

SOME RESULTS IN BAYESIAN CONFIRMATION
THEORY WITH APPLICATIONS

A Dissertation

Presented to the Faculty of the Graduate School
of Cornell University

in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

by

David Jehle

January 2009

© 2009 David Jehle
ALL RIGHTS RESERVED

SOME RESULTS IN BAYESIAN CONFIRMATION THEORY WITH
APPLICATIONS

David Jehle, Ph.D.

Cornell University 2009

This dissertation presents some formal results in Bayesian confirmation theory with novel applications to topics in contemporary epistemology, political science, and legal theory. The primary theme in each chapter is that *confirmation*, not *probability*, is the central formal notion one should adopt in approaching various topics in these three disciplines. In each chapter, I argue for this by highlighting some wide-ranging and important consequences that result from a confirmation-theoretic approach.

Chapter 1 provides an introduction to the machinery of Bayesian confirmation theory and, with a case study, demonstrates the significance of moving away from probability to confirmation.

Chapter 2 addresses the question of whether or not an epistemic closure principle is correct. Most contemporary epistemologists think that epistemic closure is obviously correct, in need of little defense or argument. I show that by shifting our focus from probability to confirmation, and exploiting the different formal properties exhibited by probability functions and confirmation functions, we are able to develop a novel case against closure. Here I also prove some new formal results that bear on the debate over the adequacy of Bayesian measures of confirmation.

Chapter 3 discusses a particular form of voting—*supermajority voting*—from a Bayesian perspective. The standard route taken to motivate supermajority voting is via the Condorcet framework and its well-known “jury theorem”. I show

that a Bayesian confirmation-theoretic approach provides a much more general and powerful approach to supermajority voting. The basic idea is that supermajority voting provides *superior evidence* to simple majority voting. From an epistemological point of view, I argue this makes supermajority voting preferable to simple majority voting. This chapter also presents some new mathematical results that improve upon the contemporary Bayesian literature on evidential support from multiple pieces of evidence.

Chapter 4 is an essay in Bayesian jurisprudence and legal epistemology. In Anglo-American criminal jurisprudence, one important and frequently used standard of criminal proof is *proof beyond reasonable doubt*. Recent work on defining the notion of proof beyond reasonable doubt, however, is rather vague and often confusing. In this chapter, I propose a new and precise account proof beyond reasonable doubt. My approach will make central use of Bayesian confirmation theory, and the guiding idea will be that proof beyond reasonable doubt is established when a certain level of confirmatory support is reached. I will also discuss the “presumption of innocence” doctrine, since it plays an important role in my account of proof beyond reasonable doubt. My discussion will make some important philosophical and mathematical advances on our understanding of both these doctrines.

BIOGRAPHICAL SKETCH

David Jehle was born in Royal Oak, MI, on August 29, 1983. He split his time growing up between Southern California and Detroit. He attended Wheaton College, IL, earning his B.A. (*magna cum laude*) in May 2004. He began his graduate studies at Cornell University in September 2004, completed his M.A. in January 2007, and from January 2007 to August 2008 was a visiting student at both UC Berkeley and Rutgers University. His dissertation defense was on September 5th, 2008.

This dissertation is dedicated to my parents, for their unwavering support.

ACKNOWLEDGEMENTS

Writing a dissertation is not something you could possibly do alone; it's something that results from the generous help, support, and encouragement of colleagues, friends, and family.

So many people have read my work and provided invaluable feedback. I can't hope to list them all here, but I owe thanks to Carl Ginet, Nicholas Silins, Tamar Gendler, Zoltan Szabo, Berit Brogaard, Mike Titelbaum, Kenny Easwaren, Nick Kroll, Raul Saucedo, David Van Bruwaene, Roald Nashi, Mark Heller, Franz Huber, Uriah Kriegel, Benj Hellie, Tomoji Shogenji, and Eric Hiddleston.

I want to highlight a few people for special thanks—without their help, I could not have completed this project. First, James Hawthorne: You're a "guru", a true expert, in all things Bayesian. Your mathematical help was absolutely invaluable at some parts of this dissertation. I simply could not have obtained many of the results here without your guidance. Thanks, too, for always being so encouraging in this project, even when I thought it was a dead end. Second, Branden Fitelson: You're a true friend and an amazing mentor. I'll always remember our weekly—sometimes bi-weekly—meetings in Berkeley. I can't tell you how much I've learned from you, both professionally and personally. Your dedication to me, to this project, has been unflagging. Your generosity, your patience, your acceptance of me—it has meant the world to me. That year at UC Berkeley was one of the fondest of my life so far. Finally, my advisor, Brian Weatherson: I've told you, and I've told every prospective and fellow student I ever come into contact with—you're the best advisor one could ever hope to have. You've spent countless hours meeting with me, emailing with me, talking with me, even at times when I was advancing some pretty ridiculous ideas. You gave me free reign to pursue my own interests, but at the crucial points you always pushed me to see what was and was not important. In so

many ways, I've tried to model—albeit inadequately—my approach to philosophy after you. You've been a great friend to me, too.

Finally, I owe thanks to my two closest friends, Adam Long and Brad Nagle, and most importantly my family (and extended family). Adam and Brad: Both of you have been there for me through everything significant (and insignificant) in my life. Your friendship, your loyalty over the years has meant everything to me, and I could never have finished this thesis without your humor and encouragement. To my uncle, Jay: You housed me for over a year, providing me—quite generously—with everything I needed to write this dissertation. You might not know it, but I wrote the vast majority of this dissertation at your house. The view of the Golden Gate Bridge and the San Francisco Bay didn't hurt. Thanks for everything. To my sister, Jessie: I can't thank you enough for listening to me and encouraging me at each stop along the way in graduate school. It's been a long—and sometimes crazy—journey, but you always insisted I keep my head up and press on. Without your words of support, I would've given up a long time ago. And last, but of course not least, my parents: I don't know where to begin with you two. Your patience with me has been incredible. You've supported me in everything—and, I mean, *everything*—in my life, and you never once questioned my goals or aspirations—words can't capture how important this has been to me. You two are my rock, my center—without you, not only could I not have completed this project, but it simply would not have been worth it. I dedicate this dissertation to you both, the very least I can do, considering everything you've done for me.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vii
1 Introduction	1
1.1 Overview	1
1.2 Bayesian Confirmation Theory	1
1.3 Case Study: Dogmatism and Bayesianism	7
1.4 What's to Come	11
2 A Bayesian Approach to Closure Failure	13
2.1 Introduction	13
2.2 Epistemic Closure	13
2.3 Some Dretskean Intuitions	16
2.4 Some Theorems	20
2.5 Closure Failure, Bayesian Style	24
2.6 Motivation	27
2.6.1 Handling Counterexamples	28
2.6.2 Applications	32
2.7 Conclusion	34
3 Supermajority Voting: A Bayesian Perspective	35
3.1 Introduction	35
3.2 Condorcet Rationale for Supermajority Requirement	37
3.2.1 The Model	37
3.2.2 Application to Supermajority Requirement	40
3.3 Results	42
3.4 Implications & Advantages	46
3.5 Conclusion	49
4 Proof Beyond Reasonable Doubt and the Presumption of Innocence: An Essay in Bayesian Jurisprudence	50
4.1 Introduction	50
4.2 Against Two Recent Accounts	51
4.2.1 First Account	52
4.2.2 Second Account	55
4.2.3 Moving Past Subjectivity	56
4.3 A New Bayesian Account	57
4.3.1 Prior and Posterior Probability	58
4.3.2 The Proposal—First Pass	60
4.3.3 The Presumption of Innocence	62
4.3.4 The Proposal—Official Version	71

4.4	Objections	72
4.5	Conclusion	76
A	Proofs	77
A.1	Chapter 2	77
A.2	Chapter 3	89
A.3	Chapter 4	96

CHAPTER 1

INTRODUCTION

1.1 Overview

This dissertation presents some formal results in Bayesian confirmation theory with novel applications to topics in contemporary epistemology, political science, and legal theory. Specifically, I will present new technical results in Bayesian confirmation theory and discuss their implications for epistemic closure (chapter 2), supermajority voting (chapter 3), and the legal notions of proof beyond reasonable doubt and the presumption of innocence (chapter 4). To adequately present the arguments in the ensuing chapters, however, it is necessary that we first present and explain the theoretical tenants of Bayesian confirmation theory. This is the task of the present chapter.

1.2 Bayesian Confirmation Theory

Typically—and informally—the term ‘confirmation’ is used whenever observational data and evidence *speaks in favor of* or *supports* some hypothesis or scientific theory. As Carnap (1962) first noted, there are two ways to explicate “evidence E confirms hypothesis H relative to background K ” in the Bayesian framework.

- E confirms H relative to K iff $P(H | E \& K) > k$, for some $k \in (0, 1)$.
- E confirms H relative to K iff $P(H | E \& K) > P(H | K)$.

The first account—“confirmation as firmness”—explicates an *absolute* notion of confirmation, while the second account—“confirmation as increase in firmness”—explicates an *incremental* notion of confirmation. To appreciate the difference between the two notions, notice it’s possible on the incremental account, but not on the absolute account, that E *confirms* H relative to K , and yet the probability of H remains quite *low*, and also that E *disconfirms* H relative to K , and yet the probability of H remain quite *high*. For example, letting $P(H) = .0000001$ and $P(H | E) = .01$ gives a case where the former scenario obtains, while letting $P(H) = .95$ and $P(H | E) = .94$ yields a case where the latter scenario obtains.¹ In this dissertation, I will only be concerned with the incremental account of confirmation.²

On the Bayesian approach, then, confirmation consists in *positive probabilistic relevance*, while disconfirmation consists in *negative probabilistic relevance*. In particular, we have the following definitions.

- E *confirms* H relative to K iff $P(H | E \& K) > P(H | K)$.
- E *disconfirms* H relative to K iff $P(H | E \& K) < P(H | K)$.
- E is *confirmationally neutral* with respect to H relative to K iff $P(H | E \& K) = P(H | K)$.

I use the notation $P(H | K)$ to denote the *conditional probability* of H given K . Throughout the dissertation, I will adopt the following definition of conditional probability.

Definition. $P(A | B) = \frac{P(A \& B)}{P(B)}$, provided $P(B) > 0$.

¹I borrow these numerical examples from Sober (1994).

²Nowadays Bayesians are almost exclusively interested in incremental confirmation. See, e.g., Earman (1992) and Fitelson (2001).

An alternative approach would be to take conditional probability as *primitive* (Popper 1959, Renyi 1955), but since nothing said throughout this dissertation depends on how we axiomatize probability, I will stick with the more traditional ratio definition.

Let W be a non-empty set of outcomes or possibilities, and let \mathbf{A} be a field over W . With $P : \mathbf{A} \rightarrow \mathcal{R}$, we say P is a real-valued, confirmation-theory suitable conditional probability function satisfying the following three constraints.

Non-Negativity: For all A , $P(A) \geq 0$.

Normalization: For any tautology A , $P(A) = 1$.

Additivity: For any logically contrary propositions A and B

(i.e. $A \cap B = \emptyset$), $P(A \cup B) = P(A) + P(B)$.³

The function P can be given many different interpretations. In the Bayesian literature, it's usually given either a "subjective" or "logical" gloss. In this dissertation, I will not take a firm stand on this interpretive issue. My only assumptions about P are that (i) it has the necessary sort of objectivity built into it to ground a robust (and normative) theory of confirmation, and that (ii) it satisfies an appropriate version of the Kolmogorov (1956) axioms.⁴ In chapter 4, I will discuss this issue a

³In addition, if the underlying algebra is taken to be a σ -field—a field that is closed under complementation and countable unions—then P satisfies the following:

- If $A_1, A_2 \dots$ is a countable sequence of events and $A_i \cap A_j = \emptyset$ ($i \neq j$), then $P(\bigcup_{j=1}^{\infty} A_j) = \sum_{j=1}^{\infty} P(A_j)$

I use *sets* in presenting this axiomatization. Propositions and sentences also have the right structure to be in the domain, but following (mathematical) orthodoxy, I will stick with sets. See Weisberg (forthcoming) for further discussion on the sets/sentences/propositions dispute concerning axiomatization.

⁴With respect to assumption (i), I want to point out that this does *not* commit me to a purely formal, syntactic criteria for the assignment of probabilities. I think the formal, syntactic program—in, e.g., Carnap (1950, 1962)—is misguided, but that doesn't imply that the only other gloss of P must be purely subjective. For arguments in favor of an objective P , see Maher (1996, 2004b), Hawthorne (2005, forthcoming), Skyrms (1986), and J. Williamson (2006).

bit further.

I have so far discussed the *qualitative* account of Bayesian confirmation. But the qualitative account can be generalized in various *quantitative* ways. In doing this, Bayesians aim to quantify or measure the *inductive strength* or *degree* to which E supports H relative to K .⁵ Quantitative Bayesian confirmation theory makes use of various *relevance measures* c of evidential support to measure degree of incremental confirmation. More specifically, $c(\cdot, \cdot | \cdot)$ is said to be a relevance measure of degree of confirmation iff the following conditions are satisfied.

$$c(H, E | K) \begin{cases} > 0 & P(H | E \& K) > P(H | K) \\ = 0 & P(H | E \& K) = P(H | K) \\ < 0 & P(H | E \& K) < P(H | K) \end{cases}$$

Many Bayesian relevance measures have been proposed and defended in the literature on confirmation.⁶ In the ensuing chapters, I will restrict my attention to the four currently most prominent and most defended relevance measures in literature: the *difference* measure d , the *log-ratio* measure r , the *log-likelihood ratio* measure l , and the *normalized difference* measure s .⁷ These four measures are defined as follows.

$$d(H, E | K) =_{df} P(H | E \& K) - P(H | K).$$

$$r(H, E | K) =_{df} \log \left(\frac{P(H | E \& K)}{P(H | K)} \right).$$

⁵See Christensen (1999) for useful discussion.

⁶See Fitelson (1999, 2001) for a survey of the various measures.

⁷Earman (1992) and Jeffery (1992) are advocates of measure d . Milne (1996) and Schlesinger (1995) are advocates of measure r . Good (1984) and Fitelson (1999, 2000, 2001) are advocates of measure l . And Joyce (1999) and Christensen (1999) are recent advocates of measure s .

$$l(H, E | K) =_{df} \log \left(\frac{P(E | H \& K)}{P(E | \sim H \& K)} \right).$$

$$\begin{aligned} s(H, E | K) &=_{df} P(H | E \& K) - P(H | \sim E \& K) \\ &= \frac{1}{P(\sim E | K)} \times d(H, E | K). \end{aligned}$$

As Fitelson (1999, 2001) has discussed in detail, a number of arguments in Bayesian confirmation theory are *sensitive to choice of measure*—meaning the argument’s validity depends on which relevant measure is assumed. Consider, for example, the so-called “tacking problem” or “problem of irrelevant conjunction” for Bayesian confirmation theory. This problem shows that Bayesian confirmation has the apparently problematic property that if E confirms H , and H entails E , then E confirms $H \& X$, for *any* X , regardless of how *irrelevant* X may be to H .⁸ Several Bayesian resolutions to this problem have been proposed.⁹ One resolution, advanced by Rosenkrantz (1999), relies crucially on the following property of the difference measure d :

$$(+) \text{ If } H \models E, \text{ then } d(H \& X \& E | K) = P(X | H \& K) \times d(H, E | K).$$

Fitelson (1999) proves, however, that neither measure r nor measure l satisfy (+). So Rosenkrantz’s solution to the tacking problem is sensitive to choice of measure.

In chapter 2, the issue of measure sensitivity will arise, since one of the theorems I prove there hold for measures d , r , l , but not s . However, I don’t have much

⁸The proof of this is pretty simple. Assuming—and here I’m dropping K — $P(E) < 1$, we have, when $H \models E$, $P(H | E \& X) = 1$, for any X . So,

$$P(H \& X | E) = P(E | H \& X) \times P(H \& X) / P(E) > P(H \& X).$$

The assumption that $P(E) < 1$ is considered harmless, since if $P(E) = 1$ evidence E would be unable to confirm anything, as Glymour (1980) pointed out.

⁹See, e.g., Fitelson (2002), Hawthorne and Fitelson (2004), Maher (2004a), and Chandler (2007).

to say about issues concerning measure sensitivity in this dissertation, since it is not our main topic. Moreover, I believe measure l is the most adequate Bayesian measure of confirmation, so I'll be concerned first and foremost with pursuing our discussion in terms of l and ensuring that our results hold for this measure.¹⁰ Along the way, though, I'll indicate when I believe our formal results are relevant to the debate over the adequacy of the Bayesian measures d , r , l , s .

Two final points. First, it is important to bear in mind that while confirmational relevance is defined in terms of probability, it itself is *not* a probability. This can be seen by noting that when E confirms H relative to K , and $c = l$, we have $l(H, E | K) > 1$.¹¹ Thus, confirmational relevance can sometimes take values greater than 1, something probability *can't* do, since P is a function onto the closed interval $[0, 1]$.¹² Second, it is worth emphasizing that Bayesian confirmation is a *three-placed* relation, between evidence E , hypothesis H , and background K . For Bayesians, the information contained in the background corpus K can significantly affect E 's evidential bearing on H . This fact will play an important role in the ensuing chapters.

Having presented the Bayesian confirmation-theoretic framework, I now demonstrate with the following case study how taking confirmation instead of probability to be the central formal notion in epistemology can have important consequences for current epistemological debates.

¹⁰See Fitelson (2000, 2001) and Fitelson and Eells (2002) for numerous considerations in support of measure l .

¹¹For example, suppose $P(E | H) = .95$, and $P(E | \sim H) = .05$. Then

$$c(H, E) = \frac{P(E | H)}{P(E | \sim H)} = \frac{.95}{.05} = 19.$$

¹²Of course, confirmational relevance can also take *negative* values, again demonstrating that (dis)confirmation is not a probability.

1.3 Case Study: Dogmatism and Bayesianism

Recently several authors have argued that the dogmatic theory of perceptual justification is incompatible with Bayesianism, and many have taken this incompatibility to motivate a rejection of dogmatism. I show that the argument against dogmatism can be resisted by focusing on confirmation instead of probability.

G. E. Moore (1962) famously maintained that he could know various propositions—e.g., that he has hands—without being able to offer any sort of proof of those propositions. *Dogmatism* is an epistemology of perception based on this (anti-skeptical) Moorean idea.

Suppose Jessie looks across the quad and sees a tree. She forms the belief that there's a tree across the quad. Suppose she doesn't have antecedent justification to think that she's not being deceived by an evil demon. Is Jessie's belief that there's a tree across the quad justified? According to dogmatism, the answer is *yes*: “when it appears it perceptually seems to you as if p is the case, you have a kind of justification for believing p that does not presuppose or rest on your justification for anything else, which could be cited in an argument (even an ampliative argument) for p ” (Pryor 2000: p. 519). “The dogmatist thinks”, says Pryor, “that the *mere having* of an experience as of p is enough for your perceptual justification for believing p to be in place” (*ibid*). Following White (2006), we can state the dogmatist position as follows.

Dogmatism: For certain contents p , if it appears to S as if p , then S is justified in believing that p , regardless of whether S is independently justified in rejecting a skeptical hypothesis like being a handless BIV, or being deceived

by an evil demon, and so on.¹³

One important consequence of this view is that an agent S can come to justifiably believe that p is true by simply seeing that it appears that p is true—even if S lacks antecedent reason to believe that visual impressions are generally reliable or that certain skeptical hypotheses don't obtain. According to the *Bayesian objection* to dogmatism, this consequence of dogmatism is in tension with the core tenants of Bayesianism, and for many that provides motivation to reject dogmatism.¹⁴

To state the Bayesian objection, it'll be useful to have some notation.

H_1 = It appears to me that I have hands.

H_2 = This is a real hand.

H_3 = This is not a fake-hand.

We'll stipulate that fake-hands aren't real hands, but that they're indistinguishable visually from real hands. Notice, too, that H_2 entails H_3 .

Assuming—reasonably enough— H_2 is evidence for H_1 , we're able to derive the following inequalities in the Bayesian framework.

$$P(H_1 | H_2) > P(H_1) \tag{1.1}$$

and

$$P(H_1 | \sim H_3) > P(H_1). \tag{1.2}$$

¹³Views similar to dogmatism have been defended by Alston (1989), Audi (1993), Burge (1993), Chisholm (1989), Peacocke (2004), and Pollock and Cruz (1999).

¹⁴See Weatherson (2007) for further discussion of the Bayesian objection to dogmatism.

By Bayes's theorem¹⁵, it follows that

$$P(H_2 | H_1) > P(H_1) \tag{1.3}$$

and

$$P(\sim H_3 | H_1) > P(\sim H_3). \tag{1.4}$$

Hence,

$$P(H_3 | H_1) < P(H_3). \tag{1.5}$$

Now, in view of the fact that $P(H_3 | H_1) < P(H_3)$, it's a short step to show that my confidence in H_2 should *increase*, while my confidence in H_3 should *decrease*. And it's at this point that dogmatism looks to be incompatible with Bayesianism. To see the problem, we note the following: (i) I can justifiably believe H_2 on the basis of undergoing a perceptual experience of hands, (ii) H_2 entails H_3 , and (iii) the principle of justification closure, which states, roughly, that if S is justified in believing p and p entails q , then S is justified in believing q . So, given (i)-(iii), and that I can *gain* justification H_3 , it follows that my confidence in H_3 should *increase*, which is clearly incompatible with the above Bayesian derivation that my confidence in H_3 should *decrease*. What can the dogmatist say in response?

In presenting this argument, we've been assuming, with White, that if one is justified in believing H_2 , then one is also justified in believing H_3 . Motivating White's assumption here is the following *principle of justification closure*:

¹⁵Bayes's theorem is as follows:

$$P(H | E \& K) = \frac{P(E | H \& K) \times P(H | K)}{P(E | K)}$$

See chapter 4 for further discussion of Bayes's theorem.

Justification Closure: If agent S is justified in believing p , p entails q , and S is aware of this entailment, then S is justified in believing q .

White says this principle is “very hard to deny” (2006: p. 529). For White, not only does the principle appear intuitively plausible, but it also looks to be supported by Bayesianism itself, since it’s a theorem of the probability calculus that if H entails H^* , then $P(H | E) \leq P(H^* | E)$, for *any* H , H^* , and E .¹⁶

So, from a Bayesian perspective, justification closure certainly looks correct, *assuming that we take probability to be the central formal notion*. But—and this is the key point—any discussion of justification closure in the Bayesian framework must keep separate the following two principles.

