



ARTICLE

Reasonable doubt and reasonable priors

Yuval Abrams 

Michigan State University College of Law, East Lansing, MI, USA
Email: yabrams@law.msu.edu

(Received 5 December 2023; revised 2 August 2024; accepted 12 August 2024)

Abstract

What is guilt beyond a reasonable doubt (BARD) for a Bayesian? Is thinking of BARD in terms of probabilities a nonstarter? I propose an account of BARD compatible with Subjective Bayesianism that rejects the view that BARD is met by a threshold probability. BARD is a judgment, not merely about the credal state the factfinder endorses as her own (i.e. not merely as one's own credence in guilt), but as about alternative possible credences, specifically those the factfinder does not endorse, but finds reasonable. To this end, I employ a Bayesian framework, expounded by Lange (1999), that permits revision of past prior probability assignments. Such a framework presupposes a point of view free from one's prior from which a prior is judged. A trier-of-fact asks whether doubt persists among any reasonable starting point one might take; if it does, acquittal is warranted.

Keywords: Subjective Bayesianism; conditionalization; problem of the priors; legal epistemology; legal proof; reasonable doubt; legal probabilism; proof paradox; probability thresholds

1. Introduction

A finding of guilt in criminal trials requires proof beyond a reasonable doubt (BARD). This gold standard of proof, the law's strictest, is widely known but poorly understood. Attempts to make the standard more precise have met with little success. Recently, attempts to replace it have been proposed.¹

An underlying issue in clarifying legal proof – both generally and for BARD, in particular – concerns the extent to which it is best understood in terms of mathematical probabilities. Can the strength of the evidence offered at trial be quantified? Is a justified finding of guilt supposed to be reflected by this quantity? Is there a threshold probability, above which guilt is proven BARD?

I offer an account of BARD, compatible with Subjective Bayesianism. The account does not reduce BARD to a threshold probability of guilt: critics of that reduction, I believe, are correct. Nor will I defend the view that naked statistical evidence suffices for conviction, a view that plausibly follows from the acceptance of threshold probabilities.

What I offer instead is an explication of BARD in Bayesian terms. The explication treats BARD as a judgment, not merely about the credal state the factfinder endorses as her own (i.e. not merely as one's own credence in guilt), but as a judgment about

¹Laudan (2006).

alternative possible credences; specifically those the factfinder does not actually endorse, but still finds reasonable. To this end, I employ a Bayesian framework, expounded by Lange (1999), that permits revision of past prior probability assignments. In determining whether the evidence proves guilt beyond reasonable doubt, a trier-of-fact asks not merely about her own credence, but whether sufficient doubt persists among any reasonable starting point one might have taken; if it does, acquittal is warranted.

The proposed understanding of BARD also aids in the elucidation of two features of proof in criminal trial, both puzzling in Subjective Bayesian terms: the presumption of innocence and the dictum that factfinders must determine guilt exclusively on the evidence formally presented at trial. The account, if successful, will not only supply the correct account of these three principles in Bayesian terms but will vindicate Bayesian or probabilistic accounts of evidence at trial generally.

Furthermore, beyond criminal legal proof, the account sheds light on reasonable doubt more generally and on the norms governing the doxastic attitude of doubt, which are of independent interest in many facets of epistemology and inquiry. BARD, in other words, is not just a legally constructed notion; the law is responding to a feature of our epistemic life.

While the account is cashed out in Bayesian terms, aspects should be exportable to non-Bayesian or even non-probabilistic accounts as well. The account treats a finding of guilt as beyond reasonable doubt when the finder of fact can state that the assumptions one would have to make to defeat a conclusion of guilt are unreasonable ones to have made. It is not enough that the factfinder herself doesn't share those assumptions, or that she finds those assumptions improbable. There must be something sufficiently defective about these assumptions such that making these assumptions would be unreasonable; otherwise, reasonable doubt is not defeated, and an acquittal is warranted.

2. Reasonable doubt

Historically, BARD grew out of the concern that a false judgment of guilt would result in damnation of the judge or factfinder (*F*), for visiting wrong upon an innocent. To remedy this, the idea was that if the factfinder were *morally certain*, she would be absolved of moral culpability for an erroneous judgment to convict.² Massachusetts' jury charge is typical.³

What is proof beyond a reasonable doubt? The term is often used and probably pretty well understood, though it is not easily defined. Proof beyond a reasonable doubt does not mean proof beyond all possible doubt, for everything in the lives of human beings is open to some possible or imaginary doubt. A charge is proved beyond a reasonable doubt if, after you have compared and considered all of the evidence, you have in your minds an abiding conviction, to a moral certainty, that the charge is true. When we refer to moral certainty, we mean the highest degree of certainty possible in matters relating to human affairs — based solely on the evidence that has been put before you in this case.

I have told you that every person is presumed to be innocent until he or she is proved guilty, and that the burden of proof is on the prosecutor. If you evaluate all

²See, e.g., Langbein (2003) and Whitman (2016). The term “moral certainty” (*certitudo moralis*) seems to have been introduced by Jean Gerson, chancellor of the University of Paris (Franklin, 2001 at 69).

³From *Commonwealth v. Russell*, 470 Mass. 464 (2015), which was actually intended to clarify the former jury instruction. The situation elsewhere is very similar. See, e.g., *Miller v Minister of Pensions 2* All E.R. 372 (1947) (Denning, J.).

the evidence and you still have a reasonable doubt remaining, the defendant is entitled to the benefit of that doubt and must be acquitted.

It is not enough . . . to establish a probability, even a strong probability, that the defendant is more likely to be guilty than not guilty. That is not enough. Instead, the evidence must convince you of the defendant's guilt to a reasonable and moral certainty; a certainty that convinces your understanding and satisfies your reason and judgment as jurors who are sworn to act conscientiously on the evidence.

This is what we mean by proof beyond a reasonable doubt.⁴

The verbal formulae used by courts, signifying nebulous standards spelled out in arcane ways is a serious problem. On the other hand, the general BARD standard has great intuitive and moral appeal. It would be a mistake to give up on it, but we must be clearer about what it requires.

As a first gesture, it is clear that BARD or moral certainty is a weaker standard than absolute certainty. BARD, whatever it is, requires less than a demonstrable proof of guilt. On the other hand, it must be stronger than mere suspicion or a hunch. In particular, we can calibrate BARD by comparing it to alternative standards the law employs, such as the *preponderance of evidence* standard (also known as the “balance of probabilities” or “more likely than not”) employed in civil cases, or *clear and convincing evidence* employed in some contexts in American law. BARD must be a higher degree of certainty or confidence than either of those, but lower than demonstrable proof.

Speaking of BARD in terms of degrees of certainty lends itself to a quantitative interpretation. One tempting path to doing this is to think of proof in terms of probabilities. This is the path recommended by legal probabilists.

3. Legal Probabilism

Legal probabilism⁵ uses probability theory to “analyze, model and improve the evaluation of evidence and the process of decision-making in trial proceedings.” (Urbaniak & Di Bello, 2021).

Of particular interest is the Bayesian interpretation of probabilism, which works with *subjective* probabilities. Bayesians represent partial beliefs or degrees of belief⁶ as credences obeying the probability calculus. This means, at the very least, that these

⁴The federal jury charge in the United States, which deals in real possibilities (as quoted in *Victor v. Nebraska*, 511 U.S. 1 (1994)), is similarly opaque:

[T]he government has the burden of proving the defendant guilty beyond a reasonable doubt. Some of you may have served as jurors in civil cases, where you were told that it is only necessary to prove that a fact is more likely true than not true. In criminal cases, the government's proof must be more powerful than that. It must be beyond a reasonable doubt.

Proof BARD is proof that leaves you firmly convinced of the defendant's guilt. There are very few things in this world that we know with absolute certainty, and in criminal cases, the law does not require proof that overcomes every possible doubt. If, based on your consideration of the evidence, you are firmly convinced that the defendant is guilty of the crime charged, you must find him guilty. If on the other hand, you think there is a real possibility that he is not guilty, you must give him the benefit of the doubt and find him not guilty.

⁵The origins of this term are in Haack (2014), who is unsympathetic to the doctrine.

