

Accuracy, optimality and the preponderance standard

EDWARD K. CHENG[†]

Vanderbilt University Law School, 131, 21st Avenue South, Nashville, TN 37203

AND

MICHAEL S. PARDO

University of Alabama School of Law, Tuscaloosa, AL 35487

[Received on 14 May 2014; revised on 2 January 2015; accepted on 7 January 2015]

Law-and-economics scholars have recently argued that the legal system should set burdens of proof on the basis of ex ante welfare considerations. In this article, we reject this welfarist approach, showing that it relies on contested normative principles, raises legitimacy concerns, and is nearly impossible to implement in practice. As an alternative, we propose a decision-theory model that we constrain to account for core legal values and the practical limitations of the trial process. Specifically, we require that the burden of proof prioritize accuracy (truth) over welfare, and that it be capable of operating without knowledge of the base rates or prior probabilities of activities. The resulting optimization problem can be solved using a minimax approach, which minimizes the maximum probability of error faced by each of the parties, and remarkably, the minimax solution turns out to be precisely the preponderance-of-the-evidence standard currently imposed by courts. We thereby not only refute recent welfare theories about the burden of proof, but also provide a new theoretical justification for the traditional preponderance standard.

Keywords: burden of proof; burden of persuasion; evidence; accuracy; welfare; minimax; error.

1. Introduction

Traditional evidentiary burdens of proof¹ such as the preponderance-of-the-evidence standard have lately come under attack from law-and-economics scholars. These welfare-based theories argue that the legal system should set burdens of proof to optimize social welfare, rather than accuracy. In accord with their focus on welfare, the theories further suggest that the burden of proof should vary with the context of each specific case, rather than being uniform across substantive areas. In this article, we reject this welfarist view, arguing that the perspective sacrifices accuracy in ways antithetical to the legitimacy of the adjudicative system.² Welfare-focused burdens would also require information about costs that would be difficult if not impossible for the legal system to produce.

[†] Corresponding author. Email: edward.cheng@vanderbilt.edu.

¹ In legal circles, the burden of proof technically consists of the burden of production and the burden of persuasion, but here we will only use it to refer to the latter, as is often done colloquially.

² We should note that we do not take issue with a welfarist approach to all procedural law. Indeed, the law-and-economics lens can often provide an illuminating perspective on evidentiary rules, especially rules designed to promote substantive goals like the prohibition on subsequent remedial measures. Fed. R. Evid. 407; see also, e.g. Gideon Parchomovsky and Alex Stein, *The Distortionary Effect of Evidence on Primary Behavior*, 124 Harv. L. Rev. 518 (2010); Chris William Sanchirico, *A Primary Activity Approach to Proof Burdens*, 37 J. Legal Stud. 273 (2008); Chris William Sanchirico, *Character Evidence and the Object of Trial*, 101 Colum. L. Rev. 1227 (2001). There are nevertheless limits to this approach, and we argue that the ex ante approach

Beyond our critique, we attempt to turn the proverbial lemon into lemonade, viewing the welfare versus accuracy debate as an opportunity to revisit the justifications for traditional, accuracy-based burdens of proof, specifically the preponderance-of-the-evidence standard. Drawing on the statistical literature, we use a decision-theory model as a framework for thinking about evidentiary burdens. We constrain the model to account for the normative and practical constraints in the legal system—specifically, we require that it prioritize accuracy over welfare and that it avoid making unreasonable informational demands. By taking a minimax approach, we can solve for an optimal burden of proof, and we show that under careful examination, this optimal burden is in fact the preponderance-of-the-evidence standard currently used by courts. We thus provide a new justification for the traditional standard.

The remainder of the article proceeds in two parts. Part 1 critiques the welfarist view of burdens of proof. It articulates the problems with welfare theory's *ex ante* approach to evidence from normative, practical and logical perspectives. Part 2 offers an *ex post*, decision-theory model for analysing burdens of proof, presents the minimax solution and argues why it is identical to the preponderance-of-the-evidence standard. The article then briefly concludes and is followed by two technical appendices with proofs and other technical details.

2. Welfare theory and its problems

2.1 *Ex ante versus ex post perspectives*

A foundational question in constructing a normative theory of burdens of proof is whether to focus on *ex ante* or *ex post* considerations.³ Traditionally, the legal system sets burdens of proof from an *ex post* perspective. The *ex post* approach focuses on factual accuracy on the theory that the substantive law can only guide and evaluate conduct when it is applied to true depictions of reality.⁴ (The legal rights you have in your property become meaningless if the law mistakenly concludes it is not your property.) The primary normative considerations are 'backward looking', such as whether and how to minimize total errors; whether to treat types of possible errors equally (in terms of their costs) or to prefer one type of error over another (because of asymmetric costs); and what standards will best implement these choices.

In contrast, the welfare-based, law-and-economics approach to burdens of proof takes an *ex ante* perspective, focusing on how well a rule directly achieves the goals of substantive law.⁵ Under an *ex*

to 'burdens of proof' should be rejected. We also empathically do not mean to suggest that the traditional perspective is without problems or puzzles. Our point is only that the welfarist approach is deeply flawed and that the accuracy perspective likely remains the more fruitful approach.

³ For an illuminating discussion of *ex ante* and *ex post* perspectives on procedure, see Lawrence B. Solum, *Procedural Justice*, 78 S. Cal. L. Rev. 181, 183–90 (2004).

⁴ See Ronald J. Allen, *From the Enlightenment to Crawford and Holmes*, 39 Seton Hall L. Rev. 1, 7 (2009) ('Rights and obligations of any sort whatsoever are meaningless without accurate fact finding.'). Solum, *supra* note 3, at 320 ('[S]ubstance cannot effectively guide primary conduct without the aid of procedure.').

⁵ See Louis Kaplow, *Burden of Proof*, 121 Yale L. J. 738 (2012); A. Mitchell Polinsky and Steven Shavell, *The Economic Theory of Public Enforcement of Law*, 38 J. Econ. Lit. 45 (2000); Ivan P.L. Png, *Optimal Subsidies and Damages in the Presence of Judicial Error*, 6 Int'l Rev. L. & Econ. 101 (1986). For an overview of the literature, see Eric L. Talley, *Law, Economics, and the Burden(s) of Proof*, in *Research Handbook on the Economic Analysis of Tort Law* (J. Arlen ed., forthcoming 2013), available at <http://ssrn.com/abstract=2170469>. We note, however, that scholars adopting this perspective do recognize the potential welfare-enhancing role of accuracy. See, e.g. Polinsky and Shavell, *supra*; Png, *supra*. The essence of this position, in other words, is not necessarily antithetical to truth; rather, its essence is the primacy of welfare over truth should the two diverge.

ante approach, the primary normative considerations are forward looking; they concern how the standard of proof will affect future behaviour and whether that behaviour will increase or decrease general welfare (defined as the aggregate of individuals' well-being, preferences or utility).⁶ The behaviour encouraged or discouraged going forward is therefore more important than determining what happened.⁷ Low proof standards discourage parties from engaging in socially harmful conduct (which creates benefits) but may also chill socially beneficial conduct (which creates costs). High standards create the opposite effects: they decrease the likelihood of being held liable (either correctly or mistakenly) and thus encourage both harmful conduct (cost) and beneficial conduct (benefit).

More specifically, *ex ante* approaches to burdens of proof typically have the following three general features, all of which are presented as normative:

- (1) In choosing the burden of proof, the only relevant considerations are incentives on future behaviour. Backward-looking considerations about the disputed facts are irrelevant, except to the extent they may be proxies for the relevant forward-looking considerations.
- (2) The sole criterion for determining a correct (or optimal) standard of proof is general welfare. *Ex post* considerations regarding accuracy and errors are irrelevant, except to the extent they affect general welfare.
- (3) Standards of proof may vary among cases and should not be 'trans-substantive'.⁸ Because the effects that a standard has on general welfare will vary among types of cases (and even from case to case), the standard of proof should vary as well.

