Dissecting Bankruptcy Frictions

Winston Wei Dou

Lucian A. Taylor

Wei Wang

Wenyu Wang*

August 20, 2019

Abstract

How efficient is corporate bankruptcy in the U.S.? Two economic frictions, asymmetric information and conflicts of interest among creditors, can cause several inefficiencies: excess liquidation, excess continuation, and excess delay. We quantify these inefficiencies and their underlying causes using a structural estimation approach. We find that the bankruptcy process is quite inefficient, mainly due to excess delay. Eliminating information asymmetries would increase average total payouts by 11%, and eliminating conflicts of interest would increase them by an additional 29%. Without these frictions, an extra 46% of cases would be resolved before going to court, and the remaining court cases would be 64% shorter. With less delay, the direct and indirect costs of bankruptcy would be much lower. In contrast, we find that inefficiencies from excess liquidation and excess continuation are quite small.

Key words: Bankruptcy, structural estimation, conflicts of interest, asymmetric information

*Dou and Taylor are at the University of Pennsylvania. Wei Wang is at Queen's University. Wenyu Wang is at Indiana University. Emails: wdou@wharton.upenn.edu, luket@wharton.upenn.edu, wwang@queensu.ca, wenywang@indiana.edu. We are grateful for comments from Utpal Bhattacharya, Vidhan Goyal, Daniel Kim, Xuewen Liu, David Seligman (Kirkland & Ellis), Randall Wright, participants at the 2019 Mitsui Summer School on Structural Estimation in Corporate Finance, and seminar participants at HEC Paris, Hong Kong University of Science and Technology (HKUST), Norwegian School of Economics (NHH), Nova School of Business and Economics, Stanford University (GSB), University of Pennsylvania (Wharton), and University of Wisconsin-Madison.

Introduction

Bankruptcy plays an important role in our economy. On average from 1998 to 2017, 95 U.S. corporations with liabilities above $100 million filed for Chapter 11 bankruptcy each year.[1] During the most recent recession, from 2008 to 2009, 379 such companies with combined liabilities of $1.3 trillion filed for bankruptcy. How efficiently does the U.S. bankruptcy system redeploy these firms' assets? The answer clearly matters for distressed firms, but it also matters for healthy firms: under the tradeoff theory of capital structure, expected bankruptcy costs affect even healthy firms' borrowing costs and leverage choices. Our goals in this paper are to quantify the efficiency of corporate bankruptcy in the U.S., and to dissect the underlying economic causes of any inefficiencies.

There are several potential bankruptcy inefficiencies. Some firms that should get reorganized instead get liquidated ("excess liquidation"). Other firms that should get liquidated instead get reorganized ("excess continuation"). There can be large direct costs, such as legal fees, as well as indirect costs, such as the loss of customers, employees, and suppliers. Long court battles can amplify these costs.

Why do these inefficiencies occur? We focus on two economic frictions that have featured prominently in the literature. The first is a conflict of interest between creditors. In recent years, equity holders are wiped out when a firm files for bankruptcy, leaving senior and junior creditors to bargain with each other.[2] During this bargaining, each creditor maximizes its "piece of the pie," which is different from maximizing the firm's value. The second friction comes from asymmetric information between creditors. Asymmetric information leads creditors to make tough, low-ball offers, which delay the case. Delay allows legal and other costs to accumulate.

Quantifying these frictions and their resulting inefficiencies is a challenge. Key factors like creditors' private beliefs and the optimal reorganization plan are inherently unobservable. Data on creditors' subjective valuations of firms' assets are not available. More important, quantifying inefficiencies requires observing a parallel, counterfactual world with no frictions. Natural

[1] This fact and the following are from Altman et al. (2019).

[2] This view is consistent with the evidence of Ayotte and Morrison (2009) and Bharath et al. (2014), for example.

experiments can help to observe that counterfactual, but they are hard to find, and their results do not generalize easily. While natural experiments help to identify causal relations in the data, quantifying the system's overall efficiency requires a model.

We overcome these challenges by structurally estimating a bankruptcy model. The model features dynamic bargaining between a senior and junior creditor, with two-sided incomplete information. The creditors must choose and agree on a business plan and a financial plan, and the judge must approve these plans. The business plan dictates whether the firm will be liquidated or reorganized. Each creditor has its own reorganization plan. The financial plan specifies how the proceeds will be split. The creditors also choose whether to reach an agreement before going to court (i.e., file a prepackaged bankruptcy) or continue negotiating in court, which can extend over multiple periods. Creditors face a tradeoff between resolving the case early, which reduces the direct and indirect costs, and delaying, which offers the possibility of finding a better reorganization plan. Due to conflicts of interest, a creditor may also delay in hopes of extracting better deal terms from the counterparty. The model includes the two frictions discussed previously: creditors maximize their own payout rather than the total payout, and they privately observe the quality of their own reorganization plans.

We estimate the model using data on 311 Chapter 11 filings (prepackaged and traditional) by large, public, non-financial U.S. firms from 1996–2014. Our sample is one of the most comprehensive in terms of having detailed, complete information on the timing of events, debt structure, estimated liquidation values, final outcome (liquidation versus reorganization), and debt recovery. We estimate the model's parameters using the simulated method of moments (SMM). Parameters estimated include the fixed cost of going to court, the direct cost per period in court, the rate of decay in going-concern value during court, the initial quality of creditors' reorganization plans, the speed at which plan quality increases, and creditors' relative bargaining power. Data on creditors' average payoffs, especially for cases resolved early, help identify creditors' initial skill levels. The way in which payoffs are split between creditors helps identify their relative bargaining power. Data on the length of court cases and likelihood or reorganization help to disentangle the speed at which reorganization values decay and the speed

of learning, which determines the amount of asymmetric information. Overall, the model does a good job fitting the distributions of creditors' recovery rates, the timing of outcomes, the negative relation between debt recovery and case duration, the frequency of outcomes (liquidation versus reorganization), as well as several inputs to the estimation: debt structure, liquidation values, and industry valuation ratios.

After estimating the model, we use it as a laboratory to quantify bankruptcy inefficiencies and the frictions that cause them. We do so by comparing simulated data from the estimated model to two counterfactual benchmark models. The first benchmark turns off the asymmetric-information friction, and the second benchmark additionally turns off the conflicts-of-interest friction. The second benchmark corresponds to a social planner who maximizes firm value and perfectly observes both creditors' reorganization skill. We find that the average total payout to both creditors, equivalent to firm value, increases by 11% if we remove asymmetric information, and it increases an additional 29% if we also remove conflicts of interest. The frictions together destroy about $23 billion per year, on average, in large U.S. bankruptcies. These results imply that the observed bankruptcy process is quite inefficient. Asymmetric information between creditors generates a significant inefficiency, and conflicts of interest among creditors generate an even larger inefficiency.

