LUND UNIVERSITY

Bayesian Argumentation: The Practical Side of Probability

Zenker, Frank

2012

[Link to publication](#)

Citation for published version (APA):

Zenker, F. (Ed.) (2012). *Bayesian Argumentation: The Practical Side of Probability*. (Synthese Library). Springer.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

Foreword

This volume collects contributions to the international workshop on Bayesian Argumentation held 22-23 October 2010 at the Philosophy Department of Lund University, Sweden.

I would like to express my gratitude to the Swedish Research Council (VR) and the Wenner-Gren Foundations for their generous financial support of this event, and to all speakers and participants for fruitful discussions.

Thanks are due to Stefan Schubert for useful comments on the introduction, to an anonymous reviewer and, last but not least, to Ties Nijssen, Ingrid van Laarhoven, and Christi Lue at Springer's Dordrecht office.

Lund, June 2012

F.Z.

Bayesian Argumentation—The Practical Side of Probability

Editor's Introduction

Frank Zenker
Lund University
Department of Philosophy & Cognitive Science
Kungshuset, Lundagård
22 222 Lund, Sweden
frank.zenker@fil.lu.se

1. Introduction

It is often found not too difficult to compare arguments with respect to their comparative strength, or their rational persuasiveness. For present purposes, this may be understood as the capacity of reasons to reduce, sustain, or enhance the extent to which a rational agent endorses a conclusion. Perhaps the easiest way of assessing argument strength consists in asking listeners for an opinion. But, clearly, what does persuade need not be coextensive with what *should* persuade.

A less simple way consists in deploying a normative standard—here: the theory of probability, particularly its notion of evidential support—in order to answer two questions: First, what part of probability theory might be used towards arriving at a *normative* account that yields a motivated separation of natural language arguments into good and less good ones? Secondly, are the assessments of argument strength performed in laymen or institutional contexts (such as a court room) to some extent consistent with that theory?

Contributions to this volume converge upon the use of Bayes's theorem (which is introduced further below) as *the* norm in the light of which evidence—"arriving" in various forms, e.g., as a third party report, a test result, or a direct observation—affects the endorsement of some hypothesis. A subset of these contributions also report data, some of which has been methodologically hardened in controlled experimental settings, on the extent to which arguers comply with this

norm, or not. Further contributions address challenges arising in (computationally) modeling rational agents and their argumentative interactions, while yet others offer novel solutions to long-standing “anomalies.”

Below follows a brief introduction to the formal apparatus, particularly as it applies to the study of natural language argumentation, meant to “pave the way” for the novice reader. An overview of the book’s content is provided in Section 3.

2. The Bayesian Approach to Argumentation

Bayes’s theorem (Bayes 1763) expresses the relations between a hypothesis, H , and evidence, E , in terms of probability, P . On a subjectivist interpretation, ‘probability’ denotes *degrees of belief*. These degrees are mapped onto the unit interval $[0, 1]$ such that $P(H)=1-P(\text{non}H)$. The relata are variously referred to as:

- $P(H)$ the *prior* or *unconditional probability* of a hypothesis (given no evidence)
- $P(E)$ the *marginal* or *a priori probability* of the evidence
- $P(H|E)$ the *direct* or *posterior* or *conditional probability* of a hypothesis given the evidence
- $P(E|H)$ the *inverse* or *conditional probability* of the evidence given the hypothesis (aka. the *likelihood of the evidence*)

To reach terms standardly used in the study of natural language argumentation (and episodes of reasoning thus suggested), ‘evidence’ will be interpreted as reason, ground or argument, and ‘hypothesis’ as conclusion or proposal. Those used to representing argumentative structures as premise-conclusion complexes of the form $P_1, \dots, P_n; \text{ ergo } C$ need to project premises into evidence, and conclusion into hypothesis.

Bayes’s theorem may count as non-controversial and takes the following form¹, where it is assumed—here and further below—that $0 < P(H) < 1$ and $0 < P(E) < 1$.