2.2 *An example*

One example of the *ex ante* approach to burdens of proof was recently proposed by Louis Kaplow.⁹ Our critique applies to *ex ante*, welfare theories more generally,¹⁰ but Kaplow's proposal helps ground the ensuing discussion.

⁶ See Talley, *supra* note 5, at 36 ('[L]aw and economics scholars have begun to deliver insights into how elements of litigation (including burdens) are likely to have feedback effects on primary behavior of plaintiffs and defendants—the subject matter than is most closely associated with substantive law (rather than procedural rules)'). On general welfare, see Louis Kaplow and Steven Shavell, *Fairness Versus Welfare* 18 (2002) ('The notion of well-being used in welfare economics is comprehensive in nature. It incorporates in a positive way everything that an individual might value.').

⁷ For example, as Louis Kaplow suggests:

[T]he preponderance rule, and other rules based on the likelihood that the individual before the tribunal is one who committed the harmful act, embody an *ex post* perspective that takes behavior as given (exogenous) and asks, in light of that behavior, what is the likelihood of one versus another characterization. By contrast, a welfare-based, optimal threshold in these settings is determined by asking how behavior (taken to be endogenous) will change as a function of a change in the evidence threshold.

Kaplow, *supra* note 5, at 748.

⁸ See generally David Marcus, *The Past, Present, and Future of Trans-Substantivity in Federal Civil Procedure*, 50 DePaul L. Rev. 371 (2012) (discussing the trans-substantive goals of the Federal Rules of Civil Procedure).

⁹ See Kaplow, *supra* note 5; see also Louis Kaplow, *Optimal Proof Burdens, Deterrence, and the Chilling of Desirable Behavior*, 101 Am. Econ. Rev. Papers & Proc. 277 (2011); Louis Kaplow, *On the Optimal Proof Burden*, J. Political Economy, forthcoming; Louis Kaplow, *Likelihood Ratio Tests and Legal Decision Rules*, American Law & Economics Review, forthcoming, available at <http://ssrn.com/abstract=2284035>. Kaplow nominally presents a theory of 'optimal evidence thresholds', Kaplow, *supra* note 5, at 756–77 (defining an evidence threshold as 'the overall strength of the evidence... required for liability'), but they serve essentially the same function as the traditional standards of proof. They provide a decision rule for fact-finders by telling them how strong the evidence must be to find a disputed factual issue as proven.

¹⁰ For an overview and additional examples, see Talley, *supra* note 5.

In accord with the *ex ante* perspective, Kaplow's approach to evidentiary burdens concentrates on 'deterrence benefits' and 'chilling costs', rather than on maximizing accuracy.¹¹ Because decreases in the evidentiary burden make liability more likely, they increase not only the number of harmful acts deterred, but also the number of socially productive acts chilled. The crux of Kaplow's theory is thus to conceptualize, measure and compare these two effects, and thereby to find the evidentiary threshold that maximizes social welfare.

Although Kaplow goes into greater detail, the deterrence benefit associated with a decrease in the burden of proof essentially consists of 'the number of individuals [or acts] who will be deterred as a consequence' multiplied by 'the net social gain per deterred act'.¹² The chilling cost calculation is essentially a mirror image, multiplying the number of benign acts chilled by the social loss suffered per chilled act.¹³ The final step is to compare the values associated with each side. Under Kaplow's theory, the optimal evidence threshold is the one that causes the marginal costs and benefits to be equal.¹⁴ Although Kaplow does not provide any examples of an optimal evidence threshold, he does note that the optimal threshold may vary dramatically from the current standards of proof in either direction.¹⁵ And just to re-emphasize the gap between *ex ante* welfare and *ex post* accuracy approaches, Kaplow further explains:

[T]he many elements that determine the magnitudes of deterrence and chilling effects do not pertain to whether it is more likely than not that the individual before the tribunal committed a harmful act rather than a benign one—and conversely.¹⁶

2.3 Problems with *ex ante* theories

The problems with the *ex ante* approach fall into three general categories: (1) theoretical problems with welfare as the sole normative criterion; (2) the empirical and conceptual infeasibility of implementing such an approach; and (3) tensions within the approach itself that may frustrate rather than serve the welfare goal. We discuss each in turn.

2.3.1 Contested normative grounds. The first limitation to the welfare theory is that it depends on a controversial normative criterion. It assumes that aggregate welfare provides the only legitimate basis for constructing a standard of proof, yet that presupposes a larger normative argument about the legitimate grounds for justifying legal rules. Kaplow's discussion relies on his work with Steven Shavell as the general philosophical grounding for his burden-of-proof analysis.¹⁷ Their book

¹¹ 'Chilling' is the process of discouraging actors in ways they desire, except that unlike 'detering', it refers to benign acts (i.e. those that are not subject to criminal or civil liability) that may otherwise provide beneficial gains to the actors and society. Kaplow, *supra* note 5, at 746.

¹² Id. at 764, 766. The number of individuals deterred is found by multiplying the 'increase in expected sanction' (since a lower threshold will increase the expected sanction) by the 'concentration of marginal harmful acts', which measures how many individuals change their behaviour in response to this increase in sanction. Id. The net social gain is the difference between the act's social harm and its private benefit (to the malefactor). Id.

¹³ Id. at 764–68. For benign acts, there is no social cost component, unlike with harmful acts. Id. ('The net cost per chilled act is straightforward: it is simply the act's private benefit because the social loss consists of the forgone benefit from the act.').

¹⁴ Id. at 769 ('The optimum will have the feature that the deterrence benefit just equals the chilling cost.').

¹⁵ Id. at 748.

¹⁶ Id. See also id. at 859 ('[C]onventional conceptions of the burden of proof do not constitute even weak proxies for welfare because they depend almost entirely on factors different from those that determine the optimal evidence threshold.').

¹⁷ Kaplow and Shavell, *supra* note 5.

argued forcefully in favour of the primacy of welfare over ‘fairness’ (broadly defined to include ‘justice, rights, and cognate concepts’), but as Kimberly Ferzan has pointed out, while their arguments may have been welcome news to economic types already predisposed to accept the conclusions, they failed to persuade those on the fairness side because they simply begged many of the critical questions.¹⁸ This is not to deny a central place for welfare in law and public policy. Far from it. But caring about welfare is one thing; caring ‘only’ about maximizing welfare another.¹⁹

Moreover, even if one accepts many of the general arguments by Kaplow and Shavell, there are reasons specific to procedure and evidence for rejecting general welfare as the sole criterion. In the procedural domain, fair process, equal treatment and rights of parties (particularly criminal defendants) affect the legitimacy (and perceived legitimacy) of the proceedings.²⁰ Consider the following passage from Kaplow:

[D]epending on the context, the optimal evidence threshold could be much more demanding or notably more lax than the preponderance rule (or a requirement of proof beyond a reasonable doubt). For example, if much benign activity might be chilled and the harmful acts in question cause little social loss and are mostly undeterrable in any event, very strong evidence should be required for liability. In contrast, if there is little prospect of chilling beneficial activity and the pertinent harmful acts impose extreme damage and might readily be deterred, a low threshold should be employed.²¹

Suppose that we have a criminal case involving a brutal rape and murder. The key issue at trial is identity—was the defendant the culprit? For *ex ante* welfare theory, whether this defendant committed this crime and how strong the evidence must be to justify a conviction should not concern us. *Ex ante* considerations should. This crime involves a ‘harmful act[] that impose[s] extreme damage and might readily be deterred’—and very little socially benign activity gets mistaken for a brutal rape and murder—therefore ‘a low threshold [or burden of proof] should be employed’. Many, we assume, would find this to be an unacceptable standard to employ in this context. The focus on social welfare comes at a serious price. It sharply diminishes the adjudicative system’s ability to reach accurate results in the individual case. If adjudication is to have a truth-seeking function, in which judgements of liability (or no liability) are expressions of blame (or no blame), a primary goal of adjudication must be on reaching the correct individual outcome. In other words, we should be focusing on accuracy, not social welfare.²²

¹⁸ See Kimberly Kessler Ferzan, *Some Sound and Fury From Kaplow and Shavell*, 23 *Law & Philosophy* 73, 78–79 (2004) (‘Kaplow and Shavell not only beg the question of how to define fairness theories but also fail to provide any argument for why welfare, as they define it, is a good thing and why we should seek to maximize it.’).

