

2012

Law, Economics, and the Burden(s) of Proof

Eric L. Talley
Columbia Law School, etalley@law.columbia.edu

Follow this and additional works at: https://scholarship.law.columbia.edu/faculty_scholarship



Part of the [Criminal Law Commons](#), [Law and Economics Commons](#), and the [Torts Commons](#)

Recommended Citation

Eric L. Talley, *Law, Economics, and the Burden(s) of Proof*, RESEARCH HANDBOOK ON THE ECONOMICS OF TORTS, JENNIFER ARLEN, ED., EDWARD ELGAR, 2013; UC BERKELEY PUBLIC LAW RESEARCH PAPER No. 2170469 (2012).
Available at: https://scholarship.law.columbia.edu/faculty_scholarship/1768

This Working Paper is brought to you for free and open access by the Faculty Publications at Scholarship Archive. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarship Archive. For more information, please contact scholarshiparchive@law.columbia.edu.

Law, Economics, and the Burden(s) of Proof

Eric L. Talley¹

Forthcoming in *Research Handbook on the Economic Analysis of Tort Law* (J. Arlen Ed, 2013)

Version 2.4, November 5, 2012

(Original version: March, 2012)

Contents

1. Introduction.....	3
2. Definitional Dark Matter.....	6
3. Law and Economics Analysis of the Burden of Proof.....	10
<i>A. The Normative/Positive Spectrum</i>	11
<i>B. Modeling Frameworks for the BoP</i>	13
1. <i>Decision Theoretic Frameworks</i>	13
2. <i>Game Theoretic Frameworks</i>	21
3. <i>Hybrids: Game Theoretic Frameworks with a Bayesian Fact-Finder</i>	27
4. <i>Mechanism Design Approaches</i>	30
<i>C. Scope of Inquiry: Ex Ante, Litigation, or Both?</i>	34
4. Moving Forward	38
References.....	40

¹ Rosalinde & Arthur Gilbert Foundation Professor of Law, Business and the Economy, UC Berkeley School of Law. Many thanks to Jennifer Arlen, Richard Holden, Andrea Roth, Chris Sanchirico, three anonymous referees, and seminar participants at the Australian Business School (UNSW), University of Bonn, and UC Berkeley for helpful comments and discussions. Samuel Edwards, Tristan Favro and Samantha Strimling provided excellent research assistance. All remaining errors are my own.

Law, Economics, and the Burden(s) of Proof

Eric L. Talley

Forthcoming in *Research Handbook on the Economic Analysis of Tort Law* (J. Arlen Ed, 2013)

Version 2.3, November 1, 2012

(Original version: March, 2012)

ABSTRACT: This chapter presents an overview of the theoretical law and economics literature on the burden of proof within tort law. I begin by clarifying core legal definitions within this topic, demonstrating that the *burden of proof* actually refers to at least five doctrinal concepts that substantially overlap but are not completely interchangeable. I then provide a conceptual roadmap for analyzing the major extant contributions to this topic within theoretical law and economics, emphasizing three key dimensions that organize them: (a) where they fall in the positive-normative spectrum; (b) what type of underlying modeling framework they employ (ranging from decision theoretic to game theoretic to mechanism design); and (c) whether they focus on litigation activity or primary activities (or both). In the aggregate, the resulting theoretical landscape is a complex one, yielding a number of interesting insights. Yet it still suffers from having no single unified theory. I conclude by offering a number of recommendations about where applied law and economics scholars interested this topic could direct their research efforts.

1. Introduction

Since nearly its very inception, law and economics (L&E) scholarship has contributed important and valuable insights about how “substantive” legal rules in general—and tort law in particular—can affect behavior and economic welfare. Like other vehicles of regulation and taxes, the contours of tort law (such as negligence standards, affirmative defenses, damages measurement, and the like) directly distort individual incentives, risk allocations, activity levels, cost realizations and wealth/income distributions. The collective observations of L&E scholars about the efficiency attributes of tort law are real and rich, and they have informed legal policy-making and reform efforts for much of the last half century.

That said, arguably a more distinct contribution of L&E scholarship on torts may lie *not* with its considerable insights about substantive law (an analytic approach common to much of welfare economics), but rather with its (arguably) less heralded insights about the *procedural rules* through which law operates. Unlike conventional mechanisms of public policy, the legal system embodies an idiosyncratic set of traditional (and highly cherished) rules and processes—many of which themselves introduce special forms of incentives, risks, and strategic behavior among litigants and other stakeholders. Consequently, the application of legal process through courts (and other quasi-judicial actors) bears a heavy hand in mediating both the delivery and consequential impact of substantive law.

It is therefore hardly surprising that the last four decades have spawned a rich sub-literature in L&E centering on how the legal process itself—including the rules, protocols and traditions to which it subscribes—operates to alter outcomes of litigation, and in so doing incentives and allocations in ways that could either promote or frustrate the operation and policy goals of substantive law. Any efficiency-minded legal reformer advocating prescriptive reform

to substantive legal rules who ignores the procedural implementation of such reforms does so at her peril.

Within the broad array of topics germane to the legal process, there is probably none so central as the allocation of the (so-called) “Burden of Proof” (BoP) in litigation. Although it may appear upon first blush to bear the marker of a “mere” procedural formality, the BoP is far from it. To the contrary, the BoP embodies the very decision-making structure that animates and defines legal order itself. Indeed, as most judges and practitioners are fond of observing, in many practical instances these seeming procedural formalities often substantially determine the outcome of a case.

This criticality of process in general—and the BoP in particular—may be especially salient in cases where the economic stakes are high, and where the informational environment is complex, opaque, and difficult to navigate. In such situations, it is plausible that no single entity—not the jury, not the judge, not attorneys, not the parties themselves—has full and complete command of all the “facts” pertinent to a legal dispute. It is here where the burden allocation may be the most influential in catalyzing information discovery, reducing verification costs, and ultimately contributing to overall welfare policy goals.

This chapter reviews the conceptual nature of the BoP, and organizes and catalogues the L&E contributions to this literature, assessing how law and economics scholarship has both sharpened and complicated our understanding of how evidentiary burdens plausibly operate in practice, how they should operate, and how economic welfare-minded policy-makers can incorporate such insights into practical prescriptive reforms.² In the end, I conclude that the literature on BoP has spawned a large number of interesting and durable insights about how

² In the balance of this chapter, I will tend to use the term “economic welfare” rather than “efficiency” under the premise that the former is a more general term, capable of incorporating (for example) distributive desiderata as well as conventional efficiency-oriented aims.

substantive and procedural law interact with one another. At yet, at the same time, the range of methodologies and objectives reflected in this literature is still so broad as to elude easy categorization or simple prescriptive sound bites. Consequently, I conjecture, the L&E literature on the burden of proof is likely to make significant contributions to practical legal policy only when L&E scholars are better able to unify, test and, for want of a better term, “market” its insights to practitioners.

One caveat deserves explicit mention before proceeding. My focus in this essay will predominantly center on *theoretical* models of the BoP within the relevant law and economics literature (rather than empirical or experimental contributions). This is in part because of unavoidable space constraints, and in part because most of the significant contributions heretofore have come through theory. It is likely that the next phases of inquiry will be more evenly divided, however, between theory, empiricism, and experimentalism. Accordingly, at the end of this essay when I chart a course for future research, my recommendations will tend to fall along all three methodological dimensions.

The remainder of this chapter is organized as follows. Section II provides a brief conceptual overview of the core legal concepts that surround the burden of proof, and how they are conceived of within traditional legal parlance. Here I illustrate that, far from being a monolithic concept, the “burden” of proof is more akin to a composite of at least five interlocking concepts.³ The precise application and interaction of these concepts, however, remain somewhat elusive notwithstanding an immense literature discussing and dissecting them—a state of play that simultaneously confounds and invites economic analysis. Section III turns to the economic analysis of the burden of proof, and catalogues major L&E contributions to

³ Those are: the burden of production, the burden of persuasion, evidentiary standards, and legal presumptions and legal assumptions. I do not classify (nor do most other commentators) the “burden of pleading” as part of this group.

the field along three different dimensions: (a) The normative/positive spectrum; (b) modeling approaches employed; and (c) scope of inquiry. Section IV concludes, offering thoughts for future work in the field.

2. Definitional Dark Matter

Before reviewing the law and economics literature on the topic, it is necessary first to understand—or at least appreciate—the rough doctrinal contours of the *burden of proof*. This is no mean feat: As recently as 2011, the U.S. Supreme Court described the BoP as “one of the ‘the slipperiest members of the family of legal terms.’ ” *Microsoft Corp. v. i4i Ltd. P’ship* (2011); *Schaffer ex rel Schaffer v. Weast* (2005).

The doctrinal slipperiness of the BoP has at least two important implications for economic analysis. First, it presents a challenge for those who aspire to model the burden of proof theoretically, and particularly those who wish to posit a model that comports (at least to some extent) with existing legal practices.⁴ The task of such efforts clearly is made more difficult by the lack of a consistent, precise verbal definition among many legal practitioners. Second, it also presents something of a constructive opportunity for law and economics scholars, whose efforts might provide insights into what a welfare-minded BoP *should* look like. Such contributions may be of significant assistance in helping scholars of torts, evidence and procedure develop sharper and more precise definitions of the relevant terms.

In practice,⁵ the term “burden of proof” bundles together at least five inter-related phenomena, which scholars, practitioners and laypeople alike tend intermittently to reorganize, redefine, contest and conflate. The first—and most intuitive—phenomenon is commonly

⁴ As discussed below, the extent to which an economic / game theoretic model of legal process should “fit” extant practices depends on the objectives of the scholarly enterprise.

