Logic, Methodology and Philosophy of Science

Proceedings of the Fifteenth International Congress

Hannes Leitgeb
Ilkka Niiniluoto
Päivi Seppälä
Elliott Sober
Editors

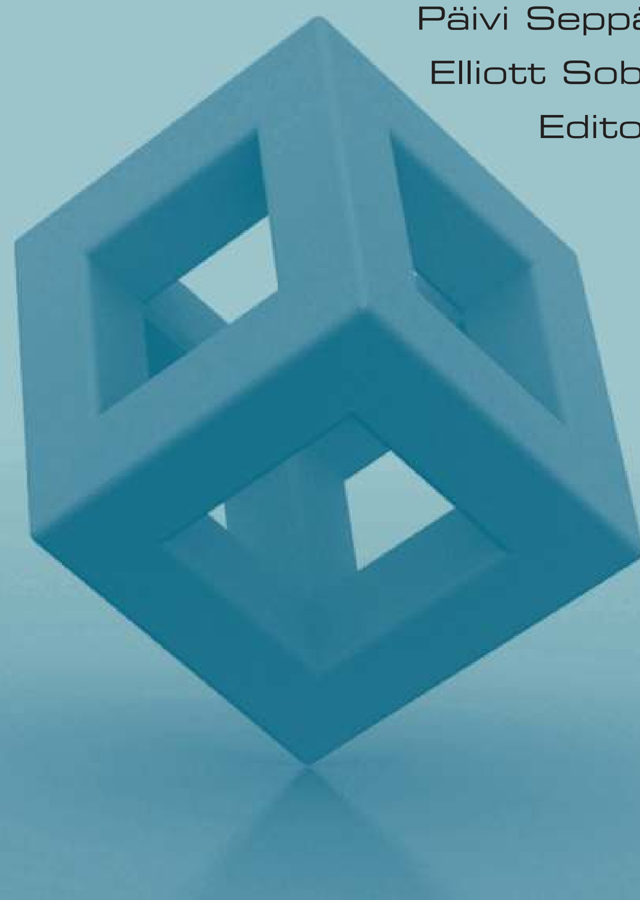

5 Logic revision: Some formal and semi-formal techniques for logic choice

EDWIN MARES *

Abstract. This paper sets out a probabilist theory of logic revision and discusses some of its consequences and the philosophical challenges that if faces. The probability theory is semantic. It is built upon a model for a very weak logic (the logic of bounded lattices) but that has sub-models for a wide variety of other logical systems – modal logics, relevant logics, linear logics, other substructural logics, fuzzy logic, and so on. The probability theory used is itself non-classical, although in those regions of the model in which classical logic holds probabilities act classically. Suggestions are made as to the treatments of debate and negotiation about logical rules and about the problems of logical omniscience and the apriority of logic.

Keywords: philosophy of probability, non-classical logic, non-classical probability, logic choice.

1 Introduction

Here is a problem. People can sensibly debate which logical system is right. They can take a purpose for logic, let's say its use in the evaluation of arguments that are presented as deductive, and argue which is the right logic for the job. These debates

*Philosophy Programme, Victoria University of Wellington, Wellington, New Zealand

appear reasonable. We don't know for certain which logic is the correct one for this purpose. Thus, we need a theory that makes sense of these debates and, more generally, about our reasoning about logic. But it would seem that such a theory needs to incorporate some sort of logic itself, because it needs to set rules of inference that can be used in such debates. Moreover, in order to allow that someone might hold (what turns out to be) the wrong logical views but be rational, it would seem that we want a theory that is neutral between all substantive logical positions.

The purpose of this paper is to outline one such theory. The theory that I construct is a form of probabilism. It uses a probabilistic mechanics. It treats ideal agents as having their degrees of beliefs, including their beliefs about logic, conform to the principles of a probability theory. This probability theory is formalised by a generalisation of the standard Kolmogorov axioms to fit with a wide range of classical and non-classical logics. The theory of belief revision presented employs a version of Jeffrey conditionalisation.

I present this theory to demonstrate how we can construe logic choice as a rational process. This is important, not only demonstrate that those who actually do choose logics – largely philosophers and mathematicians – can be construed to be rational, but also to show that the choice of such apparently foundational beliefs as those concerning logic can be understood as rational and revisable. Moreover, although I think that the probabilistic theory is defensible, the construction here also shows how other theories of belief revision, such as AGM based theories and dynamic doxastic logics can be used as bases for logic revision. In short, the probabilistic theory of logic revision is presented as a test case to show that many of our most fundamental beliefs can be treated as rational and revisable.

The plan of this paper is as follows: First, I set out the requirements on models used in the present theory and look at some logical systems for which there are semantics that can be used to construct models that fit these requirements. Second, construct a single *background model* that is used as a basis for the theory of belief revision. Third, I set out the the probability theory that corresponds to this class of models and define a Jeffrey update function. Fourth, I use a scorekeeping pragmatics as a framework for interpreting debates and negotiation about logic. Fifth, I look at two issues that arise from this theory of logic choice. One of these problems is the omniscience problem and the other concerns the apriority of logic.

I do not include any proofs of theorems in this paper. For these, see (Mares, 2014).

2 Logics and models

In order to place the logical systems in a single framework, I formulate them in a single language. This language has propositional variables $p, q, r, \ldots$, negation ($\neg$), conjunction ($\wedge$), disjunction ($\vee$), and implication ($\rightarrow$), as well as two propositional constants $\top$ and t. I define equivalence using implication and conjunction, as usual:

$A \leftrightarrow B =_{df} (A \to B) \wedge (B \to A)$. Standard formation rules apply. The logics that I can model using the techniques in this paper are ones that can be given an indexical semantics (a semantics in which truth is relativised to worlds in a frame) and for which propositions can be interpreted as sets of worlds in frames.

