



---

GW Law Faculty Publications & Other Works

Faculty Scholarship

---

2020

## Modeling Settlement Bargaining with Algorithmic Game Theory

Michael B. Abramowicz

Follow this and additional works at: [https://scholarship.law.gwu.edu/faculty\\_publications](https://scholarship.law.gwu.edu/faculty_publications)

 Part of the [Law Commons](#)

---

# Modeling Settlement Bargaining with Algorithmic Game Theory

Michael Abramowicz  
abramowicz@law.gwu.edu  
George Washington University Law School

October 12, 2020

## Abstract

Past computational models of settlement bargaining have lacked explicit game theoretic foundations. Algorithmic game theory, however, offers techniques that can find perfect Bayesian equilibria even where closed-form mathematical solutions may be intractable. Some recent mathematical models tackle two-sided asymmetric information, including evidentiary signals models, in which the judgment is a sum of both shared and independent private information, and correlated signals models, in which both parties receive noisy signals about the same information. To relax assumptions inherent in these models, this paper employs several progressively more complicated techniques, including iterative elimination of dominated alternatives, no regret learning, and counterfactual regret minimization. Although these algorithms are not guaranteed to produce Nash equilibria in general-sum games like litigation, they nonetheless succeed in producing either exact or close approximate equilibria on discrete versions of the corresponding mathematical models. A single algorithmic game theory model can incorporate a number of features that state-of-the-art mathematical models cannot handle simultaneously, such as two-sided correlated signals of both liability and damages, risk aversion, and options to concede.

## 1 Introduction

Builders of settlement bargaining models feel about computer simulations the same way that builders at construction sites feel about duct tape. Like duct tape, a computer simulation can be used to plug an occasional hole, but only temporarily and only with a little bit of embarrassment. With modern developments in algorithmic game theory, however, model builders no longer need to apologize for using simulations. Mathematical models that produce clean, exact, closed-form solutions will always have a beauty that computational models lack. But computer simulations have a corresponding advantage, their ability

in many cases to engage far richer modeling environments. In the physics literature, some specialties remain the province of mathematics (e.g., quantum mechanics), others are dominated by simulation (e.g., hydrodynamics), and in still others, the two modeling approaches exist side by side (e.g., traffic). This paper's modest suggestion is that the settlement bargaining literature should move to this last state of happy symbiosis. Computational modeling should not and will not replace mathematical modeling, but the combination of the two can advance understanding better than either alone.

Historically, the principal problem with computational models has been that the players' expectations about trial and their strategies in settlement negotiations are defined exogenously. The mathematical literature on settlement bargaining, on the other hand, seeks to *derive* the strategies that rational players will play. The output of successful modern mathematical models is often a perfect Bayesian Nash equilibrium. In such an equilibrium, the parties' beliefs and strategies are in sync. That is, each player's strategies, on issues such as what settlement offers to make, are rational given the player's beliefs about the other party's information and strategy; moreover, the game play given the strategies confirms that the players' beliefs are correct. In contrast, a computational model in the original Priest and Klein (1984) [19] model of selection effects in litigation exogenously sets litigants' beliefs, assuming that litigants will settle cases when and only when the litigants are mutually optimistic and the degree of their mutual optimism exceeds the sum of the difference between their trial and settlement costs. Some computational approaches, such as Ryll (1996) [21], derive strategies based on learning rules, but without constraining the learning rules to Bayesian rationality. Others, such as Hylton (2002) [12], first derive models mathematically and then simulate results given the formulas proven. If a computer simulation must be told what the equilibrium strategies are, it is useful only after the mathematical model has done the real work, or as a cheat when an equilibrium cannot be found but we have a hypothesis as to what it might be.

Algorithmic game theory differs from computational models traditional in the settlement bargaining literature in that it includes various techniques that seek to identify Bayesian Nash equilibria. A caveat is that these techniques seek to identify only approximate equilibria rather than exact equilibria. In zero-sum games like poker, some algorithms have convergence guarantees, so that with enough algorithm iterations, one can obtain an arbitrarily close equilibrium, such as one where no party could gain more than a dollar by deviating from that party's approximate equilibrium strategy. Unfortunately, litigation is a general-sum game; one party's losses are not the other party's gain if the losses are due to legal expenses. The algorithms give poor convergence guarantees for general-sum games. But all is not lost, and indeed we will see that this caveat does not necessarily present insurmountable problems for modeling settlement bargaining. Algorithmic game theory techniques can be used to determine just how good an approximate equilibrium is. At least in the games modeled here, we will see that the game is considerably closer to equilibrium than one could expect any human to play, even though the algorithms do not

guarantee this outcome for all general-sum games. Indeed, even without any guarantees of achieving convergence, the algorithms are often successful in finding demonstrably *exact solutions*, differing from perfection only by miniscule rounding errors. Although the steps to arrive at those solutions are different from the steps of a mathematical model, the results are no less Bayesian.

An alternative critique of computational models is that they are not necessary. The relevant features of litigation are not all that complicated, and thus mathematics can do the work without any help, thank you very much. This argument might appeal to someone unfamiliar with the literature, because one can devise simple rules about how litigants obtain information about the merits, and one might imagine that it would be easy to solve for the parties' strategies. But even with very simple models, identification of equilibrium strategies may be quite difficult. The economist who has toiled to solve settlement bargaining models and has ambitious dreams of extending such models knows that modeling settlement bargaining in full fidelity to Bayes is challenging. At the heart of the challenge is that litigation is an asymmetric information game. A plaintiff and a defendant may have different information, or they might have different beliefs about the same information. Even if the distributions from which the information is drawn are known, solving for an equilibrium strategy may be dauntingly difficult. This is especially so if the model is to incorporate some features that seem important, such as two-sided asymmetric information, risk aversion, parties' options to concede rather than fight, and the possibility of uncertainty about both liability and damages.

Still, if any computational model could be solved with equal ease mathematically, then the computational approach would be clearly inferior. We will see, however, that this is not the case. There may well be some problems that are easier to solve with a mathematical approach, and we will point out some types of mathematical modeling that algorithmic game theory cannot (yet) handle. But we will see that computational models can easily mimic some of the most demanding mathematical models of settlement bargaining, and moreover that algorithms can improve on those models by relaxing key assumptions. Many changes that would require extensive reworking of mathematical models can be effected in a computational model with a few short lines of code at no cost of algorithmic complexity. Other changes can be achieved but at a price of a slower algorithm. The computational approaches documented here cannot incorporate an arbitrary number of seemingly relevant features, but often, they can juggle more variables than mathematical models. At least sometimes, this gain may be sufficient to make them a valuable complement to closed-form mathematical models with smaller numbers of variables.

When a computational model incorporates all the features of a mathematical model plus some additional useful features that are not easily added mathematically, the results provide a strong indication of what a hypothetical mathematical model with those features would show. The results must still be treated with some caution for two reasons. First, it is possible that the discreteness of computational models may produce differences from models using continuous variables. That danger can be reduced by allowing for more strategies

(for example, allowing for more discrete signals or more discrete settlement offers). We will also see that computational models sometimes allow for *less* discreteness than some mathematical models, for example converting a binary signal in an existing model into a range of possible signals. Second, just because some pattern occurs in selected simulations with particular parameters does not guarantee that it would occur in any broader universe of simulations with other parameters. Still, simulations can be run with many different sets of parameters, and if they include more features than a corresponding mathematical model, they expand the number of dimensions of a problem that can be considered. Running various simulations produces data that can be analyzed with empirical techniques. At the very least, computational models can guide paths for future research. Surely researchers would consult an oracle that could tell them the results they might achieve if successful in solving for a unique equilibrium in a mathematical approach. Sometimes computational models can serve this function, for example by showing researchers that relaxing a seemingly innocent assumption will change fundamental results or that relaxing a seemingly significant assumption will not do so. The oracle might not be able to answer everything or might sometimes respond with less precise answers than one would like, but it is still worth consulting.

The literature on settlement bargaining is vast and contains many useful features that we will not cover here. The focus will be on three recent cutting edge models of settlement bargaining with two-sided asymmetric information. The first, Friedman and Wittman (2006) [8], provides a foundation for Dari-Mattiacci and Saraceno (2019) [6], who extend the model to allow for fee-shifting, and Klerman, Lee, and Liu (2018) [13], who focus on rigorously analyzing the selection of disputes for litigation. Each of the papers includes lengthy rigorous proofs that substantially improve understanding. Each also makes a number of simplifying assumptions, and this paper will use computational models to analyze how relaxing these assumptions affects the results. For each paper, we will use a different technique of algorithmic game theory, beginning with the elementary and familiar method of iterative elimination of dominated strategies, used to model linear strategies; continuing with no regret learning, which enables the parties to use nonlinear strategies by separately optimizing at different information sets; and culminating in counterfactual regret minimization, which rigorously allows the optimization at each information set to account for the distribution of cases that will reach that information set given other parties' current strategies. Counterfactual regret minimization is the algorithm that revolutionized computer performance in the game of poker, an asymmetric information game with an information structure that has much in common with litigation.

Part 2 will introduce each of the three mathematical models, and Part 3 will elaborate on computational extensions for each one. Part 4 concludes. The source code for all simulations reported is available at <https://github.com/mbabramo/ACESim4/tree/CRM>.

## 2 Mathematically Modeling Two-Sided Asymmetric Information

First-generation settlement bargaining models were non-Bayesian. The canonical articles of Landes (1971) [15], Posner (1973) [18], and Gould (1973) [10] assumed that cases would settle when, taking the parties' divergent expectations of the result of litigation as givens, settlement would produce a social surplus. This assumption may seem natural given bargaining in the style of Coase (1960) [5], but Myerson and Satterthwaite (1983) [16] later showed that when parties' valuations differ, bargaining cannot be guaranteed to produce efficient results absent subsidies. Moreover, models in this tradition, including Priest and Klein (1984), beg the question of how the parties formed their expectations.

### 2.1 One-Sided Information Models

Most of the Bayesian mathematical literature on settlement bargaining incorporates models of one-sided asymmetric information. One party knows the strength of the case, and the other party knows only the distribution of outcomes. Sometimes, the plaintiff is the party with the information, while at other times the defendant has the information. Sometimes, the information concerns damages (with liability conceded), while at other times the information concerns the probability of liability (with damages uncontested). Sometimes, the party with the information makes an offer, thus potentially signaling the party's type (i.e., conveying credible information about the information that the party possesses), while at other times the party without the information makes the offer, thus potentially screening the more knowledgeable party based on its type (i.e., leading the more knowledgeable party to take some action consistent with its information).

For example, Bebchuk (1984) [1] assumes that the defendant knows the probability that liability will be imposed, while the plaintiff knows only the cumulative distribution of the probability of liability across all cases. Because the cumulative distribution determines the expected total payout, the model is general enough to apply equally to uncertainty about the level of damages or even uncertainty about both liability and damages. The unknowledgeable plaintiff makes a take-it-or-leave-it offer to the knowledgeable defendant. The model is thus a screening model. The equilibrium derived is a partial pooling equilibrium, in which the defendant types with relatively strong defenses reject offers, and the defendant types with relatively weak defenses accept offers. Meanwhile, Reinganum and Wilde (1986) [20], while still imagining the plaintiff making an offer, assume that only the plaintiff is endowed with the relevant information. Over some range of plaintiff offers, the defendant plays a mixed strategy, with some positive probability less than 1 of accepting the offer. Below the range, the defendant accepts all offers; above the range, the defendant rejects all.

Early work on two-sided asymmetric information models endows the litigants with different types of information. Schweizer (1985) [22] features two types of

plaintiff and two types of defendant. The plaintiff receives either good or bad news about the probability of liability, and the defendant receives either good or bad news about the level of damages. The defendant proposes settlement terms that the plaintiff must accept or reject. The model produces multiple equilibria, including both separating and pooling, though at least one criterion for refining equilibria narrows the result down to a single outcome. Daughety and Reinganum (1994) [7] allow either party to propose a settlement and reverse the information structure, with the plaintiff having information about damages and the defendant, about liability. The model allows for continuous types, and information about damages and liability can be combined into a single variable representing litigation strength. They derive equilibria, showing that the rate of trial may differ depending on which party proposes a settlement, with the rate depending on parameter values.

Any model will have unrealistic assumptions, and the fundamental question is whether those assumptions make a difference. The problem, of course, is that one can resolve such questions only with intuition or with new models. One view of the status quo is that models in which expectations are endogenous may be no better than the first-generation divergent expectations models in which they are exogenous. Gelbach (2018) [9] offers a reduced form model that shows that the approaches may not be all that different. Gelbach shows that even the Landes-Posner-Gould model can be constructed with rational Bayesian expectations, if cost parameters are allowed to vary across cases instead of being fixed. From this perspective, perhaps the sacrifices needed to achieve Bayesian rationality are too great. Certainly, many issues of interest can be fruitfully modeled without common priors. Spier and Prescott (2019) [23], for example, model contracting between litigants and with third parties on the outcome of the litigation, a task that would be considerably more difficult in a fully Bayesian model.

Still, it may also be useful to generate more realistic Bayesian models. The challenge that we will focus on here is two-sided asymmetric information in which both parties have information of the same sort, such as information on liability or information on damages (or, as our computational models will tackle later, independent information on each). The vast majority of litigated cases, after all, likely involve differences in the information that the parties have (or in beliefs about such information), and it is not a priori obvious whether the conclusions of the models discussed above extend to the case in which both parties have estimates of the same sort. Of course, depending on the topic of interest, a model must include additional features, and indeed the intellectual history to be relayed below shows one model, the Friedman and Wittman (2006) [8] model, extended in very different directions. Dari-Mattiacci and Saraceno (2019) [6] extend the model to study various fee-shifting rules, while Klerman, Lee, and Liu (2018) [13] offer different extensions focused on placing the Priest-Klein model on a Bayesian foundation. This section reviews these models to identify their commonalities, differences, and assumptions that computational models may be able to relax. After reviewing the models, we will use a separate technique to enhance each model computationally.

## 2.2 Friedman and Wittman's Averaged Signals Model

In the one-sided information models, the structure of bargaining often affects which party receives most of the surplus. Friedman and Wittman avoid this problem by adopting the bargaining protocol of Chatterjee and Samuelson (1983) [4]. In Chatterjee-Samuelson bargaining, the plaintiff and defendant simultaneously submit offers. If the plaintiff's exceeds the defendant's, the case definitively settles at the midpoint; otherwise, bargaining has failed. Costs of trial are borne only in the event of bargaining failure. Friedman and Wittman justify this choice not on the ground that the protocol is commonly used (it is not), but on the ground that it provides a useful reduced form of a more complicated bargaining process. In contrast to divergent expectations models, with Chatterjee-Samuelson bargaining, a case may go to trial even though there is a social surplus from settlement given the parties' expectations. The reason is that the parties may shade their offers in the hope of deriving a larger portion of the settlement surplus, even at the risk of bargaining failure.

The informational structure is arrestingly simple. The plaintiff observes a signal  $\theta_p$  drawn from a known distribution, and the defendant independently observes a signal  $\theta_d$  drawn from the same distribution. The principal results of the paper apply to a "basic litigation model" in which the distribution is the uniform distribution; this extends without loss of generality to any uniform distribution between a lower bound of  $L$  and an upper bound of  $U$ . In the event that settlement fails, a judgment is entered in the amount of the average  $(\theta_p + \theta_d)/2$ . This assumption makes the model tractable. It may seem problematic for the judgment to depend on the signals, rather than for the signals to depend on the underlying truth to be revealed at judgment. We will return to that issue when we discuss the correlated signals model, but one can imagine circumstances in which this is realistic. For example, the parties might have information about different components of damages in a case in which liability is uncontested, and should trial ensue, the information will be revealed and the judgment will be the sum.