Probabilistic Entailment (PE): For any H , H^* , and E : If H entails H^* , then $P(H | E) \leq P(H^* | E)$.

Confirmational Entailment (CE): For any H , H^* , and E : If E (incrementally) *confirms* H and H entails H^* , then $c(H, E) \leq c(H^*, E)$, where c is a relevance measure of support.

This distinction is important, because, while (PE) is true by the probability calculus, it is easy to show that (CE) is *false*, confirmation does *not* transmit across logical entailment. Consider, for instance, the following lottery example from Jeffery (1992). Suppose we have a fair eight-ticket lottery. Let E = the winning ticket is either ticket 2 or ticket 3, G = the winner is either ticket 3 or 4, and H = the winner is neither ticket 1 nor ticket 2. In this setup, we have $P(G | E) = .5 > P(G) = .25$, even though $P(H | E) = .5 < P(H) = .75$.¹⁷

¹⁶*Proof.* Assume $A \Rightarrow B$. Then $A \subseteq B$, and so $B = A \cup (\sim A \cap B)$. Hence $P(B) = P(A) + P(\sim A \cap B) \geq P(A)$ [Heathcote (1971)].

¹⁷Hempel (1945) introduced principle (CE), calling it the *special consequence condition*. See

So the upshot of distinguishing between (PE) and (CE), and my claim that confirmation, not probability, is the central formal notion, is that justification closure begins to look far less obvious in the Bayesian framework. By taking a confirmation-theoretic approach to closure, in contrast to White’s probabilistic approach, it is not clear why the dogmatist would want to accept closure. And without closure the Bayesian objection to dogmatism does not go through.

My aim here is not to defend dogmatism, but rather to highlight how the confirmation-theoretic approach can affect the landscape of current debates in modern epistemology. I will be pushing this theme in the rest of the dissertation.

1.4 What’s to Come

Let me close with a brief overview of the rest of the dissertation.

In the next chapter, I will turn to the question of whether or not an epistemic closure principle is correct. The dominant attitude among epistemologists nowadays is that closure is obviously correct, in need of little defense or argument. One reason for this—though, admittedly, this is a bit speculative—is that it is being implicitly assumed that probability is the central formal notion, and since probability is closed under entailment, so too is knowledge and justification. By shifting our focus from probability to confirmation, and exploiting the different formal properties exhibited by functions P and c , we are able to develop a novel case against closure. Here I also prove some new formal results that bear on the debate over the adequacy of Bayesian measures of confirmation.

Salmon (1975) for excellent discussion of (CE). (PE) and (CE) provide just one example where P and c exhibit different properties; for further examples and discussion, see Crupi, Fitelson, and Tentori (forthcoming).

Chapter 3 discusses a particular form of voting—*supermajority voting*—from a Bayesian epistemological perspective. The standard route taken to motivate supermajority voting is via the Condorcet framework and its well-known “jury theorem”. I show that a Bayesian confirmation-theoretic approach provides a much more general and powerful approach to supermajority voting. The basic idea is that supermajority voting provides *superior evidence* to simple majority voting, which, from an epistemological perspective, makes it preferable to simple majority voting. This chapter also presents some new technical results that improve upon the contemporary Bayesian literature on evidential support from multiple pieces of evidence.

Chapter 4 is an essay in Bayesian jurisprudence and legal epistemology. In Anglo-American criminal jurisprudence, one important and frequently used standard of criminal proof is *proof beyond reasonable doubt*. According to this standard, the jurors begin trial with a “presumption of innocence” and then it’s the prosecution’s task to present a case for the defendant’s guilt to the jury in “such convincing character that a reasonable person would not hesitate to rely and act upon it in the most important of his own affairs.” (Devitt 2003: sec. 12.10). Recent work on defining the notion of proof beyond reasonable doubt, however, is rather vague and often confusing. This chapter proposes a new and precise account of proof beyond reasonable doubt. My approach will make central use of Bayesian confirmation theory, and the idea will be that proof beyond reasonable doubt is established when a certain level of confirmatory support is reached. I will also discuss the “presumption of innocence” doctrine, since it plays an important role in my account of proof beyond reasonable doubt. My discussion will make some important philosophical and mathematical advances on our understanding of both these doctrines. All proofs are given in the Appendix.

CHAPTER 2

A BAYESIAN APPROACH TO CLOSURE FAILURE

2.1 Introduction

This chapter presents a new and mathematically well-defined approach to closure failure. The account offered here is partly inspired by Fred Dretske's work on closure. In particular, I show how some of his insights can be formalized in the Bayesian confirmation-theoretic framework and used to develop a novel route to the failure of closure. My overall goal is to show that Bayesian confirmation theory has much to offer modern epistemology. Along the way, I prove theorems that concern the plurality of Bayesian confirmation measures and the debate over the adequacy of the measures.

2.2 Epistemic Closure

Very roughly, epistemic closure states that if agent S knows that p and knows that p entails q , then S knows that q . This formulation is a bit problematic, though, since it's possible that S fails to believe q , while knowing that p and knowing that p entails q . Following Williamson (2000) and Hawthorne (2004), I will adopt the following formulation of closure for the rest of this chapter.

Epistemic Closure (C): If S knows that p , competently deduces q from p , and thereby comes to believe q , then S knows that q .

Note the closure principle, as stated, is meant to be a *dynamic* principle about the expansion of knowledge by competent deduction; it's not meant to be a *static* principle about what states of knowledge are possible.

Closure plays a central and fundamental role in contemporary epistemology. One reason for this is that it's a key principle in the so-called *problem of easy knowledge* (Cohen 2002). Another reason is that it's the principle underlying the second premise, P2, in the following prominent skeptical argument (here '*p*' is some ordinary proposition, '*sh*' is some skeptical hypothesis, and '*KSp*' stands for *S* knows that *p*).¹

P1: $\sim KS \sim sh$

P2: $KSp \Rightarrow KS \sim sh$

C: $\sim KSp$

Given closure's prominence in contemporary epistemology, it is natural to ask: Is the principle true? It is widely believed that the principle is (obviously) correct. Witness Feldman (1995):

I believe some version of the closure principle is...surely true. Indeed, the idea that no version of [the] principle is true strikes me, and many other philosophers, as one of the least plausible ideas to come down the philosophical pike in recent years (p. 487).

Despite its apparent obviousness, closure looks susceptible to counterexample. Here's a classic example from Dretske (1970). Suppose I'm at the zoo looking at

¹DeRose (1995), Vogel (1999), Pritchard (2002, 2005a), Klein (1981, 2004), Brueckner (1994), Sosa (1999), Silins (forthcoming) among others advance the skeptical argument along these lines. Marian and Warfield (forthcoming) present a number of objections to the closure-based formulation of the skeptical argument.

a zebra cage. It seems that I know that there's a zebra in the cage, but it seems that I do not there's not a cleverly disguised mule in the cage, even though the former proposition entails the latter.

This chapter attempts to make some headway on whether or not closure is true. Using Bayesian machinery, I argue that the principle is false. The account of closure failure offered here is partly inspired by Dretske's (1970, 2005, 2006) work on closure. In particular, I show how some of his insights can be formalized in the Bayesian framework and used to develop a novel route to the failure of closure.²

In the next section, I present some Dretskean intuitions on closure. Section 2.4 offers a few theorems to motivate the intuitions in the Bayesian framework. Section 2.5 puts the intuitions to work in my Bayesian approach to closure failure. In section 2.6, I discuss some advantages of the Bayesian account. Section 2.7 sums up.

²While Dretske's arguments against closure aren't exhaustive, many of the other accounts of closure failure rest on very particular accounts of knowledge, or evidence, and so they tend to lack the generality of Dretske's accounts, especially his most recent account. Nevertheless, here is a brief overview of some important work on closure failure. Nozick (1981) famously argued that closure is false. His core idea was that one might *track* and hence know some ordinary proposition p (e.g., that one has hands), while fail to track and hence not know the negation of some skeptical hypothesis (e.g., that one is not a BIV). I connect my discussion up with Nozick in section 2.6. McGinn (1984) argues that closure fails because he holds that a necessary condition to know p is that one be able to *tell whether* p . We know we have hands (in part because we can tell whether we have hands) and we know having hands entails \sim BIV. But we can't tell whether or not we are in BIVs, since things would look and feel exactly the same in them. So we don't know \sim BIV. Christensen (2004) argues Makinson's (1965) *preface paradox* provides reason to think that beliefs aren't closed under entailment; for a similar style argument against closure, see Olin (2006). (Though see Weatherson (2005) and Maher (2006b) for a trenchant criticism of Christensen.) Maitzen (1998) argues the *Knower paradox* makes serious trouble for closure. For more arguments against closure, see Goldman (1986), Kvanvig (2006, forthcoming), Frances (1999), Lawor (2005), and Baumann (2006).

2.3 Some Dretskean Intuitions

In Dretske’s work on closure failure, he articulates several important intuitions regarding epistemic closure, specifically ones concerning relations of evidential support. This section discusses three of Dretske’s intuitions and shows how they can be formalized in the Bayesian confirmation-theoretic framework.³

In his essay “Epistemic Operators”, Dretske provides a counterexample to closure. The counterexample, mentioned above, claims that when I’m looking at a zebra cage, under normal lighting conditions, I know there’s a zebra in the cage, but I don’t know there’s not a cleverly disguised mule in the cage. After presenting this counterexample, Dretske writes the following important passage.

If you are tempted to say [that the agent does know it’s not cleverly disguised mule in front of her], think for a moment about the reasons that you have, what evidence you can produce in favour of this claim. The evidence you *had* for thinking them zebras has been effectively neutralized, since it does not count toward their *not* being mules cleverly disguised. Have you checked with the zoo authorities? Did you examine the animals closely enough to detect such a fraud? (1970: p. 1016)

There’s a lot going on in this passage, but two crucial intuitions about evidential support are being expressed here. First, there’s the intuition that one’s perceptual evidence provides support for there being a zebra in the cage. Second, there’s

³The reconstruction of Dretske I offer here reflects my understanding of Dretske. Since Dretske’s work on closure leaves a lot of room for interpretation, I don’t intend to defend a specific reading of Dretske. So, even if what I say below isn’t what Dretske actually had in mind, it’s still interesting in its own right and worthy of discussion.

the intuition that the perceptual evidence doesn't lend much support to there not being a cleverly disguised mule in the cage. This second intuition is a bit unclear, and there are probably several ways to understand it. One interpretation is that the perceptual evidence provides *no* evidential support for there not being a cleverly disguised mule. I do not endorse this interpretation. Rather, I think the proper—and more charitable—interpretation is that the perceptual evidence provides *less* support for there not being a cleverly disguised mule in the cage than it does for there being a zebra in the cage.

I will take these intuitions on board, but before I formalize them, let's introduce some terminology for propositions like *there are mind-independent material objects* and *the universe isn't 5 minutes old*. Following Dretske (2005, 2006), I will call these *heavyweight propositions*. As I understand Dretske, heavyweight propositions are supposed to form a class of not easily knowable propositions, and yet they are entailed by propositions we readily take ourselves to know, e.g., *there are cookies in the jar* and *I had breakfast yesterday*.⁴

For much of this paper we're going to be concerned with the following ordinary proposition *O* and heavyweight *H*.⁵

⁴Unfortunately, Dretske doesn't offer a rigorous definition of heavyweight propositions. Hawthorne (2005) offers some helpful guidance: "Let P be a heavyweight proposition just in case we all have some strong inclination to think that P is not the sort of thing that one can know by the exercise of reason alone and also that P is not the sort of thing that one can know by use of one's perceptual faculties (even aided by reason)" (p. 33). Luper (2006) mentions another feature common to heavyweight propositions: "... many of Dretske's [heavyweight] propositions refer to situations with a feature we might call *elusiveness*. A situation *sk* is elusive to me when the following is true: were *sk* not to hold, I would still have the experiences I have now" (p. 385; emphasis in original). I don't know how much of this Dretske would accept, but perhaps it gives us a better idea of what heavyweight propositions are supposed to be. It's also worth noting that a lot of what Dretske says about heavyweight propositions is rather similar to what Wittgenstein says in *On Certainty* about *hinge propositions*. According to Wittgenstein, hinge propositions "lie apart from the route travelled by enquiry" (OC §88). "There are" he says, "historical investigations and investigations into the shape and also the age of the earth, but not into whether the earth has existed during the last hundred years" (OC §138). See Pritchard (2005b) for further discussion on Wittgenstein and hinge propositions.

⁵Nothing hangs on this selection of *O* and *H*; they're used simply for convenience.

- O = There is a 20-year-old tree slice before me.
- H = The universe isn't 5 minutes old.

Our evidence here will be a typical perceptual experience, represented by E .

- E = There appears to be a 20-ringed tree slice before me.

Using this setup, we can unpack Dretske's intuitions, expressed in "Epistemic Operators", as follows. First, there were the intuitions that E supported O and H respectively to some degree. In our Bayesian framework, this is formalized as follows.

$$(T1) \ c(O, E \mid K) > 0 \text{ and } c(H, E \mid K) > 0.$$

Despite the fact that E supports both O and H , Dretske claimed that E supported H *less than* it supported O . This intuition is expressed with (T2).

$$(T2) \ c(H, E \mid K) < c(O, E \mid K).$$

In his later work on closure, Dretske (2005, 2006) considers the possibility that (T2) is false, by considering scenarios where E , conjoined with some additional evidence E' , supported H to an equal or greater degree than E 's support for H , i.e. scenarios where $c(H, E \ \& \ E' \mid K) > c(H, E \mid K)$. But, according to Dretske, such scenarios are *not* possible, since there are no other "accredited ways of knowing" that could deliver evidence E' such that it, when conjoined with E , would support H more than E already does.⁶ For Dretske, accredited ways of knowing include touching, tasting, smelling, memory. These ways of knowing, he writes,

⁶Here's the relevant passage:

“do not supply us with all (some, but not all) the information implied by what they inform us about—that, for instance, there is a real physical world distinct from our impressions of it (2006: p. 411). In support of the idea, Dretske offers the following example:

a voltmeter’s pointer pointing at the numeral ‘5’ provides information that the voltage is 5 without providing information—logically implied by the information it does provide—that the meter isn’t *misrepresenting* a voltage of 4 as 5 (ibid; emphasis in original).

So, Dretske concludes, “we cannot see (hear, smell, or feel) that [heavyweight propositions] are true” (2005: p. 20).

The basic idea, then, is that E ’s support for H can’t be strengthened by conjoining it with additional evidence E' , even when E' is from an accredited way of knowing. In our Bayesian confirmation-theoretic framework, this intuition gets formalized as follows.

$$(T3) c(H, E | K) = c(H, E \& E' | K).$$

Here $E' = E_1 \& E_2 \& \dots \& E_n$, with each E_i being the evidence recieved from an accredited source of knowledge. The formula in (T3) says that the support H receives from E is the same as the support H receives from $E \& E'$.

...*none* of our accredited ways of knowing about our material world are capable of telling us that there is a material world, none of the accredited ways of finding out what people feel and think are ways of finding out that they are not mindless zombies, and none of the accepted ways of finding out what, specifically, happened yesterday are ways of finding out that there was a yesterday (2005: pp. 22-23; emphasis in original).

I believe (T1)-(T3) capture Dretske's intuitions about relations of evidential support concerning typical perceptual evidence and ordinary and heavyweight propositions. The intuitions expressed in (T1) and (T2) strike me as fairly innocuous, and they will be taken on board in our Bayesian framework, though I will have a bit more to say about them below.⁷ The less obvious intuition here, I think, is the one expressed in (T3). Since (T3) plays a role in my Bayesian theory of closure failure, the next section will be spent (formally) substantiating it in our Bayesian framework. In section 2.5 I will put claims (T1)-(T3) to work in my Bayesian theory of closure failure.

2.4 Some Theorems

In this section, I will state four formal results that bear on (T3). I will continue to use propositions O , H , and E from above. As I will discuss below, some of the theorems below bear on the problem of measure sensitivity and are relevant to the debate over the adequacy of our four relevance measures d , r , l , and s .

(T3), recall, expressed the idea that the support H received from E was equal to the support H received from conjoining E with additional evidence E' . A fairly simple condition that will be both necessary and sufficient to prove (T3) is the following:

$$P(H \mid E \ \& \ K) = P(H \mid E \ \& \ E' \ \& \ K).$$

Informally, this condition says that what E says about H is the same as what the conjunction $E \ \& \ E'$ says about H . The condition explicates nicely the core idea behind (T3), that the boost provided by $E \ \& \ E'$ is the same as the boost provided

⁷In my (forthcoming), I discuss the intuition expressed in (T2) in more detail, and I state conditions that suffice to prove $c(O, E \mid K) > c(H, E \mid K)$.

by E . And while this condition might seem like the obvious choice to prove (T3), it turns out that it will only allow us to derive $c(H, E, K) = c(H, E \& E', K)$ for c -measures d, r, l , but, surprisingly, *not* for measure s , as the following theorem shows.

Theorem 1. *If (i) $P(H | E \& K) = P(H | E \& E' \& K)$, then*

$$c(H, E | K) = c(H, E \& E' | K)$$

where c may be either the difference measure d , the ratio measure r , or the likelihood ratio measure l , but c may not be the normalized difference measure s .

Since measures d, r, l satisfy this theorem, but s does not, there's some measure sensitivity here, and for some, particularly proponents of s , this might look problematic. One way to bypass this worry is to find conditions that will suffice to entail $c(H, E | K) = c(H, E \& E' | K)$ for all four of our measures d, r, l , and s . I will now present two such conditions. Then I will explain why I prefer some measure sensitivity in proving (T3), via Theorem 1, over the following insensitive route.⁸

Consider the following two conditions:

$$P(E | H \& K) = P(E \& E' | H \& K)$$

and

$$P(E | K) = P(E \& E' | K).$$

Informally, the first condition, $P(E | H \& K) = P(E \& E' | H \& K)$, says that what H says about E is the same as what H says about $E \& E'$, and the second condition, $P(E | K) = P(E \& E' | K)$, says that E and $E \& E'$ have the same prior probability. I observe that:

⁸For this dissertation, it's important to notice that measure l satisfies both theorems.

Theorem 2. If $(i^*) P(E | H \& K) = P(E \& E' | H \& K)$, and $(ii) P(E | K) = P(E \& E' | K)$, then

$$c(H, E | K) = c(H, E \& E' | K)$$

where c may be any measure among d , r , l , or s .

In view of Theorem 2, we can prove our desired equality, $c(H, E | K) = c(H, E \& E' | K)$, without any measure sensitivity, bypassing any sensitivity-based worries one might have with Theorem 1. However, there are at least *three* very compelling reasons to stick with condition (i) , $P(H | E \& K) = P(H | E \& E' \& K)$, over conditions (i^*) and (ii) in proving $c(H, E | K) = c(H, E \& E' | K)$.

The first reason is that conditions (i^*) and (ii) , $P(E | H \& K) = P(E \& E' | H \& K)$ and $P(E | K) = P(E \& E' | K)$, are *extremely* strong. They entail all kinds of bizarre stuff. For example, they entail that

$$P(E \& \sim E' \& H | K) = 0$$

and

$$P(E \& \sim E' \& \sim H | K) = 0$$

which in turn entail

$$P(E \& \sim E' | K) = 0.$$

In essence, this is like assuming that E and E' are *mutually exclusive*. Is this a plausible assumption in our context? I think not.

The second reason is that it can be shown, by the following theorem, that conditions (i^*) and (ii) aren't *necessary* for $c(H, E, K) = c(H, E \& E', K)$, since (i^*) itself isn't necessary.

Theorem 3. *There exists a probability model in which both: (a) $c(H, E | K) = c(H, E \& E' | K)$ and (b) $P(E | H \& K) \neq P(E \& E' | H \& K)$, for all four of our c-measures d , r , l , and s .*

I offer this theorem as a motivation for sticking with condition (i), $P(H | E \& K) = P(H | E \& E' \& K)$, in proving (T3), since it *is* necessary for (T3).⁹

The third and final reason is that there's a convincing formal rationale to not endorse $P(H | E \& K) = P(H | E \& E' \& K)$ and $s(H, E | K) = s(H, E \& E' | K)$ simultaneously. This rationale is given by the following theorem.

Theorem 4. *For $0 < P(H | K) < 1$ and $0 < P(E \& E' | K) < 1$, any pair of the following three clauses implies the remaining clause:*

1. $s(H, E | K) = s(H, E \& E' | K)$.
2. $P(H | E \& K) = P(H | E \& E' \& K)$.
3. $P(E' | E \& K) = 1$.

Now, I claim that requiring $P(E' | E \& K) = 1$ is very implausible. So, in light of this theorem, one cannot plausibly maintain both $P(H | E \& K) = P(H | E \& E' \& K)$ and $s(H, E | K) = s(H, E \& E' | K)$.

I think these three considerations basically settle the formal issues here and strengthen the case for proving (T3), via condition (i), $P(H | E \& K) = P(H | E \& E' \& K)$. As such, I think a much better approach is to adopt measures d , r , or l , and to go with the necessary and sufficient condition (i) in proving (T3).

⁹This can easily be verified by seeing what's involved in the proof of Theorem 1 given in the Appendix.

At this point, one might object to (T3) along the following lines. Suppose we say our additional evidence E' is just the ordinary proposition O , that there is a 20-year-old tree slice before me (i.e., $E' = O$). Then it would seem that, *pace* (T3), $c(H, E) \neq c(H, E \& E')$, since $c(H, E) < c(H, E \& E') = c(H, E \& O)$ appears true. In response, I would say that O itself is *not* evidence, but rather something that we acquire evidence for and believe on the basis of such evidence. Indeed, we have plenty of evidence for ordinary propositions like O , but that evidence doesn't support heavyweight propositions like H . After all, if we did allow things like O to be in our evidential base, it would look pretty question-begging against the skeptic.

Finally, I won't attempt to settle the debate over the adequacy of d , r , l , s here, but the properties of the measures described in the theorems have not yet been considered in the literature on Bayesian confirmation theory. Moreover, I believe that some of these newly discovered properties can be used to advance the debate over the adequacy of the measures. In particular, it strikes me as very odd that if $c = s$ it's possible for $P(H | E) = P(H | E \& E')$ and yet $c(H, E) \neq c(H, E \& E')$. How, exactly, could that be? Perhaps s is measuring something important, but it doesn't look to be measuring *confirmation*. In any case, I won't press this point against s here, but I do think that the burden is on proponents of the measure to say something useful here.