⁶The debate as to whether credences are partial beliefs or degrees of belief will not concern us here.

credences obey the Kolmogorov axioms. On the extreme subjective interpretation of Bayesianism, meeting this condition suffices for *rational belief*.⁷

Under Orthodox Subjectivist Bayesianism,⁸ agents begin with a prior probability;⁹ the only constraint is probabilistic coherence according to the axioms. As the agent learns new information, she updates her beliefs via the rule of conditionalization, which updates her new credence in a hypothesis (H) to what her prior probability accorded that hypothesis, conditional upon the learning of this new evidence (E): $\text{Pr}_{\text{New}}(H) = \text{Pr}_{\text{Old}}(H|E)$. This can be calculated using Bayes' theorem:

$$\text{Pr}(H|E) = \text{Pr}(H)\text{Pr}(E|H)/\text{Pr}(E)$$

If rationality merely requires obeying the axioms, as subjectivists claim, these requirements are permissive; they admit any possible states satisfying the axioms. If conditionalization is required for updating,¹⁰ subsequent posterior probabilities will be thereby constrained, but they are still a function of what the prior probability is.

Are there rational constraints on the prior probability distribution beyond adherence to the axioms? Subjectivists say “no.” Any prior satisfying the axioms is rational. Objective Bayesians, on the other hand, claim that there is a unique rational prior probability function for an agent to have.¹¹ Whether such a prior exists and how it is determined raises

⁷Adherence to the probability axioms (sometimes referred to as “Probabilism”) is motivated by several traditional arguments. The first, and the most famous, are the Dutch Book arguments (Ramsey 1926; Hajek 2008). A second set of arguments are Representation Theorems. If an agent's preferences satisfy certain consistency constraints, they can be represented as resulting from beliefs that satisfy the probability calculus (for the constraints on cardinal Utility, see Von Neumann and Morgenstern 1944). A third strand of arguments, which attempts to justify probabilism on purely epistemic, rather than pragmatic grounds, argues that credences that violate the probability axioms are dominated (in terms of their accuracy score) by probabilistic credences. Accuracy is measured by a Brier score. The accuracy of $\text{Cr}(p)$ is the distance between its value and the truth value of p (so if p is true, and the agent's degree of belief in p is 0.3, the accuracy score is 0.7). For the accuracy argument, see Joyce (1998), and Joyce (2009) as well as Pettigrew (2016). For an attempt at a non-pragmatic Dutch Book argument, see Christensen (2001).

⁸“Orthodox Bayesianism” is a general term. There is no high church of Bayesianism. I.J. Good (1971) famously claims that there are at least 46656 varieties of Bayesianism. For our purposes, I mean subjective Bayesians who constrain the prior probability function minimally, at most (for instance with certain credence-change bridge principles [see Lewis 1980] or Reflection principles governing future credence [see van Fraassen 1984]). Furthermore, Orthodox Bayesians will normally accept a constraint on updating via a rule of conditionalization. Other views typical of Orthodox Bayesians include the norm that you are certain of the evidence that you learn (such that when one's probability function is updated, the proposition that one updates on (E) receives probability 1). This is in contrast, for example, with a rule such as probability kinematics, also known as Jeffrey Conditionalization (Jeffrey 1965), allowing for multiple possibilities if the agent is uncertain of what he has learned. Consequently, for the Orthodox Bayesian, once a proposition is learned, it can no longer be “unlearned” (if $\text{Pr}(p) = 1$, there is no further proposition, conditional upon which, it can be lowered). This will feature below when discussing the incorrigibility of the prior for standard Orthodox Bayesians. For further discussion of Bayesian orthodoxy, see Titelbaum (2013) and Smith (2013)).

⁹Prior probability is used in two related ways. On the one hand the prior is the agent's starting point – it is a probability function that reflects the agent's state at the beginning of inquiry. On the other hand, at any event of updating, the prior can refer to the agent's probability before learning the new information that is to be assimilated into the agent's belief state.

¹⁰Whether conditionalization is required is contested. Strictly speaking, an agent can remain probabilistically coherent without updating via conditionalization. The arguments for conditionalization in the literature include a diachronic Dutch Book (see Paul Teller (1976), Lewis (1999), criticized in van Fraassen (1989) and Christensen (1991)), a calibration argument (Lange 1999), and arguments from maximization of Epistemic Utility (Greaves and Wallace 2006; Leitgeb & Pettigrew 2010).

¹¹Objective Bayesians ideally aim at capturing an indifference principle, along the lines of Laplace (1814), in which an agent ascribes equiprobability to outcomes over which he has no evidence. Indifference principles run into serious trouble with the Bertrand paradoxes, which demonstrate that the division of the possibility space is

serious difficulties. Regardless, jurors don't actually begin with such a prior and certainly do not all begin the trial with the same credences. On the other hand, Subjective Bayesians inherit the *problem of the priors*. The prior probability function is purely subjective: it is not subject to any rational constraint except consistency with the axioms.

If F updates using conditionalization, where F ends up after hearing the evidence is a function of where she began. Two triers-of-fact with different assumptions – and starting points – at the outset of the trial could reach different conclusions about guilt; two triers-of-fact might even rationally disagree whether the very same evidence weighs for or against H . On subjectivist views, this divergence can be entirely rational, as long as each factfinder responds to the new evidence consistent with the axioms.¹²

A Bayesian trier of fact, will, at the conclusion of trial, upon hearing the evidence, assign a probability to the question at hand: of liability (in a civil case) and of guilt (in a criminal one). A natural extension of the Bayesian model would then render the preponderance standard as met when $\Pr(\text{Liability}|\text{Evidence}) > 0.5$. A further extension might render BARD as met when $\Pr(\text{Guilt}|\text{Evidence}) > t$ where t is some threshold ($0.5 < t \leq 1$). I'll call this last extension the *Threshold view*.

Much of the debate pertaining to Bayesian analyses of factfinding at trial traces to an influential paper by Laurence Tribe on "trial by mathematics."¹³ The debate is most heated around the use of statistics to determine guilt.¹⁴ Statistics, of course, are relevant for many inferences in trial, and nobody denies that a factfinder should be attentive to them, but naked statistical evidence conjoined with the Threshold view, delivers perverse results, no matter how high we adjust the threshold (unless we require 100%).

Consider a lottery-style case, *Gatecrasher*.¹⁵ In *Gatecrasher*, we stipulate that $(10^4 - 1)$ people at a show crashed the gate, with 10^4 in attendance, such that $(10^4 - 1)$ committed a crime; exactly one ticket was sold and subsequently collected; we know that $\Pr(\neg\text{Guilt}) = 10^{-4}$ for any given defendant, quite plausibly beyond the threshold. *Gatecrasher* can be raised as an objection to threshold views more generally. For any threshold shy of 1, a lottery/gatecrasher scenario can be constructed where the threshold is met, and yet a verdict of guilt would not be warranted.¹⁶

description dependent. See discussion in Keynes (1921) and van Fraassen (1989). Laplace and Keynes are often categorized as adherents of the classical and logical interpretations of probability. Objective Bayesians differ from these camps, in interpreting probability as modeling a belief state. Contemporary attempts to develop the Objective Bayesian view include Jaynes (2003), White (2005), and Williamson (2010). Titelbaum (2013: ch. 5) criticizes the division between objective and subjective Bayesians as vague and misleading (2013 at 117). For a recent defense of the principle of indifference, see Pettigrew (2014).

¹²Under certain constraints, the differences between priors "washes out" and converge (Gaifman & Snir (1982)).

¹³Tribe (1971). Tribe's paper came after several seminal papers defending versions of legal probabilism, including Kaplan (1968), Cullison (1969), and Finkelstein and Fairley (1970). For a rich extensive discussion of this history, see Di Bello (2013). For a recent defense of probabilism in the context of this debate, see Fenton & Lagnado (2021).

¹⁴Statistics were used controversially, and, as it turned out, erroneously in *People v Collins*, 438 P 2d 33 (68 Cal 2d 319 1968).