How do these frictions generate inefficiencies? One reason is that asymmetric information and (especially) conflicts of interest result in too many cases going to court without a prepackaged agreement, and they make court cases excessively long. The fraction of cases going to court without a prepackaged agreement decreases by 8 percentage points (from 75% to 67%) when we remove asymmetric information, and further decreases by 38 percentage points (to 29%) when we remove conflicts of interest. The average duration of the remaining cases decreases from 13.7 to 7.1 months without asymmetric information, and to 4.9 months without conflicts of interest. In other words, removing these frictions would reduce court cases' duration by 64%. Less delay results in lower legal, accounting, and other direct costs, which, according to our model, explains 20-25% of the efficiency improvements. The larger share of efficiency improvements comes from improved reorganizations. While the frictions have little effect on the fraction of firms that

reorganize, they significantly reduce the value of reorganized firms. Removing the frictions increases reorganization values by making reorganizations happen sooner, which reduces decay in the going-concern value, and by ensuring that the creditor with the best plan leads the reorganization. Surprisingly, few firms switch between liquidation and reorganization when we compare the estimated model to the counterfactual benchmarks, implying that excess liquidation and excess continuation are quantitatively small problems.

Our results are intuitive. Asymmetric information causes delay by leading creditors to engage in inefficient screening. Each creditor makes tough, low-ball offers to prevent its counterparty from signaling that it has a high-quality reorganization plan. Conflicts of interest cause further delay. If the two creditors have similar ability, they have an incentive to reject each other's offers in hopes of making a counteroffer and extracting better deal terms in the future. While creditors delay in this fashion, the going-concern value decays, potentially to the point where liquidation becomes the best remaining option. More simply, by "playing tough" with each other, creditors can destroy significant value and even liquidate a firm that should have been reorganized much earlier.

Discussions of bankruptcy inefficiencies, and the agency and information frictions that cause them, date back at least to Baird (1986), Jackson (1986), Bebchuk (1988), Giammarino (1989), Gertner and Scharfstein (1991), and Aghion et al. (1992). We contribute by quantifying these inefficiencies and their economic sources. Several papers provide reduced-form evidence that bankruptcy frictions exist. In a recent contribution, Bernstein et al. (2017) compare the efficiency of liquidation and reorganization, and they show that liquidation results in lower asset utilization, mainly due to search and financial frictions. Ivashina et al. (2016) show that higher debt concentration (a proxy for low coordination frictions) is correlated with indicators of efficient Chapter 11 outcomes, namely, faster bankruptcy resolutions and higher likelihoods of survival as an independent going concern. Evidence of conflicts of interest between creditors comes from Ayotte and Morrison (2009), who show that a bankrupt firm is more likely to be sold, even at a fire-sale price, when senior creditors are oversecured, meaning they are highly likely to be paid back in full. Stromberg (2000) finds further evidence of creditor conflicts in Swedish cash

auction bankruptcies. Gilson (1990) and others show that when senior bank lenders make up a more prominent part of the firm's capital structure, pre-court restructurings are more likely, meaning legal costs are reduced. These papers provide important evidence on the economic mechanisms that generate bankruptcy inefficiencies. They do not, however, attempt to quantify these inefficiencies, which is our main goal. These papers also point out that there are bankruptcy frictions beyond those we study. We therefore do not claim to quantify all bankruptcy frictions or inefficiencies.

A related literature measures the direct and indirect costs of bankruptcy. Altman et al. (2019) provide a summary. Several studies extending from Gruber and Warner (1977) through LoPucki and Doherty (2004), Bris et al. (2006), and Lopucki and Doherty (2008) document bankruptcy's direct costs, meaning out-of-pocket expenses for lawyers, accountants, and other professionals. It is of course harder to measure bankruptcy's indirect costs, for example, from the loss of customers and employees. Opler and Titman (1994), Davydenko et al. (2012), Graham et al. (2016), and others find evidence of significant indirect costs. In contrast, Andrade and Kaplan (1998) and Maksimovic and Phillips (1998) find that while financial distress is costly, Chapter 11 itself entails few real economic costs. Other notable studies, such as Hortacsu et al. (2013), Brown and Matsa (2015), and Glover (2016), study the costs of financial distress or default, which are related to but distinct from the costs of bankruptcy. The direct and indirect costs of bankruptcy play an important role in our work, but our approach overall is quite different. For one, we seek to understand the economic frictions that generate these costs. For example, why do bankruptcy cases last so long and therefore incur such large legal costs? Also, we provide a benchmark for judging whether the observed costs are large or small. In addition to these costs, we study two other forms of inefficiency: excess liquidation and excess continuation. Finally, we take a different approach to estimating indirect costs. We infer these costs from creditors' decisions and payouts rather than from, for example, product-market variables, labor-market variables, or ex ante leverage choices.

Two other papers apply structural estimation to bankruptcy data. Like us, Eraslan (2008) estimates a dynamic bargaining model of corporate bankruptcy, but her goal is to quantify

liquidation values and the impact of mandatory liquidation. Closer to our work, Jenkins and Smith (2014) estimate the losses from inefficient liquidations in bankruptcy. They find an average loss of 0.28 percent of firm value across all bankrupt firms. Unlike Jenkins and Smith (2014), we allow asymmetric information, we model the dynamics of bankruptcy cases, and we allow not only inefficient liquidation but also inefficient reorganization and inefficient delay. Due in part to our broader notion of inefficiency, we find much larger average inefficiencies.

1 Model

This section describes the model's setup and then explains the predictions that form the basis of our estimation. The model features an insolvent firm that is considering filing for Chapter 11 bankruptcy. The firm's senior and junior creditors bargain with each other over a potentially infinite time horizon. The bargaining game features two-sided information asymmetry and combines elements from Bebchuk (1984), Chatterjee and Samuelson (1987), and Spier (1992).