The Friedman and Wittman tour de force is their derivation of a Nash equilibrium in the basic litigation game. In this equilibrium, the plaintiff will ordinarily offer  $\frac{2}{3}\theta_p - 2c + \frac{1}{2}$ , and the defendant will ordinarily offer  $\frac{2}{3}\theta_d + 2c - \frac{1}{6}$ , where  $c$  represents each party's trial cost. The word "ordinarily" signals what may seem a mild caveat: Neither party will ever make an offer beyond the range of the other party's possible offers. Thus, the plaintiff's offers are truncated above at  $\min(1, 2c + \frac{1}{2})$  and below at  $\max(0, 2c - \frac{1}{6})$ , while the defendant's offers are truncated above at  $\min(1, \frac{7}{6} - 2c)$  and below at  $\max(0, -2c + \frac{1}{2})$ .

These truncations may seem unimportant, because they will never affect whether a case will settle. But the plaintiff's lower truncation increases the plaintiff's surplus from settlement, and the defendant's upper truncation has the reverse effect. Indeed, we will see later that the piecewise linearity of the model is critical to proving a Nash equilibrium. Friedman and Wittman do not eliminate the possibility that there might be some nonlinear Nash equilibrium, but they prove that the equilibrium they derive is the unique nontrivial piecewise



linear equilibrium. There are also infinitely many trivial equilibria, in which the plaintiff's settlement demands always exceed the defendant's. We will return to the challenges inherent in filtering out these equilibria when we develop a computational model.

Friedman and Wittman's model permits them to focus on the trial rate. They derive a piecewise quadratic formula for the trial rate, and they also examine, in the tradition of Priest-Klein, how the trial rate varies near the midpoint of the decision spectrum. They show that when trial costs are sufficiently low ( $c < \frac{1}{6}$ ), the probability of a trial is higher, the farther the judgment would be from  $\frac{1}{2}$ , and when trial costs exceed this threshold, the probability of a trial is highest at the  $\frac{1}{2}$  point. The intuition is that when costs are high, the parties become more generous, and so the plaintiff's range of offers will be below the defendant's. The truncations then ensure that cases at the extremes, where either both parties receive a low signal or both parties receive a high signal, are more likely to settle. When trial costs are low, the parties are less generous, and the plaintiff's range of offers will be above the defendant's. Cases at the extremes are then less likely to settle. With the basic litigation game, the  $\frac{1}{6}$  cost threshold occurs where the parties' range of offers are equal. Friedman and Wittman also offer a graphical argument that extends to other continuous distributions, though they do not expressly consider the case where liability rather than damages is uncertain.

The Friedman and Wittman model can be extended only with considerable effort. Most fundamentally, it does not extend easily to a case in which the expected judgment is not simply the sum of two parties' signals. What if a weighted sum is used, or if the judgment is a more complex function of the signals? But even accepting the averaging of the two signals, many questions arise. One might wonder, for example, what would happen to variables such as trial rate and the distribution of tried disputes if the parties were risk averse, or if one of the parties is risk averse while the other is not. Or one might wonder what the effect would be of an offer-of-settlement rule, in which the trial court punishes parties who make offers that after the judgment seem stingy. Certainly the Friedman and Wittman model can help generate intuitions on these matters, but the real answer to these questions is that new models are needed. We will see how a computational model can easily be adapted to these situations.

### 2.3 Dari-Mattiacci and Saraceno's Evidentiary Signals Model

Dari-Mattiacci and Saraceno (2019) illustrate the challenge of building on Friedman and Wittman by successfully extending the model to fee shifting. The article includes an online appendix with 60 pages of proofs. The difficulty stems from the need to address four principal cases, depending on relative values of parameters, and within these principal cases, to make various calculations that depend on the relative values of other parameters, including in many instances five different formulas for five different ranges of a variable. The resulting product is testimony both to human ingenuity and to endurance, and it makes

breakthroughs in our understanding of the effects of fee-shifting with two-sided asymmetric information.

As in Friedman and Wittman, plaintiff and defendant receive signals, now denoted  $\theta_{\Pi}$  and  $\theta_{\Delta}$ , respectively, and the judgment is an average of the signals. Now, however, both parties have common knowledge of the true merits of the litigation, denoted by  $q$ . The signals thus do not serve the function of informing the parties of the true merits, but rather of providing the parties with evidence that they may use to convince the court. The plaintiff's signal  $\theta_{\Pi}$  is drawn from a uniform distribution on the interval  $(0, q)$ , and the defendant's signal, on the interval  $(q, 1)$ . Because the defendant's signal can be no less than  $q$ , the plaintiff's best possible evidence, where  $\theta_{\Pi} = q$ , would convince the court that the judgment must be at least  $q$ . Similarly, the defendant's best possible evidence, where  $\theta_{\Delta} = q$ , would convince the court that the judgment must be no more than  $q$ . But the litigants do not always draw the best possible evidence.

The fee shifting rule that Dari-Mattiacci and Saraceno primarily analyze is triggered based on (1) whether the final judgment is above or below  $\frac{1}{2}$  (i.e., which party "wins" in the sense of being awarded more than half of the contested damages), and (2) whether the evidence of the winning party is sufficiently strong. If the judgment is less than  $\frac{1}{2}$ , then the defendant might be able to shift its costs to the plaintiff, but only if the defendant's signal falls below some threshold, i.e.  $\theta_{\Delta} < t$ , where  $0 \leq t \leq 1$ . Likewise, if the judgment is greater than  $\frac{1}{2}$ , then the plaintiff might be able to shift its costs to the defendant, but only if the plaintiff's signal exceeds a threshold, i.e.  $\theta_{\Pi} > 1 - t$ . An intuition is that if a party wins a case merely because its opponent has produced little evidence, a court will not order fee-shifting; another is that a court will only order shifting of fees when those fees were spent on producing strong evidence. Note that when  $t = 0$ , fees will never be shifted, so this extreme is the American rule of no fee shifting, and when  $t = 1$ , fees will always be shifted to a winning party (i.e., to the plaintiff if the final judgment exceeds  $\frac{1}{2}$  and to the defendant if the final judgment is less than  $\frac{1}{2}$ ), so that extreme is the English rule of universal fee shifting. The analysis thus effectively allows for a continuum of fee shifting rules.

This information structure enables Dari-Mattiacci and Saraceno to derive the offers that the parties will make. They prove that each party's offer function is a best response to its opponent's offer function and thus that a Bayesian Nash equilibrium exists. They also derive formulas for settlement amounts, along with identification of the ranges of parameters values where such settlements occur, and accordingly of the litigation rate. They prove that the litigation rate depends only on  $c$  (now representing the combined trial cost of the two parties) and is thus independent of both case quality  $q$  and the fee-shifting rule  $t$ . This produces the surprising conclusion that the litigation rate is the same under both the American and the English rule. Finally, they offer a calculation of litigation accuracy, and they prove that when costs are below a certain threshold, the English rule produces more accuracy than the American rule, while the reverse is true when costs are above a certain threshold. The stylized fact that litigation is cheaper in England may thus help explain the choice of rule in each country.

The Dari-Mattiacci and Saraceno model adopts a number of assumptions, mostly for purposes of tractability. Let us identify many of these assumptions, so that we can assess whether they present concerns about the generality of the conclusions and whether a computational model might be able to relax them.

### Structural constraints

**Piecewise linearity** Dari-Mattiacci and Saraceno explicitly assume a linear relationship between the parties' signals and their offers. They allow, however, for discontinuities in the linear relationship. In this sense, the strategies they model are similar to those of Friedman and Wittman, and indeed Dari-Mattiacci and Saraceno similarly truncate the parties' strategies. The assumption is somewhat stronger, however, in that Friedman and Wittman demonstrated that the piecewise linear strategies they derived would be a Nash equilibrium even when nonlinear strategies are possible. On the other hand, Dari-Mattiacci and Saraceno allow for additional discontinuities at points where fee-shifting would change. This is central to the design of their model and the thrust of their analysis. Because fee-shifting depends partly on the quality of the evidence possessed by the winning party, a litigant will know whether it will be entitled to fee-shifting if it wins, and the signal values at which this fact changes are points at which Dari-Mattiacci and Saraceno are able to break the problem down into smaller pieces. Piecewise linearity thus allows for explicit modeling of the effects of changes in a fee-shifting rule, but because it is unclear how restrictive this assumption is, it is a prime candidate for relaxation in a computational model.

**Asymmetric information quality equivalence** Recall that the plaintiff receives a signal in the range  $(0, q)$  and the defendant, in  $(q, 1)$ . As a consequence, when  $q > \frac{1}{2}$ , the plaintiff's signal has a greater potential effect than the defendant's, and when  $q < \frac{1}{2}$ , the reverse is true. The single variable  $q$  thus serves two, independent functions in the model: one is to represent the "true merits" of the case, while the other is to represent the degree of information asymmetry. This greatly increases the tractability of the model, and plausibly it allows for consideration of both issues related to accuracy and issues related to information asymmetry. The problem, though, is that the issues are necessarily conflated; where a case is at an extreme of the probability distribution, there is always high information asymmetry. There is no particular reason to believe that true merits should generally track information asymmetry in this way. The question thus arises whether the results would be the same if the model allowed independent variation of true merits and information asymmetry.

### Parameter values

**Balanced asymmetric information** Meanwhile, the true merits variable is constrained so that  $\frac{1}{3} \leq q \leq \frac{2}{3}$ . The reason for this constraint is that

with more extreme values of  $q$ , the increasingly one-sided nature of asymmetric information leads the pure strategy equilibria derived by the authors to break down. This highlights once again the problematic nature of asymmetric information quality equivalence, because it means that the authors not only cannot model situations with relatively high information asymmetry, but also that they cannot model situations in which the true merits of a case are near the extremes of the probability distribution. Perhaps a computational model might be able to find an equilibrium with relatively extreme quality values and/or with relatively extreme information asymmetry, and this could help extend the understandings provided by the model.

**Low or moderate cost** The authors make an implicit assumption that the cost variable is not so high that the plaintiff's untruncated offer range is entirely below the defendant's untruncated offer range. This implicit assumption also exists in Friedman and Wittman. The truncation functions defined by Friedman and Wittman are undefined, because when their  $c$  is sufficiently high, they instruct that the plaintiff's offers should be truncated above at 1 and below at a number greater than 1, and similarly the defendant's offers are truncated below at 0 and above at a number less than 0. With sufficiently high costs, there will be many Nash equilibria; the parties will be determined not to go to trial, but neither party would deviate from any positive allocation of the surplus from settlement. A computational model may be able to handle high costs, if the equilibria are subject to additional refinements.

**Risk neutrality** The plaintiff and defendant are assumed to be risk neutral. Incorporating risk aversion (especially the possibility of asymmetric risk aversion) into the model would likely add considerable challenge, though the argument could still proceed in case-by-case fashion. Incorporating risk aversion is virtually costless to a computational model, requiring only the transformation of the parties' utilities in any game outcome. A computational model could also easily incorporate behavioral assumptions, such as regret aversion, which Guthrie (1999) [11] argues may help explain litigation behavior.

## Game structure

**Fee-shifting structure** Recall that fee shifting depends on which party wins more than half of the judgment at trial and also on the quality of the evidence produced by the winning party. One might imagine alternative bases for fee-shifting. For example, fee-shifting might depend on the evidence produced by the losing party. The intuition would be that the court is punishing the losing party for proceeding with such a weak case. Meanwhile, one might also shift fees based on both parties' evidence. Indeed, Dari-Mattiacci and Saraceno explicitly consider fee-shifting based on the margin of victory, defined by a parameter  $m$ , where  $0 \leq m \leq 1$ . With this approach, if  $\theta_{\Pi} + \theta_{\Delta} < m$ , then the plaintiff must pay the defendant's fees, and if  $\theta_{\Pi} + \theta_{\Delta} > 2 - m$ , then the

defendant must pay the plaintiff's fees. In this regime, if  $m = 0$ , no fee shifting occurs (the American rule), and if  $m = 1$ , fee shifting always occurs absent an evenly split judgment (the English rule); thus, the margin-of-victory approach converges with the other approaches at the extremes. Dari-Mattiacci and Saraceno explicitly calculate the parties' offers under this approach, but they do not prove their results related to accuracy. This raises the question whether their accuracy results are robust to the alternative specification. One might also imagine other fee-shifting rules, such as a quadratic rule that imposes a fee-shifting schedule progressive in the distance from the middle of the spectrum or one styled as an offer-of-settlement rule based on the distance between the judgment and the parties' offers. Changing the fee-shifting rule should be trivial in a computational model.

**Damages vs. liability** Dari-Mattiacci and Saraceno explicitly describe their model as one in which the parties are arguing about how to divide a disputed asset, such as in a case of divorce, and they point out that without loss of generality, this can be extended to a judicial determination of damages between some minimum and maximum value. An extension would be to consider cases where liability is at issue, i.e. where the plaintiff will receive 1 if  $\theta_{\Pi} + \theta_{\Delta} > 1$  and 0 otherwise. For example, they might generalize the model to an arbitrary cumulative distribution function mapping  $\theta_{\Pi} + \theta_{\Delta}$  onto the judgment, but this would add considerable challenge. Once again, this should be trivial in a computational model, which need only transform the judgment values in particular cases, either to 0 or 1 or based on some other distribution.

**Signal variance independent of true merits** Dari-Mattiacci and Saraceno refer to the signals that the parties receive as "evidence" of the true merits of the case, but there is a paradox: The parties are assumed to know the true merits of the case ( $q$ ) and indeed use this information in constructing their offer functions. Thus the variance in the signals that each party may receive has nothing to do with the merits. Given the fixed value of  $q$ , whether the plaintiff receives a signal slightly above 0 or slightly below  $q$  tells the plaintiff nothing about the true merits. What receipt of the signal accomplishes is to inform the plaintiff about the plaintiff's likely ability to persuade the judge about the true merits. The judge does not know the true merits, but is trying to guess the true merits. The higher  $q$ , the higher the parties' signals will tend to be, so the judge's strategy is reasonable, even if non-Bayesian. But the result is that from the perspective of the parties, for whom  $q$  is fixed, the randomness in case outcomes has to do only with who is lucky in finding promising evidence.

This point can also clearly be seen in a transformation of the model that Dari-Mattiacci and Saraceno offer. They note that the signals  $\theta_{\Pi}$  and  $\theta_{\Delta}$  can be mapped one-to-one onto signals from 0 to 1, which they label  $z_{\Pi}$  and  $z_{\Delta}$ . These signals are thus independent signals from a unit uniform distribution, and the  $\theta$  signals can be derived from them according to the formulas  $\theta_{\Pi} = qz_{\Pi}$  and  $\theta_{\Delta} = q + (1 - q)z_{\Delta}$ . This highlights that the  $\theta$  signals result from commingling

the true merits of the case and the random uniform distribution draws. With these transformations, the judgment depends on the following formula:

$$J(z_{\Pi}, z_{\Delta}) = \frac{1}{2}q + \frac{1}{2}(qz_{\Pi} + (1 - q)z_{\Delta}) \quad (1)$$

As this presentation makes clear, the judgment amount is half based on the true merits of the case, independent of any evidence presented by the parties. Meanwhile, the amount is half based on a weighted average of the parties' uniform distribution draws, with the weights equal to  $q$ . Recall that  $q$  represents the degree of information asymmetry, as well as the true merits. Thus, in effect, half of the judge's decision is based on the true merits and half of the judge's decision depends on a weighted average of signals that are entirely independent of the true merits. The only reason that this makes sense from the perspective of the judge is that the weighting variable happens to reflect not only information asymmetry, but also, as a result of the signal variance independent of true merits assumption, the true merits. From the perspective of the parties, the judge might as well be responding based on the judge's political inclinations or the breakfast cereal the judge ate that morning, with the parties' evidence concerning those matters rather than the merits. The question is whether this affects the results of the model.