2.5 Closure Failure, Bayesian Style

This section presents a theory of closure failure in the Bayesian confirmation-theoretic framework. My overall strategy to derive closure failure is to add a pretty weak evidentialist component to the analysis of justified belief and knowledge and

argue, using claims (T1)-(T3) above, that one fails to satisfy this evidentialist component with respect to heavyweight propositions. The evidentialist constraint I propose fits very naturally in the Bayesian framework.

Here are two claims one often finds in the contemporary analysis of justification and knowledge (see, e.g., Steup (2006)). First, S is justified in believing Z only if S has evidence for Z . And, second, S knows that Z only if S is justified in believing Z . The first claim about justification amounts to a broadly “evidentialist” constraint on justification,¹⁰ and the second constraint on knowledge is given to rule out knowledge by luck.

I will focus on the evidentialist condition and attempt to recast it in the Bayesian framework. A natural way to state this condition in the Bayesian framework would be talk of S having *confirming evidence* for Z . Roughly, the claim would be that S is justified in believing Z only if S has confirming evidence for Z .¹¹

Less roughly, I propose sharpening this confirming evidence condition on justified belief with a threshold conception of confirming evidence. Specifically, I will say that E is (incremental threshold) confirming evidence for H relative to K iff $c(H, E, K) > \tau$, for some threshold τ appropriate to the particular measure c . I adopt the following two theses in our Bayesian framework.

- (1) S is justified in believing Z only if S has confirming evidence for Z , where ‘confirming evidence’ is understood according to the threshold conception mentioned above.

¹⁰For more on evidentialist theories of justification, see Haack (1993), Conee and Feldman (2004), and Adler (2002).

¹¹I take no stand on what such ‘having’ of evidence amounts. My sympathy is with an *internalist* gloss, but that won’t matter for what follows.

(2) S knows that Z only if S is justified in believing Z .

(1) and (2) both seem pretty plausible to me. At the very least, I think they can be incorporated into an intuitive Bayesian epistemology. Now with claims (1) and (2), along with (T1)-(T3) from above, we are in position to present a Bayesian theory of closure failure.¹² We reason as follows.

According to (T1), E supports O and H respectively, but by (T2), the support E provides for H is less than the support E provides for O . In view of this, the following two assumptions seem pretty reasonable: (a) E supports O to a pretty high degree, enough to cross the threshold, while (b) E 's support for H is quite small, not enough to cross the threshold. Consequently, one has confirming evidence for the ordinary O . And, intuitively—unless we are skeptics—one can know O . Now, by (T3), there is no other additional E' such that it, conjoined with E , would support H over the requisite threshold. From our first thesis about justification, (1), it follows that one isn't justified in believing H . So, by (2), it follows that one doesn't know H , despite the fact that O entails H . Putting this all together, the following three claims are jointly possible in our Bayesian framework:

- E supports O at or above the requisite threshold τ .
- One justifiably believes and knows O .
- One fails to justifiably believe and know that H .

Where c is measure l , for example, it's clearly possible to have $\log[P(E | O \& K)/P(E | \sim O \& K)] > \tau$, while $\log[P(E | H \& K)/P(E | \sim H \& K)] < \tau$, even when

¹²Again, O , H , and E are as above.

$H \& K \models O$.¹³ Put simply, then, these three claims show that closure can fail in our Bayesian framework.

I have left some assumptions here undefended. One is that E supports O over the requisite threshold τ , while E doesn't support H over τ . Another is that it's possible for one to satisfy the other conditions—whatever they are—needed to know the ordinary proposition O .¹⁴ I think these assumptions are pretty innocuous, especially the latter. To those who might have some worries here, note that it certainly looks *possible* for these assumptions to be true. And that's enough to undermine closure, since contemporary epistemologists hold that closure is a *necessary* truth.

I do not offer this theory as a knockdown case against closure. Instead, I have tried to present a new (Bayesian) challenge to the principle, and in doing so, I have shown that closure can no longer be regarded as an obviously correct principle, in need of little or no defense.

2.6 Motivation

Until now the discussion has been fairly abstract. What, if anything, is the upshot of switching to a confirmation-theoretic framework when approaching the question of knowledge closure? This section suggests some answers. My focus here will be on (i) the account's ability to handle prominent counterexamples to various open theories of knowledge, and (ii) some of the important epistemological consequences that result from adopting an anti-closure stance. Throughout the ensuing discus-

¹³In the Appendix, I present some probability models in which these properties obtain.

¹⁴I'm also assuming that skepticism is false, but that's an assumption I won't argue for here, or anywhere else, in this dissertation.

sion, I will exploit the Bayesian claim that confirmation is a three-placed relation between evidence, a hypothesis, and background information.

2.6.1 Handling Counterexamples

If knowledge isn't closed in the Bayesian framework, it's natural to ask: Does the Bayesian account have advantages over the other prominent anti-closure theories of knowledge, such as the theories in Dretske (1970, 1971) and Nozick's (1981). I believe the answer is *yes*. I will try to argue for this by showing that the account can handle some prominent counterexamples to Dretske and Nozick's theories. I want to stress that the results in what follows are tentative, because the counterexamples I take up aren't exhaustive, and there is some question whether the examples really do succeed in refuting either Dretske or Nozick's theory.¹⁵ Instead, the goal is to give one a sense of the sort of work my account can do and show how adding a confirming evidence component on the analysis of knowledge might provide a more satisfactory route to an open theory of knowledge.

The first counterexample is from McGinn (1984). Suppose we live in a world—call it γ —full of material objects, which we perceive and interact with daily. In γ , there's a benevolent deity who watches over our sensory input in this sense: he has the intention to preserve our sensory input of material objects by artificial means in the event of a cataclysm in which the material objects that actually produce our sensations should suddenly go out of existence. The cataclysm is physically possible and the deity has the power to carry out his intention. The following counterfactual seems true: “If the objects around me were to go out of

¹⁵See Adams and Clarke (2005) and Adams (2006) for a defense of Dretske and Nozick from some of the counterexamples discussed here.

existence, then I would still believe I was surrounded by these material objects”. So on Nozick’s account I don’t know that there are material objects in γ . That seems wrong.

My account: the existence of material objects in the universe receives confirming evidence from visual perception and the other senses. The existence of the deity and his intentions doesn’t undermine the confirming evidence for the material objects. So on my account there isn’t any obvious reason to think that one doesn’t know there’s material objects in γ .¹⁶ However, if one had the background information that there was a deity in γ with his intentions and that a cataclysm had taken place in γ , then one wouldn’t have confirming evidence for there being material objects in γ .¹⁷ Of course when I say my account allows one to know that there are material objects in γ , I don’t (and can’t) mean that one knows the (heavyweight) proposition ‘there are material objects in γ ’. Rather, I mean that my account allows one to know that there are particular material objects—tables, chairs, books—in γ , which is entirely compatible with everything that has been said so far.

The second counterexample is from Kripke (early 1980’s). Nick is looking at a red barn in Barn Country. In Barn Country, red barns can’t be faked, though other colored barns can and often are faked. There’s good lighting in Barn Country and Nick’s perceptual system is functioning properly. On Nozick’s theory, Nick knows there’s a red barn in front of him, but he doesn’t know there’s a barn in front of

¹⁶I hedge here because I have confirming evidence a necessary but insufficient condition, which means I am only entitled to say when one lacks knowledge, not when one has knowledge. So the claim is a conditional: if the other necessary and jointly sufficient conditions are satisfied, then we have a case of knowledge. Similar remarks apply to my responses to the other counterexamples.

¹⁷Suppose one didn’t have the background information about the deity or that a cataclysm had happened. Would one have confirming evidence for there being material objects after the cataclysm? Yes. But that doesn’t mean one would know that there are material objects in γ , since post-cataclysm there aren’t any material objects in γ .

him, because if there weren't a barn in front of him, he would believe there was.

My account: seeing a red barn in good lighting and with a properly functioning visual system provides confirming evidence for there being a red barn in front of Nick and also provides confirming evidence for there being a barn in front of Nick. So on my account there isn't any obvious reason to think that Nick doesn't know there's a barn in front of him.¹⁸ However, if Nick's background information included the information that he was in Barn Country (assuming this implies knowing red barns can't be faked), then he'd still have confirming evidence for there being a red barn in front of him. Of course, things would be different if he had this background information and was looking at a green barn.

The third counterexample is from Pappas and Swain (1973). George is having a visual experience of a cup on a table in circumstance C . The lighting is good in C and George's perceptual system is functioning properly. So, in C , George knows there's a cup on the table on the basis of his visual experience, which in this case is conclusive. Assume it's physically possible that in C instead of there being a cup on the table there's an indistinguishable hologram of a cup projected on the table. In this case, "George would not be having the visual experiences as of a cup being on the table in these circumstances unless there was a cup on the table" is false. Hence we have a "case in which [George] knows that p (there is a cup on the table) on the basis of R (his visual experience plus background knowledge) in the relevant circumstances, but the corresponding conditional [R would not have been the case unless p] is false" (p. 64), *pace* Dretske's theory.

¹⁸Here I'm assuming that the confirming evidence one has for a red barn is, in part, going to be confirming evidence for a barn. That one is seeing something that appears to have the shape of a barn, looks red, is located on a farm, seems to also be confirming evidence for there being a barn in front of one. After all, the object looks like a barn, is located where barns are usually located, and so forth. So it's not that closure goes through in this case, but rather that given the set-up the confirming evidence for the one happens to be confirming evidence for the other as well. Thanks to Berit Brogaard for pressing me on this.

My account: the good lighting, the properly functioning perceptual system, it seeming as if there's a cup on the table in C —all this provides George with confirming evidence that there's a cup on the table. If George's background information included information about hologram machines nearby or in the very room he's in, he wouldn't have robust confirming evidence for there being a cup on the table.

The final counterexample is also from Pappas and Swain (1973). Fred's good friend, John, has a back-up generator installed in his basement. Although the generator is turned off, it works perfectly fine. Fred works at the local power plant and has a good amount of knowledge about how generators work. He walks home from work, looks at the various houses and sees that they have their lights on. When he gets to John's house, he sees the lights are on there as well. Fred comes to believe p = the downtown generators are powering the lights by John's house. Assume p is true. Fred isn't aware of the backup generator in John's basement that comes on smoothly and without any hitch when the power goes off in the house. He walks in John's house and continues to believe p on the conclusive basis of seeing the lights on his way to Fred's house, seeing the lights on in Fred's house, and his background knowledge of how the power plant and generators works. “[S]urely we do not want to say that the fact that his friend [John] has a generator in his basement prevents [Fred] from having knowledge that company's generators are causing the lights to be on” (Pappas and Swain 1973: p. 66).

My account: Fred seeing the lights on in the various houses in the neighborhood, seeing the lights in John's house, along with his background knowledge that comes from working at the power plant, provides him with confirming evidence for p . However, if Fred's background information included that John has a generator in

his basement that goes on without a hitch when the power goes off or had the background information that the power just went off, then he wouldn't have the robust confirming evidence for p .

No doubt there are other counterexamples to address, but hopefully I've given one a sense for how they would be dealt with.

2.6.2 Applications

Let me now discuss some interesting epistemological consequences that result from a rejection of closure. Without epistemic closure, it is possible to provide a principled response to some very hard and long-standing problems concerning knowledge and rationality. In one way or another, closure plays a central role in the skeptical argument discussed above, the problem of easy knowledge, the lottery paradox, and the preface paradox.¹⁹

Moreover, as I discussed in chapter 1, closure plays a key role in the Bayesian objection to dogmatism—the view that, for certain contents p , if it appears to S as if p , then S is justified in believing that p , regardless of whether S is independently justified in rejecting a skeptical hypothesis like being a handless BIV. Without the principle, we showed how the dogmatist could resist the Bayesian objection. In closing, I'll mention one more anti-dogmatist argument that relies on closure.

This objection to dogmatism is from White (2006), and it turns on a *dominance principle* of justification. Roughly, the dominance principle of justification states that if S is justified in believing at t_0 that she will be justified in believing p at t_1 ,

¹⁹For more on the problem of easy knowledge, see Cohen (2002). For the lottery paradox, see Kaplan (1996) and Wheeler (2007). And for the preface paradox, see Makinson (1965) and Christensen (2004).

then S is justified in believing p at t_0 (or, more formally, $J_{t_0}[J_{t_1}p] \Rightarrow J_{t_0}p$).

Consider *super-fake-hands*. Like fake-hands, these aren't real hands, but they are visually indistinguishable from real hands. In addition, they have magical powers that prevents anyone from gaining evidence that they aren't real hands. Let Hands-3* = This is not a super-fake-hand, and suppose Moore is about to offer a new proof of the external world using his hands. Suppose too that before Moore presents his proof we don't have any independent reason to think he doesn't have super-fake-hands. Before Moore takes his arms out of his pockets, there are three possibilities about how his arms will appear to us.

1. He will not appear to have hands.
2. He will appear to have hands, but we will have reason to suspect that the appearance is deceptive in some way.
3. He will appear to have hands, and we will have no reason to suspect that the appearance is deceptive in any way.

For White, the interesting case here is the third one. If this case obtains, and if dogmatism is correct, then “we will be justified in believing it's a hand on the end of his arm, and hence in turn that it's not a super-fake-hand” (White 2006: p. 538). So dogmatism has the consequence that “no matter how things may appear to us when we see what is in [Moore's] coat pocket, we will be justified in believing Hands-3*” (ibid). By the aforementioned dominance principle of justification, it follows that before Moore even takes his arms out of his pockets, we're justified in believing Hands-3*. But, according to White, that seems wrong, since we have no reason to think Moore doesn't have super-fake-hands. In fact, says White, it's quite possible that Moore *does* have super-fake-hands.²⁰

²⁰White's argument here has an implicit premise, and it's that we can justifiably tell whether

I only want to point out that this argument relies on the assumption that if we're justified in believing that Moore has hands, then we're also justified in believing that Moore doesn't have super-fake-hands. This is supposed to follow since having hands entails not having super-fake-hands. But that assumption, of course, relies on a closure principle. So by rejecting closure this anti-dogmatist argument will not work.

2.7 Conclusion

This chapter has outlined a Bayesian account of closure failure. My account is inspired in many ways by Dretske's work on closure. I've shown that his intuitions about relations of evidential support can be both substantiated rigorously in the Bayesian framework and put to use in a Bayesian account of closure failure. The arguments given here don't show that we should endorse my account. At best, they do point in the direction one might take in giving a satisfactory account of closure failure.

we're in case 2 or case 3. If we can't tell whether we're in case 2 or case 3, I think it's clear we won't be justified in believing Hands-3*. Whether or not this assumption is plausible is not something I can pursue here.

CHAPTER 3
SUPERMAJORITY VOTING: A BAYESIAN PERSPECTIVE

3.1 Introduction

Rousseau was an early advocate of the idea that a supermajority vote should be used in deliberations on particularly important or fundamental matters. Consider the following passages.

[T]he more important and solemn the matters under discussion, the nearer to unanimity the voting should be... (Rousseau, *The Social Contract*, bk. 4, ch. 2)

Between the veto and plurality...there are various propositions for which one can determine the preponderance of opinions according to the importance of the issue. For example, when it concerns legislation, one can demand at least three-fourths of the votes, two-thirds for matters of State, a simple plurality for elections and other affairs of the moment. This is only an example to explicate my idea and not a proportion that I recommend. (Rousseau, *Considerations on the Government of Poland*, ch. 9, cited in Weirich (1986))

Rousseau's idea here seems to be that there is some attribute particular to supermajority voting that makes it preferable, at least in some circumstances, to simple majority voting. Let's call Rousseau's claim here the *supermajority requirement*. Amendments to the U.S. Constitution only become valid when they are

ratified by 3/4 of the states' legislative bodies. Countries must receive a 2/3 majority vote to acquire a seat on the Security Council of the United Nations. So the idea behind Rousseau's supermajority requirement is both familiar and intuitive: when much is at stake, sizable majorities are preferable. "Narrow electoral victories", as Goodin and Estlund (2004) put it, "should pack less of a punch...than should massive majorities" (p. 140). Aside from intuition, however, what can be said in favor of the supermajority requirement? The answer isn't immediately obvious, and Rousseau's own rationale for the supermajority requirement is far from clear.¹

In this chapter, I examine the supermajority requirement from a Bayesian perspective. I give two theorems in the Bayesian framework, each demonstrating the *evidential superiority* of supermajority voting over simple majority voting. I explain why the evidential feature is significant from a Bayesian epistemological perspective, and discuss how it can be used to motivate the supermajority requirement. In many ways the Bayesian rationale outlined here is guided by the Condorcet framework and its well-known jury theorem. Below, I bring these connections out, and suggest some advantages with the approach taken here.

The structure of this chapter is as follows. In the next section, I present a Condorcet jury theorem with supermajority voting rules, and explain the relevance of the result to the supermajority requirement. Section 3.3 presents the main formal results that support the Bayesian approach to supermajority voting. I end, in section 3.4, with a discussion of the implications of the results for the supermajority requirement, and compare our Bayesian approach to the Condorcet approach.

¹See Weirich (1986) for an interesting discussion of Rousseau on supermajority voting.

Before proceeding, let me flag four issues this chapter does not address. First, I do not defend supermajority voting from various objections, e.g., that it’s “politically impractical”, or “too conservatively biased”. Second, I do not defend the claim that supermajority voting is the uniquely or maximally rational voting procedure, applicable in all possible voting scenarios. Third, while it is true that what is meant by “important” deliberation needs more careful explanation, I offer no analysis of the notion here; instead, I will simply assume that such deliberations exist. Finally, I am aware that various “political” or “practical” considerations may also be given to motivate the supermajority requirement, and while a discussion of what these conditions could (or could not) be is important, I do not pursue them here.

3.2 Condorcet Rationale for Supermajority Requirement

Fey (2003) proves an important (Condorcet) jury theorem with supermajority voting rules. This section presents the result and discusses its bearing on the supermajority requirement.

3.2.1 The Model

In presenting the Condorcet model, I follow the notation in Fey (2003) closely. Fey gives jury theorems for both *homogeneous* groups, where each voter is identically competent, and *heterogeneous* groups, where each voter is not identically competent. I restrict my attention to his theorem for homogeneous groups; nothing said below will hang on this.

Given n jurors, let p_i denote the probability that juror i votes correctly ($0 \leq p_i \leq 1$). We denote the vector of these probabilities with $p = (p_1, \dots, p_n)$. We represent the state of the world with a Bernoulli random variable $X = \sum_{i=1}^n X_i$, with parameter p_i .² Juror i is said to have made the *correct* decision if the correct alternative was actually chosen by i . So if i casts a vote for “guilty”, and the defendant is in fact “guilty”, then i ’s vote is correct.

Denote $\lceil x \rceil$ as the smallest integer greater than or equal to x . Let q be the fraction of the majority required to choose some alternative. Then when $\lceil qn \rceil$ votes are necessary to make a decision it is a *supermajority rule*, often called a *q-rule* in the literature. A majority voting rule corresponds to $q = 1/2 + 1/n$, and a unanimity rule corresponds to $q = 1$. We will use the notation $P_{q,n}(X \geq qn)$ to represent the probability that a supermajority group of size n makes the correct decision using a q -rule.³

Along with the assumption that each juror is more likely than not to arrive at the correct decision, the Condorcet model assumes that (i) each juror has the same chance of being correct, and (ii) all votes among the jurors are independent of each other. Also, since we are working in the “classic” Condorcet model, we assume that each voter acts “sincerely”, or “non-strategically”, when casting their vote.⁴

The standard Condorcet Jury Theorem concerns homogeneous groups with majority rules. The theorem has been extensively discussed in the literature, and

² X is said to have a Bernoulli distribution if $P(X = 1) = p$ and $P(X = 0) = 1 - p$.

³Following Fey, I assume that no decision is made if neither choice receives the supermajority needed for passage. Paroush (1984) and Ben-Yashar and Nitzan (1997) address supermajority voting in contexts where the “status quo” is privileged, i.e., unless there’s a supermajority to overturn a status quo, the status quo remains in place.

⁴For a discussion of jury theorems in strategic voting contexts, see Feddersen and Pesendorfer (1996) and Duggan and Martinelli (2001).

many important generalizations of the theorem have been obtained.⁵ I will state the jury theorems in terms of the “limit”, but notice the formula at work in both is given by

$$\sum_{n/2+1}^n (n!/(n/2+1)!(n-(n/2+1))!p^{(n/2+1)}(1-p)^{(n-(n/2+1))}.$$

The Condorcet Jury Theorem was first expressed in Marquis de Condorcet’s *Essay on the Application of Analysis to the Probability of Majority Decisions*.⁶ The original formulation of the theorem is as follows.⁷

Theorem 5 (Condorcet). *Assume $p_1 = \dots = p_n = p$ and $q > 1/2 + 1/n$. Let $P_{1/2+n/2,n} = P_n$.*

1. *If $1/2 < p < 1$ and $n \geq 3$, then $P_n > p$ and $P_n \rightarrow 1$ as $n \rightarrow \infty$.*
2. *If $0 < p < 1/2$ and $n \geq 3$, then $P_n < p$ and $P_n \rightarrow 0$ as $n \rightarrow \infty$.*
3. *If $p = 1/2$, then $P_n = 1/2$ for all n .*

For a sufficiently large n , Fey shows that a similar results holds for supermajority voting rules.

Theorem 6 (Fey). *Assume $p_1 = \dots = p_n = p$ and $q > 1/2$.*

1. *If $p > q$, then there exists an integer N such that for all $n > N$, $P_{q,n} > p$ and $\lim_{n \rightarrow \infty} P_{q,n} = 1$.*

⁵These generalizations include showing that version of the theorem holds when the juror competence assumption is relaxed (Grofman, Owen and Field (1983), Borland (1989)), showing that a version of the theorem holds when the symmetrical competence assumption is dropped (List 2004), showing that a version holds when certain dependences among the jurors are allowed (Hawthorne (2001), Estlund (1994), Ladha (1992)), and showing the theorem holds when the jurors’s have a plurality of choices, not just binary ones (List and Goodin (2001)).

⁶For interesting historical discussion of Condorcet’s work, see Hacking (2006).

⁷See Miller (1986) for a proof, and Black (1958) for extended discussion.

2. If $p < q$, then there exists an integer N such that for all $n > N$, $P_{q,n} < p$ and $\lim_{n \rightarrow \infty} P_{q,n} = 0$.

3. If $p = q$, there exists an integer N such that for all $n > N$, $P_{q,n} < p$ and $\lim_{n \rightarrow \infty} P_{q,n} = 0$.