For each logic, L, there is a set of models, M_L. A model is a structure,

$$\mathfrak{M} = (W, N, [\![\cdot]\!], Prop),$$

where W is a non-empty set (of points, worlds, situations, ...), N is a non-empty subset of W, $[\![\cdot]\!]$ is a function from formulas into $\wp(W)$, and $Prop$ is a lattice of subsets of W that includes N, W, and Ø.

I also stipulate that the conjunction of each logic behaves semantically as set-theoretic intersection. I do not, however, assume that disjunction acts semantically as set-theoretic union. Instead, I merely require that disjunction behave as a least upper bound in the algebra of propositions. That is, in a model $\mathfrak{M}$, for any formulas A and B, there is no proposition in $\mathfrak{M}$ that is a proper superset of $[\![A]\!]$ and of $[\![B]\!]$ and also is a proper subset of $[\![A \vee B]\!]$. These requirements entail that the algebra of propositions on any model constitute a lattice, but this lattice need not be a distributive lattice. The avoidance of distributivity is because certain of the logics that I wish to model reject the rule of distribution of conjunction over disjunction, in particular, linear logic and orthologic.

A set of models M_L is the set of models for a logic L if and only if

1. A is a theorem of L if and only if for every $\mathfrak{M} \in M_L$, $N_{\mathfrak{M}} \subseteq [\![A]\!]_{\mathfrak{M}}$;
2. for $n \geq 1$, $A_1, \ldots, A_n \vdash B$ is a derivable rule in L if and only if $[\![A_1]\!]_{\mathfrak{M}} \cap \ldots \cap [\![A_n]\!]_{\mathfrak{M}} \subseteq [\![B]\!]_{\mathfrak{M}}$ for every $\mathfrak{M} \in M_L$.

The first condition corresponds to weak semantic completeness and the second condition corresponds to strong completeness. In the present context, strong completeness does not entail weak completeness, so both need to be stated.

Here is a familiar example. The modal logic S4 (in which $\to$ is taken to be strict implication) has as its models triples, (W, R, V), where W is a non-empty set (of possible worlds), R is a reflexive and transitive binary relation, and V is an assignment from formulas into subsets of W that accords with the standard truth conditions for Kripke semantics. The truth condition for strict implication is

$$a \in V(A \to B) \text{ iff } \forall b(Rab \Longrightarrow (b \notin V(A) \vee b \in V(B))).$$

The model that I extract from this is a quadruple, $\mathfrak{M} = (W, N, [\![\cdot]\!], Prop)$ such that for all formulas A, $[\![A]\!] = V(A)$ and $N = W$.

The set N is the set of worlds in its model in which the theorems of the logic are verified. For some logics, such as classical logic, classical modal logic, intuitionistic logic, first-degree entailments (FDE), and strong Kleene (K_3), this set of so-called normal worlds is just the whole set of worlds, W. For substructural logics such as relevant logics and linear logic, N may not be the same as W.

What happens when I integrate a modal logic into the present framework is that the modal accessibility relation is no longer part of the surface grammar of the model theory. What is assumed is that there are mechanisms in the individual models that determine the truth values of the formulas at each of the points in the model. In the current framework, the model in the sense it is being used here is abstracted from the model in the standard sense.

2.1 Some logics and their models

The model theory for different logics require different treatments to produce models in the required sense. In what follows, I look at a few logics to illustrate in more depth the method used. I start with lattice logic, which is the minimal logic that fits with probability theory.

Lattice logic

The weakest logic that has a semantics of this sort is *lattice logic* (Paoli & Restall, 2005). Lattice logic is just the logic that has as its algebraic semantics the class of lattices. Alasdair Urquhart's representation theorem for lattices (Urquhart, 1978) can be employed to show that lattice logic is complete over a set of models, which has an interesting topological frame theory. But I am interested here in a more general and more abstract characterisation of a set of models that characterises lattice logic.

A model for lattice logic is a $(W, Prop, [\![\cdot]\!])$ such that W is a non-empty set, $Prop \subseteq W^2$, and $[\![\cdot]\!]$ is a function from formulas into $Prop$ such that

- $[\![A \wedge B]\!] = [\![A]\!] \cap [\![B]\!]$;
- $[\![A \vee B]\!] = [\![A]\!] \sqcup [\![B]\!]$.

where $X \sqcup Y$ is the least upper bound of X and Y in $Prop$. In models for lattice logic, the least upper bound always exists. The treatments of implication and negation are arbitrary in models for lattice logic, since neither belongs to the language of lattices.

It is easy to show that, given any lattice, such a model can be constructed. The worlds in this model are the filters of the lattice. A member of $Prop$ is a set of filters X such that there is some point a in the lattice for which $a \in X$. Let $[a]$ be the set of filters X such that $a \in X$. It can be shown that $[a \vee b] = [a] \sqcup [b]$ and $[a \wedge b] = [a] \cap [b]$.

To construct a model in the sense used in this paper, we add a set N to the definition of a model for lattice logic such that $N = W$.

An easy way to formulate lattice logic is by means of sequent calculus in which the sequents are of the form $\Gamma \vdash \Delta$ where Γ and Δ are sets. The axioms of the system are all sequents of the form $A \vdash A$ and its one structural rule is K:

$$\frac{\Gamma \vdash \Delta}{\Gamma' \vdash \Delta'}$$

where $\Gamma \subseteq \Gamma'$ and $\Delta \subseteq \Delta'$. The other classical structural rules are obtained from the fact that antecedents and succedents are sets. Lattice logic has the following connective rules:

$$\frac{\Gamma \vdash A, \Delta \quad \Gamma \vdash B, \Delta}{\Gamma \vdash A \wedge B, \Delta} \qquad \frac{\Gamma \vdash A \wedge B, \Delta}{\Gamma \vdash A(B), \Delta}$$

$$\frac{\Gamma \vdash A, \Delta}{\Gamma \vdash A \vee B(B \vee A), \Delta} \qquad \frac{A \vdash \Delta \quad B \vdash \Delta}{A \vee B \vdash \Delta}$$

Apart from the requirement that the premises be singular on the left in the disjunction on the left rule, these are standard classical meaning rules for the connectives.

