Proof Discontinuities and Civil Settlements

*Mark Spottswood**

This Article explores settlement incentives under three different burden of proof rules. The conventional burden of proof is a discontinuous step-function, jumping from no damages to full damages at the 0.5 jury confidence level. Continuous burdens of proof, by contrast, would permit sanctions to steadily increase as juror confidence rises from 0 to 1, with no discontinuity. Linear burdens, which have received extensive attention in prior literature, escalate sanctions steadily across the whole range of confidence levels, while the logistic burden takes a nonlinear form.

Using a data simulation approach guided by the empirical realities of American civil litigation, I consider the incentives that each of these rules creates for parties contemplating settlement, using a model in which parties make divergent forecasts of their expected outcomes at trial due to optimism bias. Based on this analysis, I conclude that a linear burden would likely raise our settlement rate by a modest amount, except in very large cases and in "easy" cases, in which an unbiased person would predict that a trial factfinder would have a level of confidence in liability quite close to either zero or one. I also compare the expected error rate of the settlements that each rule produces, and find that the linear rule modestly lowers the expected error rate of settlement overall, although this benefit does not hold for easy cases or those with very high damages. Lastly, I conduct a similar analysis for the logistic burden, finding that it induces a similar quality and quantity of settlements as we currently achieve using conventional burdens.

---

* David and Deborah Fonvielle Professor, Florida State University College of Law. I am deeply grateful to Ronen Avraham, Lee Fennell, Elissa Gentry, Eric Kades, Leo Katz, Shay Lavie, Saul Levmore, and Omer Pelled for their thoughtful comments on early drafts of this Article. Thanks as well to Joshua Funderburke for his research assistance. Cite as: Mark Spottswood, *Proof Discontinuities and Civil Settlements*, 22 Theoretical Inquiries L. 201 (2021).

Introduction

Burden of proof rules mediate between factfinders' judgments regarding the likelihood of guilt or liability and the imposition of sanctions. There is a vast literature on the topic of these burdens, exploring the nature of existing burdens as well as the effects that differing burden rules may have on various outcomes such as error rates at trial and deterrence of wrongdoing. Reading this literature in isolation, one might easily be led to believe that the vast majority of civil cases in our legal system receive a trial on the merits. Anyone with even a passing familiarity with the empirical realities of the American litigation system will realize, however, that this assumption is radically false. The vast majority of our cases end either in settlement or in various pretrial dispositions, such as summary judgments or voluntary dismissals. Only a tiny handful reach the trial stage.[1] As a result, any serious analysis of burden of proof rules must acknowledge that any effects these rules have on *trial* outcomes may be outweighed, at the level of social policy, by even a small effect they have on the rate or quality of settlements or other pretrial dispositions.

Of course, as an author I have been just as guilty as anyone of the tendency I describe above. In a prior article, I devoted extensive efforts towards analyzing the impacts that three different burdens of proof might have at the trial stage, mentioning pretrial impacts only in passing.[2] In the present Article I seek to remedy that oversight, by offering a tentative analysis of the ways that those same three burden of proof rules might alter the incentives that civil parties face to settle cases.

In the analysis that follows, I study the effects of three different kinds of proof burdens on the incentives that parties would have when making decisions to settle a case. The first type of burden I analyze is the traditional preponderance standard. This burden is formally equivalent to a *discontinuous step-function*, awarding $0 in damages to the plaintiff below 0.5 confidence that the defendant is liable, and full damages at higher levels of confidence. In addition to being our current default burden of proof rule in civil cases, this rule is worth studying carefully because it is the approach that, on reasonable

---

1 *See* Paula Hannaford-Agor, Scott E. Graves & Shelley Miller, *The Landscape of Civil Litigation in State Courts*, Nat'l Ctr. St. Cts. 25 (2015) [hereinafter *Landscape Report*] (reporting that only 3.5% of state court cases in their sample were resolved by bench or jury trials).

2 *See generally* Mark Spottswood, *Continuous Burdens of Proof*, Nevada L.J. (forthcoming 2021).

assumptions, minimizes the size of expected errors when it is applied by factfinders at trial.[3]

I first compare this discontinuous burden with a proof burden that has often been praised in the law and economics literature.[4] The *continuous linear* burden of proof rule assigns damages as a simple product of the factfinder's level of confidence in liability and the full amount of recoverable damages caused by a defendant's conduct. Under this rule, if a jury concludes that it is 75% likely that the defendant committed conduct that renders him liable for a plaintiff's injury and values the plaintiff's total damages at $100,000, the jury should award $75,000; conversely, if the jury thought liability was 25% likely, it should award $25,000.[5] This type of burden is attractive because it improves deterrence effects under a variety of situations in both civil and criminal cases.[6] It also has other benefits that have been less broadly recognized: it tends to diminish the magnitude of error caused by a variety of improper influences on the judicial process, such as disparities of party wealth or the impact of racial or other biases on jurors. It further has a valuable property of reducing the overall magnitude of the errors that are inflicted on individual parties, thus distributing the risk of wrongful litigation outcomes in a fashion that is somewhat analogous to risk pooling in insurance.[7] And finally, by reducing

---

3 *See generally* David Kaye, *The Limits of the Preponderance of the Evidence Standard: Justifiably Naked Statistical Evidence and Multiple Causation*, 7 Am. B. Found. Res. J. 487 (1982).

4 *See, e.g.*, Talia Fisher, *Conviction Without Conviction*, 96 Minn. L. Rev. 833 (2012); Henrik Lando, *The Size of the Sanction Should Depend on the Weight of the Evidence*, 1 Rev. L. & Econ. 277 (2005); Steven Shavell, Economic Analysis of Accident Law 115-17 (1987); David Rosenberg, *The Causal Connection in Mass Exposure Cases: A "Public Law" Vision of the Tort System*, 97 Harv. L. Rev. 849, 861-67 (1984).

5 The linear rule thus provides more compensation to plaintiffs for cases in the bottom half of the confidence range, and less in the upper half of the range, relative to the traditional rule. For explorations of the impact of such a change on the distribution of expected errors at trial, see Spottswood, *supra* note 2, at Parts I-B and II-A, and Kaye, *supra* note 3.

6 *See* Spottswood, *supra* note 2, at Part I-A; Fisher, *supra* note 4, at 857-59 (discussing under-deterrence in criminal cases); Shavell, *supra* note 4, at 115-17 (discussing over- and under-deterrence in civil tort cases).

7 Risk pooling refers to arrangements that allow multiple individuals who are subject to a risk of injury to share their losses more evenly, rather than concentrating those losses entirely on a small subset of unlucky individuals. *See* Amy B. Monahan, *Health Insurance Risk Pooling and Social Solidarity: A Response to Professor David Hyman*, 14 Conn. Ins. L.J. 325, 326 (2008).

the payoff that arises from small shifts in jury assessments of likelihood of liability, the linear rule may also disincentivize the spoliation of evidence.[8]

As will be seen below, the interaction of these rules with incentives to settle is complicated, in part because it requires us to make assumptions regarding the level of confidence with which parties forecast jury verdicts in advance of trial. To see this, note first that if parties could forecast jury confidence without uncertainty, then they would only predict outcomes of certain wins or losses under the discontinuous burden of proof. In the real world, parties obviously express uncertainty over possible outcomes, and this can best be captured by modelling trial forecasts in terms of a continuous distribution of predicted levels of jury confidence with both a mean value and a variance, each of which will influence settlement outcomes. To make a study of the effects of differing burden rules on settlement behavior tractable, I develop a simple settlement model that is one-shot, with parties negotiating based on shared information and equivalent future litigation costs, and in which party forecasts of the expected outcome at trial are distorted by optimism bias.[9] After showing that neither rule has analytically superior behavior with respect to settlement incentives in all cases, I then simulate a large number of hypothetical cases that mimic real-world civil litigation in America, and analyze both the number of settlements and the quality of settlements that arise under each alternative burden of proof rule. As shall be seen, the linear burden of proof rule modestly increases the settlement rate while modestly decreasing the expected error rate of the resulting settlements. Thus, those who find the settlement of cases to be preferable to trying them on the merits may find the linear burden particularly attractive.

Of course, some analysts might find that prospect unattractive,[10] and prefer a rule that encourages bringing slightly more cases to trial. In prior work, I have described the *logistic burden of proof*, a functional form that allows us to smoothly vary the imposition of liability in ways that can closely approximate either of the other two burdens, and which also lets us split the difference

---

8 *See generally* Spottswood, *supra* note 2, at Part II (surveying these other considerations).

9 The model builds on the Posner-Landes approach to predicting settlements, but it is elaborated to model the impacts of varying jury confidence distribution forecasts on outcomes, and then further supplemented by extensions that apply to the alternative burden rules. *See generally* Richard A. Posner, *An Economic Approach to Legal Procedure and Judicial Administration*, 2 J. Legal Stud. 399 (1973); William M. Landes, *An Economic Analysis of the Courts*, 14 J. L. & Econ. 61 (1971).

10 *See, e.g.*, Owen M. Fiss, *Against Settlement,* 93 Yale L.J. 1073, 1085-87 (1984).

between them.[11] Logistic burdens, like linear burdens, operate continuously, so that there is a steady escalation of sanctions as confidence in liability increases. Instead of a simple product, however, the logistic rule imposes the following relationship between damages (D), confidence in liability (p), and the size of an award (J):[12]

$$J(p) = \mathrm{D} * \left(\frac{A}{1 + e^{-rp+\frac{r}{2}}} + B\right)$$

This function produces a sort of "s-shape," which initially escalates slowly at very low levels of confidence, then rises more quickly through its middle range, and finally rises more slowly again as confidence in liability approaches 100%. One advantage of such a burden is that it lets us strike a balance between error costs and deterrence costs at trial, rather than make an all-or-nothing choice to optimize one kind of cost at the expense of the other. Perhaps surprisingly, a logistic burden does not simply split the difference between the settlement incentives that arise under the linear or the step-function rules. Instead, it produces settlements that closely mimic the step-function burden, in both frequency of settlements and expected error costs. As a result, analysts who appreciate the benefits of continuous trial burdens, but who wish to prevent any further reduction of our rate of trials,[13] may find the logistic burden particularly attractive.

The remainder of this Article proceeds as follows. Part I explicates the method of analysis that I have used, including both the underlying economic models and the assumptions used to generate simulated cases. In Part II, I show that the linear rule should generally increase incentives to settle cases by a modest amount, explain why this is the case, and also show some situations in which the anticipated increase in settlements would be less likely. Part III focuses on a linear rule's impact on the accuracy of the resulting settlements, showing that the linear rule would also increase settlement accuracy, but that this effect is produced largely by the way it operates in cases with unusually high litigation costs. Next, Part IV extends the analysis to include a logistic burden of proof, and shows that the logistic burden would produce an expected pattern of settlements that are much more similar to what we currently observe. To be clear, this analysis is not intended to present a complete normative case

---

11 *See* Spottswood, *supra* note 2.

12 A and B are scaling constants in the above equation.

13 *See generally* Marc Galanter, *The Vanishing Trial: An Examination of Trials and Related Matters in Federal and State Courts*, 1 J. Empirical Legal Stud. 459 (2004) (documenting a dramatic decline in trial rates across both federal and state courts during the second half of the 20th century).

for the superiority of any one burden. Not all theorists will agree, for instance, on whether an increase in the settlement rate is a net benefit or a net cost for the system as a whole.[14] My more modest goal is simply to elucidate how these rules shape parties' settlement behavior.

I. Method of Analysis

This section describes the mode of analysis that I used to explore the ways that differing burden-of-proof functions influence settlement behavior. First, I will describe the underlying model of how jury confidence forecasts interact with the decision to settle a case, under both a linear continuous burden of proof and the traditional discontinuous step-function burden. Once the basic model has been laid out, I will then give a summary of my data simulation approach to assessing the impacts that the burdens create on decisions to settle cases and settlement amounts.

A. Settlement Models

My overall approach is based on the classic Posner-Landes model of settlement, with F_p and F_d representing each party's forecast of their expected outcome if the case goes to trial, T representing the cost of going to trial, and S representing the costs of settling the case.[15] The model assumes that settlement is a one-shot decision, that the stakes and costs are symmetric, that parties are similarly informed about the facts of the case, but that they each view the expected trial outcome with varying levels of optimistic bias.[16] It also assumes that we

14 Moreover, even calculating the overall net impact of each rule on outcome accuracy (which would be a useful first step towards a normative assessment) is not feasible at present. To be even minimally credible, such an analysis would have to go beyond analyzing the net impacts on accuracy of settlement amounts and trial verdict amounts, and also include the outcomes of other important pretrial resolutions, such as summary judgment decisions. Sadly, we do not yet have a good understanding of how summary judgment should work in a system that implements continuous burdens of persuasion at trial, so any attempt at aggregate analysis would be premature.

15 *See* Posner, *supra* note 9; Landes, *supra* note 9.

16 The Posner-Landes model was later extended by Priest and Klein to analyze the relationship between underlying merit in the overall pool of cases and the expected victory rate of plaintiffs at trial. *Cf.* George L. Priest & Benjamin Klein, *The Selection of Disputes for Litigation*, 13 J. Legal Stud. 1 (1984). In the course of this exploration, Priest and Klein extended the model to accommodate

are applying the "American Rule," under which each party pays their own litigation costs and attorneys' fees, regardless of whether they prevail on the merits. Under these assumptions, settlement occurs whenever the following inequality is satisfied:

$$F_p - T + S \leq F_d + T - S$$

To get a basic sense of how this model works, assume for a moment that two parties have exchanged discovery materials and are contemplating whether to either attempt to settle a case or take it to trial. Typically we would expect the costs of bargaining towards a settlement to be less than trying a case,[17] so let's say for now that the bargaining will cost each party $5,000 and the trial $20,000. In such a case, we can simplify the condition for settlement as follows:

varying stakes for each party in a case. *See id.* at 25. An interesting (if more complex) future extension of this study might profitably incorporate varying stakes into the model, as well as the possibility of differing litigation costs for either party, although I avoided such complexities in the analysis above, due to limits on available data that could be used to estimate such quantities. Likewise, a more elaborate approach might model each party's estimate of the probability of success based on an explicit representation of variable levels of information, rather than taking the simpler route of assuming that variations are due to simple optimism bias. *Cf.* Holger Sieg, *Estimating a Bargaining Model with Asymmetric Information: Evidence from Medical Malpractice Disputes*, 108 J. Pol. Econ. 1006 (2000) (using an unusually rich dataset regarding Florida medical malpractice cases to estimate a model incorporating variable information levels and litigation costs for either party, in an unfortunately unrepresentative subset of American litigation in general). Given my simplified approach, this Article should be read as an indication of what incentives we should expect parties to have following an exchange of relevant information in discovery, in cases where each side foresees reasonably similar stakes and litigation costs. For the most part, variations in stakes or costs should be orthogonal to the effects modelled here, and thus the basic findings of this Article should still hold. However, an interesting (if challenging) extension to the current project might consider the interaction effects between burden of proof rules and the amount of discovery taken. To the extent that the amount of discovery effort could vary endogenously within a model, both costs and information levels might be influenced by the choice of a burden of proof rule, producing possible complications to the effects I show here.

17 *See* discussion at Appendix I-C.

$$F_p - \$20{,}000 + \$5{,}000 \le F_d + \$20{,}000 - \$5{,}000$$

$$F_p \le F_d + \$30{,}000$$

Thus, in this simplified form, we can see that settlements should occur whenever the plaintiff's forecast of their trial outcome does not exceed the defendant's forecast by a sufficiently large margin (representing twice the difference between the trial costs and the settlement costs). Thus, the likelihood of settlement will hinge on factors that cause the two parties to foresee different outcomes at trial. These could include variation in beliefs concerning what evidence will be admitted at trial, variations in estimates concerning a jury's likely confidence levels in liability given that evidence, and (as we shall see) variations in the *precision* with which those levels of confidence are forecasted by either party.