1.1 Setup

The model starts with a firm that is insolvent, meaning its debt exceeds its continuation value. The equity holders have been wiped out, and now the firm's senior and junior creditor are bargaining with each other, consistent with the evidence of Ayotte and Morrison (2009), Bharath et al. (2014), and Kim (2018). The senior creditor is owed D_S, and the junior creditor is owed D_J. We denote the firm's total debt as $D = D_S + D_J$. We normalize D to 1 without loss of generality, so all dollar-denominated variables should be interpreted as scaled by D. Bargaining starts at $t = 0$, which we interpret as the pre-court period. If the creditors cannot not reach an agreement in period $t = 0$, the case goes to court starting in $t = 1$. Once in court, bargaining continues in each period $t = 1, 2, ...$ until creditors reach an agreement. The firm incurs a one-time direct cost of $c_0 D$ if the case goes to court, and it incurs a direct cost of $c_1 D$ during each period the case stays in court. The direct cost includes legal costs, accounting costs, and other out-of-pocket professional fees. The creditors ultimately pay these direct costs, because the costs reduce creditors' final payoffs. Accumulated direct costs at the end of period t are denoted

$C_t = \mathbf{1}_{\{t>0\}}(c_0 + c_1 t)\, D$. The initial cost c_0 should be interpreted as the cost of going to court minus the cost of a pre-court settlement. For example, a low estimate of c_0 implies that the cost of going to court is close to the cost of coordinating creditors pre-court.

The outcome for the firm is either liquidation or reorganization, either before court or in court. In a liquidation, the firm's assets are sold for a known amount L, legal costs C_t are paid, and then any remaining proceeds are paid to the creditors. The absolute priority rule (APR) holds in liquidation, so the senior creditor collects $\min(L - C_t, D_S)$, and the junior creditor collects the residual, $L - C_t - \min(L - C_t, D_S)$. APR thus creates an asymmetry between the two creditors.

Reorganizing entails choosing a new scope and vision for the firm, and possibly replacing the management team. If a reorganization occurs in period t, the firm emerges from bankruptcy as a going concern with value between 0 and $V_{h,t}$. These lower and upper bounds represent the worst and best possible outcomes from reorganization. Being in court can erode a firm's going-concern value, for example, by causing it to lose employees, customers, brand value, and suppliers, and also by distracting the management team. We model this value erosion by assuming only a fraction $\rho < 1$ of a firm's reorganization value survives into the next period:

$$V_{h,t} = \rho^{t-1} V_{h,0}, \quad t = 1, 2.... \tag{1}$$

Leading a reorganization requires skill. This skill reflects the quality of the reorganization plan and, since creditors sometimes manage the firm post-bankruptcy, managerial ability. We allow the senior and junior creditors to have different levels of reorganization skill, and we allow these levels to change over time. Specifically, the senior and junior creditors have reorganization skill $\theta_{S,t}$ and $\theta_{J,t}$, respectively, at time t. Both θ values are in $[0, 1]$. If creditor k leads the reorganization in period t, then the total payoff upon emergence is

$$U_t(\theta_{k,t}) \equiv \theta_{k,t} V_{h,t} - C_t = \rho^{t-1} \theta_{k,t} V_{h,0} - C_t, \quad \text{with } k \in \{S, J\}. \tag{2}$$

This assumption implies that higher skill produces a higher reorganization value, but the total

payoff will always be in $[0, V_{h,t} - C_t]$.

The two creditors' initial reorganization skills are $\theta_{S,0}$ and $\theta_{J,0}$, respectively. These initial values are publicly known, but their future values are privately known. We allow creditors' skill to increase over time, which we interpret as learning. Learning could result from creditors' information gathering, analysis, and unexpected insights, all of which are reasonably known privately by each creditor.[3] We allow learning because it arguably takes time for creditors to find the best possible reorganization plan, which corresponds to $\theta = 1$. We capture these effects by assuming $\theta_{J,t}$ and $\theta_{S,t}$ follow independent, increasing Markov processes. Specifically, if a creditor's ability is θ_t at time t, then its ability next period, θ_{t+1}, is drawn randomly from the generalized beta distribution, which has the cumulative distribution function

$$F_\beta(\theta_{t+1}|\theta_t) = 1 - \frac{(1-\theta_{t+1})^\beta}{(1-\theta_t)^\beta}, \qquad \theta_t \leq \theta_{t+1} \leq 1, \qquad \beta \geq 1. \tag{3}$$

A higher value of β implies slower learning, meaning smaller average increments to ability. We choose the beta distribution for a few reasons. It guarantees that next period's ability is between current ability and the maximum ability of 1. The speed of learning slows over time, which is a feature of many learning models and captures the natural idea that "low hanging fruit" is picked early. The distribution is quite flexible, nesting the the uniform distribution as a special case when $\beta = 1$. We show later that this distribution allows us to fit the bankruptcy data quite well. Finally, the beta distribution significantly improves tractability.

Bargaining works as follows in each period, including the pre-court period. First, one creditor is given the opportunity to make a proposal. The junior creditor receives this opportunity with probability λ_J, and the senior creditor receives it with probability $1 - \lambda_J$.[4] Proposals are "take it or leave it," so a higher λ_J increases the junior's relative bargaining power. A creditor can propose reorganizing, liquidating, or waiting. In a reorganization proposal, the junior creditor

[3] Ability could also increase if a high-ability investor buys the stake of a low-ability creditor. While such a transaction would be public knowledge, the new investor's ability to restructure this firm is arguably private information.

[4] The random-proposal scheme is common in bankruptcy models. See, for example, Posner and Kordana (1999), Eraslan (2008), and Antill and Grenadier (2019). This assumption reflects that creditors cannot continuously negotiate with each other, and creditors rarely have detailed reorganization plans ready simultaneously.

(for example) proposes reorganizing the firm under his own plan and paying the senior creditor $\xi_{J,t-1}$, with the remaining value going back to the junior. The subscript $t-1$ on $\xi_{J,t-1}$ means $\xi_{J,t-1}$ depends on information up to the end of period $t-1$. A creditor can also propose liquidating the firm for the total payout of $L - C_t$, which is split according to APR. Liquidation proposals automatically end the game, but if the responding creditor prefers not to liquidate the firm, it can instead reorganize the firm under its own plan as long as it pays the proposing creditor what it would receive upon a liquidation. For example, a creditor can propose liquidating the firm, but the responding creditor can prefer to reorganize, in which case we would classify the outcome as a reorganization in our simulated data. Finally, a creditor can propose waiting by making an offer that is so low that it is rejected by the counterparty for sure. The waiting offer allows the creditor to move the game to the next period. After the proposal is made, creditors observe their updated abilities $\theta_{S,t}$ and $\theta_{J,t}$. The responding creditor then reviews the proposal and either accepts it, which ends the game, or rejects it, which moves the game to the next period.

Each creditor is risk neutral and maximizes its expected payoff. Each period, the proposing creditor optimally chooses its financial plan (i.e., ξ) and business plan, and the responding creditor optimally chooses whether to accept or reject. We assume a perfect Bayesian equilibrium (PBE). The creditors behave rationally at every node of the game given beliefs, and their beliefs are derived from the equilibrium strategies by Bayes' rule.