¹ To arrive at the theorem, one may start from a basic axiom of probability theory which states the probability of the conjunction of H and E , $P(H \& E)$, to be the probability of E multiplied by the probability of H given E , as expressed in (1).

$$(BT) \quad P(H|E) = \frac{P(H) P(E|H)}{P(E)}$$

The theorem may be said to relate the *direct* probability of a hypothesis conditional on evidence, $P(H|E)$, to the *inverse* probability of the evidence conditional on the hypothesis, $P(E|H)$. A more useful version—more useful, because $P(E)$ can be expressed in extant terms²—is the following:

$$(BT') \quad P(H|E) = \frac{P(H) P(E|H)}{P(H) P(E|H) + P(nonH) P(E|nonH)}$$

The basic idea underlying most uses of Bayes's theorem is that a hypothesis is supported by any evidence which is rendered (either sufficiently or simply) probable by the truth of that hypothesis. For the dynamic case, this entails that the probability of a hypothesis increases to the extent that evidence is more likely if the

$$(1) \quad P(H \& E) = P(E) P(H|E)$$

The same is stated in (2):

$$(2) \quad P(H \& E) = P(H) P(E|H)$$

This allows us to equate the right sides of (1) and (2) to reach:

$$(3) \quad P(E) P(H|E) = P(H) P(E|H)$$

Upon which dividing both sides of (3) by $P(E)$ yields (BT):

$$(BT) \quad P(H|E) = \frac{P(H) P(E|H)}{P(E)}$$

² To express $P(E)$ in extant terms, note that the law of total probability states the probability of E to be the probability of H and E, plus the probability of *nonH* and E (where '*nonH*' designates the complement, or the negation, of H). In terms:

$$(1) \quad P(E) = P(H \& E) + P(nonH \& E).$$

Expressed in conditional form (see the previous footnote), this yields:

$$(2) \quad P(E) = P(H) P(E|H) + P(nonH) P(E|nonH)$$

hypothesis were true than if it were false. (Novices may want to read the last sentence twice!)

The theorem allows calculating $P(H|E)$, which is the degree of belief in the hypothesis H to which (having received) evidence E does lead—or *should* lead, insofar as Bayes’s theorem is treated as a normative update rule—, provided assumptions on:

(i) The initial degree of belief in the hypothesis, $P(H)$. This is normally kept distinct from 1 or 0, for otherwise evidence will not affect $P(H)$. Hence, “good Bayesians” cannot (without “signing-off” on evidence) entertain full belief, while dogmatists cannot (without inconsistency) fully endorse the Bayesian update rule.

(ii) Qualities of the evidence, $P(E)$. As explained in footnote 2, $P(E)$ can be calculated from the values of $P(H)$ and $P(E|H)$.

(iii) The relation between E and H , particularly how (much more) likely, or expectable, evidence would be if the hypothesis were true than if it were false. This comes down to comparing $P(E|H)$ with $P(E|nonH)$.

Bayes’s theorem features two limiting cases. One captures a situation where a hypothesis *entails* the evidence—which, in logical terms, is expressed as $H \rightarrow E$, or as $P(E|H)=1$ in probabilistic terms. The other captures the converse situation where evidence entails the hypothesis, expressed as $E \rightarrow H$ or $P(H|E)=1$. Conveniently, amongst others, the following hold:

- (1) $P(H|E)=1$, if $P(E|nonH)=0$
- (2) $P(H|E)<1$, if $P(E|nonH)>0$
- (3) $P(E|H)/P(E|nonH) > 1$, iff $P(H|E) - P(H) > 0$

In words: (1) we have full belief in the hypothesis given the evidence, *if* we have a zero degree of belief in the evidence given the negation of the hypothesis. (2) We have less than full belief in the hypothesis given the evidence, *if* our degree of belief in the evidence given the falsity of the hypothesis is greater than zero. (3) The ratio of the evidence likelihoods—i.e., the degree of belief in the evidence given the hypothesis over the degree of belief in the evidence given the negation of the hypothesis—is greater than one *if, and only if*, the difference between the degree of

belief in the hypothesis given the evidence (aka. the posterior probability) and the degree of belief in the hypothesis (aka. the prior probability) is greater than zero.

Put differently, (1) and (2) assert that *fully* endorsing H (as true *vis-à-vis* E) requires E to be *completely* unexpectable under the assumption that H is false. Conversely, as long as E is somewhat expectable under the assumption that H is false, H cannot be fully endorsed (as true *vis-à-vis* E). And (3) asserts that a difference between prior and posterior probability—thus, any support lent to the hypothesis by the evidence—is mirrored by a likelihood ratio greater than 1.