One variant of lattice logic in which I am interested is the logic of bounded lattices. To create this logic we add the following axioms:

$$\Gamma \vdash \top \quad \bot \vdash \Delta \quad \Gamma \vdash t \quad f \vdash \Delta$$

where $\bot = \neg\top$ and $f = \neg t$. I also add to the definition of a model for lattice logic the following conditions:

$$[\![\top]\!] = W = N = [\![t]\!] \qquad [\![\bot]\!] = \emptyset = [\![f]\!]$$

The choice of lattice logic has to do with the theory of probability given below. It is the weakest logic that supports this theory of probability. It can be proven that any model for lattice logic can have a probability function that satisfies the generalised probability axioms imposed on it. And, as I argue, there is good reason for wanting probability functions in this sense.

Fuzzy logic

There are various fuzzy logics, but I only discuss here a simple fuzzy logic defined by a set of fuzzy models. A fuzzy model is a triple $([0,1], d, V)$ such that $[0,1]$ is the real interval between 0 and 1, d is some distinguished member of $[0,1]$, and V is a function from the formulas into $[0,1]$ such that the following clauses hold:

- $V(A \wedge B) = \min\{V(A), V(B)\}$;
- $V(A \vee B) = \max\{V(A), V(B)\}$;
- $V(\neg A) = 1 - V(A)$;
- $V(A \rightarrow B) = \max\{1 - V(A), V(B)\}$.

A formula A is said to be true on V if and only if $V(A) \geq d$.

A worlds model for fuzzy logic is a quadruple $(W, N, d, \mathbb{D}, \mathbb{V})$ such that W is a non-empty set of points, N is a non-empty subset of W, d is a distinguished point in $[0,1]$, $\mathbb{D}$ is a function that assigns to each world in W a real number in $[0,1]$ such that for all worlds $a \in N$, $\mathbb{D}(a) = d$, and $\mathbb{V}$ is a function that assigns to each world in W a fuzzy valuation.

To turn a worlds model for fuzzy logic into a model in the sense required by this paper, we define a new valuation function $[\![\cdot]\!]$ such that for all formulas A,

$$[\![A]\!] = \{a \in W : \mathbb{V}(a)(A) \geq \mathbb{D}(a)\}.$$

This says that the proposition assigned to A is the set of worlds in which A has a value that is at least as great as the designated value of the world. The reason that I include non-normal worlds in models is so that models will represent fuzzy consequence as it is intended. That is, for each fuzzy model, $\mathfrak{M}$, as usual, I set

$$A_1, \ldots, A_n \models_{\mathfrak{M}} B \text{ iff } [\![A_1]\!]_{\mathfrak{M}} \cap \ldots \cap [\![A_n]\!]_{\mathfrak{M}} \subseteq [\![B]\!]_{\mathfrak{M}}$$

where $A \models_{\mathfrak{M}} B$ if and only if, for all $a \in W$, if $a \in [\![A]\!]$ then $a \in [\![B]\!]$. Given the variation of designated values on non-normal worlds, this definition characterises the same consequence relation as one that is defined over a set of fuzzy valuations and sets $A \models B$ if and only if $V(A) \leq V(B)$ on every valuation in that set.

Supervaluations

The version of supervaluational frames that I present is a slightly simplified form of that given by Kit Fine in (Fine, 1975). I begin with a set of points W, a distinguished subset C of W, and a partial order $\leq$ on W. The set C is the set of *complete* points in the frame. We add the condition that if $c \in C$, then for any point a if $c \leq a,\ a = c$.

A valuation V on a supervaluational frame is a function from propositional variables to sets of worlds closed upwards under $\leq$. Thus, for example, if $a \leq b$ and $a \in V(p)$, then $b \in V(p)$. Moreover, we constrain valuations such that if for any propositional variable p and any point a, if all complete points $c \geq a$ are such that $c \in V(p)$, then $a \in V(p)$.

Each valuation determines a satisfaction relation, $\models$, according to the following recursive definition:

- $a \models p$ iff $a \in V(p)$;
- $a \models A \vee B$ iff $\forall c \in C(a \leq c \ \Rightarrow\ (c \models A \ \vee\ c \models B))$;
- $a \models A \wedge B$ iff $a \models A$ and $a \models B$;
- $a \models \neg A$ iff $\forall c \in C(a \leq c \ \Rightarrow\ c \not\models A)$;
- $a \models A \rightarrow B$ iff $\forall c \in C(a \leq c \ \Rightarrow\ (c \not\models A \ \vee\ c \models B))$.

We can prove easily that $a \models A$ if and only if, $\forall c \in C(a \leq c \Rightarrow c \models A)$, for all a and all formulas A. Truth and falsity at each of the complete points behaves in a completely classical manner. It is is also easy to show that $[\![A \vee B]\!]$ is the least upper bound of $[\![A]\!]$ and $[\![B]\!]$ in the set of propositions in a super valuational model, and hence these models can be turned into models in the sense required by the present project.

Putting together models for supervaluations, fuzzy logic, and epistemic modal logics, we can create a background model that can be used to represent reasoning about a choice between approaches to vagueness. The way in which models for various logics are to be combined to create background models is treated in the next section.