It might appear that a simple change in conceptual approach could make irrelevant the problem of the independence of the signal variance and the merits. Suppose we conceive of the independent  $z$  signals as being part of the true merits, rather than separate from the true merits. On this formulation,  $q$  represents knowledge that the parties share about the true merits, and the  $z$  signals represent private information about the true merits. The judge adds these together, weighing the  $z$  signals by the information asymmetry (which happens to equal  $q$ ) to obtain the true merits. With this reformulation, the model remains a two-sided asymmetric information model, but the signals are genuinely signals of the merits, even if they are weighted by some variable other than the true merits. Critically, if the weighting variable happens to be  $q$ , this conceptual reformulation has no impact at all on which cases settle. The parties' offers remain the same functions of their signals as before. We simply need to measure accuracy relative to  $J$ , considering the judgment to be the correct answer and comparing it to the case outcome, including fee shifting and costs. This would indeed eliminate any lingering concern about the extent to which the independence of signal variance and the true merits is driving model results. The problem is that, as their title reflects, one of Dari-Mattiacci and Saraceno's principal results relates to accuracy, and they define accuracy relative to true merits, defined in turn as  $q$ . Let us thus turn to their definition of accuracy.

**Accuracy definition** Dari-Mattiacci and Saraceno define inaccuracy in their appendix as “the square distance between the expected outcome  $E_t$  and the merits  $q$ .” Let us analyze this from right to left, beginning with  $q$ .

**Definition of true merits** As noted above, the use of  $q$ , known to both parties, highlights that the signals that the parties receive tell them nothing about the actual merits of the case, only about how accurately the judge will be able to compute the true merits. It is worth stressing that there is nothing inherently wrong with Dari-Mattiacci and Saraceno's approach. The approach merely leaves the reader wondering whether their results apply to a world in which the evidence that parties find serves the function of informing them of the true merits. Given the Dari-Mattiacci and Saraceno conceptualization, the use of  $q$  is correct. The definition assesses the inaccuracy of the non-Bayesian judge for any  $q$ , and one could consider developing measures of accuracy that aggregate over hypothetical distributions of  $q$ . If one were to adopt the alternative conceptualization above, then the definition of true merits must be changed from  $q$  to one that incorporates also the parties' weighted  $z$  signals.

The definition of  $E_t$  introduces its own complications. The definition is complex, involving double integrals over both costs and the parties' signals. But the essence is that it is a measure of the expected outcome of a dispute, taking into account both the settlements and the trials. The accounting for settlements is unproblematic, but the outcome in the event of trial that they calculate is represented by  $G$ , which "captures both the decision on the merits and fee shifting." For example, if the judgment is for 0.45 and the plaintiff pays costs of 0.10 to the defendant, then  $G = 0.35$ . The inclusion of fee shifting costs reflects that imposition of fee shifting not only affects settlement negotiations, but also affects the amount that the plaintiff must pay to the defendant at trial. There are, however, two arguable problems with using the expectation of  $G$  in the event a case goes to trial as the measure of outcome accuracy.

**Accounting for costs** The first problem is that this measure ignores the pre-fee shifting costs that the parties pay. Dari-Mattiacci and Saraceno note "that the plaintiff receives  $G - \frac{c}{2}$  and the defendant pays  $G + \frac{c}{2}$ ." Imagine a case with very high costs and no fee shifting, where each party spends a million dollars and the court arrives at precisely the correct conclusion that the defendant owes the plaintiff 50 cents. From this definition's perspective, this outcome counts as a perfectly accurate result. That is a plausible definition of accuracy, but one that offers no comfort to the parties. An alternative definition of accuracy would consider any amounts actually spent at trial, for example counting the outcome from the plaintiff's perspective as  $G - \frac{c}{2}$ . A similar definition could measure accuracy from the defendant's perspective. Either of these two approaches captures three distinct aspects of costs: (1) the costs impact settlement negotiations; (2) when trial occurs, the costs are deadweight losses to society at large; and (3) costs may reduce (or perhaps in some cases increase) the accuracy of adjudication viewed as a black box from the perspective of each individual litigant. The second consideration may produce activity-level effects (e.g., one may not open a business in a litigious industry), while the third consideration may produce primary conduct effects (e.g., a business might take

excessive precautions to avoid injuries). Thus, while Dari-Mattiacci and Saraceno's measure may be interesting as an abstract inquiry into whether the court system achieved the correct answer, it would also be useful to assess accuracy by considering the total amount of money actually paid or received by each of the parties, including costs and the judgment itself, relative to the amount that would be paid in a hypothetical perfect judicial system. This change would impose no demands on a computational model, which can calculate performance statistics based on any definable metric.

**Outcome expectation** There is also an additional problem with the use of  $E_t$  as the measure of the outcome. The (in)accuracy measure focuses on the expectation of settlement or trial results, rather than on the actual result in particular cases. It is a comparison of the expectation of the result with the true merits, not a measure of the error. If, for example, there are two scenarios in which the correct result based on the true merits would be for the defendant to pay the plaintiff 0.50, and in one scenario the defendant pays 0 and in the other scenario the defendant pays 1, then this measure would count the legal system as perfectly accurate. A counterargument would be that only the expected judgment matters because, as discussed above, the parties are risk-neutral, and risk-neutral parties would be indifferent between receiving perfectly accurate results and results that are correct on average. Ideally, however, the definition of accuracy would not be dependent on an assumption about risk neutrality.

An alternative definition would aggregate the distance between the actual outcome and the ideal outcome in each case. In more technical terms, ideally, instead of calculating a measure of inaccuracy that is a function of  $E_t$ , the authors might have calculated a measure of expected inaccuracy in which the inaccuracy is calculated within each case rather than based on an average across cases. Easier said than done, of course. This would require moving a minus  $q$  term and a squared term within the double integrals in the current  $E_t$  definition. But improving on this aspect of the accuracy definition should be straightforward for a computational model.

**Squared vs. absolute value** Finally, one might quibble about the use of a squared term rather than an absolute value. It is conventional to measure accuracy using the  $\ell_2$  norm rather than the  $\ell_1$  norm. The convention reflects the dominance of ordinary least squares regression over least absolute deviation regression, but that dominance stems at least in part from the greater tractability of the former. Portnoy and Koenker (1997) [17] note that computational power mitigates this advantage, and that an advantage of the  $\ell_1$  norm is that it is more robust to outliers. Ultimately, the choice between them should depend on some broader model of social welfare. For example, in a model like Hylton (2002), disputes arise endogenously as a result of primary activity, such as decisions to take care, allowing derivation of direct social welfare measures. One might then compare hypothetical legal systems that perform differently on  $\ell_1$  and  $\ell_2$  measures of accuracy to determine which measure makes more sense in



that legal context. That is beyond the scope of this paper, though endogenizing disputes may be a fruitful direction for future research on computational models. In the meantime, a task for the computational models here is to be able to switch between the norms.

## 2.4 Klerman, Lee, and Liu's Correlated Signals Model

While Dari-Mattiacci and Saraceno improve on the Friedman and Wittman model by adding support for fee-shifting, which Friedman and Wittman did not consider, Klerman, Lee, and Liu seek to improve the Friedman and Wittman model's usefulness in addressing an issue that Friedman and Wittman directly considered: the distribution of tried cases given selection effects from settlement. Recall that Friedman and Wittman found that with low costs, tried cases will tend to be in the middle of the distribution, while with high costs, tried cases will tend to be at the extremes of the distribution. This result may depend on the structure of signals in the model, in particular the relationship between the plaintiff's and defendant's signals, drawn from separate distributions. The signals are correlated conditional on the judgment, because the judgment derives from the signals, but the signals are unconditionally independent. What Klerman, Lee, and Liu endeavour to accomplish is to model litigation with the signals as functions of the true state, rather than the other way around. If the true state of the world is that the plaintiff's case is strong, both parties are more likely to receive relatively strong signals than they would receive if the true state is that the case is weak.

Before focusing on the structure of and assumptions in their model, it is worth addressing which information structure is more realistic. The information structure in Friedman and Wittman, continued as well in Dari-Mattiacci and Saraceno, can be seen as one in which the parties have entirely independent information, and this information is averaged to the judgment. It may seem unrealistic that the parties do not have shared information, but this assumption is without loss of generality. One can conceive of the plaintiff and defendant as having much shared information, allowing them to narrow down the range of the judgment. At least in the most general version of the Friedman and Wittman model, the signals are from arbitrary cumulative distributions. Meanwhile, the information structure in Klerman, Lee, and Liu may be seen as one in which the parties have different beliefs based on essentially the same information. These beliefs may themselves be a form of information. For example, the plaintiff may have beliefs based on the analysis performed by the plaintiff's lawyer, who could reach a different conclusion based on the same facts from the defendant's lawyer. Or, the plaintiff and defendant may have different recollections of events. For our purposes, we can consider the signals to be information, but informally, the beliefs construct may provide a better intuition for why signals might be correlated.

In assessing which information structure makes more sense, one might ask two questions: (1) Does the case concern damages or liability? In a damages case, it is more plausible to see each piece of evidence adding independently

to the amount of money that the plaintiff can expect to receive. The pain that the plaintiff experiences, the lost wages, the effect of the accident on the plaintiff's relationships with others: each of these can be seen as a separate form of injury entitling the plaintiff to independent relief. If each of these pieces of information is largely private until trial, then a model that treats the parties as having independent information may well be accurate. In a liability case, however, the question is whether the evidence meets the decision standard, such as preponderance of the evidence. As noted above, outcomes are thus highly nonlinear in evidence. Thus, even if parties have separate evidence, it seems unrealistic that the judgment would be equal simply to the sum of the values of the evidence.

(2) Does the failure of bargaining mean litigation or trial? The literature on settlement bargaining generally disregards pre-filing bargaining. Pre-filing bargaining can be seen as reflecting similar logic to post-filing bargaining. If no deal is reached pre-filing, then a litigation will result (with some average outcome depending primarily on how long settlement negotiations drag on and what the ultimate settlement is), while if no deal is reached post-filing, a trial will result (with some average outcome depending on the decision of judge and jury). Thus, the same general models apply, but the information structure may be different. Shortly after the possibility of litigation becomes apparent, there will be some shared information but also some independent information, and so at least at very early stages, the information structure of Friedman and Wittman may be closer. Discovery will tend to equalize the parties' information, though incompletely and leaving room for different beliefs about the information, so after discovery, the information structure of Klerman, Lee, and Liu may be more appropriate.

Ideally, one model might be embedded within the other, or a unifying model might explicitly model the interactions among various pieces of information, some of which are inherently common knowledge, some of which can be shared easily, some of which can be partly shared through discovery, some of which can never be shared, and some of which there might be varying levels of disagreement about. A computational model may well be able to achieve that, though it is beyond the scope of the current project. In the meantime, the fact that different models may be appropriate for different situations highlights a limitation of mathematical models that the computational modeling in this paper may partially address. Different models are built to address different issues, such as fee shifting for Dari-Mattiacci and Saraceno and dispute selection for Klerman, Lee, and Liu, but the modeling choices that they make for tractability in addressing these issues may also have implications for information structure. Thus, it is difficult to determine whether, for example, Dari-Mattiacci and Saraceno's fee shifting results would be the same in a correlated signals model. A computational model will use a particular information structure, but it facilitates examination of a larger number of outcome variables. Similarly, Klerman, Lee, and Liu assume the American rule, and it would be difficult to modify their model to incorporate fee-shifting. (See Chang and Hubbard (2018) for empirical evidence on the Priest-Klein model outside the United States.) In a

computational model, this again becomes a trivial change, because it requires only changing the code that determines the judgment at trial.

In part because they do not focus on fee-shifting and because they ignore the possibility that there could be disputes about damages as well as about liability, the model of Klerman, Lee, and Liu is admirably general. A random variable  $Y$  (with realization  $y$ ) determines the disputes merits. If  $y > 0$ , the plaintiff has a probability  $\theta_H$  of winning, and if  $y \leq 0$ , the plaintiff's win probability is  $\theta_L$ , where  $0 \leq \theta_L \leq \theta_H \leq 1$ . (Note that while  $\theta$  was used to represent the signals in the Dari-Mattiacci and Saraceno model, it is used to represent win probabilities here. Because this paper's purpose is to point out how computational models can enhance mathematical ones, rather than proposing its own grand unified theory, the paper uses each paper's notation in discussion of that paper.) The Priest-Klein model has a similar setup but requires that  $\theta_L = 0$  and  $\theta_H = 1$ . The more significant improvement lies in the correlated signals, with plaintiff receiving a signal  $y_p$  from the distribution  $Y_p = Y + \epsilon_p$ , and the defendant,  $y_d$  from the distribution  $Y_d = Y + \epsilon_d$ , where the error terms are normally distributed, i.e.  $\epsilon_p, \epsilon_d \sim N(0, \sigma^2)$ . Each party is assumed to be a sophisticated Bayesian who knows the distribution  $Y$  and of the other party's signal based on its own signal, though the analysis also allows litigants to be naive and construct estimates as if the probability distribution were flat.

Klerman, Lee, and Liu include models in which the plaintiff makes a take-it-or-leave-it offer to the defendant, where the defendant makes such an offer to the plaintiff, and where the parties engage in Chatterjee-Samuelson bargaining. We will focus on the last of these to make our analysis here as close to our analysis above as possible and also because this approach, though stylized, seems a more realistic reduction of an extended bargaining process than one with take-it-or-leave-it offers. A computational model can easily switch between these approaches; simultaneous offers can even be used to model take-it-or-leave-it bargaining, for example if the code provides that a defendant's offer will be deemed accepted by the other plaintiff if the plaintiff's submission is greater than or equal to the offer. A more difficult challenge, which we defer to a later paper, is bargaining over multiple rounds.

Another division in the Klerman, Lee, and Liu analysis allows for two different approaches to modeling the distribution of cases. One approach is to model the distribution,  $g_Y(x)$ , as an improper uniform distribution over all real numbers, so that  $g_Y(x) = 1$  for all  $x$ . Although  $g_Y(x)$  cannot then be considered a probability density function, because it does not integrate to 1, Klerman, Lee, and Liu are able to apply earlier literature on such distributions. A virtue of mathematical models not easily available in computational ones is the ability to play with mathematical abstractions involving infinity. A computational simulation uses finite numbers and so it cannot engage this trick, though it can allow for case quality to be distributed over an arbitrary finite range. For this reason, we will focus on the portion of the Klerman, Lee, and Liu analysis in which they assume that  $g_Y(x)$  is a continuous, strictly positive probability distribution.

Given such a distribution, Klerman, Lee, and Liu identify two types of equilibria. Both types of these equilibria are limit equilibria, in the sense that they

are defined as  $\sigma \rightarrow 0$ . This helps establish versions of the Priest-Klein 50% limit hypothesis, namely that the plaintiff will win 50% of litigated cases regardless of the distribution of meritorious cases overall. a hypothesis that applies only in the limit as the parties' error approaches zero. The first type of equilibrium is a symmetric limit equilibrium. By symmetric, the authors mean that the plaintiff's offer given a signal  $x$  would be equal to the defendant's offer given a signal  $-x$ . They establish equilibrium in the usual way, by showing that the plaintiff's offer given its signal maximizes the plaintiff's utility, and similarly that the defendant's offer given its signal minimizes a loss function. They show that the plaintiff trial win rate in this equilibrium is  $\frac{\theta_H + \theta_L}{2}$ , confirming Priest-Klein where  $\theta_L = 0$  and  $\theta_H = 1$ . The plaintiff calculates a threshold, offering 0 below this threshold and  $\theta_H + \theta_L$  above; the defendant does the same thing, with the negative of the same threshold function.

Klerman, Lee, and Liu also identify another pair of equilibria for each  $\sigma$ , which they call obstinate limit equilibria. One features an obstinate plaintiff, and one features an obstinate defendant. The obstinate plaintiff always sets its offer at the same relatively high constant value given a particular  $\sigma$ , thus screening the defendant, and the obstinate defendant does the opposite. The Priest-Klein 50% limit hypothesis does not apply to these equilibria, except in the coincidence of one of the win probability parameters being set to 50% or obstinate plaintiff and obstinate defendant equilibria occurring in the right proportions. The plaintiff trial win rate is  $\theta_L$ , and the defendant trial win rate is  $\theta_H$ .