⁵ The description in the text comes from the Model Code of Evidence’s definitions, but in some states—such as California—the term “Burden of Proof” denotes what is more widely identified as the *Burden of Persuasion*.

referred to the *Burden of Persuasion*. This burden articulates how, once all evidence is submitted, courts will weigh the offered evidence from either side, deciding remaining uncertainties about contested claims. The party charged with the burden of persuasion legitimately carries a “burden,” because it is she who must present enough probative evidence to convince the fact-finder that her case is sufficiently “strong” to warrant a verdict in her favor.⁶ Should the plaintiff fail in this regard, by implication, she has failed to satisfy her Burden of Persuasion, and the other side prevails.⁷

As a default rule in most civil cases within Anglo-American courts, burden of persuading the fact finder that liability is present generally rests with the plaintiff. Even without economic analysis to justify it, this default allocation of the burden of persuasion makes great intuitive sense as a litigation- and enforcement-cost saving device, since it is the plaintiff who (usually) advocates a change from the status quo ante (at least before litigation).⁸ As a general matter, the party bearing the burden of persuasion remains constant during litigation (though affirmative defenses raise interesting analytic puzzles, as discussed below).

The second related term is the “Standard of Proof.”⁹ This term specifies criteria for applying the burden of persuasion. In other words, it specifies how far the party bearing the burden of persuasion must push the judicial fact-finder’s assessment in order to prevail in her legal claim. (See *Microsoft Corp. v. i4i Ltd. P’ship*, at n.4). Borrowing on a Bayesian/Quasi-Bayesian account (to which I return again in later sections),¹⁰ there are many different standards

⁶ In a canonical negligence claim in torts, this involves proving multiple elements, including the existence of a duty, its breach by lack of due care, actual causation of injury, proximate causation of the injury, and resulting damages. Whether the plaintiff must establish these elements one by one versus jointly is a topic I take up below.

⁷ The party bearing the burden of persuasion is sometimes said analytically to also bear the concomitant risk of non-persuasion. See, e.g., Winter (1971).

⁸ *Id.*

⁹ This is sometimes referred to alternatively as the “evidence threshold” (Kaplow 2012).

¹⁰ The term “Bayesian” refers to the process by which a statistically-minded decision maker marshals available information to update her probabilistic beliefs about the world. In the present context, the relevant decision maker is

of proof in litigation, ranging—in ascending order—from “preponderance of evidence” common in civil litigation (something akin to a 51% Bayesian confidence assessment), to “clear and convincing evidence” (perhaps on the order of 70-80%), to “beyond a reasonable doubt” common to criminal proceedings (perhaps on the order of 95% or more).

The third related term is the *Burden of Production*. This is a procedural rule identifying which party is required to submit information or evidence about the case if she wishes to alter a (possibly interim) legal conclusion in a she desires. At least initially, in most civil litigation cases, the burden of production and the burden of persuasion both fall on the plaintiff.

Nevertheless, the Burden of Production differs from the Burden of Persuasion in at least two important respects. First, unlike the Burden of Persuasion, which explicitly weighs previously submitted evidence of one party against that of the other, the Burden of Production is what regulates the (usually) sequential process by which those pieces of evidence are provided to the fact-finder. When a plaintiff bearing the Burden of Persuasion brings forth evidence which—if wholly uncontroverted by the defendant—would satisfy the Standard of Proof, then the plaintiff is frequently said to have made out a *prima facie* case, satisfying her burden of production, and allowing the case to be decided by the fact-finder. However, that is not the end of the case because the defendant now has the opportunity to refute the plaintiff’s evidence or

a judicial fact-finder (either jury or judge). The process of updating follows what is known as “Bayes’ rule,” which states (for the case of discrete random variables) that the probability of an event A occurring, conditional on knowing that some other event B has occurred (or $Pr[A|B]$) can be derived from a combination of the respective “base rate” probabilities of A and B (or $Pr[A]$ and $Pr[B]$) and the “reverse conditional” probability that B occurs, conditional on knowing that A has occurred (or $Pr[B|A]$). These four probabilities are related to one another according to the following expression:

$$Pr[A|B] = \frac{Pr[B|A] \times Pr[A]}{Pr[B]}$$

For example, suppose that one were attempting to use the outcome of a diagnostic test to infer whether a medical patient was carrying a deadly virus. And suppose further that when administered to people who are known to have the virus, the test yields a positive result 50% of the time. Moreover, among the general population, the test yields a positive result 25% of the time, and 10% of the general population carries the virus. Using Bayes’ rule, it is possible to calculate the probability that an individual who has tested positive also carries the virus is equal to 20%, twice the unconditional base rate in the population (but still less than 50% likely that she has the disease).

adduce contrary evidence of her own. Thus, while the Burden of Persuasion as to a particular claim remains largely constant through litigation, the Burden of Production is sometimes said to oscillate in contingent fashion as each side proffers its evidence.¹¹

The second major difference is that unlike the Burden of Persuasion (where the ultimate question of satisfying the burden is up to the fact-finder), the assignment and assessment of the Burden of Production is solely the province of the judge. Failure of a plaintiff to adduce evidence satisfying her initial Burden of Production, for example, can justify a dismissal or directed verdict in the defendant's favor without ever being put to a jury. A similar implication frequently may work in the reverse direction, particularly when a defendant asserts affirmative defense (such as claiming contributory negligence by the plaintiff in a torts case). As to this affirmative defense, the defendant would bear (at least initially) the burdens of persuasion and production. Should she fail to produce sufficient evidence to substantiate the defense, the judge will instruct the jury to ignore it.

The final "nearby" definition concerns the overlapping terms of *legal assumptions* and *presumptions*. These are perhaps the most slippery concepts in the field, and their precise interpretations (and particularly that of *presumptions*) remain inconsistent and heavily contested. Under the simplest definitions, assumptions and presumptions are respectively the default conditions – and thus the analytic duals – for the burdens of persuasion and production. *Assumptions* fix initial conditions in light of the Burden of Persuasion. In typical cases (civil and criminal), the initial assumption is that the defendant is not liable, and the act of dislodging this assumption is

¹¹ See, e.g., Wigmore (1940) §§ 2489; 2494. This "shifting" view of the burden of production is frequently disputed among evidence scholars, with some viewing the burden of production as nothing more than a lower bound for what evidentiary showing is necessary for the petitioning party's case to go to the fact-finder. Accordingly, under this account, should the petitioning party satisfy her burden of production, the "burden" does not shift, as much as the non-petitioning side now bears a risk of liability if she does not endeavor to counter the petitioner's evidence. See, e.g., Wright et al. (2012) § 5122. Declaring the victor of this battle of the evidence-wonk bands is, thankfully, beyond the scope of this chapter.

coterminous with carrying one's burden of persuasion (which, as noted above, remains fixed through trial¹²). *Presumptions*, in contrast, fix initial conditions in light of the Burden of Production at various stages of trial, identifying which party must bear burden of going forward with evidence. "Rebutting" a presumption is therefore equivalent under this view to satisfying one's burden of production as to a disputed issue. Thus, as with the Burden of Production, presumptions are often said to shift back and forth as trial proceeds.¹³

As should be obvious from this brief overview, there is a rich and longstanding debate within evidence and civil procedure regarding the definition (and redefinition) of the elements that collectively constitute the burden of proof. And with that debate is a significant literature on the topic, which takes a number of different perspectives ranging from formalistic to historical to expressive to positive to normative.¹⁴ The balance of this essay cannot do justice to the totality of this literature, and thus concentrates solely on positive and policy-oriented approaches that utilize economic analysis.

3. Law and Economics Analysis of the Burden of Proof

As noted above, the contested and sometimes elusive understanding of the Burden of Proof in legal discourse presents both a challenge for law and economics scholars writing in the field, as well as an opportunity for them (and others) to attempt to clarify the landscape. The last four decades have borne witness to an interesting theoretical literature in law and economics on the topic. Perhaps channeling the topic's complex doctrinal landscape, L&E contributions in the

¹² It is important to distinguish statement that the defendant is "assumed not liable" with the (stronger) Bayesian assertion that one's prior probability assessment of the defendant's liability is zero. As described in greater detail *infra* at n.24, the latter statement would produce a pathological result in which – because of an inflexible prior probability assessment – *no amount* of evidence could ever be persuasive enough to lead to a finding of liability.

¹³ As noted above, the term "presumption" often is invoked to sweep in concepts even broader than what is stated in the text. For example, the term is at times employed to describe the burden of persuasion in an affirmative defense, or alternatively as tie-breaking decisional rules in cases where the evidence is unclear. See Allen (1981) for an overview, as well as a critique of the highly elusive definition of *presumption*.

¹⁴ Stein (2005) provides an excellent overview.

field have similarly tended to be multifaceted and heterogeneous, undertaking distinct and sometimes divergent approaches, depending on the goals of the research, the modeling decisions of the researcher, and the nature of the inquiry. This section considers that heterogeneity along three different dimensions: (1) the contribution's place on positive/normative spectrum; (2) the contribution's theoretical modeling approach; and (3) the contribution's place on the substantive law/legal process spectrum.

A. The Normative/Positive Spectrum

Like many other endeavors in the economic analysis of law, most contributions can be classified as either *positive* or *normative*. Positive contributions within this area tend to take the existing BoP structures as given, analyzing how they shape various types of behavior (both during litigation and before), as well as generating falsifiable predictions about how various “shocks” or changes in BoP structures might alter that behavior. Normative contributions, in contrast, tend to focus on desirable behavior (usually against a posited economic welfare measure), and the consequent social desirability of the status quo against a host of potential alternative institutional designs.