3 Background models

Given a set of models for the various different logics that one wishes to choose between, we can construct a single model to act as the basis for a belief revision structure. Crudely put, the single model is a fusion (largely a union) of the models.

Let's make this more precise. I assume that a set of models M has been selected. The background model based on M is a quadruple $(W, N, Prop, [\![\cdot]\!])$ that is defined as follows. The set of worlds is just the union of the sets of worlds of the models in M, that is, $W = \bigcup_{\mathfrak{M} \in M} W_{\mathfrak{M}}$ and similarly the set of normal worlds is just the collection of all the normal worlds in M, i.e. $N = \bigcup_{\mathfrak{M} \in M} N_{\mathfrak{M}}$. The construction of $Prop$ needs a bit more care. For each formula, A, $[\![A]\!] = \bigcup_{\mathfrak{M} \in M} [\![A]\!]_{\mathfrak{M}}$. $Prop$ is just the set of $[\![A]\!]$ for all formulas A.

For each logic L that is being considered, I also add two constants $\top_L$ and t_L. Let M_L be the set of all the models in M for the logic L. The interpretations of the new constants are as follows:

$$[\![\top_L]\!] = \bigcup_{\mathfrak{M} \in M_L} W_{\mathfrak{M}} \qquad [\![t_L]\!] = \bigcup_{\mathfrak{M} \in M_L} N_{\mathfrak{M}}$$

These two constants give us a means in the language to talk about the logic L. T_L represents all of the formulas and rules that are true or valid in every world in every model for L and t_L represents the theorems of the language and the rules under which the theorems are closed. (The sets of worlds that these constants represent are members of $Prop$.)

In the next section, I continue to examine the different ways in which we can understand formulas and rules of inference.

3.1 Representing theorems and rules of inference

In accepting or rejecting a formula, we may accept or reject it as a simple truth or we may accept or reject that it is a logical truth. To accept or reject a formula as a logical truth (as an axiom or theorem), what one accepts or rejects is not the proposition expressed by the formula but rather the normal sets that verify that proposition as a logical truth. In other words, a formula, taken as an axiom, is the union of the sets of normal worlds that make true that formula. In formal notation:

$$[\![A]\!]^{Th} = \bigcup\{N_{\mathfrak{M}} : \mathfrak{M} \in M \ \wedge \ N_{\mathfrak{M}} \subseteq [\![A]\!]\}$$

I add to *Prop*, for each formula A in the language, the union of sets of sets of normal worlds that verify A.

The treatment of rules is more complicated. There are various types of rules. Two sorts of rules that are, in my opinion, most important are the ones that Dana Scott calls "horizontal" and "vertical" rules of inference (Scott, 1971). A horizontal rule of inference, can be represented as a single conclusion sequent of the form

$$A_1, \ldots, A_n \vdash B$$

This is to be understood in the semantics as saying that if any world makes true all of the premises, $A_1, \ldots, A_n$, then it also makes true the conclusion B. A vertical rule is a rule that sets a closure condition for the set of theorems of a logic. For example, weaker relevant logics such as B and MC obey the following version of the rule of modus ponens:

$$\frac{\begin{array}{c}\vdash A \to B \\ \vdash A\end{array}}{\vdash B}$$

These logics, however, do not obey the horizontal version of this rule, i.e. $A \to B, A \vdash B$.

It also might be desirable to distinguish between types of vertical rule in order to provide the correct semantic interpretations of them. Consider the rule of universal generalisation for first order classical logic:

$$\frac{\vdash A(a)}{\vdash \forall x A(x)}$$

where a is free for x. This rule is quite different from the modus ponens rule presented above. There are models for classical first order logic in which the normal worlds, either individually or as a set, do not satisfy this rule. Consider a one world model of the natural numbers. The domain of quantification is the natural numbers and in it the formula $\neg 0 = 1$ is true. But the formula $\forall x \neg x = 1$ fails to be true. What is true, however, is that where a is an individual constant, if $A(a)$ is true in every world in every model for classical first order logic then $\forall x A(x)$ is also true in every world. We can represent a rule like this as the union of sets of normal worlds of all the logics that are closed under this rule.

In a more involved version of the theory, we might add a ternary relation R to models in order to represent *intensional sequents*. An intensional sequent is a sequent of the form $\Gamma \vdash A$ where Γ is a *structure* of formulas and A is a formula. The set of structures of formulas is the smallest set such that: if A is a formula, then A is a structure; if Γ and Δ are structures then so is $(\Gamma;\Delta)$ (Restall, 2000, ch 2). The semi-colon is used to express an intensional connection between premises, often formalised in the object language of substructural logics by fusion ($\circ$). An intensional sequent $A;B \vdash C$ is valid in a model $\mathfrak{M}$ if and only if for all worlds a,b,c in $\mathfrak{M}$, if $Rabc$, $a \models A$, and $b \models B$, then $c \models C$. This idea is extended to treat all intensional sequents using products of the ternary relation with itself. The notation $R^2(ab)cd$ means that $\exists x(Rabx \wedge Rxcd)$. And we can have products of R any finite power. Thus, for example, we have $(A;B);C \vdash D$ valid in $\mathfrak{M}$ if and only if for all a,b,c,d in $\mathfrak{M}$, if $R^2(ab)cd$, $a \models A$, $b \models B$, and $c \models C$, then $d \models D$. In models for substructural logics the ternary relation is also used to give a truth condition for implication: $a \models A \rightarrow B$ iff $\forall x \forall y((Raxy \wedge x \models A) \Rightarrow y \models B)$. This connection between the representation of intensional sequents and implication yields the following form of the deduction theorem:

$$\Gamma;A \models B \text{ iff } \Gamma \models A \rightarrow B$$

In this paper, however, I do not make any use of this ternary relation except in this digression.

