A Decision Procedure for Probability Calculus with Applications[a]

Branden Fitelson

Department of Philosophy
Group in Logic and the Methodology of Science
&
Cognitive Science Core Faculty

University of California–Berkeley

branden@fitelson.org

http://fitelson.org/

[a]The paper behind this talk was recently published in the *Review of Symbolic Logic* [14].

Overview of Today's Talk

- Motivation: The Problem(s)
  - Bayesian confirmation theory (overview)
  - The Problem of Measure-Sensitivity
  - Lots of questions about the validity of Bayesian arguments
  - Leading to lots of problems in the probability calculus
  - Solutions needed!
- The Solution: `PrSAT` – A decision procedure for Probability Calculus
  - Probability calculus: axiomatic and algebraic approaches
  - Algebra, quantifier elimination, and the probability calculus
  - Implementation of `PrSAT` in *Mathematica*
  - Demonstration of `PrSAT` on some examples
  - Future Work: optimizing and scaling-up `PrSAT`

Some Bayesian Background I

- Orthodox Bayesianism (*i.e.*, Bayesian *epistemology*) assumes that the degrees of belief (or credence) of rational agents are (*Kolmogorov* [25]) *probabilities*.
- $\Pr(H \mid K)$ is the degree of belief that a rational agent with background knowledge K (S_K) assigns to H. This is S_K's *prior* probability of H.
- $\Pr(H \mid E \,\&\, K)$ is the degree of belief S_K assigns to H, *on the supposition that* E (*i.e.*, the d.o.b. that S_K *would* assign to H *upon learning* E). This is S_K's *posterior* probability of H (on E). I'll drop the "K"s now, for simplicity.
- A simple toy example (just to help fix our ideas): Let H be the hypothesis that a card (drawn at random from a standard 52-card deck) is a spade, and let E be the (evidential) proposition that the card is the ace of spades.
- Given standard assumptions about random card draws, $\Pr(H) = 1/4$ and $\Pr(H \mid E) = 1$. So, learning E raises the probability of (indeed, *verifies*) H.

Some Bayesian Background II

- In (contemporary) Bayesian confirmation theory, evidence E *confirms* (or *supports*) a hypothesis H if learning E *raises the probability of* H.
- If learning E *lowers* the probability of H, then E *dis*confirms (or *counter*-supports) H, and if learning E *does not change* the probability of H, then E is confirmationally *neutral* regarding H. This is a Pr-*relevance* theory.
- Within (*Kolmogorov*! [10], [13]) probability theory, there are many *logically equivalent* ways of saying that E confirms H. Here are a few:
  - E confirms H if $\Pr(H \mid E) > \Pr(H)$.
  - E confirms H if $\Pr(E \mid H) > \Pr(E \mid {\sim}H)$.
  - E confirms H if $\Pr(H \mid E) > \Pr(H \mid {\sim}E)$.
- By taking differences, (log-)ratios, *etc.*, of the left and right sides of these (or other equivalent) inequalities, a *plethora* of candidate *relevance measures* of *degree of confirmation* can be formed. Relevance measures are such that:

  $(\mathcal{R})\ \mathfrak{c}(H, E) \gtreqless 0$ iff $\Pr(H \mid E) \gtreqless \Pr(H)$.

Four Popular and Representative Relevance Measures

- *Dozens* of Bayesian relevance measures have been proposed in the philosophical literature (see [26] for a survey). Here are four popular ones.[a]
  - *Difference*: $d(H, E) =_{df} \Pr(H \mid E) - \Pr(H)$
  - Log-*Ratio*: $r(H, E) =_{df} \log\left[\frac{\Pr(H \mid E)}{\Pr(H)}\right]$
  - Log-*Likelihood-Ratio*: $l(H, E) =_{df} \log\left[\frac{\Pr(E \mid H)}{\Pr(E \mid {\sim}H)}\right]$
  - "Normalized Difference": $s(H, E) =_{df} \Pr(H \mid E) - \Pr(H \mid {\sim}E) = \frac{1}{\Pr({\sim}E)} \cdot d(H, E)$
- Logs are taken to ensure easy satisfaction of relevance criterion ($\mathcal{R}$). They are merely a useful convention (they're inessential, but they simplify things).
- The first part of our story concerns the *disagreement* exhibited by these measures, and its ramifications for Bayesian confirmation theory ...

[a] Users of d include [9], [8], and [22]. Users of r include [19], [27], and [20]. Users of l include [17], [32], and [12]. Users of s include [23] and [5]. See [10], [11], and [12] for further references.