1.2 Discussion

We choose not to explicitly model the bankruptcy judge, for a few reasons. Foremost, the judge is not the source of the inefficiencies we study. Instead, creditor conflicts and information asymmetries determine the inefficiencies.

Second, in the large bankruptcies that we study, judges mainly facilitate the process without intervening actively. Judges in reality do not negotiate or bargain directly with creditors. Judges respond to motions made by the debtor or creditors. When deciding whether to approve a plan and send it out to all creditor classes for a vote, judges apply a fairly lenient feasibility

standard.[5] Judges typically do not confirm a plan unless all creditor classes vote in support of it. The bankruptcy system therefore values consensus among creditors, consistent with our assumption that a case is resolved when and only when all creditors accept the plan. However, judges can, at times, intervene actively by "cramming down" a plan, meaning they force a plan onto one or more unwilling creditor classes, when at least one creditor class votes for the plan. We choose not to model such direct intervention by judges, however, because cram down is rare in large Chapter 11 cases.

Related to the previous point, there is no evidence that judge-specific preferences matter in the large bankruptcy cases that we study. Small bankruptcy cases are quite different. For example, Bernstein et al. (2017) study a sample of mostly small bankruptcy filings, and they find that judges have significant biases regarding conversion from Chapter 11 reorganization to Chapter 7 liquidation. However, when they limit their analysis to firms with more than 1,000 employees, they find no significant evidence of judicial bias.[6] Given the lack of evidence in large cases, we choose not to model judge-specific preferences.

Although we do not explicitly model the judge, the parameter estimates we find may reflect the judge's role. Judges in reality must approve a creditor's proposal before sending it out for a vote. We can therefore think of our parameter λ_J as depending in part on the judge's willingness to approve plans from the junior versus the senior creditor. Also, judges in reality have some limited control over the speed of a case. If judges pressure cases to move quickly, this pressure could show up in our parameter estimates as faster creditor learning (i.e., lower β) or shorter periods (e.g., one period could correspond to one month rather than one year).

1.3 Model Solution

Next, we describe a few features of the model solution that are important for our estimation approach. Appendix A contains technical details on the solution.

[5] A plan is considered feasible if it makes it unlikely that a firm will fall back into bankruptcy or piecemeal liquidation in the near future.

[6] As Bernstein et al. (2017) explain, "presumably the stakes are large enough in these cases that judicial preferences are of less consequence." Some of the earliest evidence on judge fixed effects come from Chang and Schoar (2013). The cases in their sample are orders of magnitude smaller than ours. Like us, Iverson et al. (2018) study large bankruptcy filings. They show that judge fixed effects explain little to no variation in Chapter 11 outcomes.

We solve the model numerically via dynamic programming. Solving the model entails finding the two creditors' value functions and policy functions. The state variables are the creditor's true ability, the counterparty's perception of the creditor's ability, the counterparty's ability as perceived by the creditor, and the period (t). The case is guaranteed to be resolved by some period T defined by $\rho^T V_{h,0} < L$. This means that eventually so much going-concern value has been lost that liquidation becomes optimal, and there is no benefit of further delay.

To start, we illustrate the model's assumptions about learning and reorganization values. The top panel of Figure 1 illustrates how creditors' skill levels increase over time. Specifically, we use the estimated model parameters from Section 3.2 of the paper, and we plot the median simulated values of $\theta_{S,t}$ and $\theta_{J,t}$ versus t. With these parameter values, skill levels increase quite slowly. Not shown in the figure, the shocks to ability generate randomness around these medians. The bottom panel translates skill levels into reorganization values. The top line shows the decay in maximum reorganization value, $V_{h,t}$, which this figure normalizes to 1 at $t = 1$. The lower lines show the creditors' median reorganization value, which equal $V_{h,t}$ times each creditor's median simulated ability. We see that learning and value decay combine such that median reorganization values are roughly constant in the initial periods, and then they gradually decline. An important implication is that the option to wait and learn is not very valuable, with these parameter values. Therefore, any large observed delays are likely to be inefficient.

Next, we describe the creditors' optimal offers, starting with their business plans. Figure 2 plots the offers for different combinations of ability and different points of time. The horizontal axis denotes the creditor's true ability, and the vertical axis denotes said creditor's perception of the counterparty's ability. The red areas represent the regions in which the creditor makes waiting offers, the gray areas represent the regions of quitting offers, and the blue areas represent the region of reorganization offers. The top two subplots show the offers made by the senior and junior creditor at the first period (t=0), and the bottom two subplots show an example of offers made in a later period (t=2).

We see a few interesting patterns. A creditor with very high reorganization ability always proposes reorganization, regardless of the other creditor's perceived ability. The value of waiting

is low for a high-ability creditor, because waiting is costly and ability does not have much more room to grow. If a creditor's ability is lower, the value of waiting increases, because the creditor's ability has more room to grow. The senior creditor (for example) is especially willing to wait if it believes the junior has high ability, because the junior could lead a highly valuable reorganization in some future period, creating a large surplus to be shared. Important for our findings, if the creditors' ability levels are similar— either both low or both high— then both creditors often prefer to wait until next period, which delays the bargaining process. Creditors prefer to wait in these situations because their bargaining power is quite similar, and each creditor hopes to gain greater bargaining power or a higher ability level in the future. Finally, the top-left corner of the first panel shows "bluffing" offers. Even though these senior creditors have low reorganization ability, they propose leading a reorganization in hopes of tricking the junior creditor into believing their ability is high. The senior knows these offers will almost surely be rejected, which again delays bargaining.

The right panel of Figure 3 describes creditors' optimal financial offers and the counterparty's optimal response. Specifically, it plots $\xi_{J,0}$, the optimal payout offered by the senior to the junior creditor at $t = 0$. We vary the senior's true skill and hold all else equal. If the senior's skill is low, it makes a waiting offer by offering a low payout that will be rejected for sure. For higher skill levels, the offered payout increases with the senior's own skill. Intuitively, if the senior's skill is high, he can make the pie larger and thus has more to share. Also, higher ability makes the senior creditor more eager to have its offer accepted, so it offers more to the junior. An important implication, which we prove formally in the appendix, is that reorganization offers perfectly reveal a creditor's true skill. For example, if the junior sees the senior offering a payout of 0.13, corresponding to the red dot, then the junior can infer that the senior's true ability is 0.5. It is important to note, however, that only the "morning" level of skill can be inferred from an offer. Neither creditor knows how much skill will increase in the afternoon of the same period, so the ultimate level of skill in a reorganization remains uncertain.