4 Belief revision

Once a background model is constructed it can be used as the basis for a model of a theory of belief revision. For instance, we can impose on it a system of spheres in the sense of Grove's semantics for the AGM theory of belief revision (Grove, 1988; Gärdenfors, 1988). We could also place a binary preference relation on worlds and construct a model for some form of dynamic epistemic logic or dynamic doxastic logic (van Benthem, 2011; Segerberg, 1995). Instead of these, I use the background model as a basis for a model of probability, and develop a probabilist theory of belief revision.

4.1 Probability functions

In (Mares, 2014), the notion of a probability function is generalised to fit the range of logics treated here. In this section, I briefly state the axioms of the resulting class of functions and discuss some of their properties that are salient for the current project.[1]

A probability function P on a model $\mathfrak{M} = (W,N,V,Prop)$ is a function from the formulas into $[0,1]$ such that

[1] In (Mares, 2014), I do not have the axioms $P(W) = 1$ and $P(\emptyset) = 0$, but rather slightly more complicated axioms. The simpler axioms will do for my purposes here.

1. $P(W) = 1$ and $P(\emptyset) = 0$;
2. If $X \subseteq Y \cup Z$, then $P(X) \leq (P(Y) + P(Z)) - P(Y \cap Z)$;
3. If $Y \cup Z \subseteq X$, then $(P(Y) + P(Z)) - P(Y \cap Z) \leq P(X)$;
4. If $X \subseteq Y$, then $P(X) \leq P(Y)$.

The complex forms of additivity expressed by postulates 2 and 3 are needed because the standard form of additivity, viz. $P(A \cup B) = (P(A) + P(B)) - P(A \cap B)$, assumes that the probability of unions of members of *Prop* are always defined. This isn't the case in models for logics without distribution, such as linear logic, or in my background model.

I adopt the traditional definition of conditional probability:

$$P(Y/X) \;=\; \frac{P(X \cap Y)}{P(X)}$$

It can be easily shown that if in a model the propositions are closed under union and boolean complement, any probability function over this model also satisfies the standard finite additivity axiom:

$$P(X \cup Y) = (P(X) + P(Y)) - P(X \cap Y)$$

For all logics that have classical models (such as classical and modal logics), the associated probability functions are all traditional Komolgorov functions. This is just an example of how the probability functions adapt to be appropriate to the logics, given both the base definition of a probability function (given above) and the models for the logics. The logics incorporated into the present framework satisfy the following "equation":

Base Probability Theory + Logic L = Appropriate Probability Theory for L

The class of probability functions in which I am most interested in are those that are defined over the background model. It is clear, however, that for each probability function P defined over the background model and for each sub-model $\mathfrak{M}$ the function $P \restriction \mathfrak{M}$ is a probability function meeting the postulates given above.

Using this notion of probability, I define a version of Jeffrey conditionalisation as the method of updating agents' probability functions in response to new information. Given a new piece of information $P_{New}(X) = r$, for any proposition Y,

$$P_{New}(Y) \;=\; [\frac{P(X \cap Y)}{P(X)} \times P_{New}(X)] \;+\; [\frac{P(Y) - P(X \cap Y)}{1 - P(X)} \times (1 - P_{New}(X))].$$

This is just Jeffrey's definition adapted to the present context. I replace the negations in the original definition because not all of the logics considered here have negations that express boolean complement on sets in their model theory. The problem with this form of updating is that it does not treat cases in which an agent has given a proposition a prior probability of zero and then later discovered that it should be given a positive value. To deal with this difficulty, I suggest that a move be made to a more complicated structure in which a class of probability functions on a model are placed in a linear order, using something like Wolfgang Spohn's ranking theory (Spohn, 2012).

4.2 Beliefs and rejections

In thinking about logic, we both decide which principles to accept and which to reject. Formulating rejection can be problematic. Some logics do not have a negation that can be understood as representing rejection. Paraconsistent negations, for example, allow one to accept both a proposition and its negation without commitment to every proposition. These logics have non-trivial models and are used to represent, among other things, inconsistent belief states. Thus I think of a cognitive state of an agents as including both a set of beliefs and a set of propositions that he or she rejects. Rejection is a primitive of the present theory (as in (Mares, 2002)). The implementation of the probability theory here is to explain both how an individual organises his or her beliefs and his or her rejections.

One's beliefs and rejections do not just require internal organisation, they need to bear certain relations to one another. Let's call the pair (Γ, Δ) of an agent's beliefs and rejections, respectively, her *content*. This content is said to be *pragmatically consistent* if and only if Γ does not entail Δ, that is,

$$\bigcap \Gamma \not\subseteq \bigsqcup \Delta.$$

Where the generalised join operation used here is defined as:

$$\bigsqcup \Delta = \bigcap \{X \in Prop : \forall Y (Y \in \Delta \Rightarrow Y \subseteq X)\}$$

Note that whereas the infinite meets (intersections, in this case) and joins are definable, an infinite join or meet of propositions need not belong to *Prop* itself.

This notion of pragmatic consistency can be modified to produce a probabilistic notion in the same way as the notion of classical consistency is altered by classical Bayesians to a notion of probabilistic consistency. On the classical notion, an agent's beliefs are probabilistically consistent if and only if the strength of her beliefs accord with the probability calculus. Similarly, an agent's content can be said to be *probabilistically pragmatically consistent in a context* if and only if the strength of beliefs and rejections accord with the probability calculus and, in the context there is a threshold of belief b and a threshold of rejection r that do not overlap. This requires some explanation.