¹⁵*Gatecrasher* is based on similar cases. See Cohen (1977), Nesson (1979: 1192–1193), Smith (2018), and Moss (2018: 204–205). See also the discussion of "naked" statistical evidence in Kaye (1980). A more recent example discussed in the literature is due to Redmayne (2008). For an extensive discussion of lottery cases, see Hawthorne (2004).

¹⁶Nelkin (2000) following Harman (1968) took it to be rationally impossible to believe claims on the basis of purely statistical evidence (such as that one's ticket in a fair lottery is a loser) since an appropriate causal nexus forming the abductive basis for such an inference is absent. That and other attempts (see, e.g., Pollock (1990)) sounding in the impossibility of forming impermissible rational belief has been thought wanting; see Douven & Williamson (2006). Others – such as Lin & Kelly (2012) and Leitgeb (2017) – relied on contextual limitation of possible propositions to delineate the contexts in which an agent can rationally believe of a ticket in a fair lottery

3.1. What legal probabilism requires

We should, following (Di Bello 2013), distinguish between two claims:¹⁷

QUANTIFICATION: a probabilistic quantification of the case for defendant's guilt can be given through an appropriate weighing of all the fallible evidence available (that is, of all the evidence against, and of all the evidence in defense of, the accused);

THRESHOLD: an appropriately high threshold guilt probability should be the decision criterion for criminal convictions.

In defending Bayesianism, I endorse Quantification but reject Threshold. I will also resist the sufficiency of naked statistical evidence for conviction, although I only devote a little bit of space to it here.

But while I reject Threshold, it will not do to simply state that one should assign a credence of guilt according to Bayes' rule. For probabilism to be illuminating, it should go beyond how evidence is weighted, to how, ultimately, that weighting determines a verdict of guilt.

The basic commonsense story of the factfinder's job is that she begins with a presumption of innocence, hears the admissible evidence at trial, and forms a judgment based on that evidence as to whether guilt has been proven. A judgment that guilt has been proven is not exactly the same as a judgment that the defendant is guilty. It is a judgment that the weight of the evidence is sufficient to sustain a judgment of guilt. A judgment of guilt requires BARD. The judgment of the evidence is an evidential probability. A probabilist will capture it with $\Pr(\text{Guilt}|\text{Evidence})$. If our account of factfinding at trial is a Subjective Bayesian one, using subjective probabilities, the story presumably is as follows: the factfinder begins the trial with some prior probability of guilt, she then learns new information (the evidence at trial) and subsequently updates her credences via Bayes' rule, arriving at a posterior probability of guilt.

The above is a tempting story of what a Bayesian factfinder does. But it raises three challenges for the Bayesian:

1. **Presumption of Innocence.** What is the presumption of innocence for a Bayesian? It is not obvious what a prior that respects this presumption is. Some have suggested a very low probability of guilt (how low?), and others that the probability of guilt is $1/n$, where n denotes the other possible suspects (which creates a reference class problem).¹⁸ But this will not suffice, since adjusting one's probabilities in this manner to accommodate the presumption has spillover effects. One cannot just isolate credence in one particular proposition and adjust it. All other evidentially relevant propositions must be adjusted as well. This can be done, of course, via conditionalization, but the resulting "prior" will not be the epistemic state of the actual factfinder. An account that adjusts the prior in this

that it will lose. However, this contextual limitation is highly controversial and possibly ad hoc; see Douven & Rott (2018). Since epistemic contextualism is also a highly controversial thesis – see, e.g., Adler (2012) – lottery cases are very plausibly still an outstanding philosophical problem for many approaches to rational belief and knowledge.

¹⁷Some writers take "probabilism" to refer to the conjunction of these two claims (e.g. Di Bello, 2013). Others take *Threshold* to be the defining feature of probabilism (Ross, 2024). This is unfortunate, in my view, because the two hypotheses can be distinguished. Denial of *Threshold* does not imply that legal proof resists quantification or that it is outside of probability theory.

¹⁸For the $1/n$ principle, see Dahlman et al. (2021). For the reference class criticism, see Allen & Pardo (2007, 122).

way has already strayed from the standard subjectivist story.¹⁹ Additionally, these adjustments require “unlearning” propositions that the subject has already learned and assigned probability 1. This cannot be undone by conditionalization.

2. **Restriction to Admitted Evidence.** How are we to understand the injunction to conditionalize only on the evidence at trial? Real factfinders learn and infer many things. In what sense is the probability of guilt a *subjective* probability, if it neither emanates from the subject’s own prior (due to the presumption of innocence) nor conditionalizes exclusively on the information the subject actually learns (due to the restriction to evidence at trial)? One could formalize this as an evidential probability of E, conditional on some artificially constructed prior probability, but this would neither correspond to the factfinder’s own posterior probability of guilt nor would it truly be free of the subject’s priors, since, that conditional probability itself will be informed by other features of the subject’s priors. Unless, that is, the prior is entirely artificially constructed, at which point, the account is no longer one of subjective probability at all.
3. **BARD.** The threshold picture is too simplistic, even if we ignore problems 1 and 2. But these issues play off each other as well. Ultimately, the question is not simply what credence the factfinder has in guilt. It is a question of what credence an ideal fact finder who began with the presumption of innocence and only updated on the evidence would have in guilt. Even if such a credence could be established, the BARD question remains: what facts about a credal state or otherwise determine whether BARD has been achieved?

3.2. Thresholds

It is tempting to assign a particular value of credence as a threshold for BARD. Substantial disagreement exists as to what the threshold probability actually is.²⁰ While it must clearly be above 50%, no obvious reason favors selection of any particular value in the closed interval between 50% and 100%.

Suppose the threshold was set at 100%. This would avert the charge of arbitrariness. But would arguably be too strong, since guilt is rarely (if ever) established at that level. Furthermore, it would be contrary to the spirit of the BARD/moral certainty standard, which was to allow for a lesser standard than certainty itself. Arguably, it would also be too weak. Couldn’t a factfinder harbor no doubts but still find it reasonable for someone else to doubt? I’ll come back to this objection, which characterizes all threshold accounts of BARD, later.

Suppose instead we settled on a threshold less than 100%, but higher than 50%. This causes several problems. For one, we get the gatecrasher/lottery problem issue, discussed

¹⁹There are also normative ramifications of the presumption, involving the defendant’s rights. But these are not at issue. We are concerned merely with the epistemic presumption of innocence.

²⁰Distinguished jurists, like Judge Weinstein, have called for $\Pr(\textit{Guilt}) \geq 95\%$ in criminal cases, see Weinstein & Dewsbury (2006). Weinstein’s survey of his colleagues revealed they thought the threshold $\Pr(\textit{Guilt})$ when BARD was the applicable standard could be set elsewhere too. Out of ten judges, one reported that he could not assign a number (for reasons that are unclear), while the other nine reported: 76%, 80%, 85% (four judges), 90% (two judges), and 95% (Weinstein himself). See *United States v. Fatico*, 458 F. Supp. 388, 410 (E.D.N.Y. 1978).

A more comprehensive survey circulated to all federal judges in the United States by McCauliff (1982: 1325) had similar results although with one further major absurdity: one judge felt there was a 50% threshold, one judge a 60% threshold, one judge a 70% threshold, eight judges liked 75%, 14 judges went for 80%, a score of them for 85%, 56 for 90%, one for 92%, one for 93%, one for 94%, 31 thought 95% the threshold for BARD, one preferred 97%, six judges opted for 98%, eight for 99%, and an 21 judges felt the threshold probability is 100%.

earlier.²¹ No threshold, shy of 100% avoids that problem. Secondly, the threshold seems arbitrary: why, e.g. 0.9 and not 0.91?