The right panel describes the junior's optimal response to an offer from the senior. Specifically, it plots the junior's reservation value as a function of its own ability. If the junior is

offered a payout below its reservation value, it rejects the offer, otherwise it accepts the offer. The junior's reservation values increase with its own ability, because the junior's outside option is to propose a reorganization, under its own ability, in a future period. The red and blue dots describe two example scenarios. The red dots have the same y-axis value in both panels, as do the blue dots. The dots on the right therefore define a critical value for junior ability, above which the junior rejects the offer, and below which the junior accepts it. If the senior's skill is 0.5, then he makes an offer worth 0.13 (red dot on left panel), and this offer is rejected only by juniors with skill above 0.6 (see red dot on right panel). If the senior's skill is instead 0.95, he makes an offer worth 0.21 (blue dot on left panel), and this offer is rejected only by juniors with skill above 0.8 (see blue dot in right panel). Intuitively, strong creditors make strong offers, and strong offers will be rejected only by strong opposing creditors, because strong opposing creditors have strong outside options. The decision to accept or reject an offer therefore signals that the responding creditor's skill is on one side of a critical threshold. Our model takes this signaling and the resulting belief updating into account.

Finally, to illustrate the types of inefficiencies that arise in the model, Figure 4 shows four cases simulated from the model. The four simulations differ only in the choice of proposer each period and the shocks to ability. The left column shows each creditor's ability, and the right column shows the type of offer made as well as its value. Circles indicate waiting offers, squares represent reorganization offers, and triangles represent quitting offers. Red represents the senior creditor, blue the junior.

The top row illustrates the situation where the senior and junior creditor have similar and relatively low ability along the path, and the senior gets the chance to propose in early rounds. As Figure 2 shows, when the creditors' ability levels are similar, the senior creditor prefers making a waiting offer. These waiting offers continue until period 6, at which point $\rho^t V_{h,0} < L$, meaning the reorganization value has decayed so much that liquidation is now the optimal choice. This case is inefficient in the sense that, with the benefit of hindsight, the firm would have been better off liquidating immediately and avoiding the direct costs C_t.

In the second simulation (row 2), we keep everything the same but just change the proposer

in the first round from senior to junior. From Figure 2, we know the junior makes a liquidation offer, which ends the case. The senior, faced with low current ability, agrees to liquidate.

In the third simulation (row 3), we keep the realized path of the junior's ability and the proposers the same as in simulation 1, but we increase the senior's ability along the path. The senior creditor still proposes waiting in the first three rounds. By the fourth round, though, the senior's ability is quite high and does not have much more room to grow, so the senior's value of waiting is low, which leads the senior to make a reorganization offer. The junior still has low skill and therefore accepts the reorganization offer, and the case ends at $t = 4$. This case is efficient in the sense that, with the benefit of hindsight, the firm was better off for giving the senior's ability a chance to grow over time, which allowed a relatively high-ability reorganization to occur.

The fourth row shows a highly inefficient case. We start from simulation trial #3 but make the junior's ability the same as the senior's ability. With higher ability, the junior now rejects the senior's reorganization proposal made in the fourth period. The bargaining now ends at $t = 6$, when so much reorganization value has decayed that liquidation is the best remaining choice. This is an example of a firm that should have been reorganized (in an earlier period, and by either creditor) but instead was liquidated. Comparing simulations 3 and 4, we notice that when the proposer has high ability, it is not necessarily good to have a responder with high ability. Fierce competition between two high-ability creditors may lead to delay, which entails high direct costs and decay in the going-concern value.

Delay in our model comes from both asymmetric information and conflicts of interest. Similar to Admati and Perry (1987), asymmetric information leads to strategic delay. The mechanism in our model is different, however, and more in line with that of Chatterjee and Samuelson (1987). In our model, asymmetric information leads creditors to make tough, low-ball offers that are typically rejected, which delays the case. The benefit of making a low-ball offer is that it prevents the responding creditor from signalling high ability. After all, almost any creditor—even one with low ability—would reject a low-ball offer. By keeping the responding creditor's perceived ability low, the proposing creditor improves its bargaining power in future periods. Conflicts of

interest lead to further delay. Each creditor in our model maximizes its share of the surplus, not the total surplus. A creditor's share of the surplus depends on its bargaining power, which is greater when a creditor is proposing a deal compared to responding to one. Delay occurs when a creditor rejects a "good" proposal in hopes of being chosen to make its own proposal next period, which would allow the creditor to capture a bigger share of the surplus. Rejecting a "good" proposal can be privately optimal even if the creditor knows that delay will destroy part of the total surplus. With either asymmetric information or conflicts of interest, creditors play tough with each other, delaying the case and potentially destroying part of the total surplus.

2 Estimation Method

This section describes our data, SMM estimator, and intuition behind the estimation method.

2.1 Data and Empirical Measures

Our sample consists of 311 Chapter 11 filings by large, public, non-financial U.S. firms from 1996–2014. To construct this sample, we first retrieve all business bankruptcy filings (Chapter 7 and Chapter 11) by U.S. firms from 1996 to 2014 from the UCLA LoPucki Bankruptcy Research Database. This database contains bankrupt U.S. firms that have assets above $100 million in constant 1980 dollars and must have filed financial reports with the SEC within three years of their bankruptcy. This step produces 752 filings, which includes 733 Chapter 11 filings and 19 Chapter 7 filings.[7] The status and outcome of the Chapter 11 cases, including reorganization, liquidation, converted to Chapter 7, sold as a going-concern, dismissed, or still pending, are cross-checked and verified with New Generation Research's bankruptcydata.com as of March, 2016. After removing dismissed cases and pending cases, we have 705 Chapter 11 filings and 19 Chapter 7 filings in the sample. We further remove filings by financial institutions (SIC 6000-6999), due to their unique capital structure and debt structure, which results in 626 bankruptcy filings, only 2 of which are Chapter 7 filings. From the LoPucki database, we collect each

[7]The fraction of Chapter 7 filings in our sample is small compared to that reported by U.S. court systems because our sample consists of the largest U.S. firms. Chapter 7 is typically filed by small businesses that often have no going-concern value (Altman et al., 2018).

case's basic information: the firm's book assets and liabilities at filing; whether the case has a prepackaged/pre-negotiated filing; the confirmation date and effective date of the reorganization or liquidation plan, or the conversion date for Chapter 11 cases converted to Chapter 7; and whether there are asset sales through Section 363 or the reorganization plan.[8]

Next, from New Generation Research and Public Access to Court Electronic Records (PACER), we retrieve the final reorganization or liquidation plans and disclosure statements that are confirmed by the bankruptcy court. Since the majority of U.S. bankruptcy courts started to maintain electronic case dockets on PACER only in 2002, we must rely on other sources before 2002. We obtain these documents for a large fraction of our pre-2002 cases from National Archives at various locations and U.S. bankruptcy courts for various districts. This comprehensive data retrieval process allows us to obtain the final bankruptcy plan and/or disclosure statements confirmed by the court for 520 of our sample cases. (This step drops the two remaining initial Chapter 7 filings.) We use these documents to identify two pieces of information. First, these documents contain each claim class's recovery rate, meaning the fraction of the debt that is repaid at the resolution of the case.[9] Second, the plans and disclosure statements provide a detailed classification of claim/debt classes and the estimated amount owed or outstanding of each class of claims, which we use to measure the total debt D as well as D_S and D_J. We have enough information to determine the type of claim classes, priority of the claim, and claim amount for 439 Chapter 11 cases. Given the focus of our study, we require that a debtor firm have at least two debt claim classes to be included in the sample. This step eliminates 128 cases with a single class of debt, resulting in the final sample of 311 Chapter 11 filings for our study.