Each of the approaches described above comes entail advantages and disadvantages, as well as different modeling desiderata. Scholars interested in making purely (or primarily) positive contributions tend to focus on models whose details correspond (at least loosely) with the key institutional features of the BoP as articulated and practiced by courts. For example, all else constant, a positive model would preferably capture the central characteristics of burdens of persuasion and production, standards of proof, assumptions and presumptions as outlined in Section 2. This makes perfect sense, since the goal of positive theory is to make predictions

about how existing institutions operates, thereby placing a premium on fidelity to institutional detail.

For those interested in making purely (or predominantly) normative contributions, in contrast, institutional fidelity may be singularly unattractive, since it significantly constrains the task of designing and implementing wholesale reform. Normative analyses, therefore, often tend to eschew modeling frameworks that unduly fetishize prevailing practice at the expense of prescriptive possibility. Indeed, a key strength of normative analysis is its ability to evaluate longstanding institutional processes without assuming those traditions to be sacrosanct.¹⁵

Both pure positive and pure normative approaches also have weaknesses, however. Pure positive analysis may be unsatisfying to reform minded readers, while pure normative approaches may be too fanciful or abstract to implement in practice. Perhaps recognizing these limitations, a third strand of the law and economics literature on burdens attempts (rather successfully) to occupy a prominent middle ground. Such hybrid contributions, while seeking ultimately to deliver normative policy prescriptions, are also practically oriented around incremental improvements, suggesting tweaks to and reforms of existing institutional practices rather than a wholesale reinvention of litigation systems. Consequently, theoretical contributions within this tradition necessarily must balance incorporating some core institutional details from evidence and civil procedure, leaving some degrees of freedom to consider the designs of central “parameters” of that system. Because each approach described above may be appropriate in certain situations, it should not be surprising that each—the purely positive, purely normative, and hybridized—is now quite common within the law and economics literature on the BoP.

¹⁵ Sometimes the object may also be to rationalize current practices by demonstrating that they are “optimal” in an economic welfare sense (or approximately so). But even here, to be convincing at rationalizing current practice, the underlying model must be capable to generalize practices beyond the status quo ante.

B. Modeling Frameworks for the BoP

A second significant dimension of heterogeneity in the law and economics literature on the BoP concerns the general conceptual framework that scholars adopt to analyze its operation. Although myriad variations exist, they most naturally fall into four distinct clusters: Decision-theoretic analyses, Game-Theoretic frameworks, hybrid Game-Theoretic frameworks involving a Bayesian Fact-finder; and Mechanism-Design approaches.

1. Decision Theoretic Frameworks

Perhaps the simplest and most intuitive branch of the L&E scholarship on burdens adopts what one might call *decision-theoretic* perspective (Kaplan 1968; Easterbrook 1984; Hylton & Salinger 2001). This approach trains exclusive (or near exclusive) focus on the decision problem facing the finder of fact, weighing the previously-submitted evidence in a case. Under this approach, the fact-finder's role is tantamount to a welfare-minded social planner, using adduced evidence to learn and make decisions within an information-constrained environment. The evidence produced in a case enables an uncertain fact-finder to "update" her assessment of the case (possibly in a Bayesian manner) in light of that evidence (Feess et al. 2009).

For example, suppose a court were attempting to assess whether a defendant had exercised due care in an automobile accident, and that the substantive negligence standard was triggered only if the defendant's precautions, denoted as $v \in [0, \infty)$, fell below some articulated negligence standard $v^* > 0$.¹⁶ This would, of course, be an easy task to accomplish if the fact-finder could directly observe the driver's precautions. The more interesting case occurs when the driver's precautions are not directly observable, so that the fact finder must resort to looking towards observable "evidence" about v , such as the driver's evident speed, whether his cell phone was in use in the time leading up to the accident, his driving record, and the like.

¹⁶ For simplicity, assume that v is a scalar value. The analysis generalizes to vector-valued precautions.

Accordingly, denote this observable evidence as $x \in X$, and suppose this evidence provides noisy information about v . Formally, suppose that v and x are distributed according to a prior density function $f(v,x)$, with associated marginal densities of $g(v)$ and $h(x)$ for v and x , respectively.¹⁷ Having observed evidence x , the fact-finder can update its assessment in any number of ways. Perhaps the most intuitive updating process might be Bayes rule (e.g., Feess et al. 2009), under which the fact-finder uses the evidence observed to generate a posterior conditional probability density on the defendant’s actual behavior, $f(v|x) = \frac{f(v,x)}{h(x)}$. (Note that so long as v and x are not completely statistically independent of one another, the evidence represented by x is statistically “probative,” so that this updated density will generally not coincide with the marginal distribution on v , or $g(v)$).

The decision-maker could then utilize this information to assign liability according to the underlying substantive rule, perhaps weighing the relative costs associated with false positives (finding that the defendant was liable when in fact she met the negligence standard) and false negatives (finding that the defendant was not liable when in fact she fell short of the standard). For concreteness, suppose the decision maker cares about making “accurate” judgments, and that she receives a payoff normalized to be 0 whenever she correctly finds that the defendant was negligent (a “true positive”) or when she correctly finds the defendant was not negligent (a “true negative”). On the other hand, the decision maker incurs a welfare loss of e_l whenever she

¹⁷ A *density function*, denoted $f(y)$, is a standard way to express the probabilistic behavior of some continuously distributed random variable y , whose realizations – unlike discrete random variables – do not fall into countable, discrete outcomes. Because there are infinitely many realizations for a continuous random variable, the probability of realizing any specific realization is effectively zero. Nevertheless, one can characterize the probability of the outcome falling in some range $[a,b]$: This is simply the area under the density function between a and b . It is common to express the area under the density function to the left of some prescribed cutoff b using a *cumulative distribution function*, which is often denoted as $F(b) = \int_{-\infty}^b f(y)dy$. Consequently, the probability of the outcome falling in range $[a,b]$ could be expressed using this notation as $F(b)-F(a)$. Perhaps the most well-known continuous density function is the Gaussian, or so-called “normal” distribution of a variable with mean μ and variance σ^2 . The density function for a normally distributed variable z is given by: $f(z) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(z-\mu)^2}{\sigma^2}}$.

incorrectly finds negligence when the defendant took due care (a “false positive”) and a loss of e_2 whenever she incorrectly finds non-negligence when the defendant was actually negligent (a “false negative”). If the decision maker finds negligence, her *expected* payoff is the probability-weighted value of her “true positive” and “false positive” contingent payoffs:

$$\Pr[v < v^*|x] \cdot 0 + \Pr[v \geq v^*|x] \cdot (-e_1) = -(1 - F(v^*|x)) \cdot e_1 \quad (1)$$

where $F(v^*|z)$ denotes the cumulative distribution¹⁸ associated with density function $f(v|z)$ evaluated at $v=v^*$. If the decision maker finds no negligence, in contrast, her expected payoff is the probability weighted value of her “true negative” and “false negative” payoffs:

$$\Pr[v \geq v^*|x] \times 0 + \Pr[v < v^*|x] \times (-e_2) = -F(v^*|x) \cdot e_2 \quad (2)$$

Combining expressions (1) and (2), it follows that a loss-minimizing decision maker will find liability if and only if the odds ratio of negligence to non-negligence is greater than the ratio of losses associated with Type 1 and Type 2 errors¹⁹:

$$\frac{\Pr\{D \text{ was Negligent}|x\}}{\Pr\{D \text{ was Not Negligent}|x\}} \equiv \frac{F(v^*|x)}{1 - F(v^*|x)} > \frac{e_1}{e_2}, \quad (3)$$

In the special case where the decision-maker places equal weight on Type 1 and Type 2 errors ($e_1 = e_2$), this condition simplifies down into requiring that the evidence be strong enough to satisfy a more-likely-than-not condition:

$$\Pr\{D \text{ was Negligent}|x\} > \Pr\{D \text{ was Not Negligent}|x\} \quad (4)$$

Significantly, this condition corresponds to the description that commentators frequently identify with the *preponderance of the evidence* standard. In contrast, when Type 1 errors are more costly than Type 2 errors ($e_1 > e_2$), the evidentiary standard would weigh evidence differently too. When false positives are three times as costly as false negatives, for example, this critical

¹⁸ See note 17, *supra*.

¹⁹ Recall the convention from probability theory that “Type 1” errors are false positives, and “Type 2” errors are false negatives. (A “Type 3” error, in contrast, is the failure to keep straight which is which.)

cutoff increases from 50% to 75% (something akin to a clear and convincing standard). When false positives are nineteen times more costly, the critical cutoff increases further to 95% (something akin to beyond reasonable doubt).

The above framework, moreover, is general enough to admit numerous generalizations. For example, suppose the observed realizations of evidence x were systematically related to underlying precaution levels v in an ordered fashion – that is, higher realizations of x predicted higher values of v .²⁰ In such settings, the condition in (3) would imply the existence of a “cutoff” threshold for evidence -- say x^* -- such that an evidentiary signal of x^* or less would imply that the burden of persuasion has been met, and the court should find negligence. Alternatively, one might imagine expanding the dimensionality of the framework, allowing for one dimension to capture the defendant’s precautions (and evidence about it), and another to capture the plaintiff’s precautions (and evidence about it). Such generalizations would also allow the decision-theoretic framework to capture both causes of action and affirmative defenses.

A decision-theoretic account also lends itself to analyzing (and in some ways complicating) the long-standing quandary about how to apply the BoP when the underlying cause of action has multiple conjunctive elements. In both civil and criminal law settings, the traditional (albeit sometimes controversial) requirement is that in order to prevail, the petitioning party must prove “every element” of a legal claim by the applicable evidentiary standard.²¹ A common impression is that the every-element requirement favors the defendant, raising the degree of difficulty for the petitioning party to procure a victory. While this impression possibly has merit, a decision-theoretic account can also complicate this understanding. Consider, for example, a hypothetical cause of action that requires proof of two conjunctive elements, A and

²⁰ Such might be the case, for example, if x constituted a “noisy” but unbiased observation of v , so that $x = v + \varepsilon$, where ε denotes an error term with zero and a strictly positive variance (representing the observational noise).