Fey's jury theorem shows that if the individual competencies of the voters is greater than the fraction q needed for passage, then a group decision is more likely to be correct than the decision of a randomly chosen individual. The theorem is qualified for a sufficiently large n , because, as Fey points out, if n is not sufficiently large, the theorem fails. For example, by letting $n = 100$, $q = 2/3$, and $p = .673$, it can be shown that $P_{2/3,100} = .5752 < .673 = p$.⁸

3.2.2 Application to Supermajority Requirement

Clearly, the Condorcet model is concerned primarily with *limit* results and *likelihoods*, which theorists in turn standardly use to motivate various majority voting procedures.⁹ In examining the supermajority requirement, we are after a *synchronic* (epistemic) rationale that will tell us, *right here* and *right now*, why supermajority voting is preferable to simple majority voting when the proposition or referendum being voted on is important.

⁸It turns out that the "formula" being used here is somewhat more complicated than Fey's discussion suggests. However, by taking advantage of the fact the binomial distribution is very nearly a normal distribution (for $n > 12$), the salient formula can be stated as follows:

(*) For $n > 12$, where z is the point at which the area under the standard normal distribution (from minus infinity up to point z) is p is the closest integer just larger than $p(1-p)[z/(p-q)]^2$.

Using Fey's example, we find that $p(1-p)[z/(p-q)]^2 = 1122.8$, so N is around 1123. Notice that this formula is somewhat sensitive to small differences in p , but we need around $N = 1000$ in any case.

⁹See, e.g., Spitz (1984).

In view of Fey’s result, we can state the following (Condorcet) rationale for the supermajority requirement. The Condorcet model shows that, for a sufficiently large n —say, 10,000—a supermajority of 7,500 votes for choice H implies that H will be much more probably correct than before the vote. The model also implies that when H is correct it is harder to get a supermajority to vote for it, and when H is not correct, it is very unlikely that a supermajority will vote for H . So, when H is important, it is preferable to opt for supermajority voting, because not only does a supermajority vote for H make it much probable that H is correct, but it also has a very high false negative rate. Call this the *Condorcet rationale* for the supermajority requirement. Section 3.4 will discuss this rationale in a bit more detail.

My Bayesian approach in what follows will be similar to the Condorcet approach in that I will attempt to provide a synchronic (epistemic) rationale for the supermajority requirement. But, instead of focusing on limits and likelihoods, I will be concerned with the confirmation-theoretic “core” of the Condorcet model. In particular, I will adopt confirmation-theoretic analogues of the Condorcet assumptions mentioned here and give results highlighting the evidential significance of supermajority voting. While the Bayesian rationale is not incompatible with the Condorcet rationale, it is shown, in closing, to have some important advantages.

In the next two sections, I will present my Bayesian approach to supermajority voting. Section 3.3 presents the main theoretical results that motivate the Bayesian approach, and section 3.4 discusses the implications of these results for the supermajority requirement.

3.3 Results

This section presents two new results that illustrate the confirmatory power of multiple evidential statements taken together. My operative assumptions here are the confirmation-theoretic analogues of the Condorcet assumptions listed above, though some of them are shown to be a bit weaker (logically) than the Condorcet ones.¹⁰ As usual, we will be working in the Bayesian confirmation-theoretic framework outlined in chapter 1. And, because measure sensitivity is not our concern here, I restrict my attention to relevance measure l , the log-likelihood ratio measure.

In his paper on independent evidence, Fitelson (2000) describes a notion of independence, called *confirmational independence*, that is weaker than (Condorcet) probabilistic independence. Following Pierce (1878), Fitelson says that if two pieces of evidence, E_1 and E_2 , provide independent inductive support for a hypothesis H , then the degree to which E_1 supports H should not depend on whether E_2 is part of our background evidence (and vice versa). Letting $c(H, E_i | E_j)$ be the degree E_i supports H , on the supposition that E_j is true, and $c(H, E_i)$ be the degree E_i supports H unconditionally (i.e., relative only to *tautological* background information), Fitelson defines the notion of confirmational independence as follows.

Definition. E_1 and E_2 are (mutually) *confirmationally independent* regarding H according to c iff both $c(H, E_1 | E_2) = c(H, E_1)$ and $c(H, E_2 | E_1) = c(H, E_2)$.¹¹

Now it can be shown that confirmational independence is a weaker notion of independence than (Condorcet) probabilistic independence. Specifically, I will find an

¹⁰For related results, see List (2004) and List and Pettit (2004).

¹¹See Appendix A.2 for technical discussion of this definition.

E_1 and E_2 such that both (α) and (β) are true, which will suffice to demonstrate that confirmational independence is a weaker notion than probabilistic independence.

$$(\alpha) l(H, E_i | E_j) = l(H, E_i) [i \neq j].$$

$$(\beta) P(E_1 \& E_2 | H) \neq P(E_1 | H) \times P(E_2 | H).$$

To this end, I will describe a class of probabilities spaces which have properties (α) and (β) . These probability spaces contain three events— H , E , E' —where the eight atomic events in the space each have a specified probability. I borrow the following probability model \mathcal{M} from Fitelson (2001).

$P(H \& \sim E \& \sim E') = \frac{683}{3800}$	$P(H \& E \& \sim E') = \frac{1}{50}$
$P(\sim H \& E \& \sim E') = \frac{1169}{1102000}$	$P(H \& \sim E \& E') = \frac{1}{50}$
$P(H \& E \& E') = \frac{1}{456}$	$P(\sim H \& E \& E') = \frac{2}{29}$
$P(\sim H \& \sim E \& E') = \frac{922}{1305}$	$P(\sim H \& \sim E \& \sim E') = \frac{15179}{9918000}$

Setting $E = E_1$ and $E' = E_2$, it is easy to verify that properties (α) and (β) obtain in \mathcal{M} (proof omitted). Thus, confirmational independence is a weaker notion than probabilistic independence.

Fitelson (2000) proves a number of results concerning confirmational independence, one of which is the following.¹²

Theorem 7 (Fitelson). *If $c(H, E_1) > 0$, $c(H, E_2) > 0$, and $c(H, E_2 | E_1) = c(H, E_2)$, then $c(H, E_1 \& E_2) > c(H, E_1 \& \sim E_2)$.*

¹²My statement of this theorem is a bit different than it appears in Fitelson's paper.

This result shows that two pieces of independent confirmatory evidence will always provide better confirmational support when taken together than when considered individually.

Before we discuss the significance of this theorem for the voting contexts we are interested in, we present our first result, which is a generalization of Fitelson’s theorem. Some notation:

E = The conjunction of all of the past outcomes, some of which may be positive, some negative (the exact makeup doesn’t matter).

N = The new (next) outcome, which may be positive or negative (the exact makeup doesn’t matter, except where explicitly stated below).

H = The hypothesis under evaluation.

In the following theorem, I “weaken” the antecedent of Fitelson’s theorem, and so “strengthen” the overall result.

Theorem 8. *If $c(H, E | N) = c(H, E)$ or $c(H, N | E) = c(H, N)$, then*

1. *If $c(H, N) > 1$, then $c(H, E \& N) > c(H, E) > c(H, E \& \sim N)$.*
2. *If $c(H, N) = 1$, then $c(H, E \& N) = c(H, E) = c(H, E \& \sim N)$.*
3. *If $c(H, N) < 1$, then $c(H, E \& N) < c(H, E) < c(H, E \& \sim N)$.*

Remark. It is easy to show that

$$c(H, E | N) = c(H, E) \iff c(H, N | E) = c(H, N)$$

so one could use either of these in the antecedent, depending on what makes most sense for the given context. Notice that by letting $E = E_1$ and $N = E_2$ we have shown Fitelson’s assumption that $c(H, E_1) > 0$ is unnecessary.

In the voting contexts we're interested in, the result tells us that two “guilty” votes are better, evidentially speaking, than one “guilty” vote. This is the general line we will take in motivating the supermajority requirement. But the scope of this result is somewhat limited, since it only implies that *two* pieces of independent confirmatory evidence provide better support than one individual piece of confirmatory evidence. We desire a theorem that generalizes to the case of *three or more* pieces of evidence.

As the following theorem shows, we can derive the desired result for three or more pieces of evidence, and at the same time give the exact factor to which the additional evidence supports H . Here I assume, in effect, that each vote is confirmationally independent, and also that each vote is “equally weighted”. I use the notation $\pm E_i$ to stand for E_i being either positive or negative, where E_i is the i^{th} alternative possible outcome of n experiments.

Theorem 9. *Let there be n distinct experiments, each with an alternative possible outcome, denoted by E_k and $\sim E_k$ for the k^{th} one. Let E be the conjunction of $m < n$ distinct outcomes or their alternatives, and suppose that for any possible distinct outcome E_k such that neither it nor its negation is a conjunct in E . Assume:*

1. For all E_i , $c(H, \pm E_i) > 0$.
2. $c(H, \pm E_1 \ \& \ \dots \ \& \ \pm E_m) = \prod_{i=1}^m c(H, \pm E_i)$.
3. $c(H, \sim E_k \mid E) = c(H, \sim E_k) = c(\sim H, E_k)$.
4. For all E_j distinct from E_k , $c(H, E_j) = c(H, E_k)$.

Then for any conjunction E consisting of m outcomes of the form E_i and $(n-m)$ outcomes of the form $\sim E_j$,

$$c(H, E) = (m - (n - m)) \times c(H, E_k) = (2m - n) \times c(H, E_k).$$

Remark. Above I said that I was in “in effect” assuming confirmational independence in this theorem. I made this qualification because while condition 2, the so-called “additivity” condition, does *not* entail confirmational independence for cases involving exactly n statements, it *does* entail confirmational independence for each possible subset of n statements (proof omitted). So, from a logical point of view, condition 2 is weaker than confirmational independence—hence I adopt it here. But note that a theorem very much like this one can be proved with the following confirmational independence condition in place of condition 2:

$$c(H, E_k | E) = c(H, E_k).$$

Second, I laid down an “equal weight” assumption in stating this theorem. The motivation for doing this is to rule out cases where “your vote” may overwhelm cumulative votes of others, by somehow being “more important”—stronger evidence—than the cumulative votes of many of the other people.

Having presented the formal results, I now discuss their implications for the supermajority requirement.

3.4 Implications & Advantages

What, exactly, do the results show, and how do they bear on the supermajority requirement? In essence, my results show that all we have to do is count the number of positive reports among a complete set of reports to tell how it compares in its joint confirmational power regarding H with any other complete set of reports.

With this in mind, I am in position to offer the following Bayesian rationale for the supermajority requirement. For the voting contexts we are interested in, I

have shown that supermajority voting, compared to simple majority voting, provides the *best evidence* for the truth of the proposition or referendum H being voted on. This is significant from a Bayesian epistemological perspective because one's credences in the truth of a hypothesis H should be, as I argued in chapter 2, influenced by the amount of confirming evidence one has (lacks) for H . That is, from a Bayesian epistemological perspective, the more confirming evidence one has for hypothesis H , the more one is *justified* in believing that H is true.¹³ Consequently, since supermajorities provide *superior evidence* for H , it follows from our Bayesian epistemological thesis that a supermajority vote on H gives us stronger justification for thinking that H is the correct choice. And, when much is at stake, such strong justification certainly seems desirable. This is our Bayesian rationale for the supermajority requirement, and while the main idea is simple, it provides a natural and elegant explanation for why supermajority voting is preferable, at least in some circumstances, to simple majority voting.

What advantages, if any, does the Bayesian approach have over the Condorcet approach described above? I suggest two, each resulting from limitations with the Condorcet approach.

First, Fey's jury theorem for special majority rules holds only if the size of the voting group n is sufficiently large. When n is not sufficiently large— $n = 100$, for example—it was shown not to be the case that a special majority is more likely to arrive at the correct decision than a randomly chosen individual. This leads to a problem with the Condorcet rationale not shared by our Bayesian rationale. To see this, suppose a group of 100 is deliberating on an important proposition or referendum H . Then, since H is important, it follows by the special majority requirement's "more is better" intuition that a 2/3 majority of 67 votes for H

¹³See Hawthorne (2005) and Fitelson (2007) for related discussion.

is preferable to a simple majority of 51 votes for H . Similarly, according to the requirement, 38 votes for H out of 50 is better than a simple majority of 26 votes for H . Given the size of n in these cases, Fey's jury theorem doesn't apply, and so it's not clear what the Condorcet approach can say about the epistemological significance of supermajority voting in these cases. Our Bayesian approach, by contrast, is *not* sensitive to the size of n . Whether it be for a jury of size 12, or a legislative body of size 1000, the Bayesian approach has an account for why "more is better" when voting on an important H . Since an adequate rationale for the supermajority requirement should be able to accommodate all different sizes of voting groups, this difference highlights the first advantage of the Bayesian approach.

Second, Fey's jury theorem holds only when the individual competences are greater than the fraction q needed for passage, i.e., $p > q$. For most q rules, this implies that individual competences need to be quite high, significantly greater than .5. Assuming that individual competences are slightly greater than .5 seems reasonable enough. But what reason do we have for thinking that individual competences are *much greater* than .5? Very little, I'd say. In fact, as the size of the voting group increases, the assumption looks quite implausible.¹⁴ What's more, Fey says nothing substantive to motivate this assumption.¹⁵ In contrast, the Bayesian approach only requires, along with confirmational independence, that

¹⁴Interestingly, Condorcet himself would seem to agree. He writes:

A very numerous assembly cannot be composed of very enlightened men. It is even probable that those comprising this assembly will on many matters combine great ignorance with many prejudices. Thus there will be a great number of questions on which the probability of the truth of each voter will be below 1/2. It follows that the more numerous the assembly, the more it will be exposed to the risk of making false decisions. (cited in Waldron 1999: p. 32)

¹⁵To be fair, Fey is concerned solely with technical issues, and so the lack of motivation for the assumption is somewhat understandable in his context. For related discussion, see List (2004).

each vote is given equal evidential weight. And, *prima facie*, the equal weight assumption looks far less controversial than assuming $p > q$; indeed, the assumption seems quite natural in any democratic framework.

So there are at least two limitations of the Condorcet rationale not shared by the Bayesian approach. This implies, I suggest, that the Bayesian approach is preferable to the Condorcet approach.¹⁶

3.5 Conclusion

In this chapter, I have examined the supermajority requirement from an epistemological perspective. In particular, I have given results to explain why, from a Bayesian epistemological point of view, supermajority voting is preferable to simple majority voting in some circumstances. I have identified an alternative Condorcet rationale for the requirement, based a jury theorem, and discussed some advantages of my approach over the Condorcet approach.

¹⁶See List (forthcoming) for additional worries with the epistemic significance of supermajority voting in the Condorcet model.

CHAPTER 4
PROOF BEYOND REASONABLE DOUBT AND THE
PRESUMPTION OF INNOCENCE: AN ESSAY IN BAYESIAN
JURISPRUDENCE

4.1 Introduction

In Anglo-American criminal jurisprudence, *proof beyond reasonable doubt* is an important and frequently used standard of criminal proof. According to this standard—as usually conceived—the jurors begin trial with a “presumption of innocence”, and then it’s the prosecution’s task to present a case of the defendant’s guilt to the jury in “such convincing character that a reasonable person would not hesitate to rely and act upon it in the most important of his own affairs.” (Devitt 2003: sec. 12.10).¹

As Justice Lemuel Shaw notes, though, ‘reasonable doubt’ is a “term often used, probably pretty well understood, but not easily defined” (*Commonwealth vs. Webster*, 59 Mass. 295, at 320, 1850). Indeed, recent work on defining the notion of proof beyond reasonable doubt is vague and often confusing.² This chapter works to remedy the situation. I develop a new and precise account of proof beyond reasonable doubt. My approach will make central use of Bayesian confirmation theory, and the idea—very roughly—will be that proof beyond reasonable doubt is established when a certain level of confirmatory support is reached. I will also discuss the “presumption of innocence” doctrine, since it will play an important

¹See also Firth Circuit Criminal Jury Instruction 1. 06, 1990, and Mueller and Kirkpatrick (2003).

²I’m not alone in thinking this. Laudan (2006), for example, says recent work on proof beyond reasonable is “obscure, incoherent, and muddled” (p. 30).

role in my account of proof beyond reasonable doubt.³

Let me pause here to clarify the scope and aim of this chapter. This is an essay in Bayesian jurisprudence, in which I'll advance a particular (Bayesian) understanding of two important legal concepts. There has, of course, been some discussion on the various ways the Bayesian apparatus might be brought to bear on legal theory, but so far the discussion—especially on proof beyond reasonable doubt—has been pretty schematic.⁴ So what I'm trying to do here is present a more detailed and nuanced account of both doctrines in the Bayesian framework. Moreover, Bayesian legal theorists have so far made *no* use of confirmation theory, and here I'll try to show that confirmation theory can—and should—play an important role in Bayesian jurisprudence. So this is an essay for Bayesian legal theorists, pushing them to take seriously the importance of confirmation theory.

4.2 Against Two Recent Accounts

In this section, I discuss two currently prominent accounts of proof beyond reasonable doubt. I present some problems with each account, and then turn to my own account in section 4.3. Some of the criticisms given here are discussed elsewhere in more detail, but going through them is worthwhile because they'll point us in the

³In what follows, I won't enter the debate over whether mathematical machinery—probability theory, in particular—can usefully be applied to the law. I find the arguments against applying mathematical tools to law rather unconvincing, but I won't argue for that here; instead, I will simply assume that such tools can be usefully applied. Below, I will discuss this issue a bit further. See Robertson and Vignaux (1993), Kaye (1979), and Finkelstein (1978). For a nice example of the fruits of applying formal machinery to legal theory, see Lillquist (2002).

⁴For example, Fienberg and Kadane (1983) present a Bayesian interpretation of proof beyond reasonable doubt, but aside from mentioning the need for a high posterior probability of guilt they provide very few details.

direction of a more adequate account of proof beyond reasonable doubt.⁵

4.2.1 First Account

The first account we'll consider focuses on the connection between belief and action.

According to this account:

A reasonable doubt is a doubt based on a reason and common sense—the kind of doubt that would make a reasonable person hesitate to act. Proof beyond reasonable doubt must, therefore, be proof of such convincing character that a reasonable person would not hesitate to rely and act upon it in the most important of his own affairs.

[Devitt (1987)]

So, for this account, if, at the end of a trial, the jurors find themselves with the sorts of doubts that would cause them to hesitate in acting in their own personal and important affairs, then the prosecution has failed to establish the defendant's guilt beyond reasonable doubt.

This account is pretty popular in the literature, but it faces a number of problems. I'll discuss four.

First, in ordinary life, we often take action, despite having considerable doubt. For instance, I'll marry my fiancé even though I don't believe beyond a reasonable doubt that our marriage will work out, or I'll move across the country to take

⁵The remarks in this section are heavily indebted to Lauden (2003, 2006). For further discussion, see Tillers and Gottfried (2007), Newman (2007), Franklin (2006), and Weinstein and Dewsbury (2006).

a more prestigious job, despite some doubts about the company's recent financial performance. Since "even sizable doubts cause little or no hesitancy to act" (Laudan 2006: p. 38), this account's claim that reasonable doubt would induce hesitancy to take action on such important affairs is mistaken.⁶

Second, whether or not one would hesitate to act on important affairs doesn't seem very relevant to determining if the prosecution has established proof beyond reasonable doubt or not. Laudan (2006) brings the worry out with the following example.

Try to imagine a mathematician saying that he has proved a new theorem and that the proof consists in the fact that he believes the theorem without the slightest hesitation. His colleagues would be aghast since what establishes a mathematical theorem as a theorem is the robustness of its proof, *not* the confidence of its discoverer. (p. 52; emphasis in original)

Some might object that this example relies on the idea that when experts believe p without the slightest hesitation that's no evidence for p . On the contrary, says the objection, when experts have such beliefs that's often *good evidence* for p . I think that's probably right, but I don't think it can be the whole story. Even if the jurors believe without hesitation that the defendant is guilty, and we take that

⁶Related to this objection, Laudan says the account is problematic because there are cases where we hesitate to act, despite there being an absence of reasonable doubt. "Many people", he says, "confronted with major life decisions, fidget and fret even when it is wholly clear and beyond doubt what course of action they should take" (2006: p. 37). His example is a "battered wife, knowing perfectly well that her husband will continue beating her when money is short in the household, will nonetheless often hesitate about extricating herself from the situation" (pp. 37-38). This doesn't strike me as a very compelling counterexample. It's hard to see how the battered wife is being *rational* in staying with her husband. Also, suppose S is presented with the following two options: perform A and get \$100 or perform B and get \$20 dollars. It would be extremely odd for S to say, "I believe beyond reasonable doubt that I should do A , but maybe I shouldn't perform A , and instead opt for B ".

to be some evidence for her guilt, I think an adequate account of proof beyond reasonable doubt should have something to say about how strong and compelling the prosecution's case needs to be, about the kind of evidence they need to bring to bear throughout the trial. I'll return to this point below.

Third, notice that there's something odd about trying to give an account of proof beyond reasonable doubt in terms of our willingness to take action in our most important life decisions. This is because legal decisions don't seem anything like the (important) decisions we make in personal lives. The Judicial Conference of the United States puts the point this way.

Indeed, decisions we make in the most important affairs of our lives—choosing a spouse, a job, a place to live, and the like—generally involve a very heavy element of *uncertainty* and *risk-taking*. They are wholly unlike the decisions jurors ought to make in criminal cases. (Federal Judicial Center, Pattern Criminal Jury Instructions 18-19, commentary on instruction 21, cited in Laudan (2006); emphasis mine)

It should be mentioned, I think, that in both this objection and the belief-action account of proof beyond reasonable doubt being considered here, the operative decision-theoretic assumption seems to be that we act on the basis of p . A proponent of the belief-action account might try to modify this assumption by saying that it's not p that we act upon, but rather *probably, p* or *very probably, p*. Indeed, such a move might fit well with the idea that in any criminal standard of proof there's bound to be some uncertainty. I don't have any knockdown objection to this proposal, but I would point out that it requires a pretty radical decision theory, something similar to Jeffery's (1992) "radical probabilism". I think it'd be preferable if we could give an account of proof beyond reasonable doubt that

didn't require countenancing such a view.

Finally, notice this account doesn't lay down any set of *normative constraints* on the jurors and their decision making procedures. This is problematic for a number of reasons, one being that irrational people act without the slightest hesitation all the time, and it's clear this account doesn't want to allow for this—or any other related—possibility. So there needs to be some story here about the rationality of the jurors's beliefs, about how they *should* respond to the prosecution's evidence, about how they *ought* to deliberate in light of the total evidence, and so forth.

4.2.2 Second Account

On our second account, proof beyond reasonable doubt is glossed in terms of an “abiding conviction” in the defendant's guilt. More specifically, here's the proposal.