On the present theory, an ideal agent's cognitive state in a context can be represented by a triple (P, b, r) where P is a probability function over a background model and $b, r \in [0, 1]$ such that $b - r > 0$. Her content in that context is (Γ, Δ) where Γ is the set of propositions X such that $P(X) \geq b$ and Δ is the set of propositions Y where

$P(Y) \leq r$. The requirement that $b - r$ be strictly greater than 0 entails that $\Gamma \cap \Delta = \emptyset$. This gives us a very weak notion of consistency.

The notion of pragmatic consistency requires that the intersection of the believed propositions is always non-empty. This, I argue below, is a reasonable constraint for debates and negotiations concerning logic, but not a reasonable principle to govern an individual's beliefs. The lottery paradox shows that there are cases in which one should believe a collection of propositions but not believe in their conjunction. A virtue of probabilism is that it allows one to hold pragmatically inconsistent beliefs, even about logic, and to have a very nuanced attitude towards logical principles.

5 Positions and scorekeeping

In this section, I use a version of score-keeping pragmatics to act as a foundation for a social-epistemological version of the theory belief revision that I have outlined above.

According to David Lewis's score-keeping theory, in any well-run conversation there is a running score. At each point in the conversation, the score represents various parameters that are essential to understanding the conversation, such as, the degree of vagueness that is acceptable, and what presuppositions are being made (D. Lewis, 1979). The element of a score that is of most interest in the present context is what David Ripley and Greg Restall call the common *position* of the participants of the conversation. The position at a particular point in a conversation is the social counterpart of what I call in the previous section an agent's content, that is, a position is a pair (Γ, Δ) such that Γ is the set of propositions that are accepted by all the participants in the conversation and Δ is the set of propositions that are jointly denied by all of them.

Ideally, positions should be pragmatically consistent. When our aim is to decide (by debate, negotiation, or other means) which logic to apply, we need to accept a class of models. In debates and negotiation, the aim to is to accept a set of rules as the ones to apply. In one's own content, one might hedge one's acceptance of a particular logic with less certain beliefs in other principles, perhaps in some cases with weak rejections of related principles, and so on. But if the aim is to find one logic to accept, these nuanced attitudes are to be avoided.

It is a necessary condition of the acceptance of a proposition by a participant in a conversation that her degree of belief in that proposition be above the threshold that she has set for belief. It cannot be a sufficient condition because of conflicts with others in the conversation and because of the rejection of the sorts of nuances that I have discussed above.

In thinking about debates concerning logic, I adopt some of the connections between probabilist reasoning and argumentation theory developed by researchers working in Bayesian argumentation theory (Betz, 2012, 2013; Zenker, 2013), and combine these with scorekeeping pragmatics. Probability theory can be used to analyse the strength of evidence and arguments used in debates and to give norms for when participants in the debate and members of the audience should accept or reject claims made in the debate. At a particular time in a debate, the logic that is acceptable is given by the region of the background model constructed from largest class of models that all contain all the propositions that are accepted at that time by the participants of the debate and fail to contain any of the propositions that are rejected at that time.

5.1 Circular justification

The key to understanding debate on this theory is that evidence is understood as supporting a contrast between logical systems. This is particularly useful because evidence for logical systems is often seen as circular. Paul Boghossian has us consider a logical system with modus ponens as its only primitive rule. Then any rule of inference that we use in its justification must either be modus ponens or derived from modus ponens (Boghossian, 2001, p. 10).

Boghossian thinks of the epistemic justification of logical rules in terms of deductive arguments. Probabilists, on the other hand, generally treat epistemic justification inductively or contrastively. Here is a Bayesian analysis of allegedly circular justification by Tomoji Shogenji:

> Here is a general procedure for avoiding epistemic circularity. Suppose the suspicion of epistemic circularity arises because the evidence that we hope to use for confirming the hypothesis is useful only if the truth of the hypothesis is assumed. To avoid epistemic circularity, assume the hypothesis only in the sense of envisioning its truth. ... Assume next the negation of the hypothesis – again in the sense of envisioning its truth. Then proceed to compare the degree of coherence between the evidence and the hypothesis and the degree of coherence between the evidence and the negation of the hypothesis, using the background assumptions when necessary. (Shogenji, 2013, p. 180)

Shogenji's analysis is available for justification of basic logical principles and rules, such as modus ponens, in the present framework. The only principles and rules that are excluded are those of lattice logic. The idea is that, given the position that the participants occupy in a debate, they take a principle and contrast the probabilities we give to the union of the class of models that accept the principle or rule with the union of the class of models that do not accept it.

Tonk

Let us consider Paul Boghossian's example of a circular justification of a logic that contains Prior's tonk rules. Boghossian shows that if we allow circular justification, then anyone who accepts the tonk rules can justify these very same rules (Boghossian, 2001). The probabilist solution to this problem is to look at the models for the tonk logic. The tonk logic only has models that are trivial in the sense that every world is in every proposition. Trivial models are excluded by the definition of 'model' that I am using. Every model, on this definition, contains at least one world. In addition the empty set is a proposition in every model. Hence in no model is there is a world is in every proposition. There are good reasons from the theory of probability for these structural constraints on models. If we have an empty model, then the probability of the set of worlds will be the same as the probability of the empty set, but the former is one and the latter zero. Thus, the background model cannot be a model for the tonk logic, nor can any submodel of the background model be a model for the tonk logic.