The arbitrariness concern is, perhaps, not hopeless. A similar debate surrounds the analysis of full belief for Bayesian and the relationship between full beliefs and credence. Not all Bayesians have need for a category of “full belief.”²² Among those that do are those that identify belief with probability 1.²³ Others advocate versions of the Lockean Thesis,²⁴ identifying belief as some probability in the range of 0.5–1. The Lockean Thesis can be interpreted either descriptively (what credence constitutes belief?) or normatively (what credence justifies belief?); as an objective standard (setting the same threshold for all believers, sometimes referred to as the *Strong Lockean Thesis*) or a subjective standard (determining what threshold a particular believer believes at, referred to as the *Weak Lockean Thesis*), possibly incorporating contextualist or pragmatic considerations as well.²⁵ I don’t intend to delve into the debates concerning the Lockean Thesis, except to note that many of the issues that arise with threshold arise there as well. Perhaps it’s plausible to locate a threshold for belief.²⁶

What about doubt? Can doubt be identified with a threshold? Doubt might be scalar: whatever credence the agent has in the negation of the hypothesis considered (doubt of guilt = credence of innocence). Alternatively, it could be categorical: doubt is any non-zero credence in innocence or credence in innocence above a particular threshold. Alternatively, doubt is a doxastic attitude, distinct from credence. Still, if this attitude connects to a credal state, a threshold would still be possible, involving bridge principles (descriptive or normative) between one’s credence in $\sim p$ and doubt that p . Perhaps there is no threshold for doubt, but, at the very least, if one disbelieves the hypothesis in question, then one at least doubts it.

Reasonable doubt, on the other hand, is a normative judgment on doubting. The norms that govern BARD are the norms beyond which doubt is no longer warranted. Just as with the normative Lockean principle for belief, a normative principle for doubt cashed out exclusively in terms of a threshold looks unpromising. A better understanding of BARD is one that goes beyond one’s own credence.

4. What is BARD for a Bayesian?

My suggestion is that judging doubt as beyond reasonable requires going beyond the actual credence you have in guilt (which may be a necessary condition but not a sufficient one for conviction), referring to other credences it would be reasonable to have. This requires a slightly different perspective on Subjective Bayesianism; one that involves stepping outside of your own credences. This stepping outside is limited, in that it retains subjectivism’s insight that you are always judging probabilities from your own

²¹For a developed criticism of this line, applying the norms of blame to lottery cases, see Buchak (2014) and Staffel (2015).

²²Jeffrey (1970); Kaplan (1996).

²³What Roorda (1995) calls the received view. Advocates of this view include Tang (2009), Wedgwood (2012), Clarke (2013), Greco (2015), and Dodd (2016).

²⁴Foley (1993), Bouvens and Hawthorne (1999), Christensen (2004), Sturgeon (2008), Lin and Kelly (2013), Locke (2014), Easwaran (2016), Fitelson & Shear (2019), Dorst (2019), Thorn (2020). Critical of the Lockean Thesis: Friedman (2013), Buchak (2014), Staffel (2015), Smith (2010, 2016), Kelp (2017), Jackson (2019).

²⁵Many of defenses of the Lockean thesis include decision theoretic or contextualist elements. This will not do for an explication of BARD, however, since the standard of proof must be uniform across cases and defendants.

²⁶A recent proposal by Leitgeb (2017) identifies belief with the property of stability. This is perhaps a Lockean view, but it is not a threshold view.

subjective point of view. But it affords you the ability to look at alternative credal functions as reasonable or relevant alternatives.

Rejecting thresholds as an account of BARD, need not imply rejecting them as an account of *doubt*. It is one thing to *doubt* a claim, it is another to judge doubt as *unreasonable*. It is the latter that BARD is after. Thresholds are inadequate at capturing BARD and are at best, a necessary, but not a sufficient condition for BARD. Moreover, the problem of the priors exacerbates that inadequacy: credence in guilt (or innocence) is a function of one's prior, i.e., of where one began.

Perhaps *doubt* can be defined adequately in terms of the degree of belief in innocence. Our question is whether such doubt, however defined, is *reasonable*. For this, *F* doesn't determine merely whether *F* doubts, which is a function of where *F* began; instead, *F* must ask about other starting points as well, thinking about which would be reasonable and which unreasonable. My suggestion is that, for BARD to obtain, doubt must be ruled out from *all* reasonable starting points. A judgment of guilt consists in the judgment that all such reasonable starting points lead beyond doubt.

Except for cases of deductive certainty, there will *always* be priors, starting from which, and conditional upon evidence adduced at trial, doubt would be warranted. In other words, on any body of evidence, there is a coherent prior probability (consistent with the axioms) for which, conditional on the evidence, $\Pr(\text{guilt})$ is low. If BARD meant ruling out such rational doubt, one could never convict, since there will always have been at least one other rational possible starting point that would have led to doubt. BARD cannot mean the absence of any *rational* doubt, at least not in a Subjective Bayesian framework.

4.1. Reasonable doubt v. rational doubt

What we require is a distinction between *rational doubt* and *reasonable doubt*.²⁷ Rational doubt just is doubt that would be rationally permissible per the axioms to have. In Orthodox Bayesianism, this just means doubt that is consistent with updating on some prior probability. This is too easy. BARD does not require ruling out rational doubt, but it does require ruling out reasonable doubt. This is more demanding than ascertaining that $\Pr(\sim\text{Guilt})$ is low but less demanding than ascertaining that doubt is rationally impermissible: it requires that *F* judge the (rationally possible) alternative conclusion (that $\Pr(\text{Guilt})$ is low) as *unreasonable*. In what can this unreasonableness consist? It cannot be in the updating itself, since this follows mathematically via conditionalization. There is nothing unreasonable about the alternative conclusion given the alternative prior. Rather, the reasonableness must be in the prior itself. It is this, rationally permissible prior that the finder of fact needs to be able to rule out as an unreasonable one to have had.

Accordingly, the BARD test traces back to all possible starting points (priors) and asks: *of all the possible starting points that lead to doubt, is there at least one that is reasonable?* If there is, then guilt is not proven, since doubt is reasonable by the factfinder's own lights; even if the factfinder herself doesn't doubt the defendant's guilt.

On the view being put forth, some alternative priors are reasonable alternatives, whereas others are unreasonable (even though they are rationally permissible). Only if doubt emanates from reasonable priors is doubt reasonable; if it only results from unreasonable priors, it is not reasonable doubt.

²⁷The rational/reasonable distinction has an illustrious history in moral philosophy. In that context too, the reasonable is a subset of the rational. Rationality means largely the same thing in that context as it does here (a coherence or consistency requirement), but reasonableness is typically defined in terms of moral constraints or mutual justifiability constraints over and above pure consistency. See Rawls (1993), Scanlon (1982).

4.2. Reasonable credence functions and conditionalization

In developing this model, I am employing an account of conditionalization due to Lange (1999). Lange views the requirements of Bayesianism as steps in a justificatory argument from a set of reasonable priors to an argued-for conclusion, rather than as a description of how an agent must update her views given some actual past credal state. Lange shows how an agent can step outside of her particular priors and ask if and how a belief state that would have been reasonable at time t_1 could have rationally evolved, given the evidence, into a different credal state at t_2 , regardless of whether the agent possessed that credal state at t_1 . What matters is the calibration of credal states across time relative to the available evidence at that time. On this view, an agent can recognize at some later time that her earlier credal state was unreasonable and adjust it.

Traditionally, Bayesianism represents a subject's epistemic state via her credence function. In Orthodox Subjective Bayesianism, credences are determined by the prior probability, a historic, prior belief state. Subsequent belief states are permutations of this function via conditionalization. Subjective Bayesians afford the subject no recourse to any vantage point outside her own subjective probability function by which to judge whether her beliefs are reasonable: as long as they are probabilistically coherent and obey the axioms, they are rational. This means that subjectivists allow a degree of arbitrariness as to what the subject believes: the rational thing to believe at t_2 depends on what the subject believed at t_1 , which depends on the prior probability, of which there is no unique rational answer as to what the subject should have believed. Each step is just the product of conditionalization, ultimately tracing back to the subject's arational determined prior.

All subsequent beliefs that the subject has are determined by this prior. This creates trouble as an account of factfinding at trial, both because it suggests a high degree of dependence on the idiosyncrasies of the particular factfinder, as well as because it fails to account for the factfinder's duty to be impartial and afford defendants the presumption of innocence.²⁸ A factfinder who begins with a highly biased prior against the defendant could rationally end up with a high probability of guilt through no epistemic fault of her own.²⁹

The view I am proposing does not treat the prior as an incorrigible starting point to which the subject is bound for all time. It does not deny the subjectivist insight that one's credences at any moment rationally involve a degree of subjectivity, and hence that rational agents might disagree. It does, however, assume that agents can rethink their own credences and adjust them. A prior, on this view, is not necessarily a historic actual subjective probability that the agent affirmed at a prior time; nor is it a randomly assigned credal state that is entirely free of judgment. Rather, it is a probability distribution that an agent *now* affirms as the correct probability distribution to have held at a previous time, given what was in fact known at that time.