Thus, in both equilibria, the offer function for at least one party either is entirely flat or consists of two flat segments with a jump discontinuity. This presents a contrast with Dari-Mattiacci and Saraceno. Although they also accepted the possibility of discontinuities, the Dari-Mattiacci and Saraceno offer functions include regions in which offers monotonically rise as a result of the signal received. Klerman, Lee, and Liu do not prove that the equilibria they identify are exclusive, however, so their model leaves open the possibility that other equilibria exist that include monotonically increasing segments.

An assumption in the model that might seem problematic is that although a continuous range of case types exist, there are only two possible probabilities of plaintiff victory,  $\theta_L$  and  $\theta_H$ . The intuition that this is not a problem is that the parties' signals,  $Y_p$  and  $Y_d$ , can take on values across the real line, and thus the expected value of the probability of victory can take on any value between  $\theta_L$  and  $\theta_H$ . But the question remains whether the division of cases into two types contributes to the derived equilibria in each of which a party's offer takes on either of two values rather than a continuous range.

Assessing whether this is true is important for assessment of the validity of the Priest-Klein hypotheses, but also more broadly. If litigants really are inelastic with respect to the signals they receive, except at a single jump discontinuity, then one might worry that settlement outcomes will necessarily be highly inaccurate. Thus, the Klerman, Lee, and Liu results invite questions about the accuracy of litigation. That is a question that they do not directly address and moreover that is difficult to address within the context of their model. After all,

their model has no concept of a “correct answer.” Perhaps one could postulate that perfect accuracy would be achieved if of all cases in which  $y > 0$ , the proportion  $\theta_H$  are resolved with the plaintiff winning, and similarly  $\theta_L$  for  $y < 0$ , but this begs the question of whether the proportion of plaintiffs who win are those whom one would want to win. A model that began by designating some cases as those in which the defendant is truly liable, distributing case strength as a function of true merits, could address the concern.

The broader critique of the two case types limitation may seem not only unfair (the proofs are quite demanding even with two types), but also ironic. The computational models to be discussed in the next section are inherently discrete models, allowing for only some finite, relatively small number of discrete signals. Yet the critique makes two points. First, even many mathematical models that include continuous distributions also include discrete elements. Thus, computational models should not necessarily be rejected for their own (in)discretions. Second, and more importantly, computational models may help relax discretizations in mathematical models. Instead of two case types, for example, a computational model might allow three, four or five. Arguably, it is better to have a handful of case types and a handful of signal types than to have just two case types and an infinite range of signal types. The handfuls may leave for the possibility of artificially imposed discontinuities, but at least they can clarify whether binary outcomes in parties’ strategies stem from a model limitation or arise as a fundamental aspect of the information and bargaining setup.

The existence of only two case types improves the model’s tractability in another way that highlights an important assumption. Klerman, Lee, and Liu assume that  $\theta_L \geq C_p$  and  $\theta_H \leq 1 - C_d$ , where  $C_p$  and  $C_d$  represents the plaintiff’s and defendant’s costs, respectively. The reason for this assumption is that it guarantees that a plaintiff will always have a credible threat to take a lawsuit to trial, and the defendant will always have a credible threat to defend. Thus, they do not need to worry about whether a plaintiff either might not file a case in the first place or might abandon a case after settlement fails, and similarly about whether a defendant might default, either immediately or after bargaining failure. Of course, with a range of case types, they might similarly assume that even with the extreme case types, it will be in the interest of both litigants to proceed. But this assumption becomes increasingly implausible if the most extreme case types have sufficiently low and high probability, as for example in the original Priest-Klein model where  $\theta_L = 0$  and  $\theta_H = 1$ . The authors justify their assumption by noting that allowing parties not to take a case to trial “would further complicate the model and would distract from the key point of the analysis here, which is selection, not nuisance suits.” (p. 393 n.10) That is reasonable, but a computational model might allow for assessing whether the possibility of noncredible threats might affect selection rates at trial. If, for example, plaintiffs with low-quality cases drop out after bargaining failure, that might increase the win rate of plaintiffs in the cases that remain.

An additional simplification in the model that may be of concern is the assumption that damages are fixed. There may well be many cases in which

damages are more or less known and liability is unknown, but there are also cases in which both liability and damages are contested. One might have the intuition that this should not matter because each party, focusing on the liability question, calculates the expected damages, making appropriate adjustments for risk aversion. On the other hand, if there is two-sided asymmetric information about both damages and liability, bargaining failure could result from disagreement about damages, disagreement about liability, or both, and so a prediction of 50% victory does not mean a 50% win rate on liability. There are effectively two ways for the defendant to win, but the win rate applies to only one of them. Klerman, Lee, and Liu's project, in conjunction with Klerman and Lee (2016), is to place Priest-Klein on sounder Bayesian footing, and at that they succeed spectacularly, but a computational model may help flesh out the further implications of correlated signals.

### 3 Computationally Modeling Two-Sided Asymmetric Information

This section is designed more as primer than as critique. The challenge, given the models reviewed in the previous section, is to determine to what extent discrete computational models can mimic the mathematical models, producing similar Bayesian equilibria, and how computational models might be able to relax some assumptions. At the same time, it is important to highlight aspects of the existing models that the computational models cannot easily replicate. The section uses distinct modeling techniques for each of the three articles presented in the prior section. This is not because one modeling technique happens to be the appropriate one to use for the particular mathematical model. To the contrary, the last and most sophisticated of the techniques, counterfactual regret minimization, could easily be applied to all three problems, and the other two techniques could also be applied to all three, albeit with more constraints on the players' strategies than would be ideal. Indeed, one of the great virtues of the computational approach is that one can use the same algorithmic game theory techniques (indeed, even the same code base, such as the open source code used for this paper) to model many different games.

The organization reflects the pedagogical goal of providing some intuition for the different methods. The simplest, but least efficient, technique for finding Nash equilibria is simply to enumerate all of the plaintiff's and all of the defendant's possible strategies and then to examine each pair of strategies to determine if it is a Nash equilibrium. The algorithm is the familiar iterated elimination of dominated alternatives that game theory students learn in the first week or so of class. We start with this section at some peril; the technique is both widely known and of limited power. But the discussion highlights two points. First, the settlement bargaining literature has heretofore ignored the possibility of employing a technique with which virtually everyone writing in the literature is familiar, yet the approach makes it easy to derive interesting

results or extensions. Second, with this technique, the atomic unit of decision is a litigant's strategy writ large. That is, a single strategy specifies what the litigant will do in all circumstances.

The more advanced techniques succeed by using smaller atomic units. With the technique of "no regret" learning, the atomic unit can be considered to be information provided to a litigant by what the game theory literature refers to as Nature or Chance. For example, in litigation, this information could consist of one or more signals of the litigation's value that each litigant receives at the beginning of the game. The approach allows the litigant to treat the problem of determining appropriate responses independently for each information result provided by Chance. A litigant applying the "no regret" learning techniques need not worry about the consistency of the parties' own decisionmaking at different states of nature, or more importantly, even about how the litigant's adversary might react to the evolution in the party's strategy. Because each information condition is treated separately, "no regret" learning avoids the exponential blowup in number of strategies with iteration of dominated alternatives that occurs when complexity increases.

The final technique, counterfactual regret minimization, narrows the atomic unit to a specific information set, including not only the information provided by Chance, but also information learned in the course of the game as a result of the other parties' moves. For example, bargaining may proceed over multiple rounds, with litigants in later rounds accounting for information such as the offers made by opponents in early rounds. Similarly, after bargaining fails, a litigant may decide whether to concede the case rather than go to trial in part based on the offers made. Counterfactual regret minimization structures learning among related decisions, so that each party's strategy at each information set is informed by updated distributions of variables unknown to that party given that the information set is reached.

### **3.1 Iterative Elimination of Strictly Dominated Strategies**

We will begin by developing a computational model using the Friedman and Wittman setup, though the same approach can be applied easily to either of the other two articles. The approach is to enumerate a list of many possible strategies for each player and then systematically search for a Nash equilibrium by considering each pair of opponent strategies. The goal of this analysis is not to persuade the reader that this is the best way to implement a computational model or that this approach is better than the one that Friedman and Wittman develop. Indeed, the goal is in part to show how constrained this approach is, because allowing for greater strategic complexity makes the process unwieldy. At the same time, however, the approach is very easily implemented and it allows for great flexibility in modifying some aspects of the game. In particular, it allows, at no additional cost of computational complexity, for any modification to the game affecting the parties' utility payouts.

How can we enumerate the strategies? Recall that each player must announce a proposed offer given a signal received of the strength of the plaintiff's

case. Thus, we might constrain the parties' strategies by requiring each player to submit a strategy in the form of a linear function. To reduce further the number of strategies, we can restrict the strategies to nondecreasing lines, and we can require that each line either intersect or be tangent to a square with corners at (0,0) and (1,1). The bottom of this square represents the axis for the parties'  $\theta$  signals, and the left of the square represents the axis for the parties' offers, which we label the  $\sigma$ -axis. The height of the line at the left of the box thus represents the offer that a party will give for the lowest possible signal, and the right of the box, the offer for the highest possible signal. Let  $n_{\text{endpoints}}$  represent the number of available endpoints and  $n_{\text{angles}}$  represent the number of angles. Figure 1 illustrates available strategies if  $n_{\text{endpoints}} = 10$  and  $n_{\text{angles}} = 10$ . The simulations that follow use  $n_{\text{endpoints}} = 50$  and  $n_{\text{angles}} = 50$ , thus 25 times as many strategies as are pictured in Figure 1.




Figure 1: A fraction of the available strategies

The strategy matrix thus includes  $n_{\text{endpoints}}^2 n_{\text{angles}}^2$  entries, or in each of our simulations, 6,250,000. We arbitrarily designate the plaintiff as row player (so that the plaintiff plays the same strategy in any given row of the matrix) and the defendant as column player. For each entry in the matrix, we must approximate the parties' utility, which is accomplished by taking permutations of  $n_{\text{signals}}$  for each of the plaintiff and defendant, where the signals are drawn evenly from the unit uniform distribution. Thus, for each matrix entry, the game results must be calculated  $n_{\text{signals}}^2$  times, with each game result calculating the plaintiff's and defendant's strategies given their offers, determining whether the case settles or not, and applying the game rules to determine utility. Our simulations used  $n_{\text{signals}} = 100$ , so the total number of game plays for each simulation was 62,500,000,000. With this strategy matrix calculated, the algorithm for iterative elimination of strictly dominated strategies may be applied.

The algorithm is straightforward: For each row, identify the column in which the defendant earns the highest utility. Likewise, for each column, identify the row in which the plaintiff earns the highest utility. For every cell in the matrix, if the defendant's utility is lower than the maximum defendant's utility in the row or if the plaintiff's utility is lower than the maximum plaintiff's utility in the column, then mark that cell as permanently eliminated. This reflects that a strategy is a Nash equilibrium if and only if neither player would change



strategies given the other's choice of strategies. Repeat this entire process for as many times as leads to further elimination of matrix entries. Once there is a pass in which no further entries are eliminated, the process is complete, and the remaining matrix entries represent Nash equilibrium strategies in the game, given the constraint that strategies must be linear and given the discretization of only considering a finite number of possible lines.

This process is not guaranteed to find a unique Nash equilibrium. It may find many Nash equilibria. This may occur either because the underlying game with continuous signals features multiple equilibria or because the discretization introduces multiple equilibria, which may for example be clustered in some portion of the matrix, none corresponding perfectly to a pure equilibrium in the corresponding nondiscretized game. Ideally, one would analyze all multiple equilibria that arise and describe the family of Nash equilibria. Our approach is simply to choose the Nash equilibrium that minimizes the square distance between the parties' utilities; this is admittedly arbitrary, but it tends to produce more symmetric equilibria and thus may be easier to interpret. One implication is that where any settlement between 0 and 1 would be a Nash equilibrium, as is the case if the cost of trial is sufficiently high, the equilibrium selected will be one in which each litigant receives  $\frac{1}{2}$ . This is thus consistent with the approach in the Landes-Posner-Gould framework of assuming that parties share the surplus from settlement, but it still leaves open the possibility of trial even where the social surplus from settlement is positive.

Iterative elimination of dominated strategies also may not find any Nash equilibria. Nash equilibria are guaranteed to exist only in mixed strategies, and even where a Nash equilibrium exists in a game with continuous signals and infinitely many lines, it may not exist in the discrete game. In this case, the combination of strategies closest to equilibrium was selected. Distance from equilibrium is measured as the sum of the square of the utility gain that the plaintiff could obtain from switching rows and the square of the utility gain that the defendant could obtain from switching columns.

The results are illustrated in Figure 2. Each box represents a different simulation, corresponding to a different level of costs  $c$ . In fact, the simulation for each box did find at least one equilibrium. The blue lines with square marks represent the plaintiffs' offers; the orange lines with no marks represent the defendants' offers. The dashed lines represent the lines that Friedman and Wittman identified as furnishing equilibria. On one hand, the solid and dashed lines reflect a similar phenomenon that accords with common sense. As costs rise, plaintiffs become more generous by lowering the amount that they demand for any given signal and defendants become more generous by raising the amount that they offer for any given signal. At some point, costs are high enough that the defendant's offer line will be higher than the plaintiff's offer line. When this first occurs, some cases may still go to trial, but only when the plaintiff's signal is sufficiently higher than defendant's. In the last panel, the minimum defendant offer is greater than the maximum plaintiff offer given the Friedman and Wittman calculations, a situation in which the truncated lines would be undefined. In the same panel, the computational model produces a Nash equi-

equilibrium in which the plaintiff and defendant always settle at a constant value of  $\frac{1}{2}$  to avoid the oppressive risk of litigation; recall that this is the result of the equilibrium refinement.




Figure 2: Equilibrium offers for different costs

On the other hand, the differences between the Friedman-Wittman lines and the computational models' equilibria are stark. Plaintiffs' offers are considerably more elastic with respect to costs in the Friedman-Wittman formulation than in the computational model. Note that in the third panel, where  $c = 0.2$ , the computational model has the plaintiff's offer line above the defendant's, but the reverse is true in the Friedman-Wittman model. What accounts for this? The answer is simple. The computational model does not include the Friedman-Wittman truncations. This highlights that the truncations are not mere mathematical conveniences but in fact are central to identifying the relevant equilibria.

It might appear that it would be straightforward to modify the computational model to mimic the truncations. After all, in the Friedman-Wittman equilibrium, each litigant's line is always truncated based on the litigant's opponent's line. Thus, we can modify the simulation so that the game outcome is calculated based on the truncated version of each line. The only truncations that affect the outcome in the Friedman and Wittman model are the plaintiff's lower truncation and the defendant's upper truncation, where the defendant's offer is higher than the plaintiff's, because the other truncations designate regions in which the parties go to trial anyway. Figure 3 illustrates the results. In the case where the truncations are undefined, we use the midpoint of the defendant's minimum strategy and the lower plaintiff's maximum strategy.




Figure 3: Equilibrium offers for different costs, with truncations

Why do the results still not match up with those of Friedman and Wittman? The answer is that the truncations are central to the Nash equilibrium. When the plaintiff truncates its offers, ensuring that it never offers an amount below what the defendant is offering, it not only improves its utility against that particular defendant strategy at the defendant's expense, but also may affect utilities as against other strategies. In other words, a strategy in Friedman and Wittman represents a schedule of offers for particular signals, not the combination of such a schedule with a promise to truncate offers based on whichever opponent happens to come along. The computational model thus highlights that the truncations of Friedman and Wittman are not simply maneuvers to make their mathematical analysis tractable but are in fact central to the equilibria they develop.