Perhaps instructively for those more familiar with classical logic, when the relation between evidence and hypothesis is neat enough for the truth of one to ensure the truth of the other, Bayes's theorem “degrades” to the bi-conditional of deductive logic, $H \leftrightarrow E$.³ Therefore, some hold that much of what can be done in classical logic may be done with (limiting case instances of) Bayes's theorem. Where such neat relations are not at hand, the theorem traces the effect that

³ Floridi (2009) provides an instructive example, here slightly adapted. Unrealistically(!), it assumes *perfect* antivirus software: “Maggie thought that, *if* ‘the email was infected’ [abbreviated as the hypothesis H], *then* ‘the antivirus blocked it’ [abbreviated as evidence E], and since the quarantine folder contains only emails blocked by the antivirus, then all the emails in it must be infected. More formally, she reasoned that: ‘ $H \rightarrow E$, E *therefore* H’. Jill explains to Maggie that the previous inference is a typical fallacy (called: *affirming the consequent*), but that she should not feel silly at all. For, consider Bayes's theorem once again. Look at the formula $P(E|nonH)$ [...] which indicates the probability that the antivirus blocks the email (E) when the email is not infected (*nonH*). *Suppose we have perfect, infallible antivirus software*. This will generate no false positives (no mistakes). But if there are no false positives, that is, if $P(E|nonH)=0$, then $P(H|E)=1$ and Bayes's theorem is degraded to a double implication: ‘ $H \leftrightarrow E$, E *therefore* H’, which is what Maggie perhaps had in mind. On the other hand, if there are some false positives, that is, if $P(E|nonH)>0$, then $P(H|E)<1$ and the formula bears a strong family resemblance to the fallacy of *affirming the consequent*: ‘ $H \rightarrow E$, E *therefore* H’, which is what Maggie might also have had in mind” (Floridi 2009: 399).

(receiving) evidence makes (or should make) on the degree to which a hypothesis is supported *vis-à-vis* endorsing the negation of the same hypothesis, provided a prior degree of belief in the hypothesis and an estimate of the likelihood of the evidence.

Since, formally, this is very much as it should be, challenges arise “only” in interpreting the Bayesian terms, and in reasonably choosing the numerical values by which these terms express the kinds of evidential considerations featured in natural language arguments. Take, for instance, the contention that hypotheses should be submitted to *severe* tests, and that arguments which report severe tests are stronger (or more persuasive) than those reporting less severe tests. In Bayesian terms, this means the probability of obtaining evidence (in case the hypothesis is true) should be comparatively low. So, assuming two tests to choose from, the *more severe* should sport a lower value for $P(E|H)$.

As pointed out above, evidence supports (“strengthens”) or undermines (“weakens”) a hypothesis, or else is irrelevant to it. The degree of confirmation (or support) that evidence lends to a hypothesis may be expressed as the difference between the posterior and the prior probability of a hypothesis: $P(H|E) - P(H)$. A second and equally defensible measure of support is the likelihood ratio: $P(E|H) / P(E|nonH)$.⁴ These measures have been proposed as estimates of the *force of an argument*, since they express the magnitude of a change of conviction which receiving evidence brings about *vis-à-vis* a prior conviction. Relatedly, the absolute degree of conviction to which receiving an argument leads—i.e., the posterior probability—may then be used to estimate the *strength of an argument*.

This coinage goes back to Oaksford and Hahn who observe that, consequently, the *relative* degree of conviction which evidence brings about may be the same, even if two agents differ with respect to the prior probabilities they endorse. This holds as long as discussants do not disagree about properties of the evidence such as being obtained from a trustworthy source or by a reliable method. Conversely, arguers would have to differ in ways that point beyond the priors for their disagreement to remain rationally reconstructable. In this limited but important sense, Bayes’ theorem promises to provide a normative theory of argument strength, given assumptions on the quality of reasons.

⁴ For other support measures, see Fitelson (2001) and Pfeifer (this volume).

With a view to natural language argumentation, the Bayesian approach recommends itself for its precise expressions, and—other than by successfully explaining the quality of arguments/fallacies *vis-à-vis* various contexts and audiences—also receives support from recent results in the study of human reasoning. According to authors such as Evans (2002), humans may generally not count as “deduction machines.” Rather, much of man-made reasoning appears to be consistent with some of the assumptions made in probabilistic modeling. Nevertheless, for reasons of computational complexity alone (Korb 2004), it is clear that humans are not “Bayesian machines” either.

The Bayesian approach to natural language argumentation is a quasi-natural choice, firstly, for the study of any argument which seeks to support, or undermine, a claim on the basis of statistical data. After all, on the Bayesian approach, the standards appealed to—i.e., those of inductive logic—will, in one way or another, be part of the reconstructive apparatus, and thus be available in argument evaluation.