Upon learning new evidence, an agent might reconsider her prior, just as, in non-probabilistic contexts, an agent might reconsider her premises upon discovering their unwelcome conclusion. It is possible that at t_1 , the agent actually believed $\Pr(H) = p1$ and $\Pr(H|E) = p2$, but upon learning E , instead of adjusting her credence to $p2$ (as orthodox subjectivists require), the agent may come to realize that she *should not have had those priors*. Thus, suppose the agent adjusts her posterior $\Pr(H)$ to $p3$. This is

²⁸Cf. Tribe (1971).

²⁹A related desideratum is to take into account the weight of the evidence, in addition to its balance. Balance is reflected in the probability, but weight measures the resilience of the evidence. See Keynes (1921) and also Joyce (2005), who charges Keynes with conflating weight and specificity of evidence. Lange's account is not an account of weight, but what it captures is arguably similar in terms of resilience. See Gardiner (2019) for a recent paper defending the importance of weight over balance, specifically in the context of legal proof.

permissible, so long as she simultaneously affirms that her prior probability $\Pr(H|E)$ *should have been* p_3 as well. The agent may adjust her prior or her posterior, but the adjustments must be in tandem, keeping her well-calibrated. Rationality, on this view, requires conditionalization, not because you are bound by your actual priors; but rather because affirming that something would have been rational constrains you to believe what follows from it conditional on the new evidence.

In other words, the requirements of probabilistic rationality involve determining what a reasonable thing to believe *would have been* and showing how, given such a starting point, where the agent ended up was just a stepwise process of updating on the evidence (via conditionalization) from what is now taken as a reasonable starting point. Such an agent is not a prisoner of her prior, since she can learn that her actual prior was, in fact, unreasonable and retroactively adjust it accordingly.

A similar process can occur between jurors when deliberating. Their deliberation is not restricted to how to conditionalize on their own priors, but on what the appropriate priors should have been.

More importantly, for our purposes, the ability to judge what a reasonable prior is involves a modal judgment, not only of the agent's actual prior but of the relevant reasonable alternatives to the agent's priors. When an agent judges an alternative prior as reasonable, she is still doing so by her own lights. But this judgment need not involve *endorsing* that prior. She still represents the world via her own credences, which are determined by her own prior (i.e., the prior she currently endorses, and which she judges now as *most reasonable*), but judging any prior as reasonable involves the ability to make (some) comparisons with priors other than her own. This is the sense of reasonableness we are after: not merely by the lights of one's own prior, but by the lights of other priors one deems as reasonable alternatives to one's own. If an agent can rule out *doubt* on all reasonable priors, she is *beyond reasonable doubt*.

What counts as a reasonable prior? It is not algorithmic like the constraints of rationality. The details can be filled in by how objectivist or subjectivist your favorite account of Bayesianism is. On the objective end of this spectrum, the range of reasonable priors is narrow, perhaps constrained by indifference principles; on the subjectivist end, we can say at least this: the reasonableness of the prior is determined by the subject herself. This is subjective, indeed, but no more so than the selection of a prior probability function in the first place. Unlike the single prior, however, it allows for a range of possibilities, an assumption that seems not unrealistic: Don't we all distinguish between those with whom we disagree and those whom we find unreasonable?³⁰ By the lights of one's own actual credences, all other priors are simply mistaken. The ability to consider alternative priors means that one has the ability to consider a prior from a vantage point outside of one's own actual specific commitments and consider whether such a prior is intelligible or acceptable.³¹ Absent this ability, there is no sense in which an agent can consider whether a particular prior is worthy of adoption.

Some alternative priors might seem wrong but reasonable; others are ruled out as not worthy of consideration. It is important to understand that a rejected but reasonable prior is different from a mere credence. I might have credence C in a particular proposition, leaving a remaining credence $(1-C)$ to express my doubt. This is not the same as judging an alternative credence as reasonable. I might doubt the testimony

³⁰This distinction has parallels in the discussion of Peer Disagreement. But I am not restricting reasonable disagreement to the standard epistemic peer.

³¹The idea that the agent's belief state might be represented by more than one credence function is also endorsed by those who advocate imprecise probabilities or mushy credences (Levi (1974), Van Fraassen (1984), Joyce (2005), Sturgeon (2008). For criticism, see White (2010), Elga (2010).

I hear in a case, while still judging someone who takes that testimony to be credible as reasonable; contrast this with the claim that I might see how a particular theory, if true, would undermine my confidence in guilt, but that lending credence to such a theory would not merely be mistaken but unreasonable. This involves judging a prior that gave high credence to this claim as itself unreasonable.

The reasonable doubt standard is one in which you look to those alternative credences that it would have been reasonable to have adopted, even though they are not your own, and ask whether any of them lead to doubt. It is not enough to remove doubt on your own credences – that would just be a threshold view – rather, you must do so along all credence functions within the set of reasonableness. When this happens, guilt is proven BARD.

The reasonable set is a proper subset of the rationally permissible set: the set of all consistent credence functions. BARD is a property of the *reasonable set as a whole*, rather than of a particular credence function. In other words, BARD does not supervene on your credences. Two agents might agree on all aspects of their credence function (particularly, they might agree on the probability of guilt) but still disagree on whether guilt is proven beyond reasonable doubt because they disagree about which alternative credences are reasonable.

This lack of direct dependence of BARD on probability is a feature of any view that rejects thresholds. It is brought out nicely in a recent paper by Moss (2022), which shares an important insight with the view presented here: namely that BARD requires ruling out relevant alternative hypotheses. In Moss' account, this is captured by a knowledge requirement. The present account takes no stand on whether knowledge itself is required, but shares with Moss the idea that it is the relevant salience of alternatives (which are always conceivable) that makes or breaks legal proof. Moss' account, as an externalist knowledge-based account, treats the context dependence of knowledge as the central guiding criterion. On my subjective Bayesian account, the relevance of an alternative is captured in the subjective assessment of the priors themselves and requires no externalist features.

For this account to work, all that is required is for the agent to be able to judge priors other than her own as reasonable alternatives. Presumably, this would also involve treating some alternatives as unreasonable as well. The sense of reasonableness here is by the agent's own lights. This goes beyond orthodox Subjective Bayesianism in granting coherence to the judgment that some alternatives are consistent but unreasonable, without committing to the Objectivist Bayesian claim that there are objective constraints on priors (although such objective constraints are consistent with the view).

On the other hand, if you like Moss' account of credences, according to which one can have knowledge about the probability of a claim (e.g. knowledge that the probability of guilt = x or is larger than x), that can work with this account as well. On such an account, the agent must look not only at her own credence in the defendant's guilt; she must be in a position to rule out alternative credences that would undermine her judgment. This means that she must look to alternative priors, on the basis of which, $\text{Pr}(\text{guilt})$ is low, and be in a position to rule them out. On Moss' account, this means she must *know* these priors are false. To defeat BARD, in Moss's view, it is not enough to know that there are internally consistent alternative priors. Rather, one must know that those priors are irrelevant and thus dismissible. The basic structure, according to which BARD is a matter of proof along a range of permissible credence functions, rather than merely following from the properties of one's own credence (even if justified), remains the same.

Suppose at the beginning of trial, F has high credence in medical testimony. This could be for various reasons: perhaps she has had good experience with doctors, or perhaps she is related to (or is) a doctor. A second factfinder might have low credence in such testimony. If the evidence at trial involves a doctor's testimony, when each

factfinder conditionalizes on this evidence, they will arrive at different posterior probabilities. On the standard subjectivist view, as long as the conditionalization was carried out correctly, both credences can be rational. End of story.