Comparing the likelihood of settlement for a given case using either a linear or a conventional burden of proof will require us to determine what the values of F_p and F_d will be under each burden. To determine these values, we must first model the parties' assumptions regarding the likely confidence levels that a jury would feel with respect to the facts establishing liability following a trial on the merits. To begin with, we can describe the way that an *unbiased* observer would forecast the probability, p, that a future jury will place on the likelihood of liability. Since the range of possible values of p can vary continuously between 0 and 1, we will define these predictions as the probability density function of a beta distribution, C_u:

$$C_u = beta(\alpha, \beta) = f(x, \alpha, \beta) = x^{\alpha-1}(1-x)^{\beta-1}\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}, with$$

$$\alpha \; and \; \beta > 0,$$

$$\alpha \; or \; \beta \ge 1, and$$

$$\alpha \; and \; \beta \le \varphi$$

The listed constraints on the values of α and β ensure that the forecasts are unimodal, with the constant φ providing a ceiling on the level of confidence that parties may have in their predictions. Thus, a typical unbiased confidence forecast, involving a case that is very likely to be viewed as close, could be characterized as follows:

Predictions of Fact-Finder's Confidence in Liability

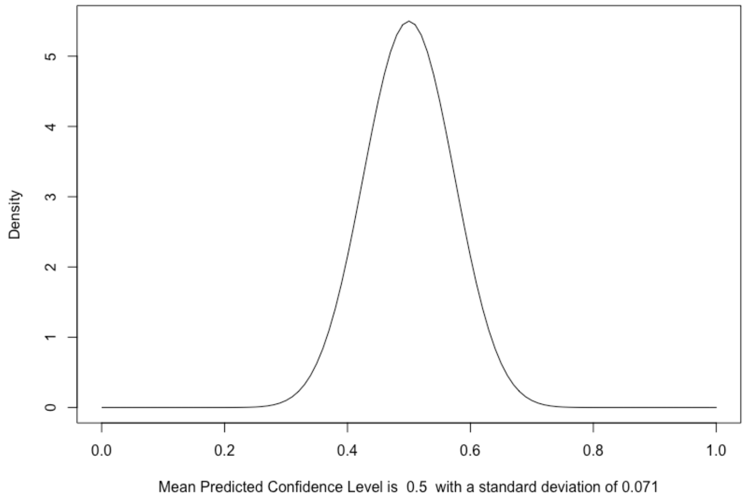

We can further assume that each party will have their own forecasts, C_p and C_d, which take the same form but tend to have distributions that are shifted in either party's favor based on randomly varying degrees of optimism bias.[18] The model permits the *average* extent of optimistic bias to be varied using the constant θ, ranging from just above 0 (low optimistic bias) to 1 (very high levels of optimistic bias).

---

18 *See generally* Oren Bar-Gill, *The Evolution and Persistence of Optimism in Litigation*, 22 J.L., Econ. & Org. 490 (2006) (surveying empirical evidence in favor of widespread optimism bias among litigators and suggesting that such biases may in fact be an adaptive means of obtaining more favorable settlements by making threats to take cases to trial more believable to opponents). *Cf.* James A. Shepperd, William M.P. Klein, Erika A. Waters & Neil D. Weinstein, *Taking Stock of Unrealistic Optimism*, 8 Persp. Psychol. Sci. 395 (2013) (surveying the broader literature on optimism bias and showing that such tendencies are shown across a wide range of human endeavors, especially in settings where individuals have partial control over outcomes and feedback on predictive accuracy occurs only long after the initial prediction is made).

$$C_p = Beta\ (\alpha + \varepsilon_1, \beta - \varepsilon_2),$$
$$with\ \varepsilon_1 = \ \theta * unif(0, \varphi - \alpha + 0.001)$$
$$and\ \varepsilon_2 = \ \theta * unif(0, \beta - 0.001)$$
$$C_d = Beta\ (\alpha - \varepsilon_3, \beta + \varepsilon_4),$$
$$with\ \varepsilon_3 = \ \theta * unif(0, \alpha - 0.001)$$
$$and\ \varepsilon_4 = \ \theta * unif(0, \varphi - \beta + 0.001)$$

The results of these varying levels of bias can be visualized as follows, for two randomly generated cases:

Parties' Optimistic Forecasts

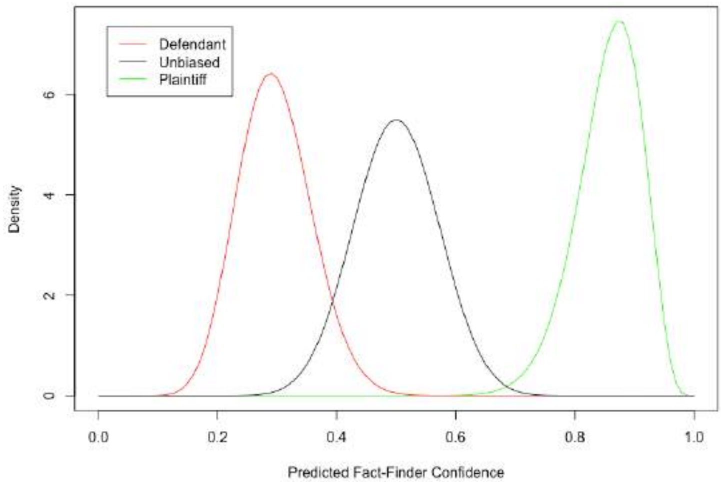

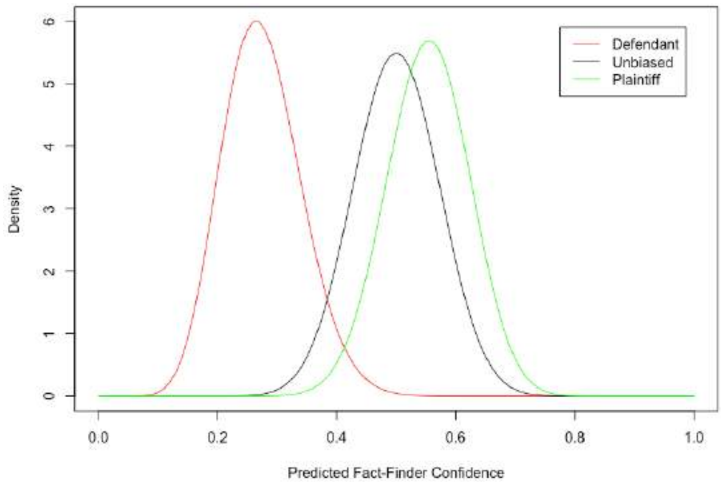

The values of F_p and F_d can then be derived for both burdens using each parties' biased forecast of possible jury confidence levels. First, here are the expected outcomes and associated victory probabilities under a discontinuous burden, in which a plaintiff receives J if the jury finds her case to be more than 0.5 probable and receives 0 when the jury finds her case to be less than 0.5 probable.

$$F_p = 0 * \int_0^{0.5} beta(x;\ \alpha_p, \beta_p)dx + J * \int_{0.5}^{1} beta(x;\ \alpha_p, \beta_p)dx$$

$$F_d = P_d J = 0 * \int_0^{0.5} beta(x;\ \alpha_d, \beta_d)dx + J * \int_{0.5}^{1} beta(x;\ \alpha_d, \beta_d)dx$$

and we can therefore redefine the expected outcomes in terms of the expected probability of victory, P_p and P_d

$$F_p = J \int_{0.5}^{1} beta(x;\ \alpha_p, \beta_p)dx = P_p J,$$

$$where\ P_p = \int_{0.5}^{1} beta(x;\ \alpha_p, \beta_p)dx$$

$$F_d = J \int_{0.5}^{1} beta(x;\ \alpha_d, \beta_d)dx = P_d J$$

$$where\ P_d = \int_{0.5}^{1} beta(x;\ \alpha_d, \beta_d)dx$$

Now we can consider expected outcomes under a linear continuous burden. For clarity, let us denote these using F'_p and F'_d. Under this rule, for a given level of jury confidence p, a plaintiff can expect to receive $J * p$, and so the expected outcome under that burden and the associated values of F'_p and F'_d can be derived as follows:

$$F'_p = \int_0^1 beta(Jx;\ \alpha_p, \beta_p)x dx \ = \ M_p\, J$$

$$where\ M_p = \int_0^1 beta(x;\ \alpha_p, \beta_p)x dx \ = \frac{\alpha_p}{a_p + \beta_p}$$

$$F'_d = \int_0^1 beta(Jx;\ \alpha_d, \beta_d)x dx \ = \ M_d\, J$$

$$where\ M_d = \int_0^1 beta(x;\ \alpha_d, \beta_d)x dx \ = \frac{\alpha_d}{a_d + \beta_d}$$

We can now specify the settlement conditions for each rule in more detail as follows:

$$V_p - T + S \ = \ P_p J - T + S \ \leq P_d J + T - S = V_d + T - S$$

$$with\ P_p = \int_{0.5}^1 beta(x;\ \alpha_p, \beta_p)dx \ and\ \ P_d = \int_{0.5}^1 beta(x;\ \alpha_d, \beta_d)dx$$

and

$$V'_p - T + S \ = \ M_p J - T + S \ \leq M_d J + T - S = V'_d + T - S$$

$$where\ M_p = \frac{\alpha_p}{a_p + \beta_p}\ and\ M_d = \frac{\alpha_d}{a_d + \beta_d}$$

These conditions for settlement under each rule can be rewritten in the following form, showing that settlement is more likely whenever the values of P_p and P_d (each party's prediction of the plaintiff's probability of victory under a discontinuous rule), or M_p and M_d (each party's mean prediction regarding the jury's likely confidence level in liability following trial), are close.

$$(P_p - P_d)J \ \leq 2(T - S)$$

$$(M_p - M_d)J \ \leq 2(T - S)$$

Analyzing these inequalities, one notes that for a given pair of distributions C_p and C_d, the values of M_p and M_d will normally lie closer to 0.5 than the values of P_p and P_d, because the definite integral of the beta distribution between 0.5 and 1 will normally have a more extreme value than the distribution's overall mean. The following figures provide a graphical illustration of this tendency, with the location of the red line indicating each distribution's mean, while the size of the shaded area illustrates the probability that the jury would find the case to be more than 0.5 likely. The first graph shows a case where an unbiased observer would say that the most likely level of jury confidence in liability following trial would be 0.67. As can be seen, the area of the shaded region (which is equivalent to the probability of victory for the plaintiff under a discontinuous rule) is higher, at 0.79:

Chance of Success Based on Predicted Fact-Finder Confidence Levels

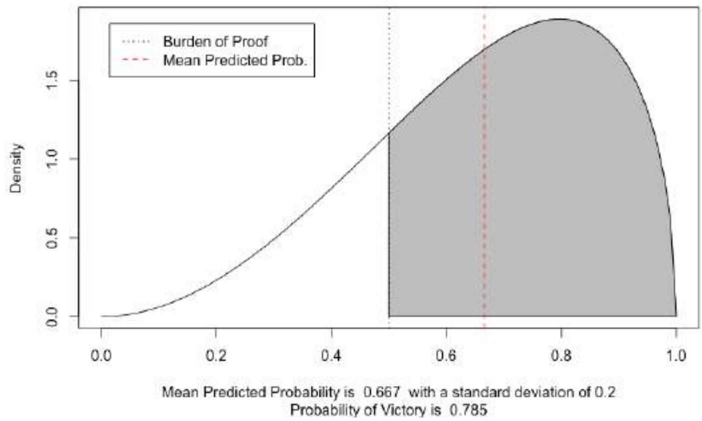

The next graph shows how the dichotomous probability of victory trends rapidly towards the extremes of the distribution as the precision of the confidence forecast grows. This time, an unbiased observer still predicts a most likely confidence level of 0.67, but now they are also much more confident that the value will not be much more or much less than that. As a result, they would estimate a 98% probability of victory under the traditional, all-or-nothing rule:

Chance of Success Based on Predicted Fact-Finder Confidence Levels

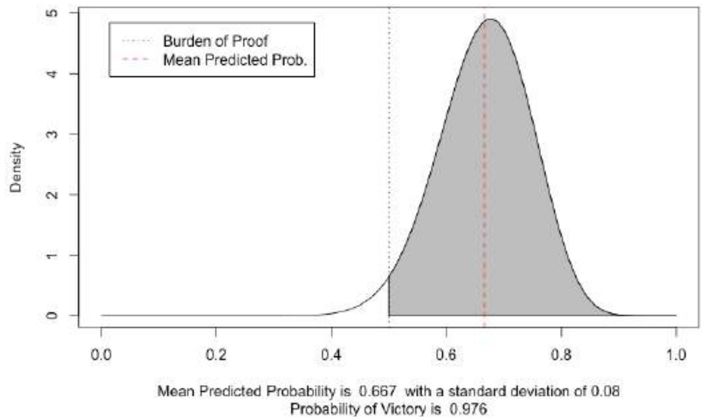

In general, the divergence decreases[19] as the case gets closer and the forecast grows cloudier. The third graph shows a case where the factfinder is almost totally unsure of the outcome. Their best guess as to the jury's confidence level is 0.55, but it could easily be much higher or much lower. As a result, the predicted probability of plaintiff victory is only a tiny bit higher, at 0.58.

Chance of Success Based on Predicted Fact-Finder Confidence Levels

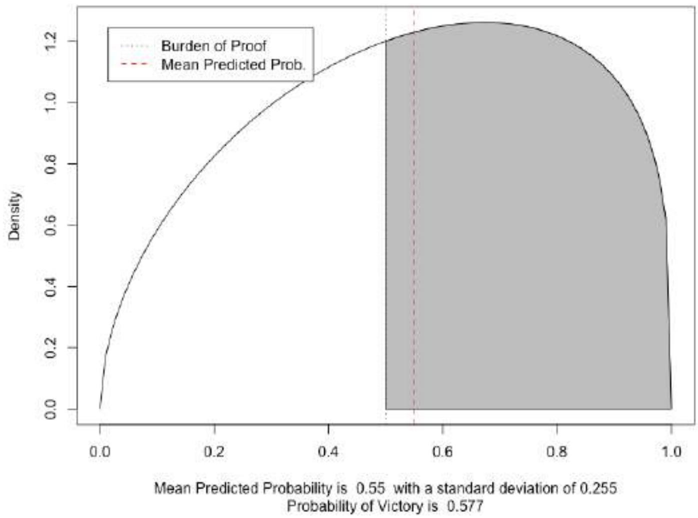

The inequalities shown above will be sufficient to predict, for a given case, whether the parties would reach a settlement under either the linear or step-function burden rules. But we must still provide a model for the size of the settlement, in order to make it possible to analyze the comparative error rate of settlements that would arise under either rule. Under the conventional burden, parties' willingness to settle under the basic model will be contingent on the possibility of making a deal that satisfies the following inequalities, with the plaintiff's lower bound on acceptable offers represented by O_p, the

19 In the limiting case, the integral will approach the mean confidence level as alpha and beta grow closer to zero. Unfortunately, such cases involve an extremely bimodal confidence prediction function. In essence, it amounts to saying, not that you think a typical decision-maker will have a certain level of confidence with some level of uncertainty in the prediction, but rather that you think the decision-makers will be evenly split into groups who are certain of guilt and certain of innocence, with the mean confidence level simply reflecting the proportion between the two groups.

defendant's upper bound on acceptable offers represented by O_d, and the settlement value itself represented by V:

$$O_p = P_p J - T + S \le V$$

$$O_d = P_d J + T - S \ge V$$

Meanwhile the range of possible settlements under a continuous burden is similarly constrained, but the range is defined instead by each party's predicted mean of jury confidence, M_p and M_d. In the equation below, O'_p and O'_d represent the lower and upper bounds on the settlement range for a continuous burden, with V' representing the settlement value:

$$0'_p = M_p J - T + S \le V'$$

$$0'_d = M_d J + T - S \ge V'$$

The model then picks a value lying somewhere within the indicated ranges[20] as the settlement that the parties would actually reach following bargaining. I use a beta distribution to represent the likelihood that negotiations will typically result in a value towards the center of the range, but that the specific results will vary based on bargaining power and bargaining ability. For example, if a plaintiff would gain by settling for any value greater than $5,000, and the defendant would be willing to settle for any amount less than $10,000, the model will predict that the most likely settlement is $7,500. At the same time, either the plaintiff or the defendant might achieve a better outcome at the negotiating table, which could lead to values that lie closer to $10,000 or $5,000, respectively. In the model, the varying quantity B accounts for these variations, with high values indicating that the plaintiff captured more of the settlement surplus and low values indicating that the defendant struck the better bargain.

$$B = Beta(12{,}12)$$

$$V = O_p + B\left(O_d - O_p\right)$$

$$V' = O'_p + B\left(O'_d - O'_p\right)$$

20 In some cases, the left side of these inequalities will be negative, indicating that the plaintiff has learned (perhaps through discovery or legal research) that the potential damages recovery is less than the expected benefit of going to trial. Since plaintiffs can simply choose to drop their suit to avoid such costs, all such values are converted to a $0 lower bound in the model.

One complication that emerges from this model is that both the likelihood of settlement and the size of the resulting range of possible settlement outcomes hinge on the distances between P_p and P_d (for the conventional burdens) or M_p and M_d (for the linear continuous burden). Unfortunately, the relationship between the sizes of these two ranges varies depending on both the mean and the variance of the parties' forecasts of jury confidence, C_p and C_d. When M_p is greater than 0.5 and M_d is less than 0.5, then the extremizing tendency of the transformation from M to P will generally increase the divergence between the parties' expectations, decreasing the probability of settlement.[21] I will explore the variations that may arise under the two rules in more detail in Part II below, but for now a pair of examples will help to illustrate the difficulty. In the first example, an unbiased forecast would yield an expected factfinder confidence level of 0.5 in liability. Each party makes a slightly biased forecast, with the defendant expecting the factfinder to think that liability is 0.38 likely, and the plaintiff forecasting a confidence level of 0.63. The first of the two charts below shows the resulting impact on the defendant's predicted probability of a plaintiff victory under a conventional rule, while the second chart shows the impact on the plaintiff's prediction:

21 This extremizing tendency is similar to what has been found in models of the British rule, in which the loser pays the opposing side's litigation costs and attorneys' fees. *See* John J. Donohue, *Opting for the British Rule, Or If Posner and Shavell Can't Remember the Coase Theorem, Who Will?*, 104 Harv. L. Rev. 1093, 1096-99 (1991) (summarizing model results from Posner and Shavell, which show that expected outcomes tend to diverge more under the British rule as parties' level of optimism increases). As will be seen below, however, this similarity only exists for a subset of the cases; for others, it will be the linear burden that produces more divergent forecasts.