¹⁹ We might compare the welfare theory in this regard to arguments about extreme forms of act utilitarianism. One may concede to the utilitarian that an act will maximize utility and still maintain that it is not morally permissible. See Judith Jarvis Thomson, *Turning the Trolley*, 36 *Phil. & Public Affairs* 359 (2008) (arguing that some of the crude utilitarian solutions to so-called ‘Trolley Problems’ are not morally permissible).

²⁰ See Solum, *supra* note 3, at 277–81 (discussing the relationship between legitimacy and procedural values).

²¹ Kaplow, *supra* note 5, at 748.

²² Sometimes the accuracy concerns will be asymmetric and focus on one type of accurate outcome (e.g. true convictions) at the expense of another (e.g. true acquittals), rather than maximizing total accuracy, because of the costs associated with different types of mistakes. Perceptions about asymmetric costs, as well as the need for strong proof when condemning citizens under the criminal law, provided the basis for the Supreme Court’s enshrining the ‘beyond a reasonable doubt’ standard with constitutional status in criminal cases. *In re Winship* 397 U.S. 358, 364 (1970); see also Fredrick E. Vars, *Toward a General Theory of Standards of Proof*, 60 *Catholic Univ. L. Rev.* 1 (2010) (arguing on the same grounds for the ‘clear and convincing evidence’ standard for will contests alleging mental incapacity).

The reply ‘but it maximizes welfare’ begs the question, and no new substantive normative arguments are forthcoming about why we should accept the welfare conclusion in the burden-of-proof context at the expense of accuracy or the constitutional rights of criminal defendants. The economist’s ‘counter-intuitive finding’ may in this context be everyone else’s *reductio ad absurdum*.²³

Similar considerations apply to less extreme examples. The welfare approach is committed to different evidence thresholds for civil parties in similar types of cases.²⁴ Imposing different proof requirements on plaintiffs bringing similar types of claims, or the same claim filed weeks apart, is also likely to undermine the perceived legitimacy of the decisions. To the extent that the appearance of fairness requires that courts impose a single set of rules for all litigants, the presence of shifting burdens of proof will almost certainly lead to charges of political manipulation or ideological bias by judges.

Again, none of this is meant to suggest that welfare is not an important consideration for law.²⁵ We are merely pointing out that one limitation of welfare theory is that its fate depends on a contested and highly controversial position in political philosophy. We do not intend to rehash these philosophical debates, and nothing in our critique depends on a general rejection of welfare as a normative criterion. Even if one is persuaded by *ex ante* welfare theory as a general approach to law, there are counter-vailing normative considerations that apply specifically in the procedural context. These considerations raise doubts about the desirability of applying welfare theory to burdens of proof. Moreover, even if one endorses welfare theory as a normative matter—both generally and in the domain of procedure—there are additional problems with *ex ante* welfare approaches to burdens of proof. The following two sections discuss these problems. We note these problems exist regardless of where one stands on the normative issues.

2.3.2 Practical feasibility. A second problem with the *ex ante* welfare approach is that it is simply impractical to implement. In the law of evidence ‘ought’ implies ‘can’.²⁶ This slogan expresses the general methodological principle that normative advice must require actions that the relevant actors are actually capable of performing. An evidence rule telling jurors to travel back in time and watch the events that gave rise to the litigation would obviously run afoul of this constraint, and so do *ex ante* welfare theories. In order to implement *ex ante* theory for any category of cases, one would need a staggering amount of empirical data, including the costs, benefits and frequency (base rates) of the relevant harmful and benign acts, as well as answers to the counter-factual questions of how behaviour is likely to change if different evidence thresholds were employed. Courts are certainly not in a position to know this information from case to case, nor, for that matter, are legislatures and agencies in many contexts. Virtually any kind of human behaviour can become the subject matter of litigation,

²³ Compare, once again, arguments about crude forms of act utilitarianism. See *supra* note 19.

²⁴ Kaplow, *supra* note 5, at 786 n.86 (explaining that optimal evidence thresholds ‘will vary greatly by context, even at fairly refined levels’).

²⁵ If lawmakers decide that certain antisocial behaviours are more difficult to detect or cause more harm, then the substantive law can change to perform that social accounting, much as traditional law-and-economics scholarship has attempted to explain in tort law. Indeed, although maintaining this strict separation between substantive law and evidentiary proof removes some flexibility, the separation arguably simplifies the optimization problem. Evidentiary rules can focus on getting things right. Allowing bleed between the two spheres sets up moving targets, making optimization more difficult. See Edward K. Cheng, *The Perils of Evidentiary Manipulation*, 93 Va. L. Rev. In Brief 191 (2007).

²⁶ See Ronald J. Allen and Brian Leiter, *Naturalized Epistemology and the Law of Evidence*, 87 Va L. Rev. 1491, 1499 (2001) (quoting Alvin I. Goldman, *Epistemics: The Regulative Theory of Cognition*, 75 J. Phil. 509, 510 (1978)).

and to require that we know all these answers before we can formulate a decision rule is a fool's errand.²⁷

There is a deeper, related problem lurking here as well. The above paragraph assumed that we know for any given case the 'relevant category of cases' to which it belonged. But this is precisely where the problem of 'reference classes' emerges.²⁸ The level of generality and the boundary conditions for any given category of cases are not something fixed by the world. Any two cases may be alike in some respects, but they will also be different in many respects. For these reasons, any given case will belong to many different reference classes (indeed, infinitely many), and each one will give different numbers. Take, e.g., a case involving something simple like littering. Assuming we had the time, resources and energy to collect the needed data to search for our optimal evidence threshold, is the relevant category: littering; littering in parks; littering near highways; littering with Styrofoam, plastic, or food; or littering in the summer? And so on. Each category will provide different numbers and a different optimal threshold.²⁹

Kaplow's proposal recognizes these two problems, but does not trace out their implications. For example, he notes that the 'high information requirements for [his theory's] application' present an 'obvious practical challenge'.³⁰ Anticipating the reference-class problem, he notes that the high informational requirements apply to 'a broad class of cases, a narrower cluster, or a given case'³¹ and that the 'optimal level of generality of burden of proof rules is beyond the scope of the present investigation'.³² Responding to a possible objection regarding the theory's demands, he explains:

[P]ossible differences in information requirements have little to do with the merits of one or another way of setting the burden of proof. The optimal evidence threshold is that which maximizes social welfare. If welfare is the objective, it would make no sense to substitute some other formulation because it happened to be easier to employ when that alternative has essentially no relationship to what matters. (Easiest of all would be to flip a coin, but its ease offers little to commend it.)³³

True enough. But there is a great deal of conceptual space between a coin flip and the impossible (and it is to this space that we turn in the second half of this article). Fact-finding based on time travel would be a pretty good normative theory from an epistemological perspective, if it could be done. Nevertheless, we are still quite right to ignore the normative advice of the time-travel theory, at least until its defenders provide some reason to think it is within reach.³⁴

²⁷ In a forthcoming article, Ronald Allen and Alex Stein also note the insurmountable information costs underlying the ex ante framework. See Ronald J. Allen and Alex Stein, *Evidence, Probability, and the Burden of Proof*, 54 *Arizona L. Rev.* 557 (2013).

²⁸ See generally Ronald J. Allen and Michael S. Pardo, *The Problematic Value of Mathematical Models of Evidence*, 36 *J. Legal Stud.* 107 (2007); Edward K. Cheng, *A Practical Solution to the Reference Class Problem*, 109 *Colum. L. Rev.* 2081 (2009).