On the proposed view, however, *F*'s job does not end here. The question factfinders must ask is not merely whether their own credence of guilt is rationally attained, but whether a reasonable factfinder might have, on the very same evidence at trial, arrived at a different conclusion. For this, they must ask: what credences must one have had at the outset of the proceedings to rationally arrive at a different conclusion? This will yield a set of possible prior probabilities, each of which is rational. Among this set of priors, they must then ask: are any of these reasonable? In other words, even if the alternative prior probability distribution is not the factfinder's own, the account still assumes that a factfinder can distinguish between a reasonable and unreasonable alternative point of view: if the alternative prior needed would be one that involves slightly higher or lower credence in medical testimony, the answer would presumably be "yes": one could reasonably have lower or higher confidence in such testimony than one has. If this is so, that one reasonably believes in guilt or innocence based on such testimony does not place the matter beyond reasonable doubt, no matter how high the credence in guilt. Importantly, this judgment of a reasonable alternative is not captured in my current credences, or even in the weight of my evidence. I may have overwhelming evidence for a proposition and believe with very high credence, and yet still recognize the reasonableness of another starting point.

On the other hand, suppose that one's high credence in guilt or innocence would only be defeated by very high credence in a conspiracy theory or some alternative belief to which one not only attaches low credence but also judges that attaching a high credence would be wholly unreasonable. In such a case, while one recognizes that *rational* doubt is possible (because one cannot rule out the conspiracy theory), the conditions under which such doubt could be possible, i.e. conditional upon a prior one deems unreasonable, are such that one reasonably rules it out, and thus one can deem such doubt unreasonable.

4.3. Applied to specific accounts of doubt

Turning back to the various candidates for doubt discussed earlier, my account is neutral between these, as it is not meant as an analysis of doubt per se. But any account of doubt can be plugged back into BARD. This is easiest to see with threshold accounts. But, even without a threshold, the account can still shed light.

1. **Doubt as credence > 0 in innocence.** Doubt of guilt just means lack of certainty in guilt. But certainty of guilt, while necessary, is not sufficient for BARD. *F* might believe in guilt with credence 1, but nevertheless judge that a reasonable set of priors would have led to a credence lower than 1. This would be a very demanding standard.
2. **Doubt as a threshold between 0.5 and 1.** *F* must ascertain that all reasonable priors meet the threshold. This will hold regardless of whether the threshold is a strong/objective one or a weak/subjective one.
3. **No threshold for doubt.** If there is no threshold for doubt, the test will run on whatever else, other than a threshold, determines doubt. For example, if the additional factor is something akin to resilience,³² then resilience must be met in every relevant alternative. On the other hand, if the relation between doubt and credence is one of complete independence, the account is no longer Bayesian in

³²Logue (1997), Skyrms (2011), Di Bello (2013), Dahlman et al (2015), Jellema (2024).

any interesting sense. Notice that, on my account, credences still determine a finding of guilt, it's just that *F* looks beyond her own credences towards alternative credences she deems reasonable. Lacking even that, the credences lose much of their traction.

Still, there are lessons here for non-Bayesian accounts as well. Whatever it is that merits doubt must be met from *every* reasonable starting point. The injunction to go beyond your own free assumptions (assumptions you could have with reason modified) in determining whether something is reasonably doubted can be generalized.

I'll add that even on views that eschew a threshold for doubt, another possibility presents itself:

4. **BARD as reasonable belief in innocence.** In other words, in the examination of alternative credences, *F* asks whether, in any of them, belief in innocence is warranted, i.e., whether in any of them $\text{credence}(\text{not-guilt}) > 0.5$ (or any other range up to the Lockean threshold for belief). This is a far more permissive standard than one that rules out doubt, but, as a floor, it is difficult to resist, even for those who reject thresholds.

5. Applying the analysis to the other norms

What about the presumption of innocence? Notice that a satisfactory solution falls right out of our BARD standard: In order to satisfy BARD, proof must be established from all reasonable priors. This means that even if *F*'s own prior is too suspicious of the defendant, she must factor in priors that are not. So long as it is reasonable to presuppose such a prior, *F* must include it in the set, even if in her own credence she assigns it low probability. Even if such a set is not within the trier's reasonable set, we can stipulate it as a member of the set, for normative reasons, without compromising the subjective Bayesian nature of the calculation. *F* needn't abandon her priors, she must merely accept the restriction on the set that includes the presumption.

As for restriction to evidence at trial, that the actual trier of fact knows more than what was introduced at evidence no longer matters. We can exclude information not formally presented as evidence, because in order to establish BARD, the relevant question is not "what do I believe about the defendant's guilt?", it is whether all admissible priors (including those that lack my information) would have arrived at a judgment of guilt, conditional on the evidence presented. The subjectivist nature of this judgment is not in the updating, it is in the endorsement of priors as reasonable. Once the priors are established, the rest of the procedure is mechanical, conditional on the evidence. As long as we can determine what the admissible priors are, the rest of the procedure presents no special difficulties. The question is now properly framed as whether guilt was established on *all* of them, conditional on the evidence at trial.

6. A Note on Naked Statistical Evidence

As stated, the insufficiency of naked statistical evidence was presupposed. In fact, it formed part of the motivation for rejecting thresholds. Still, one might wonder whether my account excludes proof on the basis of naked statistical evidence after all. Couldn't we construct a gatecrasher/lottery case large enough to establish "doubt" on every suitable prior? Would such a case suggest that BARD can be established on the basis of naked statistical evidence alone?

I think not. While it is true that for any threshold there is a large enough sample of gatecrashers to place the probability that the defendant crashed the gate above that threshold, if we assume that each defendant was equally likely to crash the gate (i.e. if we assume something like an indifference principle), this is insufficient for BARD. BARD requires that for *every* reasonable prior, doubt can be eliminated. The factfinder in a naked statistical evidence case has nothing herself to go on beyond the statistics. But when contemplating what the evidence shows, she must consider all reasonable alternative priors. Suppose that an alternative factfinder had a high prior that this defendant would never crash the gate. In this case, the statistical evidence would barely move the needle, if at all. The large number of gatecrashers would not put any pressure on the hypothesis that D did not crash the gate, so long as the other participants are deemed sufficiently likely to gatecrash. In other words, the larger number of participants renders the original hypothesis (that D would not crash the gate) hardly less likely or more surprising than it was before.³³ This would not be true if D had similar views about more than one defendant. In such a case, the evidence would put pressure on that original hypothesis and point to revision. But for BARD, the hypothetical prior that treats this particular defendant as special seems unimpeachable in terms of reasonableness, even if the factfinder herself has no particular reason or inclination to adopt it.

This account is similar to Moss' in this case. On her telling of Gatecrasher, you do not *know* that this defendant crashed the gate, even if you have highly justified credence that he did. You do not know because you cannot rule out the relevant possibility that this defendant's character distinguishes him from the rest of the group.

Still, if this doesn't satisfy you, all the standard responses to statistical evidence remain available. In other words, one can accept my account, including thresholds for doubt, and still reject naked statistical evidence as inadmissible because, e.g., it cannot result in belief,³⁴ cannot result in knowledge,³⁵ it is uncaused³⁶, or due to normalcy³⁷ or other moral or policy reasons.³⁸

7. A Note on Applicability

This account is not meant as jury instructions. Nor do I advocate that juries necessarily sit with a calculator. I am not, after all, offering a threshold for BARD. Rather, this is a model of what a factfinder is aiming at. Perhaps some factfinders would best achieve this aim by actually using the model. For (most) others, heuristics may be preferable. Perhaps in cases with explicit probabilities, such calculation should be actively encouraged. But even then, factfinders must be careful to proceed from the entirety of the set of reasonable priors.

In practice, a reverse form of inference might be more effective. Ask yourself: what need I assume to make doubt salient? Then ask whether that is an assumption you'd rule out as unreasonable. This in practice will matter far more than settling the precise thresholds.

³³It is true that the larger n is, the more evidence there is that D gatecrashed and the probability that D is liable would increase ever so slightly (unless the probability that D is innocent = 1). But it is also true that however large the sample, there could be an arbitrarily high prior probability that would still render D innocent (at whatever threshold).

³⁴Buchack (2014).