Predicted Confidence v. Victory Probability, Close Case

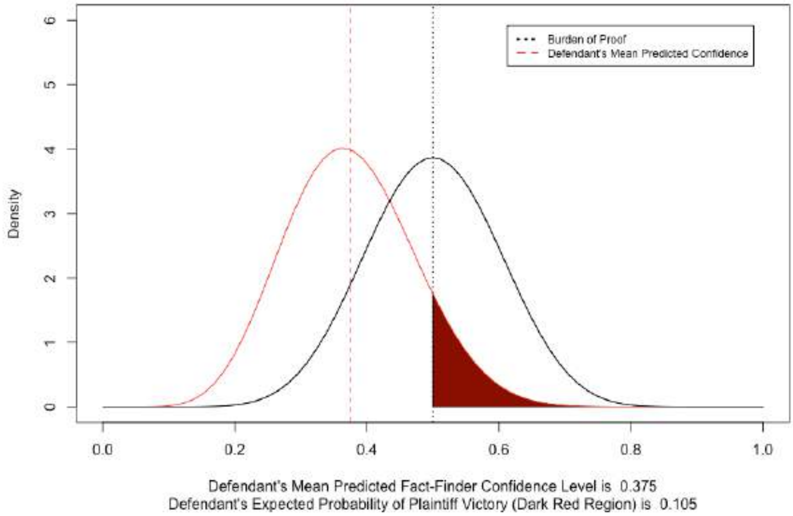

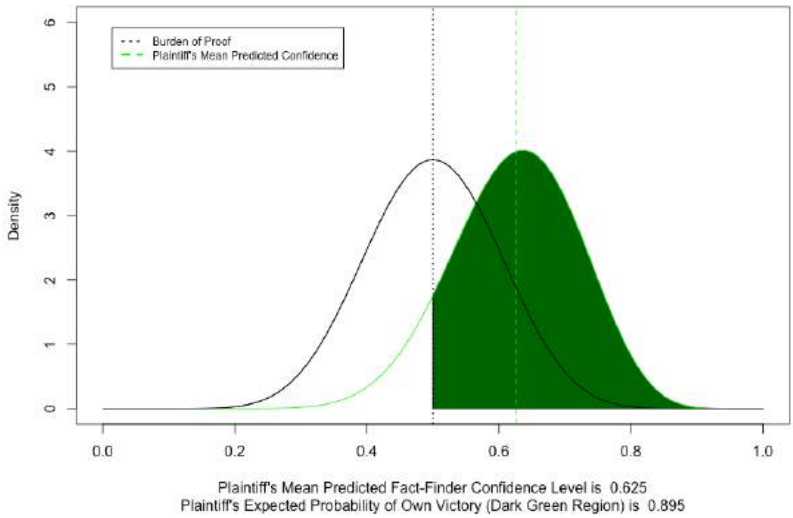

As can be seen, the distance between the predicted victory probabilities under the conventional rule is quite large (0.90 – 0.11 = 0.79), while the gap between the mean predictions of factfinder confidence is much narrower (0.63 – 0.38 = 0.25). Thus, in this example the case would be more likely to settle under a linear burden, as the divergence in the parties' expected result would only be .25 the damages, whereas under a conventional burden the divergence would be a much larger .79 times the damages, which is more likely to make the cost of a trial worthwhile.

In the second example, both parties foresee that the jury is likely to favor the plaintiff. The graphs below show a case where an unbiased prediction would yield a mean expected jury confidence level of around 0.75. The defendant optimistically predicts a mean factfinder confidence level of only 0.64, which yields a dichotomous victory probability for the plaintiff of 0.88:

Predicted Confidence v. Victory Probability, Case Favoring Plaintiff

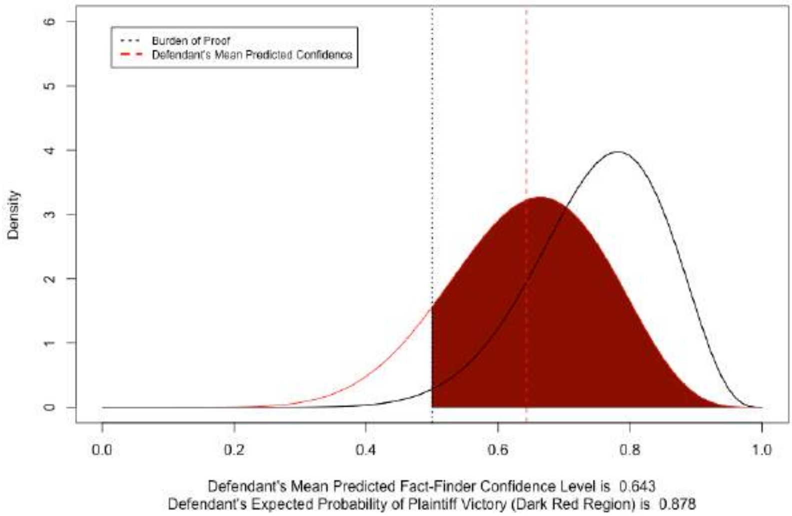

The plaintiff, meanwhile, predicts a mean fact-finder confidence level of 0.83, which results in a predicted dichotomous victory probability of 1 under the conventional rule:

Predicted Confidence v. Victory Probability, Case Favoring Plaintiff

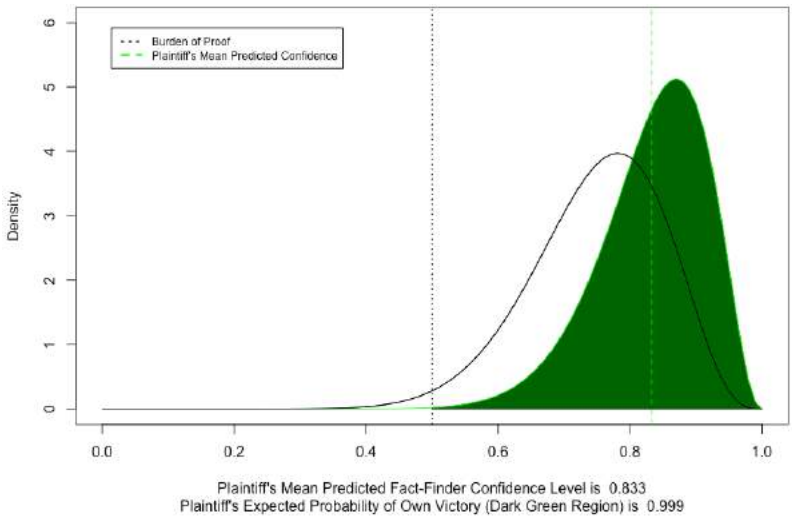

Now the parties would expect more divergent results under the linear burden (a difference of .19 times the damages) versus under the conventional burden (with a difference of only .12 times the damages). For this reason, it is hard to determine which rule might have more favorable properties in terms of the settlements they would produce through analytic methods, as the answers will depend on the values one assumes for varying parameters in the model. The next section describes the design of a simulation model that provides some illumination regarding the actual settlement incentives that these rules might create based on variations in case strength, damages amounts, and litigation costs.

B. Simulation Design

In my simulation, cases are generated at random, based on certain assumptions about the distribution of merit, damages, and litigation costs in American civil litigation, as well as varying rates of precision in an unbiased observer's forecast regarding case merit. The details by which case parameters were generated are summarized in Appendix I.

Damages and cost values were generated to match observed values for a representative selection of state court civil cases, drawing on data collected by the National Center for State Courts. As in real life,[22] the vast majority of simulated cases had amounts in controversy between $100 and $10,000, although there were also a small fraction of cases with much higher stakes. Expected damages if cases were taken to trial followed an approximately log-normal distribution, although with a slightly fatter right tail:

22 *See* Landscape Report, *supra* note 1, at 27 (showing a very large number of bench trials, with a mean verdict for prevailing plaintiffs of $6,408, and a much smaller fraction of jury trials, with a mean verdict for prevailing plaintiffs of $1,468,554).

Expected Damages If Case Is Tried

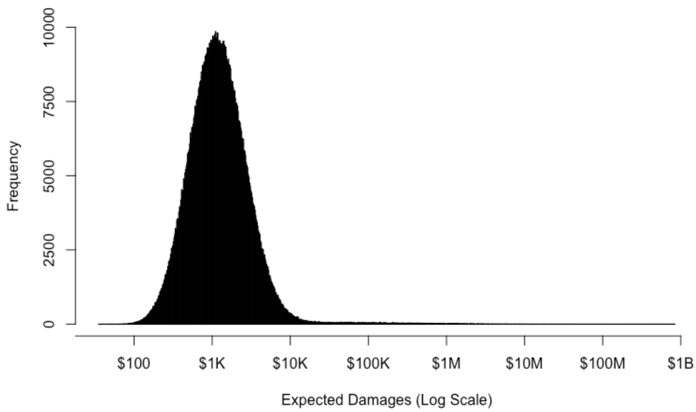

I then generated values for overall litigation costs for each litigated case. In the absence of representative data across all types of cases tried in state courts, these values were based on log-normal distributions that were fitted to reported overall litigation costs for three types of cases for which cost data was available: a large proportion of debt collection cases with low costs, a smaller but still substantial fraction of automobile accident cases for which costs were moderate, and a much smaller fraction of professional malpractice cases with very high costs.[23] The data were then further adjusted to produce a .25 log-log correlation with the expected damages values, as follows, with L_{naive} representing the raw values and $L_{adjusted}$ representing the values after adjustment to produce the required correlation.

$$L_{adjusted} = L_{naive} + \frac{unif(0, J - L_{naive})}{2}, if\ J > L_{naive}$$

$$L_{adjusted} = L_{naive},\ if\ J = L_{naive}$$

$$L_{adjusted} = L_{naive} - \frac{unif(0, L_{naive} - J)}{2}, if\ J < L_{naive}$$

23 *See* Paula Hannaford-Agor & Cynthia G. Lee, *Utah: Impact of the Revisions to Rule 26 on Discovery Practice in the Utah District Courts*, Nat'l Ctr. St. Cts. 132 (2015) [hereinafter *Utah Report*] (providing median and interquartile range cost estimates for these case types, as well as others whose proportions could not be identified in the overall report of state cases provided in the Landscape Report).

The resulting overall litigation costs had the desired correlation with damages amounts,[24] and had the following "one-shouldered" overall distribution, with a similar mode but a much more prominent right tail, and a narrower overall range:

Expected Overall Litigation Costs If Case Is Tried

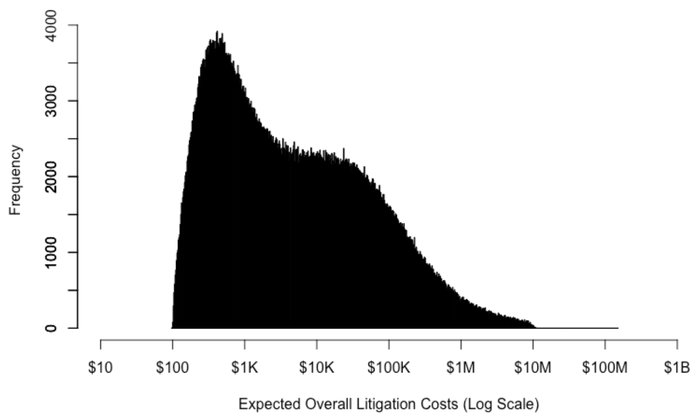

Specific values for settlement and trial costs are then produced by taking mean proportions of these costs from the overall NCSC data, and then adding random variation around those means:[25]

$$T = 0.375\, L_{adjusted} + \mathcal{N}(0, (\frac{0.375 L_{adjusted}}{5})^2)$$

$$S = 0.065 L_{adjusted} + \mathcal{N}(0, (\frac{.65 L_{adjusted}}{5})^2)$$

The model then generates values for an unbiased jury confidence forecast for each case, C_U, according to the following constraints, which are designed to ensure that the cases involve an even spread over levels of true case strength, while preserving the unimodality of the forecast of jury confidence:

---

24 Log-log correlation was .2499 in the overall dataset, with p < 2.2e-16. The distribution was also adjusted to remove cases with implausibly small or large costs that were generated by the basic method. *See* Appendix I-C for details.

25 *See* Appendix 1 for discussion of relevant assumptions.

$$\mu = unif(0.1, 0.99)$$

$$var = \frac{unif(0.001, \mu)}{10}, if\ \mu \leq 0.5$$

$$var = \frac{unif(0.001, 1-\mu)}{10}, if\ \mu > 0.5$$

$$\alpha = (\frac{1-\ \mu}{var} - \frac{1}{\mu})\mu^2$$

$$\beta = \alpha(\frac{1}{\mu} - 1)$$

$$C_u = beta(\alpha, \beta)$$

With these values[26] for the parameters of C_U, the simulation then applies the formulas used by the model specified above to provide values for the remaining constants required to solve the settlement condition inequalities and to determine settlement values (when applicable) under each rule. With all the ingredients of both settlement functions in hand, we can simulate cases and determine which ones will settle, and for how much, under the differing burden of proof rules.

Using the determinants for each of the random variables as well as the models for determining settlement outcomes outlined above, I generated a dataset of 1,000,000 simulated cases. The descriptive statistics for these cases mirror the range one would expect to see in typical American civil cases.[27] The next two Parts of this Article analyze these data in order to gain insight regarding the type and frequency of settlements that discontinuous and linear continuous burdens tend to produce.

---

26 In practice, the simplified approach included above generated some cases with implausible confidence forecasts. These were algorithmically pruned to generate the datasets used for analysis here. I describe the rationale for such pruning and its impact on the overall universe of cases analyzed in the Appendix.

27 *See generally* Appendix I (summarizing the constructed dataset and explaining the rationales chosen for its parameters).

II. Impact on Settlement Rates

Within the universe of simulated cases, the first apparent fact is that applying a linear burden produced a slight increase in the rate at which cases resulted in settlement. Under the model for settlements with a standard, discontinuous burden, a total of 612,595 cases settled, which is 61% of all cases.[28] Conversely, using the linear burden resulted in an uptick to 645,673 settlements, which represents 65% of the cases, or an increase of 5% from the rate under the current rule.[29]

Looking at the simulated cases more closely, we can refine our understanding of the ways in which the two rules differ in terms of providing settlement incentives. As discussed above, the relationship between predicted mean confidence levels and predicted probabilities of victory under a dichotomous rule behaves differently in cases that straddle the proof burden discontinuity, as compared with cases where both values lie on the same side of the discontinuity. This implies that it should be easier to settle cases under a linear burden when the predicted mean confidence levels are close to 0.5, and easier to settle cases under a discontinuous burden when the predicted confidence levels are close to 0 or 1. That is exactly what we observe in the simulated cases:

---

28 The degree-of-optimism bias parameter was tuned to a value of 0.75 so that the observed rate of settlement under the traditional burden tracks the observed rate of settlement of state court cases. *See* Landscape Report, *supra* note 1, at 7 (summarizing the 1992 Civil Justice Survey of State Courts, which showed an overall settlement rate of 62% using a detailed case-file review method). *But cf. id.* at 20 (reporting a lower rate of settlements in a newer dataset that better accounts for small cases, while acknowledging that several other outcome coding categories might include settled cases, due to variations in reporting choices among state courts). Note that although the comparative behavior of the two rules remains similar at varying levels of optimism bias in the settlement model, the specific rates of settlement under the two rules will vary, with higher rates of optimism generally resulting in a lower rate of settlement. If one wishes to more precisely generalize these results to systems with a known rate of settlement that is higher or lower, the optimism parameter can be adjusted accordingly.

29 This difference in proportions was estimated extremely precisely, with $p < 2.2e{-16}$.

Settlement Rate Under Each Rule, by Case Strength

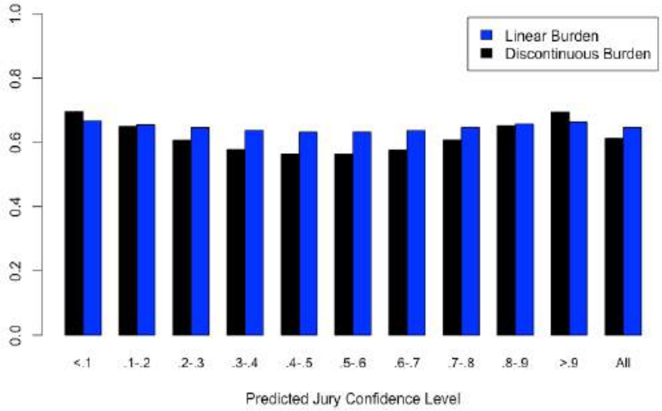

As can be seen, a discontinuous burden makes settlement more common for cases with an unbiased predicted confidence level below 0.1 or above 0.9, while cases in the middle range are more likely to be settled under a continuous burden.[30] In other words, the traditional burden provides a modestly larger settlement incentive for the small subset of cases where a neutral observer would predict that a trial jury would conclude that liability is either nearly certain or almost certainly false. If a case was only moderately hard (such that an unbiased observer would predict a mean jury confidence level of 0.75 or 0.25, for example), the linear rule starts to have a modest advantage in terms of its ability to encourage settlements. Finally, that effect grows larger for very hard cases, where the observer would predict a likely jury confidence level close to 0.5.