Reasonable doubt is defined as follows: It is not mere possible doubt; because everything relating to human affairs...is open to some possible doubt or imaginary doubt. It is that state of the case which after the entire comparison and consideration of all evidence, leaves the mind of the jurors in that condition that they cannot say they feel an abiding conviction...of the truth of the charge. (Jury Instructions, California Courts, 114 S. Ct, at 1244; cited in Laudan (2006)).

Now I take it that this proposal does *not* require the prosecution to convince the jury, over a sustained period of time, that the defendant is guilty. Rather, the idea is that the prosecution's task is to present such a compelling case that each juror comes away with a firm conviction in the defendant's guilt after deliberation.

To assess this proposal, we need to clear up an ambiguity concerning the necessity and sufficiency of this proposal for proof beyond reasonable doubt. On one hand, if this proposal is saying that an abiding conviction among the jurors is both necessary *and* sufficient for proof beyond reasonable doubt, then the proposal fails on the sufficiency front, since the strength or depth of one’s belief B often has no bearing on whether or not B is justified or evidentially well-founded. Dick Cheney had a deep and abiding conviction that there was an Iraq-Al-Qaeda connection, but that belief wasn’t justified, and it wasn’t based on any evidence. On the other hand, if the account is meant to state a necessary, but not sufficient, condition for establishing proof beyond reasonable doubt, then the account is incomplete. I have no objections to saying a firm conviction is a necessary condition for proof beyond reasonable doubt, but that doesn’t give us a very illuminating account of what proof beyond reasonable doubt amounts to.

Another worry with this proposal is that it’s unclear what an “abiding conviction” is supposed to be. Is it meant to just be a high degree of belief in the defendant’s guilt, or is there supposed to be something more—a phenomenal feel, perhaps—that comes along with such a conviction?

4.2.3 Moving Past Subjectivity

Before I present my account of proof beyond reasonable doubt, I want to note that each account discussed here relies importantly on some *subjective* factor. The first account is concerned with what actions an individual would or would not perform, while the second is concerned with the strength of one’s convictions. It seems to me, however, that a more satisfactory account of proof beyond reasonable doubt should move away from such subjective factors, and instead focus on more *objective*

factors, such as how well the prosecution's evidence *supports* or *speaks in favor* of the hypothesis that the defendant is guilty. That is, a more adequate account of proof beyond reasonable doubt should discuss the *kind of evidence* needed for the prosecution to establish proof beyond reasonable doubt. The focus on subjective states of the jurors threatens to ignore the quality and strength of the prosecution's evidence, and, by my lights, that's precisely the issue we should be focused on.

In seeking a more objective account of proof beyond reasonable doubt, then, I'm agreeing with Laudan that "[t]he principal question is not whether the jurors...are convinced by the prosecution" (2006: p. 39). "The issue", he says, "is whether the evidence they [i.e., the jury] have seen and heard *should be* convincing in terms of the level of support it offers to the prosecution's hypothesis that the defendant is guilty" (ibid; emphasis in original).

The rest of this chapter works to develop an objective account of proof beyond reasonable doubt, one that focuses on the kind and quality of the prosecution's evidence. In view of the accounts discussed here and their problems, Laudan concludes that "[s]hort of some form of radical surgery, [proof beyond reasonable doubt]'s day has come and gone" (p. 62). My goal in what follows will be to show that, *pace* Laudan, proof beyond reasonable doubt *can* be rigorously defined.

4.3 A New Bayesian Account

Since my account of proof beyond reasonable doubt will make central use of Bayesian probability and confirmation theory, it'll be useful to have some preliminaries on the Bayesian notions of prior and posterior probability in place.

4.3.1 Prior and Posterior Probability

For Bayesians, *posterior probability*, denoted $P(H | E \& K)$, represents the *net plausibility* of a hypothesis H , given evidence E and background corpus K . This probability is used by Bayesians to assess the *evidential impact* of E on H relative to K .

By Bayes’s theorem, we have

$$P(H | E \& K) = \frac{P(E | H \& K) \times P(H | K)}{P(E | K)}$$

which tells us that the value of $P(H | E \& K)$ is influenced by two factors—*likelihoods*, written $P(E | H \& K)$, and *priors*, written $P(H | K)$.

Likelihoods are the *evidential* factors that influence the posterior $P(H | E \& K)$. In particular, likelihoods are probabilities that describe what H says about E ; that is, they describe the *empirical import* of H . Likelihoods have a special status for Bayesians, since they’re considered “objective”, quasi-logical in character, and there’s usually a good deal of inter-subjective agreement on their values.

Prior probabilities, on the other hand, represent the *non-evidential* factors that influences the value of $P(H | E \& K)$. In particular, prior probabilities capture the plausibility of H by taking into account all the relevant factors the agent can bring to bear on H , *except for the evidence E* . Unlike likelihoods, priors are considered more “subjective”, and there’s usually little inter-subjective agreement on their values.⁷

The following example shows how priors and likelihoods interact to influence posterior probability. Let H be the hypothesis that “disease D is present”, and

⁷See Jaynes (2003), Hawthorne (2004), Mellor (2005), and Maher (forthcoming) for a more in depth discussion of Bayesian probability.

E be the evidence “the test for D is positive”. Suppose one tests positive on a medical test with a .05 false positive rate, and a .01 false negative rate (i.e. a .99 true positive rate). Then the relevant likelihoods here are $P(E | H) = .99$ and $P(E | \sim H) = .05$, and the evidence will strongly favor the hypothesis that you have the disease. Suppose now that the prior that you have the disease is very low, say $P(H) = .001$. Then the probability that you have the disease, given the positive test result, is only $P(H | E) = .0194$.

A few remarks about function P and its interpretation are in order here. Bayesians usually give the function a “subjective” gloss, meaning it measures one’s *degree* or *strength of belief* in a hypothesis H .⁸ So on this interpretation $P_\alpha(H | E)$ measures α ’s belief strength in H given evidence E , while $P_\beta(H | E)$ represents β ’s belief strength in H given E . How subjective one wants to push this interpretation will depend on what sort of constraints, if any, are placed on the agent. Radical subjectivists place *no* constraints on the agent, while more modest—and more plausible—brands of subjectivism hold that the agent should at least be logically consistent and in conformity with the axioms of probability.

“Objective” interpretations of P , by contrast, maintain that the function measures some objective (statistical) frequency, or degree of logical entailment, or something closely related. Why move toward a more objective interpretation of P ? Intuitively, there seems to be a difference between an agent *thinking* that E is evidence for H and E *in fact* being evidence for H , and for many it’s unclear how a subjective interpretation could capture this distinction. That is, if P is given a subjective gloss, it’s hard to see how $P_\alpha(H | E)$ represents anything more than what α *thinks* the relationship between H and E might be. Since there’s something very intuitive about the notion of evidential support, independent of

⁸For sustained discussion and defense of subjective Bayesianism, see Colin and Howson (1993).

what anyone believes, and it's unclear how a subjective P could capture this, it's argued that P must be glossed in a more objective way.⁹

I think the distinction between thinking E supports H and E in fact supporting H is important, and I agree that any (at least plausible) attempt to capture it with a probability measure requires a more objective understanding of the function. So, in what follows, I'll be understanding P in a more objective vein. Taking this line, of course, does *not* commit me to any sort of Carnapian, i.e. purely syntactic, understanding of P (Carnap 1950, 1962). What do I mean by an objective understanding of P , then? Here I have in mind something similar to Williamson's (2000) notion of *evidential probability*. On this approach, $P(H)$ measures "something like the intrinsic plausibility of $[H]$ prior to investigation" (Williamson 2000: p. 211), and $P(H | E)$ represents how strongly E in fact supports H .

4.3.2 The Proposal—First Pass

I'll begin with a rough statement of my account of proof beyond reasonable doubt, and then I'll work to tie up loose ends. For now, I'll use the notation G to stand for the hypothesis that the defendant is *guilty*, and E to stand for the *total evidence* of both the prosecution and defense. The meaning of G and E will be sharpened in section 4.3.3.

In any Anglo-American criminal trial, the accused is "presumed innocent" and then it's up to the prosecution to present a strong case, in which the total evidence—evidence for and against guilt, also taking the defense's rebuttal into account—supports the hypothesis that the defendant is guilty. What's required by

⁹For further discussion of objective Bayesianism and its motivation, see Rosenkrantz (1977) and Seidenfeld (1979). Also relevant here are Williamson (2000) and Hajek (2003).

the prosecution to establish proof beyond reasonable doubt?

Let's start by getting clear on the *kind* of evidence the prosecution should bring to bear in the trial. I say the prosecution's evidence should be *confirmatory evidence*. Confirmatory evidence seems like a natural choice, because it's evidence that's concerned with degree of support. So on this picture it's the prosecution's task to present the jury with evidence that supports or speaks in favor of the hypothesis that the defendant is guilty.

Proof beyond reasonable doubt, however, is a very demanding standard of criminal proof, so it's not enough to simply specify the kind of evidence needed in the prosecution's case. In addition to being confirmatory, the prosecution's evidence needs to be both extremely strong and highly relevant to the defendant's guilt. In confirmation-theoretic terms, I will take this to mean that the total evidence, E , must support the hypothesis that the defendant is guilty, G , to a very high degree. In particular, I will say that E must confirm G over a *high threshold* τ , which I will write as $c(G, E) > \tau$.¹⁰ One constraint I place on the threshold τ is that it must be high enough to yield a high posterior probability $P(G|E)$. Without this requirement, it would be possible for $c(G, E)$ to be relatively high, and yet the posterior $P(G|E)$ to be quite low.¹¹ And I think it'd be pretty odd if our account had the consequence that proof beyond reasonable doubt could be established even though $P(G|E)$ remained quite low, especially given the importance of posteriors

¹⁰The degree to which evidence E confirms a hypothesis H is measured by some relevance c measure, such as the ones discussed in the Introduction. As usual, I assume $c = l$ here and in what follows.

¹¹Example: Let H be "disease present", and let E be "the test is positive". Assume one tests positive on a medical test with a .05 false positive rate, and a .01 false negative rate (i.e. a .99 true positive rate). Then

$$c(H, E) = \frac{P(E|H)}{P(E|\sim H)} = \frac{.99}{.05} = 19.8$$

but the posterior $P(H|E) = .0194$ is quite low.

for Bayesians. Below, I'll have more to say about the values of τ and $P(G | E)$.

Roughly, then, the proposal is that proof beyond reasonable doubt is established by the prosecution in criminal trial C iff $c(G, E) > \tau$ is true in C , where τ is high enough to yield a high $P(G | E)$ in C . This isn't the whole story, though, since we haven't said anything yet about the "presumption of innocence" doctrine alluded to above, and it turns out that this doctrine places some important constraints on our account of proof beyond reasonable doubt.

4.3.3 The Presumption of Innocence

In any Anglo-American criminal case each juror must begin trial with a "presumption of innocence" (hereafter (PI)).¹² What does it mean to presume innocence at the start of a trial? Consider the following remarks from Wigmore (1981).

The 'presumption of innocence' is in truth merely another form of expression for a part of the accepted rule for the burden of proof in criminal cases, i.e., the rule that it is for the prosecution to adduce evidence...and to produce persuasion beyond a reasonable doubt...the 'presumption' implies what the other rule says, namely, that the accused...may remain inactive and secure, until the prosecution has taken up its burden and produced evidence and effected persuasion; i.e., to say in this case, as in any other, that the opponent of a claim or charge is presumed not to be guilty is to say in another form that the proponent of the claim or charge must evidence it. However, in a criminal case the term does convey a special and perhaps useful hint, over and

¹²I'm extremely grateful to Jim Hawthorne for discussion in this section.

above the other form of the rule about the burden of proof, in that it cautions the jury to put away from their minds all the suspicion that arises from the arrest, the indictment, and the arraignment, and to reach their conclusion solely from the legal evidence adduced. (pp. 530-531)

So the picture is that the jurors must begin trial with an open mind, free of biases and suspicions, and then it's their task to determine the defendant's guilt or innocence based solely on considerations from the total evidence. To give this doctrine a Bayesian gloss, the natural idea is to require each juror to begin the trial with a *low prior probability for guilt* (i.e., the value of $P(G)$ should be low). Indeed, most Bayesian legal theorists accept some form of the "low prior" gloss of (PI), and I'll adopt it here.¹³

Question: To respect (PI), how low should $P(G)$ be? At a minimum, we know that $P(G) \neq 0$; for if $P(G) = 0$, then $P(G | E) = 0$, making it impossible for the prosecution to adduce any evidence in favor of guilt. And splitting the difference by setting $P(G) = .5$ seems way too high. So what options remain?

I think we can make some headway on specifying a value for $P(G)$ by considering the following proposal from Lindley (1977).

[t]he presumption of innocence is itself an approximation: if a crime has been committed in Britain by someone in these islands then on that evidence alone, the probability that a given person, for example myself, is guilty is about N^{-1} where N is the population of these islands. (p. 218)

¹³See Rawling (1999) and Fienberg and Kadane (1983).

By letting the population on the islands be 10000, for example, it follows on this approach that the prior $P(G)$ should be

$$P(G) = 10000^{-1} = \frac{1}{10000} = .0001.$$

Notice, however, that when N is quite large, it'll force the jurors's priors to be *extremely low*. And, as shown in the medical example, such priors can make it quite hard to drive up the posterior probability. More concretely, here are two ways in which low priors could make trouble for the Bayesian approach to (PI).

The first worry is from Rawling (1999). He argues that a “low prior” gloss of (PI) makes it almost impossible for a Bayesian juror to ever convict. His argument is as follows. Say that the posterior $P(G | E)$ should be around .95 for establishing proof beyond reasonable doubt. Suppose N is quite large—1000000, say. Then a simple calculation shows that the likelihood ratio $P(E|G)/P(E|\sim G) \approx 1.9 \times 10^7$.¹⁴ But, according to Rawling,

the ratio $[P(E | G)]/[P(E | \sim G)] \approx 1.9 \times 10^7$ is so large that a Bayesian juror will rarely if ever vote to convict...even assuming that the prosecution is so persuasive as to push the juror's estimate of $P(E | G)$ to one (in fact $P(E | G)$ would be far smaller than this), $P(E | \sim G)$ would still be a minute 5.3×10^{-8} (p. 121).

¹⁴To see this, just observe that

$$\frac{P(E | G)}{P(E | \sim G)} = \frac{P(\sim G)}{P(G)} \times \frac{P(G | E)}{P(\sim G | E)}.$$

So when $P(G | E) = .95$ and $P(G) = .000001$, we have

$$\begin{aligned} \frac{P(E | G)}{P(E | \sim G)} &= \frac{.999999}{.000001} \times \frac{.95}{.05} \\ &= 18999981 \\ &\approx 1.9 \times 10^7. \end{aligned}$$

So, if Rawling is right, a Bayesian “low prior” gloss of (PI) appears to have a pretty unappealing consequence for the Bayesian juror.

Another worry with a “low prior” gloss of (PI) is that it can be shown that if we make $P(G)$ too low, it’ll be impossible for $c(G, E)$ —no matter how high—to drive up the posterior $P(G | E)$. That is, given any (high) value of $c(G, E)$, it’s possible to make $P(G | E)$ as low as you want, by simply selecting a small enough value for $P(G)$. The proof of this is given in the Appendix.

For the remainder of this section, I’ll try to show that Bayesians shouldn’t be too worried about using low priors—even extremely low priors—in the account of (PI). In particular, I’ll try to show that even with very low priors it’s quite possible to obtain a high posterior $P(G | E)$ at the end of the trial. I’ll also show that my proposal is not susceptible to the two aforementioned worries. It takes some work to unpack the formal machinery underlying my proposal, but the ensuing discussion is interesting enough, I think, to warrant this brief (formal) digression.

For what follows, it will prove useful to put Bayes’s theorem in its “odds” or “odds against” form:

$$\begin{aligned}
 \Omega(\sim G | E \& K) &=_{df} \frac{P(\sim G | E \& K)}{P(G | E \& K)} \\
 &= \frac{P(E | \sim G \& K)}{P(E | G \& K)} \times \frac{P(\sim G | K)}{P(G | K)} \\
 &= \frac{P(E | \sim G \& K)}{P(E | G \& K)} \times \Omega(\sim G | K).
 \end{aligned}$$

We get posterior probabilities back from odds via the following equation:

$$P(G | E \& K) = \frac{1}{1 + \Omega(\sim G | E \& K)}.$$
¹⁵

For notational simplicity, I'll drop the background K for the remainder of the discussion. Notice that when $\Omega(\sim G | E)$ gets very small, $P(G | E)$ approaches 1, and as $\Omega(\sim G | E)$ gets very large, $P(G | E)$ approaches 0. Also, observe that as the ratio $P(E | \sim G)/P(E | G)$ goes to 0, so does posterior odds against G , and as $P(E | \sim G)/P(E | G)$ gets very large, so does posterior odds against G . All of this is important because it shows that the “prior odds against G ” only has the effect of “slowing down” or “speeding up” (by a constant factor) what the likelihood ratios are doing.

Now, instead of having the jurors assess the likelihoods $P(E|G)$ and $P(E|\sim G)$, at one time, with a “total evidence” E , we'll say that the jurors break E down into “working parts” in their assessment of guilt. I'll first explain this notion formally, and then show how it helps deal (formally) with the problem of low priors. In closing, I'll describe in less formal terms how the idea would work in practice.

Assume we “chunk” or “split up” E into to n parts, denoted E_1, \dots, E_n . Abbreviate E_1, \dots, E_n by E^n . Relative to an appropriate background K , which contains general knowledge, but nothing about the particular defendant, notice that E is equivalent to E^n (i.e. $K \models E \leftrightarrow (E_1 \& E_2 \& \dots \& E_n)$). With E^0 taken to be some

¹⁵The proof of this is given in the Appendix. We make use of “odds” here, because when $c = l$ it is easiest to see the role c plays in figuring posterior probabilities. From Bayes's theorem, we get

$$\Omega(H | E) = \frac{P(H | E)}{P(\sim H | E)} = c(H, E) \times \frac{P(H)}{P(\sim H)} = c(H, E) \times \Omega(H)$$

So the likelihood ratio version of c can be thought of as directly updating to posterior odds from prior odds by multiplying the prior odds. See Hawthorne (2004, forthcoming) for useful discussion of Bayes's theorem and its “odds against” version.

tautology, it's a theorem of probability that for $k = 1$ through n ,

$$P(E^k | G) = P(E_k | G \& E^{k-1}) \times P(E^{k-1} | G).$$

So—and this will be important for below—we have

$$\begin{aligned} P(E^n | G) &= \prod_{k=1}^n P(E^k | G \& E^{k-1}) \\ &= P(E_1 | G) \times P(E_2 | G \& E^1) \\ &\quad \times \dots \times P(E_n | G \& E^{n-1}). \end{aligned}$$

Similarly,

$$\begin{aligned} P(E^n | \sim G) &= \prod_{k=1}^n P(E^k | \sim G \& E^{k-1}) \\ &= P(E_1 | \sim G) \times P(E_2 | \sim G \& E^1) \\ &\quad \times \dots \times P(E_n | \sim G \& E^{n-1}). \end{aligned}$$

Understanding the exact content of ‘ $\sim G$ ’ is a bit tricky, but I think it’s useful to break up $\sim G$ into a bunch of different hypotheses about the crime, i.e., how the juror could be faced with the same evidence and yet the defendant is innocent. So, for example, we could have G_1 be that the defendant is guilty, G_2, \dots, G_m be specific alternative hypotheses the jurors are able to construct, and G_c be the “catch-all” that says the defendant didn’t do it, but none of G_2, \dots, G_m are true either ($G_c = (\sim G_1 \& \sim G_2 \& \dots \& \sim G_m)$). One reason for breaking up G this way is that the assessment of likelihoods becomes easier relative to more specific hypotheses. Done this way, the relevant likelihoods for us become:

$$P(E^n | G_j) = \prod_{k=1}^n P(E_k | G_j \& E^{k-1}).$$

In any case, I’ll use the notation G for “guilty” ($G = G_1$) and H for $\sim G_1$, any specific alternative G_2, \dots, G_m , or the catch-all G_c .

By “chunking” or “splitting up” E into working parts, notice what we do to the likelihood ratios:

$$\begin{aligned}\frac{P(E^n | H)}{P(E^n | G)} &= \frac{\prod_{k=1}^n P(E_k | H \& E^{k-1})}{\prod_{k=1}^n P(E_k | G \& E^{k-1})} \\ &= \prod_{k=1}^n \left(\frac{P(E_k | H \& E^{k-1})}{P(E_k | G \& E^{k-1})} \right)\end{aligned}$$

In other words, by splitting up E , we chunk the likelihood ratios as well.

So what we now know is that when we assess how much more or less plausible, given all the evidence, G is to H , we take products of likelihoods of some chunks on hypotheses together with conjunctions of other relevant chunks:

$$\begin{aligned}\frac{P(H | E^n)}{P(G | E^n)} &= \frac{P(H)}{P(G)} \times \frac{P(E^n | H)}{P(E^n | G)} \\ &= \frac{P(H)}{P(G)} \times \prod_{k=1}^n \left(\frac{P(E_k | H \& E^{k-1})}{P(E_k | G \& E^{k-1})} \right)\end{aligned}$$

This applies, of course, to any pair of alternatives, whether it be G_3 to G_1 or G_c to G_1 . And if the jurors were working with *multiple alternatives*, we would have

$$\begin{aligned}\Omega(\sim G_1 | E^n) &= \frac{P(\sim G_1 | E^n)}{P(G_1 | E^n)} \\ &= \sum_{k=2}^m \left(\frac{P(G_k)}{P(G_1)} \times \prod_{k=1}^n \frac{P(E_k | G_k \& E^{k-1})}{P(E_k | G_1 \& E^{k-1})} \right)\end{aligned}$$

This completes the formal digression. We’ll now put all this to work in solving the “too low of priors” problem for a Bayesian gloss of (PI).

For concreteness, let the population $N = 100,000,000$. By Lindley’s proposal, the prior probability of guilty, $P(G_1)$, is as follows:

$$P(G_1) = 10^{-7} = \frac{1}{100,000,000} = .00000001.$$

So there are one hundred million suspects at the beginning of the trial. The prosecution must (significantly) narrow down the field and show that the defendant

on trial is indeed guilty of the crime. If the prosecution can't pin the crime down on the defendant, the jury should acquit. How does the prosecution go about pinning the guilt on the defendant?