²⁹ Indeed, some reflection may lead one to the conclusion that the relevant reference class is (1) the event itself or, alternatively, (2) all the primary behaviour in the world preceding and following the decision. The event is a member of both classes and many in between.

³⁰ Kaplow, *supra* note 5, at 786.

³¹ *Id.*

³² *Id.* at 786 n.86.

³³ *Id.* at 787.

³⁴ See Allen and Leiter, *supra* note 26, at 1521 (criticizing a focus on the elegant implications of ideals and formalisms without connection to reality as 'rootless theorizing').

2.3.3 A self-defeating proposition? A final conceptual problem arises from two premises undergirding ex ante theories: (1) welfare is the sole normative criterion, and (2) primary behaviour will change based on changes in the standard of proof. We noted limitations with the first premise above. We doubt the second is true in many contexts, except perhaps for sophisticated repeat players or those making decisions about primary behaviour in consultation with legal counsel, but we will grant the assumption for purposes of this discussion.

To illustrate the problem, we need to distinguish between the time an actor decides to engage in (or forgo) a harmful or benign act (T_1) and the time a factual decision is made by a legal fact-finder (T_2).

Suppose we could overcome all of the limitations noted above, collect all the data we would need and run the relevant calculations. The numbers will likely be different when we run them at T_1 or T_2 . So, which do we use to determine the standard of proof? Each is a problem for ex ante theories.

From the ex ante perspective, T_2 appears to be the proper time period. The relevant issue is future behaviour, and so data at the time of decision would be most pertinent to evaluating how the standard will influence behaviour immediately following the decision. Those not operating within the grip of an ex ante theory might think that there are serious fairness problems with a decision rule not available to the parties until the point of decision, and they would be right. As the Supreme Court has explained:

[T]his Court never has approved case-by-case determination of the proper standard of proof for a given proceeding. Standards of proof, like other “procedural due process rules[,] are shaped by the risk of error inherent in the truth-finding process as applied to the generality of cases, not the rare exceptions.” Since the litigants and the factfinder must know at the outset of a given proceeding how the risk of error will be allocated, the standard of proof necessarily must be calibrated in advance.³⁵

But this objection is no matter for ex ante welfare theory. Notice, however, what follows from the second premise (i.e. actors adjust their behaviour based on the threshold). Sophisticated actors at T_1 realize that the decision rule that will apply to them at trial will be calculated at T_2 . Moreover, given the highly contingent nature of the data, the applicable standard could be virtually anything at all by the time they get to T_2 . This uncertainty is likely to dilute or even destroy the incentive structure provided by the standards in the first place.

As an alternative, we can push the moment of calculation back to T_1 . This may provide more notice and a fairer rule for affected individuals (though again that is no matter for any theory that rejects non-welfare considerations). But reliance on T_1 calculations will not maximize welfare going forward (violating the first premise). This is because in the time between T_1 and T_2 , the relevant variables will surely change: thus the standard used in either case will be sub-optimal.

To be more concrete, let us consider the following example: suppose that there is an activity that is illegal, but because there is a high proof threshold at T_1 , people freely engage in it because the probability of liability is very low (e.g. illegal music sharing). At T_1 , many people make this calculation, so there is subsequently lots of illegal activity and resulting social cost. Now suppose some individuals are sued. Because the problem is so prevalent, welfare theory now requires that a low threshold be employed at trial (i.e. T_2). Notice, however, that the T_2 threshold differs dramatically from the T_1 threshold. Putting aside the fairness problems with this, as soon as the low T_2 threshold is imposed, future individuals will know that they cannot rely on the ‘optimal’ standard at the time of

³⁵ *Santosky v. Kramer*, 455 U.S. 745, 757 (1982) (quoting *Mathews v. Eldridge*, 424 U.S. 319, 344 (1976) (emphasis added in *Santosky*)).

action to calculate their expected sanctions, thus destroying the incentive structure the theory is meant to advance.³⁶ In contrast, if we use the original T_1 standard at the time of trial, it will be extremely sub-optimal, because the pervasiveness of the harmful act now justifies a lower standard of proof than previously.

2.4 *Taking stock*

This part discussed theoretical, empirical, practical and conceptual impediments to an ex ante welfare-based approach to standards of proof. Normative theories that focus on ex ante consequences, take welfare as their goal and contemplate shifting burdens of proof face these problems because they (1) rely on controversial normative premises, (2) make unreasonable empirical and conceptual demands and (3) may not produce the effects they predict.³⁷

For these reasons, we argue that burdens of proof are best conceptualized from an ex post perspective. The soundness of our conclusion, however, depends on defending an ex post approach that avoids these problems (as well as any other problems that are as significant). Towards this end, we extract two general lessons from our critique of the ex ante approach that inform our subsequent discussion of ex post approaches to standards of proof:

- (1) A theory should rely on plausible normative assumptions and accord with procedural values. Reliance on controversial normative premises is, of course, possible, but doing so requires compelling normative arguments justifying the controversial stance. This limitation undermines ex ante approaches, and ex post approaches should avoid a similar pitfall.
- (2) A theory should comply with the 'ought implies can' principle. Unlike ex ante theories, an ex post theory should avoid making impossible-to-meet informational demands.

3. Finding an ex post, accuracy-based alternative

Having criticized and rejected the ex ante, welfare-based approach, what is the ex post alternative? In this part we present an alternative framework for thinking about burdens of proof based on decision-theory principles.³⁸ Using this framework, we can find an optimal decision rule using a 'minimax' approach that respects accuracy, avoids unreasonable informational demands and establishes a uniform burden of proof for civil cases. We further show that this optimal decision rule is indeed the preponderance-of-the-evidence standard currently used by courts.

³⁶ One possible rejoinder is that individuals could simply account for the 'feedback' loop created by the (changing) burden of proof and the (changing) levels of activity, thus leaving the incentive structure intact. Yet to do so would create still another formidable data-collection obstacle to implementation. See Allen and Stein, *supra* note 27.

³⁷ It is of course possible for an ex ante theory to relax some of its premises to respond to some of these problems. But relaxing all the premises to resolve all of the problems would turn it into an ex post theory (i.e. focus on ex post considerations, consider goals beyond welfare, and require some substantial degree of trans-substantivity).

³⁸ Beginning with John Kaplan, *Decision Theory and the Factfinding Process*, 20 *Stan. L. Rev.* 1065 (1968), a substantial literature has applied a decision-theory framework to legal proof. See Michael S. Pardo, *The Nature and Purpose of Evidence Theory*, 66 *Vand. L. Rev.* 547, 549–51 (2013) for a brief overview and citations. Decision theory of course has a long history, including classic statistical texts such as Abraham Wald, *Statistical Decision Functions* (1950), and David Blackwell and M.A. Girshick, *Theory of Games and Statistical Decisions* (1954).

3.1 *A decision-theory perspective*

Under the decision-theory perspective, a legal fact-finder faces two possible truths: plaintiff's story, which we will label H_π , and defendant's story, which we will label H_Δ . At the same time, the fact-finder can render two possible decisions: a finding for the plaintiff or a finding for the defendant. There are therefore two possible kinds of error: finding for the plaintiff when the defendant's story is true, labelled as a 'false positive', and finding for the defendant when the plaintiff's story is true, labelled as a 'false negative'. There are also two possible types of correct results: finding for the plaintiff when the plaintiff's story is true, labelled as a 'true positive', and finding for the defendant when the defendant's story is true, labelled as a 'true negative'. Note that the perspective is distinctly *ex post*—decisions are evaluated as 'true' or 'false' based on consistency with past events.