Thus, assuming that filed cases are evenly distributed over the range of mean juror confidence levels, we should naturally expect the continuous burden to increase the settlement rate, because it is advantaged over 80% of the range of possible confidence levels. Moreover, one can see from the chart that its advantages are *larger* over the central range of cases than the corresponding advantage that the discontinuous burden has in the easy cases lying to either

30 The overall proportions of settlements for the moderate-to-hard cases (0.1 – 0.9 unbiased predicted mean jury confidence levels) was 54% for the discontinuous rule versus 50% for the linear rule ($p < 2.2e-16$). By contrast, the settlement proportions for the easy cases (less than 0.1 or above 0.9 unbiased predicted mean jury confidence levels) was 70% for the discontinuous rule and 67% for the linear rule ($p<2.2e-16$).

extreme.[31] Thus, unless filed cases are disproportionately tilted towards easy wins for the plaintiff or the defendant, the continuous burden would tend to create greater incentives towards settlement overall.

Of course, it is also quite clear from the models that the overall size of cases will affect the desirability of settlement for parties. As cases grow larger, even quite small differences in probability of victory will result in large differences in expected outcomes. Therefore, it will increasingly make sense for parties to bear the costs of trial in order to obtain their optimistic forecast, rather than accept a less favorable settlement. This becomes clearest when we reorganize the conditions for settlement as follows:

$$P_p - P_d \leq \frac{2(T-S)}{J}$$

$$M_p - M_d \leq \frac{2(T-S)}{J}$$

Intuitively, for very large values of potential damages, J, we should only expect settlements to arise when trial costs are unusually severe or else for very small values of $P_p - P_j$ and $M_p - M_d$. Thus, we might wonder whether the burdens behave differently at different scales. In the simulated cases, it is apparent that they do:

Settlement Rate Under Each Rule

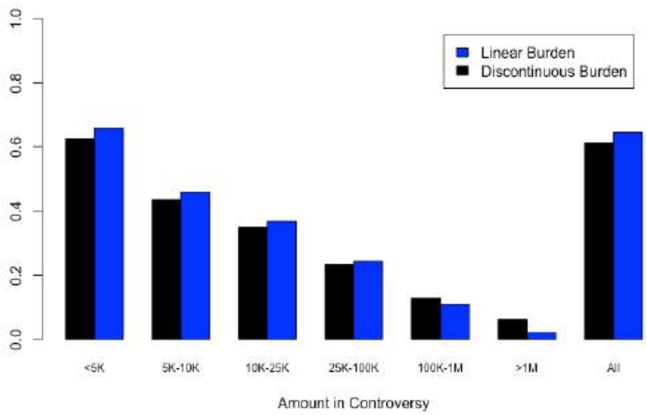

31 For the very hardest cases (0.4-0.6 unbiased predicted mean jury confidence levels), the linear burden settled 63% of the cases, versus only 56% under the traditional rule ($p < 2.2e{-16}$). This 12% increase is almost three times as large as the 4% decrease that is seen for the cases over which the traditional rule shows its advantage.

As can be seen, the general tendency of the continuous burden to favor settlement holds in all but the largest cases. In cases where the potential damages exceed $100,000, the discontinuous burden performs better.[32] This might seem surprising at first, given the previous observation that the discontinuous burden generally results in larger spreads of expected outcomes than the continuous burden does.[33] After all, in very large cases, we should only expect the costs of trial to outweigh the value of getting one's preferred outcome when the two parties' expected outcomes are unusually close together.[34] However, a closer look at the distributions of values of $P_p - P_d$ and $M_p - M_d$ over the entire set of cases helps to resolve the mystery:

32 Settlements are rare in these large cases, occurring in 11% of cases under the traditional rule and only 9% of cases under the discontinuous rule (p = 0.00001). The cases themselves are also rare, representing only 0.6% of the overall universe of cases analyzed here, so the overall systemic impact of this difference on state court systems would be minimal. Nonetheless, this result may have particular relevance for court systems that try an unusual volume of larger cases, such as the federal courts in the United States. *Cf.* 28 U.S.C. § 1332 (2006) (permitting state court cases between parties of diverse citizenship to be tried in federal court if their amounts-in-controversy exceed $75,000). The divergence is even more dramatic in very large cases; in the 0.1% of cases where the predicted trial damages exceed $1M, the discontinuous rule settles three times as many cases (6% versus 2%, p=4e-8) as the linear rule.

33 The mean of the difference $M_p - M_d$ among the whole set of simulated cases was 0.53 (0.77-0.23), while the mean of the difference $P_p - P_d$ was 0.76 (0.88 – 0.12).

34 This model assumes risk neutrality. If parties are risk averse and their degree of risk aversion derives at least partially from expected outcome variance, then we might expect to see more settlements of high-value cases than appear in this model. *Cf.* W. Kip Viscusi, *Product Liability Litigation with Risk Aversion*, 17 J. Legal Stud. 101, 103-05 (1988).

Linear Rule

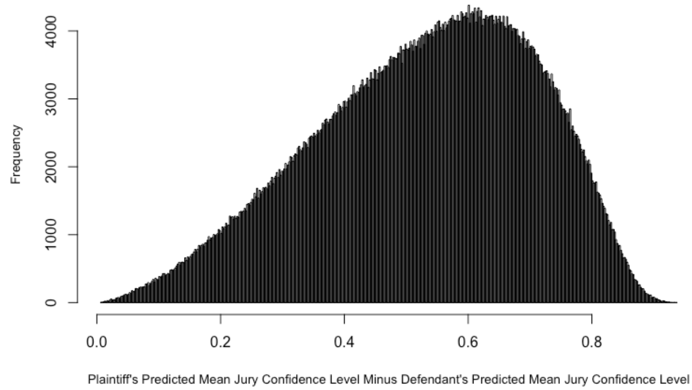

Discontinuous Rule

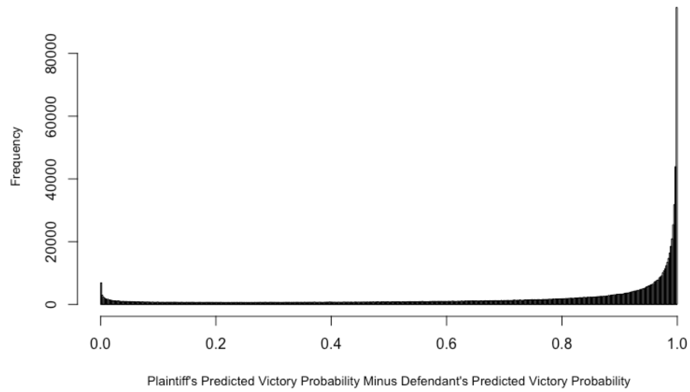

As can be seen, the differences between the parties' predictions of the mean level of jury confidence in liability range widely over the full range from 0 to 1, with moderate values being more likely than extreme values. By contrast, the transformation from confidence levels to dichotomous probabilities of success under a discontinuous burden yields a bimodal distribution of differences with one large mode close to 1 (indicating almost perfect divergence in predicted outcomes), but also a second, smaller mode close to 0 (indicating almost perfect agreement in the outcome at trial). Thus, even though the discontinuous burden results in larger differences in expected outcomes on average, it still includes a much larger fraction of cases where the difference is extremely small.[35] To see why this would occur, consider the following pattern of confidence and victory probability forecasts:

Plaintiff's Forecasts

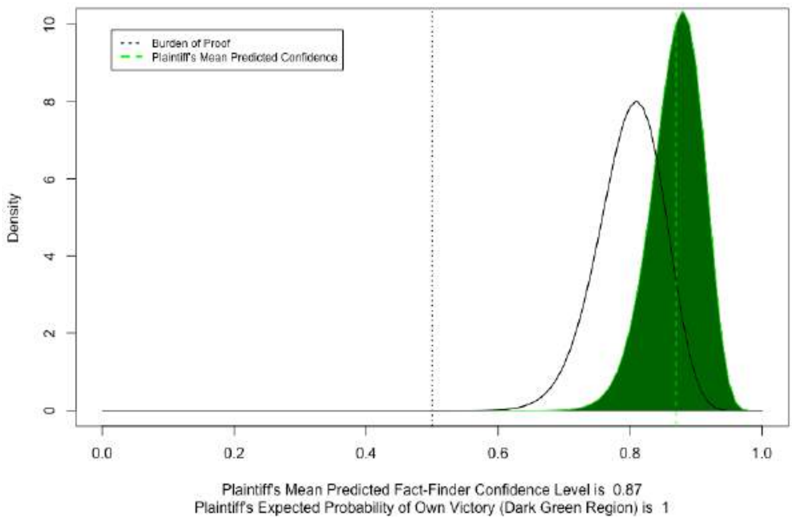

Defendant's Forecasts

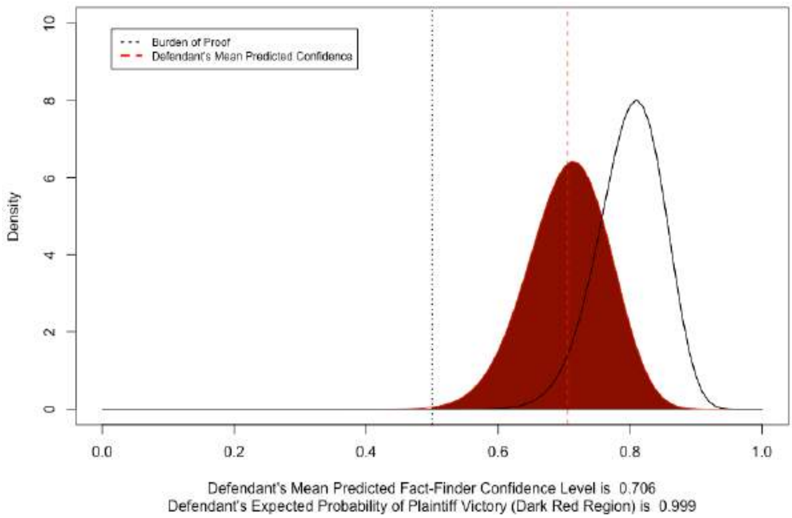

35 1.6% of cases had a divergence in predicted victory probabilities under the traditional rule that were less than 0.01, while only 0.001% of cases had a similarly small divergence between the parties' mean predicted levels of jury confidence.

Here, the unbiased forecast would have predicted factfinder confidence of 0.8, but with relatively high precision, such that there is almost no probability of a juror thinking that the probability was only 0.5 following trial. The parties each make optimistic adjustments to this forecast, with the plaintiff thinking the mean predicted confidence level is 0.87 and the defendant thinking it is 0.706, leading to a significant difference in expected outcomes if we were using a continuous burden of proof. By contrast, the expected probabilities of victory under a dichotomous burden are almost identical, with the plaintiff forecasting essentially certain liability and the defendant thinking that liability will be found at a probability of 0.999. Thus, in cases exhibiting high-precision forecasts towards the extremes of the confidence scale, the discontinuous burden will sometimes prompt the parties to essentially agree on the outcome, even while disagreeing on the most likely confidence levels that jurors might have regarding liability.

And in fact, a closer look at the high-value cases reveals that this is exactly what is going on. If we examine just the pattern of confidence forecasts in the cases with predicted damages amounts above $100,000 that did yield a settlement under the discontinuous burden, we see that very low or very high unbiased mean confidence forecasts were significantly overrepresented, relative to the number of cases with more moderate values.

Unbiased Mean Confidence Forecasts for Large Cases That Settle

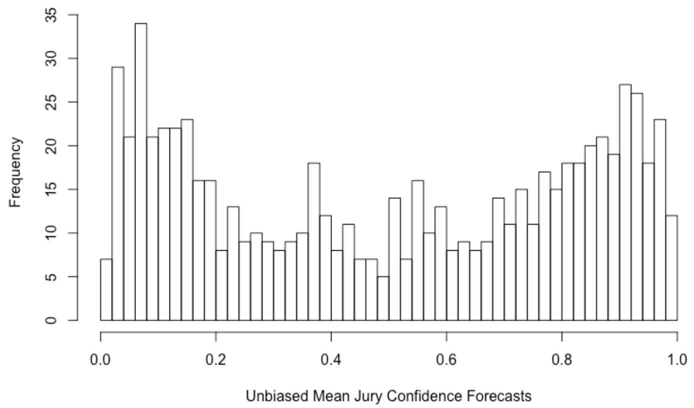

Thus, there is good reason to think that the discontinuous burden is uniquely conducive to settlement when amounts in controversy are particularly large.[36]

36 More precisely, we could say that we should expect this effect to occur in cases where the amount-in-controversy is large, but the litigation costs remain small to moderate. The data make it quite clear that when cases have trial costs

Moreover, although my simulation does not take risk aversion into account, we might expect this tendency to be magnified if we think that defendants will be particularly risk averse when potential losses loom large.[37] After all, under the discontinuous burden parties face a variable risk of either very large damages or nothing, while under the continuous burden they face a risk of certain damages that vary by smaller amounts. If that is right, then defendants should be willing to offer even more generous terms to plaintiffs when negotiating settlements in large cases, which should yield some extra quantity of settlements beyond what the simulation predicts.

Thus, our initial investigation into the settlement-incentivizing properties of standard versus continuous burdens yields three major findings. First, over the general universe of cases the continuous burden seems to incentivize more settlements than we would expect under the discontinuous burden. Second, the continuous burden provides this advantage mainly in moderate-to-hard cases, in which an unbiased prediction of the factfinder's mean expected confidence levels varies between 0.1 and 0.9. The discontinuous burden incentivizes more settlements among the easy cases outside of that range. Finally, the continuous burden's settlement-producing benefits seem to arise mainly in cases involving low to moderate damages amounts. In very large cases (those with an amount in controversy greater than $100,000), the traditional rule incentivizes more settlements than the linear alternative.

III. Impact on Expected Settlement Size and Calibration

This Part will first consider the impact of each rule on the expected size of settlements. Then, we will consider the expected amount of error that resulting settlements would typically have, relative to an ideal world where each deserving plaintiff received exactly what they were owed and each undeserving plaintiff received $0.

A. Impacts on Settlement Size

Recall that the expected settlements (V and V') for each case lie, on average, at the mean of the following ranges, subject to noise induced by variations in parties' bargaining ability and bargaining power:

---

substantially larger than the expected judgment, settlement will almost always occur, regardless of one's choice of burden of proof rule at trial.

37 *Cf.* Daniel Kahneman & Amos Tversky, *Prospect Theory: An Analysis of Decision Under Risk*, 47 Econometrica 263 (1979).

$$P_p J - T + S \leq V \leq P_d J + T - S$$

$$M_p J - T + S \leq V' \leq M_d J + T - S$$

As litigation costs shrink, one straightforward result of this model is that settlement values will tend toward the following limited range as costs shrink down towards zero:

$$If\ T = 0\ \&\ S = 0, then$$

$$P_p J \leq V \leq P_d J$$

$$M_p J \leq V' \leq M_d J$$

Since the parties' optimistic forecasts will randomly vary around the fixed bound of an unbiased forecast of either predicted discontinuous victory probability or predicted mean jury confidence level in liability, we would naturally expect that the expected settlements might closely track the expected outcome of each burden at trial, assuming that litigation costs are small in relation to damages. Analysis of the simulated cases reveals that such effects do in fact occur. The first graph shows the expected settlement outcomes under either trial burden in the subset of cases where trial costs are less than 0.1 of the expected damages if the case were tried.

Discontinued v. Linear Burdens – Predicted Settlements (Trial Costs < 10% of AIC)

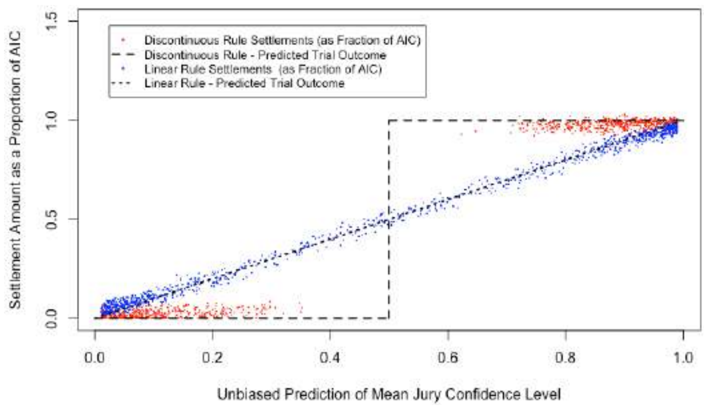

Thus, in a low-cost world we would expect settlement amounts to closely track expected outcomes at trial, and thereby to mirror whatever shape our burden of proof rule takes.[38] Sadly, we do not live in such a world. According to data collected by the National Center for State Courts, the median verdict for a prevailing plaintiff in a bench trial case (by far the most common kind, representing 98.4% of the dataset) was $1,131, and the 75% result was only $2,028.[39] Even restricting our attention to the rare cases that were tried by a jury, the median outcome for a successful plaintiff is only $31,097 in damages.[40] It is harder to find representative data for typical total litigation costs, but an NCSC survey of attorneys practicing in Utah revealed that a typical debt collection case (the single most common case-type)[41] had a median overall litigation cost value of $2,968 per side, with attorneys estimating that 14% of their efforts were devoted to trial work in the median case, for an estimated $416 in typical trial costs.[42] This represented the lowest estimated litigation cost for any case type in the Utah report, but it still results in trial costs that are more than a third of the typical actual damages amount awarded following a trial in a contract case.[43]

That understates the extent of the problem, however. Many categories of cases have litigation costs that on average greatly exceed the amount-in-controversy if the case is taken to trial. Consider automobile tort cases. These represent the plurality of tort claims[44] while having smaller amounts in controversy than many other tort claims, with one recent study reporting a median damages award at trial of just $15,000 in jury trials and $17,000 in bench trials (as compared with medians for tort claims overall of $24,000 and $21,000, respectively).[45] These values are absolutely dwarfed by the typical

---

38 Given the inherent nature of the settlement model, we would also expect settlements to be rare in such a world, as there would be fewer cases where the parties predicted outcomes would differ by more than 2(T-S)/J. Indeed, in this dataset only 6% of these low-cost cases settled using the discontinuous burden (and even fewer settled under the linear burden, for the reasons explored at the end of the preceding Part).