Typically, they introduce bits of evidence—here is where chunking starts to play a role—placing the defendant at the crime scene, showing that the defendant had the right motive(s) to have committed the crime, demonstrating that the defendant had access to the victim at the time of the crime, and so forth. So, for example, they might introduce some witness reports placing the defendant at the crime scene—call this evidence E_1 —which will make $P(E_1 | G_1)$ pretty high, leaving $P(E_1 | H)$ for most other hypotheses quite low.¹⁶ For most alternative G_k , then, $P(E_1 | G_k)$ is either 0 or very near 0. To see why this is important, suppose E_1 is useful enough to eliminate a massive amount of potential suspects, leaving us with 1000 suspects that still have the right motive, access, and whereabouts to have committed the crime. Assuming that the $P(E_1 | G_k)$ are all approximately the same, we know that

$$\begin{aligned}\Omega(\sim G_1 | E^1) &= \frac{P(\sim G_1 | E^1)}{P(G_1 | E^1)} \\ &= \sum_{k=2}^{100000000} \left(\frac{P(G_k)}{P(G_1)} \times \prod_{k=1}^n \frac{P(E_k | G_k)}{P(E_k | G_1)} \right)\end{aligned}$$

which gives us

$$\sum_{k=2}^{100000000} \left(\frac{P(G_k)}{P(G_1)} \times \prod_{k=1}^n \frac{P(E_k | G_k)}{P(E_k | G_1)} \right) = \sum_{k=2}^{1000} \left(\frac{P(G_k)}{P(G_1)} \right) \approx \frac{999}{1}.$$

What this means is that the prosecution's E_1 took us from 100000000 to 1 odds down to 1000 to 1 odds.¹⁷ Of course, this isn't enough to convict, but it gives us

¹⁶In fact, most $P(E_1 | H)$ will be close to 0. The ones that aren't 0 will be alternatives that place the person at the scene of the crime.

¹⁷Indeed, it can be shown as a theorem that (proof omitted)

$$\Omega(\sim G_1 | E^n) = \frac{P(\sim G_1 | E^n)}{P(G_1 | E^n)} = \sum_{k=2}^m \left(\frac{P(G_k | E^1)}{P(G_1 | E^1)} \times \prod_{k=2}^n \frac{P(E_k | G_k \& E^{k-1})}{P(E_k | G_1 \& E^{k-1})} \right).$$

an idea how the prosecution methodically builds their case, bit by (evidential) bit, against the defendant. As the trial goes on, for example, the prosecution might introduce another piece of evidence, $E_2 =$ the defendant's finger prints on a knife at the crime scene, and $E_3 =$ the defendant's hair at the crime scene, which, when taken together— $E_1 \& E_2 \& E_3$ (i.e. E^3)—continues to eliminate potential suspects from the remaining 1000.

Formally, the process I'm describing here can be thought of a sequence of updates in this sense:

$$\begin{aligned} \Omega(\sim G_1 | E^n) &= \frac{P(\sim G_1 | E^n)}{P(G_1 | E^n)} \\ &= \sum_{k=2}^m \left(\frac{P(G_k | E^{h-1})}{P(G_1 | E^{h-1})} \times \prod_{k=h}^n \frac{P(E_k | G_k \& E^{k-1})}{P(E_k | G_1 \& E^{k-1})} \right). \end{aligned}$$

In other words, each new piece of evidence the prosecution subsequently introduces—at least, ideally—will make most ratios $P(G_k | E^{h-1})/P(G_1 | E^{h-1})$ very small or 0. It does this by making the successive likelihood ratios— $P(E_{h-1} | G_k \& E^{h-2})/P(E_{h-1} | G_1 \& E^{h-2})$ —small enough to keep eliminating the remaining alternatives, until we are left with only G_1 , the guilty hypothesis, having a posterior of any significant value.

So if the prosecution builds their case with multiple bits of evidence that has likelihood near 0 on the alternative hypotheses, but not near 0 on the guilty alternative, they can make the posterior $\Omega(\sim G | E^n)$ approach 0, which will make the posterior $P(G | E^n)$ approach 1. In short, then, I've show that even if priors get extremely low, it's still possible for the prosecution to drive up the posterior probability of guilt to a value near 1.

With this analysis, we can now see that Rawling’s worry—Bayesian jurors would almost never convict—is unfounded, since I’ve adopted an even lower prior than the one in Rawling’s objection and shown how jurors could rationally convict. And for the worry that if $P(G)$ gets too low, it’ll be impossible for $c(G, E)$ to drive up $P(G | E)$, we note that by adopting Lindley’s proposal for selecting a prior, no realistic population—even of size one hundred million—will be large enough to cause problems here.

4.3.4 The Proposal—Official Version

With these pieces in place, I’m now in position to state my account of proof beyond reasonable doubt.

- (†) Proof beyond reasonable doubt is established by the prosecution in criminal trial C iff $c(G, E) > \tau$ is true in C , where τ is high enough to yield a high posterior $P(G | E)$ in C , and each juror begins the trial with a low prior for guilt.

The details behind the account are somewhat involved, but the idea is intuitive enough. Each juror starts out with a low prior probability for guilt. Then, over the course of the trial—the prosecution’s case, the defense, and the rebuttals—it’s the prosecution’s task to present the jury with highly confirmatory evidence for the defendant’s guilt. When their evidence is strong enough, we’ll have high values for both $c(G, E)$ and $P(G | E)$ —and that, I say, is at the heart of establishing proof beyond reasonable doubt. So this account is “objective” (in Laudén’s sense), because it focuses on the kind and quality of the prosecution’s evidence, not being concerned with the jurors’s subjective mental states. Notice we obtain objectivity

in this account, because of the underlying probability function and our objective understanding of it.

At both the beginning and end of the criminal trial, we have probabilities: low priors at the start, high posteriors at the end. But the middle part of the trial—the back and forth between the prosecution and defense—is all about likelihood ratios and confirmation. Indeed, it’s confirmation that pushes us to high posterior at the end of the trial. So my approach has both a probabilistic and confirmation-theoretic component, and I’ve tried to show that confirmation is central to establishing proof beyond reasonable doubt.

An important, but distinct, question here is how the jurors *determine* whether or not the prosecution has established proof beyond reasonable doubt. My view on this is simple: assuming the prosecution has successfully established proof beyond reasonable doubt, I say the jurors render the correct verdict when they accurately judge the evidential import of the prosecution’s case.

4.4 Objections

The theory sketched in the previous section seems intuitive and compelling to me. Here I present and respond to a number of objections to my proposal.

Objection: Your Bayesian account says that $c(H, E) > \tau$ must be true to establish proof beyond reasonable doubt. This assumes that proof beyond reasonable doubt can be *quantified*. But, as the Appeals Court in Massachusetts puts it, “[t]he idea of reasonable doubt is not susceptible to quantification; it is inherently qualitative” (*Massachusetts vs. Sullivan*, 482 N. E. 2d 1198, 1985).

Reply: It's pretty hard to tell, exactly, what the objection to quantifying reasonable doubt amounts to here. I suspect it has to do with the fact that quantification requires the use of mathematical machinery in legal theory, which has long been looked upon with suspicion by some prominent legal theorists.¹⁸ But—and though they are overstating the point a bit—Tillers and Gottfried (2006) are correct when they point out that “it borders on the ridiculous to argue that the use of numerical probabilities or odds in criminal trials is harmful simply because the use of numbers (or quantitative expressions) in criminal trials is harmful” (p. 143). Now if the worry with quantification is that it rules out the possibility of uncertainty in the jurors's minds, observe that quantification and uncertainty aren't necessarily incompatible; indeed, my quantified account explicitly allows for some uncertainty.¹⁹

Objection: What level should threshold τ be set at, and what level should the posterior $P(G | E)$ be? It's not obvious how to answer these questions, and any answer is going to seem pretty *arbitrary*.

Reply: To deflect any worry of arbitrariness here, observe that proof beyond reasonable doubt is an exacting standard of criminal proof. So one thing we know is that the values of τ and $P(G | E)$ are going to be quite high. Plausibly, the value of $P(G | E)$ should be somewhere around .95 or .99. I'm inclined to have $P(G | E) \approx .99$ —especially in capital punishment cases—and depending on what the prior is (as given by Lindley's proposal), we'll be able calculate what value we need for τ . For example: Suppose we want $P(G | E) \approx .99$ and our prior $P(G) \approx .01$. Then it's easy to verify that we need a value of $c(G, E) \approx 10000$ to

¹⁸Tribe (1970a, 1970b) are two classic papers expressing skepticism of “trial by mathematics”; also relevant are Sharivo (1989) and Finkelstein and Fairley (1970, 1971).

¹⁹Weinstein and Dewsbury (2006) suggest some arguments in favor of a quantified approach to proof beyond reasonable doubt, and Tillers and Gottfried (2006) offer a defense of quantifying proof beyond reasonable doubt.

convict. So, in short, I don't see much arbitrariness here.

Objection: On your account of proof beyond reasonable account, the likelihood ratio, $c(G, E)$, must be at above a threshold τ together with a high posterior $P(G | E)$ (.99, say). This means that as sufficient condition for acquittal there would be either a likelihood ratio with a value less τ or a posterior with a value less than .99. Consequently, in the following situation the jury should acquit because the likelihood ratio $c(G, E)$ is too low.

$$\begin{aligned} 99 &\leq \frac{P(G | E)}{P(\sim G | E)} \\ &= \frac{P(E | G)}{P(E | \sim G)} \times \frac{P(G)}{P(\sim G)} \\ &< \tau \times \frac{P(G)}{P(\sim G)}. \end{aligned}$$

But this can happen iff $P(G) > 99/(99 + \tau)$.²⁰ So, even if the posterior is high enough, on your account it looks like all the threshold τ is doing is requiring that we acquit *unless* the value of the posterior got its (high) value from updating on a low enough prior—i.e., a prior less than equal to $99/(99 + \tau)$.

Reply: This objection assumes that in glossing (PI) with low priors, they can't take extremely low values. But I've shown how to allow for extremely low priors in the gloss of (PI). So on my account you can still have a threshold for $c(G, E)$ that allows for acquittals even when the posterior is high enough for conviction *only because* the prior was so *high* that the prosecution's evidence didn't have to do much work at all.

Objection: Suppose a juror judges the prosecution's evidence to be highly confirmatory, and suppose she judges the evidence to be good enough to yield a high

²⁰See the Appendix for the proof.

posterior probability for guilt, and convicts the defendant. Still, “most jurors who had decided the accused committed the crime would be hard pressed to determine nonarbitrarily whether their confidence in guilt did or did not satisfy the demands imposed by [my account of proof beyond reasonable doubt]” (Laudan 2006: p. 76).²¹

Reply: This objection rests on a dubious second-order access principle, requiring the jurors to not only accurately judge the import of the prosecution’s evidence, but to also judge that they have accurately judged the import of the prosecution’s evidence. I reject this principle, and Laudan doesn’t supply any argument for it. I only require that the jurors’s beliefs be properly influenced by the prosecution’s evidence, not that they also have some sort of higher-order awareness of whether their beliefs are in fact properly influenced by the prosecution’s evidence.

Objection: By not requiring that the prosecution’s total evidence, E , to entail G or $\sim G$, your account openly accepts the possibility of *wrongful convictions*. And any account of criminal proof that explicitly acknowledges such a possibility undermines the confidence we should (and need to) have in our judicial system.

Reply: First, this objection assumes that 100 percent certainty is necessary to convict a defendant. But why think that? “No one”, writes Laudan, “believes that kind of proof is available” (2006: p. 46). Second, the objection assumes that a necessary condition for having confidence in our judicial system is believing that it’s infallible, not susceptible to mistakes. That assumption strikes me as pretty dubious. I certainly have confidence in our judicial system, even though I’ll openly admit it’s not infallible. As long as jurors take seriously their job, honestly evaluating and deliberating on the prosecution’s and defendant’s case, it’s hard

²¹Laudan discusses this objection in the context of any account of criminal proof that makes use of degrees of belief and probabilities.

to see why we shouldn't have confidence in our judicial system and its ability to render correct verdicts.

4.5 Conclusion

In this final chapter, I have developed a new Bayesian account of proof beyond reasonable doubt, making heavy use of confirmation theory. I've also discussed some important technical issues in the Bayesian "low prior" gloss of the presumption of innocence doctrine. I've made confirmation the cornerstone of my proposal, and tried to show that such an approach has a number of attractive advantages for the Bayesian.

APPENDIX A

PROOFS

This appendix proves all theorems, claims, and observations from chapters 2-4. We begin with proofs for chapter 2.

A.1 Chapter 2

Four preliminaries to start this section. First, I here assume that P is a *regular* probability measure. Second, I take logarithms of the ratios $P(H|E \& K)/P(H|K)$ and $P(E|H \& K)/P(E|\sim H \& K)$. Since the logarithm is an isotone function, no generality is lost in doing this. Third, throughout most of the Appendix, I will suppress the contents of background K . Finally, all probability models given below were obtained and verified using the *Mathematica* program **PrSAT**.¹

Proof of Theorem 1

Theorem 1. *If (i) $P(H|E) = P(H|E \& E')$, then*

$$c(H, E) = c(H, E \& E')$$

where c may be either the difference measure d , the ratio measure r , or the likelihood ratio measure l , but c may not be the normalized difference measure s .

Proof. The proof has four parts. The proofs for measures d and r are easy.

¹This program can be downloaded at www.fitleson.org/PrSAT.

1. For the case $c = d$, by definition $d(H, E) = P(H | E) - P(H)$. To show that $d(H, E) = d(H, E \& E')$, we must show that

$$[P(H | E) - P(H)] = [P(H | E \& E') - P(H)]$$

It suffices to show that

$$P(H | E) = P(H | E \& E')$$

which holds by assumption (i).

2. For the case $c = r$, by definition of r and our assumption $P(H | E) = P(H | E \& E')$, we have

$$\begin{aligned} r(H, E) &= \log\left(\frac{P(H | E)}{P(H)}\right) \\ &= \log\left(\frac{P(H | E \& E')}{P(H)}\right) \\ &= r(H, E \& E'). \end{aligned}$$

3. For the case $c = l$, notice that from our assumption $P(H | E) = P(H | E \& E')$, it follows that $P(\sim H | E) = P(\sim H | E \& E')$. So

$$\frac{P(H | E)}{P(\sim H | E)} = \frac{P(H | E \& E')}{P(\sim H | E \& E')}$$

Since

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

and

$$\frac{P(H | E \& E')}{P(\sim H | E \& E')} = \frac{P(E \& E' | H)}{P(E \& E' | \sim H)}$$

it follows that

$$\log\left(\frac{P(E | H)}{P(E | \sim H)}\right) = \log\left(\frac{P(E \& E' | H)}{P(E \& E' | \sim H)}\right).$$

Hence $l(H, E) = l(H, E \& E')$.

4. For the case $c = s$, we will show that condition (i), $P(H|E) = P(H|E \& E')$, is neither necessary nor sufficient for $s(H, E) = s(H, E \& E')$. That is, we will establish (α) and (β) .

$$(\alpha) \ s(H, E) = s(H, E \& E') \text{ and } P(H | E) \neq P(H | E \& E').$$

$$(\beta) \ s(H, E) \neq s(H, E \& E') \text{ and } P(H | E) = P(H | E \& E').$$

To this end, I will describe a class of probabilities spaces which have properties (α) and (β) . These probability spaces contain three events, H, E, E' , where the eight atomic events in the space each have a specified probability.

To establish (α) , consider the following probability model \mathcal{M}_1 .

$P(\sim H \& \sim E \& \sim E') = \mathbf{a} = .000984925$	$P(\sim H \& E \& \sim E') = \mathbf{b} = .14433$
$P(\sim H \& \sim E \& E') = \mathbf{c} = .001001$	$P(H \& \sim E \& \sim E') = \mathbf{d} = .000518997$
$P(\sim H \& E \& E') = \mathbf{e} = .567164$	$P(H \& E \& \sim E') = \mathbf{f} = .26$
$P(H \& \sim E \& E') = \mathbf{g} = .025$	$P(H \& E \& E') = \mathbf{h} = .001001$

To verify that \mathcal{M}_1 has property (α) , we first verify that $P(H|E) \neq P(H|E \& E')$ obtains in \mathcal{M}_1 , by the following computational proof.

$$P(H | E) = \frac{\mathbf{f} + \mathbf{h}}{\mathbf{e} + \mathbf{b} + \mathbf{f} + \mathbf{h}} = \frac{.261001}{.972495} = .2683829$$

$$P(H | E \& E') = \frac{\mathbf{h}}{\mathbf{e} + \mathbf{h}} = \frac{.001001}{.568165} = .00176181$$

Now we verify that $s(H, E) = s(H, E \& E')$ obtains in \mathcal{M}_1 .

$$s(H, E) = \frac{f + h}{e + b + f + h} - \frac{g + d}{1 - e - b - f - h} \approx -0.659$$

$$s(H, E \& E') = \frac{h}{e + h} - \frac{g + d + f}{1 - e - h} \approx -0.659$$

To establish (β) , consider the following probability model \mathcal{M}_2 .

$P(\sim H \& \sim E \& \sim E') = a = .145456$	$P(\sim H \& E \& \sim E') = b = .131579$
$P(\sim H \& \sim E \& E') = c = .131148$	$P(H \& \sim E \& \sim E') = d = .0294118$
$P(\sim H \& E \& E') = e = .0220871$	$P(H \& E \& \sim E') = f = .435897$
$P(H \& \sim E \& E') = g = .03125$	$P(H \& E \& E') = h = .0731707$

To verify that \mathcal{M}_2 has property (β) , we first verify that $P(H|E) = P(H|E \& E')$.

$$P(H | E) = \frac{f + h}{e + b + f + h} = \frac{.5090677}{.6627338} \approx .768$$

$$P(H | E \& E') = \frac{h}{e + h} = \frac{.0731707}{.0952578} \approx .768$$

Now we verify that $s(H, E) \neq s(H, E \& E')$

$$s(H, E) = \frac{f + h}{e + b + f + h} - \frac{g + d}{1 - e - b - f - h} \approx .591$$

$$s(H, E \& E') = \frac{h}{e + h} - \frac{g + d + f}{1 - e - h} \approx .228$$

This completes the proof of Theorem 1.

□

Proof of Theorem 2

Theorem 2. *If (i*) $P(E | H) = P(E \& E' | H)$, and (ii) $P(E) = P(E \& E')$, then*

$$c(H, E) = c(H, E \& E')$$

where c may be any measure among d , r , l , or s .

Proof. This proof has four parts.

1. For the case $c = d$, it suffices to show that

$$P(H | E) = P(H | E \& E').$$

By Bayes's theorem,

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

By assumptions (i*) and (ii),

$$\frac{P(E | H) \times P(H)}{P(E)} = \frac{P(E \& E' | H) \times P(H)}{P(E \& E')}$$

By another application of Bayes's theorem,

$$\frac{P(E \& E' | H) \times P(H)}{P(E \& E')} = P(H | E \& E').$$

Hence $d(H, E) = d(H, E \& E')$.

2. For the case $c = r$, by definition of r ,

$$r(H, E) = \log \left(\frac{P(H | E)}{P(H)} \right) = \log \left(\frac{P(E | H)}{P(E)} \right)$$

By assumptions (i*) and (ii),

$$\log \left(\frac{P(E | H)}{P(E)} \right) = \log \left(\frac{P(E \& E' | H)}{P(E \& E')} \right)$$

Hence

$$\log\left(\frac{P(E | H)}{P(E)}\right) = \log\left(\frac{P(E \& E' | H)}{P(E \& E')}\right) = r(H, E \& E').$$

3. For case $c = l$, we use the fact from case $c = d$ that $P(\sim H | E) = P(\sim H | E \& E')$. So

$$\frac{P(H | E)}{P(\sim H | E)} = \frac{P(H | E \& E')}{P(\sim H | E \& E')}$$

Since

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

and

$$\frac{P(H | E \& E')}{P(\sim H | E \& E')} = \frac{P(E \& E' | H)}{P(E \& E' | \sim H)}$$

it follows that

$$\log\left(\frac{P(E | H)}{P(E | \sim H)}\right) = \log\left(\frac{P(E \& E' | H)}{P(E \& E' | \sim H)}\right).$$

Hence $l(H, E) = l(H, E \& E')$.

4. For the case $c = s$, to show that $s(H, E) = s(H, E \& E')$, we must show that

$$[P(H | E) - P(H | \sim E)] = [P(H | (E \& E')) - P(H | \sim(E \& E'))]$$

So we must show two things, namely:

$$P(H | E) = P(H | (E \& E')) \tag{A.1}$$

and

$$P(H | \sim E) = P(H | \sim(E \& E')). \tag{A.2}$$

(A.1) is true by the proof in case $c = d$.

For (A.2), we will make use of the following theorem of the probability calculus: $P(X | \sim Y) = 1 - P(X | Y)$. Call this theorem (\dagger). By assumption (i^*), (\dagger), and some algebra, it follows that

$$P(\sim E | H) = P(\sim(E \& E') | H).$$

Now

$$P(\sim E | H) \times P(H) = P(\sim(E \& E') | H) \times P(H).$$

By (\dagger), assumption (ii), and some algebra, we have

$$\frac{P(\sim E | H) \times P(H)}{P(\sim E)} = \frac{P(\sim(E \& E') | H) \times P(H)}{P(\sim(E \& E'))}.$$

By Bayes's theorem, it follows that

$$P(H | \sim E) = P(H | \sim(E \& E')).$$

This establishes (A.2) and completes the proof. □

Proof of Theorem 3

Theorem 3. *There exists a probability model in which both: (a) $c(H, E) = c(H, E \& E')$ and (b) $P(E | H \& K) \neq P(E \& E' | H \& K)$, for all four of our c -measures d , r , l , and s .*

Proof. I will prove this theorem by describing a class of probability models in which (δ) obtains:

$$(\delta) \ c(H, E) = c(H, E \& E') \text{ and } P(E | H) \neq P(E \& E' | H).$$

For considerations of space, I will not present all the computational details necessary to verify that each probability model has property (δ).

1. For cases $c = d$, $c = r$ and $c = l$, the following probability model has property (δ) and so does the trick for these three measures.

