finite odds. However, propagating an interval $(0, y)$ of valid real substitutions for ε (as discussed in section 7) will partly solve this problem by clarifying how small we're assuming ε to be for any given inference.

A technical issue is whether it is possible to bring proof (\vdash^*) and truth (\models^*) closer together. However we can't strengthen \vdash^* by adding extra inference rules: a consequence of Finite Completeness is that any sound inference rule is a derived inference rule of \vdash^* , so any additional inference rules (or axioms) are either unsound or redundant. An alternative idea is to weaken \models^* by increasing the set of models; a natural way to do so is to increase the range of the probability distributions: i.e., to allow probability values in some larger set \mathbb{R}^{**} (without increasing the language \mathcal{L}^*). Possible choices for \mathbb{R}^{**} include rational func-

tions in $\sqrt{\varepsilon}$ or even power series in $\sqrt{\varepsilon}$. Soundness is maintained even though the associated semantic entailment relations are weakened. However, it is not yet clear whether or not Almost-Completeness can be achieved in such a manner.

The close relationship (see e.g., [9]) between order of magnitude probabilities and possibility theory [3] means that this logic can be used as a possibilistic logic. Upper bounds on the possibility of A can be translated to sets of gambles of the form $\{A - N\varepsilon^n : N = 1, 2, \dots\}$, for a particular value of n ; lower bounds get translated to sets of gambles of the form $\{\frac{1}{N}\varepsilon^n - A : N = 1, 2, \dots\}$. Because these are infinite sets, we lose the completeness of the proof theory. However, a trick can be used to represent each of these infinite sets as a single statement, but this complicates the logic and it is not clear that completeness is restored.

Acknowledgements

Thanks to Peter Walley and the referees for their helpful comments.

References

- [1] E. Adams. Probability and the Logic of Conditionals. In *Aspects of Inductive Logic*, J. Hintikka and P. Suppes (eds.). North Holland, Amsterdam, 265–316, 1966.
- [2] C. Boutilier. The Probability of Possibility: Adding Uncertainty to Default Rules. *Proceedings of the Ninth Conference of Uncertainty in Artificial Intelligence (UAI93)*, David Heckerman and Abe Mamdani (eds.). Morgan Kaufmann Publishers, San Mateo, California, 461–468, 1993.
- [3] D. Dubois and H. Prade. *Possibility Theory: An Approach to Computerized Processing and Uncertainty*. Plenum Press, New York, 1988.
- [4] S. Kraus, D. Lehmann and M. Magidor. Non-monotonic Reasoning, Preferential Models and Cumulative Logics. *Artificial Intelligence* 44: 167–207, 1990.
- [5] S. MacLane and G. Birkhoff. *Algebra* (second edition). Macmillan, New York, 1979.
- [6] D. E. Nering. *Linear Algebra and Matrix Theory* (second edition). John Wiley and Sons, 1970.
- [7] W. Spohn. A General Non-Probabilistic Theory of Inductive Reasoning. *Uncertainty in Artificial*

Intelligence 4, R. D. Shachter et al. (eds.). Elsevier Science Publishers, 1990.

- [8] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.
- [9] N. Wilson. An Order of Magnitude Calculus. *Proceedings of the Eleventh Conference of Uncertainty in Artificial Intelligence (UAI95)*, Montreal, Canada, August 95. Morgan Kaufmann Publishers, San Mateo, California, 1995.
- [10] N. Wilson. *Extended Probability*. Unpublished report, available from the author on request, 1995.
- [11] N. Wilson. Extended Probability. *Proceedings of the 12th European Conference on Artificial Intelligence (ECAI96)*. John Wiley and Sons, 667-671, 1996.
- [12] N. Wilson and S. Moral. A Logical View of Probability. *Proceedings of the 11th European Conference on Artificial Intelligence, (ECAI-94)*, Amsterdam, The Netherlands, August 8-12, 386-390, 1994.