²¹ See, e.g., *In re Winship* 397 U.S. 358 (1970).

B, and an applicable standard of proof consisting of a preponderance rule (i.e., a judicial Bayesian posterior of strictly more than 50%).

		<u>Element A</u>	
		<u>Present</u>	<u>Absent</u>
<u>Element B</u>	<u>Present</u>	α	β
	<u>Absent</u>	γ	$1-\alpha-\beta-\gamma$

Table 1: *Hypothetical Joint Posterior Probabilities of Presence / Absence of Required Elements*

Suppose that the available evidence were such that the fact-finder’s posterior assessment of the presence/absence of Elements A and B were as shown in Table 1 above. Thus, the evidence produced at trial suggests that the contingency where Element A is present but Element B is absent (i.e., the southwest cell) occurs with probability γ , and similarly for all other combinations, so that each entry in the table represents a value falling between zero and one, and the cells collectively sum to 1.0 (or 100%).

One plausible interpretation of an every-element requirement is that the petitioning party must prove each of (1) the presence of Element A in isolation by a preponderance (the left column of the table), and also prove (2) the presence of Element B in isolation by a preponderance (the top row). Under this interpretation, from Table 1 above the plaintiff will succeed if both $\Pr[A] = \alpha + \gamma > 0.5 = 50\%$ and $\Pr[B] = \alpha + \beta > 0.5 = 50\%$, a burden that is equivalent to the condition that $\alpha + \text{Min}\{\beta, \gamma\} > 0.5$.

Now consider the plaintiff’s task were she subject *not* to an element-by-element requirement, but rather something else. One plausible “something else” would be that the plaintiff must demonstrate the *joint* occurrence of Element A and Element B (or formally, $A \cap B$) by a preponderance of the evidence. Here, Table 1 suggests that this would be equivalent to

showing $\Pr[A \cap B] = \alpha > 0.5$, a condition that is *more* (not less) restrictive than the element-by-element test. In other words, relative to this alternative, an every-element test would make for a smoother (not bumpier) road to a plaintiff's victory (and a defendant's loss).²²

As the discussion above indicates, the decision-theoretic framework has a number of ostensibly appealing characteristics. It is intuitive. It uses foundational concepts from probability theory. It yields interpretations that tend to “fit” with at least some existing practices, and particularly the Burden of Persuasion, Assumptions, and Standards of Proof. It can be used to analyze and frame various types of positive and normative analysis about the BoP.²³ And it is capable of engaging with interesting current debates about evidentiary burdens. For these reasons, perhaps, the decision-theoretic model has remained a popular modeling choice for evidentiary burdens among lay people and practitioners, garnering numerous adherents among law and economics scholars in the process.

And yet, for all its intuitive appeal, the decision-theoretic framework also suffers from a number of potentially worrisome shortcomings that ultimately limit its utility as a tool for economic analysis. First, the decision theoretic approach gives no guidance to the fact-finder about the appropriate “prior” beliefs to adopt at the onset of litigation, and how to update such prior beliefs with evidence. It is well known in probability theory that “improper” prior beliefs

²² Much in this discussion turns on the “compared to what” question—specifically, how one portrays the alternative to an every-element test. Another plausible candidate might be a “sliding scale” standard, where the plaintiff can prevail if she proves only a subset of the substantive elements according to the standard of proof (or is allowed to prove some subset according to a discounted standard). It is readily shown that the sliding-scale test as described here is indeed more plaintiff-friendly than the every-element test. Nevertheless, as this discussion suggests, the decision-theoretic frame helps to isolate important (and possibly confounding) intuitions in the debate.

²³ For example, while Kaplan (1968) and others use a decision-theoretic framework as a means of “gaining insights into our factfinding processes,” (*id.* at 1091), both Polinsky & Shavell (2000) and P’ng (1986) employ a decision-theoretic framework as a backdrop to analyzing the optimal design of substantive legal rules in the presence of jurisprudential errors. Both Polinsky & Shavell (2000) and P’ng (1986) provide a normative argument for equal weighting of false negatives and false positives, in the case of zero litigation costs and risk neutral parties. The game-theoretic literature discussed below retrains focus on how evidentiary standards affect litigation costs.

can be unreliable guideposts for statistical inference and Bayesian learning.²⁴ Leaving this as an “open parameter” to be filled in by the fact finder may be a recipe for practical disaster. Many who apply a decision theoretic frame tend to neglect this issue even as they implicitly make assumptions about its size.²⁵

Second, even if a sophisticated decision maker had access to informed priors on available evidence and justiciable conduct, the decision-theoretic model provides an unsatisfactory account of how the evidence (x) makes its way to the fact-finder. To the contrary, evidence falls from the metaphorical sky, like manna (or at least manila file folders) from heaven. The process by which each side uses, produces, authenticates, cross-examines, or rebuts evidence is substantially black-boxed within this approach (or at least kept offstage). Such a modeling choice has obvious problems for those interested in positive theory, since the framework delivers few predictions about how litigants are likely to behave (and at what economic cost), or how that behavior might change if one (say) inverted the underlying burdens or changed requisite standard. Moreover, the strategic means by which parties search for and present evidence in practice almost certainly implies that the fact-finder will observe a biased draw of evidence from the proverbial Bayesian evidence urn – a factor generally suppressed for in the decision-theoretic frame (cf. Daughety & Reinganum (2000)). As a result, despite its intuitive allure, the decision-

²⁴ A classic puzzle/paradox in Bayesian reasoning, originally attributed to Kraitchik (1953), provides an entertaining example of this problem. You and an “opponent” are each allocated (at random) one of two sealed envelopes that contain monetary prizes. You are both told nothing about the magnitudes of the prizes in the two envelopes, other than the fact that one prize is twice as large as the other. You open your envelope, revealing M dollars. You reason that your counterpart’s envelope is equally likely to have either $2M$ or $\frac{1}{2}M$ dollars in it, and accordingly you compute the expected value of her envelope to be $1.25M = \frac{1}{2}(2M + \frac{1}{2}M)$. Sensing a profit opportunity, you coyly ask your opponent if she wishes to trade envelopes. To your surprise, she has been doing exactly the same calculation, and she eagerly agrees. The reasoning described above presents a paradox, because neither contestant has any information about whether her envelope contains the most, and they both know that exchanging envelopes is zero sum game; yet each of them appears to have a strict preference for switching. The paradox occurs because the reasoning of both players (described above) depends erroneously on the Bayesian updating of an “improper” prior distribution on $(0, \infty)$. For a longer explanation of the paradox and its resolution, see Christensen & Utts (1992).

²⁵ For example, Kaplan’s (1968) seminal paper utilizing the decision theoretic approach recognizes the criticality of specifying prior beliefs, but offers little in the way providing guidance other than to suggest that they should be somewhere between $\frac{1}{2}$ and 1 in 200 million.

theoretic framework is poorly equipped to analyze how evidentiary burdens operate through *litigation*, since the underlying approach is simply not a theory that has much to do with litigation in the first instance.²⁶

Finally, even putting the foregoing objections aside, the decision-theoretic framework provides only limited traction about how one should weigh (or even think about) Type 1 and Type 2 errors in a broader sense. Although there are some notable exceptions,²⁷ judicial errors frequently appear in decision-theoretic frameworks as unadorned costs whose value is determined outside the model, with little context or evaluation behind them from an underlying welfare economics theory. In fact, there is some reason to believe that the decision-theoretic framework described above is a strange bedfellow with robust economic theories of welfare, particularly the frequently-made assumption that all Type 1 errors (or Type 2 errors) have equal intensities within their own categorization ranks. For instance, there is little reason to believe that the nature and consequences of a false positive when the defendant had “barely” complied with the negligence standard have the same welfare implications as in situations where the defendant’s precautions were highly supererogatory. Consequently, the decision theoretic approach provides a helpful—but somewhat incomplete—vehicle for making normative assessments of the burden of proof within law and economics.

²⁶ As an anonymous referee correctly points out, the critique that the decision theoretic model is incomplete may also be applicable to the other modeling tools below. Indeed, as I discuss in the concluding section *infra*, the project of developing a unified, all-inclusive, and tractable model of evidence production and primary behavior is still a significant challenge to this literature. *Accord* Kaplow (2012).

²⁷ The most notable of these are Polinsky & Shavell (2000) and P’ng (1986), who embed a decision theoretic account of litigation in an ex ante model of deterrence. I discuss these papers at greater length in Section 3C, *infra*.

2. Game Theoretic Frameworks

More recent contributions have focused on developing alternative frameworks that are better able to capture the incentives and behavior of litigants (focusing less on the decision theoretic issues for the fact-finder). One particularly influential framework in the literature conceives of evidentiary burdens as setting the rules for a non-cooperative litigation game of “evidence production,” where the litigants, acting simultaneously or sequentially, adduce (potentially conflicting) evidence with the intent to influence the ultimate judgment of the court. This literature draws considerable influence from other branches of economics that study redistributive “contests,” including the literatures on political rent-seeking, lobbying, and international conflict (e.g., Tullock 1980; Hirschleifer 1991; Skaperdas 1996).