39 Landscape Report, *supra* note 1, at 27.

40 *Id*.

41 In the Landscape Report, contract actions made up 64% of the overall dataset, followed by small claims actions (16%), "other civil" actions (9%) and then tort claims (7%). Debt collections made up the plurality of contract actions (37%), followed by landlord/tenant claims (29%) and then foreclosure actions (17%). *Id.* at 18-19.

42 Utah Report, *supra* note 23, at 82.

43 Landscape Report, *supra* note 1, at 27.

44 They made up 40% of tort cases filed in the Landscape dataset. *Id.* at 19.

45 Thomas H. Cohen, *Civil Justice Survey of State Courts: Tort Bench and Jury Trials in State Courts*, 2005, Bureau Just. Stat. 1, 5 (2009).

cost of fully litigating these cases, however. The Utah study estimated a median per-side litigation cost of $46,375, with 42% of attorney effort being devoted to trial work in the typical case. Thus, even though few would describe automobile torts as the apex of legal complexity, each side typically foresees trial costs in excess of the likely damages award. To see how this changes our expected settlement size under each burden, consider the following graph, which now includes all cases with expected trial costs that are between half and double the expected damages should the plaintiff prevail at trial.

Predicted Settlements (Trial Costs Ranging from .5*AIC to 2*AIC)

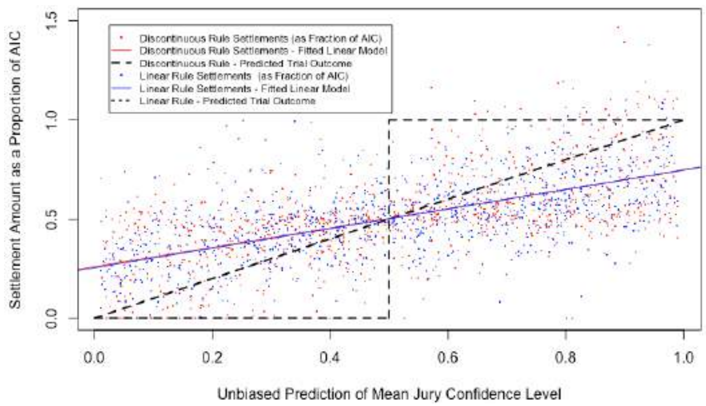

Notice that as trial costs rise to a level that is comparable to the stakes of the case, there is a dramatic increase in the noisiness of resulting settlement amounts. The resulting pattern of discontinuous settlements ceases to be a close approximation to the step-function, and in fact both burdens lead to settlements that are well-modelled by simple linear regressions with extremely similar parameters.[46] Moreover, the expected settlement outcome of *both* rules more closely approximates the result that we would expect at trial under a continuous burden, while neither leads to a pattern of results that is anywhere near a step-function.

Lastly, we must consider what happens when we also include cases with very high costs. Recall that damages and amounts-in-controversy are only

46 The best-fit linear regression models for each pattern of settlements in this range had nearly identical parameters. The model for discontinuous settlements was $y = 0.260 + 0.487x$, while the model for linear settlements was $y = 0.255 + 0.493x$. Both models were significant at $p<2e-16$, but there was slightly more variance among the discontinuous burden settlement amounts (sd of .24 vs. .20 for the linear rule settlements), leading to an R-squared of 0.35, versus a value of 0.42 for the settlements seen under the linear burden.

weakly related to one another, with one study estimating log-log correlations of 0.25 between stakes and litigation costs for both plaintiffs and defendants.[47] Faithfully capturing such a spread in the model meant that a significant fraction of cases have trial costs that are not just similar in scale to the stakes of the case, but in fact are *substantially* higher. For instance, if we consider just the subset of cases in the dataset with stakes that are plus or minus $50 from the median value of $1168, we find the following distribution of expected trial costs:

Expected Trial Costs (AIC Ranging from $1119 to $1219)

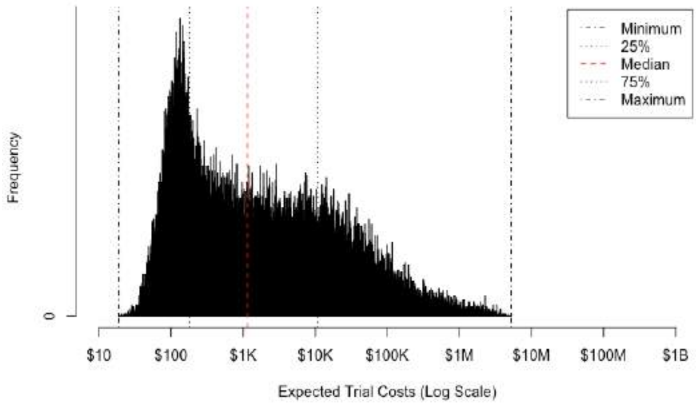

Thus, given what we know about real-world cases, we should expect there to be a substantial fraction of cases that are filed whose trial costs (if realized) would be ten or even a hundred times greater than the damages that a prevailing plaintiff could hope to achieve. This occurs for the simple reason that some misconduct is small in scale but extraordinarily difficult to prove. Once we include these cases in the analysis, we see the following pattern of expected settlements.

47 *See* Emery G. Lee III & Thomas E. Willging, *Litigation Costs in Civil Cases: Multivariate Analysis: Report to the Judicial Conference Advisory Committee on Civil Rules*, Fed. Jud. Ctr. 1, 13 (2010).

Discontinued v. Linear Burdens – Predicted Settlements Ratios (All Cases)

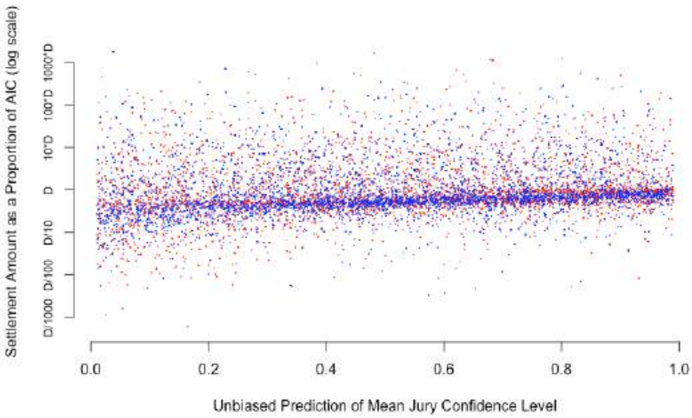

To show the full range of settlement sizes on one plot, a log-scale needed to be used on the y-axis. (Settlements of $0 are omitted from this plot as their logarithms are undefined.) As can be seen, once high-cost cases are included, the bulk of the settlements still lie within one order of magnitude of the amount in controversy, but a substantial fraction of outliers exist. As a result, the expected ratio of settlement amounts to amounts-in-controversy for all cases ends up being quite a bit higher than we saw when only lower-cost cases were included in the analysis.

Discontinued v. Linear Burdens – Predicted Settlements (All Cases)

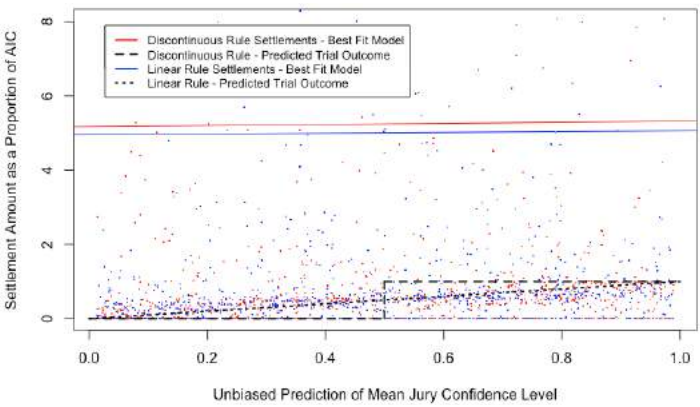

As can be seen above, there are still some cases with trial costs that are much smaller than the stakes of the case, whose settlements generally lie close to the expected trial outcome, as well as many cases with trial costs that are within one order of magnitude of the stakes, which lead to slightly noisier results that approximately mimic the linear trial burden. But there is also a small fraction of cases with trial costs (and therefore bargaining ranges) that are much larger than the size of the disputed damages. To further understand this dynamic, it helps to derive what the settlement bounds would be under that assumption that the quantity (T-S) is greater than the expected judgment amount should the plaintiff prevail:

$$O_p \leq V \leq O_d,$$
$$with\ O_p \leq (P_p - 1)J$$
$$and\ O_d \geq (P_d + 1)J$$
$$O_p' \leq V' \leq O_d',$$
$$with\ O'_p \leq (M_p - 1)J$$
$$and\ O'_d \geq (M_d + 1)J$$

Examining the simplified form that these inequalities take for high-cost cases, a few things become clear. First, in these cases the plaintiff's theoretical lower bound on acceptable offers is below zero, because the trial costs less settlement costs exceed the maximum possible result that they could win at trial. But plaintiffs need not pay to walk away from a case; by simply dismissing their own claims or ceasing to prosecute them, they may avoid any further costs. This is represented in the model above as a "settlement" of zero dollars. Conversely, the maximum upper bound that a defendant should be willing to pay to avoid trial will be larger than the damages at stake in the case. Moreover, since some cases involve very high costs relative to damages, in some cases they might rationally pay a large multiple of damages to avoid trial. Since the results are capped at zero at the lower end but have no intrinsic upper bound, this results in the expected settlement being pulled well above the expected trial outcome in all these cases, under both burdens. And since the mean of the difference $M_p - M_d$ among the whole set of simulated cases was 0.53 (0.77-0.23), while the mean of the difference $P_p - P_d$ was 0.76 (0.88 – 0.12), the upper bounds (and hence the expected settlements) are a bit higher on average using the traditional rule than they would be if we shifted to the linear burden.[48] Thus, even though both rules produce a median settlement that

48 This difference would be mitigated to the extent that other factors work to prevent defendants from rationally agreeing to pay large sums to avoid very expensive trials in low value cases. One such mechanism is the "offer of judgment," which permits the defendant to make an offer of what they believe to be the full

is exactly half the expected damages should the plaintiff prevail at trial, they both produce a *mean* settlement that is much higher, due to the influence of a small number of very high-cost cases where a plaintiff is able to successfully capture an outsized portion of the bargaining surplus. And when these cases are taken into account, the overall result is that both rules lead to settlements that overcompensate plaintiffs on average, with the discontinuous rule doing so to a slightly greater extent.

B. Impacts on the Expected Quantity of Error Per Settlement

Given these variations in the expected sizes of settlements produced by each rule, it seems worthwhile to analyze the impact that either rule would have on the expected quantity of error that arises when a case is settled. For a given settlement *s*, the expected error *e* can be determined as a function of the actual damages *d* and the probability of liability *p*, as follows.[49] First, imagine that the amount of the settlement is smaller than the damages. In that case, the

damages in a case and then penalizes the plaintiff if they reject the offer and recover less than that amount following a trial. Although the federal provision, Fed. R. Civ. P. 68, is fairly toothless, only penalizing the plaintiff through an award of court costs (which are typically a small share of overall expenses), a few states have provisions which create a much stronger incentive to accept such offers. *See e.g.*, Alaska Stat. § 09.30.065 (providing that, if the plaintiff receives a verdict that is less than 95% as large as the offer of judgment, they must pay any reasonable attorney's fees incurred by the defendant after the offer was rejected). Such provisions, if properly utilized, could cap the maximum offer that a rational defendant would accept, and thereby reduce the extent to which settlements in high-cost cases are expected to exceed expected trial outcomes. Other factors, such as a defendant's ability to pay, might similarly constrain the magnitude of these high-cost settlements. Another option is to go into default, *cf.* Fed. R. Civ. P. 55, but this strategy carries large potential downsides if the defendant does not contest the amount of damages, as plaintiffs might obtain significantly larger judgments if their proof of damages is not contested. If the defendant *does* contest the damages amount following a default, the trial costs are only partially avoided. As a result, such strategies (if adopted widely) would still leave room for plaintiffs to obtain some settlements that exceed the true case value through bargaining, and the underlying mechanism for escalated errors under the traditional rule would persist (albeit with smaller average error magnitudes under all rules).

49 In this portion of the discussion, the probability p represents an unbiased epistemic forecast of the likelihood of liability, rather than any single party's biased estimate of that probability.

expected errors are as follows. For each settlement, there is a *p* probability that the plaintiff deserved *d*, so that the error harming the plaintiff is:

$$e_{pl} = (d - s)p$$

Meanwhile for each settlement there is a $1 - p$ chance that the defendant is not liable, in which case the entire settlement was erroneously paid:

$$e_{def} = s(1 - p)$$

So in such cases the total error is the sum of these two terms:

$$e_{total} = (d - s)p + s(1 - p)$$

When the settlement equals the damages, any deserving plaintiffs will have been perfectly compensated, so the total error is simply $s(1 - p)$.

Finally, in cases where the settlement is larger than the true damages, a different dynamic applies. By definition, in such cases all plaintiffs are fully compensated (and more)! But now there is an error affecting the defendant in every case. If the defendant was innocent, the total amount of the settlement *s* is an error. Conversely, if the defendant is liable, just the overpayment, $s - d$, is an error. Thus the total expected error is as follows:

$$e_{total} = (s - d)p + s(1 - p)$$

Happily the above equations can be simplified to yield the following form that will apply in all cases:

$$e_{total} = |(d - s)|\, p + s(1 - p)$$

To clarify the operation of this formula, consider three simple examples involving a case where the plaintiff was injured in the amount of $10,000, but there is only a 2/3 chance that the defendant wrongfully caused those injuries. Let us first imagine that the plaintiff receives a modest settlement of $6,666.67 (which is exactly what an unbiased observer would predict is the expected value of the case to the plaintiff). In such a case, there is a 2/3 chance that the defendant is at fault, so there is likewise a 2/3 chance that the plaintiff has received too little for her injury, by an amount of $3,333.33. At the same time, there is a 1/3 chance that the defendant did not wrongfully injure the plaintiff, in which case the $6,666.67 actually paid to the plaintiff are an error. The total expected error is the sum of these two possibilities, 2/3 * $3,333.33 + 1/3 * $6,666.67 = $2,222.22 + $2,222.22 = $4,444.44.

Next, consider the same case but with a settlement that precisely equals the plaintiff's damages. Since there is a 2/3 chance that the defendant is at fault for the plaintiff's injury, 2/3 of the time this case involves $0 error. 1/3 of the time, however, it involves a $10,000 error, for a slightly lower expected error rate of $3,333.33. The total expected error is therefore $3,333.33.[50]

Lastly, consider the same case but with an especially generous settlement of $15,000. There is a 2/3 chance that the plaintiff was truly injured, in which case the $5,000 premium above the true damages was wrongfully transferred. There is also a 1/3 chance that the whole settlement is an error, in which case the whole $15,000 transfer is an error. The total expected error is now 2/3 * $5,000 + 1/3 * $15,000 = $5,000 + $3,333.33 = $8,333.33.

With the formula for calculating expected errors in hand, we can now analyze the quality of the settlements produced by either rule in a more precise way. When we analyze the expected errors we should expect from each rule, we first find that the conventional burden yields settlements with a mean expected error rate of $4,879, while the linear burden yields a lower mean expected error rate of $4,625 (a 5% decrease).[51] Nor was this merely an artifact of the discontinuous rule's ability to produce more settlements in very large cases (where we should expect higher amounts of absolute error).[52] If we normalize error rates by dividing them by the expected damages should

---

50 In discussions at the Tel Aviv University Conference on Legal Discontinuity, a question was raised as to why the expected error of settlement was not simply the degree of variance of the settlement amount from the total damages multiplied by the probability that the plaintiff would obtain them, as described by the following formula: $|S - p * D|$. This measure would, in effect, treat the trial award given by a linear continuous burden as normative. The first two examples illustrate the problem with that approach. Each case has a true "best" outcome in fact, even if we cannot know it epistemically. Our limited information will not permit us to detect *which* cases involve *which* kind of error, but we can see analytically that settling for the precise expected value of the case should create more overall error than settlements that lie closer to either full damages (for cases where liability is more probable than not) or no damages (for cases where liability is less than 50% likely). *See generally* Kaye, *supra* note 3.