$P(\sim H \ \& \ \sim E \ \& \ \sim E') = .00053098$	$P(\sim H \ \& \ E \ \& \ \sim E') = .0217391$
$P(\sim H \ \& \ \sim E \ \& \ E') = .001001$	$P(H \ \& \ \sim E \ \& \ \sim E') = .001001$
$P(\sim H \ \& \ E \ \& \ E') = .00101607$	$P(H \ \& \ E \ \& \ \sim E') = .930233$
$P(H \ \& \ \sim E \ \& \ E') = .001001$	$P(H \ \& \ E \ \& \ E') = .0434783$

2. For case $c = s$, the following probability model has property (δ).

$P(\sim H \ \& \ \sim E \ \& \ \sim E') = .109124$	$P(\sim H \ \& \ E \ \& \ \sim E') = .119048$
$P(\sim H \ \& \ \sim E \ \& \ E') = .0961538$	$P(H \ \& \ \sim E \ \& \ \sim E') = .209242$
$P(\sim H \ \& \ E \ \& \ E') = .12766$	$P(H \ \& \ E \ \& \ \sim E') = .121212$
$P(H \ \& \ \sim E \ \& \ E') = .001001$	$P(H \ \& \ E \ \& \ E') = .097561$

This completes the proof of Theorem 3. □

Proof of Theorem 4

Theorem 4. *For $0 < P(H) < 1$ and $0 < P(E \ \& \ E') < 1$, any pair of the following three clauses implies the remaining clause:*

1. $s(H, E) = s(H, E \ \& \ E')$.
2. $P(H \mid E) = P(H \mid E \ \& \ E')$.
3. $P(E' \mid E) = 1$.

Proof. First, we establish that $s(H, E) =$

$$P(H | E) - P(H | \sim E) = \frac{P(H | E) - P(H)}{P(\sim E)}.$$

and $s(H, E \& E') =$

$$P(H | E \& E') - P(H | \sim(E \& E')) = \frac{P(H | (E \& E')) - P(H)}{P(\sim(E \& E'))}.$$

I establish equality $s(H, E) = P(H | E) - P(H | \sim E) = \frac{P(H|E)-P(H)}{P(\sim E)}$ only. The proof for equality $s(H, E \& E') = \frac{P(H|(E\&E'))-P(H)}{P(\sim(E\&E'))}$ is similar.

By definition of conditional probability and some algebra, notice

$$P(H | \sim E) = \frac{P(H \& \sim E)}{P(\sim E)} = \frac{P(H) - P(H \& E)}{P(\sim E)}.$$

Then,

$$\begin{aligned} \frac{P(H) - P(H \& E)}{P(\sim E)} &= \frac{P(H)}{P(\sim E)} - \frac{P(H \& E)}{P(\sim E)} \\ &= \frac{P(H)}{P(\sim E)} - \frac{P(H | E) \times P(E)}{P(\sim E)} \\ &= \left(\frac{P(H)}{P(\sim E)} \right) - \left[(P(H | E)) \times \left(\frac{1 - P(\sim E)}{P(\sim E)} \right) \right] \\ &= \left[\left(\frac{P(H)}{P(\sim E)} \right) - \left(\frac{P(H | E)}{P(\sim E)} \right) \right] + P(H | E) \\ &= \frac{P(H) - P(H | E)}{P(\sim E)} + P(H | E). \end{aligned}$$

So we have

$$P(H | \sim E) = \left(\frac{P(H) - P(H | E)}{P(\sim E)} \right) + P(H | E)$$

giving us

$$\begin{aligned}\frac{P(H | E) - P(H)}{P(\sim E)} &= \frac{P(H) - P(H | E)}{P(\sim E)} \\ &= P(H | E) - P(H | \sim E).\end{aligned}$$

Now, we will proceed as follows. We will first show that if we assume claim 1, then claim 2 and claim 3 imply each other. Then we will show that claims 2 and 3 imply claim 1.

Assume claim 1 holds: $s(H, E) = s(H, E \& E')$. Then

$$\frac{P(H | E) - P(H)}{P(\sim E)} = \frac{P(H | (E \& E')) - P(H)}{P(\sim(E \& E'))}$$

From this observe that

$$P(E) = P(E \& E') \iff P(H | E) = P(H | (E \& E')).$$

Assume claims 2 and 3 hold: $P(E) = P(E \& E')$ and $P(H | E) = P(H | (E \& E'))$.

Then $s(H, E) =$

$$\frac{P(H | E) - P(H)}{P(\sim E)} = \frac{P(H | (E \& E')) - P(H)}{P(\sim(E \& E'))} = s(H, E \& E').$$

This completes the proof of Theorem 4.

□

Probability Models and Thresholds

In section 2.5, I claimed, by way of example, that it was possible to have $\log[P(E | O \& K)/P(E | \sim O \& K)] > \tau$, while $\log[P(E | H \& K)/P(E | \sim H \& K)] < \tau$,

with $H \& K \models O$. That means what we need are models with the following three properties:

$$P(O \mid H \& K) = 1 \tag{A.3}$$

$$\frac{P(E \mid O \& K)}{P(E \mid \sim O \& K)} > \tau \tag{A.4}$$

$$\frac{P(E \mid H \& K)}{P(E \mid \sim H \& K)} < \tau \tag{A.5}$$

To avoid triviality, I will assume that

$$\frac{P(E \mid H \& K)}{P(E \mid \sim H \& K)} > 0. \tag{A.6}$$

Notice that the logs don't matter, since they are inequality-preserving functions, so they can be dropped. Finding probability models with the properties outlined here is no problem, and in fact there are models for τ as large (or as small) as you like. I note that the ratios can be much bigger than 1 here, so the thresholds should be allowed to be arbitrarily large. I will now provide models in (A.3)-(A.5) obtain. For the sake of generality, I will select values 10 and 100 for our threshold τ and provide models in which properties (A.3)-(A.5) obtain.²

First, we set our threshold $\tau = 10$. Now I will describe a probability model in which the properties expressed in (A.3)-(A.5) obtain.

²I found probability models in which $\tau = .1$ and $\tau = 1$ and the properties expressed in (A.3)-(A.5) obtained as well.

$P(H \& K \& O \& E) = \frac{2}{33}$	$P(H \& K \& O \& \sim E) = \frac{1}{274}$
$P(H \& K \& \sim O \& E) = 0$	$P(H \& K \& \sim O \& \sim E) = 0$
$P(H \& \sim K \& O \& E) = \frac{1}{12}$	$P(H \& \sim K \& O \& \sim E) = \frac{1}{1374}$
$P(H \& \sim K \& \sim O \& E) = \frac{1}{52}$	$P(H \& \sim K \& \sim O \& \sim E) = 0$
$P(\sim H \& K \& O \& E) = \frac{2}{37}$	$P(\sim H \& K \& O \& \sim E) = 0$
$P(\sim H \& K \& \sim O \& E) = \frac{1}{52}$	$P(\sim H \& K \& \sim O \& \sim E) = \frac{15}{44}$
$P(\sim H \& \sim K \& O \& E) = \frac{39}{98}$	$P(\sim H \& \sim K \& O \& \sim E) = \frac{1}{1374}$
$P(\sim H \& \sim K \& \sim O \& E) = \frac{1}{565}$	$P(\sim H \& \sim K \& \sim O \& \sim E) = \frac{163616496331}{9191117845910}$

In this model, we have

$$P(O \mid H \& K) = 1$$

$$\frac{P(E \mid O \& K)}{P(E \mid \sim O \& K)} = 18.1496 > \tau = 10$$

and

$$\frac{P(E \mid H \& K)}{P(E \mid \sim H \& K)} = 5.33082 < \tau = 10.$$

This suffices to show that the properties expressed in (A.3)-(A.5) obtain in this model.

Second, we set our threshold $\tau = 100$. Now I will describe a probability model in which the properties expressed in (A.3)-(A.5) obtain

$P(H \& K \& O \& E) = \frac{2}{119}$	$P(H \& K \& O \& \sim E) = \frac{1}{88}$
$P(H \& K \& \sim O \& E) = 0$	$P(H \& K \& \sim O \& \sim E) = 0$
$P(H \& \sim K \& O \& E) = \frac{3}{140}$	$P(H \& \sim K \& O \& \sim E) = \frac{2}{59}$
$P(H \& \sim K \& \sim O \& E) = 0$	$P(H \& \sim K \& \sim O \& \sim E) = \frac{3}{100}$
$P(\sim H \& K \& O \& E) = \frac{1}{166}$	$P(\sim H \& K \& O \& \sim E) = \frac{1}{72}$
$P(\sim H \& K \& \sim O \& E) = \frac{1}{13977}$	$P(\sim H \& K \& \sim O \& \sim E) = \frac{3}{166}$
$P(\sim H \& \sim K \& O \& E) = \frac{1}{126}$	$P(\sim H \& \sim K \& O \& \sim E) = \frac{22}{197}$
$P(\sim H \& \sim K \& \sim O \& E) = \frac{1}{46}$	$P(\sim H \& \sim K \& \sim O \& \sim E) = \frac{4784145766102817}{6765914845385850}$

In this model, we have

$$P(O \mid H \& K) = 1$$

$$\frac{P(E \mid O \& K)}{P(E \mid \sim O \& K)} = 120.412 > \tau = 100$$

and

$$\frac{P(E \mid H \& K)}{P(E \mid \sim H \& K)} = 3.72481 < \tau = 100.$$

This suffices to show that the properties expressed in (A.3)-(A.5) obtain in this model.

A.2 Chapter 3

I drop the logarithmic function from l in the proofs here. This is done for notional simplicity, and no generality will be lost by doing this.

Proof of Theorem 8

Theorem 8. *If $c(H, E | N) = c(H, E)$ or $c(H, N | E) = c(H, N)$, then*

1. *If $c(H, N) > 1$, then $c(H, E \& N) > c(H, E) > c(H, E \& \sim N)$.*
2. *If $c(H, N) = 1$, then $c(H, E \& N) = c(H, E) = c(H, E \& \sim N)$.*
3. *If $c(H, N) < 1$, then $c(H, E \& N) < c(H, E) < c(H, E \& \sim N)$.*

Proof. Let $c = l$ and assume $c(H, E | N) = c(H, E)$ or $c(H, N | E) = c(H, N)$. We begin by establishing an important equality:

$$\begin{aligned} \frac{P(N \& E | H)}{P(N \& E | \sim H)} &= \frac{P(N | E \& H)}{P(N | E \& \sim H)} \times \frac{P(E | H)}{P(E | \sim H)} \\ &= \frac{P(E | N \& H)}{P(E | N \& \sim H)} \times \frac{P(N | H)}{P(N | \sim H)} \end{aligned}$$

Call this equality (*). We will make use of it several times in the proof.

From $c(H, E | N) = c(H, E)$, it follows that

$$\frac{P(E | N \& H)}{P(E | N \& \sim H)} = \frac{P(E | H)}{P(E | \sim H)}.$$

By equality (*),

$$\begin{aligned} \frac{P(N \& E | H)}{P(N \& E | \sim H)} &= \frac{P(E | N \& H)}{P(E | N \& \sim H)} \times \frac{P(N | H)}{P(N | \sim H)} \\ &= \frac{P(E | H)}{P(E | \sim H)} \times \frac{P(N | H)}{P(N | \sim H)} \end{aligned}$$

Alternatively, from assumption $c(H, N | E) = c(H, N)$, we have

$$\frac{P(N | E \& H)}{P(N | E \& \sim H)} = \frac{P(N | H)}{P(N | \sim H)}$$

By equality (*),

$$\frac{P(N | H)}{P(N | \sim H)} \times \frac{P(E | H)}{P(E | \sim H)} = \frac{P(N \& E | H)}{P(N \& E | \sim H)}$$

and

$$\frac{P(N \& E | H)}{P(N \& E | \sim H)} = \frac{P(N | E \& H)}{P(N | E \& \sim H)} \times \frac{P(E | H)}{P(E | \sim H)}$$

Hence, in either the case of $c(H, E | N) = c(H, E)$ or $c(H, N | E) = c(H, N)$, we have

$$\frac{P(N \& E | H)}{P(N \& E | \sim H)} = \frac{P(E | H)}{P(E | \sim H)} \times \frac{P(N | H)}{P(N | \sim H)}$$

To complete the proof, we have three cases.

Case 1. Assume $c(H, N) > 1$, i.e., $\frac{P(N|H)}{P(N|\sim H)} > 1$. Then,

$$\begin{aligned} \frac{P(N \& E | H)}{P(N \& E | \sim H)} &= \frac{P(E | H)}{P(E | \sim H)} \times \frac{P(N | H)}{P(N | \sim H)} \\ &> \frac{P(E | H)}{P(E | \sim H)}. \end{aligned}$$

Since

$$\frac{P(N \& E | H)}{P(N \& E | \sim H)} > \frac{P(E | H)}{P(E | \sim H)}$$

we have

$$\frac{P(N | E \& H)}{P(N | E \& \sim H)} > 1$$

So

$$\begin{aligned} P(N | E \& H) &> P(N | E \& \sim H) \\ \Rightarrow P(\sim N | E \& H) &< P(\sim N | E \& \sim H) \\ \Rightarrow \frac{P(\sim N | E \& H)}{P(\sim N | E \& \sim H)} &< 1 \\ \Rightarrow \frac{P(\sim N \& E | H)}{P(\sim N \& E | \sim H)} &< 1 \end{aligned}$$

It follows that

$$\begin{aligned}
\frac{P(\sim N \& E \mid H)}{P(\sim N \& E \mid \sim H)} &= \frac{P(\sim N \mid E \& H)}{P(\sim N \mid E \& \sim H)} \times \frac{P(E \mid H)}{P(E \mid \sim H)} \\
&< \frac{P(E \mid H)}{P(E \mid \sim H)} \\
&\Rightarrow c(H, E \& \sim N) < c(H, E)
\end{aligned}$$

Since

$$\begin{aligned}
\frac{P(N \& E \mid H)}{P(N \& E \mid \sim H)} &> \frac{P(E \mid H)}{P(E \mid \sim H)} \\
&\Rightarrow c(H, E \& N) > c(H, E)
\end{aligned}$$

we have established

$$c(H, E \& N) > c(H, E) > c(H, E \& \sim N)$$

as desired.

Case 2. Assume $c(H, N) = 1$. By replacing ‘>’ and ‘<’ everywhere by ‘=’, the claim $c(H, E \& N) = c(H, E) = c(H, E \& \sim N)$ is established by the same line of reasoning in case 1.

Case 3. Assume $c(H, N) < 1$. By interchanging ‘>’ with ‘<’ everywhere, the claim $c(H, E \& N) < c(H, E) < c(H, E \& \sim N)$ is established by the same line of reasoning in case 1. □

Proof of Theorem 9

Theorem 9. *Let there be n distinct experiments, each with an alternative possible outcome, denoted by E_k and $\sim E_k$ for the k^{th} one. Let E be the conjunction of $m <$*

n distinct outcomes or their alternatives, and suppose that for any possible distinct outcome E_k such that neither it nor its negation is a conjunct in E . Assume:

1. For all E_i , $c(H, \pm E_i) > 0$.
2. $c(H, \pm E_1 \& \dots \& \pm E_m) = \prod_{i=1}^m c(H, \pm E_i)$.
3. $c(H, \sim E_k | E) = c(H, \sim E_k) = c(\sim H, E_k)$.
4. For all E_j distinct from E_k , $c(H, E_j) = c(H, E_k)$.

Then for any conjunction E consisting of m outcomes of the form E_i and $(n-m)$ outcomes of the form $\sim E_j$,

$$c(H, E) = (m - (n - m)) \times c(H, E_k) = (2m - n) \times c(H, E_k).$$

Proof. Terms of the form $c(H, E) = \frac{P(E|H)}{P(E|\sim H)}$. So putting assumptions 1-4 in their likelihood ratio forms, we have:

For all E_i ,

$$\frac{P(\pm E_i | H)}{P(\pm E_i | \sim H)} > 1 \tag{A.7}$$

$$\frac{P(\pm E_1 \& \dots \& \pm E_m | H)}{P(\pm E_1 \& \dots \& \pm E_m | \sim H)} = \prod_{i=1}^m \left(\frac{P(\pm E_i | H)}{P(\pm E_i | \sim H)} \right) \tag{A.8}$$

$$\frac{P(\sim E_k | E \& H)}{P(\sim E_k | E \& \sim H)} = \frac{P(\sim E_k | H)}{P(\sim E_k | \sim H)} = \frac{P(E_k | \sim H)}{P(E_k | H)} \tag{A.9}$$

For all E_j distinct from E_k ,

$$\frac{P(E_j | H)}{P(E_j | \sim H)} = \frac{P(E_k | H)}{P(E_k | \sim H)}. \tag{A.10}$$

Notice that from the second equality in 3 together with 4 we get, for all j and k ,

$$\begin{aligned} \frac{P(\sim E_j | H)}{P(\sim E_j | \sim H)} &= \frac{P(E_j | \sim H)}{P(E_j | H)} \\ &= \frac{P(E_k | \sim H)}{P(E_k | H)} \\ &= \frac{P(\sim E_k | H)}{P(\sim E_k | \sim H)} \end{aligned}$$

Now, in likelihood ratio form we need to show that

$$\frac{P(E | H)}{P(E | \sim H)} = \left(\frac{P(E_k | H)}{P(E_k | \sim H)} \right)^{2m-n}$$

With no loss of generality, we can suppose that the members of E are labeled in the order $(\pm E_n \& \dots \& \pm E_1)$. Abbreviate the part of E from E_1 down to E_k (i.e., $\pm E_k \& \dots \& \pm E_1$) by E^k .

Then,

$$\begin{aligned} \frac{P(E | H)}{P(E | \sim H)} &= \frac{P(\pm E_n | \pm E_{n-1} \& \dots \& \pm E_1 \& H)}{P(\pm E_n | \pm E_{n-1} \& \dots \& \pm E_1 \& \sim H)} \\ &\times \frac{P(\pm E_{n-1} \& \dots \& \pm E_1 | H)}{P(\pm E_{n-1} \& \dots \& \pm E_1 | \sim H)} \\ &= \frac{P(\pm E_n \& \dots \& \pm E_1 | H)}{P(\pm E_n \& \dots \& \pm E_1 | \sim H)} \\ &= \prod_{i=1}^n \left(\frac{P(\pm E_i | H)}{P(\pm E_i | \sim H)} \right) \\ &= \frac{P(\pm E_n | H)}{P(\pm E_n | \sim H)} \times \prod_{i=1}^{n-1} \left(\frac{P(\pm E_i | H)}{P(\pm E_i | \sim H)} \right) \end{aligned}$$

We can carry out this process for each E^j , yielding the following equation

$$\left(\frac{P(\pm E_n | H)}{P(\pm E_n | \sim H)} \right) \times \prod_{i=1}^{n-1} \left(\frac{P(\pm E_i | H)}{P(\pm E_i | \sim H)} \right)$$

= \dots =

$$\frac{P(\pm E_n | H)}{P(\pm E_n | \sim H)} \times \dots \times \frac{P(\pm E_1 | H)}{P(\pm E_1 | \sim H)}$$

From 3 and 4 together, we have $c(H, \sim E_j) = c(H, \sim E_k)$, for all E_j distinct from E_k . By 4, it follows that

$$\left(\frac{P(E_k | H)}{P(E_k | \sim H)} \right)^m \times \left(\frac{P(\sim E_k | H)}{P(\sim E_k | \sim H)} \right)^{n-m}$$

since precisely m of the terms $\left(\frac{P(\pm E_j | H)}{P(\pm E_j | \sim H)} \right)$ are of form $\left(\frac{P(E_j | H)}{P(E_j | \sim H)} \right)$ and the remaining $n - m$ are of form $\left(\frac{P(\sim E_j | H)}{P(\sim E_j | \sim H)} \right)$. From assumption 3, it follows that

$$\left(\frac{P(E_k | H)}{P(E_k | \sim H)} \right)^m \times \left(\frac{P(E_k | \sim H)}{P(E_k | H)} \right)^{n-m} = \left(\frac{P(E_k | H)}{P(E_k | \sim H)} \right)^{m-(n-m)}$$

Since

$$\left(\frac{P(E_k | H)}{P(E_k | \sim H)} \right)^{m-(n-m)} = \left(\frac{P(E_k | H)}{P(E_k | \sim H)} \right)^{2m-n}$$

we have

$$\left(\frac{P(E | H)}{P(E | \sim H)} \right) = (2m - n) \times \left(\frac{P(E_k | H)}{P(E_k | \sim H)} \right).$$

Hence,

$$c(H, E) = (2m - n) \times c(H, E_k).$$

□

Observation: We show that if $c = l$ in Fitelson's definition of confirmational independence we don't need both $c(H, E_1 | E_2) = c(H, E_1)$ and $c(H, E_2 | E_1) = c(H, E_2)$.

Definition. E_1 and E_2 are (mutually) *confirmationally independent* regarding H according to c iff both $c(H, E_1 | E_2) = c(H, E_1)$ and $c(H, E_2 | E_1) = c(H, E_2)$.

Proof. From $c(H, E_1 | E_2) = c(H, E_1)$, it follows that

$$\frac{P(E_1 | H \& E_2)}{P(E_1 | \sim H \& E_2)} = \frac{P(E_1 | H)}{P(E_1 | \sim H)}$$

Thus,

$$\begin{aligned} \frac{P(E_1 | H)}{P(E_1 | \sim H)} \times \frac{P(E_2 | H)}{P(E_2 | \sim H)} &= \frac{P(E_1 | H \& E_2)}{P(E_1 | \sim H \& E_2)} \times \frac{P(E_2 | H)}{P(E_2 | \sim H)} \\ &= \frac{P(E_1 \& E_2 | H)}{P(E_1 \& E_2 | \sim H)} \\ &= \frac{P(E_2 | H \& E_1)}{P(E_2 | \sim H \& E_1)} \times \frac{P(E_1 | H)}{P(E_1 | \sim H)} \end{aligned}$$

Observe that the term “ $\frac{P(E_1|H)}{P(E_1|\sim H)}$ ” occurs at the beginning and the end of this list of equalities. So we have

$$\frac{P(E_2 | H)}{P(E_2 | \sim H)} = \frac{P(E_2 | H \& E_1)}{P(E_2 | \sim H \& E_1)}$$

Hence,

$$c(H, E_2) = c(H, E_2 | E_1).$$

□

A.3 Chapter 4

We prove three claims from Chapter 4.

Claim 1: We can get probabilities back from odds $\Omega(\sim G | E \& K)$ by the following equation

$$P(G | E \& K) = \frac{1}{1 + \Omega(\sim G | E \& K)}$$

Proof. It will suffice to show that

$$P(A) = \frac{1}{1 + \Omega(\sim A)}.$$

Observe that by plugging in the definition of odds to the right side we have

$$\begin{aligned} \frac{1}{1 + \Omega(\sim A)} &= \frac{1}{1 + \frac{P(\sim A)}{P(\sim\sim A)}} \\ &= \frac{1}{1 + \frac{P(\sim A)}{P(A)}} \\ &= \frac{P(A)}{P(A) + P(\sim A)} \\ &= \frac{P(A)}{1} \\ &= P(A). \end{aligned}$$

□

Claim 2: No matter how high the value of $c(H, E)$ it will always be possible to make $P(H | E)$ as small as you want, by simply choosing a small enough value for $P(H)$.