We then classify whether a debt claim is senior or junior. This is an easy task for about 60% of our sample cases that have only two classes of debt claims. For cases with more than two classes of debt claims, we classify them using the following guidelines. First, when a firm has both secured and unsecured debt, we classify secured as senior and unsecured as junior. Second,

[8] See Ma et al. (2018) for a description of Section 363 asset sales.

[9] The documents contain comprehensive information on whether a claim class is impaired and how it is treated in terms of compensation and the recovery. For example, a debt claim can be unimpaired, in which case the debt claim is paid off with 100% recovery. If a debt claim is deemed impaired, the firm will pay the claim holders with cash, new debt, new equity, or a combination of these securities, but typically the expected recovery, based on the estimated enterprise valuation, is less than 100%.

we group debt claims that have similar recovery rates into one class. This procedure allows us to estimate both the amount and recovery rates of both senior and junior claims.

Next, by searching court dockets of a large fraction of our sample cases via PACER, we are able to determine whether there are intermediate bankruptcy plans or disclosure statements filed before the final plans and disclosure statements are confirmed. We also record when these documents were filed. With these data, we can measure the number of months between observed reorganization proposals.

We merge our sample of Chapter 11 firms with Compustat to retrieve firm-level financial information for each firm as of the fiscal year-end within 12 months before a Chapter 11 filing.

We map each sample case into one of the model's four possible outcomes: pre-court reorganization, pre-court liquidation, in-court reorganization, and in-court liquidation. Pre-court reorganization occurs if the firm files a prepackaged plan at Chapter 11 filing, the firm successfully reorganizes or sells all assets either through Section 363 or the plan, and the whole reorganization process (from Chapter 11 filing date to plan confirmation date) takes less than six months.[10] Pre-court liquidation occurs if the firm files a prepackaged plan with an intent to liquidate, and it is liquidated in Chapter 11 or converted to Chapter 7, regardless of how long the process is. Note our final sample includes no initial Chapter 7 filings. In-court reorganization occurs if either (1) the case is non-prepackaged and the firm is reorganized or sold as a whole through either Section 363 or a plan; or (2) the case is prepackaged, the firm is successfully reorganized or sells all assets either through Section 363 or the plan, and the whole reorganization process takes more than six months. In-court liquidation occurs if either (1) the case is non-prepackaged and the firm is liquidated piecemeal or converted to Chapter 7; or (2) the firm files a prepackaged plan with an intent to reorganize at Chapter 11 filing, yet the firm is liquidated piecemeal in Chapter 11 or converted to Chapter 7.

We measure firms' liquidation values, which correspond to L in our model, as follows. To emerge from bankruptcy reorganization, the debtor firm must pass the "best interest" test for a

[10] We classify sales of all assets (i.e., an M&A outcome) as a reorganiziation rather than a liquidation, because the going concern remains intact. Part of reorganizing a firm involves finding the best management team for the firm's assets, regardless of whether that team is part of another firm or not. This classification also agrees with our model's assumption that reorganization requires skill. It is plausible that reaching a good M&A outcome requires skill. For example, many CEOs are compensated based on their M&A activities.

bankruptcy judge to confirm the plan. As part of this test, the debtor firm must perform a hypothetical liquidation analysis, which includes an estimated proceeds from liquidating the firm's assets. The party that performs such analysis is typically the independent financial advisors that are retained by the debtor firm. We search our sample cases' dockets for independent liquidation analyses for all sample cases filed from 2003 to 2014, the period with electronic records. We measure L as the total gross liquidation proceeds, from the initial liquidation analysis report. We are able to find this measure for 228 of our 311 sample firms. For the remaining firms, we estimate L as the fitted value from a regression of observed L values on firm and creditor characteristics; details are in Appendix B.

Estimation also uses a proxy for $V_{h,0}$, the firm's highest possible initial reorganization value. We estimate $V_{h,0}$ following the method of Edmans et al. (2012). Their goal (and ours) is to measure firms' maximum potential value absent managerial inefficiency and mispricing. In a first step, we estimate each firm's potential Tobin's Q as the 50th percentile Q among firms in the same industry and year. In a second step, we obtain our estimate of $V_{h,0}$ by multiplying the potential Q by the firm's pre-filing book assets. Appendix B contains additional details.

2.2 Simulated Method of Moments Estimator

We estimate the model using SMM, which chooses parameter estimates that minimize the distance between moments generated by the model and their sample analogues. The following section defines our moments and explains how they identify our parameters. We estimate six model parameters: $\theta_{S,0}$, the senior creditor's initial ability; $\theta_{J,0}$, the junior creditor's initial ability; λ_J, the junior's probability of proposing each period; c_0, the fixed cost of going to court; ρ, which controls the rate of decay in $V_{h,t}$; and β, which controls the speed of creditors' learning. To map the model to the data, we also need to define the length of one model period. We therefore estimate a seventh parameter, defined as the number of months per period.

One remaining model parameter is c_1, the direct costs per period during court cases. We calibrate c_1 to 0.15% of total debt value. We choose this number because it makes our estimated model produce direct costs during court, averaged across cases making it to court, equal to 1.5%

of total debt value. This value is close to the 1.4% average legal costs estimated by LoPucki and Doherty (2004).

Three model parameters are directly observed and therefore do not need to be estimated by SMM. These parameters are D_J, the amount of debt held by the junior; $V_{h,0}$, the initial maximum reorganization value; and L, the firm's liquidation value.[11] The previous subsection explains how we measure these three parameters for each case. When simulating data, we feed into the model these three parameters' values, allowing heterogeneity across sample cases. Additional details on this step and the overall SMM procedure are in Appendix C.