Disagreement Between Alternative Relevance Measures

- What kind of disagreement between relevance measures is important?
- Mere *numerical* differences between measures are not important, since they need not affect *ordinal* judgments of what is more/less well confirmed than what (by what). Think about C *vs* F temperatures (numbers *vs* comparisons).
- *Ordinal* differences are crucial. Such *comparative* differences affect the cogency of many arguments surrounding Bayesian confirmation theory.
- For instance, it is part of Bayesian Lore that the observation of a black raven (E_1) confirms the hypothesis (H) that all ravens are black *more strongly than* the observation of a red herring (E_2) does (given "actual corpus" K).
- Given the standard background assumptions (K) in Bayesian accounts of Hempel's ravens paradox, this conclusion [$\mathfrak{c}(H, E_1) > \mathfrak{c}(H, E_2)$] follows for *some* measures of confirmation $\mathfrak{c}$, but it *fails to follow* for some choices of $\mathfrak{c}$.
- Such arguments are said to be *sensitive to choice of measure* [11].

Tabular Summary Some Measure-Sensitive Arguments

	Valid *wrt* relevance measure:			
Argument	d?	r?	l?	s?
Horwich [19] *et al.* on Hempel's Ravens Paradox	Yes	Yes	Yes	No
Eells [9] on Goodman's "Grue" Paradox	Yes	No	No	Yes
Sober [33] on Goodman's "Grue" Paradox	Yes	No	Yes	Yes
Rosenkrantz [30] on Irrelevant Conjunction	Yes	No	No	Yes
Earman [8] on Irrelevant Conjunction	Yes	No	Yes	Yes
Horwich [19] *et al.* on the Variety of Evidence	Yes	Yes	Yes	No
Christensen [5] on the Old Evidence Problem	No	No	Yes	Yes
Popper-Miller's [29], [16] *Critique* of Bayesianism	Yes	No	No	Yes

Axiomatic Probability Calculus

- All well-known problems in Bayesian confirmation theory can be represented in small, finite Kolmogorov probability models $\mathcal{M} = \langle \mathcal{B}, \Pr \rangle$.
- Such $\mathcal{M}$'s are just small, finite Boolean algebras $\mathcal{B}$ (of propositions), with a function Pr: $\mathcal{B} \mapsto [0, 1]$ satisfying the following three axioms [25]:
  1. For all $X \in \mathcal{B}$, $\Pr(X) \geqslant 0$.
  2. $\Pr(\top) = 1$, where $\top$ is any tautological proposition in $\mathcal{B}$.
  3. For all $X, Y \in \mathcal{B}$, if X and Y are mutually exclusive, then:
  $$Pr(X \vee Y) = \Pr(X) + \Pr(Y)$$
- Conditional probabilities $\Pr(\cdot \mid \cdot)$ are then *defined* in terms of $\Pr(\cdot)$ as:

  Definition. $\Pr(X \mid Y) = \frac{\Pr(X \,\&\, Y)}{\Pr(Y)}$
- Thus, Pr is a finitely additive measure over some $\mathcal{B}$ of propositions. This allows Pr's to be interpreted (roughly) as "areas" in Venn Diagrams ...

Algebraic Probability Calculus I

- Almost all problems in Bayesian confirmation theory are expressible in Kolmogorov probability models with just *three* atomic propositions.
- Such $\mathcal{M}$'s can be interpreted using Venn Diagrams ([28], [1]) in a simple way. For the remainder of the talk, I will focus on the 3-proposition case.

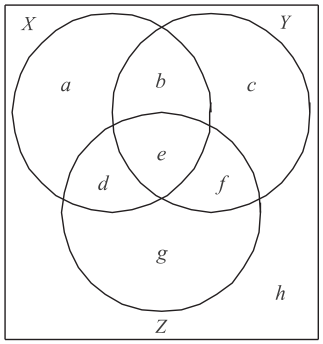

$$h = 1 - (a + b + c + d + e + f + g)$$

$$\Pr(X) = a + b + d + e$$

$$\Pr(Y) = b + c + e + f$$

$$\Pr(Z) = d + e + f + g$$

$$\Pr(X \mid Y) = \frac{\Pr(X \,\&\, Y)}{\Pr(Y)} = \frac{b + e}{b + c + e + f}$$

$$\Pr(Y \mid {\sim}Z) = \frac{\Pr(Y \,\&\, {\sim}Z)}{\Pr({\sim}Z)} = \frac{b + c}{1 - (d + e + f + g)}$$