The computational model thus might seem quite disappointing, this section serving as the equivalent of Elon Musk's demonstration of his truck's unbreakable glass that resulted in the glass shattering. That is a reasonable reaction. It's worth noting, though, that the problem is potentially fixable. One could define a universe of strategies consisting not simply of a start point and an angle, but also containing a truncation distance. The plaintiff's truncation distance, if nonzero, would indicate the point at which the plaintiff's offers start increasing from the plaintiff's minimum to maximum value, and the defendant's truncation distance, where the defendant's offers reach their maximum and no longer increase. This exercise, however, will be left to the reader. The real payoff of the paper will come with the more advanced models of later sections, and these models do not include linearity constraints. Moreover, the example highlights the primary practical difficulty with use of iterative elimination of dominated strategies. The utilities matrix grows with the square of the number of strategies. Thus, if each player's strategy were allowed to incorporate one of 50 truncation values, in addition to the existing endpoint and angle values, the matrix would be 2500 times as large. And a truncated line is still a fairly simple strategy. A quadratic or cubic strategy, or a strategy that is a function of multiple variables, will be cumbersome to enumerate.

Still, one should not leave with the conclusion that this tool is worthless. Many models include linearity constraints; recall that Dari-Mattiacci and Saraceno's model incorporates piecewise linearity constraints. One may still be able to learn something with a settlement bargaining model in which each party's settlement is constrained to a line. The Friedman and Wittman model is obviously superior to this model as applied to the exact same problem; a proof that identifies each party's strategy is much more useful than an algorithm that approximates it. But the algorithm has a benefit of greater flexibility. Earlier, for example, this article promised that it would be trivial for a computational model to incorporate risk aversion. We won't waste space by making good on every such promise, but let's make good on that one. Figure 4 adds a fourth dimension to the three dimensional graph in Figure 2.

The  $RA$  axis in Figure 4 adjusts for risk aversion. Each party's utility  $u'$  is redefined as a function of its risk neutral utility  $u$ , according to the formula  $u' = \ln(2^{4-RA} + u)$ . The result may look repetitive, but that is itself a bit




Figure 4: Equilibrium offers, based on cost and risk aversion

surprising. At least in this model, risk aversion does not have much effect on the parties' bargaining. The computational model also makes it straightforward to change the parties' levels of risk aversion independently, and we can examine aggregate statistics under each equilibrium. Figure 5 plots in each mini-graph a measure of accuracy (specifically, the absolute difference between the amount the plaintiff received and the amount the plaintiff would receive if trial costs were zero and all cases went to trial) with a solid blue line, as well as the trial rate with a dashed orange line. What makes the results exciting is that they are even more boring than those in Figure 4. Squinting, one can see that trial rates are ever so slightly lower when the plaintiff and the defendant have the same level of risk aversion (along the panels diagonally from the bottom left to the top right). But the central message is that even significant levels of risk aversion have very little effect on variables like trial rates. Once both parties are shading their bids in an effort to maximize individual surplus, the benefits that

would accrue to a risk-averse litigant from being less aggressive do not seem to matter much at the margin. At least, this is a possibility that deserves further exploration.

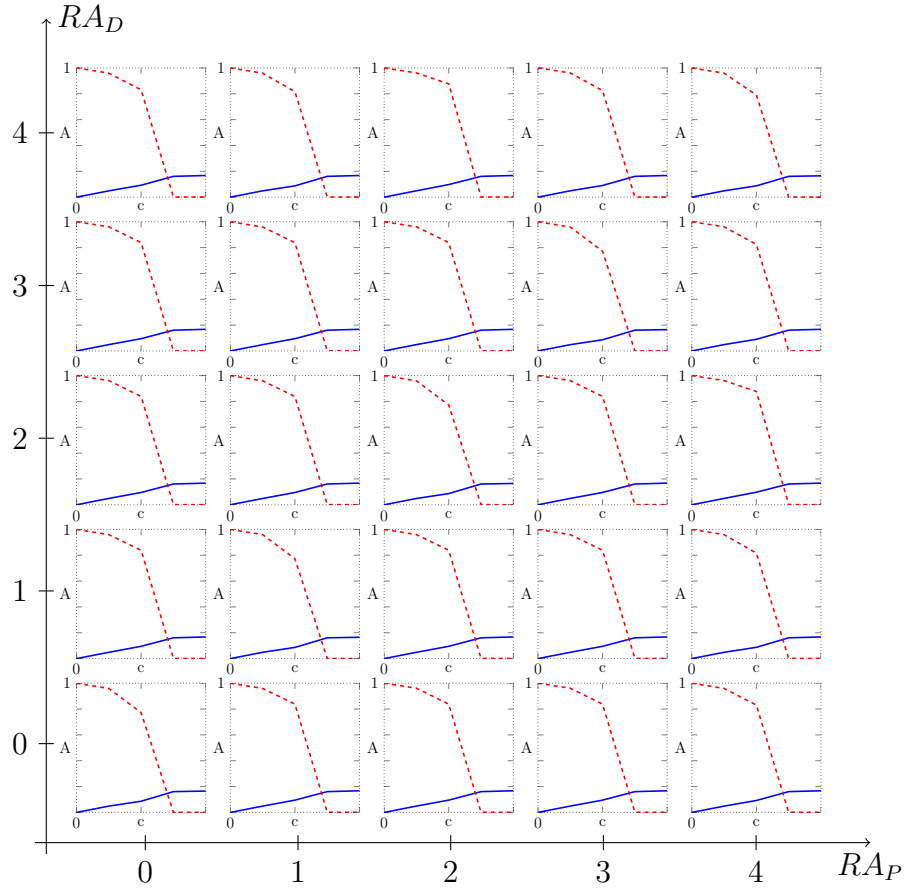


Figure 5: Accuracy and trial rates as a function of costs and risk aversion

We will make good also on the promise that the computational model can easily incorporate an offer-of-settlement rule. Figure 6 does this. The  $F$  axis represents a parameter determining the strength of the offer-of-settlement rule. When a case goes to trial, the plaintiff must pay the defendant an amount proportional to the difference in the parties' aggressiveness relative to the judgment, i.e.  $F\{(\sigma_P - J) - (J - \sigma_D)\} = F(\sigma_P + \sigma_D - 2J)$ . This figure also includes a dotted red line indicating the trial rate (using the same scale as the vertical axis).

As the figure demonstrates, the offer-of-settlement rule leads to considerably less aggressive offers by the parties, increasing settlement for any particular level of costs. There are decreasing returns from the final offer rule, however, as most




Figure 6: Equilibrium offers with offer-of-settlement rules

of the trial rate reduction arises in raising  $F$  from 0 to 0.25. Of course, to tell a more complete story about offer-of-settlement rules, one would want to look at more variables. Ideally, moreover, one would like to lift the linearity constraint or add additional complexity to the parties' strategies, for example by giving the parties options to concede rather than face trial. This paper's primary aims are methodological, so we will return to these methodological challenges but will not comment further here on offer-of-settlement rules.

### 3.2 No Regret Learning

The principal problem with the technique employed in the prior section is that it constrained the parties' strategies to being linear, and that constraint could be relaxed only with an exponential blowup in the strategies matrix. Preferable would be a technique that required enumeration of moves in discrete situations, but that allowed each move to be optimized separately from all of the other moves. For example, one might enumerate all of the offers that a defendant can make when confronted with the signal 0.30 and, separately, all the offers when confronted with 0.35 and so forth. The technique that we will use to accomplish this is "no regret" learning. See Cesa-Bianchi and Lugosi (2006) [3].

Focus first on just one decision that a player faces, such as the decision that the defendant faces when confronted with the signal 0.30. Even with no knowledge of the litigation game, the defendant can learn a strategy if the defendant plays repeatedly at time steps  $t \in 1, 2, \dots, T$ . At each time step  $t$ , let  $a \in A_\Delta$  represent an action that the defendant may select. The defendant's task is to assign a probability weight for each action. In effect, the defendant rewinds time and tries each action every time to learn what the utility for each action would have been. Define the defendant's external regret for not taking  $a$  as  $r_a^t = U_a^t - U^t$ , where  $U^t$  is the utility the defendant in fact received and  $U_a^t$  is the utility that the defendant would have received had the defendant chosen  $a$  with probability 1. Note that a positive value for regret means that the corresponding action is a relatively good one; the defendant is regretful for not having taken it with higher probability. When the regret is zero, the player either did choose  $a$  with probability 1 or chose some other action or combination of actions that happened to produce the same average utility. Finally, where regret is negative, that means that the player is pleased to have not chosen  $a$  with higher probability.

We then define the defendant's accumulated external regret as follows:

$$R^T = \max_{a \in A_\Delta} \sum_{t=1}^T (U_a^t - U^t) \quad (2)$$

The accumulated external regret represents the maximum sum of external regrets over all iterations for a single strategy  $a$ , that is, the sum of the maximum regrets that the player would have received if the player had played the same action in every iteration, assuming that the plaintiff played the same actions that the plaintiff in fact played. Critically, this is not the same as the sum of the maximum utility that the defendant could have earned in each iteration, by choosing the best action in each iteration given what the plaintiff in fact played. We can then define the average regret as  $\bar{R}^t = \frac{R^t}{t}$ . If the defendant plays a strategy that ensures that  $\lim_{t \rightarrow \infty} \bar{R}^t = 0$ , then external regret is said to be minimized. The "no regret" learning literature identifies algorithms that adjust the player's mixed strategy based on the accumulated external regrets and that succeed in minimizing regret. A critical point is that these algorithms are guaranteed to work even against hostile adversaries, for example against an adversary who observes the defendant's style of play and then seeks to counter it. This is possible because regret minimization requires only that the average regret approach zero, not the expected regret. Thus, the algorithms may recommend that the defendant mix strategies against the plaintiff. The algorithm's goal at each iteration  $t$  is to choose an action  $a$  that minimizes accumulated external regret for all iterations up to and including  $t$ , not necessarily in iteration  $t$ .

Among the most popular of the "no regret" algorithms is "regret matching." See Hart and Mas-Colell (2000). This algorithm chooses among the actions that have positive external regrets in proportion to those positive external regrets, and it eschews all actions with zero or negative regret. Formally, let  $R^{t,+} =$

$\max(0, R^t)$ . Then,

$$\sigma_{\Delta}^{t+1}(a) = \begin{cases} \frac{R^{t,+}(a)}{\sum_{b \in A_{\Delta}} R^{t,+}(b)}, & \text{if } \sum_{b \in A_{\Delta}} R^{t,+}(b) > 0 \\ \frac{1}{|A_{\Delta}|}, & \text{if } \sum_{b \in A_{\Delta}} R^{t,+}(b) = 0 \end{cases} \quad (3)$$

The Blackwell approachability theorem can be used to show that regret matching is sublinear in time and thus meets the requirements of being a “no regret” algorithm.

All of this analysis applies not only to the defendant, but also to the plaintiff. Thus, plaintiff and defendant may play adversarially, each selecting a single mixed strategy at any given iteration. The regrets that each player calculates for each action are then based on the opponent’s mixed strategy at that iteration, i.e. determining the utility that the player receives at  $t$  for  $a$  as the weighted average of the utilities that the player would receive for each action the opponent might play, weighed by the probabilities the opponent chose for that iteration. Of course, the analysis applies to each and every decision that each player must make. Just as the defendant need not worry that the plaintiff is adversarial (and vice-versa), so too need the defendant not worry about coordinating the defendant’s own strategies in two different situations, such as the strategies to be played when the signal is 0.30 and when the signal is 0.35. In effect, the defendant can assign a separate employee to tally the regrets for each of those separate decisions. So long as each employee learns the regrets that arise from playing each action, the employee need not worry what the other employees are choosing to do in the same iteration for different decisions.

The procedure that we will follow is thus simple. We will repeatedly, over 1,000 iterations, play the “no regret” learning game separately for each discrete strategy signal that the plaintiff or the defendant may receive. Does this guarantee that the mixed strategies played in the last iteration will be close to a Nash equilibrium? No, it does not. The only guarantees in a general sum game like those explored here stem directly from the guarantees of “no regret” learning. Regret minimization implies, for a particular decision optimization, that for any  $\epsilon > 0$ , there will exist a  $T_0$  such that for every  $t > T_0$ , the accumulated external regrets will be within  $\epsilon$  of the best possible fixed strategy. Aggregating this principle over all the decisions optimized by a player implies that there will exist a  $T_1$  such that for every  $t > T_1$ , the sum of the accumulated external regrets for all of the player’s decisions will be within  $\epsilon$  of the best possible fixed strategy, where a strategy now dictates a probability distribution for every decision being optimized.

This in turn implies that for  $t > T_1$ , the set of actions played by the parties in each iteration up to that point form an  $\epsilon$ -coarse correlated equilibrium. When there exists a coarse correlated equilibrium of  $T_1$  strategy sets (corresponding to the strategies of all players), if a centralized agent were to choose a strategy set, thus dictating a strategy for each player, then a player who believes that other players will follow the centralized agent’s instructions will be at least as well off committing to follow the centralized agent’s instructions (not knowing which strategy set will be chosen) than choosing any other strategy. In an



$\epsilon$ -coarse correlated equilibrium (a concept that we will formalize in the next section), a player can gain no more than  $\epsilon$  from choosing some strategy rather than committing to follow the one recommended by the centralized agent. This equilibrium concept is far weaker than that of Nash equilibrium. Note that it is weaker even than the concept of an  $\epsilon$ -correlated equilibrium, in which each player will agree to the centralized agent's instructions even after hearing which strategy is played.


Still, the "no regret" algorithm may perform better on some tasks than the  $\epsilon$ -coarse correlated equilibrium guarantee promises. After all, while plaintiffs and defendants are adversaries, if each is employing no regret learning simultaneously, neither is actively seeking to take advantage of the patterns that the other is employing. It is possible that as the defendant gradually converges to a strategy, the actions of the plaintiff that contributed to that convergence will be at least as much in the interests of the plaintiff as before. If the parties do converge on a Nash equilibrium, they may well continue to play the Nash equilibrium strategy. Indeed, one of the primary goals of this paper is to demonstrate that at least for some models of litigation games, "no regret" learning (as well as the more detailed counterfactual regret minimization algorithm to be described in the next section) does converge to Nash equilibria. It remains possible, of course, that different choices in modeling litigation might prevent a Nash equilibrium from being attained. Further work beyond this paper may help clarify the types of settlement bargaining modeling decisions that will lead to Nash equilibria, or to very strong approximate Nash equilibria with very low benefits from switching strategies.

To demonstrate this empirically, we return to the model offered by Dari-Mattiacci and Saraceno. We will start with a version that relaxes three of their assumptions. First, using no regret learning will inherently relax the assumption of piecewise linearity. Second, we will allow for the full range of values of  $q$ , representing the true merits and also the degree of information asymmetry. And third, we will allow for costs beyond the level where an equilibrium would be defined, stretching the costs axis all the way to 0.6 per party. We will execute separate simulations for various permutations of  $q$ ,  $c$ , and  $t$ .

If at the end of the simulation, the learning algorithm dictates that an action be played with greater than a 0.995 probability, we round this up, on the theory that the algorithm may be converging to playing that strategy with certainty. This is admittedly an ad hoc maneuver, but empirically, this improved performance. With such rounding, in each simulation, the results converged strongly to a Nash equilibrium. A best response function (to be discussed in the next section) was used to measure the maximum amount that a party could gain from switching strategies, and the amounts were generally in the range of  $1.0 \times 10^{-17}$  to  $1.0 \times 10^{-15}$ , essentially an accumulation of rounding errors in double precision numbers. In other words, a party could gain no more than a quadrillionth of the amount at stake by changing strategies. That the technique effectively identified an exact equilibrium here and also in the remaining games described in this section does not guarantee that it will do so with alternative means of modeling litigation, but there is also no guarantee that any particular proof

technique will succeed at finding an equilibrium. This finding establishes the central thesis of this paper, that computational models can at least sometimes identify Bayesian equilibria in settlement bargaining models.