The intuitive idea behind such litigation games is relatively straightforward. Litigation, under this view, serves as a performative venue where litigants engage in what amounts to a costly arms race to send signals about the strength of their case to the fact-finder. The more evidence a litigant can muster, the framework posits, the better his chances of prevailing. While it may seem wasteful on first blush to design a system that encourages this type of costly litigiousness, it turns out that such mechanisms can be welfare enhancing *so long as* the production of evidence serves as a valuable signal about the merits of the party’s position that are not already known to the fact-finder. Such a signaling role is possible, in turn, if parties with the most righteous cases also find it relatively cheaper (on the margin) to add evidence into the record.

In such situations, a judicial rule that invites costly evidence production as a strategic maneuver may serve as an effective verification device to separate deserving from non-deserving parties. Moreover, knowing that such a costly verification device exists – and knowing that their

costs of participating in it will be lower if they behave lawfully now, the prospective parties may also be deterred from socially wasteful activity *ex ante*.

To understand how litigation games play out more concretely, suppose a plaintiff (π) and defendant (Δ) were contemplating litigating an issue (which could pertain to either civil or criminal liability). Should litigation occur, the outcome will be a function of underlying procedural rules, the evidence produced by each side, and possibly a set of external factors. Accordingly, let the ultimate judgment (or “judgment function”) be denoted as $J(x,y,z)$ where the arguments $x \in X$ and $y \in Y$ denote measures of evidence presented by the plaintiff and the defendant (respectively). The argument $z \in Z$ denotes a vector of other extraneous state-contingent considerations that are readily observable to the court, and that may (depending on policy and design decisions) be relevant to the judgment. (These arguments may conceivably be vector values, as may the judgment function itself²⁸).

Embedded within the judgment function is a host of potential practical considerations in litigation, including substantive legal rules, policy considerations, and (significantly) BoP protocols for weighing the evidentiary productions of the parties. Consequently, contributions within this literature diverge considerably with respect to the key ingredients of the judgment function. However, virtually all of these approaches commonly posit (with little loss of generality) that $J_x > 0$ and $J_y < 0$, so that larger amounts of evidence produced by the plaintiff generally enhance the judgment, while larger amounts of evidence produced by the defendant

²⁸ The judgment function may be vector valued for a number of reasons. For example, in many situations, the judgment may involve more than a mere transfer from the defendant to the plaintiff. For example, a criminal prosecution might result in both a monetary sanction and a non-monetary punishment. Another example is a civil case in which damages paid by the defendant are “decoupled” from those received by the plaintiff, with part of the damages going to the state. See Che & Polinsky (1991); Choi & Sanchirico (2004). Nevertheless, to keep things simple, I will constrain the discussion to scalar-valued terms.

dampen it. (See, e.g., Daughety & Reinganum (2000) Bernardo et al. (2000), and Sanchirico (2008) for general discussions).

Formal Analysis

Consider, for example, a slight generalization of the litigation contest studied in Bernardo et al. (2000), who identify J with a function that assigns a liability probability to a case with known damages D , so that expected liability is given by $J(x, y, z) \cdot D$ where:

$$J(x, y, z) = \frac{x}{x + b(z) \cdot y} \tag{5}$$

The term $b(z) > 0$ represents the (possibly state-contingent) “weight” that the defendant’s evidentiary showing is accorded relative to the plaintiff’s – what Bernardo et al. (2000) analogize to the strength of an initial assumption / presumption.

Note that this framework presumes that any state-contingent considerations reflected in z must be fully observable by the court, reflecting any number of general policy commitments. There may also be a host of other considerations that are relevant, yet not directly observable to the court. Most saliently, judicial actors are likely to be unable to observe the level of care undertaken by the defendant. To reflect the distinction between easily observable facts and opaque ones, suppose the “complete” state of the world is denoted by Ω which can be partitioned into $\Omega \equiv \{Z, W\}$, corresponding to components that are observable and non-observable to the court, respectively. A key feature in evidence production games (not present in the decision-theoretic framework) is that the parties’ strategic choices about how much evidence to produce is that such evidence can provide signals about the unobservable aspects of Ω (as discussed in greater detail below).²⁹

²⁹ See Sanchirico (2000 & 2001) for general discussions.

Evidence production models typically assume that evidence production is costly to both parties. Denote the cost to the plaintiff and defendant, respectively, as $c_\pi(x, \omega_\pi)$ and $c_A(y, \omega_A)$, where $(\omega_\pi, \omega_A) \in \Omega^2$. Significantly, note that the cost of providing evidence for each party is allowed to turn both on publicly observable information as well as private information of the parties. While this assumption introduces some complexity, it is often a critically important ingredient of the evidence production contest approach. In particular, if parties with different information (say, about their prior negligence) face differential costs of providing evidence, their performative behavior in court may signal credible evidence about that information—and in so doing allow the judgment function to provide a mechanism to screen (at least partially) for liability.

To see more concretely how this screening mechanism might work, suppose that the parties are risk neutral, that the judgment function is given as in (5) above, fixing $b(z) = 2$, and that damages (conditional on liability) are fixed by substantive law at \$100. Suppose further that the plaintiff and the defendant submit their evidence to the court simultaneously, and that the plaintiff faces a linear cost of evidence production given by $c_\pi(x, \omega_\pi) = x$. The defendant's cost of evidence, in contrast, turns on whether she had previously acted negligently, and in particular assume that $c_A(y, \text{Negligent}) = 2y$ and $c_A(y, \text{Not Negligent}) = y$. Finally, suppose that only the defendant knows with certainty whether she has previously behaved negligently, and that is common knowledge that negligent and non-negligent defendants are equally prevalent.

It is easily verified in this setting that the equilibrium levels of evidence production consists of $x = 23.789$ by the plaintiff, $y_N = 12.492$ by the negligent defendant, and $y_{NN} = 22.594$ by the non-negligent defendant. Note that in equilibrium, the non-negligent defendant brings substantially more evidence into court than her negligent counterpart, since it is relatively cheap

for her to adduce evidence in her favor. Consequently, the equilibrium probability of liability that ensues is approximately 49% for the negligent defendant and 22% for the non-negligent defendant. That is, the court delivers a harsher expected penalty to the negligent defendant than it does to the non-negligent one. This is significant for anyone interested in linking the role of evidence production to ex ante incentives for primary behavior (taken up in Section 4 below).

Discussion

An obvious advantage of the game theoretic approach is that the former explicitly provides a representation of the process by which litigants bring evidence into the courtroom. Consequently, it facilitates a number of insights into how their efforts and ultimate outcomes are affected by altering legal processes (through distortions to the judgment function). Moreover, the game-theoretic account allows one to analyze more completely the relationship between post-injury litigation activity and pre-injury deterrence. This richer account of litigation has some clear advantages over the decision-theoretic framework, where such litigation/deterrence tradeoffs are largely (though not completely) suppressed. For example, as Bernardo et al (2000) and Sanchirico (2000; 2001) both demonstrate, the signaling/screening advantages of a litigation contest come at the cost of inducing significant evidence production costs by the parties. In some circumstances, these litigation inefficiencies can grow so large as to swamp the beneficial signaling and deterrence benefits (such as when the difference in evidence production costs of negligent and non-negligent parties is small). In such situations, it may be optimal to push the applicable presumption / assumption in the direction of an irrebutability.

At the same time, the evidence production game approach to evidentiary burdens some distinct drawbacks too. First, the most tractable functional forms for $J(.)$ do not have obvious correspondences to the various institutional features of the BoP analyzed in Section 2. For

example, in the judgment function explored above, the term $b(z)$ is effectively a proxy for burdens, presumptions/assumptions, and standards of proof all rolled into one. While it is possible to interpret an increased burden, a stronger presumption/assumption, or an enhanced standard through a change in the $b(z)$ function, they tend not to have effects that are distinguishable from one another.

Second, while the evidence production approach generally does a much better job than the decision-theoretic model at capturing the litigants' incentives and behavior, it does a *worse* job at capturing the behavior of the fact-finder. Indeed, in this class of models, the fact-finder is identified largely as an automaton, technocratically applying a judgment function. This representation can present some obvious difficulties, particularly if one concedes that the fact-finder may be a strategic participant in the game. For example, in the evidence production game analyzed above, evidence production decisions generated a "separating" equilibrium, where negligent defendants presented a relatively "low" amount of evidence when compared to non-negligent defendants. As a result, the posited judgment function assigned a greater equilibrium likelihood of liability to negligent defendants than to non-negligent ones. But it still committed errors—negligent defendants were exonerated 51% of the time, and non-negligent defendants were erroneously held liable 22% of the time. Were this judgment function applied by a sophisticated, Bayesian judge/jury, it would almost certainly understand the separating nature of the equilibrium. In turn, the fact-finder would plausibly be tempted to abandon the posited judgment function, and instead assign liability if and only if the defendant had produced a "low" amount of information according the separating equilibrium. Anticipating that the fact-finder might change the rules, of course, the litigants themselves might change their own behavior, potentially upsetting the separating nature of the equilibrium. The basic setup in the evidence

production framework (as presented above) does not easily allow for such strategic interplay between the litigants and a sophisticated fact-finder.

Finally, as with the decision-theoretic approach, the continued necessity of committing to an express functional form for the judgment function may prove to be unduly restrictive for those interested in purely normative contributions, such as the optimal design of a system of procedure from the ground up. To be sure, one could still use this approach to formulate modest normative prescriptions within (say) a particular family of forms. While such an approach may work well for pragmatically grounded normative inquiries, the limited domain of the resulting prescriptions would not be particularly attractive in a more purely normative analysis.

3. Hybrids: Game Theoretic Frameworks with a Bayesian Fact-Finder

One potential approach for dealing with at least some of the objections noted above is to re-introduce the fact-finder as a strategic player who behaves in a sequentially rational fashion. Such a setup tends to resemble the general setup from the last subsection, but it imposes an additional “commitment” constraint on the fact finder, ruling out at least certain judgment functions that would require her to make sub-optimal or inaccurate decisions when called upon to decide the outcome of the case.