Secondly, empirical research on message persuasiveness should find the Bayesian approach natural when explaining the differential persuasiveness of messages *vis-à-vis* various sources, contexts, and audiences. For instance, when receiving identical messages from a reliable *versus* an unreliable source, in the first case, $P(E|H)$ may reasonably be taken to exceed $P(E|nonH)$, providing one way of modeling the impact of trust on the degree to which a message is believed, or not.

Finally, the Bayesian approach challenges the view that the reconstruction and evaluation of natural language argumentation must either rely on informal means, or else be unrealistically confined to a comparatively small class of arguments which instantiate deductively valid structures. Surely, branding the Bayesian approach as *the* solution appears equally unrealistic. Nevertheless, as the contributions to this volume demonstrate, formal tools such as Bayes’ theorem have a rightful place both in argument reconstruction and argument evaluation.

3. Chapter Overview

Comprising theoretical backgrounds and disciplines as diverse as social psychology, jurisprudence, and philosophy, contributions remain unified with respect to the normative standard to which they relate. While chapters are self-contained, readers

may appreciate the division into: The Bayesian Approach to Argumentation; The Legal Domain; Modeling Rational Agents; and Theoretical Issues.

3.1 The Bayesian Approach to Argumentation

Ulrike Hahn, Mike Oaksford, and Adam J.L. Harris lay out a Bayesian perspective on testimony and argument. Normally treated as independent variables or alternative routes to audience persuasion, the authors build on empirical studies which validate a rather precise connection between source and content considerations of argument. This yields an alternative to standard default and plausibility treatments, and crucially differs from the latter with respect to the principle employed for evaluating linked argumentation structures. Rather than the so-called MAXMIN principle, their contribution deploys Bayes's theorem, and extends the investigation to cases of witnesses who differ not only with respect to the reliability of testimony, but also with respect to the argumentative strength of its content. After contrasting possible ways of modeling source reliability either exogenously or endogenously, the authors proceed to a "rehousing" of argumentation schemes within a Bayesian framework, resulting in a network representation of Walton's well-known critical questions for the appeal to expert opinion. As they content, "reasoning appropriately about source reliability in any given context involves Bayesian inference within a suitable model. How complex that model needs to be depends on the context, such as the relevant dimensions of variation within that context."

Mike Oaksford and Ulrike Hahn report on an empirical study that investigates why *ad hominem* argumentation is found convincing, and contrast their findings with recent empirical results obtained from the perspective of the normative pragma-dialectical model. In the latter, the reasonableness of an *ad hominem* argument is construed as a function of the discussion stage in which it occurs. The argument form is deemed an illegitimate move (aka. a fallacy) only in the opening stage of a critical discussion, where it violates the "Freedom Rule" (according to which discussants may not prevent each other from forwarding a standpoint). In contrast, Oaksford and Hahn deploy a Bayesian model which, amongst others, respects considerations of source reliability. Moreover, they also vary further conditions such as the initial degree of belief in a conclusion, and whether the *ad*

hominem appears as a pro or a con argument. Although further empirical investigation into this and other argument forms is deemed necessary, their study partly fails to corroborate—and thus challenges—the explanation offered on the pragma-dialectical model, insofar as they find “no differences between different types of argumentum ad hominem, where the freedom rule was violated, and a control, which introduced no violation of the freedom rule.”

3.2 The Legal Domain

Matthias Grabmair and *Kevin D. Ashley* distinguish four kinds of interdependent uncertainties that lawyers must plan for in litigation cases, then illustrate these vis-à-vis the claim of trade-secret-misappropriation recently brought against Facebook founder Mark Zuckerberg. Typically, in deciding what legal claim to assert, lawyers will profit from a structured assessment of what is normally called the “strength of a case.” It is for this purpose, then, that factual, normative, moral and empirical uncertainties become relevant, insofar as winning or losing a case may be understood as a function of having estimated these correctly, or not. Employing an extension of the Carneades argumentation model, the authors use probability values to model that audiences (here, judges and juries) may accept or reject some statement or principle to some degree, or else be undecided about it. This inference is a function of probability values representing the audiences’ belief in certain assumptions and assessments of an argument’s persuasiveness. The resulting Bayesian model, which uses argument weights, can represent pro and con arguments for or against some claim. Moreover, it allows for dynamic weights. The authors further suggest and illustrate that weights can be subject to argumentation as well. Thus, a formal model of legal argument mandates amongst other things “the moral assessment of a case [e.g., by a judge] to influence the probability with which certain legal arguments will prevail over others.”