- Using this technique, we can translate any equation, inequation, or inequality in (3-event) probability calculus into a simple algebraic formula in terms of the 7 variables a, b, c, d, e, f, and g.
- Moreover, a probability function over such a 3-event space is simply an assignment of real numbers on $[0, 1]$ to $a, \ldots, h$ such that $a + \cdots + h = 1$.
- I prefer *stochastic truth tables* to Venn Diagrams for representing probability models (easier to generalize to $n > 3$). Example:

X	Y	Z	States	$\Pr(s_i)$
T	T	T	s_1	$\Pr(s_1) = e$
T	T	F	s_2	$\Pr(s_2) = b$
T	F	T	s_3	$\Pr(s_3) = d$
T	F	F	s_4	$\Pr(s_4) = a$
F	T	T	s_5	$\Pr(s_5) = f$
F	T	F	s_6	$\Pr(s_6) = c$
F	F	T	s_7	$\Pr(s_7) = g$
F	F	F	s_8	$\Pr(s_8) = h$

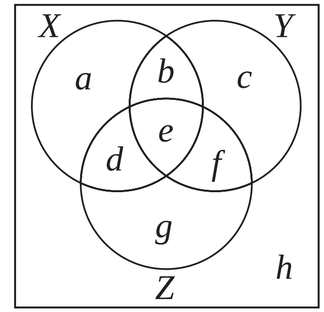

Algebraic Probability Calculus II

- The class of questions I've been discussing are of the following form:
  - Do there exist real numbers $a, \ldots, g$ satisfying all members of a set of n simple algebraic formulae $\mathcal{S} = \{P_1(a, \ldots, g), \ldots, P_n(a, \ldots, g)\}$?
  - In other words, we're interested in claims of the following form:

    $$(\exists a \in \mathbb{R}) \cdots (\exists g \in \mathbb{R})[P_1(a, \ldots, g) \,\&\, \cdots \,\&\, P_n(a, \ldots, g)]$$

    where the P_i are simple algebraic statements in terms of $a, \ldots, g$.
  - The first eight statements in $\mathcal{S}$ will state that $a, \ldots, h$ are on $[0, 1]$, and that they sum to 1 — *i.e.*, that the $a, \ldots, h$ are *basic probabilities*.
- Example. Are there $a, \ldots, h \in (0, 1)$ such that $a + \cdots + h = 1$, *and*
  1. $\Pr(X \,\&\, Y) = \Pr(X) \cdot \Pr(Y)$ $[b + e = (a + b + d + e) \cdot (b + c + e + f)]$,
  2. $\Pr(X \,\&\, Z) = \Pr(X) \cdot \Pr(Z)$ $[d + e = (a + b + d + e) \cdot (d + e + f + g)]$,
  3. $\Pr(Y \,\&\, Z) = \Pr(Y) \cdot \Pr(Z)$ $[e + f = (b + c + e + f) \cdot (d + e + f + g)]$,
  4. $\Pr(X \,\&\, (Y \,\&\, Z)) \neq \Pr(X) \cdot \Pr(Y \,\&\, Z)$ $[e \neq (a + b + d + e) \cdot (e + f)]$?

Algebraic Probability Calculus III

- Surprising (less trivial) Example: Do all relevance measures $\mathfrak{c}$ satisfy (†)?

(†) $\Pr(H \mid E_1) \geqslant \Pr(H \mid E_2) \Longrightarrow \mathfrak{c}(H, E_1) \geqslant \mathfrak{c}(H, E_2)$.

- Bayesians have *assumed* that (†) holds for all four of our $\mathfrak{c}$'s (and many others). But, using `PrSAT` (see below), we have shown this to be false.
- How can we express (†) *algebraically* (let $X = H$, $Y = E_1$, and $Z = E_2$)?
  - First, translate the left hand side of (†):

    (1) $\dfrac{b + e}{b + c + e + f} \geqslant \dfrac{d + e}{d + e + f + g}$

  - Second, pick a measure $\mathfrak{c}$, and translate the right side (here, $\mathfrak{c} = s$):

    (2) $\dfrac{b + e}{b + c + e + f} - \dfrac{a + d}{1 - b - c - e - f} \geqslant \dfrac{d + e}{d + e + f + g} - \dfrac{a + b}{1 - d - e - f - g}$

- Question: Do there exist $a, \ldots, h \in (0, 1)$ which sum to one and which make (1) true and (2) false? This is the algebraic equivalent of asking whether (†) is false (for $\mathfrak{c} = s$) in some 3-event model $\mathcal{M}$. Answer: YES!