For example, consider a similar framework as above, but suppose the fact-finder were a Bayesian motivated by judicial accuracy, and was sufficiently sophisticated to “decode” equilibrium play among the parties. Such a fact finder would by definition respond to a completely separating equilibrium by assigning liability to the defendant if and only if the equilibrium reveals her to be negligent. Consider how such a rule would play out in the framework developed above, where damages are fixed at \$100, plaintiffs and non-negligent

defendants have marginal evidence production costs of \$1, and negligent defendants face marginal evidence production costs of \$2. As demonstrated above, the judgment function posited in equation (5) would no longer be viable if the fact-finder always inverted the equilibrium. However, other forms of judgment function may be consistent with the fact-finder's hypothesized commitment to accurate adjudication. Consider, for example, the following judgment function:

$$J(x, y, z) = \begin{cases} 1 & \text{if } x \geq 20 \text{ and } y < 85 \\ 0 & \text{otherwise} \end{cases}. \quad (6)$$

It can be easily confirmed that this judgment function supports a fully separating Bayesian perfect equilibrium in which plaintiffs produce evidence of $x=20$, non-negligent defendants produce evidence of $y_{NN}=85$, and negligent defendants produce no evidence at all. Moreover, such a framework gives rise to liability if, and only if, defendants are negligent, thereby ensuring that the fact-finder can commit to the decision rule ex post.³⁰

The addition of a commitment constraint for the fact-finder provides an intuitive way to combine the decision-theoretic and evidence production models. In fact, the example above even lends itself to an interpretation of the burdens of persuasion and presumption. (That is, the plaintiff's initial burden of production is satisfied by producing at least 20 units of evidence, at which point the burden of production switches to the defendant, who can rebut it by producing 85 units of evidence).

Nevertheless, it is important to note the limitations of this approach. First, by definition, the imposition of a commitment condition narrows the range of judgment functions that are available to implement. This can have real welfare costs: As is well known in the principal/agent

³⁰ This does not imply, of course that the judgment function posited in the text is optimal, and there exist others that are less costly to implement but equally informative. (Indeed, the function is not even particularly robust to various equilibrium refinements common to the signaling literature). That said, the example in the previous section was not assumed to be optimal either.

optimal contracting literature, good incentive devices tend to reveal information that the principal would ideally wish later to exploit, but her obvious temptation to do so necessarily would distort the very behavior that revealed that information in the first instance. Here too, non-commitment by the fact-finder necessarily limits the types of institutional procedures that are possible, and as a result the set of achievable outcomes is reduced (possibly in a welfare reducing way). One interesting and provocative implication of this observation is that a “good” system of evidence from a welfare perspective may be one that utilizes relatively “dumb” (or at least non-reflective) fact-finding processes—which are neither designed nor function to make nuanced assessments of truth, but are instead gauged to incentivize behavior, distribute risk, and affect other allocative consequences between defendants, plaintiffs, and broader society.

To the extent this last observation is valid, it may raise difficult questions about how (and even why) the commitment issue comes about in the first place. If the fact-finder were truly motivated by welfare concerns, then the fact that she may wish to commit to making classification errors in the process should not itself be objectionable, so long as the system she is applying induces optimal behavior. Hence, perhaps the best motivation for this approach is to assume that the fact-finder may also be an agent, and may therefore be motivated by factors that are either extraneous to social welfare or enter differently in social welfare calculus (such as accuracy per se, minimizing workload, and so forth). This possibility is taken up below.

Finally, just as with the evidence production approaches with commitment, the simple addition of a commitment constraint does not ensure that the class of judgment function being studied is sufficiently broad to satisfy purely normative scholars interested in “optimal” process design. To engage in this enterprise requires a more general approach, less wedded to a particular class of processes or functional forms. It is to this last literature I now turn.

4. Mechanism Design Approaches

A final significant strand in the BoP literature endeavors to engage more directly in broad institutional design. That is, rather than focusing on a particular set of functional forms for designing legal procedure, the mechanism design (MD) approach endeavors to derive the characteristics of an “optimal” set of fact-finding processes BoP from a set of first principles about social objectives, individual incentives, and information constraints. This approach can yield powerful insights useful to those interested in normative analysis and is increasingly popular. (See, Sobel 1985; Shin 1998; Sanchirico 1995, 2000, 2001; Bull & Watson 2004; Demougin & Fluet. 2006. Cf. Cooter & Emons 2004).

The kernel of the MD approach focuses on structuring legal process from the “ground up”, but with consistent attention to how design choices alter and affect the incentives of the parties both (a) to participate in the process; and (b) to “play by the rules” – i.e., reveal truthful information to the fact-finder – when litigation occurs. These two key design constraints are sometimes referred to (respectively) as the parties’ “Individual Rationality” and “Incentive Compatibility” conditions. Unlike the game-theoretic approach described above, the MD approach does not commit to a particular judgment function. Rather, it endeavors to derive an “optimal” one in the sense of maximizing economic welfare according to *any* possible set of institutional rules, but without violating to the two aforementioned types of constraints.

A handy tool that many MD approaches employ is to concentrate on (so-called) “direct revelation mechanisms,” which are processes where the litigants are simply asked to report to the fact-finder all of relevant information they privately know about the case (e.g., their extent of precautions, level of risk preferences, extent of damage, etc.), and the fact-finder delivers a binding judgment and instructed behavior based on those reports. While there is nothing (in

principle) that would prevent the litigants from attempting to deceive the fact-finder when making their reports to the fact-finder, the “trick” of the MD approach is to design a set of decision functions that anticipate these incentives, account for them, and alter payoffs to make truth-telling optimal for all parties.³¹

More formally, and utilizing the notation from the previous sections, under the MD approach the parties are asked to report their own private information $(\omega_\pi, \omega_\Delta)$ to the fact finder. Recall that this information is not directly verifiable to the mechanism designer, and thus nothing necessarily precludes misrepresentation by the parties. Thus, let (μ_π, μ_Δ) denote the actual reports (truthful or not) submitted by the parties to the court. Once the parties have submitted (μ_π, μ_Δ) , the mechanism announces (i) quantities of evidence that the parties are required to produce based on their reports $(x(\mu_\pi, \mu_\Delta; z), \text{ and } y(\mu_\pi, \mu_\Delta; z), \text{ respectively}),$ along with (ii) an expected final judgment $j(\mu_\pi, \mu_\Delta; z)$, which the defendant damages to the plaintiff.³² Together, these three decisions constitute the decision mechanism of the fact finder, collectively denoted as $\Gamma(\omega_\pi, \omega_\Delta | z)$. As before, assume that the cost to the plaintiff and defendant of producing these levels of evidence are $c_\pi(x, \omega_\pi)$ and $c_\Delta(y, \omega_\Delta)$, respectively.

Using this notation, any adjudication mechanism that induces both truth-telling and participation by both the plaintiff and defendant must satisfy the following four conditions:

$$\begin{aligned}
 (\text{IC}_\pi): \quad & \omega_\pi \in \operatorname{argmax}_{\mu_\pi} E_{\omega_\Delta} \{J(\mu_\pi, \omega_\Delta; z) - c_\pi(x(\mu_\pi, \omega_\Delta; z), \omega_\pi)\} & (7) \\
 (\text{IC}_\Delta): \quad & \omega_\Delta \in \operatorname{argmax}_{\mu_\Delta} E_{\omega_\pi} \{-J(\omega_\pi, \mu_\Delta; z) - c_\Delta(y(\omega_\pi, \mu_\Delta; z), \omega_\Delta)\} \\
 (\text{IR}_\pi): \quad & E_{\omega_\Delta} \{J(\omega_\pi, \omega_\Delta; z) - c_\pi(x(\omega_\pi, \omega_\Delta; z), \omega_\pi)\} \geq 0
 \end{aligned}$$

³¹ The ability to limit one’s attention solely to mechanisms that induce truth-telling in equilibrium is sometimes referred to as the “revelation principle.” It is generally applicable in most litigation settings, but perhaps not universally so.

³² .The actual judgment may be stochastic, and can reflect both liability and damages assessments.

$$(IR_{\Delta}): \quad E_{\omega_{\Delta}}\{-J(\omega_{\pi}, \omega_{\Delta}; z) - c_{\Delta}(y(\omega_{\pi}, \omega_{\Delta}; z), \omega_{\pi})\} \geq -K$$

Conditions (IC_{π}) and (IC_{Δ}) , the “incentive compatibility” constraints, essentially require that truthful disclosure is a Nash equilibrium under the posited mechanism (assuming that truthful disclosure is a normative goal). Conditions (IR_{π}) and (IR_{Δ}) the individual rationality constraints, requires that it be rational for the plaintiff and defendant to participate in the game. Thus, Condition (IR_{π}) requires that the plaintiff’s expected payoff in equilibrium be no lower than simply choosing not to litigate (and received payoff of zero). A similar interpretation applies to (IR_{Δ}) , which requires that the defendant’s expected payoff be no lower than what he would expect to pay in a default judgment.³³

Finally, and common with the analysis of the previous section, a direct revelation mechanism may permit the fact-finder to learn private information from the litigants. If the fact-finder herself is unable to commit to any type of judgment mechanism, then there may be another constraint on the set of admissible mechanisms in light of what the fact-finder is able to commit to given her own preferences (e.g., truth-seeking, effort minimizing, etc.). Denoting the set of mechanisms to which a fully-informed fact-finder could commit as $\Gamma^*(\omega_{\pi}, \omega_{\Delta})$, this would suggest a fifth constraint on the design problem:

$$(C_{FF}): \quad \Gamma(\omega_{\pi}, \omega_{\Delta}|z) \subseteq \Gamma^*(\omega_{\pi}, \omega_{\Delta}) \quad (8)$$

As should be clear by this point (perhaps painfully so), the MD approach is easily the most abstract and technically demanding framework discussed in this chapter. Moreover,

³³ Condition (IR_{Δ}) is not a mirror image of Conditions (IR_{π}) because the defendant does not voluntarily appear in court. Instead, (IR_{Δ}) simply places an upper limit (K) on the damages the defendant can lose in a default judgment, which in civil cases is plausibly the defendant’s bankruptcy threshold. When the defendant is well capitalized, (IR_{Δ}) need not constitute a binding constraint. In criminal cases, in contrast, the situation is frequently a bit more complicated, since the judgment might accord non-zero-sum payoffs to the prosecutor and defendant.

characterizing the solution to an MD problem is also a technical and intricate enterprise, and therefore it is beyond the scope of the current inquiry.