51 This difference in mean expected error rates was significant at p=0.002. The median error rate, by contrast, was increased from $575 to $585 (a rise of 2%).

52 An analysis of those cases which settled under the discontinuous rule but not under the linear rule (which represented just 1.6% of the cases in the dataset) did reveal that the mean expected settlement error was $4,325. This was quite a bit larger than the average quantity of expected error for the cases that only settled using the linear burden, which was just $1,154. There were also far more of the latter type of case, as 4.9% of the simulated cases settled under the linear burden but not under the discontinuous burden.

the plaintiff prevail at trial, we see a similar pattern: the expected error rate of settling cases under the discontinuous burden as a proportion of the amount in controversy per case is 5.19, while the linear rule yields settlements with an expected error of 4.96 times the damages (a 4% decrease).[53]

An examination of the overall dataset revealed that the linear burden's error rate advantage was mostly confined to the large number of typically sized cases in which a prevailing plaintiff would be expected to recover less than $5,000.[54] From there until $100,000 in controversy, the two burdens performed similarly well,[55] and in the tiny fraction of cases above $100,000,

---

53 This difference in mean expected error rates over damages was significant at p=0.005. Once normalized in this manner, the median error ratios were essentially identical, with median expected settlement error ratios of 0.500 for discontinuous settlements and 0.496 for settlements under the linear rule. Some readers may naturally wonder, at this point, if the overall impact of this decrease might be offset by a corresponding increase in the accuracy of trial verdicts under the discontinuous rule. *Cf.* Kaye, *supra* note 3, at 496-501. Sadly, this question cannot be answered in a rigorous way within the scope of this Article's methodology. A substantial portion of litigation outcomes arise by way of pretrial dispositions, such as default judgments and summary judgments. But there is currently no ready way to estimate the impact of the alternative burden forms on these other dispositions, in part because doing so requires a theory of how summary judgment should operate in a world of continuous burdens (or possibly, whether summary judgments would be eliminated in such a world). But for those who wish to get a rough sense of the balance between differing sources of error, the expected error rates at trial under the two rules are .25 times the damages for the traditional rule and .33 for the linear rule, or an increase of 0.08 under the linear rule. *See* Spottswood, *supra* note 2, at Part I-B. That 0.08 increase for tried cases is smaller than the 0.23 decrease in expected error that the linear rule produces for settled cases. Moreover, we would expect settlements to exceed trials by at least an order of magnitude in a real-world court system. As a result, it is quite likely that the linear rule decreases the overall cost of error in litigated cases, although this cannot be shown rigorously until we devise, and then model, entirely new modes of pretrial procedure that take the linear burden of proof rule into account.

54 95% of the simulated cases had less than $5,000 at stake. In this subset, the mean expected error rate, expressed as a proportion of each case's amount-in-controversy, was 5.3 for the settlements that arose under the discontinuous rule and 5.08 for the settlements that arose under the linear rule (p = 0.004).

55 No significant differences in error rates were found in the cases with between $5,000 and $100,000, either as a whole group or when broken down into subsets.

the discontinuous rule produced lower expected errors.[56] Of course, given the loose correlation between stakes and litigation costs in these simulated cases, the overall error rates were also driven down under both rules as the amount in controversy rose. The following chart summarizes these trends:

Expected Errors Under Each Rule

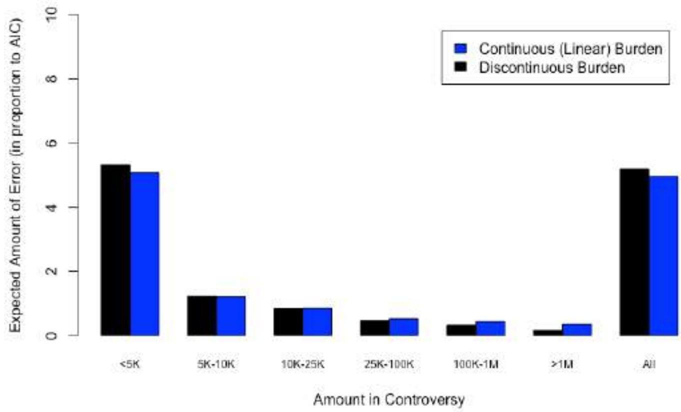

To understand the impact of this change for a system as a whole, two things are worth keeping in mind. First, one should note that the magnitude of the linear rule's settlement error advantage in small cases is greater (+0.23 of AIC) than its disadvantage in very large cases (-0.13 of AIC). Second, one should note that for a typical state court system, there are *far* more small cases than large ones. This is why the mean expected errors overall favor the linear rule, even when expressed in absolute terms (rather than normalized by the amounts at stake in each case).

To understand why this pattern of error advantage arises, recall the earlier discussion of how expected settlements varied based on the ratio of trial costs to the amount-in-controversy. When the stakes are many times greater than the cost of trying a case, recall that both burdens produce settlements that closely track the outcomes they would generate at trial. As has been previously shown by Kaye, the discontinuous burden produces a lowered expected error rate at trial,[57] so in these cases it likewise produces higher-quality settlements. And because of the loose correlation between stakes and litigation costs, we

56 For cases with stakes above $100,000, the settlements produced by the discontinuous rule had an average expected error ratio of 0.30, while the linear rule's settlements produced an expected error ratio of 0.43 ($p<2e-16$).

57 *See* Kaye, *supra* note 3 (exploring expected trial error rates under discontinuous and linear proof burdens).

should expect many more high-stakes cases to have comparatively affordable trial costs.

Next, recall that as the trial costs grow closer to the amount-in-controversy, the discontinuous rule settlements no longer cluster close to either zero damages or full damages, but instead spread out across the intervening space, so that the expected outcome is essentially the same as what we see with a continuous burden, but with higher variance. Given this pattern, it makes sense that the overall pattern of expected errors under either rule is very similar in mid-sized cases, where the stakes and the costs of trial are more likely to be in equipoise.

Lastly, consider the small cases for which the linear rule is advantaged. Given the loose correlation between stakes and costs, we should now expect to find many more cases where the cost of trying the case is quite a bit larger than the amount-in-controversy. This creates a very large bargaining range, which means that some unfortunate defendants may end up paying quite a bit more than the case is worth to walk away. And since the spread between expected outcomes is generally larger under the discontinuous rule, one generally expects slightly larger errors in these cases.

Now, let us consider how expected errors of settlements vary depending on case strength. An initial look suggests a chaotic pattern, with the linear rule generally showing its advantage but no clear pattern emerging:

Expected Errors Under Each Rule, by Case Strength

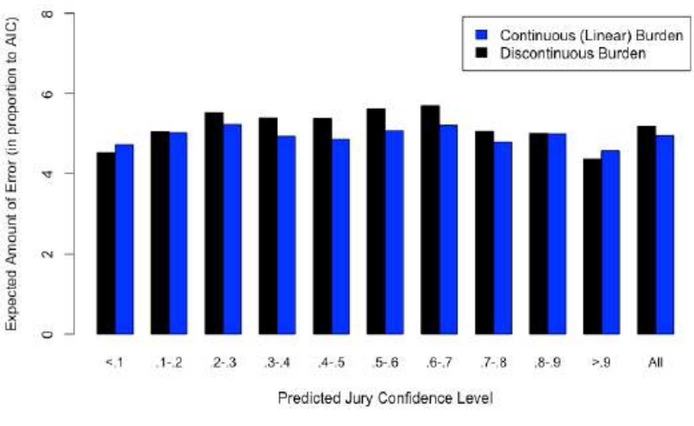

In part, this apparent confusion seems to be driven by the strong effect that a small number of very high-cost cases have on the resulting error-ratio means. If we limit our attention to the vast majority of cases that have trial costs that are less than 100 times greater than the stakes of the case, a much clearer picture emerges:

Expected Errors Under Each Rule, by Case Strength

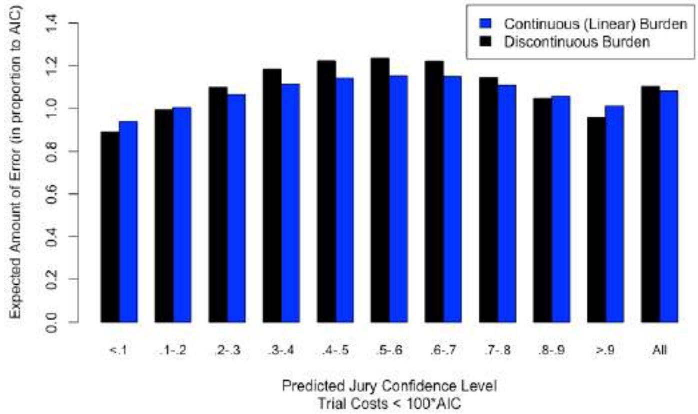

Here, it can be seen quite clearly that we should expect to see more accurate settlements from the linear rule in the middle range of hard cases (meaning the cases where an unbiased observer would forecast typical mean jury confidence levels in liability between 0.2 and 0.8 following a trial).[58] Conversely, we should expect more accurate settlements from the traditional rule in the easy cases (meaning the cases with mean unbiased forecasts below 0.2 or above 0.8), although the possibility that this result was merely due to chance could not be excluded in the overall dataset, only in the subset that excluded the cases with exceptionally high trial costs.[59] In other words, over the terrain

58 For cases lying between 0.2 and 0.8 unbiased predicted mean jury confidence levels, the discontinuous rule produced an average error of 5.45 times the stakes of each case, while the linear rule produced an error ratio of only 5.02 (an 8% decline, with $p = 2e\text{-}5$). In the subset of cases where the trial costs were less than 100 times the stakes, the linear rule's advantage was smaller but still present, with an error rate of 1.12 versus 1.19 under the traditional rule (a 5% decline, with $p<2e\text{-}16$). Thus, although the full spectrum of cases induces more noise into the analysis and introduces more settlement error in general, including the highest-cost cases actually results in a larger benefit for the linear rule.

59 In the subset of cases where trial costs were less than 100 times the stakes, the discontinuous rule produced settlements with an average error of 0.98 times the stakes of each case, while the linear rule had a slightly higher error-over-damages ratio of 1.01, which represents a 3.1% increase ($p=2e\text{-}5$). When the same analysis was done for the dataset as a whole, a similar increase occurred when moving from the traditional rule to the linear one, with the discontinuous rule producing settlements with an error-over-damages ratio of 4.75 and the linear rule resulting in a higher ratio of 4.84. This 1.9% increase was smaller in magnitude than what was seen in the restricted sample, and there is a substantial

where each rule produces an increased number of settlements, those settlements are either more accurate than what the other rule generates on average, or at least produce no measurable decrease in accuracy.

These results follow straightforwardly from the way we calculated expected trial outcomes for cases under each rule. Recall that the traditional rule has an extremizing tendency, typically producing estimates of likelihood of all-or-nothing victory that are closer to 0 or 1, versus the linear burden, for which expected outcomes or more likely to lie in the middle ranges. For cases close to the 0.5 decision threshold, this naturally leads to wider settlement ranges for the discontinuous rule, given that optimism bias will often pull such cases to either side of the threshold, at which point the extremizing tendency will produce a larger spread between each parties' expected outcome. But for cases in which unbiased forecasts of jury confidence lie closer to the extremes, the tendency to pull expected outcomes closer to full damages or nothing can result in both parties agreeing more strongly in the outcome, and thus having a narrower bargaining range.

In summary, we have seen that the types of settlements produced by each rule closely track the results we can expect at trial when trial costs are sufficiently low in relation to the stakes of the case. As a result, when damages are substantially greater than costs, the traditional rule also yields a lower amount of expected error per settlement. In cases with trial costs that are close in size to the stakes of the case, both rules yield a very similar pattern of settlements, although the traditional rule gives settlements with slightly higher variance in this range. For these kinds of cases, there appears to be no difference in expected error rates for settlements produced by either rule. Unfortunately, a small number of cases with trial costs much greater than the expected damages can yield cases with very high (and very erroneous settlements). The linear rule moderates this tendency somewhat, and thus yields a smaller expected error rate overall. Finally, the rules also behave differently depending on the size of each case, with the traditional rule yielding fewer expected errors when cases are easy to decide, while the linear rule reduces expected error rates for the harder cases in between.

---

possibility that it was merely an artifact of random sampling (p=0.48). Restricting the analysis to the easiest cases (below 0.1 or above 0.9) also did not produce a significant result, despite a larger difference in means (4.45 vs. 4.65).

IV. Settlements Under a Logistic Burden of Proof

The foregoing analysis suggests that the policy choice between a linear continuous burden and a traditional discontinuous burden may partly depend on our preferences regarding the desirability of settling cases versus letting them go to trial. On one hand, settlements allow the parties to conserve resources and achieve an option that they would prefer to trial. On the other hand, society as a whole might prefer to have more cases decided on the merits, as a means of publicizing dangerous wrongdoing and encouraging the development of the law over time. Moreover, although most cases should settle for amounts that are, in expectation, quite close to what a continuous burden would yield at trial, the variable influences of party optimism and bargaining power create a substantial amount of variance around this central tendency, and in cases with very high trial costs the expected error rate shoots up dramatically. Accordingly, a shift towards more settlements comes at a general cost to the system's accuracy, by raising our expected error rate well above what either burden would produce at trial. Such results might undermine our desire to optimize levels of deterrence or to fairly penalize wrongdoers via the civil litigation system.

A natural question, when one sees such a tradeoff, is whether there is some way of splitting the difference and striking a middle course.[60] In the context of the effect of burdens at trial, I have previously shown that such a middle course does exist, in the form of a *logistic* burden of proof.[61] Such burdens produce judgment amounts J as a function of the jury's determination of a probability of liability p and the damages D, as given by the following formula:

$$J(p) = \mathrm{D} * \left(\frac{A}{1 + e^{-rp+\frac{r}{2}}} + B\right)$$

The constants A, B, and r are shaping constants, which allow us to design a logistic burden that closely approximates either the step-function or the linear burden, as well as design shapes of the burden that can steer middle courses[62] in between them:

60 *See* Spottswood, *supra* note 2, at Part III.

61 *See id.* (defining and explicating the logistic burden of proof in more detail).

62 The moderate steepness graph has values r = r=8.155, A=1.03 and B=-0.07. These parameters are tuned to minimize the expected sum of squared errors, which allows the burden to optimally avoid inflicting large errors on either party. *See generally* Spottswood, *supra* note 2, at Part III.

Logistic Burden of Proof

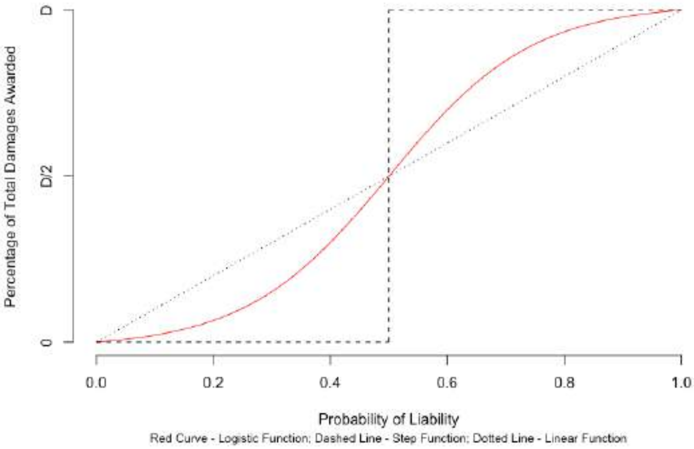

Since this burden takes an intermediate form between linear and step-function burdens, it allows us to strike a balance between the strengths of either rule. This is desirable because the step-function burden has the lowest expected error rate at trial,[63] while the linear burden is advantageous from the standpoint of deterrence,[64] avoids *concentrating* the risk of error excessively on one party,[65] and minimizes the impact of various factors that might affect a jury's decision for illegitimate reasons, such as a litigant's wealth, race, or social status.[66] We might likewise wonder if such an intermediate burden similarly lets us split the difference in terms of the kinds of settlements it incentivizes. Interestingly, however, forms of the burden that seem to strike an intermediate course in terms of trial policy do not also lead to a balanced

63 *See* Kaye, *supra* note 3, at 496-500.

64 *See* Shavell, *supra* note 4 (discussing the over- and under-deterrence produced by the discontinuous rule in tort cases); Rosenberg, *supra* note 6 (same).

65 *See* Spottswood, *supra* note 2, at Part III.

66 *See id.* at Part II-B (explaining and illustrating the ways that the linear rule minimizes these effects). The basic mechanism results from the likelihood that biasing influences move a jury's credence with respect to liability by small amounts, along with the likelihood that cases close to the 0.5 threshold will be disproportionately selected for trial under the traditional rule. This results in a significant proportion of cases in which we might expect the small nudge to change the outcome, leading to a variation due to bias that is the full size of the damages. The linear rule, by contrast, leads a small change in credence to produce only a small change in the size of a damages award, across all levels of case strength. As a result, we should expect fewer cases where a biasing factor leads to a large change in the outcome. *See id.*

approach to settlement incentives. Instead, they incentivize a very similar quantity and quality of settlements as would be seen under the discontinuous burden, so that one can only strike a middle course in pretrial by choosing a logistic burden that is very close to the linear burden of proof.