6 Negotiating logic

In actuality, the problem of alternative logics has been confined to conversations amongst philosophers, computer scientists, and mathematicians. But what if it broke out of these disciplines and became more widely discussed? In particular, what if in wider society there were a debate about what sorts of rules of inference we could use in political discussions and other social activities? C.I. Lewis posed this problem in 1921:

> Now whoever enters a discussion, pragmatically assumes that the logical sense of those engaged is the same with his. The pursuit of common enterprises, regarded as rational, rests at the bottom upon a similar assumption. But in making this assumption – as we are frequently aware – we take a certain risk. In the interest of our rational enterprise we must take this risk. ... The ideal of a universal logical sense is one strongly demanded by its importance to all social enterprises, and is more closely approximated in fact than most of our ideals. But sticking to the facts, in the spirit of the facts, we are obliged to admit that it does not completely exist and probably never will. It is easy to define "rationality" in one's own terms. But that can only lead to the familiar conclusion, "All the world is strange save thee and me – and thee's a little strange." With respect to our ideals, we all of us stand in the egocentric predicament; we can only assert our own and hope for agreement. (C. Lewis, 1921, p. 379)

Lewis thinks that there is *in fact* widespread agreement about which rules of reasoning should be used. I'm not sure that he is right about this, but even if there is agreement it is an important question to ask whether we could find rational support for the sort of agreement we have.

Different social enterprises may require different methods of reaching agreement. Consider, for example, engineers who wish to build a large and potentially dangerous structure, who are in disagreement about the logic should underly the mathematics that they use. Some wish to use calculus based on classical logic, others will accept only intuitionist mathematics. In such cases what is appropriate is to determine what the relative risks involved are in making one choice over the other. If the people involved are rational in the sense of the present paper, then they will be able to construct subjective assessments of the risks involved and then a discussion can begin between them in a rational manner.

With regard to political discussions, the values that need to be satisfied include political ideals such as fairness. In order to understand how to implement the theory in dealing with disagreement about the logical structure of political discourse, we would have to become clear about what is at stake for individuals with regard to different choices of logic. If we can determine what the gains and costs are of allowing certain rules of inference, we can then deal with the participants as Bayesian agents negotiating to gain advantage. Although there has been some work on the oppression of women and logic (Plumwood, 1993), there has been overall very little work on this, but it is an interesting avenue of research.

7 Problems concerning logical omnicience

Surely one sort of logical evidence concerns what can be proven given a particular logic. People do often compare logics in terms of their theorems and in terms of how much mathematics we can do using them. This sort of evidence is easily dealt with on the present theory. We compare two (or more) classes of models with one another, and set our preference for one or the other in terms of the probabilities we assign to the models.

A problem arises, however, when one discovers, against her prior beliefs, that a contested formula is or is not a theorem of a given logic. For example, suppose that we wish to model an agent who is trying to decide between two logics, L_1 and L_2. She would like it if A were a theorem of the logic chosen, but does not know whether it is. In fact, A is a theorem of L_1 but not of L_2. So, in a background model representing her space of alternative logics, $[\![A]\!]^{th} = [\![t_{L_1}]\!]$. So, $P([\![A]\!]^{th}) = P([\![t_{L_1}]\!])$. But,

$$P([\![t_{L_1}]\!]/[A]\!]^{th}) = \frac{P([\![t_L]\!] \cap [\![A]\!]^{th})}{P([\![A]\!]^{th})}.$$

Therefore,

$$P([\![t_{L_1}]\!]/[A]\!]^{th}) = \frac{P([\![t_{L_1}]\!])}{P([\![A]\!]^{th})} = 1.$$

So, representing the fact that the agent doesn't know that A is a theorem of L_1 and not of L_2 seems impossible on the present theory as does treating the discovery that A is a theorem of L_1 but not of L_2 in a non-trivial manner. This problem is a variant of the familiar problem of logical omniscience that is faced by most epistemic and doxastic logics.

In order to deal with this difficulty, I complicate the structure somewhat. I add a finite set of agent states AS each of which can be thought to be considering a background model with two logics, but logics that are perhaps slightly different from L_1 and L_1. In some of these background models, the logic labelled 'L_1' has A as a theorem and in some it does not and in some the logic called 'L_2' has A as a theorem and in some it does not. This range of states represents the ambiguity produced by an agent's thinking about a logic but attributing to it the wrong theorems or failing to attribute to it the right theorems.

I also add to each agent state a set of formulas that in that state the agent attributes as theorems to L_1 and a set of formulas that the agent in that state thinks are theorems of L_2 (the same can done with rules, of course). In addition, I assume a function $P_{AS} : AS \longrightarrow [0,1]$ such that $\sum_{a_1 \in AS} a_1 = 1$. The P_{AS} function represents the weighting of the various states in terms of their approximation to the actual cognitive state of the agent who is being modelled. I write $[\![A]\!]_i$ to refer to the proposition expressed by A in the background model assigned to a_i. Now we can represent the probability P^* of a *formula* as a statistical average of probabilities over the states in AS:

$$P^*(A) \;=\; \sum_{a_i \in AS(a)} (P_{AS}(a_i) \times P_{a_i}([\![A]\!]_i))$$

We can represent the probability of A's being a theorem – $P^*(\vdash A)$ – as an average of $P_{a_i}([\![A]\!]_i^{th})$s, and similarly with regard to rules.

I am, of course, not presenting this solution as a general solution to the problem of logical omniscience in the different ways it appears in a probabilist context. Rather, I think as do others working in probabilist theories, that various modelling tasks require different tweaks to the basic model. This is a central feature of the way in which scientific modelling operates in practice.

8 Half-baked remarks: The apriority of logic

Anti-exceptionalism about logic is the doctrine that logics are to be thought of as being on a par with scientific theories. On this view, logic is up for revision and logical knowledge (or logical beliefs) are not a priori (Hjortland, forthcoming). What I say in this paper is compatible with some form of anti-exceptionalism – it presents a formal theory that can be taken to be an anti-exceptionalist methodology. The view of this paper, however, is not committed to an anti-apriorism in the epistemology of logic, nor of course is it committed to apriorism.