Conventional decision theory then assigns costs (or benefits) to each of these four outcomes, and tries to find a decision rule that minimizes the expected costs.³⁹ (For a derivation of the optimal decision rule under the conventional model and further details, see Appendix A). The problem with this approach, however, is that it is subject to many of the same critiques that were levelled against the *ex ante* welfare approach. By accounting for the costs and benefits of decisions, the model prioritizes social welfare over accuracy. It also raises practical feasibility concerns, because the values to be assigned to each of the possible trial outcomes are invariably data-intensive, contingent and contested. And finally, because the costs change by context (or by case), courts cannot impose a single threshold for all litigants, thereby creating legitimacy problems.

We thus need to constrain the more general decision-theory framework to accord with the legal system's concern for accuracy. Specifically, the goal becomes not to find the decision rule that minimizes expected cost, but rather the rule that minimizes the number of errors, or more usefully, the probability of error.⁴⁰ To do this, let us examine the consequences associated with the two decision options that a legal system faces. If the system finds for the plaintiff, then it makes an error if the defendant's story is true. Therefore, the probability of error is the probability that the defendant's story is true given the evidence presented, or $P(H_\Delta|E)$. Similarly, if the system finds for the defendant, then the probability of error is the probability that the plaintiff's story is true given the evidence presented, or $P(H_\pi|E)$.

To minimize the probability of error, we should decide in favour of the plaintiff only if finding for the plaintiff generates a lower probability of error than finding for the defendant—in other words, if $P(H_\Delta|E) < P(H_\pi|E)$. Rearranging this algebraically, we have the following decision rule:

$$\text{Choose } H_\pi \text{ if } \frac{P(H_\pi|E)}{P(H_\Delta|E)} > 1. \quad (1)$$

Now, the problem is that we do not really know the probability of the defendant's or plaintiff's story given the evidence, or at a minimum, the decision rule does not offer much guidance on how to use the

³⁹ The costs considered can be quite general, and can include *ex ante* as well as *ex post* considerations. For example, the cost of a false positive can include the injustice of holding an innocent defendant liable (*ex post*) as well as the social welfare costs of failing to deter future antisocial behaviour (*ex ante*).

⁴⁰ Here, in decision theory parlance, we have chosen a particular loss function, and one could imagine others—e.g. one that prioritized false positives to false negatives. The symmetry, however, 'is appealing if the legal system does not [favor] one type of litigant over the other'. D.H. Kaye, Clarifying the burden of persuasion: what Bayesian decision rules do and do not do, 3 *Int'l J. Evid. & Proof* 1, 7 (1999). We have also placed aside the trivial (and unhelpful) solution not to decide cases at all, which would also reduce errors to zero. Courts, whether by rule or by tradition, are plainly required to decide cases for better or worse.

evidence to reach a decision. To determine these probabilities, one can (at least in theory) use Bayes's Theorem,⁴¹ which after some substituting and simplifying, gives us as a decision rule:

$$\text{Choose } H_{\pi} \text{ if } \frac{P(E|H_{\pi})}{P(E|H_{\Delta})} > \frac{P(H_{\Delta})}{P(H_{\pi})}. \quad (2)$$

In this equation, $P(H_{\pi})$ and $P(H_{\Delta})$ are known as the 'prior probabilities' for the plaintiff's and defendant's stories, respectively. They are the baseline probabilities that the stories are true 'prior' to seeing the evidence.⁴² Note also that $P(E|H_{\pi})$, which is the probability of observing the presented evidence assuming the plaintiff's story is true, is emphatically not the same as $P(H_{\pi}|E)$, which is the probability of plaintiff's story being true given the evidence.⁴³

This decision rule is a reasonable foundation for evidentiary burdens. At its core (as seen on the left side of the inequality), it requires that the fact-finder compare the probability of the evidence given the plaintiff's theory with the probability of the evidence given the defendant's theory. The comparative aspect of this model accords with the story model of jury decision-making as well as the general structure of legal proof, which involves the epistemic evaluation of competing explanations.⁴⁴ The comparison under this model, however, is not 'straight-up', but is instead weighted by the priors on the right side. If the defendant's story is highly implausible as an initial matter (e.g. aliens caused his car to swerve), we would expect it to be difficult for the defendant to win, and the rule operates accordingly.

3.2 The base rate problem

The decision rule in (2), however, still has a problem. To be optimal, it requires precise knowledge of the priors for the plaintiff's story and the defendant's story. One way to estimate priors is to consider the base rates for the two stories, but as with ex ante welfare theory, the relevant base rates are highly data-dependent and raise reference class problems. For example, if the defendant claims that plaintiff's car ran the red light, what is the base rate or prior probability that is true? And what reference class is appropriate? Any car? Cars that get into accidents? Cars during rush hour? In addition, because base rates are context-dependent, courts still could not impose a single threshold for civil litigants. To implement the decision rule, courts would need to impose different standards from case to case, thereby creating legitimacy problems.

⁴¹ By Bayes's Theorem,

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

and

$$P(H_{\Delta}|E) = \frac{P(E|H_{\Delta})}{P(E)} \times P(H_{\Delta}).$$

⁴² For example, when presented with a pair of dice, the prior probability that it is fair is quite high. However, after rolling 20 times and getting only sevens, one might conclude that the probability given the evidence of the dice being fair (known as the 'posterior probability') is much lower.

⁴³ For example, the probability a defendant was negligent given evidence that he had been drinking is not the same as the probability that, if the defendant were negligent, there would be evidence that he had been drinking.

⁴⁴ See Pardo, *supra* note 38, at 596–612; Michael S. Pardo and Ronald J. Allen, *Juridical Proof and the Best Explanation*, 27 *Law & Philosophy* 223, 233–42 (2008); Nancy Pennington and Reid Hastie, *A Cognitive Theory of Juror Decision Making: The Story Model*, 13 *Cardozo L. Rev.* 519 (1991).

But how does one construct a test when base rates are unknown? Two approaches stand out in the statistical literature, which we will consider in turn.⁴⁵

3.2.1 The classical (Neyman–Pearson) solution. One path for deciding between two stories or hypotheses without requiring base rates is offered in the classical Neyman–Pearson test.⁴⁶ Under classical hypothesis testing, one holds the false positive probability below a preset level (usually called α) and then finds the test that minimizes the false negative probability. The famous result that Neyman and Pearson proved was that the optimal test under these constraints has the following form:⁴⁷

$$\text{Choose } H_\pi \text{ if } \frac{P(E|H_\pi)}{P(E|H_\Delta)} > \tau_{\text{NP}}.$$

where τ_{NP} , the Neyman–Pearson threshold, depends on α and the probability distributions involved.⁴⁸

Although the Neyman–Pearson solution provides a reasonable approach to testing that does not require knowledge of base rates, it does not solve the legal system’s problem. First, calculating τ_{NP} is no easy task, since the probability distributions for $P(E|H_\pi)$ and $P(E|H_\Delta)$ are complex and unknown in the legal context. In more conventional statistics problems, the probability distribution for the evidence is known or modelled based on assumptions, making the calculation (sometimes) more straightforward. For example, the distribution for height is assumed to be approximately normally distributed. Evidence in a legal case typically lacks such convenient structure.

Even more fatal to the Neyman–Pearson approach is that it prioritizes one kind of error (false positives), a state of affairs that, although perhaps desirable in the criminal context, is potentially unfair in most civil contexts. The approach pegs the false positive rate—the risk of error against the defendant—at all costs, and minimizes the risk of error against the plaintiff only after that constraint is satisfied. Such a procedure will always result (except in rare coincidences) in different risks of error for plaintiffs and defendants, which is at odds with the equal treatment expected in civil cases.

3.2.2 The minimax solution. Another base-rate-free approach to hypothesis testing involves minimax optimization.⁴⁹ The minimax approach rests on a relatively simple idea. Since we do not have a good sense of the base rate, then rather than guessing blindly and basing our calculations on those guesses, we should instead operate conservatively on a worst-case scenario basis.⁵⁰ More specifically, we should develop tests that minimize the maximum error, ‘guarantee[ing] a minimum level of test

⁴⁵ One additional approach not explored here is to use existing data to form the prior, an ‘empirical Bayes’ approach. Contexts naturally involving statistical studies, such as pharmaceutical liability litigation, lend themselves especially to such a data-driven approach, as illustrated in Joseph L. Gastwirth, Should Law and Public Policy Adopt ‘Practical Causality’ as the Appropriate Criteria for Deciding Product Liability Cases and Public Policy?, 12 L. Prob. & Risk 169 (2013).