Algebraic Probability Calculus IV

- Problems like ours, on the satisfiability of sets of statements in probability calculus are expressible in the *theory of real-closed fields* (TRCF).
- A *real-closed field* (RCF) is a structure $\langle F, 0, 1, +, -, *, ^{-1}, < \rangle$ such that:
  1. $\langle F, 0, 1, +, -, *, ^{-1} \rangle$ is a field.
  2. (a) $x \not< x$
     (b) $x < y \;\&\; y < z \Rightarrow x < z$
     (c) $x < y \Rightarrow x + z < y + z$
     (d) $x, y > 0 \Rightarrow x * y > 0$ [$x > 0$ iff $0 < x$]
     (e) $x > 0 \vee x = 0 \vee 0 > x$
  3. Every positive element of F has a square root in F and every odd degree polynomial in $F[x]$ has a root in F. [The set $\mathbb{R}$ forms an RCF.]
- TRCF is overkill for us. We only need the fragment of TRCF that involves *existentially quantified, simple algebraic claims over* $\mathbb{R}$.

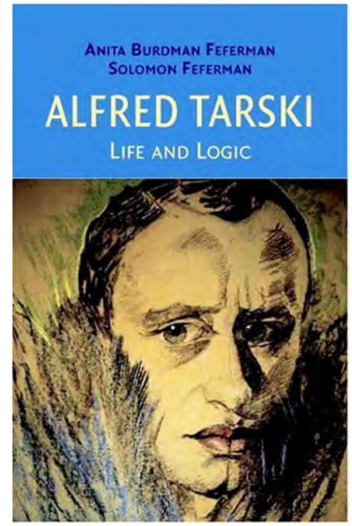

A Decision Procedure for the Probability Calculus I

- Tarski [35] described a decision procedure for the theory of real-closed fields. Thus, in principle, Tarski's method gives us a way to determine whether an arbitrary argument in probability calculus is valid.
- Tarski's idea is called *elimination of quantifiers*. He showed that a formula $(\exists x)P(x, a)$, where $P(x, a)$ is quantifier-free, is equivalent to a quantifier free formula $Q(a)$. [a can stand for several variables $a_1, \ldots, a_n$]
- Simple Example: $P(x, a)$ might have the form $f(x, a) = 0 \;\&\; g(x, a) = 0$, where f and g are polynomials, so we are asking for the condition(s) on a under which the polynomials f and g have a common root.
- It is a classical result of algebra that there is a polynomial $Q(a)$, called the *resultant* of f and g, which vanishes exactly when f and g have a common zero. Tarski's method is a (vast) generalization of this result.
- Tarski not only showed Q exists, he showed how to *compute* Q from P.

- Applying this procedure again and again, we can strip off one quantifier after another (from the inside out), eliminating all the quantifiers in a formula with nested quantifiers (note: we can write $\forall$ as $\sim\exists\sim$).
- The crux of the matter, then, is the elimination of a single ($\exists$) quantifier.
- This is a piece of algebra whose historical roots go back to Sturm's theorem, which counts the number of roots of a polynomial in an interval in terms of the alternations of signs in the coefficients.
- I won't get into the details of Tarski's method. But, I will say something about its complexity. Unfortunately, it is *very* complex. Its complexity cannot be bounded by any tower of exponentials – see [7] for discussion.
- Even for our simple class of problems (strings of 7 $\exists$'s binding 7 variables), Tarski's method is (in general) not really feasible.
- The intuitive reason why Tarski's method is so inefficient is that it eliminates "one quantifier at a time," and the formula "expands doubly-exponentially" each time a quantifier is eliminated.

A Decision Procedure for the Probability Calculus II

- George Collins invented an improved quantifier-elimination method known as *cylindric algebraic decomposition* (CAD) [6].
- The method decomposes a set defined in Tarski's language into a disjoint union of finitely many cells [a CAD], such that the polynomials involved in the definition of the original set do not change sign on any cell.
- Geometrically, $\exists$-quantification corresponds to projection onto a subspace with fewer variables. The projections of a set defined by a CAD are also defined by a CAD, as is the complement of a set defined by a CAD.
- Collins's CAD algorithm is "only" double exponential in the # of variables (polynomial in #/degree of polynomials, bit length of coefficients, and # of atomic formulas). This is the *lower bound* on the complexity of quantifier elimination in the (*general*) TRCF [7].
- Unlike Tarski's method, Collins's CAD method in some sense "eliminates all the quantifiers at once." This is why it is so much more efficient.