Amit Pundik presents a case study on the use of statistical evidence in criminal courts, and defends the standpoint that, under certain assumptions, courts can have good reasons to refrain from the use of statistics, and experts delivering them. His case involves the correct diagnosis of sudden infant death syndrome (SIDS), and addresses the question whether public and scholarly attention has been

wrongly directed to the statistics—which, as it happened, were seriously flawed. Pundik reviews and rejects extant explanations that purport to show why the SIDS statistics have led to a wrongful charge against an innocent mother, then draws on the theory of contrastive explanation to provide an alternative account. According to this theory, the statistics on spontaneous child death would have made sense in the context of this trial only if they had been compared with statistics on the preferred contrast class (here: particular acts of murder). Pundik argues that, regardless of whether a comparison of probabilities between contrasting explanations is in fact possible, it should not be conducted as part of criminal proceedings. He concludes that his case study “should serve as a warning against any attempt to prove the fact of causation using statistical evidence about the rate of potential exonerating causes.”

3.3 Modeling Rational Agents

Erik Olsson presents a computer-based simulation environment, called *Laputa*. This provides a Bayesian model of group interaction which may be interpreted as exchanges of arguments for or against some proposition p among agents/inquirers. Provided certain constraints are met—amongst others, that message sources are independent; that inquirers cooperate in the sense of providing arguments that are novel to their interlocutors; that arguing pro (con) p entails personally endorsing (not) p to a reasonable degree; that proponents deem their own arguments sound—agents can be modeled to update both their degree of belief in p , and their degree of trust in the message source. From this basis, Olsson proceeds to show that, over time, agents in the group will polarize in the sense of endorsing ever greater degrees of belief in p , while assigning ever less trust to those agents endorsing not p (or vice versa). Consequently, he can suggest that “polarization on the belief level is accompanied, in a sense, with polarization on the trust level.” Moreover, he can demonstrate that, on this model, seemingly minute differences between degrees of belief and degrees of trust will, over time, lead to polarization. Consequently, these simulations support the claim “that even ideally rational inquirers will predictably polarize or diverge under realistic conditions.”

Gregor Betz provides an outline of his recent theory of dialectical structures, and investigates the relation between such structures and truth-likeness (*verisimilitude*). Amongst others, his theory serves to reconstruct attack and support relations between premises and conclusion at the level of complex debates (aka. “controversies”). In turn, this provides the basis for a structured evaluation of the positions endorsed by proponents and opponents. In particular, the pre-theoretic idea of a “(comparative) strength of justification” can be rendered formally precise by means of the notion “degree of justification”—which obeys Bayes’s rule—and that of “degree of partial implication.” Given assumptions on background knowledge, Betz’s account yields ways of measuring the inferential density of a given state of a debate. Consequently, hypothetical debate progressions may be subjected to computer simulation, in order to investigate the relation between the degree of justification, the inferential density, and the proportion of true sentences. As he demonstrates, “[a]dopting positions with a high degree of justification fosters the achievement of a basic epistemic aim [i.e., to acquire true beliefs], and that’s why degrees of justification should guide the belief-forming process of rational beings.”

Robert van Rooij and *Kris de Jaegher* address argumentation from the perspective of game theory. Their starting point is the observation that “[w]e communicate more than standard game theory predicts.” Moreover, standard game theory is seemingly unable to explain deception and the strategic manipulation of messages. This suggests that standard game theory suffers from one or the other idealized assumption which reduces its explanatory power in application to natural language argumentation. Of these assumptions they identify, and then significantly relax, the following three: what game is being played is common knowledge; agents are completely rational; the assessment of probabilities and subsequent decision making is independent of the way alternatives and decision problems are stated (framing effects). In each case, they demonstrate in a formally precise way that explanatory scope increases when assumptions are deemed false. For instance, that sub-optimal outcomes are chosen, or dominated strategies nevertheless played, may be explained by participants’ lack of insight into the full breadth of options. In line with recent criticisms of the ideal rational agent assumption, their general conclusion

is that “we talk and argue so much because we believe others are bounded rational agents.”