The offers generated with this approach are illustrated in Figures 7-8. There is no reason to squint this time, as all one needs to appreciate is the overall shape of the curves. There are three graphs this time because the Dari-Mattiacci and Saraceno model results in different offer functions with different values of  $q$ . Thus, the graphs include a middling value of  $q$  (0.4), an extreme value of  $q$  (1.0), and a combination averaging the results of all  $q \in \{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$ . In each of the three panels, the large horizontal axis represents costs, and the vertical axis represents the fee shifting threshold  $t$ , with the English rule in the top row and the American rule in the bottom row. The blue line (in all cases here, the top line) is the plaintiff's offer curve, and the orange line is the defendant's offer curve. (For those unable to see the color, the defendant's line is always below the plaintiff's.) Note that in the middle panel, only the plaintiff's strategies are shown, because the defendant's signal has no bearing on the judgment and is entirely irrelevant.



Focusing on the left two panels, the results suggest that although piecewise linearity was a constraint in their original model, even with the restriction removed, strategies that are more-or-less piecewise linear often result. Indeed, it is striking that for any particular value of  $q$ , the offer curves include long flat sections (placing aside the leftmost column in each panel, where there are many Nash equilibria, as all cases go to trial with zero costs). The higher the costs, the longer the flat sections. This calculation of optimal offers thus provides general support for the proposition that the Dari-Mattiacci and Saraceno equilibria may be reasonably close to the equilibria that obtain in the absence of a piecewise linearity assumption.

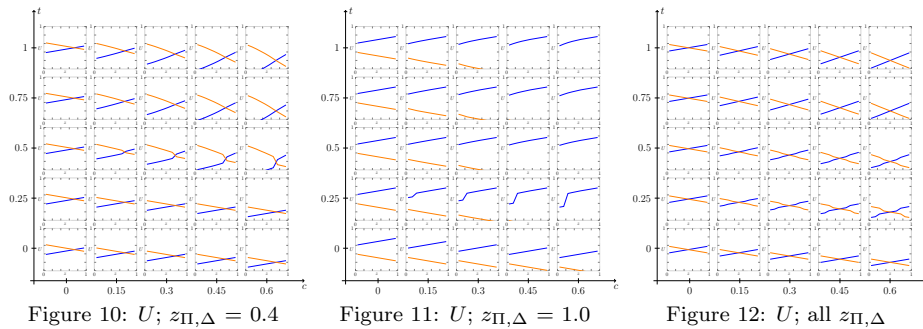
The “no regret” learning technique can also be used to generalize the Dari-Mattiacci and Saraceno model along the lines described above. Let us change the judgment  $J$  as follows:  $J = \alpha_Q q' + (1 - \alpha_Q) z_N$ , where  $q' = \alpha_\Pi z_\Pi + \alpha_\Delta z_\Delta + \alpha_{\Pi,\Delta} z_{\Pi,\Delta}$ , where  $\alpha_\Pi + \alpha_\Delta + \alpha_{\Pi,\Delta} = 1$  and where  $z_\Pi, z_\Delta, z_{\Pi,\Delta}, z_N \sim U$ . In this formula  $q'$  replaces  $q$  as a definition of the true merits, the sum of three components, one resulting from the plaintiff's private information, one from the defendant's private information, and one from information shared between the

parties (formerly labeled  $q$ ). Recall that in the original model, the parties know the true merits, and their signals provide them no evidence of the true merits. This generalization also allows for different degrees of information asymmetry, while decoupling information asymmetry from the true merits. Meanwhile, the  $z_N$  random variable represents random noise that is truly independent of the true merits. Dari-Mattiacci and Saraceno highlight that their model reflects an anti-winner bias, as the randomness in the component of the judgment attributable to the parties' normalized signals will on average regress judgments to the mean, and this noise term, of variable size, can serve a similar function.

We also generalize the fee-shifting rule. Let  $FS_{\Pi} = (\alpha_{\Pi,\Delta} z_{\Pi,\Delta} + \alpha_{\Pi} z_{\Pi}) / (\alpha_{\Pi,\Delta} + \alpha_{\Pi})$ , where  $\alpha_{\Pi,\Delta} + \alpha_{\Pi} > 0$ . This thus provides, consistent with Dari-Mattiacci and Saraceno's model, a measure of the quality of the plaintiff's evidence, and the plaintiff's fees will be shifted to the defendant for a given fee-shifting threshold  $t$  when  $J > \frac{1}{2}$  and  $FS_{\Pi} > 1 - t$ . We can define  $FS_{\Delta}$  analogously to consider the parties' shared evidence and the defendant's unique evidence, and fees will be shifted to the plaintiff when  $J < \frac{1}{2}$  and  $FS_{\Delta} < t$ . We will use this definition in this paper, though the computational model also supports calculating fee shifting in other ways, such as based on the margin of victory.

With this approach, we can assess whether Dari-Mattiacci and Saraceno's findings on accuracy generalize. Rather than offer plots illustrating accuracy immediately, we will make the ultimate charts more comprehensible by offering a series of charts. We will start by showing how the parties' utilities change in an alteration of the model in which all cases go to trial because settlement is disallowed. This can help make sense of the parties' offer functions, and that in turn can help clarify how trial rates vary at different parts of the continuum of true merits cases. From there, it is a short step to assessing accuracy. This journey thus analyzes the selection of disputes as part of its findings of accuracy, thus uniting concerns that ordinarily would require an entirely separate model. Except where stated otherwise, we will use  $\alpha_Q = 0.75$ , so that a quarter of the judgment will be the result of pure noise;  $\alpha_{\Pi,\Delta} = 0.5$ , so that half of the total information is shared; and  $\alpha_{\Pi} = \alpha_{\Delta} = 0.25$ , so that asymmetric information does not systematically favor either party.

Figures 10-12 show the utility that parties will receive when trial is guaranteed. On the left panel, note that there are kinks in the utility functions only where  $t = 0.5$ . At  $z_{\Pi,\Delta} = 0.4$ , fee shifting will never occur when  $t = 0$  or even  $t = 0.25$  (because the highest possible value of  $FS_{\Pi}$  is then 0.6, which is less than 0.75). Fee shifting can occur when  $t = 0.5$  if  $z_{\Pi} \geq 0.7$ , and this is where the plaintiff's offer curve is kinked. Fee shifting will always occur at  $t = 0.75$  and  $t = 1$ , so there are no kinks in any other row. Similar reasoning shows why there are kinks only at  $t = 0.25$  in the middle panel when the parties' shared information is as good as possible for the plaintiff, and the right panel combining results where  $z_{\Pi,\Delta} \in \{0, 0.2, 0.4, 0.6, 0.8, 1\}$  includes wavy lines that partially smooth out various kinks. In each panel, the primary effect of trial cost is to reduce parties' utilities; with high fee shifting, it is the party that has the weaker case whose utility rapidly declines, as this party can expect to bear the brunt of the trial cost.



Consider now Figures 13-15, which illustrate the parties offers as a function of their signals, much as Figures 7-9 did for the original Dari-Mattiacchi and Saraceno model. In each mini-graph, the plaintiff's offer curve (in blue) is either over or overlapping the defendant's. As in the original model, in the left column of each panel, where trial is costless, all cases go to trial and thus the particular offer functions are essentially arbitrary. With  $c = 0.15$ , there are dramatic jump discontinuities. In the left panel, for low fee shifting, the jump occurs for the plaintiff only where the plaintiff receives a very strong signal, but with fee shifting, both parties' offers jump where private information suggests that the plaintiff may be able to overcome the deficit indicated by the shared information, thus leading to the possibility of fee shifting to the plaintiff or at least a reduced possibility of fee shifting to the defendant. In the middle panel, where the shared information is all in the plaintiff's favor, there are still discontinuities, but not all in the middle of the probability spectrum. Meanwhile, with high costs, the flat portions of the offer curves dominate and overlap. Where the shared information is all in the plaintiff's favor and  $c \geq 0.3$ , each graph represents a pooling equilibrium, because avoiding trial is in the parties' mutual interest and the risk from placing too much weight on private information when the direction of any fee shifting is fairly clear becomes great. Higher fee shifting in these cases simply shifts up the offer curves, and where fee shifting and costs are maximal, the defendant always concedes full liability. When we combine a range of  $z_{\Pi,\Delta}$  values in the right panel, we see each party with a region of poor signals in which the party's offer is flat (effectively screening on the other party's type), and a region of strong signals in which higher signals make the party more aggressive. Fee shifting tends to make this region of ascending signals longer, at least if costs are relatively low. For each level of fee shifting, high costs tend to reduce the impact of the signal, flattening the offer curves.

Given these offers, the data on trial rates, represented by Dari-Mattiacchi and Saraceno's notation  $L$ , are unsurprising. Figures 16-18 show trial rates given both the plaintiff's (blue) and the defendant's (orange) signals. In regions where plaintiff offers increase, trial rates increase as well, and where defendant offers increase, trial rates fall. Some of the changes in trial rates are very steep. As the plaintiff becomes sufficiently confident that it will benefit from fee shifting, the plaintiff becomes willing to accept a much higher rate of trial than in the




Figure 13:  $\phi$ ;  $z_{\Pi,\Delta} = 0.4$




Figure 14:  $\phi$ ;  $z_{\Pi,\Delta} = 1.0$




Figure 15:  $\phi$ ; all  $z_{\Pi,\Delta}$

absence of fee shifting. With very high costs, however, trial is relatively rare, though it still occurs with some probability if there is fee shifting and parties receive strong offers.




Figure 16:  $L$ ;  $z_{\Pi,\Delta} = 0.4$




Figure 17:  $L$ ;  $z_{\Pi,\Delta} = 1.0$





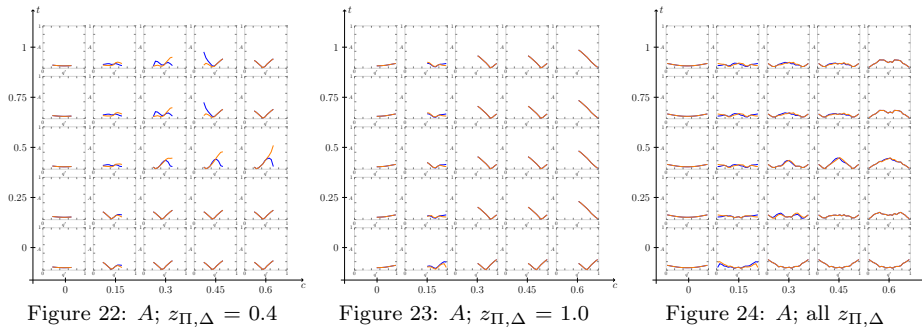
Figure 18:  $L$ ; all  $z_{\Pi,\Delta}$

In turn, Figures 16-18 enable understanding of Figures 19-21. In these latter figures, the horizontal axis in each mini-graph represents  $q'$ , the true merits of the case. For any particular value of  $q$ , the true merits may only be defined over a narrow region; if, for example,  $z_{\Pi,\Delta} = 1.0$ , then  $q' \geq 0.50$ . The overall selection pattern here is quite complex, with considerable changes in shape based on changes in  $q$ ,  $c$ , and  $t$ . But an examination of the corresponding figure for trial rates makes each individual pattern comprehensible. For example, consider for the left panel the mini-graph with  $c = 0.15$  and  $t = 1$ . This dip in trial rates corresponds to the signal region where the trial rates fall, which correspond to the region in Figure 13 where the parties' offers rise steeply. Focusing on the rightmost panel, these results suggest that a wide variety of selection patterns are possible, including the familiar Priest-Klein pattern in which the cases that go to trial are clustered in the middle of the true merits spectrum, the competing pattern predicted by Bebchuk in which extreme cases go to trial, and also bimodal patterns in which there are either two valleys or two hills defining the trial rate. Consistent with Friedman and Wittman, there appears to be a tipping point based on the level of costs.

Finally, we can now consider some measures of inaccuracy. Figures 22-24 plot in each minigraph the level of inaccuracy against the true merits. Two



measures of inaccuracy are used, one from the perspective of the plaintiff and one from the perspective of the defendant. Thus,  $A_{\Pi} = |q' - U_{\Pi}|$ , the absolute difference between the true merits and the net amount the plaintiff recovers, and similarly,  $A_{\Delta} = |1 - q' - U_{\Delta}|$ . Lower curves thus correspond to greater accuracy. As before, the large horizontal axis is for costs, and the large vertical axis is for fee shifting. Once again, blue is for the plaintiff and orange is for the defendant, but the curves are heavily overlapping in any event. Once again, the previous diagrams can help make sense of these. For example, consider the panel for  $z_{\Pi,\Delta} = 0.4$ ,  $c = 0.45$ , and  $t = 0.5$ . Note the anomaly that the defendant's inaccuracy increases rapidly as  $q'$  becomes sufficiently high, while the plaintiff's inaccuracy shrinks in the same region. This is because there is a point at which the plaintiff's signal becomes high enough that the plaintiff's offer spikes (in the corresponding mini-graph of Figure 13), causing trial rates to go from 0 to 1 (in the corresponding mini-graphs of Figure 16 and 19). The plaintiff performs well enough so that fee shifting will generally occur, in turn placing the burden of accuracy attributable to fees on the defendant rather than on the plaintiff.



In this generalization of the model with revised definitions of accuracy, Dari-Mattiacci and Saraceno's accuracy claims still retain some of their power, though the pattern is sufficiently complex that generalization is difficult. Recall that in their model, with sufficiently low costs, fee shifting will improve accuracy, while for sufficiently high costs, fee shifting will reduce accuracy. Certainly at the last column of the third panel, where costs are high, fee shifting appears to increase

inaccuracy in the broad middle of the true merits spectrum. These cases are all settled but with a pooling equilibrium that gives everyone the same result. One might think that would lead to good accuracy for cases in the middle of the probability spectrum; after all, Figure 15 appears to show that these cases receive middling settlement amounts. This, however, is a misimpression that arises from averaging results for various levels of  $q$  in that right panel. Recall that in fact with high costs and fee shifting, the offers will tend to settle at extreme values, such as 0 where  $z_{\Pi,\Delta} = 0$  and 1 where  $z_{\Pi,\Delta} = 1$ . The parties' private information is thus ignored, and inaccuracy is high. On the extreme ends of the  $q'$  true merits spectrum, accuracy is much better, because for these values to occur, the private information in fact must be consistent with the shared information. But the regions with high accuracy are shorter, and so it is fair to say that inaccuracy is high overall. With high costs but no fee shifting, the shape is inverted, but overall accuracy appears marginally better. Meanwhile, when costs are sufficiently low, the case is unclear, though accuracy appears to be slightly better without fee shifting for cases in the middle range.

The pattern of how accuracy changes with costs and fee-shifting is complex, and further analysis is warranted, including for example measuring average inaccuracy with various definitions and examining alternative definitions of fee shifting, as well as the possibilities of partial or multiple fee-shifting. For example, Figure 25 demonstrates how inaccuracy changes with shifts in the amount of information that is shared between the parties. The red single-dashed line is the case in which the parties have no shared information; the orange double-dashed line, 25% of the total information is shared; yellow triple-dashed, 50%; and green, 75%. As one would expect, more information sharing generally promotes accuracy. But this is not clearly true in some circumstances, such as where costs and fee-shifting are both high. A paper focusing on accuracy might explore these patterns in more detail. The goal of this paper is simply to demonstrate how a computational model, in this case using no regret learning, can allow for principled inference about the possible effects of changing a mathematical model.

### 3.3 Counterfactual Regret Minimization

The previous section illustrated that “no regret” learning algorithms allow the parties to optimize their decisions by repeatedly playing the modeled game against one another, assigning a probability distribution to all actions available at each decision point. A limitation is that the decision points, as described so far, are all decisions that each party makes at the beginning of the game after receiving its signal. Indeed, the game requires a single action for each player. The probability distributions that one player chooses will affect the other player's utilities and thus that player's mixed probabilities in later iterations. But no player's action affects whether another player has a chance to choose an action in the first place, or even whether the first player will need to make another decision. In a more complex game, we must confront this difficulty. For example, a decision by a plaintiff not to file a lawsuit means that the defendant




Figure 25: Inaccuracy with different levels of information sharing

will not face any decisions at all. More importantly, assuming that the plaintiff plays a mixed strategy in which the plaintiff sometimes files and sometimes does not, the distribution of cases that face the defendant will depend on the mixed strategy that the plaintiff chooses. Because “no regret” learning is robust to adversaries, this is not inherently a problem. At each iteration, the defendant’s utilities simply must be calculated based on the distribution of cases filed as a result of the plaintiff’s strategy.