2.3 Identification and Selection of Moments

Since we conduct an SMM estimation, identification requires choosing moments whose predicted values move in different ways with the model's parameters, and choosing enough moments so there is a unique parameter vector that makes the model fit the data as closely as possible. It is important to exclude moments contaminated by forces outside the model. We use eight moments to identify our seven parameters.

Next, we define our moments and, to show how the identification works, we explain how the predicted moments vary with our parameters. Each moment depends on all parameters, but we explain below which moments are most important for identifying each parameter. To illustrate, Table 1 shows the sensitivity of our eight simulated moments to our seven parameters.

The first moment is the average log number of months between observed reorganization proposals, for in-court cases. Table 1 shows that this moment mainly pins down the nuisance parameter Months per Period. In our model, one period consists of one proposal by a creditor. By measuring the average months between observed proposals, the first moment is highly informative about the typical duration of a single model period. Some proposals in the model are waiting proposals, which an econometrician would not observe, so this moment is computed using only observed reorganization proposals, both in the actual and simulated data.

Moment two is the average log duration of court cases, in months. Table 1 shows that, once

[11] Since the model normalizes total debt, D, to one, we scale D_J, $V_{h,0}$, and L by the value of D before taking these parameter values to the model. Doing so makes D_S redundant.

Months per Period is pinned down, this second moment mainly helps identify ρ. Specifically, longer court cases indicate a higher value of ρ. A high ρ means that not much reorganization value decays from one period to the next. Intuitively, if there is not much cost for waiting another period, then court cases tend to last longer.

The third moment is the fraction of cases that go to court. As expected, once the previous parameters are pinned down, this moment is most informative about c_0. A higher c_0 means higher fixed costs of going to court, so we expect fewer cases to go to court.

The fourth (fifth) moments is senior (junior) creditor's average recovery rate among cases that result in a pre-court reorganiziation. These moments are highly informative about the creditor's initial ability, $\theta_{S,0}$ ($\theta_{J,0}$), as we see in Table 1. This result is expected. If a reorganization occurs in period 0 in the model, then the reorganization value is $\theta_{k,0}V_{0,t}$ for whichever creditor k leads the reorganization. If the senior has higher initial ability, it is both more likely to lead a reorganization in period 0, and the resulting reorganization value will be higher, leading to a higher recovery rate for the senior creditor. Interestingly, higher initial ability for the senior creditor (for example) leads to lower recovery rates for the junior. This result also makes sense. If the senior credit has higher ability, then the total "pie" is bigger, but there is a stronger force in the opposite direction: the senior has higher bargaining power and can give the junior a smaller slice of the pie.

Moment six is the fraction of cases that result in a reorganization, conditional on the case going to court. This moment is most informative about β, which controls the speed of learning. In Table 1, we see that more reorganizations in court indicate a lower β, meaning faster learning. If creditors can learn faster, they reach the maximum reorganization value $V_{h,t}$ sooner, which makes reorganization more attractive than liquidation in the typical case.

Arguably the toughest identification challenge is disentangling ρ and β, because both influence the costs and benefits of waiting. Table 1 confirms that moments two and six move in different directions with ρ and β, as we require for identification. The intuition is the following. Slow value decay and slow learning both tend to make court cases last longer (moment two). Slow value decay also results in more reorganizations among in-court cases, because it makes

reorganization remain an attractive option for longer. But slow learning results in *fewer* reorganizations. For example, in the limit where learning is infinitely slow, reorganization values $\theta_{k,t}V_{h,,t}$ will always be low, so reorganization will typically remain unattractive compared to liquidation.

The seventh moment is the junior creditor's average relative gain, conditional on an in-court reorganization. We define the junior's relative gain in a given case as the junior's dollar amount recovered minus the senior's dollar amount recovered, scaled by the total dollar amount recovered. This moment is designed to be highly informative about λ_J, the junior's probability of proposing. Table 1 confirms that once the previous parameters are pinned down, this moment helps pin down λ_J. The relative gain measures how the junior and senior split up the bankruptcy proceeds. As such, it depends strongly on their relative bargaining power. As discussed in Section 1, a higher λ_J gives the junior creditor more bargaining power. It therefore makes sense that a higher relative gain for the junior indicates a higher λ_J. Our approach to identifying bargaining power is similar in spirit to the approach that Ahern (2012) and others use to identify bargaining power in the context of mergers and acquisitions.

The last moment is the total recovery rate averaged across all in-court reorganizations. We define a case's total recovery rate as the total dollar payout to both creditors scaled by their total debt, D. Whereas the previous moment measures how the pie is split, this moment measures the total size of the pie. This extra moment provides some over-identification and is informative about several parameters. For example, the pie is bigger if either creditor's initial ability is higher.

3 Estimation Results

We begin by assessing how the model fits the data, and then we present the parameter estimates.

3.1 Model Fit

Table 2 shows how the model fits the eight moments that are targeted in the SMM estimation. The t-statistics test whether each data moment matches its model counterpart. The average log

months between plans is 1.86 in the model, 1.77 in the data. The average log duration (months) of in-court cases is 2.62 in the model, 2.57 in the data.[12] The model fits these features of the data quite well. For cases that go to court, the fraction that results in a reorganization is 0.945 in the model but only 0.877 in the actual data. This is the dimension on which the model has the most trouble fitting the data ($t = -3.10$). Next, we see that the model, by taking into account APR, is able to capture that senior creditors typically recover more than juniors, and the model matches the magnitudes fairly well. Looking at pre-court reorganizations, the senior creditor's average recovery rate is 80.8% in the model, 87.8% in the data. This difference between the model and data is marginally significant, with a t-statistic of 2.12. The junior creditor's recovery rate is matched much better: 20.6% in the model, 22.1% in the data. The model also does a decent job matching the fraction of cases that go to court: 78.2% in the model, 73.3% in the data ($t = -1.97$). Finally, looking at in-court reorganizations, we see that the model does a good job matching how the pie is split (i.e., the junior's relative gain) and the pie's total size (i.e., total recovery rate). Junior's relative gain is slightly negative in both the model and data, consistent with APR favoring the senior. Aggregating the creditors, the total recovery rate is only around 37%, both in the model and the data.

Our SMM estimation targets averages. As an out-of-sample test, we check how well the model can match the full distribution of key variables. Results are in Figure 5. The model fits these distributions surprisingly well. In Panel A, we see that both in the model and data, the senior creditor most often recovers 100%, but occasionally the recovery is mediocre or even quite bad. The junior's recovery distribution also matches quite well (Panel B). Both distributions are bi-modal, both in the model and data. Panel C shows that the model does a good job of matching not only the mean of court case duration, but also its variance and the shape of the distribution. The number of months between observed proposals also matches well (Panel D). Overall, Figure 5 suggests that the distributions and functional forms assumed in the model are reasonable.