Proof. From Bayes's theorem, we have the following equation:

$$P(H | E) = \frac{1}{\left[1 + \left(\frac{1}{c(H, E)}\right) \times \left(\left[\frac{1}{P(H)}\right] - 1\right)\right]}.$$

Let $c(H, E) = r$, where r is some extremely large number, and set

$$P(H) = \frac{1}{(r \times r)} = \frac{1}{r^2}$$

From this, it follows that

$$\left(\left[\frac{1}{P(H)} \right] - 1 \right) = (r \times r) - 1$$

Notice that for a really big r , r is almost the same as $r \times r$. So,

$$\left(\frac{1}{c(H, E)} \right) \times \left(\left[\frac{1}{P(H)} \right] - 1 \right) = r - \left(\frac{1}{r} \right)$$

Since $1/r$, where r is really huge, is almost the same as r , it follows that

$$\begin{aligned} P(H | E) &= \frac{1}{\left[1 + \left(\frac{1}{c(H, E)} \right) \times \left(\left[\frac{1}{P(H)} \right] - 1 \right) \right]} \\ &= \frac{1}{\left[1 + r - \left(\frac{1}{r} \right) \right]} \end{aligned}$$

which is very nearly just $1/r$. This holds for any r you choose as the value of $c(H, E)$, regardless of how big. \square

Claim 3: The following can obtain iff $P(G) > 99/(99 + \tau)$, where $\tau = P(E | G)/P(E | \sim G)$.

$$\begin{aligned} 99 &\leq \frac{P(G | E)}{P(\sim G | E)} \\ &= \frac{P(E | G)}{P(E | \sim G)} \times \frac{P(G)}{P(\sim G)} \\ &< \tau \times \frac{P(G)}{P(\sim G)}. \end{aligned}$$

Proof.

$$\begin{aligned} 99 &< \frac{P(G)}{P(\sim G)} \times \tau \\ &\Leftrightarrow [1 - P(G)] \times \frac{99}{\tau} < P(G) \\ &\Leftrightarrow \frac{99}{\tau} - \left(\frac{99}{\tau} \times P(G) \right) < P(G) \\ &\Leftrightarrow \frac{99}{\tau} < \left(\frac{99}{\tau} + 1 \right) \times P(G) \\ &= \left(\frac{99 + \tau}{\tau} \right) \times P(G) \\ &\Leftrightarrow \frac{99}{99 + \tau} < P(G). \end{aligned}$$

□

BIBLIOGRAPHY

- Adler, Jonathan (2002) *Belief's Own Ethics*. Cambridge, MA: MIT Press.
- Adams, Fred (2006) "Tracking Theories of Knowledge", *Veritas* 50.4: 11-35.
- Adams, Fred and M. Clarke (2005) "Resurrecting the Tracking Theories", *Australasian Journal of Philosophy* 83: 207-221.
- Alston, William P. (1989) *Epistemic Justification*. Ithaca, NY: Cornell University Press.
- Audi, Robert (1993) *The Structure of Justification*. Cambridge: Cambridge University Press.
- Baumann, Peter (2006) "Information, Closure, and Knowledge: On Jager's Objection to Dretske", *Erkenntnis* 64: 403-408.
- Ben-Yashar, Ruth C. and Shmuel I. Nitzan (1997) "The Optimal Decision Rule for Fixed-Size Committees in Dichotomous Choice Situations: The General Result", *International Economic Review* 38: 175-186.
- Black, Duncan (1958) *The Theory of Committees and Elections*. Cambridge: Cambridge University Press.
- Boland, Philip J. (1989) "Majority Systems and the Condorcet Jury Theorem", *The Statistician* 38: 181-189.
- Brueckner, Anthony (1994) "The Structure of the Skeptical Argument", *Philosophy and Phenomenological Research* 54: 827-35.
- Burge, Tyler (1993) "Content Preservation", *Philosophical Review* 102: 457-488

Carnap, Rudolf (1950) *Logical Foundations of Probability*. Chicago, IL: University of Chicago Press.

——(1962) *Logical Foundations of Probability*, 2nd edition. Chicago, IL: University of Chicago Press.

Chandler, Jake (2007) “Solving the Tacking Problem with Contrast Classes”, *British Journal of the Philosophy of Science* 58: 489-502.

Chisholm, Roderick (1989) *Theory of Knowledge*. Prentice-Hall Publishing.

Christensen, David (1999) “Measuring Confirmation”, *Journal of Philosophy* 96: 437-461

——(2004) *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press.

Cohen, Stewart (2002) “Basic Knowledge and the Problem of Easy Knowledge”, *Philosophy and Phenomenological Research* 65: 309-329.

Conee, Earl and Richard Feldman (2004) *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.

Crupi, V., and Fitelson, B., and Tentori, K (forthcoming) “Probability, Confirmation, and the Conjunction Fallacy”, *Thinking and Reasoning*.

DeRose, Keith (1995) “Solving the Skeptical Problem”, *Philosophical Review* 104: 1-52.

Devitt, Edward J. et al (1987). *Federal Jury Practice and Instructions*, section 12.10, at 354.

- Dretske, Fred (1970) "Epistemic Operators", *Journal of Philosophy* 67: 1007-1023.
- (1971) "Conclusive Reasons" *Australasian Journal of Philosophy* 49: 1-22.
- (2003) "Skepticism: What Perception Teaches", in *The Sceptics*, ed. S. Luper. Aldershot: Ashgate Publishing Limited.
- (2005) "The Case Against Closure", in *Contemporary Debates in Epistemology*, eds. E. Sosa and M. Steup. Oxford: Blackwell Press. pp. 13-26.
- (2006) "Information and Closure" *Erkenntnis* 64: 409-413.
- Duggan, John and Cesar Martinelli (2001) "A Bayesian Model of Voting in Juries", *Games and Economic Behavior* 37: 259-294.
- Earman, John (1992) *Bayes or Bust?: A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
- Eells, Ellery and Branden Fitelson. (2002) "Symmetries and Asymmetries in Evidential Support", *Philosophical Studies* 107: 129-142.
- Estlund, David (1994) "Opinion Leaders, Independence, and Condorcet's Jury Theorem", *Theory and Decision* 36: 131-162.
- Feddersen, Timothy J. and Wolfgang Pesendorfer (1998) "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts", *American Political Science Review* 92: 23-35.
- Feldman, Richard (1995) "In Defence of Closure", *Philosophical Quarterly* 45: 487-494.

Fey, Mark (2003) “A Note on the Condorcet Jury Theorem with Supermajority Voting Rules”, *Social Choice and Welfare* 20: 27-32.

Fienberg, Stephen and Joseph Kadane (1983) “The Presentation of Bayesian Statistical Analyses in Legal Proceedings”, *The Statistician* 32: 88-98.

Finkelstein, Michael (1978) *Quantitative Methods in Law: Studies in the Applications of Mathematical Probability and Statistics to Legal Problems*. The Free Press.

Finkelstein, Michael O. and William B. Fairley (1970) “A Bayesian Approach to Identification Evidence”, *Harvard Law Review* 83: 489-517.

——(1971) “The Continuing Debate over Mathematics in the Law of Evidence: A Comment on ‘Trial By Mathematics,’” *Harvard Law Review* 84: 1801-1809.

Fitelson, Branden (1999) “The Plurality of Bayesian Confirmation Measures and the Problem of Measure Sensitivity” *Philosophy of Science* 66: 362-378.

——(2000) “A Bayesian Account of Independent Evidence with Applications”, *Philosophy of Science* 68: S123-S140.

——(2001) “Studies in Bayesian Confirmation Theory”, Ph.D. dissertation, University of Wisconsin, Madison.

——(2002) “Putting the Irrelevance Back Into the Problem of Irrelevant Conjunction”, *Philosophy of Science* 69: 611-622.

——(2007) “Logical Foundations of Evidential Support”, *Philosophy of Science* 73: 500-512.

- Frances, Bryan (1999) “Contradictory Belief and Epistemic Closure Principles”, *Mind & Language* 14: 203-226.
- Franklin, James (2006) “Case comment—*United States v. Copeland*, 369 F. Supp. 2d 275 (E.D.N.Y. 2005): Quantification of the ‘Proof Beyond Reasonable Doubt’ Standard”, *Law, Probability, and Risk* 5: 159-165.
- Friedman, Richard (2000) “A Presumption of Innocence, Not of Even Odds”, *Stanford Law Review* 52: 873-887.
- Fumerton, Richard (1995) *Metaepistemology and Skepticism*. Lanham, MD: Rowman and Littlefield.
- Glymour, Clark (1980) *Theory and Evidence*. Cambridge, MA: MIT Press.
- Goldman, Alvin (1986) *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Good, I. J. (1984) “The Best Explicatum for Weight of Evidence”, *Journal of Statistical Computational and Simulation* 19: 294-299.
- Goodin, Robert and David Estlund (2004) “The Persuasiveness of Democratic Majorities” *Politics, Philosophy & Economics* 3: 131-142.
- Grofman, B., G. Owen and S. L. Feld (1983) “Thirteen Theorems in Search of the Truth”, *Theory and Decision* 15: 261-278.
- Haack, Susan (1993) *Evidence and Inquiry: Towards Reconstruction in Epistemology*. Oxford: Blackwell.

Hacking, Ian (2006) *The Emergence of Probability: A Philosophical Study of Early Ideas About Probability Induction and Statistical Inference*, 2nd Edition. Cambridge: Cambridge University Press.

Hajek, Alan (2003) “Interpretations of Probability”, in *Stanford Encyclopedia of Philosophy*, ed. E. Zalta.

Hawthorne, John (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.

——(2005) “The Case for Closure”, in *Contemporary Debates in Epistemology*, eds. E. Sosa and M. Steup. Oxford: Blackwell. pp. 26-42.

Hawthorne, James (2001) “Voting in Search of the Public Good: The Probabilistic Logic of Majority Judgements”, unpublished manuscript.

——(2004) “Inductive Logic”, in *Stanford Encyclopedia of Philosophy*, ed. E. Zalta.

——(2005) “Degree-of-Belief and Degree-of-Support: Why Bayesians Need Both Notions”, *Mind* 114: 277-320.

——(forthcoming) “Confirmation Theory”, in *Handbook of the Philosophy of Science: Philosophy of Statistics*, eds. P. Bandyopadhyay and M. Forster. Elsevier.

Hawthorne, James and Branden Fitelson (2004) “Re-solving the Problem of Irrelevant Conjunction”, *Philosophy of Science* 71: 505-514.

Heathcote, C. R. (1971) *Probability: Elements of the Mathematical Theory*. Dover, DE: Dover Publications.

- Hempel, Carl (1945) “Studies in the Logic of Confirmation”, *Mind* 54: 1-26, 97-12.
- Howson, Colin and Peter Urbach (1993) *Scientific Reasoning: The Bayesian Approach*. Open Court Press.
- Jaynes, T. J. (2003) *Probability Theory: The Logic of Science*. Oxford: Oxford University Press.
- Jeffery, Richard (1992) *Probability and the Art of Judgement*. Cambridge: Cambridge University Press.
- Joyce, James (1999) *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Kaplan, Mark (1996) *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Kaye, David (1978) “The Laws of Probability and the Law of the Land”, *University of Chicago Law Review* 47: 34-56.
- Klein, Peter D. (1981). *Certainty: A Refutation of Skepticism*. Minneapolis: University of Minnesota Press.
- (2004) “Closure Matters: Academic Skepticism and Easy Knowledge”, *Philosophical Topics* 14: 165-184.
- Kolmogoroff, Andrej N. (1956) *Foundations of the Theory of Probability*, 2nd edition. New York: Chelsea Publishing Company.
- Kvanvig, Jonathan (2006) “Closure Principles”, *Philosophy Compass* 1: 256-267.

——(forthcoming) “Closure and Alternative Possibilities”, in *The Oxford Handbook of Skepticism*, ed. J. Greco. Oxford: Oxford University Press.

Ladha, Krishna (1992) “The Condorcet Jury Theorem, Free Speech, and Correlated Votes”, *American Journal of Political Science* 36: 567-576.

Laudan, Larry (2003) “Is Reasonable Doubt Reasonable?”, *Legal Theory* 9: 295-331.

——(2006) *Truth, Error, and Criminal Law: An Essay in Legal Epistemology*. Cambridge: Cambridge University Press.

Lawor, Krista (2005) “Living Without Closure”, *Grazer Philosophische Studien* 69: 93-117.

Lillquist, Erik (2002) “Recasting Reasonable Doubt: Decision Theory and the Virtues of Variability”, *U.C. Davis Law Review* 36: 85-197.

Lindley, D. V. (1977) “Probability and the Law”, *The Statistician* 26: 203-220.

List, Christian (2004) “On the Significance of the Absolute Margin”, *British Journal for the Philosophy of Science* 55: 521-544.

——(forthcoming) “The Epistemology of Special Majority Voting”, *Social Choice and Welfare*.

List, Christian and Philip Pettit (2004) “An Epistemic Free-Riding Problem?”, in *Karl Popper: Critical Appraisals*, eds. P. Catton and G. Macdonald. London: Routledge.

- List, Christian and Robert Goodin (2001) “Epistemic Democracy: Generalizing the Condorcet Jury Theorem”, *Journal of Political Philosophy* 9: 277-306.
- Luper, Steven (2006) “Dretske on Knowledge Closure”, *Australasian Journal of Philosophy* 84.3: 379-394.
- Maher, Patrick (1996) “Subjective and Objective Confirmation”, *Philosophy of Science* 63: 149-174.
- (2004a) “Bayesianism and Irrelevant Conjunction”, *Philosophy of Science* 71: 515-520.
- (2004b) “Probability Captures the Logic of Scientific Confirmation”, in *Contemporary Debates in the Philosophy of Science*, ed. Christopher Hitchcock. Oxford: Blackwell. pp. 69-93.
- (2006a) “The Concept of Inductive Probability”, *Erkenntnis* 65: 185-206.
- (2006b) “Review of David Christensen, *Putting Logic in its Place*”, *Notre Dame Journal of Formal Logic* 47: 133-149.
- (forthcoming) “Bayesian Probability”, *Synthese*.
- Maitzen, Stephen (1998) “The Knower Paradox and Epistemic Closure”, *Synthese* 114: 337-354.
- Makinson, D.C. (1965) “The Paradox of the Preface”, *Analysis* 25: 205-7.
- Marian, David and Ted Warfield (forthcoming) “Knowledge-Closure and Skepticism”, in *Epistemology: New Philosophical Essays*, ed. Q. Smith. Oxford: Oxford University Press.

McGinn, Colin (1984) “The Concept of Knowledge”, *Midwest Studies in Philosophy*. Minneapolis: University of Minnesota Press. (Reprinted in his *Knowledge and Reality: Selected Essays*, Oxford: Oxford University press. pp: 7-35. Page references to this version.)

Mellor, D. H. (2005) *Probability: A Philosophical Introduction*. London & New York: Routledge.

Miller, Nicholas R. (1986) “Information, Electorates, and Democracy: Some Extensions and Interpretations of the Condorcet Jury Theorem”, in *Information Pooling and Group Decision Making*, eds. Grofman, Bernard and Guillermo Owen. Greenwich, CT: JAI Press.

Milne, Peter (1996) “ $\log[\Pr(h|eb)/\Pr(h/b)]$ is the One True Measure of Confirmation”, *Philosophy of Science* 63: 21-26.

Mueller, Christopher and Laird Kirkpatrick (2003) *Evidence* 3rd edition. Aspen Publishers.

Newman, Jon O. (1993) “Beyond ‘Reasonable Doubt’”, *NYU Law Review* 979, 984.

——(2006) “Quantifying the Standard of Proof Beyond a Reasonable Doubt: A Comment on Three Comments”, *Law, Probability, and Risk* 5: 267-269.

Nozick, Robert (1981) *Philosophical Explorations*. Cambridge, Mass: Harvard University Press.

Olin, Doris (2006) “A Case Against Closure”, *Veritas* 50.4: 235-247.

Pappas, George and Marshall Swain (1973) "Some Conclusive Reasons Against 'Conclusive Reasons'", *Australasian Journal of Philosophy* 51: 72-76. (Reprinted in Pappas and Swain (eds.) *Essays on Knowledge and Justification*. Ithaca, NY: Cornell University Press. Page references to this version).

Paroush, Jacob (1984) "Are Qualified Majority Rules Special?", *Public Choice* 42: 257-272.

Peacocke, Christopher (2004) *The Realm of Reason*. Oxford: Oxford University Press.

Peirce, C. S. (1878) "The Probability of Induction", *Popular Science Monthly* 12: 705-718.

Pollock, J. and Cruz, J. (1999) *Contemporary Theories of Knowledge*, 2nd edition. Totowa, N. J.: Roman and Littlefield.

Popper, Karl (1959) "The Propensity Interpretation of Probability", *British Journal of the Philosophy of Science* 10: 25-42.

Pritchard, Duncan (2002) "Resurrecting the Moorean Response to Scepticism", *International Journal of Philosophical Studies* 10: 283-307

——(2005a) *Epistemic Luck*. Oxford: Oxford University Press.

——(2005b) "Wittgenstein's *On Certainty* and Contemporary Anti-Scepticism", in *Investigating On Certainty: Essays on Wittgenstein's Last Work*, eds. D. Moyal-Sharrock and W. H. Brenner. Palgrave Macmillan.

Pryor, James (2000) "The Skeptic and the Dogmatist", *Nous* 34: 517-549

Rawling, Piers (1999) “Reasonable Doubt and the Presumption of Innocence: The Case of the Bayesian Juror”, *Topoi* 18: 117-126.

Renyi, Alfred (1970) *Foundations of Probability*. Holden-Day, Inc.

Robertson, Bernard and G. A. Vignaux (1993) “Probability—The Logic of the Law”, *Oxford Journal of Legal Studies* 13: 457-478.

Rosenkrantz, Roger (1977) *Inference, Method and Decision*. Reidel: Dordrecht
——“Bayesian Confirmation: Paradise Regained”, *British Journal for the Philosophy of Science* 45: 467-476.

Rousseau, Jean-Jacques (1762) *du Contract Social*, in *Social Contract: Essays by Locke, Hume, and Rousseau*, ed. Sir E. Barker. Oxford: Oxford University Press.

Salmon, Wesley (1975) “Confirmation and Relevance”, in *Minnesota Studies in the Philosophy of Science*, vol. 6, eds. M. Grover and R. Anderson. Minneapolis: University of Minnesota Press. pp. 3-36

Schlesinger, George (1995) “Measuring Degrees of Confirmation”, *Analysis* 55: 208-212.

Schum, D. (1994) *The Evidential Foundations of Probabilistic Reasoning*. New York: JohnWiley & Sons.

Seidenfeld, Teddy (1979) “Why I Am Not An Objective Bayesian: Some Reflections Prompted by Rosenkrantz”, *Theory and Decision* 11: 413-440.

Shaviro, Daniel (1989) “Statistical-Probability Evidence and the Appearance of Justice”, *Harvard Law Review* 103: 530-554.

Silins, Nicholas (forthcoming) “Basic Justification and the Moorean Response to the Skeptic”, *Oxford Studies in Epistemology* vol. II, eds. T. Gendler and J. Hawthorne. Oxford: Oxford University Press.

Sobel, Elliot (1994) “No Model, No Inference: A Bayesian Primer on the Grue Problem”, in *Grue! The New Riddle of Induction*, ed. D. Stalker. Chicago, IL: Open Court.

Sosa, Ernest (1999) “How to Defeat Opposition to Moore”, *Philosophical Perspectives* 13: 141-154.

Skyrms, Brian (1986) *Choice and Chance*, 3rd edition. Belmont, CA: Wadsworth.

Spitz, Elaine (1984) *Majority Rule*. Chatham, NJ: Chatham Publishing House.

Tillers, Peter and Jonathan Gottfried (2006) “Case Comment—*United States v. Copeland*, 369 F. Supp. 2d 275 (E.D.N.Y. 2005): A Collateral Attack on the Legal Maxim That Proof Beyond A Reasonable Doubt Is Unquantifiable?” *Law, Probability, and Risk* 5: 135-157.

Tribe, Laurence (1971a) “Trial by Mathematics: Precision and Ritual in the Legal Process”, *Harvard Law Review* 84: 1329-1393.

—(1971b) “A Further Critique of Mathematical Proof”, *Harvard Law Review* 84: 1810-1820.

Waldron, Jeremy (1999) *The Dignity of Legislation*. Cambridge: Cambridge University Press.

Weatherson, Brian (2005) “Can We Do Without Pragmatic Encroachment?”, *Philosophical Perspectives* 19: 417-443.

——(2007) “The Bayesian and the Dogmatist”, *Proceedings of the Aristotelian Society* 107: 169-85.

Weinstein, Jack B. and Ian Dewsbury (2006) “Comment on the Meaning of ‘Proof Beyond a Reasonable Doubt’” *Law, Probability, and Risk* 5: 167-173.

Weirich, Paul (1986) “Rousseau on Proportional Majority Rule”, *Philosophy and Phenomenological Research* 47: 111-126.

Weisberg, Jonathan (forthcoming) “Varieties of Bayesianism”, in *Handbook of the History of Logic*, vol. 10, eds. D. Gabbay, S. Hartmann, and J. Woods. Elsevier.

White, Roger (2006) “Problems for Dogmatism”, *Philosophical Studies* 131: 525-557.

Wheeler, Gregory (forthcoming) “A Review of the Lottery Paradox”, in *The Place of Probability in Science*, eds. G. Wheeler and Harper. London: King’s College Publications.

Wigmore, J. H. (1981) *Evidence in Trials at Common Law*, vol 9. Boston, MA: Little, Brown & Co.

Williamson, Timothy (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, Jon (2006) “Philosophies of Probability: Objective Bayesianism and its Challenges”, in *Handbook of the Philosophy of Mathematics*, Volume 4 of the *Handbook of the Philosophy of Science*, ed. A. Irvine. Elsevier.

Wittgenstein, Ludwig (1969) *On Certainty*, eds. G. E. M. Anscombe and G. H. von Wright, (tr.) D. Paul and G. E. M. Anscombe. Oxford: Blackwell.

Vogel, Jonathan (1999) "The New Relevant Alternatives Theory", *Philosophical Perspectives* 13: 155-180.