Nevertheless, the mechanism design approach can be a powerful tool, particularly for normative analysis of the BoP. An “optimal” mechanism generally represents the best possible outcome that one could hope to achieve in a non-cooperative setting. Consequently, its characteristics represent an enormously useful benchmark for comparing, assessing, and rationalizing current institutional practices. Moreover, the MD approach can sometimes deliver helpful suggestions about what should or should not matter in designing a “better” evidentiary system.

Yet the strengths of the MD framework are also its limitations. Since it is focused on deriving optimal rules, it is not particularly effective at analyzing the economic effects of existing BoP institutions. More to the point, mechanism design approaches generally do not produce a complete description of processes by which their outcomes can be achieved. They merely characterize what sorts of outcomes are possible under ideal circumstances. Consequently, even if one were confident about the desirability of an optimal mechanism, it might be very difficult in practice to build a set of institutions capable of generating those outcomes.

Second, even the normative application of the mechanism design approach requires one to articulate exactly what the appropriate welfare function *is*. There are a number of factors that may go into that decision, and ultimately this is a dimension on which many existing contributions to the literature (both within the mechanism design tradition and outside it) tend to diverge. It is to this last substantive distinction I now turn.

C. Scope of Inquiry: Ex Ante, Litigation, or Both?

Perhaps one of the most substantial (and in some ways vexing) challenges to normative applications of law and economics to the BoP is that of articulating what, precisely, the appropriate social welfare goal (or goals) should be. While most economists tend to embrace welfare measures that correspond (roughly) with various measures of efficiency, the appropriate domain of an efficiency measure is itself somewhat elusive within this literature.

Part of the reason for this elusiveness stems from the traditional distinction that the law (and many legal scholars) have placed on procedural versus substantive rules. One way to appreciate this distinct (at least heuristically) appears in Figure 1:




Figure 1: Elements of Welfare and Substance vs. Process

The figure represents five relevant economic considerations for the function of law: (A) Transaction/Bargaining costs; (B) Investment, Reliance and/or Precaution Costs, (C) Incidence / Likelihood of Harm; (D) Adjudication / Litigation costs; and (E) Sanctions and Enforcement Activity. Although these categories are almost certainly not exclusive, they are all clearly

relevant for economic analysis, and are intuitive components of virtually any economic welfare measure. (For example, a standard Calabresian efficiency analysis might ask what sets of substantive and procedural rules would tend to minimize the total expected costs from each of the five categorical dimensions illustrated in the figure).

At the risk of simplification, it is largely accurate to state that economic analysis of most *substantive* legal rules—such as liability or negligence standards, remedies, affirmative defenses, and the like have typically focused on how such rules affect the activities in categories (A), (B), and (C). And, at the same time, most economic analysis of the legal process—such as pleading standards, filing fees, and, notably, burdens of proof—tends to fall most naturally into steps (D) and (E). In other words, at least initial scholarship on burdens of proof largely concentrated on how they affected the incidence and distribution of costs in litigation, as well as the cost and accuracy of trial outcomes.

This characterization, of course, is also somewhat of a gross over-simplification. Even moderately well-read legal scholars will recognize immediately that most claims about the distinction between form and process have serious conceptual problems. Procedural rules, they will assert, can profoundly (even if indirectly) affect deterrence, proof, and available remedies, and vice versa. Indeed, plausibly since the very introduction of the Substance/Process distinction (largely thought to have come from Blackstone), legal commentators have expressed frustration at it, describing it in turns as “vague,” “imprecise,” “chameleon-like,” and “superbly fuzzy” (Main 2010).

Nevertheless, the distinction is one that to this day animates legal practice, and the trajectory of economic analysis of legal burdens has largely followed a similar course. Early contributions in this literature tended to focus predominantly on welfare measures that pertained

solely to litigation and beyond. Within this domain, a number of considerations are at the forefront, such as how the BoP interacts with expended litigation costs, delay costs, judicial accuracy, and the likelihood / terms of settlement. (Rubinfeld & Sappington 1987; Easterbrook 1984).

More recently, however, law 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 that is most closely associated with substantive law (rather than procedural rules). One set of important contributions to the literature (most notably, Polinsky & Shavell (2000), and P'ng (1986)) study a decision-theoretic model of adjudication in a standard deterrence framework to analyze the effects of judicial error. Within this framework, judicial errors tend to dilute deterrence incentives, reducing prospective liability risk from bad behavior and increasing risk from good behavior. Consequently, it is generally welfare improving to attempt to reduce the incidence of error. Perhaps more interesting is the observation that under a number of relatively general conditions, the efficiency costs of false positives and false negatives are symmetric, giving some justification to a preponderance standard (which equally weighs Type 1 and Type 2 errors).

Another cluster of papers in law and economics beginning soon began to push that analytic step a notch further, explicitly studying how strategic behavior during litigation interacts with primary activities. (See, e.g., Sanchirico 1995; 2008; Bernardo et al. 2000; Demougian & Fluet 2006). These papers tend to tie game theoretic or mechanism design frameworks to models of deterrence. This is an area that has been particularly fertile during the last decade or so. By way of example, Bernardo et al (2000) develop a game-theoretic model of evidence production that explicitly links *ex ante* deterrence and the standard of proof at trial. Using this model, they

illustrate how pro-defendant reforms can have two countervailing effects. First, the hypothesized reform makes it less likely (all else constant) for plaintiffs to prevail in trial and thus less attractive for them to litigate; but second, defendants have an incentive to engage in fewer precautions *ex ante*, because of the protection the procedural reform affords them. In fact, Bernardo et al (2000) demonstrate that this second effect can easily be the dominant one – so that following a defendant-friendly procedural reform, one would predict an *increase* in the number of lawsuits filed as well as an *increase* in plaintiffs’ overall win rate.

Sanchirico (2008) provides a cogent analysis of burdens of persuasion and production in a plausible game-theoretic framework linking *ex ante* litigation to strategic evidence production in trial. He offers an insightful justification for the possibly curious practice of usually placing the burden of persuasion (and initial production) on the party who is *not* the target of the law’s deterrence efforts (e.g., the injured plaintiff and not the defendant). Specifically, his model posits that the exercise of due care *ex ante* by a prospective defendant is likely to have two effects on later evidence production costs: (a) It will reduce the prospective defendant’s costs of producing exculpatory evidence should injury occur; and (b) it will increase the plaintiff’s cost of producing inculpatory evidence. The combined effects of this shift enhance economic welfare by making it both more likely that an injured plaintiff will file suit only against negligent defendants, and more likely that non-negligent defendants can prevail (cheaply) should the plaintiff does sue.³⁴

³⁴ Although Sanchirico (2008) does not analyze the issue squarely, his framework would also seem useful in understanding why in some situations tort law places the burden of production on the defendant, such as in application of the *res ipsa loquitur* doctrine. Such cases are plausibly typified by plaintiffs whose access to inculpatory evidence is both highly limited (e.g., they were not conscious) and likely unaffected by the defendant’s care level. In such situations, the sorting / screening benefits of burdening the plaintiff highlighted by Sanchirico (2008) tend to diminish substantially, if not completely.

Other interesting contributions outside the tort setting also link primary activity to evidence production. Lando (2006), for example, develops a model in which a large number of wrongful convictions might be efficient if there is high serial correlation across offences. Bernhardt and Nosal (2004) study the bankruptcy question of which distressed firms to liquidate or preserve. They demonstrate that even when a reliable technology exists for discerning “high quality” firms (who would efficiently be bailed out), it may make sense to eschew such technologies since the knowledge of a future bailout would induce high quality firms to take excessive risk ex ante. Demougin & Fluet (2006; 2008) work within a mechanism design framework, demonstrating within a special setting (where costs of producing evidence are of second order importance) that a preponderance standard can implement the allocation of an optimal mechanism.

4. Moving Forward

Although the theoretical literature on the BoP is now quite mature, it is in many ways still in a conceptual growth spurt. There are at least four challenges for this literature going forward. The first challenge is categorical. As noted above, most contributions within the field tend to fall into one of four modeling classifications: decision-theoretic; game-theoretic; hybrid; and mechanism design. Each of them generates particular and helpful insights about how to understand and design evidentiary procedures. However, until recently, there has been little work to synthesize them. (This chapter and contemporaneous and important work by Kaplow (2011; 2012) may go some distance to address this issue). It is likely that mechanism design frameworks have the most promise for making the literature more parsimonious, but (as noted above) this approach faces significant challenges to implementation. A significant amount of

theoretical work has yet to be done on deriving extensive form implementations of optimal mechanisms within familiar procedural environments. Moreover, although behavioral economics and prospect theory have penetrated this area of L&E just as elsewhere, there are likely many more opportunities for theoretically minded scholars to develop models of burdens incorporating prospect theory, anchoring effects, and bounded rationality.