C.I. Lewis held that all choice between logical systems is pragmatic (C. Lewis, 1932). That is, we must appeal to theoretical virtues such as simplicity, intuitiveness, and strength to justify the acceptance of one logical system over another. This sort of justification is a priori in the sense that it does not appeal to empirical data about the things that logic is about (i.e. everything) but it is also defeasible. We may come across new evidence that another logical system, perhaps that we have not considered, better exemplifies these virtues. In general, a priori justification need not be conclusive.

What is perhaps more interesting is the problem of defining the notion of apriority in the current theory. In (Mares, 2011) I claim that one cannot define the notion of a priori outside a more general epistemological framework. A notion of absolute apriority can be defined easily in any probabilisitic theory. The proposition that is just the total set of worlds is absolutely a priori in an agent's belief state and a sentence that expresses the complete set of worlds in every belief state is absolutely a priori is a very strong sense.

There is, however, room in the theory for weaker senses of 'a priori'. For example, I can adopt a notion of *relative apriority*. A formula A is a priori relative to a logic L if and only if $N_L \subseteq [\![A]\!]$. On this definition, all theorems of L are a priori relative to L. If we strenghten this notion to say that $[\![A]\!]$ needs to be a superset of the theorems of L on all models, then only theorems of L would be a priori relative to L, but the weaker version is more interesting. It is an agent-relative notion of relative apriority.

Clearly, other definitions of 'a priori' are possible in this framework, but I will leave this topic until after I have the opportunity to give it further thought.

9 Concluding remarks

In this paper I have set out a formal epistemology for logic choice. The theory in this paper locates the choice between alternative logics within a widely accepted and well-understood epistemological framework – probabilism. This seems right. Reasoning about logic and debates concerning logic seem unremarkable in the sense that they are commonplace and seem to be extensions of our standard belief-forming practices. Whether probabilism or some other form of belief revision is correct is not the real issue here. Rather, the fact that very little logic needs to be presupposed in order to construct a theory of belief revision that can treat logic choice. This bodes well for the creation of theories that will represent as rational reasoning about other apparently fundamental beliefs and values.

Acknowledgements. I am grateful to Greg Restall and the other members of the programme committee for inviting me to address the CLMPS and for organising this session on philosophical logic. I also thank the members of the audience for their helpful comments, especially Raymundo Morado, Francesco Paoli, Arnon Avron, and Giovanna Corsi.

Bibliography

Betz, G. (2012). On degrees of justification. *Erkenntnis*, *77*, 237–272.

Betz, G. (2013). Degrees of justification, Bayes' rule, and rationality. In F. Zenker (Ed.), (pp. 135–146). Dordrecht: Springer.

Boghossian, P. (2001). How are objective epistemic reasons possible? *Philosophical Studies*, *106*, 1–40.

Fine, K. (1975). Vagueness, truth and logic. *Synthese*, *30*, 265–300.

Gärdenfors, P. (1988). *Knowledge in flux: Modelling the dynamics of epistemic states*. Cambridge, MA: MIT Press.

Goheen, J., & Mothershead, J. (Eds.). (1970). *Collected papers of Clarence Irving Lewis*. Stanford, CA: Stanford University Press.

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

Hjortland, O. (forthcoming). Anti-exceptionalism about logic. *Philosophical Studies*.

Lewis, C. (1921). The structure of logic and its relation to other systems. *The Journal of Philosophy*, *18*, 505–516. Read at the conference of the American Philosophical Association, December 1920. Reprinted in Goheen, & Mothershead, 1970, pp 371-382. Page references are to the reprinted version.

Lewis, C. (1932). Alternative systems of logic. *The Monist*, *17*, 481–507. Reprinted in Goheen, & Mothershead, 1970, pp 400-419. Page references are to the reprinted version.

Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, *8*, 339–359.

Mares, E. (2002). A paraconsistent theory of belief revision. *Erkenntnis*, *56*, 229–246.

Mares, E. (2011). *A priori*. Durham: Acumen-Routledge.

Mares, E. (2014). Belief revision, probabilism, and logic choice. *Review of Symbolic Logic*, *7*, 647–670.

Paoli, F., & Restall, G. (2005). The geometry of non-distributive logics. *The Journal of Symbolic Logic*, *70*, 1108–1126.

Plumwood, V. (1993). The politics of reason: Towards a feminist logic. *Australasian Journal of Philosophy*, *71*, 436–462.

Restall, G. (2000). *Introduction to subsructural logics*. London: Routlege.

Scott, D. (1971). On engendering an illusion of understanding. *Journal of Philosophy*, *68*, 787–807.

Segerberg, K. (1995). Belief revision from the point of view of doxastic logic. *Bulletin of the IGPL*, *3*, 535–553.

Shogenji, T. (2013). Reductio, coherence, and the myth of epistemic circularity. In F. Zenker (Ed.), (pp. 165–184). Dordrecht: Springer.

Spohn, W. (2012). *The laws of belief: Ranking theory and its philosophical applications*. Oxford: Oxford University Press.

Urquhart, A. (1978). A topological representation theorem for lattices. *Algebra Universalis*, *8*, 45–58.

van Benthem, J. (2011). *Logical dynamics of information and interaction*. Cambridge: Cambridge University Press.

Zenker, F. (Ed.). (2013). *Bayesian argumentation: The practical side of probability*. Dordrecht: Springer.

Author biography. Edwin Mares is a professor of philosophy at Victoria University of Wellington. He has written more than sixty articles and chapters, largely on philosophical logic, and he written the books, Relevant Logic (Cambridge, 2004), A Priori (Routledge, 2011), and (with Stuart Brock) Realism and Antirealism (Routledge, 2007).