- Hong [18] improved further on Collins's CAD [*partial* CAD `qepcad`].
- This program has subsequently been further improved upon by many others, and is now publicly available on the Web [3] (linux/unix/PC). [`http://www.cs.usna.edu/~wcbrown/research/qebycad/Tutorial/Tutorial.html` is a nice tutorial on quantifier elimination algorithms, including CAD.]
- Some of `qepcad`'s functionality was implemented (by Strzebonski [34]) in *Mathematica* 4.1 (`Experimental`), and has since become part of the main Kernel in versions 5^+. It's pretty good on our class of problems.
- *Mathematica* now contains a suite of CAD functions, including `CylindricalDecomposition` (which computes CADs), `Resolve` (which eliminates quantifiers using CAD), and `FindInstance` (which finds assignments to variables that satisfy a set of formulae in TRCF).
- In light of these developments (& Moore's Law), Michael Beeson [2] remarks "there is some hope of solving interesting, even open, problems that are too hard for humans, before the exponential behavior of the algorithm takes its toll."

A Decision Procedure for the Probability Calculus III

- This last function `FindInstance` is particularly useful for our purposes, since it is designed, specifically, for settling pure $\exists$-questions, like ours.
- `FindInstance` takes as its argument a finite set $\mathcal{S}$ of equations, inequations, inequalities (in TRCF) over a finite set of real variables $\mathcal{V}$.
- `FindInstance` outputs an assignment of real numbers (if one exists) to the variables in $\mathcal{V}$, which satisfies all of the members of $\mathcal{S}$. If $\mathcal{S}$ is unsatisfiable, then `FindInstance` outputs "{ }".
- Using the translation procedure above (and the `FindInstance` function), I developed a *Mathematica* function `PrSAT`, which takes as input a set $\mathcal{S}$ of equations, inequations, or inequalities in probability calculus.
- If $\mathcal{S}$ is satisfiable (in TRCF), then `PrSAT` returns a probability model satisfying all the members of $\mathcal{S}$. If not, `PrSAT` returns "{ }".

A Decision Procedure for the Probability Calculus IV

- As you might imagine, `PrSAT` is quite useful for a practicing Bayesian confirmation theorist! [I used an early prototype in my dissertation [12].]
- `PrSAT` is particularly useful for finding models of complicated sets of probabilistic equations, inequations, and inequalities. While `PrSAT` does not generate *proofs*, it does *verify* theoremhood (which is also useful!).
- With J. Alexander's and B. Blum's help, `PrSAT` is now a *Mathematica package* [15], with additional functionality (*e.g.*, a random search algorithm — see below).
- **Demo**. The first example is quite well-known ([31], p. 85). Problem: show that X can be independent of each of Y and Z, but *dependent* on Y & Z.
- The second example is the "surprising" one I mentioned above. The model I found using `PrSAT` came as a shock to Bayesians [24, *fn*. 11].
- `PrSAT` can now easily decide all questions mentioned here. I use it to teach, and to find new results (some mentioned in my HPS talk this morning).

A Decision Procedure for the Probability Calculus V

- PrSAT is highly effective for problems involving $\leqslant 3$ propositions (I've not seen any problems of this size that I couldn't solve w/PrSAT).
- PrSAT becomes inefficient for larger spaces. But, Ben Blum has developed a random search add-on to PrSAT, which is included in the latest version. This algorithm finds models (generally) in *linear* time.
- Galen Huntington just finished a thesis [21] on decision procedures for the $\exists$-fragment of TRCF. There are single-exponential algorithms for $\exists$TRCF, and Galen is (as far as we know) the first to fully implement one.
- At some point, I need to port my own *Mathematica* code (which is interpreted and slow), and merge it with Galen's (Haskell) code.
- My interest in such methods is largely *instrumental*. It grew out of a desire to be able to reconstruct, understand, and improve upon arguments involving the probability calculus that appear in the PoS literature.
- I think this is a promising area for collaborative research in several fields.

References

[1] E. Adams, *A primer of probability logic*, CSLI Lecture Notes, vol. 68, 1998.

[2] M. Beeson, *The mechanization of mathematics*, Alan Turing: Life and Legacy of a Great Thinker (C. Teuscher, ed.), Springer-Verlag, 2004.

[3] C. Brown, *Quantifier elimination by partial cylindrical algebraic decomposition*, `http://www.cs.usna.edu/~qepcad/B/QEPCAD.html`.