There is, however, a more complex type of dependency between players’ actions. One or more pieces of information in a player’s “information set” (defined below) may be information about the choices that a player made earlier in the game. Assuming that a defendant makes a settlement offer only when the plaintiff files a case, it is meaningless to suggest that the fact of filing is information within the defendant’s information set, because the plaintiff will



always have filed when the defendant makes a settlement offer. But suppose that bargaining fails, and each party must decide whether to proceed to trial or to concede defeat (nonprosecution by the plaintiff or acceptance of default by the defendant). That party will want to consider all available information in assessing the plaintiff's probability of victory, including the offer that the other party made in settlement negotiations. The opponent's offer will be part of a litigant's information set when deciding whether to continue to trial or concede. Thus, "no regret" learning must apply separately at each information set.

More importantly, the game must be played for every permutation of possible moves, including moves by both Chance and the players. Any given information set may thus be reached through multiple paths in the same iteration. Suppose the defendant's information set following bargaining failure contains a quality signal (say, 2 of a maximum 5) and the plaintiff's offer level (say, 4 of 5 offer levels). This information set is consistent with different possible signals that the plaintiff received. It may be that for some such signals, given the defendant's strategy and information set, the game would have settled with 100% probability, so those situations should receive no weight in formulating the defendant's response. But it may be that if the plaintiff received a quality signal of 3, 4, or 5, there was still some positive probability that settlement would not occur and that the defendant would face the default decision. Thus, the defendant must determine what weight to assign to the utilities that the defendant receives with each of these quality signals that the plaintiff may have received. The resolution to this conundrum is intuitive, albeit with a small twist: the weight should be the probability of arriving at the information set through a particular path, but ignoring the defendant's own contributions to that probability. That is, if Chance makes a particular path unlikely (by assigning a signal a low probability), then the utilities received from that path to the information set should receive a low weight, and if a plaintiff will file a lawsuit given that signal with only a low probability, then the weight should be lower still. But if the defendant had an earlier opportunity to default, it makes no difference what the probability was that the defendant would do so.

This is the essence of counterfactual regret minimization, an algorithm developed by Zinkevich et al. (2008) [25]. We will provide some formalism here. Let  $A$  represent the set of all actions that players, including the chance player, may take. A game history consists of a sequence of actions from the beginning of the game. Let  $H$  represent the set of all possible game histories, and let  $Z \subset H$  represent the set of all possible completed game histories, including a separate entry for each sequence of chance and player moves leading to the resolution of the litigation. An information set consists of a subset of the actions in a game history, representing the actions of which a player has knowledge at the corresponding point in the game. Let  $I$  represent a particular information set, and let  $Z_I$  designate the set of all possible game histories in which the player to whom the information set belongs chooses an action at that information set. Let  $z_I \in Z_I$  represent a sequence of actions (whether or not known to the player to whom  $I$  belongs) that includes an opportunity to make a decision at that information set. We restrict our attention to games with perfect recall (i.e., a player

remembers each move she has made), so a player can encounter a particular information set no more than once in any game history.

Then, let  $\sigma$  represent the current strategy profile, i.e. the set of mixed strategies that each player currently plays at each information set. Let  $\pi^\sigma(z_I)$  represent the probability that the players, including Chance, will make the set of moves in the history  $z_I$  leading up to information set  $I$ ; the probability of moves at and after information set  $I$  are thus ignored. For player  $\phi$ , let  $\pi_\phi^\sigma(z_I)$  represent player  $\phi$ 's own contribution to this probability, i.e. the product of the probabilities of actions that player  $\phi$  must choose to arrive at information set  $I$  along history  $z_I$ . Then, we can define  $\pi_{-\phi}^\sigma(z_I) = \pi^\sigma(z_I)/\pi_\phi^\sigma(z_I)$ , i.e. the probability that other players, including Chance, play to the information set. For any terminal history  $z \in Z_I$ , let  $\pi^\sigma(z_I, z)$  represent the conditional probability, given that player  $\phi$  reaches information set  $I$  at history  $z_I$ , that the completed game is  $z$ , i.e. the probability of the actions at and after information set  $I$  by all players along terminal history  $z$ . Finally, let  $U_\phi(z)$  represent the utility that player  $\phi$  receives at this terminal history.

With these definitions, we can formally define the counterfactual value  $v_\phi$  for player  $\phi$  at information set  $I$  as

$$v_\phi(\sigma, I) = \sum_{z \in Z_I} \pi_{-\phi}^\sigma(z_I) \pi^\sigma(z_I, z) U_\phi(z) \quad (4)$$

Reading the equation from right to left after the summation, the counterfactual value averages utilities that the player earns at terminal histories stemming from the information set by the probabilities that the players will play from the information set to each terminal history and the probabilities that the other players will have played to the information set. Two-player poker may provide some useful intuition. Suppose that a player is choosing what to bet. The player's information set includes all of the player's cards, plus all the other player's cards that are face up at the table, plus all betting moves previously made by the player or the player's opponent. But there are many possible permutations of face down cards that the other player may have, and for each of these permutations, there may be many different ways that the game can play out from here. Counterfactual regret minimization requires weighting the different possible permutations of face down cards by the probability that those cards could have been dealt and that the opponent would have taken any actions that the opponent took with those face down cards. For example, for each possible set of face down hands, there is some probability that the opponent would have made a large bet (either because the face down cards were good for the opponent or because the opponent was bluffing). Since each player knows its opponent's strategy, these calculations can be made explicitly. Similarly, the player can compute, for each permutation of face down cards, every permutation of how the game might proceed. This highlights why asymmetric information games like poker are difficult to model: the player must not only consider all possible future game paths but also all past game paths that might have led the player to the current point.

The counterfactual value  $v_\phi(\sigma, I)$  is analogous to the  $U_a^t$  in the discussion of no regret learning in the previous section. Thus, we define the regret value  $r^t(I, a) = v_\phi(\sigma_{I \rightarrow a}^t, I) - v_\phi(\sigma^t, I)$ , where the notation  $\sigma_{I \rightarrow a}^t$  represents a strategy identical to  $\sigma^t$  except that player  $\phi$ , to whom information set  $I$  belongs, chooses action  $a$  at that information set with probability 1. Then, the cumulative counterfactual regret at information set  $I$  for action  $a$  through iteration  $T$  can be defined as  $R^T(I, a) = \sum_{t=1}^T r^t(I, a)$ . “No regret” learning can then be applied, for example as before in the form of regret matching:

$$\sigma_\phi^{t+1}(I, a) = \begin{cases} \frac{R^{t,+}(I, a)}{\sum_{b \in A_\phi(I)} R^{t,+}(I, b)}, & \text{if } \sum_{b \in A_\phi(I)} R^{t,+}(I, b) > 0 \\ \frac{1}{|A_\phi(I)|}, & \text{if } \sum_{b \in A_\phi(I)} R^{t,+}(I, b) = 0 \end{cases} \quad (5)$$

The counterfactual regret minimization algorithm proceeds by walking the entire game tree once at each iteration  $t$  for each player whose actions are to be optimized, in this context the plaintiff and the defendant. Utilities are recorded on the completion of the walk back from the tree.

There are many variations on the vanilla version of the counterfactual regret minimization algorithm, designed to minimize the need to play information sets that the players are highly unlikely to reach. See Lanctot (2013) [14] for detailed algorithms. For very large games, a technique that integrates counterfactual regret minimization with neural networks can allow for extrapolation. See Steinberger (2019) [24]. We will not need to employ these here, as the settlement bargaining games in this introductory article are manageable with full walks through the tree.

We will, however, adopt two variations to the standard vanilla algorithm. In experimentation, both of these proved helpful in generating better approximated equilibria in the general sum litigation game. The first is an innovation offered by Farina et al. (2017), who suggest defining a perturbation  $\mathcal{P}_T$  in iteration  $T$  where  $0 < \mathcal{P}_T < 1$ , requiring each action to be played with at least probability  $\mathcal{P}_T$ . Farina et al. find that this improves strategy refinement. Even where a strategy is a Nash equilibrium, some information sets may not be reached at all, in which case the player blindly plays each action with equal probability. The strategy thus may be susceptible to trembles, slight deviations by an opponent. The perturbation ensures that every information set is visited, and thus we use a small vanishing perturbation term, falling from  $\frac{1}{1,000}$  in the first iteration to 0 in the last iteration.

The second is a variation of an approach offered by Brown and Sandholm (2018) [2]. They describe the approach of discounting regrets, so that recent regrets count more in defining a player’s current strategy than regrets in early iterations. They show that such discounting can have a dramatic effect on performance. With settlement bargaining, for example, if in early iterations, a defendant would offer to pay the maximum amount even in a very weak case, without discounting it will take a long time before this effect is diluted. The choice to use this technique, like the choice to employ a perturbation, was based on informal experimentation showing better convergence with discounting. The

particular discounting algorithm employed here is simple: All previous regrets are multiplied by 0.99 after each iteration.

Brown and Sandholm provide convergence guarantees in zero-sum games, but these do not apply here, and as noted above, the counterfactual regret minimization algorithm does not offer strong convergence guarantees. However, it retains the  $\epsilon$ -coarse correlated equilibrium guarantee from no regret learning. With the notation developed here, we may formalize this. For any  $\epsilon > 0$ , there exists a  $T_0$  such that for  $T > T_0$ ,

$$\frac{1}{T} \sum_{t=1}^T U_\phi(\sigma_\phi^t, \sigma_{-\phi}^t) + \epsilon \geq \frac{1}{T} \sum_{t=1}^T U(\sigma'_\phi, \sigma_{-\phi}^t), \forall \sigma'_\phi \in \Sigma \quad (6)$$

where  $\Sigma$  represents the set of all possible strategies. Once again, however, we will see that although the technique we use provides no guarantee of converging to an approximate Nash equilibrium, the algorithm will produce strategies quite close to equilibria for strategies that we test here. As before, we can measure how close a strategy is to a Nash equilibrium by determining a best response to the strategy for each player. An algorithm for determining a best response is discussed in Johanson et al. (2011). This algorithm proceeds by walking forward through the game tree to calculate  $\pi_{-\phi}^\sigma$ , and then walking backward from the leaves of the game tree to choose at each information set  $I$  the action producing the highest utility given the distribution of prior histories, i.e.  $\arg \min_{a \in A(I)} \sum_{z \in Z_I | a \in z} u(z) \pi_{-\phi}^\sigma(z) \pi^\sigma(ha, z)$ , where  $\pi^\sigma(ha, z)$  represents the probability that the players will play from  $I$  to  $z$  where possible, excluding the probability of choosing  $a$  at information set  $I$ . This expression accounts for the relative probabilities of different histories leading to the information set, as well as later information sets chosen earlier in the algorithm. Given a calculated best response, one can compute the maximum amount that either player could gain by changing strategies.

With counterfactual regret maximization thus defined, we can now define the litigation game that will extend the model of Klerman, Lee, and Liu. We will start with a simple version and add additional features afterward. We assume that there are two types of cases, those in which the defendant is not truly liable ( $TL = 0$ ) and those in which the defendant is truly liable ( $TL = 1$ ). A case may have one of  $n_{LQ}$  quality strengths with respect to liability. When the defendant is truly liable, the case is more likely to be a strong one for the plaintiff, but there may be some cases of true liability in which the case is weak, and some cases in which the defendant is not truly liable but the case is nonetheless strong. The distribution of the various case strengths depends on standard deviation parameters  $\sigma_{LQ}$  and  $\sigma'_{LQ}$ , where  $\sigma_{LQ}$  indicates how accurately litigants can gauge litigation strength and  $\sigma'_{LQ}$  indicates how closely litigation quality reflects the true merits. If the legal rule were that the party first alphabetically wins, then  $\sigma_{LQ}$  would be low, because the rule is easy to apply, but  $\sigma'_{LQ}$  would be high, because the rule bears virtually no connection to true liability.

Specifically, a mixture distribution is defined by aggregating in equal measure two normal distributions with standard deviation  $\sigma'_{LQ}$ , one with mean 0

(representing not truly liable cases) and one with mean 1 (representing truly liable cases). The values in this mixture distribution are then partitioned into  $n_{LQ}$  litigation quality ranges that collectively span the real line. If  $\sigma'_{LQ}$  is sufficiently small, then each litigation quality range will consist virtually entirely of a single true liability value, while if  $\sigma'_{LQ}$  is sufficiently large, then each litigation quality range will consist of both true liability values in almost equal measure. Given a particular true liability value (either 0 or 1), the probability that the case is of litigation quality  $i$  is the probability that the sum of the true liability value (0 or 1) and a draw from a normal distribution with standard deviation  $\sigma_{LQ}$  lies in the  $i$ th litigation quality range.

The plaintiff and the defendant each receive a signal of liability quality strength, specifically one of  $n_{sL}$  signals. The approach is akin to that for determining the probability of litigation quality based on truly liable values, with the mixture distribution aggregating distributions for each litigation quality value. The signal ranges, however, are defined so that if each liability quality strength were of equal probability, there would be an equal probability of obtaining each signal. Given an actual liability quality value (determined as above), the probability of signal  $i$  is the probability that the sum of that value and a draw from a normal distribution with standard deviation  $\sigma_{PL}$  for the plaintiff or  $\sigma_{DL}$  for the defendant lies in the  $i$ th signal range.

The parties may attempt to settle using Chatterjee-Samuelson bargaining, and if settlement fails, a judge resolves the case. The judge will rule for the plaintiff if the case is in the stronger half of case strengths or for the defendant if the case is in the weaker half of case strengths.

Initially, let us start with  $n_{LQ} = 2$ , mapping all truly liable cases to the stronger point and all not truly liable cases to the weaker point. This approach will allow us to begin with an approach truer to the Klerman, Lee, and Liu approach, before we begin to extend it. Meanwhile, we set  $n_{sL} = 20$ ,  $\sigma_{LQ} = 0.4$ ,  $\sigma'_{LQ} = 0.5$ ,  $\sigma_{PL} = 0.1$ , and  $\sigma_{DL} = 0.1$ . The last two numbers are more consequential than the others and represent relatively small values for the noise. Recall that the 50% hypothesis applies only in the limit as error becomes close to 0. If we make the noise sufficiently small in a discrete model, then both litigants will receive the same signal every time. This reveals a limitation of computational models, but it also reveals the artificiality of the hypothesis, since as the error rate goes to zero, so does the trial rate. Of course, the point is to suggest that win rates should be close to 50% when error is relatively low, and so we use a somewhat arbitrarily set relatively low error, embodied by  $\sigma_{PL}$  and  $\sigma_{DL}$ .

We also set various parameters of the litigation game. We normalize damages to 1 as before, and we set the parties' trial costs  $C_P = C_D = 0.15$ . We also impose on each party a fee of 0.05 for contesting a case; that will become relevant later when the players are given an option to concede the litigation. We ran a separate simulation for each value of  $p_{TL}$  in  $\{0.05, 0.10, \dots, 0.95\}$ , where  $p_{TL}$  represents the probability that the defendant is truly liable. As in the "no regret" learning simulations, each simulation strongly converged, with the calculated best response giving neither party more than approximately a  $1.0 \times$

$10^{-15}$  increase in utility from changing strategies. Indeed, in some cases, the measured benefit from changing strategies to the best response was negative, reflecting tiny rounding errors either in the learning function or in the best response.