First, we can extend the model above to describe how settlements should occur using a logistic burden of proof. We will denote the parties' expected outcome at trial using F''_p and F''_d. Under the logistic rule, for a given level of jury confidence p, a plaintiff can expect to receive J_p and so the expected outcome under that burden and the associated values of F''_p and F''_d can be derived as follows:

$$F''_p = \int_0^1 beta(x;\ \alpha_p, \beta_p) J(x) dx$$

$$F''_d = \int_0^1 beta(x;\ \alpha_d, \beta_d) J(x) dx$$

In addition to the parts of the simulation already described, I also computed values for expected logistic burden settlement rates, settlement amounts, and settlement error rates, within the same universe of cases that was previously described. First, I defined the conditions for settlement in terms of the expected outcome under one version of the logistic burden of proof:[67]

$$F''_p - T + S \ \leq F''_d + T - S$$

$$J(p) = \ \mathrm{D} * (\frac{1.03}{1 + e^{-8.155p + \frac{8.155}{2}}} - 0.07)$$

Next I implemented bounds on acceptable offers based on what parties could expect to obtain under the logistic burden of proof, and allowed for varying offers to be accepted within that range:

$$0''_p = F''_p - T + S \ \leq V''$$

$$0''_d = F''_d + T - S \geq V''$$

$$B = Beta(12{,}12)$$

$$V'' = O''_p + B\big(O''_d - O''_p\big)$$

67 The constants are those I used in the graph above, which implement a logistic burden that was designed to minimize the expected sum of squared errors among the parties when used at trial. Running the simulation with other variations of the burden that followed intermediate courses between the burdens' extremes did not dramatically alter the results reported below, however.

The results show that the logistic burden produces results that closely mimic the settlement profile of the traditional step-function rule. The step-function rule and the logistic rule produced a very similar overall number of settlements, with just a modest increase under the logistic rule (612,595 and 620,164, which is 61.2% and 62.0% of all cases, respectively, or just a 1.3% increase).[68] When broken down by amounts-in-controversy, we can see that the logistic burden still retains a modest advantage in settling the smallest cases, but its tendency to induce settlements falls off more quickly than the linear burden, underperforming the traditional rule for all cases above $25,000 at stake.[69] But since the larger cases are comparatively rare, the overall effect is still a modest gain in the rate of settled cases.

Settlement Rate Under Each Rule

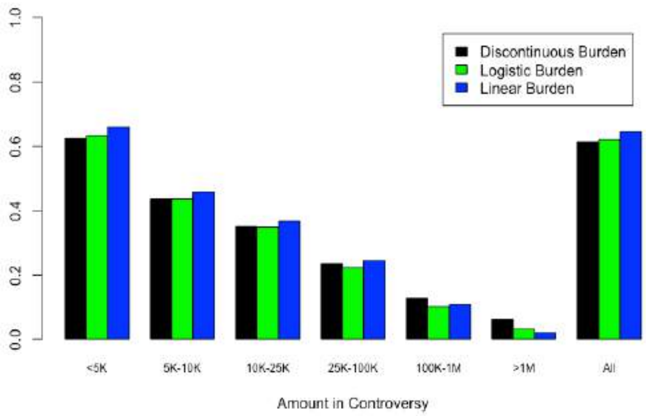

The interaction with case strength is also moderated from what was seen under the linear burden. Recall that in the prior comparison, the linear rule incentivized more settlements between .1 and .9 levels of jury confidence in liability. By contrast, when compared with the logistic rule, the discontinuous burden settles more cases above 0.8 and below 0.2 levels of unbiased predicted jury confidence:

68 The difference in proportions was highly statistically significant ($p < 2e-16$).

69 For the cases above $25,000 at stake, the traditional rule settled 16% while the logistic rule settled only 14.% ($p=6e-5$).

Settlement Rate Under Each Rule, by Case Strength

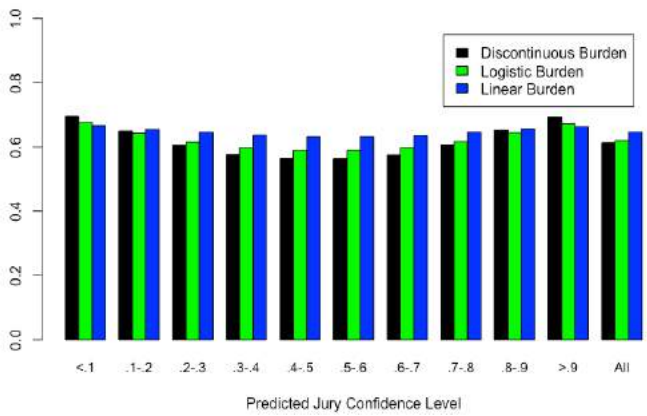

The similar behavior of the two rules in terms of settlement incentives extended into the quality of the settlements as well. The logistic burden yielded a mean expected error of settlement of $4,789, which was just a bit below the $4,879 we expect for settlements under the discontinuous rule. This is less than half of the decrease in expected errors we saw when using the linear burden of proof rule (2% reduction vs. 5% reduction), and we cannot rule out the possibility that the apparent difference arises merely due to random chance in the overall simulation.[70] When we normalize the comparison, by focusing on the ratios of expected errors over true damages amounts for each rule, we see mean expected errors of 5.19 of the damages for the step-function burden and 5.13 of the damages for the logistic burden, but again, this difference is so small that it may merely be an artifact of randomness elsewhere in the simulation.[71] An examination of various subsets showed that the similarity in expected errors was widely shared across amounts in controversy and levels of case strength.

70 The difference in means had a p-value of 0.14, which represents a significant possibility of observing these results due to random chance if we assume that there is in fact no difference between the groups.

71 This difference in means had a very large p-value of 0.53.

Expected Errors Under Each Rule, by Amount in Controversy

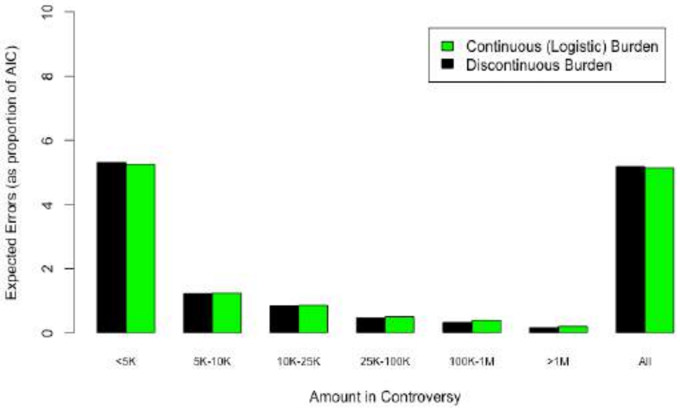

As can be seen, the two rules yield settlements of very similar quality across most levels of amounts in controversy. The only significant difference between the two rules was that the traditional rule showed a slightly lowered expected error ratio in cases over $1M of 0.30 vs. 0.38 for the logistic rule.[72] As before, the picture of error ratios across case strength levels was somewhat obscured by noise due to the effect of a small number of high-cost outlier cases, but if we remove the small fraction of cases in which the trial costs are more than 100 times greater than the amount at stake, we can see a clearer picture:

Expected Errors Under Each Rule, by Case Strength

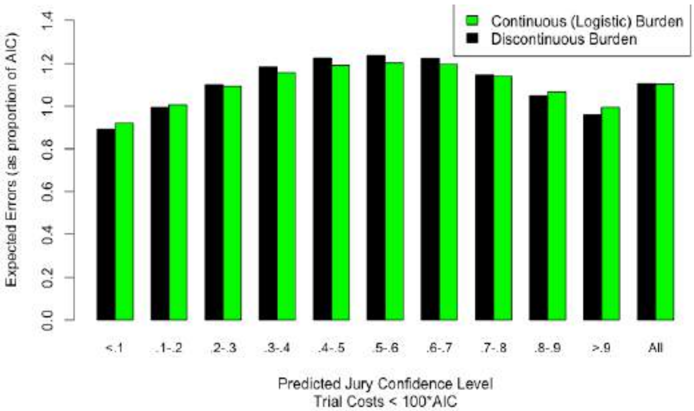

72 This difference in means was statistically significant (p=7e-7), while the apparent differences shown on the chart for the cases below $5,000 and between $25,000 and $100,000 were not.

Within this subset, we see the same familiar pattern of expected error ratios that arose when comparing the linear burden with the traditional one, except that both the advantages and the disadvantages of the logistic rule are more moderate. In the easy cases (unbiased prediction of mean jury confidence levels below 0.2 or above 0.8) the traditional rule yields smaller expected errors, giving an average error of 0.98 of the damages, versus an expected error that equals the damages for the logistic rule.[73] In the hard cases that lie in the remainder of the range, the logistic rule yielded a small benefit, reducing expected errors from 1.18 of the damages to 1.16.[74]

Thus, it appears that a logistic burden that strikes an intermediate approach between the linear and the step-function liability rule at trial does not simply split the difference in terms of its effects on settlement behavior. Instead, the logistic burden produces settlement outcomes that are much closer to what we expect under the step-function rule, with just a 1% increase in settlement rates and no overall change in expected errors for each settlement generated. This may not be enough to make the rule seem worthwhile for those who would find *any* increase in the settlement rate to be intolerable. However, those who find the prospect of more settlements worrisome but also appreciate the benefits of continuous trial burdens may find the logistic burden a useful way of balancing those considerations.

Conclusion

In previous work, I have argued that *trial* outcomes are generally fairer in important ways under continuous burdens than under the discontinuous burden of proof that is conventionally used in our court system.[75] The present Article extends this analysis by considering the impacts that such rules might have on the quantity and quality of settlements, if implemented in jurisdictions that implement the American rule, in which each party pays their own litigation costs. Clearly, more goes into a party's decision to settle a case than can be captured in a simple economic model, but such models do provide useful insight into the incentives that differing rules may create. And as seen above, a transition to a continuous burden alters those incentives, encouraging more settlements (especially in harder cases, or cases with relatively small amounts in controversy), as well as slightly reducing the overall expected error rate for the settlements that do arise. For easy cases or very high-stakes cases, by contrast, the discontinuous burden dominates, producing more (and more

73 This 2% increase was significant (p=0.001).

74 This 2% decrease was also significant (p=0.0003).

75 *See generally* Spottswood, *supra* note 2.

accurate) settlements. Since small cases are much more common than large ones, the linear rule therefore dominates overall by a modest amount, in terms of its ability to incentivize more and more accurate settlements by parties.

More generally, this Article serves to remind scholars and policymakers that changes to the trial process can have complicated impacts on pretrial practice that may be hard to predict, such as the tendency of the continuous burden rules to incentivize more and more accurate settlements. Whether these changes should be counted as a benefit or a detriment for the court system could, of course, be the subject of further debate. Some may find the idea of a continuous trial burden attractive, but not wish to pay the price of increasing the overall rate of settlements. One interesting result that arises from this study is that a logistic burden of proof, which splits the difference between the linear and the step-function rules at trial, preserves the slightly lower settlement rate that we see under the traditional rule. Thus, those who wish that our current system imposed sanctions at trial in a smoother manner, but who wish to prevent what few trials still occur from vanishing, may find the logistic burden particularly attractive.

Finally, it must be stressed that the foregoing analysis is tentative. For tractability, this initial study made several simplifying assumptions that might cause its results to deviate from real-world settlement rates. This leaves open several interesting avenues for future research. It is possible, for instance, that changing the proof burden might impact the amount of discovery that parties seek to take, or change the level of effort they wish to expend at trial. Thus, a future study might model those influences as well, and consider how they might alter the settlement incentives I describe above, which arise solely from differences in expected outcomes under each rule. The current approach likewise analyzes a static universe of potential cases which are designed to track the kinds of civil claims that are currently filed in our legal system. Continuous burdens might also influence the selection of claims for filing, inducing some new claims to be filed that would not have existed before, or deterring others from being filed, and so change the overall distribution of merit, amount-in-controversy, or litigation costs from what I assume to be true in this Article. Such selection effects would not alter the basic incentive structures that the varying rules create for an existing matrix of claims, but if they could be reliably predicted, we might further sharpen our estimates of the impact that continuous burdens would be likely to produce on settlement quantity and quality. And of course, parties' real-world estimation of expected outcomes may be biased or otherwise inaccurate, and if such errors could be modelled that might further enrich our understanding of these matters. The current study, therefore, is offered merely as a first step towards a better understanding of these interesting phenomena.

Appendix

This appendix provides information regarding the design choices that were used to perform the case simulations described above. Full R code that can be used to produce similar datasets, as well as spreadsheets containing the specific data analyzed in this paper, can be found in the following repository: https://github.com/Mark-Spottswood/PDandCS-code

A. Generating Confidence Forecasts

In order to simulate settlement decisions, the model outlined above required several inputs. First, I needed varying levels of true case strength, which would then be estimated by the parties with some optimism bias. As described above, the model generated case strength levels as follows:

$$\mu = unif(0.1, 0.99)$$

$$var = \frac{unif(0.001, \mu)}{10}, if\ \mu \leq 0.5$$

$$var = \frac{unif(0.001, 1 - \mu)}{10}, if\ \mu > 0.5$$

$$\alpha = (\frac{1 - \mu}{var} - \frac{1}{\mu})\mu^2$$

$$\beta = \alpha(\frac{1}{\mu} - 1)$$

$$C_u = beta(\alpha, \beta)$$

In practice, this procedure required some pruning as it produced a small number of cases with unrealistic confidence forecasts. In particular, at very extreme values of μ the forecasts become bimodal, such that the most likely jury confidence level is either extremely low or high, with almost no probability mass between them. This tends to occur whenever a beta distribution is defined with both $\alpha < 1$ and $\beta < 1$, so the dataset of cases was populated in a manner that skipped such cases whenever they were generated. Given the other parameters used to create the simulated cases, this exclusion turned out to be inconsequential, however. In a sample run of 100,000 cycles of this portion of code, no potential cases were omitted due to involving a bimodal confidence forecast.

A more consequential form of pruning occurred to eliminate cases that, in my judgment, involved implausible levels of confidence on the part of an unbiased observer. Lacking data to draw on, I intuitively set the maximum allowable parameters to be α <= 25 and β <= 25, which corresponds to a maximum variance of unbiased prediction of 0.00038, or a standard deviation of prediction of 0.019. This corresponds to a maximum 95% credible interval of prediction of a jury's confidence level in liability of +/- .038, which seemed as much as any person could reasonably say for a real-world case, factoring in variations in juries and their perceptions of evidence. The data in these simulations were therefore pruned to remove such cases, which constituted 22% of what would be generated using the simple parameters above. In order to make it easier for other researchers to explore the impact of varying such confidence levels, the data simulation program includes a user-tunable variable, "forecast precision," which permits any other maximum values of α and β to be substituted for what was used here.

One initially unforeseen consequence of such pruning (although it was obvious in retrospect) was the elimination of some cases with very high and very low mean confidence levels. This arises mathematically due to the inability to specify a beta distribution with a very high mean without making it either bimodal or implausibly narrow in its variance. More intuitively, one cannot both say that one is .999 confident in something while also stating that the same event could easily be .95 at a nonnegligible probability; the first state of confidence excludes the latter. As a result, pruning the very high confidence cases also had the side-effect of removing cases with unbiased mean predicted probability of liability levels above .99 or below .01. More generally, the nature of beta distributions meant that randomly generated parameters were more likely to have low variance as the mean of the distributions approached 0 or 1. As a result, the unbiased confidence forecasts produced by the simulations were distributed as follows once the pruning of overconfident predictions occurred:

Distribution of Mu Values After Pruning

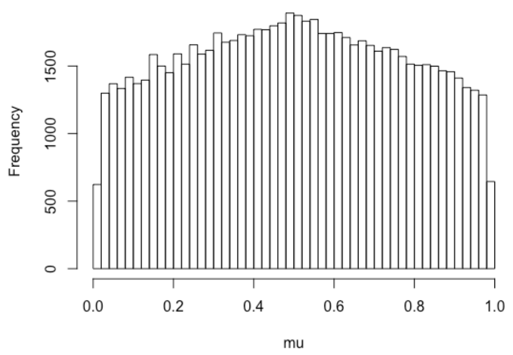

A careful reader of this Article should therefore note that the specific reported fractions of cases settled under each rule depend, in part, on these assumptions regarding the implausibility of extremely precise unbiased confidence forecasts of jury verdicts.[76] One might flatten the above distribution by discarding this assumption, and this would reduce (but not eliminate) the continuous burden's tendency to settle a greater share of cases, due to the increase of cases in the extreme ranges where the traditional burden is better at incentivizing settlements.