3.4 Theoretical Issues

Tomoji Shogenji addresses the long standing contention that circular argumentation or reasoning is defective or fallacious. For instance, it is widely considered bad reasoning to invoke perceptual evidence to support the reliability of our sense apparatus since the reasoning already assumes the reliability of our sense apparatus when it invokes perceptual evidence. To rebuke what he calls the “myth of epistemic circularity,” Shogenji distinguishes two senses of the term ‘assume’—viz. to presuppose the truth of the hypothesis and to envision the truth of the hypothesis. According to Shogenji, assuming the truth of the hypothesis in the second sense is no more problematic than assuming the negation of the hypothesis in the reasoning of *reductio ad absurdum*. In *reductio ad absurdum*, we establish the truth of the conclusion by envisioning the falsity of the hypothesis and deriving a contradiction from it. In a similar way, Shogenji proposes a procedure of envisioning the truth of the hypothesis and then the negation of the hypothesis, to compare their respective degrees of coherence with the evidence. He demonstrates in the Bayesian framework that the evidence raises the probability of the hypothesis if the evidence is more coherent with the truth of the hypothesis than it is with the negation of the hypothesis. Applying this procedure to the perceptual support of the reliability of our sense apparatus, Shogenji contends that when the perceptual evidence is more coherent with the hypothesis of reliability than it is with the negation of the hypothesis, the evidence raises the probability of the reliability hypothesis without epistemic circularity.

Niki Pfeifer contrasts probabilistic with deductive logical treatments of natural language arguments that contain conditionals and argues for a combination of both. He stresses the importance of conditionals which are uncertain and which allow for exceptions such as the classic “Birds can fly” from default logic. Working within the framework of coherence based probability logic—a combination of subjective probability theory and propositional logic—, probability values are attached directly to the conditional event, and probabilities are conceived as degrees

of belief. Pfeifer's account traces "how to propagate the uncertainty of the premises to the conclusion" in a deductive framework, and yields a formal measure of argument strength that depends on two factors: the "location of the coherent probability interval" on the unit interval and the "precision" of the probabilistic assessment, i.e., the distance between the tightest coherent lower and the upper probability bounds of the conclusion. Thus, on Pfeifer's account, standard problems incurred when working with traditional measures of confirmation (e.g., How to connect premises containing conditionals?; How to conditionalize on conditionals?) are avoided, while the intuition that strong arguments should be those that imply precise assessments of the conclusions with high probability is recovered.

Jonny Blamey presents a novel solution to the preface-paradox by invoking considerations of stake size. The paradox pivots on the tension that arises when the degree of belief that is assigned to the conjunction of a set of propositions compares with the degree of belief assigned to each conjunct in a manner that lets the *whole* (the conjunction) come out as different from the sum of its *parts* (the conjuncts). This is normally accounted for by our intuitive tendency to be certain of each conjunct in a set of statements forming a conjunction, but to be less than certain of the conjunction seen as a whole. Working within the framework of evidential probability, Blamey builds on the idea that "the same evidence can fully justify a belief at low stakes but not fully justify the same belief at high stakes." More precisely, he lends himself of a betting model of belief, equates the value of knowledge with the value of the stake, and pairs this with the idea that a conjunction has greater informational content than the conjuncts such that "conjoining the propositions escalates the informational content exponentially by 1 over the conditional probability between the conjuncts." Consequently, Dutch books (i.e., bets resulting in sure losses) can be avoided provided, amongst others, that Blamey's *minimum constraint* is assumed to hold, according to which one "cannot prefer a bet at smaller stakes to a bet at larger stakes for the same price." Thus, it is shown not to be incoherent to remain less than evidentially certain at high stakes, while one may very well be evidentially certain at low stakes.

References

- Bayes, T., & Price, R. (1763). An Essay towards solving a Problem in the Doctrine of Chance. By the late Rev. Mr. Bayes, communicated by Mr. Price, in a letter to John Canton, M. A. and F. R. S. *Philosophical Transactions of the Royal Society of London*, 53, 370–418.
- Evans, J.S.T.B. (2002). Logic and Human Reasoning: An Assessment of the Deduction Paradigm. *Psychological Bulletin*, 128(6), 978–996.
- Floridi, L. (2009). Logical fallacies as informational shortcuts. *Synthese*, 167, 317–325.
- Fitelson, B. (2001). *Studies in Bayesian confirmation theory*. Dissertation, University of Wisconsin at Madison. <http://fitelson.org/thesis.pdf> (accessed 22 June 2012).
- Korb, K. (2004). Bayesian Informal Logic and Fallacy. *Informal Logic*, 23(2), 41–70.