⁴⁶ Bernard C. Levy, Principles of Signal Detection and Parameter Estimation 29–32 (2008).

⁴⁷ E.L. Lehmann and Joseph P. Romano, Testing Statistical Hypotheses 59–63 (2005).

⁴⁸ This discussion simplifies the exact contours of the Neyman–Pearson Lemma. For example, there is technically a randomized test when the likelihood ratio is equal to τ_{NP} . For additional details, see Lehmann and Romano, *supra* note 47.

⁴⁹ Levy, *supra* note 46, at 39–45; H. Vincent Poor, An Introduction to Signal Detection and Estimation 19–31 (1988); James O. Berger, Statistical Decision Theory and Bayesian Analysis 354–57 (2d ed. 1985); see also John Rawls, A Theory of Justice 154–58 (1971) (famously employing the related ‘maximin’ solution to justify unequal distributions of resources in some circumstances [i.e. when it maximizes benefits for those worst off]).

⁵⁰ See, e.g. Wald, *supra* note 38, at 18 (‘[A] minimax solution seems, in general, to be a reasonable solution of the decision problem when an a priori distribution . . . does not exist or is unknown to the experimenter.’).

performance independently of the actual prior [or base rate] values'.⁵¹ With regard to adjudicative errors, we can think of our problem in the following way. Some proportion of plaintiffs who go to trial have true claims and deserve to win, but we have no idea what proportion. It could be zero, all, or anywhere in between. Likewise, there is some proportion of falsely accused defendants. But it could be zero, all, or anywhere in between. Based on this uncertainty, what decision rule should you employ if your primary goal is total accuracy?

Philosophically, the minimax approach also respects the need for the legal system to treat the parties equally, though equality here is defined in a particular way.⁵² The approach minimizes the maximum error rate faced by either party. Thus, it is 'fair' in the sense that the system focuses on the error rate of the party who is worst off, whether that is the plaintiff or the defendant.

3.3 *Finding the optimal burden of proof*

Thus far, the minimax approach suggests a promising way to construct a burden of proof that focuses on accuracy and does not require knowledge of base rates. Our task now is to consider what an optimal (minimax) burden of proof might look like in the legal context.

For most problems, the minimax test turns out to be tricky and difficult to find. Constructing the minimax test depends on the probability distributions involved, and the precise probability distribution for evidence at trial is neither obvious nor likely to be simple—indeed, it will often be unknown. We can, however, get a sense as to what a minimax test for legal proof might look like by modelling the proof process in the following way. Suppose that the truth in adjudication is a 'message' sent from nature. This message is either a 1 (plaintiff's story) or a 0 (defendant's story), but it is sent through a noisy channel that we call evidence—more specifically, the production, collection and presentation of evidence at trial. We can model this noise as 'Gaussian white noise', as is typically done in the signal processing literature.⁵³ The evidence presented to the fact-finder is therefore the original message obscured by white noise, as illustrated in Fig. 1.⁵⁴

As shown in Appendix B, under a Gaussian white noise model, the minimax rule has the following simple form:

$$\text{Choose } H_{\pi} \text{ if } \frac{P(E|H_{\pi})}{P(E|H_{\Delta})} > 1. \quad (3)$$

Now to re-emphasize, this minimax solution is strictly speaking only for a message passed through a Gaussian white noise channel, which the evidentiary process may or may not be empirically. But further consideration suggests that the model is a reasonable one, at least as a

⁵¹ Levy, *supra* note 46, at 39. Minimax testing embodies what is often called a 'robust' approach to testing. If we knew the base rate, the test described in (2) would result in less error than the minimax test. However, if we are wrong about that base rate, the performance of the 'optimal' test could go horribly awry. Rather than take this risk, we instead construct a minimax test, which guarantees a minimum level of performance regardless of the base rate.

⁵² As an anonymous reviewer rightly pointed out, the other methods previously discussed also treat the parties 'equally', though in different ways. For example, the posterior-probability ratio test in (1) treats plaintiff's and defendant's errors equally, though it may ultimately result in one party bearing a higher burden of errors than the other.

⁵³ Gaussian noise has all of the 'nice' characteristics that one might expect: it is normally distributed, so that extreme values are rare; it is symmetric; if there are multiple dimensions, deviations in one direction are independent from deviations in another; and it is independent of the message or signal involved.

⁵⁴ The signal and Gaussian white noise added under this model can be multi-dimensional, but for simplicity, we present only the one-dimensional model.




FIG. 1. Model of the evidentiary process.

global, rough approximation. Gaussian noise is symmetric, meaning that it favours neither party, and it has a bell-shaped distribution, meaning that it rarely creates major distortions. To the extent that we are designing a general-purpose burden-of-proof framework, assuming a fair and unbiased evidentiary process with relatively low distortion is natural, if not mandatory. If it turns out that resource and other inequities render the distortions large or asymmetric, then the legal system should adjust the evidentiary process itself (e.g. by making legal counsel more available or modifying procedural rules that benefit the well-heeled), not alter the burden of proof. Thus, (3) is arguably a good approximation for the minimax rule applicable to legal proof.

3.4 Preponderance as optimal

Notably, the decision rule offered above in (3) does not immediately look like the preponderance-of-the-evidence standard, but it in fact is. Courts of course often define the preponderance standard as ‘more likely than not’ or ‘ $p > 0.5$ ’. However, as previous work has argued, the ‘ $p > 0.5$ ’ formulation mischaracterizes the preponderance standard.⁵⁵ It fails to capture the comparative nature of evidentiary proof,⁵⁶ and also spawns a variety of undesirable probabilistic paradoxes.⁵⁷

Rather than viewing the preponderance standard as ‘ $p > 0.5$ ’, we should instead view it as a probability ratio, in which the fact-finder compares the probability of the plaintiff’s story with the defendant’s story:⁵⁸

$$\text{Choose } H_\pi \text{ if } \frac{P(H_\pi|E)}{P(H_\Delta|E)} > 1. \quad (4)$$

As we previously showed, however, through Bayes’s Rule, we can rewrite this decision rule as:

$$\text{Choose } H_\pi \text{ if } \frac{P(E|H_\pi)}{P(E|H_\Delta)} > \frac{P(H_\Delta)}{P(H_\pi)}. \quad (5)$$

⁵⁵ Edward K. Cheng, *Reconceptualizing the Burden of Proof*, 122 Yale L.J. 1254, 1278 (2013).

⁵⁶ See Pardo, *supra* note 38, at 592–94; Pardo and Allen, *supra* note 44, at 256–57; see also Pennington and Hastie, *supra* note 44.

⁵⁷ Cheng, *supra* note 55, at 1263–65, 1269–73.

⁵⁸ The probability ratio provides an analytical tool for describing the burden of proof and is consistent with inferential processes at trial that do not compare explicit probabilities (such as evaluating the epistemic strength of explanations or the plausibility of stories).

Previous work argues that the legal system normatively sets $P(H_\pi)$ and $P(H_\Delta)$ to be equal (out of fairness concerns), matching the optimal decision rule found in (3).⁵⁹ However, our minimax analysis suggests that the legal system need not normatively ‘set’ the prior probabilities at all.⁶⁰ Rather, the threshold value of 1 arises from the minimax criterion itself.⁶¹

4. Summary and conclusion

Having covered a great deal of ground in a short space, it may be useful to summarize our findings briefly. We began with a critique of *ex ante*, welfare-based approaches to burdens of proof. These *ex ante* approaches raise fundamental concerns. They rest on controversial and heavily disputed normative assumptions, are at odds with the adjudicative system’s commitment to accuracy and have insurmountable informational requirements. Having rejected *ex ante* approaches, we then developed a decision-theory framework to find the optimal burden of proof given an *ex post*, accuracy-focused perspective.