B. Simulating Amounts in Controversy

Data on the typical amounts-in-controversy in a representative sampling of state court civil cases are surprisingly scarce, especially given the fact that such cases are by far the most common form of dispute handled by our court systems. Many sources focus only on tort cases, which represent a small minority of overall claims and involve unusually high stakes.[77] Moreover, it is routine to exclude small claims cases, even though they represent a substantial fraction of real-world disputes,[78] just as it is rare to include cases brought in courts of limited jurisdiction, even though this represents the bulk of state court civil disputes.[79] Finally, one must be cautious in drawing on data representing aggregate award sizes in datasets including pretrial judgments, as such judgments may simply be embodying the outcome of a negotiated settlement into a consent decree or other official pronouncement.

With these considerations in mind, I determined that the National Center of State Courts' *Landscape of Litigation in State Courts* report was the most representative recent study of state court civil case processing. This study dissected data from ten randomly chosen counties nationwide, including small claims and cases brought in courts of limited jurisdiction. Given that

---

76 In addition to the basic reasoning above, someone who attends to the details of civil procedure will note that extremely low and high probability cases will often be filtered out before a settlement versus trial decision must be made, via mechanisms such as motions to dismiss, motions for judgment on the pleadings, and summary judgment. *See*, *e.g.*, Fed. R. Civ. P. 12, 56.

77 *See* Landscape Report, *supra* note 1, at 17.

78 *See id.* at 29 (showing 110,274 small claims cases out of an overall sample of 820,893 cases studied, representing 13% of the total state-court caseload); *id.* at 26 (showing that small claims make up 19% of state court cases that are tried to the bench).

79 *See id.* at 17 (showing that limited jurisdiction courts handled almost double the number of cases as general jurisdiction courts, at least within the ten counties sampled by the Landscape of State Courts study).

settlement values may underestimate or overestimate amounts in controversy for a particular claim in potentially unpredictable ways, I drew on the study's summary of trial outcomes for prevailing plaintiffs as a way to generate the amounts at stake in my simulated case dataset. The *Landscape* report reported the following figures for typical trial judgments for victorious plaintiffs:

	N	25%	Median	75%	Mean
Bench Trials	11,481	$679	$1,131	$2,028	$6,408
Jury Trials	194	$7,962	$31,097	$201,896	$1,468,554

As can be seen, bench trials usually involve much smaller amounts-in-controversy than jury trials, and both distributions have means well above their 75th percentile value, implying long right tails to the overall distribution. Accordingly, I fit log-normal distributions to the given quantile values for each type of trial,[80] and then populated the dataset with cases drawn from each distribution as follows:

$$X_{stakes} = Unif(1{,}11675)$$

$$If\ X_{stakes} < 195, then\ D \approx Lognormal(\mu = 656855,\ \sigma^2 = 7824328^2)$$

$$If\ X_{stakes} \geq 195, then\ D \approx Lognormal(\mu = 1623,\ \sigma^2 = 1599^2)$$

The resulting distribution mimicked state-court reality by featuring a very large number of cases with between $34 and $10,000 at stake, along with a small percentage of cases with much higher stakes.[81] Given the long right tail, this is best visualized with a log scale on the x-axis:

80 The initial lognormal fit for the bench trial parameters included an implausible number of cases whose stakes were less than even the minimal filing fees for small claims in most states. I therefore set a floor on the states' distribution of $34, which represents the mean value for a small-claims filing fee across the five largest states in the U.S. This was implemented by subtracting 34 from each quartile value before fitting the lognormal distributions, and then adding 34 to each generated value.

81 It should be noted that the best-fit lognormal parameters for the quartiles tended to undershoot the mean values, so that these overall distributions likely include fewer very large cases than we would expect to see in real life. However, since very large cases tend to settle rarely in this model, and since when they do settle the expected error is typically low under any burden of proof, the impact of this limitation in the simulation on the ultimate results is likely quite small.

Expected Damages If Case Is Tried

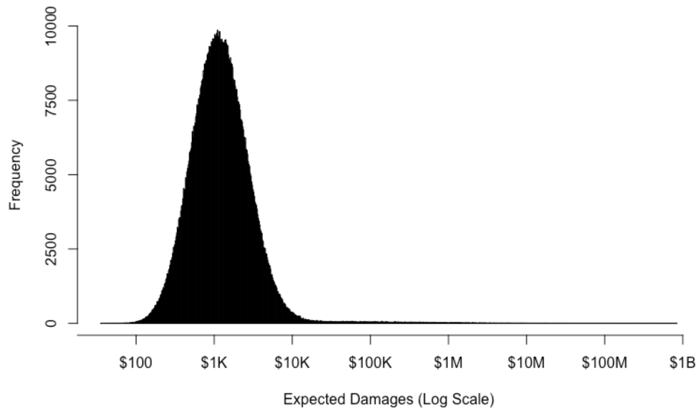

C. Simulating Litigation Cost Parameters

The final piece needed in order to simulate settlement behavior across a realistic universe of cases is data concerning litigation costs. Unfortunately, the Landscape report did not collect information concerning litigation costs across sampled cases. Nor has there been any study that attempted to broadly survey litigation costs over the whole universe of typical state-court cases. Data collection efforts in this vein have typically proceeded by sending surveys to attorneys, and of course such efforts will not permit us to estimate the typical amount of effort expended by self-represented parties. And the best national survey that I was able to locate was almost certainly biased towards the high end of the cost spectrum, with its cheapest cost category being the relatively high-value and high-effort category of automobile tort cases, rather than the sort of simpler and smaller actions that are far more common in state court systems.[82]

In order to better estimate costs in smaller cases, I drew on a recent report prepared for the Utah bar, which included detailed survey results from Utah practitioners concerning typical litigation costs.[83] This report helpfully included smaller claims such as debt cases. It also reported the subfractions of attorney effort devoted to differing tasks throughout the trial process. Unfortunately, in some respects it formed an imperfect match to the Landscape dataset, even beyond being limited to attorneys practicing in a single state. First, it included some categories of cases (such as family law matters) that were excluded from

---

82 *See generally* Paula Hannaford-Agor, Measuring the Cost of Civil Litigation: Findings from a Survey of Trial Lawyers, Voir Dire 22 (2013).

83 *See generally* Utah Report, *supra* note 23, Appendix E.

the Landscape dataset, while failing to provide data for other kinds of claims that represented a large fraction of the Landscape cases, such as landlord-tenant or foreclosure actions.[84] Second, it subdivided other categories of cases in ways that could not easily be matched to the Landscape data.[85] This presented a challenge, as ideally we would wish to simulate appropriate cost amounts that correlate with the frequency of different sorts of cases in the overall dataset. Third, it did not separately break down cost information for plaintiffs and defendants, limiting our ability to realistically model scenarios with substantial cost asymmetries.

Unfortunately, the only kinds of cases for which *both* claim frequency and litigation costs could be clearly identified were debt collection cases, automobile tort cases, and medical malpractice claims.[86] The following chart lists the relative frequencies of each of these case types in state courts, along with information concerning the distributions of per-side litigation costs that were summarized in the Utah Report for cases taken to trial.[87]

	Percentage of Total Cases	**25% of Overall Costs/Side**	**Median of Overall Costs/Side**	**75% of Overall Costs/Side**
Debt Cases	23.7%	$260	$2,698	$14,208
Automobile Tort Claims	2.8%	$19,888	$45,375	$122,163
Medical Malpractice Claims	0.35%	$56,880	$135,950	$333,275

84 *Compare id. with* Landscape Report, *supra* note 1, at 17-19.

85 *Compare* Utah Report, *supra* note 23, at 77-80 (surveying attorneys about "Business/Commercial Litigation" and "Employment Disputes") *with* Landscape Report, *supra* note 1, at 17-19 (failing to indicate what fraction of contract or tort actions might match these labels).

86 The last category was imperfectly matched, with the Landscape Report only identifying a subcategory of medical malpractice claims, while the Utah Report gave cost figures for the more generic category of professional malpractice cases. I decided to apply the Utah figures for the subfraction of medical malpractice claims anyway, given the importance of including some of these unusually high-cost cases in the model.

87 For the case frequency data, see Landscape Report, *supra* note 1, at 17-19. For the litigation cost information, see Utah Report, *supra* note 23, at 71-72, 75-77, 81-82.

In order to generate cost parameters for the whole dataset, I then followed a similar procedure as before. First, I fit log-normal distributions to the quartiles of each of these three categories of case-cost distributions. Second, for each generated case, I randomly selected from one of three cost distributions, with the frequency of sampling from each being adjusted to match their frequency in real-world cases, in order to produce each value of L_{naive}, the expected (unadjusted) per-side litigation costs if a case were taken to trial. I also adjusted the generated parameters to avoid generating cases with implausibly small[88] or large[89] litigation costs.

$$X_{costs} = Unif(1{,}100)$$
$$If\ X_{costs} < 88, then\ L_{naive} \approx Lognormal(\mu = 644{,}875.7,\ \sigma^2 = 88{,}845{,}016^2)$$
$$If\ 88 \leq X_{costs} < 99, then\ L_{naive} \approx Lognormal(\mu = 117{,}887,\ \sigma^2 = 272{,}013.8^2)$$
$$If\ X_{costs} \geq 99, then\ L_{naive} \approx Lognormal(\mu = 323{,}362,\ \sigma^2 = 692{,}337.5^2)$$

The next step was to induce the needed correlation between litigation costs and amounts-in-controversy. Lee and Willging reported a .25 log-log correlation between stakes and per-side litigation costs, for both plaintiffs and defendants, in a recent study of federal cases.[90] Lacking a similar estimation for state-court cases, I used the following adjustment to the overall costs distribution to induce the same relation in my simulated cases:

---

88 The best-fit lognormal parameters for the debt case-cost distribution yielded an implausible number of cases with total costs close to $0, given that in the real world there are certain minimum tasks that must be accomplished to secure a verdict in even the simplest case. I therefore adjusted the distribution to have a minimum of $100. This was meant to stand in for a $34 filing fee plus approximately 8 hours of effort devoted to research, evidence collection, and evidence presentation, if that effort was compensated at the federal minimum wage rate.

89 The high variance of the fitted debt cost distribution introduced a second problem, in that it also included some cases with per-side litigation costs over a billion dollars. These cases were very few in number, but given the strong effect of high-cost outliers on the settlement error rate estimation analysis it was problematic to leave them in. I settled for capping the highest per-side total cost amount at $11,184,989, which was the maximum value that any Fortune 200 company reported paying in legal fees per-case over a five-year period. *See* Lawyers for Civil Justice, *Litigation Cost Survey of Major Companies*, CONF. LITIG. DUKE. SCH. 14 (2010).

90 Lee & Willging, *supra* note 47, at 11, 13.

$$L_{adjusted} = L_{naive} + \frac{unif(0, J - L_{naive})}{2}, if\ J > L_{naive}$$

$$L_{adjusted} = L_{naive},\ if\ J = L_{naive}$$

$$L_{adjusted} = L_{naive} - \frac{unif(0, L_{naive} - J)}{2}, if\ J < L_{naive}$$

This resulted in the following distribution for expected per-side litigation costs for each case, assuming the parties decided to take all cases to trial.

Expected Overall Litigation Costs If Case Is Tried

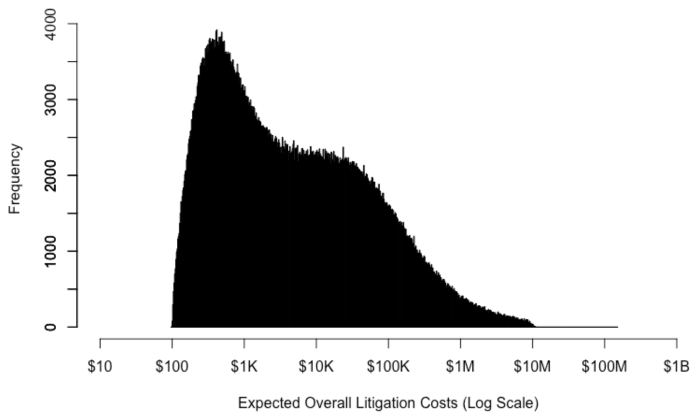

The final step was to generate values for trial costs and settlement costs for each case, based on the overall distribution. I lacked sufficient data to fit reliable models that would allow for a varying relationship between total costs, trial costs, and settlement costs, so I proceeded more simply by taking the mean portion of attorney effort devoted to trial and settlement across all categories of cases included in the Utah report. In that survey, the mean for the fraction of attorney effort devoted to settlement was 6.5% of the total work that would be needed to take the case from filing to a trial on the merits. The mean reported fraction of effort for trials, by contrast, was 37.5%. Since each case no doubt presents some variation in these effort fractions, I then simulated settlement and trial costs for each case as follows:

$$T = 0.375\ L_{adjusted} + \mathcal{N}(0, (\frac{0.375L_{adjusted}}{5})^2)$$

$$S = 0.065L_{adjusted} + \mathcal{N}(0, (\frac{.65L_{adjusted}}{5})^2)$$

The resulting relationship between settlement and trial costs is as follows:

Simulated Settlements and Trial Costs, by Total Costs

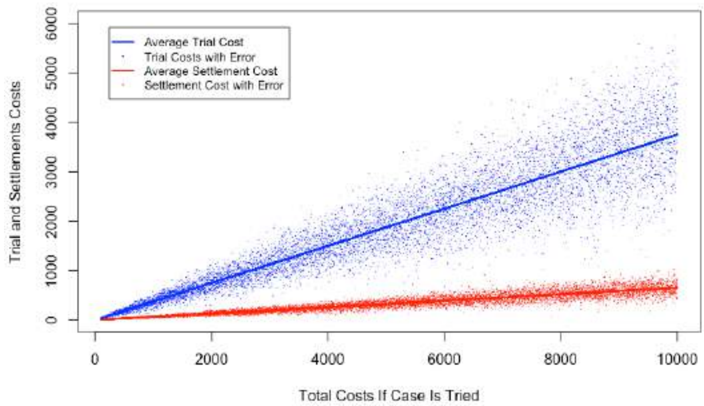

And the resulting overall distributions of trial and settlement costs mirror the overall cost distribution, just shifted by varying degrees to the left.

Expected Settlement Costs

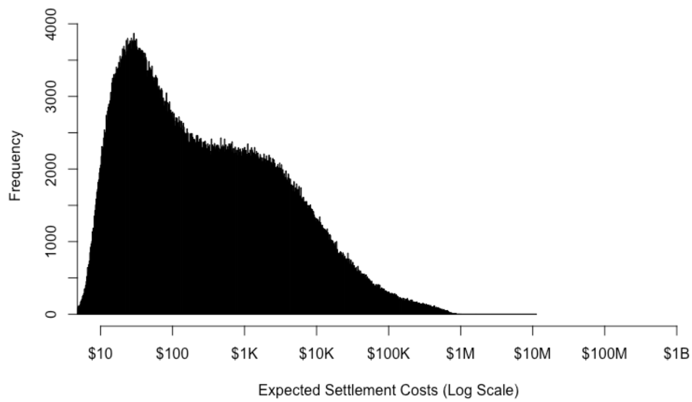

Expected Trial Costs

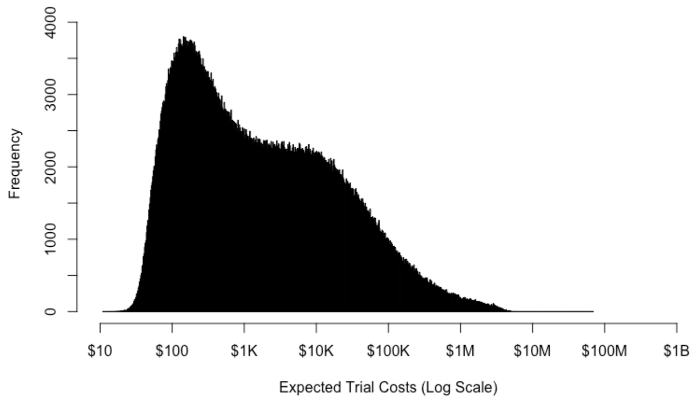

This final histogram shows the ratio of trial costs divided by stakes across the entire dataset.

Expected Trial Costs Over Damages

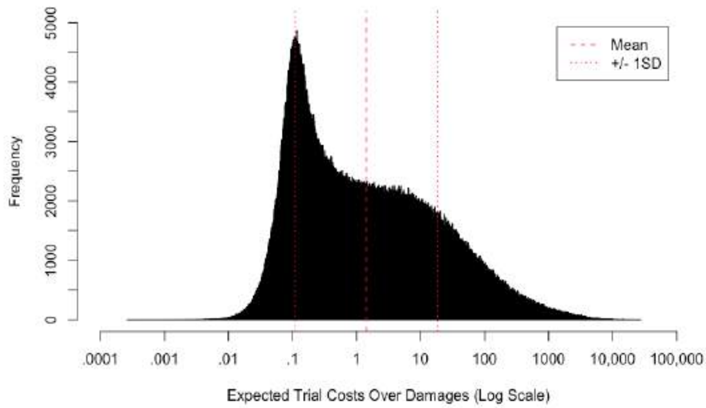

As can be seen, trial costs and damages are very close in the mean case, with a pronounced mode of cases where trial costs are about 0.1 of the stakes of the case, and a long right tail in which trial costs are multiple orders of magnitude larger than the damages.