The results are illustrated in Figure 26. The vertical axis indicates the plaintiff's win rate, and the horizontal axis the exogenous probability that the plaintiff is truly liable, which in turn affects the distribution of liability strength. The results hover around 0.5, though deviating a significant distance. The deviations may be due to discreteness, as a result of rounding to particular discrete signal values. (We ran several simulations not pictured here with larger and smaller values for  $n_{sL}$  and found that the values were generally closer to 0.5 with higher values of  $n_{sL}$ .) These errors are larger near the ends of the probability continuum, systematically erring in the direction of a higher win rate for the party less likely exogenously to have a strong case. This may well be a result of having a finite number of signals, in contrast to the Klerman, Lee, and Liu model, which has no worst signal. Meanwhile, the dashed red line indicates the trial rate using the same axis scale. The cases in the middle of the probability spectrum are more likely to go to trial (though still quite unlikely to do so). The win rate seems to deviate from 50% precisely where trial becomes extraordinarily rare, where there are very few tried cases, magnifying the effect of rounding signals to fit into discrete categories. Not visible from the graph is another important observation: in each of the simulations, although the simulation allowed for strategies with monotonically increasing strategies, instead each party's offer function was flat except at a single jump discontinuity. This is consistent with the Klerman, Lee, and Liu model.




Figure 26: Win (and trial) rates for exogenous probability truly liable

We can now assess the effect of adding a greater number of liability strength points, setting  $n_{LQ} = 20$ , instead of  $n_{LQ} = 2$ . This will help address the question whether Klerman, Lee, and Liu's results are driven by the existence of just two strength points. Figure 27 illustrates the results, providing some support for the proposition that Klerman, Lee, and Liu's limitation to two liability strength points is innocuous. The win rate line is in fact flatter than before, hovering closer to 0.5. More liability strength points appear to smooth out the inaccuracies associated with discontinuities. This provides support for

the proposition that even where there are many case strength levels, the plaintiff will win approximately 50% of cases regardless of the distribution of truly liable cases. In addition, once again in each simulation each party's offer function was flat with a single jump discontinuity. The discontinuity point varied based on the signals received, but the wider range of case types did not lead to an offer function with more than two points. This is not proof that Klerman, Lee, and Liu's results extend to a model with many or infinite liability strength points, but it suggests that the assumption of only two such points is not likely to be the problem. The example highlights that discrete computational models may be less useful when the goal is to prove an exact result (such as a plaintiff win rate of exactly a particular percentage), but that even in such circumstances they still may help rule out the possibility that results would change drastically with a more complex model. The example also may give some encouragement to a modeler interested in taking on the challenge of extending Klerman, Lee, and Liu's results to a larger or infinite number of case strengths.




Figure 27: Multiple strength points

We will now assess the impact of providing the players options. The sequencing is as follows: First, the plaintiff must decide whether to file a case. If the plaintiff does not file, the plaintiff receives 0 damages and no fees are assessed. Second, if the plaintiff does file, then the defendant must decide whether to answer. If the defendant chooses not to answer, the defendant receives utility  $-1$  and the plaintiff receives utility  $+1$ . Third, if the defendant answers, then each party provides a settlement offer subject to the Chatterjee-Samuelson mechanism. If settlement occurs, then the defendant pays that amount of money to the plaintiff, and each side bears its filing costs. Fourth, if settlement does not occur, the plaintiff and the defendant simultaneously make choices about whether to concede the litigation. Fifth, if one party concedes, the lawsuit ends, and either the plaintiff receives nothing or the defendant pays full damages to the plaintiff, while each side bears its costs. If both parties concede, then Chance chooses one party to concede to the other with equal probability. Sixth, if settlement does not occur and neither party concedes, a trial is held. The judge estimates the case liability quality with perfect precision. If the case quality is in the lower half of case qualities, then the defendant wins; otherwise, the plaintiff wins. Either way, each side bears its own costs.

We set each party's filing costs to 0.3 for this simulation, because with the original value of 0.05, neither party ever conceded. The value of 0.3 is high but might be plausible in a case with low stakes, such as a small claims case. The simulations did not converge as strongly as in all the earlier simulations in the paper. By switching to a best response, each side would be able to gain approximately 0.002. For example, with the exogenous probability set to 0.5, the plaintiff would be able to gain 0.0024 by switching to a best response, and the defendant would be able to gain 0.00097. As this suggests, the greater interactivity among decisions makes full convergence more difficult in the settlement bargaining game. Increasing the number of iterations, even into the millions, cannot cause arbitrarily good convergence, even in average strategies, as occurs in a zero-sum game like poker. Still, these are fairly strong equilibria for practical purposes. A litigant reliably told that a lawyer was good enough at settlement negotiations to be within 0.2% of the best possible negotiator would think that fairly impressive. Humans are likely to be considerably further from the equilibrium than that.

The results are illustrated in Figure 28. While there is still some zig-zagging resulting from the discreteness of the model (particularly at the endpoints), a relatively clear trend emerges. Contrary to the results above, it indicates that the Priest-Klein 50%-limit hypothesis does not apply in the model with options. When the exogenous probability of true liability falls below 0.5, the plaintiff win rate drops, and when the exogenous probability of true liability rises above 0.5, the plaintiff win rate rises.

The computational model points in a potentially promising area for future research, and close examination provides at least a suggestive explanation of this exception to Priest-Klein. When  $p_{TL} = 0.2$ , the plaintiff files a complaint in 32.3% of all possible cases, and the defendant answers in 11.8% of all possible cases (about one-third of the cases that the plaintiff files). The asymmetry exists because the fact that the plaintiff files provides the defendant information that the plaintiff's information suggests that the plaintiff has a strong case. Thus, the defendant continues to settlement negotiations in fewer cases than if the defendant had to make its decision simultaneous with the plaintiff. (The defendant also gives up at higher rates after the failure of settlement negotiations, quitting 1.0% of the time, compared to the plaintiff's 0.4%.) The recognition that the defendant will be prone to give up reasonably often may make the plaintiff even more likely to pursue cases. The result is that the plaintiff will have lower quality cases, on average, producing win rates lower than 50%.

A slightly different story can be told at the other side of the probability spectrum. With  $p_{TL} = 0.8$ , the plaintiff files 85.0% of the time, and the defendant answers 12.4%. On the one hand, the fact that the plaintiff has filed is a less powerful signal with a distribution that includes more favorable cases for the plaintiff, but on the other hand, more such cases make litigation less attractive for the defendant. The parties quit after settlement bargaining failure at very low rates. Because the defendant is choosing approximately the same proportion of cases to fight as before but with a less powerful signal from the plaintiff, the cases on average that the defendant contests will be less strong. This con-



sideration is powerful enough to produce win rates greater than 50%. At least, the data suggest that this is part of the story. A more complete account would assess the impact of various changes in information structure and the litigation game rules. A particular question of interest is what the outcome would be if the plaintiff and defendant made their decisions whether to contest litigation simultaneously. This would help clarify whether it is only timing driving the results or whether additional factors are at work.




Figure 28: Options to concede

Our last question is whether uncertainty about damages affects win rates. We proceed similarly as before. We allow for  $n_{DQ} = 6$  damages quality levels, drawn evenly from the unit uniform distribution. We assume that these are assigned to each case with equal probability. Then, each player receives one of  $n_{sD} = 6$  signals of damages quality. We also set  $n_{LQ} = 6$  and  $n_{sL} = 6$ , a lower number from before because the algorithm must now run through 36 damages signal-quality combinations for each liability signal-quality combination. The damages signal is determined by adding the damages quality and a draw from a normal distribution, with standard deviation  $\sigma_{PD} = 0.3$  for the plaintiff and  $\sigma_{DD} = 0.3$  for the defendant. Note that we use a higher noise for damages than for liability, because the claim that 50% win rates occur only with small error applies to error about liability, not error about damages. The signal that each player receives is determined by assigning each signal to a number evenly drawn from the unit uniform distribution and identifying the value closest to the sum of the actual damages quality and the noise term.

In these simulations, for analytic tractability in the small space remaining, we returned to litigants with no option to concede the litigation (though the computational model can easily incorporate damages uncertainty and options to concede). Figure 29 illustrates the perhaps surprising results. One important result is that trial rates are much higher, especially when the distribution of cases is balanced with respect to liability, approaching 0.4. When the parties have two types of uncertainty, each litigant must worry that if the opponent has bad information on either liability or damages, the opponent will be especially willing to settle. This lemons problem may make both parties relatively hesitant to settle. This suggests a point about selection that appears to be overlooked in the literature: Cases in which there are two forms of uncertainty may be much

more likely to go to trial than cases in which there is only one issue. In addition, where the overall distribution of cases aggregates a heterogeneous collection of case types, some drawn from distributions featuring high probabilities of liability and some from distributions featuring low probabilities of liability, the cases with middling probabilities may be most likely to settle, producing a tendency toward 50% win rates.

The final observation may appear paradoxical: When the exogenous probability of true liability is high and the distribution thus includes many strong cases on liability for the plaintiff, the plaintiff's win rate is very low, and when there are few strong liability cases, the plaintiff's win rate is very high. When the population includes mostly strong cases, plaintiffs appear to be very choosy about which cases they are willing to settle, preferring to increase their share of the settlement surplus in cases that settle at the expense of losing at trial most of the few cases in which the defendant's signals suggest that the result is not as favorable to the plaintiff as first appears. The reverse is true on the other side of the probability continuum. This is not the place to explore this finding fully. The results reinforce the computational model's flexibility, highlighting opposite effects with different changes in assumptions, in each case finding strong approximations to Bayesian equilibria.




Figure 29: Uncertainty in damages as well as liability

## 4 Conclusion

The case for computational modeling in some area of inquiry becomes relatively strong when mathematical modeling encounters severe limitations and when computational algorithms can achieve the same goals as mathematical modeling. Settlement bargaining is difficult to model mathematically in a rigorously Bayesian way with two-sided asymmetric information. The literature has made significant strides in doing so, tackling even issues such as fee shifting and the selection of disputes for litigation with fully Bayesian models that incorporate two-sided asymmetric information. The mathematical literature may well make continued progress, either with articles featuring much longer proofs or perhaps with the development of new mathematical modeling techniques that simplify analysis. At least for now, however, the mathematical modeler of settlement

bargaining faces severe trade-offs. The empirical results of the computational extensions of existing models establishes that in at least some cases, computational modeling can identify exact solutions to the settlement models, even when adding additional features that mathematical modelers would not easily be able to incorporate. Computational modelers cannot prove that particular relationships exist among variables, but they can explore many different parameter values. Computation should not replace mathematical models, but where computational techniques are available, the combination of mathematical proofs involving simplified assumptions and computational extensions may be more powerful than proofs alone.

Even the computational models described here are quite limited, however, in what they can accomplish, at least in comparison to the complexity of adjudication. A challenge for future research is to extend the algorithmic game theory techniques used here to larger games. A significant challenge is that of multiple equilibria. When the computational modeling techniques described here succeed in identifying an equilibrium, they provide no way of assessing whether other equilibria exist. Mathematical modelers also sometimes prove the existence of an equilibrium without establishing its uniqueness, so the flaw is not fatal. But it may be useful for computational techniques to be able to identify families of equilibria rather than just individual equilibria. This may be particularly important with more complex games, for example involving multiple rounds of bargaining, gradual learning by the players of relevant evidence, and endogenous disputes within models of particular areas of substantive law, where multiple equilibria may be more likely to exist. In the meantime, further work on existing techniques is needed to assess to what types of law-and-economics models and extensions computational techniques are well suited.

## References

- [1] Lucian Ayre Bebchuk. “Litigation and Settlement Under Imperfect Information”. In: *RAND Journal of Economics* 15 (1984), pp. 404–415.
- [2] Noam Brown and Tuomas Sandholm. “Solving Imperfect-Information Games via Discounted Regret Minimization”. In: *CoRR* abs/1809.04040 (2018). arXiv: 1809.04040. URL: <http://arxiv.org/abs/1809.04040>.
- [3] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. 2006.
- [4] Kalyan Chatterjee and William Samuelson. “Comparison of Arbitration Procedures: Models with Complete and Incomplete Information”. In: *IEEE Transactions on Systems, Man and Cybernetics* 31 (1983), pp. 835–851.
- [5] Ronald H. Coase. “The Problem of Social Cost”. In: *Journal of Law and Economics* 3 (1960), pp. 1–44.
- [6] Giuseppe Dari-Mattiacci and Margherita Saraceno. “Fee Shifting and Accuracy in Adjudication”. In: *International Review of Law and Economics* (2020).

- [7] Andrew F. Daughety and Jennifer F. Reinganum. “Settlement Negotiations with Two-Sided Asymmetric Information: Modal Duality, Information Distribution, and Efficiency”. In: *International Review of Law and Economics* 14 (1994), pp. 283–298.
- [8] Daniel Friedman and Donald Wittman. “Litigation with Symmetric Bargaining and Two-Sided Incomplete Information”. In: *Journal of Law, Economics, and Organization* 23 (2006), pp. 98–126.
- [9] Jonah P. Gelbach. “The Reduced Form of Litigation Models and the Plaintiff’s Win Rate”. In: *Journal of Law and Economics* 61 (2018), pp. 125–157.
- [10] John P. Gould. “The Economics of Legal Conflicts”. In: *Journal of Legal Studies* 2 (1973), pp. 279–300.
- [11] Chris Guthrie. “Better Settle Than Sorry: The Regret Aversion Theory of Litigation Behavior”. In: *University of Illinois Law Review* (1999), pp. 43–90.
- [12] Keith N. Hylton. “An asymmetric-information model of litigation”. In: *International Review of Law and Economics* 22.2 (2002), pp. 153–175.
- [13] Daniel Klerman, Yoon-Ho Alex Lee, and Lawrence Liu. “Litigation and Selection with Correlated Two-Sided Incomplete Information”. In: *American Law and Economics Review* 20 (2018), pp. 382–459.
- [14] Marc Lanctot. “Monte Carlo Sampling and Regret Minimization for Equilibrium Computation and Decision-Making in Large Extensive Form Games”. PhD thesis. University of Alberta, 2013.
- [15] William M. Landes. “An Economic Analysis of the Courts”. In: *Journal of Law and Economics* 14 (1971), pp. 61–107.
- [16] Roger B. Myerson and Mark A. Satterthwaite. “Efficient Mechanisms for Bilateral Trading”. In: *Journal of Economic Theory* 29 (1983), pp. 265–281.
- [17] Stephen Portnoy and Roger Koenker. “The Gaussian Hare and the Laplacian Tortoise: Computability of Squared-Error Versus Absolute-Error Estimators”. In: *Statistical Science* 12 (1997), pp. 279–300.
- [18] Richard A. Posner. “An Economic Approach to Legal Procedure and Judicial Administration”. In: *Journal of Legal Studies* 2 (1973), pp. 399–458.
- [19] George L. Priest and Benjamin Klein. “The Selection of Disputes for Litigation”. In: *Journal of Legal Studies* 13 (1984), pp. 1–56.
- [20] Jennifer F. Reinganum and Louis L. Wilde. “Settlement, Litigation, and the Allocation of Litigation Costs”. In: *RAND Journal of Economics* 17 (1986), pp. 557–566.
- [21] Wolfgang Ryll. *Litigation and Settlement in a Game with Incomplete Information: An Experimental Study*. Springer, 1996.

- [22] Urs Schweizer. “Litigation and Settlement Under Two-Sided Incomplete Information”. In: *Review of Economic Studies* 56 (1989), pp. 163–177.
- [23] Kathryn E. Spier and J.J. Prescott. “Contracting on Litigation”. In: *RAND Journal of Economics* 50 (2019), pages=391-417).
- [24] Eric Steinberger. “Single Deep Counterfactual Regret Minimization”. In: *arXiv e-prints*, arXiv:1901.07621 (2019).
- [25] Martin Zinkevich et al. “Regret Minimization in Games with Incomplete Information”. In: *Proceedings of the 20th International Conference on Neural Information Processing Systems*. NIPS’07 (2007), 1729–1736.