³⁵Blome-Tillman (2017), Levanon (2019), Littlejohn (2020) and Moss (2022). For specific application to sensitivity (as necessary for knowledge) see Enoch et al. (2012), for Safety, see Pritchard (2018). For criticism of knowledge as the relevant standard, see Papineau (2021).

³⁶Thomson 1986).

³⁷Smith (2018).

³⁸Enoch et al. (2012).

8. Conclusion

I have defended a Bayesian approach to proof BARD, both as a concept in the law and in epistemic practice more generally. This sheds light on what Bayesians should say about the presumption of innocence and about the restriction to admitted evidence. My account rejects thresholds as an analysis of BARD. Rather, it calls for a judgment that *all* reasonable paths, not just one's own, lead to a guilty verdict.

By emphasizing the modal epistemic profile, the approach makes the BARD standard clearer and, as such, enables evaluation of whether the standard was met (e.g. via appellate review). One tries to identify a prior that a reasonable factfinder could have had that would have led to acquittal.³⁹

References

- Adler J.E. (2012). 'Contextualism and Fallibility: Pragmatic Encroachment, Possibility, and Strength of Epistemic Position.' *Synthese* **188**, 247–72.
- Allen R.J., & Pardo M. S. (2007). 'The Problematic Value of Mathematical Models of Evidence.' *The Journal of Legal Studies*, **36**, 107–40.
- Blome-Tillmann M. (2017). 'More Likely than not'—knowledge first and the role of bare statistical evidence in courts of law.' In A. Carter, E. Gordon and B. Jarvi (eds), *Knowledge First—Approaches in Epistemology and Mind*. Oxford: Oxford University Press.
- Bouvens L., & Hawthorne J. (1999). 'The Preface, the Lottery, and the Logic of Belief.' *Mind* **108**, 241–64.
- Buchak L. (2014). 'Belief, Credence, and Norms.' *Philosophical Studies* **169**, 1–27.
- Christensen, D. (1991). 'Clever Bookies and Coherent Beliefs.' *The Philosophical Review*, **100**, 229–47.
- Christensen D. (2001). 'Preference-Based Arguments for Probabilism.' *Philosophy of Science* **68**(3), 356–76.
- Christensen D. (2004). *Putting Logic in Its Place*. Oxford: Oxford University Press.
- Clarke R. (2013). 'Belief Is Credence One (in Context).' *Philosopher's Imprint* **13**, 1–18.
- Cohen L.J. (1977). *The Probable and the Provable*. Oxford: Oxford University Press.
- Cullison A.D. (1969). 'Probability Analysis of Judicial Fact-Finding: A Preliminary Outline of the Subjective Approach.' *Toledo Law Review* **1**, 538–98.
- Dahlman C., & Kolflaath E. (2021). 'The Problem of the Priors in Criminal Trials.' In C. Dahlman, A. Stein, and G. Tuzet (eds), *Philosophical Foundations of Evidence Law*. Oxford: Oxford University Press.
- Dahlman C., Wahlberg L., & Sarwar F. (2015). 'Robust Trust in Expert Testimony.' *Humana Mente* **28**, 17–37.
- Di Bello M. (2013). *Statistics and probability in criminal trials: The good, the bad and the ugly*. Doctoral dissertation, Stanford University.
- Dodd D. (2016). 'Belief and Certainty.' *Synthese* **194**(11), 4597–621.
- Dorst K. (2019). 'Lockeans Maximize Expected Accuracy.' *Mind* **128**(509), 175–211.
- Douven I., & Rott H. (2018). 'From Probabilities to Categorical Beliefs: Going Beyond Toy Models.' *Journal of Logic and Computation* **28**, 1099–124.
- Douven I., & Williamson T. (2006). 'Generalizing the Lottery Paradox.' *British Journal for the Philosophy of Science* **57**, 755–79.
- Easwaran K. (2016). 'How I Learned to Stop Worrying and Love Bayesian Probabilities.' *Noûs* **50**, 1–38.
- Elga A. (2010). 'Subjective Probabilities Should Be Sharp.' *Philosophers' Imprint* **10**, 2–11.
- Enoch D., Spectre L., & Fisher T. (2012). 'Statistical Evidence, Sensitivity, and the Legal Value of Knowledge.' *Philosophy and Public Affairs* **40**, 197–224.
- Fenton N., & Lagnado D. (2021). 'Bayesianism: Objections and Rebuttals.' In C. Dahlman, A. Stein, and G. Tuzet (eds), *Philosophical Foundations of Evidence Law*. Oxford: Oxford University Press.
- Finkelstein M.O., & Fairley W.B. (1970). 'A Bayesian Approach to Identification Evidence.' *Harvard Law Review* **83**(3), 489–517.
- Foley R. (1993). *Working Without a Net*. Oxford: Oxford University Press.

³⁹For comments thanks to Marcelo Di Bello, Julien Dutant, Richard A. Epstein, Branden Fitelson, Otto Lehto, David Papineau, Graham Priest, Mario Rizzo, and Cosim Sayid, as well as to the anonymous reviewer, and to audiences at CUNY Graduate Center and Birkbeck College, London.

- Franklin J.** (2001). *The Science of Conjecture: Evidence and Probability before Pascal*. Baltimore: Johns Hopkins University Press.
- Friedman J.** (2013). 'Rational Agnosticism and Degrees of Belief.' In T.S. Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology 4*. Oxford: Oxford University Press.
- Gaifman H., & Snir M.** (1982). 'Probabilities Over Rich Languages, Testing, and Randomness.' *Journal of Symbolic Logic* **47**, 495–548.
- Gardiner G.** (2019). 'The Reasonable and the Relevant: Legal Standards of Proof.' *Philosophy & Public Affairs* **47**(3), 288–318.
- Good I.J.** (1971). '46656 varieties of Bayesians.' Letter in *American Statistician* **25**, 62–63. Reprinted in I.J. Good (1982). *Good Thinking*, pp. 20–21. Minneapolis: University of Minnesota Press.
- Greaves H., & Wallace D.** (2006). 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility.' *Mind* **115**(459), 607–32.
- Greco D.** (2015). 'How I Learned to Stop Worrying and Love Probability 1.' *Philosophical Perspectives* **29**, 179–201.
- Haack S.** (2014). 'Legal Probabilism: An Epistemological Dissent.' In *Evidence Matters: Science, Proof, and Truth in the Law*, pp. 47–77. Cambridge: Cambridge University Press.
- Hajek A.** (2008). 'Arguments for-or Against-Probabilism?' *British Journal for the Philosophy of Science* **59**, 793–819.
- Harman G.** (1968). 'Knowledge, Inference, and Explanation.' *American Philosophical Quarterly* **5**, 164–73.
- Hawthorne J.** (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Jackson E.** (2019). 'Belief, Credence, and Faith.' *Religious Studies* **55**(2), 153–68.
- Jaynes E.T.** (2003). *Probability Theory: The logic of Science*. Oxford: Oxford University Press.
- Jeffrey R.C.** (1965). *The Logic of Decision*. New York: McGraw Hill.
- Jeffrey R.C.** (1970). 'Dracula Meets Wolfman: Acceptance vs. Partial Belief.' In M. Swain (ed.), *Induction, Acceptance, and Rational Belief*. Dordrecht: Synthese Library.
- Jellemma H.** (2024). 'Reasonable Doubt, Robust Evidential Probability, and the Unknown.' *Criminal Law and Philosophy* **18**, 451–70.
- Joyce J.** (1998). 'A Non-Pragmatic Vindication of Probabilism.' *Philosophy of Science* **65**, 575–603.
- Joyce J.** (2005). 'How Probabilities Reflect Evidence.' *Philosophical Perspectives* **19**, 153–79.
- Joyce J.** (2009). 'Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief.' In F. Huber and C. Schmidt-Petri (eds), *Degrees of Belief*. Dordrecht: Synthese Library.
- Kaplan J.** (1968). 'Decision Theory and the Fact-Finding Process.' *Stanford Law Review* **20**(6), 1065–92.
- Kaplan M.** (1996). *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.
- Kaye D.** (1980). 'Naked Statistical Evidence.' *Yale Law Journal* **89**, 601–11.
- Kelp C.** (2017). 'Lotteries and Justification.' *Synthese* **194**, 1233–44.
- Keynes J.M.** (1921). *A Treatise on Probability*. London: Courier Dover.
- Langbein J.** (2003). *The Origins of Adversary Criminal Trial*. Oxford: Oxford University Press.
- Lange M.** (1999). 'Calibration and the Epistemological Role of Bayesian Conditionalization.' *Journal of Philosophy* **96**, 294–324.
- Laplace P.S.** (1814). *A Philosophical Essay on Probabilities*. London: Dover.
- Laudan L.** (2006). *Truth, Error, and Criminal Law*. Cambridge: Cambridge University Press.
- Leitgeb H.** (2017). *The Stability of Belief*. Oxford: Oxford University Press.
- Leitgeb H., & Pettigrew R.** (2010). 'An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy.' *Philosophy of Science* **77**, 236–72.
- Levanon L.** (2019). 'Statistical Evidence, Assertions and Responsibility.' *The Modern Law Review* **82**(2), 269–92.
- Levi I.** (1974). 'On Indeterminate Probabilities.' *Journal of Philosophy* **71**, 391–418.
- Lewis D.K.** (1980). 'A Subjectivist's Guide to Objective Chance.' In R.C. Jeffrey (ed.), *Studies in Inductive Logic and Probability*, vol. II, pp. 263–293. Berkeley: University of California Press.
- Lewis D.K.** (1999). *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.
- Lin H., & Kelly K.T.** (2012). 'A Geo-Logical Solution to the Lottery Paradox.' *Synthese* **186**, 531–75.
- Lin H., & Kelly K.T.** (2013). 'Propositional Reasoning that Tracks Probabilistic Reasoning.' *The Journal of Philosophical Logic* **41**(6), 957–81.
- Littlejohn C.** (2020). 'Truth, Knowledge, and the Standard of Proof in Criminal Law.' *Synthese* **197**(12), 5253–86.
- Locke D.** (2014). 'The Decision-Theoretic Lockean Thesis.' *Inquiry* **57**, 28–54.