This decision-theory model, however, runs into a problem with base rates, because they too create unrealistic informational demands in most cases. We thus took a minimax approach, which obviates the need for base rates as a general matter by operating on a worst-case scenario basis and striving to minimize the maximum probability of error for either party. By modelling the evidentiary proof process as equivalent to introducing Gaussian white noise, we are able to solve the minimax problem and suggest that the optimal burden of proof is:

$$\text{Choose } H_\pi \text{ if } \frac{P(E|H_\pi)}{P(E|H_\Delta)} > 1.$$

Although it may not appear initially, this test turns out to be precisely the preponderance-of-the-evidence standard. We have thus not only shown a way to solve the optimization problem posed by legal proof, but we have provided a new justification for the traditional preponderance standard in minimax terms.⁶² If we focus on accuracy and want to have a rule that guarantees a minimum level of performance over a wide range of contexts (and base rates), then the preponderance standard may be the best we can do.⁶³

⁵⁹ Cheng, *supra* note 55, at 1267–68.

⁶⁰ This accords with a structure of proof at trial as ‘inference to the best explanation’. See Pardo and Allen, *supra* note 44, at 233–42. Each of the competing explanations proffered or developed at trial are evaluated on their own terms, based on their comparative epistemic strength, in explaining the evidence and underlying events.

⁶¹ We can make an additional interesting observation about the preponderance standard and these prior probabilities. Recall that the minimax rule is a worst-case scenario approach that minimizes the maximum error probability. The worst-case scenario, however, arises from a particular prior distribution, called the ‘least favorable prior’. Our analysis shows that the least favourable prior in the legal context is 0.5, which is equivalent to setting $P(H_\pi)$ and $P(H_\Delta)$ to be equal.

⁶² Kaplow, *supra* note 5, at 742–43 (‘[T]he primary argument for the preponderance rule seems to be a lack of any apparent reason to do otherwise.’). Other prominent defences of the preponderance rule, and alternative decision rules, have focused on full decision-theoretic considerations (as opposed to our constrained model), see, e.g. Neil Orloff and Jerry Stedinger, A Framework for Evaluating the Preponderance-of-the-Evidence Standard, 131 U. Penn. L. Rev. 1159 (1983); David Kaye, The Limits of the Preponderance of the Evidence Standard: Justifiably Naked Statistical Evidence and Multiple Causation, American Bar Foundation Research Journal 487 (1983), or on other normative considerations, see, e.g. Mike Redmayne, Standards of Proof in Civil Litigation, 62 Modern L. Rev. 167, 171 (1999) (‘Fact-finding in the typical civil case should therefore be governed by a principle of equality, and this provides a justification for the 0.5 rule.’).

⁶³ Our focus has been on the optimal burden of proof when the goal is ‘total’ accuracy and given the various limitations the legal system faces. A further question is when the system should adopt a higher standard—such as ‘beyond a reasonable doubt’ or ‘clear and convincing evidence’—because of the asymmetric interests at stake (e.g. those of criminal defendants and civil parties facing involuntary commitment or termination of parental rights). This question has been, for the most part, outside the scope of

		Conclusion	
		Γ_π	Γ_Δ
Truth	H_π	C_{TP}	C_{FN}
	H_Δ	C_{FP}	C_{TN}

FIG. 2. Social costs of decisions.

Acknowledgements

Thanks to Ron Allen, Samuel Bray, Jenny Diamond Cheng, Paul Edelman, Joe Gastwirth, Jennifer Mnookin, Meredith Render, Alex Stein, Fred Vars, the UCLA Law, Science, and Evidence Colloquium, and an anonymous reviewer for helpful comments, and Bill Walker for library assistance.

Technical Appendices

A. General Decision-Theory Model

To review the discussion in Part 2, a legal fact-finder confronts two possible truths: plaintiff's story, H_π , and defendant's story, H_Δ . Based on the evidence, the fact-finder makes one of two decisions: a finding for the plaintiff, Γ_π , or a finding for the defendant, Γ_Δ . As seen in Fig. 2, there are therefore four possible scenarios: (1) finding for the plaintiff when the defendant's story is true, labelled as a 'false positive' (FP); (2) finding for the defendant when the plaintiff's story is true, labelled as a 'false negative' (FN); (3) finding for the plaintiff when the plaintiff's story is true, labelled as a 'true positive' (TP); and finding for the defendant when the defendant's story is true, labelled as a 'true negative' (TN). Figure 2 depicts the costs associated with the four possible scenarios using these labels. Note that under this scheme, true positives and true negatives presumably have negative costs, i.e. $C_{TP}, C_{TN} < 0$.

Broadly speaking, the goal in decision theory is to minimize the expected costs of the legal decision given the evidence presented. If we find for the plaintiff (Γ_π), then the expected costs are the probability of being wrong times the costs of being wrong, plus the probability of being right times the (negative) 'cost' of being right. In other words,

$$E[\text{Cost}(\Gamma_\pi)] = P(H_\Delta|E)C_{FP} + P(H_\pi|E)C_{TP},$$

where $P(H|E)$ is the probability of a particular story or hypothesis being true having observed the evidence. Similarly, if we find for the defendant (Γ_Δ), then the expected costs are

$$E[\text{Cost}(\Gamma_\Delta)] = P(H_\pi|E)C_{FN} + P(H_\Delta|E)C_{TN}.$$

And thus to minimize expected costs, we should find for the plaintiff when the expected error costs of finding for the plaintiff are less than the expected error costs of finding for the defendant.

our analysis, but we briefly mention two general points that follow from our analysis. First, such a decision is a decision to sacrifice total accuracy in lieu of these asymmetric (and often justified) goals. Second, the law has largely been correct to do it in a coarse-grained, categorical approach, given the many limitations we have discussed with more fine-grained approaches. This categorical approach is justified even when doing so deviates from the normative prescriptions of welfare theory. But more detailed discussion of these other standards is a topic for future work.

In other words, when

$$P(H_{\Delta}|E)C_{FP} + P(H_{\pi}|E)C_{TP} < P(H_{\pi}|E)C_{FN} + P(H_{\Delta}|E)C_{TN}$$

$$P(H_{\Delta}|E)(C_{FP} - C_{TN}) < P(H_{\pi}|E)(C_{FN} - C_{TP})$$

$$\frac{P(H_{\pi}|E)}{P(H_{\Delta}|E)} > \frac{C_{FP} - C_{TN}}{C_{FN} - C_{TP}}.$$

Now, the problem is that we do not know the probability of the defendant's or plaintiff's story given the evidence. To determine these probabilities, one can, at least in theory, use Bayes's Theorem, which states that:

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

and

$$P(H_{\Delta}|E) = \frac{P(E|H_{\Delta})}{P(E)} \times P(H_{\Delta}).$$

Substituting these expressions and simplifying, we see that to minimize the expected error costs, our decision rule is:

$$\text{Choose } H_{\pi} \text{ if } \frac{P(E|H_{\pi})}{P(E|H_{\Delta})} > \frac{C_{FP} - C_{TN}}{C_{FN} - C_{TP}} \times \frac{P(H_{\Delta})}{P(H_{\pi})}. \quad (6)$$

Note that from this general optimal decision rule, one can derive the accuracy-focused optimal decision rule in the main text (Equation (2)). If the judicial system values truth and accuracy over broader social benefits and cost, then the cost function can be assumed to have important symmetries. Specifically, the 'cost' of an error, whether false positive or false positive, should be the same—an error is an error regardless of its direction. Similarly, the 'benefit' of true positives and true negatives should be equivalent. Under these constraints, the cost-weighting factor in the general decision rule cancels out,⁶⁴ leaving us with

$$\text{Choose } H_{\pi} \text{ if } \frac{P(E|H_{\pi})}{P(E|H_{\Delta})} > \frac{P(H_{\Delta})}{P(H_{\pi})}.$$

B. *Minimax Test Derivation*⁶⁵

In this Appendix, we will show that the minimax decision rule for legal proof approximates the preponderance standard defined in (3), namely:

$$\text{Choose } H_{\pi} \text{ if } \frac{P(E|H_{\pi})}{P(E|H_{\Delta})} > 1. \quad (3)$$

⁶⁴ Specifically, $\frac{C_{FP} - C_{TN}}{C_{FN} - C_{TP}} = 1$.