The second issue is empirical. Particularly when one considers how the BoP interacts with primary activities and deterrence, the strategic dynamics quickly become complex, depending critically on deeply embedded modeling parameters. Consequently, it is difficult to know what to make of the growing theoretical literature in the absence of high quality experimental and empirical work to help calibrate the underlying theoretical models. That work is only just beginning, but is decidedly underway.³⁵

The final issue is pragmatic and political. Particularly for those interested in normative prescriptions that emanate from this literature, it is now relatively clear that within a number of realistic factual settings, an “optimal” burden of proof likely diverges (perhaps wildly) from the familiar legal procedures that have been part and parcel of legal practice for decades if not centuries. In such cases, it seems highly implausible that legal institutions are likely to embrace reform proposals with open arms. Perhaps a much better tack for those interested in normative prescriptions is to advocate the use of non-standard burdens in arbitration proceedings, where judicial processes are more of a matter of contract (Hadfield & Talley 2006). Not only would such applications constitute a ready market for new innovations, but they would also possibly act as important field experiments for them.

³⁵ See, e.g., Zamir & Ritov (2012) who present an intriguing set of experiments that suggest omission and status quo biases may alter representative jurors’ proof requirements in a preponderance case to be well over 50%.

References

- Allen, Ronald J. (1981). *Presumptions in Civil Actions Reconsidered*, Iowa L. Rev. 66: 843-67.
- Baum, Ido, Eberhard Feess & Ansgar Wohlschlegel (2009), *Reporter's Privilege and Incentives to Leak*, REVIEW OF LAW & ECONOMICS, 5, 701–15.
- Bernardo, Antonio, Eric Talley & Ivo Welch (2000), *A Theory of Legal Presumptions*, JOURNAL OF LAW, ECONOMICS & ORGANIZATION, 16, 1–49.
- Bernardo, Antonio & Eric L. Talley (1999), *A Note on Presumptions with Sequential Litigation* University of Southern California Law School, Olin Working Paper No. 99-9.
- Bull, Jesse & Joel Watson (2004), *Evidence Disclosure and Verifiability*, JOURNAL OF ECONOMIC THEORY, 118, 1–31.
- Burnhardt, Dan & Ed Nodal, (2004), *Near-Sighted Justice*, JOURNAL OF FINANCE, 6, 2655–84.
- Che, Yeon-Koo & A. Mitchell Polinsky (1991), *Decoupling Liability: Optimal Incentives for Care and Litigation*, Rand J. Econ. 22(4): 562-70.
- Choi, Albert & C. Sanchirico (2004), *Should Plaintiffs Win What Defendants Lose? Litigation Stakes, Litigation Effort, and the Benefits of Decoupling*, JOURNAL OF LEGAL STUDIES, 33, 323–54.
- Cooter, Robert. & W. Emons (2003), *Truth-Revealing Mechanisms for Courts*, JOURNAL OF INSTITUTIONAL AND THEORETICAL ECONOMICS, 159, 259–79.
- Christensen, Ronald & Jessica Utts (1992), *Bayesian Resolution of the 'Exchange Paradox'*, THE AMERICAN STATISTICIAN, 46, 274–76.
- Daughety, Andrew & Jennifer Reinganum (2000), *On the Economics of Trials: Adversarial Process, Evidence, and Equilibrium Bias*, JOURNAL OF LAW, ECONOMICS & ORGANIZATION, 16, 365–94.
- Demougin, Dominique & Claude Fluet (2006), *Preponderance of Evidence*, EUROPEAN ECONOMIC REVIEW, 50, 963–76.
- Demougin, Dominique & Claude Fluet (2008), *Rules of Proof, Courts, and Incentives*, RAND JOURNAL OF ECONOMICS, 39, 20–40.
- Dewatripont, Mathias & Jean Tirole (1999), *Advocates*, JOURNAL OF POLITICAL ECONOMY, 107, 1–39.
- Easterbrook, Frank H. (1984), *The Limits of Antitrust*, TEXAS LAW REVIEW, 63, 1–40.

- Fees, Eberhard, Gert Muehlheusser & Ansgar Wohlschlegel (2009), *Environmental Liability Under Uncertain Causation*, EUROPEAN JOURNAL OF LAW & ECONOMICS, 28, 133–48.
- Hadfield, Gillian K. & Eric Talley (2006), *On Public versus Private Provision of Corporate Law*, J. LAW ECON. & ORG. 14: 137-63.
- Hirshleifer, Jack (1991), *The Paradox of Power*, ECONOMICS & POLITICS, 3, 177–200.
- Hylton, Keith N. & Salinger, M. (2001), *Tying Law and Policy: A Decision Theoretic Approach*, ANTITRUST LAW JOURNAL, 69, 469–526.
- In re Winship 397 U.S. 358 (1970).
- Kaplan, John (1968), *Decision Theory and the Factfinding Process*, STAN. L. REV. 20:1065-92.
- Kaplow, Louis (1994), *The Value of Accuracy in Adjudication: An Economic Analysis*, JOURNAL OF LEGAL STUDIES, 23, 307–401.
- Kaplow, Louis (2011). “Optimal Proof Burdens, Deterrence, and the Chilling of Desirable Behavior,” *Am. Econ. Rev. Papers & Proc.* 101:277-80.
- Kaplow, Louis (2012), *Burden of Proof*, YALE LAW JOURNAL 121 Yale L.J. 738-859.
- Katsoulacos, Yannis S. (2008), *Optimal Legal Standards for Refusals To License Intellectual Property: A Welfare-Based Analysis*, JOURNAL OF COMPETITION LAW & ECONOMICS, 5, 269–95.
- Katsoulacos, Yannis & David Ulphz (2009), *On Optimal Legal Standards for Competition Policy: A General Welfare-Based Analysis*, JOURNAL OF INDUSTRIAL ECONOMICS, 57, 410–37.
- Kraitchit, Maurice (1953), MATHEMATICAL RECREATIONS (New York: Dover Publications, 2nd ed.).
- Lando, H. (2006), *Does Wrongful Conviction Lower Deterrence?*, JOURNAL OF LEGAL STUDIES, 35, 327–37.
- Main, Thomas (2010), *The Procedural Foundation of Substantive Law*, WASH. U. LAW REV. 87:801-41.
- Microsoft Corp. v. i4i Ltd. P’ship, 131 S. Ct. 2238 (2011).
- Milgrom, P. & Roberts, J. (1986), *Relying on the Information of Interested Parties*, RAND JOURNAL OF ECONOMICS, 17, 18–32.
- Polinsky, A. Mitchell & Steven Shavell (2000). *The Economic Theory of Public Enforcement of Law*, J. Econ. Lit. 38: 45-76.

- Posner, Richard (1999), *An Economic Approach to the Law of Evidence*, STANFORD LAW REVIEW, 51, 1477–1546.
- Png, Ivan P. L. (1986). Optimal Subsidies and Damages in the Presence of Judicial Error Int'l Rev. L. & Econ. 6: 101-5.
- Rubinfeld, Daniel & David Sappington (1987), *Efficient Awards and Standards of Proof in Judicial Proceedings*, RAND JOURNAL OF ECONOMICS, 18, 308–15.
- Sanchirico, Chris William, “Enforcement by Hearing: How the Civil Law Sets Incentives,” Columbia Law and Economics Discussion Paper No. 9596-03 (1995).
- Sanchirico, Chris William (1997), *The Burden of Proof in Civil Litigation: A Simple Model of Mechanism Design*, INTERNATIONAL REVIEW OF LAW & ECONOMICS, 17, 431–47.
- Sanchirico, Chris William (2000), Games, Information and Evidence Production: With Application to English Legal History, AM. L. & ECON. REV. 2:342-80.
- Sanchirico, Chris William (2001), Relying on the Information of Interested--and Potentially Dishonest--Parties, 3 AM. L. & ECON. REV. 320-57.
- Sanchirico, Chris William (2008), *A Primary Activity Approach to Proof Burdens*, J. LEGAL STUD. 37: 273-313.
- Schaffer *ex rel.* Schaffer v. Weast, 546 U.S. 49 (2005).
- Shin, H.S. (1998), *Adversarial and Inquisitorial Procedures In Arbitration*, RAND J. OF ECON. 29: 378–405.
- Skaperdas, Stergios (1996), *Contest Success Functions*, ECONOMIC THEORY, 7, 283–90.
- Sobel, Joel (1985), *Disclosure of Evidence and Resolution of Disputes: Who Should Bear the Burden of Proof?*, in Alvin E. Roth, ed., GAME THEORETIC MODELS OF BARGAINING (New York: Cambridge University Press).
- Stein, Alex (2005), THE FOUNDATIONS OF EVIDENCE LAW (Oxford University Press).
- Strong, John William (1999), MCCORMICK ON EVIDENCE §342 (St. Paul, Minn.: West Group, 5th ed.).
- Tullock, Gordon (1980), *Efficient Rent Seeking*, in J.M. Buchanan, R.D. Tollison & G. Tullock, eds., TOWARD A THEORY OF THE RENT-SEEKING SOCIETY (College Station: Texas A&M University Press).
- Wigmore, John H. (1940), *Wigmore on Evidence*, (3d ed.).
- Winter, Ralph, (1971), *The Jury and the Risk of Non-Persuasion*, Law & Society Rev. 5: 335-44.

Wright, Charles A., Kenneth W. Graham, Jr., Victor James Gold, Michael H. Graham (2012),
21B Federal Practice & Procedure (2nd Ed.).

Zamir, Eyal & Ilana Ritov (2012), *Loss Aversion, Omission Bias, and the Burden of Proof in
Civil Litigation*, J. LEGAL STUD., 41(1) 165-207.