[4] B. Caviness and J. Johnson (eds.), *Quantifier elimination and cylindrical algebraic decomposition*, Texts and Monographs in Symbolic Computation, Springer-Verlag, 1998.

[5] D. Christensen, *Measuring confirmation*, Journal of Philosophy **XCVI** (1999), 437–61.

[6] G. Collins, *Quantifier elimination for real closed fields by cylindrical algebraic decomposition*, reprinted in [4], pp. 85–121.

[7] J.H. Davenport and J. Heintz, *Real quantifier elimination is doubly exponential*, Journal of Symbolic Computation **5** (1988), no. 1-2, 29–35.

[8] J. Earman, *Bayes or bust?*, MIT Press, 1992.

[9] E. Eells, *Rational decision and causality*, Cambridge University Press, 1982.

[10] R. Festa, *Bayesian confirmation*, Experience, Reality, and Scientific Explanation (M. Galavotti and A. Pagnini, eds.), Kluwer Academic Publishers, 1999, pp. 55–87.

[11] B. Fitelson, *The plurality of Bayesian measures of confirmation and the problem of measure sensitivity*, Philosophy of Science **66** (1999), S362–S378, `http://fitelson.org/psa.pdf`.

[12] ———, *Studies in Bayesian confirmation theory*, Ph.D. thesis, University of Wisconsin–Madison (Philosophy), 2001, `http://fitelson.org/thesis.pdf`.

[13] ———, *Review of Joyce's [23]*, Mind **112** (2003), no. 447, `http://fitelson.org/joyce.pdf`.

[14] ———, *A decision procedure for probability calculus with applications*, Review of Symbolic Logic **1** (2008), 11–25.

[15] B. Fitelson, J. Alexander, and B. Blum, *PrSAT*, 2009, *Mathematica* Package. PrSAT can be downloaded from the following URL: `http://fitelson.org/PrSAT/`.

[16] D. Gillies, *In defense of the Popper-Miller argument*, Philosophy of Science **53** (1986), 110–113.

[17] I. Good, *Weight of evidence: a brief survey*, Bayesian Statistics 2, North-Holland, 1985.

[18] H. Hong, *An improvement of the projection operator in cylindrical algebraic decomposition*, reprinted in [4], pp. 166–173.

[19] P. Horwich, *Probability and evidence*, Cambridge University Press, 1982.

[20] C. Howson and P. Urbach, *Scientific reasoning: The bayesian approach*, Open Court, 1993.

[21] G. Huntington, *Decision procedures for the existential theory of the reals*, Ph.D. thesis, University of California–Berkeley (Group in Logic & The Methodology of Science), 2008.

[22] R. Jeffrey, *Probability and the art of judgment*, Cambridge University Press, 1992.

[23] J. Joyce, *The foundations of causal decision theory*, Cambridge University Press, 1999.

[24] ———, *Bayes's theorem*, The Stanford Encyclopedia of Philosophy (Fall 2003 Edition) (E.N. Zalta, ed.), 2003, URL = `http://plato.stanford.edu/entries/bayes-theorem/`.

[25] A. Kolmogorov, *Foundations of probability*, second english ed., AMS Chelsea Publishing, 1956.

[26] H. Kyburg, *Recent work in inductive logic*, Recent Work in Philosophy, 1983, pp. 87–150.

[27] P. Milne, $\log[p(h/eb)/p(h/b)]$ *is the one true measure of confirmation*, Phil. Sci. **63** (1996).

[28] T. More, Jr., *On the construction of Venn diagrams*, J. Symbolic Logic **24** (1959), 303–304.

[29] K. Popper and D. Miller, *The impossibility of inductive probability*, Nature **302** (1983), 687–688.

[30] R. Rosenkrantz, *Bayesian confirmation: Paradise regained*, The British Journal for the Philosophy of Science **45** (1994), 467–476.

[31] S. Ross, *A first course in probability*, fourth ed., Prentice Hall, 1994.

[32] D. Schum, *The evidential foundations of probabilistic reasoning*, John Wiley & Sons, 1994.

[33] E. Sober, *No model, no inference: A bayesian primer on the grue problem*, Grue! The New Riddle of Induction (D. Stalker, ed.), Open Court, 1994.

[34] A. Strzebonski, *Solving algebraic inequalities with version 4*, The Mathematica Journal **7** (2000).

[35] A. Tarski, *A decision method for elementary algebra and geometry*, UC Press, 1951.