- Logue J.** (1997). 'Resiliency, Robustness and Rationality of Probability Judgments.' *International Studies in the Philosophy of Science* **11**(1), 21–34.
- McCauliff C.M.A.** (1982). 'Burdens of Proof: Degrees of Belief, Quanta of Evidence, or Constitutional Guarantees.' *Vanderbilt Law Review* **35**, 1293–335.
- Moss S.** (2018). *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Moss S.** (2022). 'Knowledge and Legal Proof.' In T.S. Gendler, J. Hawthorne, & J. Chung (eds) *Oxford Studies in Epistemology* vol. 7, pp. 176–213.
- Nelkin D.K.** (2000). 'The Lottery Paradox, Knowledge, and Rationality.' *Philosophical Review* **109**, 373–409.
- Nesson C.** (1979). 'Reasonable Doubt and Permissive Influence.' *Harvard Law Review* **92**, 1187–225.
- Papineau D.** (2021). 'The Disvalue of Knowledge.' *Synthese* **198**, 5311–32.
- Paul Teller, P.** (1976) 'Conditionalization, Observation, and Change of Preference.' In W. L. Harper and C.A. Hooker (eds) *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, vol. I, pp. 205–53. Dordrecht: D. Reidel Publishing Company.
- Pettigrew R.** (2014). 'Accuracy, Risk, and the Principle of Indifference.' *Philosophy and Phenomenological Research* **92**(1), 35–59.
- Pettigrew R.** (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.
- Pollock J.** (1990). *Nomic Probability and the Foundations of Induction*. Oxford: Oxford University Press.
- Pritchard D.** (2018). 'Legal Risk, Legal Evidence and the Arithmetic of Criminal Justice.' *Jurisprudence* **9**(1), 108–19.
- Ramsey F.P.** (1926). 'Truth and Probability.' In R.B. Braithwaite (ed.) (1931), *The Foundations of Mathematics and Other Logical Essays*. London: Routledge.
- Rawls J.** (1993). *Political Liberalism*. New York: Columbia University Press.
- Redmayne M.** (2008). 'Exploring the Proof Paradoxes.' *Legal Theory* **14**, 281–309.
- Roorda, J.** (1995). 'Revenge of Wolfman: A Probabilistic Explication of Full Belief.' (Unpublished manuscript available at <https://www.princeton.edu/~bayesway/pu/Wolfman.pdf>).
- Ross L.** (2024). *The Philosophy of Legal Proof*. Cambridge: Cambridge University Press.
- Scanlon T.M.** (1982). 'Contractualism and Utilitarianism.' In A. Sen and B. Williams (eds), *Utilitarianism and Beyond*. Cambridge: Cambridge University Press.
- Shear T., & Fitelson B.** (2019). 'Two Approaches to Belief Revision.' *Erkenntnis* **84**, 487–518.
- Skyrms B.** (2011). 'Resiliency, Propensities and Causal Necessity.' In A. Eagle (ed), (3 d.) *Philosophy of Probability: Contemporary Readings*, pp. 529–536. London: Routledge.
- Smith M.** (2010). 'A Generalised Lottery Paradox for Infinite Probability Spaces.' *The British Journal for the Philosophy of Science* **61**, 821–31.
- Smith M.** (2013). Review of Michael Titelbaum *Quitting Certainties*. <https://ndpr.nd.edu/news/quitting-certainties-a-bayesian-framework-modelling-degrees-of-belief/>
- Smith M.** (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford: Oxford University Press.
- Smith M.** (2018). 'When does Evidence Suffice for Conviction?' *Mind* **127**, 1193–218.
- Staffel J.** (2015). 'Beliefs, Buses, and Lotteries: Why Rational Belief Can't Be Stably High Credence.' *Philosophical Studies* **173**(7), 1721–34.
- Sturgeon S.** (2008). 'Reason and the Grain of Belief.' *Noûs* **42**(1), 139–65.
- Tang W.H.** (2009). *Belief and Credence*. Australian National University, Ph.D. Dissertation. Canberra, ACT.
- Thomson J.J.** (1986). 'Liability and Individualized Evidence.' In J.J. Thomson and W. Parent (eds), *Rights, Restitution and Risk: Essays in Moral Theory*. Cambridge: Harvard University Press.
- Thorn P.D.** (2020). 'The Joint Aggregation of Beliefs and Degrees of Belief.' *Synthese* **187**, 5389–409.
- Titelbaum M.** (2013). *Quitting Certainties: A Bayesian Framework for Modelling Beliefs*. Oxford: Oxford University Press.
- Tribe L.H.** (1971). 'Trial by Mathematics.' *Harvard Law Review* **84**, 1329–93.
- Urbaniak R., & Di Bello M.** (2021). Legal probabilism, *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), E. N. Zalta (Ed.), <https://plato.stanford.edu/archives/fall2021/entries/legal-probabilism/>
- van Fraassen B.C.** (1984). 'Belief and the Will.' *The Journal of Philosophy* **81**, 235–56.
- van Fraassen B.C.** (1989). *Laws and Symmetry*. Oxford: Oxford University Press.
- von Neumann J., & Morgenstern, O.** (1944). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.
- Wedgwood R.** (2012). 'Outright Belief.' *Dialectica* **66**, 309–29.
- Weinstein J.B., & Dewsbury I.** (2006). 'Comment on the Meaning of 'Proof Beyond a Reasonable Doubt.' *Law, Probability, and Risk* **5**, 167–73.
- White R.** (2005). 'Epistemic Permissiveness.' *Philosophical Perspectives* **19**, 445–59.

- White R.** (2010). 'Evidential Symmetry and Mushy Credence.' In S. Gendler and J. Hawthorne (eds), *Oxford Studies in Epistemology* vol. 3, pp. 161–86. Oxford: Oxford University Press.
- Whitman J.Q.** (2016). *The Origins of Reasonable Doubt*. New Haven: Yale University Press.
- Williamson J.** (2010). *In Defense of Objective Bayesianism*. Oxford: Oxford University Press.

Yuval Abrams is Assistant Professor of Law at Michigan State University. He received his PhD in Philosophy from CUNY Graduate Center. He works in legal philosophy, private law theory, evidence law, epistemology, and the philosophy of science.