⁶⁵ The derivation presented here is adapted and simplified from the proofs found in Poor, *supra* note 49, at 19–27, and Levy, *supra* note 46, at 39–45.

B.1 Minimax definition. Before beginning the derivation, let us first unpack the definition of a minimax test. Section 3.2.2 asserted that the minimax decision rule minimizes the maximum expected error (when we do not know the applicable base rates). It also noted that the minimax test minimizes the maximum error faced by either party. In this Section, we define these criteria more precisely and show that they are in fact equivalent.

Let us start with the first criterion, which is that the minimax test minimizes the maximum expected error over all possible priors. Consider the expected error risk for a given test δ :

$$r(p, \delta) = pR_{\pi}(\delta) + (1 - p)R_{\Delta}(\delta),$$

where p denotes the prior probability that H_{π} is true, $R_{\pi}(\delta)$ is the (conditional) risk of error when H_{π} is true and $R_{\Delta}(\delta)$ is the (conditional) risk of error when H_{Δ} is true. We of course do not know the value of p (i.e. the whole point of minimax testing), but it is an important placeholder in the calculations.

Now let us consider this expected error risk for different values of the prior p . When $p = 1$, then $r(1, \delta) = R_{\pi}(\delta)$. When $p = 0$, then $r(0, \delta) = R_{\Delta}(\delta)$. So for a given test δ , the maximum value for $r(p, \delta)$ either occurs at $p = 1$ or $p = 0$.⁶⁶ Further, this maximum value equals $\max(R_{\pi}(\delta), R_{\Delta}(\delta))$.⁶⁷

Thus, if we minimize $\max_p r(p, \delta)$ over all possible tests δ , we also minimize $\max(R_{\pi}(\delta), R_{\Delta}(\delta))$, which is the maximum (conditional) risk of error faced by either party. The two minimax criteria are therefore equivalent.

B.2 Finding the minimax rule for a Gaussian white noise model. In accord with the discussion above, in seeking the minimax rule, we are looking for:

$$\delta_M = \arg \min_{\delta} (\max_p r(p, \delta)).$$

As described in Appendix A, for every prior $p = P(H_{\pi})$, there is an optimal (Bayes) test, δ_p^* , which minimizes the expected error. Each of these tests has an expected error risk, which we will label $\mathcal{R}(p)$:

$$\mathcal{R}(p) = r(p, \delta_p^*) = pR_{\pi}(\delta_p^*) + (1 - p)R_{\Delta}(\delta_p^*).$$

With that preliminary definition in hand, we can now present a well-known result in minimax theory, known as the Equalizer Rule. It states the following:

THEOREM 1 (Equalizer Rule): Suppose p_{LF} is the least favourable prior, namely that $\mathcal{R}(p_{LF}) = \max_p \mathcal{R}(p)$. Suppose further that either $p_{LF} = 0$, $p_{LF} = 1$, or $R_{\pi}(\delta_{p_{LF}}^*) = R_{\Delta}(\delta_{p_{LF}}^*)$. Then $\delta_{p_{LF}}^*$ is a minimax rule.⁶⁸

The proof of this theorem is rather involved and omitted for brevity. However, the practical implication of the Equalizer Rule is that it often provides a handy construction method for minimax rules. Assuming that one knows the form and distribution of the optimal (Bayes) test, δ_p^* , one can set the conditional risks of error to be equal and solve for the threshold that equalizes the risks.⁶⁹ Notably,

⁶⁶ Any other prior value will result in a value between $R_{\pi}(\delta)$ and $R_{\Delta}(\delta)$, which obviously cannot be the maximum.

⁶⁷ Poor, supra note 49, at 19–20.

⁶⁸ Poor, supra note 49, at 22.

⁶⁹ Strictly speaking, as seen in the theorem, the minimax rule can also occur at the endpoints (i.e. $p = 0$ and $p = 1$), but these situations are rare as a practical matter. See Poor, supra note 49, at 22–27 for a full treatment.




FIG. 3. Model of the evidentiary process (revisited).

having the conditional risk of error given H_π equal to the conditional risk of error given H_Δ is a desirable (and fair) attribute for a legal decision rule.

To find the minimax test for the Gaussian white noise model, recall the setup from Section 3.3, as seen in Fig. 3.⁷⁰ Under this stylized setup, we have two hypotheses. If plaintiff's story is true, then the evidence will be distributed around 1. If the defendant's story is true, then around 0:

$$H_\pi : Y \sim \mathcal{N}(1, \sigma^2)$$

$$H_\delta : Y \sim \mathcal{N}(0, \sigma^2).$$

With these hypotheses in mind, we can solve for the form of the optimal (Bayes) test, which after some algebraic manipulation becomes:

$$\frac{P(Y|H_\pi)}{P(Y|H_\delta)} = \exp\left\{\frac{Y}{\sigma^2} - \frac{1}{2\sigma^2}\right\} \underset{H_\Delta}{\underset{H_\pi}{\geq}} \tau,$$

which is equivalent to

$$\underset{H_\pi}{Y} \underset{H_\Delta}{\geq} \eta,$$

where $\eta = \frac{1}{2} + \sigma^2 \ln(\tau)$.

Now, given that we have 'nice' Gaussian distributions, we can directly calculate the conditional risks of error given the threshold τ , or equivalently, η .

$$R_\pi(\delta) = \int_{-\infty}^{\eta} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(y-1)^2}{2\sigma^2}\right) dy$$

$$R_\Delta(\delta) = \int_{\eta}^{\infty} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{y^2}{2\sigma^2}\right) dy.$$

⁷⁰ This discussion is derived from Levy, *supra* note 46, at 34–36, 45.

Using the symmetry of Gaussian distribution and a change of variable, these expressions can be re-written in terms of the standard Gaussian Q function:

$$R_{\pi}(\delta) = Q\left(\frac{1}{2\sigma} - \sigma \ln(\tau)\right)$$

$$R_{\Delta}(\delta) = Q\left(\frac{1}{2\sigma} + \sigma \ln(\tau)\right)$$

and now we can solve for the desired threshold, τ , using the Equalizer Rule.⁷¹

$$R_{\pi}(\delta_{p_{LF}}^*) = R_{\Delta}(\delta_{p_{LF}}^*)$$

$$Q\left(\frac{1}{2\sigma} - \sigma \ln(\tau)\right) = Q\left(\frac{1}{2\sigma} + \sigma \ln(\tau)\right)$$

$$2\sigma \ln(\tau) = 0$$

$$\tau = 1.$$

So the minimax test has the form

$$\frac{P(Y|H_{\pi})}{P(Y|H_{\delta})} = \exp\left\{\frac{Y}{\sigma^2} - \frac{1}{2\sigma^2}\right\} \underset{H_{\Delta}}{\overset{H_{\pi}}{\geq}} 1,$$

and we therefore have a minimax decision rule in the Gaussian white noise context of:

$$\delta_{p_{LF}}^* : \text{Choose } H_{\pi} \text{ if } \frac{P(E|H_{\pi})}{P(E|H_{\Delta})} > 1.$$

⁷¹ In the case of Gaussian white noise, it turns out that $\mathcal{R}(0) = \mathcal{R}(1) = 0$, which means that the least favourable prior, p_{LF} occurs between the endpoints. We can thus look for the minimax test by equating the conditional risks. For more details, see Poor, supra note 49, at 27–28.