As an extra out-of-sample test, we check whether the model matches the relation between average total recovery rates and case duration. These results are in Figure 6. The first bin

[12] Jensen's effects are quite large, so while $exp\,(2.57) = 13$, the average duration (not logs) is 17 months.

contains cases that resolve pre-court. In both the model and the data, cases that take longer to resolve yield lower average payouts to creditors. Model fit is not perfect, but simulated values are within the 95% empirical confidence interval. We confirm that the negative relation between total recovery rates and duration is statistically significant in the data.[13] The negative relation, which is new to the literature, suggests that prolonging a case destroys value. This descriptive, reduced-form result foreshadows the main result from our counterfactual analysis: asymmetric information and conflicts of interest destroy value, in large part by inefficiently prolonging cases.

3.2 Parameter Estimates

Table 3 contains parameter estimates from SMM.

Months per period, the nuisance parameter, is estimated at 1.31.[14] This value provides a mapping between model periods and calendar time. The average months between *observed* proposals, discussed above, is much greater, because many model proposals are unobserved waiting offers.

The estimated initial abilities of senior and junior creditors ($\theta_{S,0}$ and $\theta_{J,0}$) are 0.46 and 0.34, respectively. These estimates imply that the senior creditor would initially produce a reorganization value that is 46% of the firm's maximum potential value, and the junior would produce 34%. Although the point estimates suggest the senior is more skilled than the junior, these values are not significantly different from each other. It is plausible that senior creditors are more skilled on average, because senior debt is more often held by hedge funds and private equity funds, which tend to be more sophisticated (Jiang et al., 2012).

Parameter β, which controls the speed of creditor learning, is estimated at 12.43. Figure 1 shows how to interpret this value. Panel A of the figure simulates creditor ability over time using the estimated values of β, $\theta_{S,0}$, and $\theta_{J,0}$. With these values, it takes 4 periods (roughly 5 months) for the median junior creditor's ability to increase from its initial value of 0.34 to 0.5.

[13] A regression of total recovery rate on the log of 1+duration, with cluster fixed effects, yields a slope coefficient with a $t-$statistic of -2.4. The cluster fixed effects control nonparametrically for differences across cases' $\{D_J, L, V_{h,0}\}$. Details on computing these clusters are in Appendix C.

[14] This parameter's standard error is somewhat large. The reason is that the main moment identifying this parameter, the average number of months between observed plans, is measured off only a fraction of our sample cases, hence the moment is measured relatively imprecisely.

Even after 14 periods (roughly 18 months), creditors' ability levels are still bounded away from their maximum value of one. Learning does occur, in other words, but it is rather slow.

The estimate of ρ is 0.934. This value implies that 6.6% of the firm's reorganization value decays each period that the case stays in court. Panel B of 1 illustrates the implications. The solid black line shows how the maximum reorganization value, $V_{h,t}$, decays if $\rho = 0.934$. We see that 16% of reorganization value decays after 4 periods (roughly 5 months), and 48% decays after 12 periods (roughly 16 months). The remaining lines show that learning and value decay combine to make creditors' median reorganization value roughly constant at first, then decreasing. There is randomness around these medians, however, and creditors hope to receive positive shocks to their ability and (hence) reorganization value.

The fixed cost of going to court, c_0, is estimated to be roughly 5% of the firm's total debt value. Going to court entails a sizeable cost, which the model requires to explain why less than 80% of cases go to court. Given how we identify c_0, its estimate may capture not just direct legal costs of going to court, but also indirect costs coming from the loss of employees, customers, suppliers, etc. It makes sense that taking a bankruptcy case to court is a highly visible event that imposes real operating costs on the firm.

The junior's probability of proposing, λ_J, is estimated at 89.8%. This high value implies that junior creditors have relatively strong bargaining power. The model needs this high value of λ_J to fit the slightly negative observed junior relative gain. With a lower value of λ_J, the model would produce an unrealistically low (i.e., more negative) relative gain for the junior. A high value of λ_J is plausible. Junior debt typically makes up the majority of total debt. This fact, combined with APR, implies that juniors typically have more to lose than seniors, hence juniors have a stronger incentive to make proposals.

4 Quantifying Inefficiencies and Their Causes

Now that we have estimated the model, we can use it as a laboratory to quantify bankruptcy inefficiencies and the frictions that cause them. We focus on our model's two main frictions, asymmetric information and conflicts of interest. We compare the estimated model to two

counterfactual benchmark models in which we turn off one or both frictions.

The first counterfactual model turns off the asymmetric-information friction but retains the conflicts of interest. This counterfactual model is identical to the estimated model, except each creditor can perfectly observe the other creditor's ability at all times. Creditors still face uncertainty about future ability. The opposing creditor's true ability replaces its perceived ability as a state variable, and a creditor's own ability as perceived by the opposing creditor is no longer a state variable.

The second counterfactual model turns off not just asymmetric information but also conflicts of interest. In this model, a social planner maximizes the firm's value, which is equivalent to maximizing the expected total payout to both creditors. The social planner can perfectly observe both creditors' current ability but not future ability, as in the previous counterfactual model. Each period, the social planner chooses whether to wait, liquidate, reorganize under the senior's ability, or reorganize under the junior's ability. The tradeoff between pre-court settlement and in-court learning is only determined by the comparison between the fixed cost c_0 and the option value of learning. This benchmark is more efficient than the previous, but it is not frictionless. There are still the direct fixed cost c_0 of going to court, the direct per-period cost c_1 during court, in-court value decay captured by $\rho < 1$, and slow creditor learning captured by $\beta \gg 1$.

The columns of Table 4 compare simulated statistics from the estimated model and the two counterfactual models. The changes across columns represent the causal effects of adding or removing frictions, because all other model features and parameter values (apart from the frictions) are held equal. The main advantage of this approach is that we can perfectly enforce the "all else equal" assumption— we impose exogenous variation. The obvious limitation is that exogenous variation comes not from some feature of the data but rather from changing model assumptions, so results depend more than usual on the model's structure and assumptions. Another limitation is that results are subject to the Lucas critique, in the sense that it may be unrealistic to assume that a friction could be removed without altering other parameter values or model features.

The statistic that summarizes bankruptcy's efficiency is the average total recovery rate, which

is the total dollar payout to both creditors, scaled by the total amount of debt and averaged across simulations. This statistic's numerator equals the firm's expected value once bankruptcy is resolved, because the two creditors fully own the firm. In the top row of Table 4, we see that removing the asymmetric-information friction increases the average total recovery rate from 0.389 to 0.432, an 11% increase. In the social-planner benchmark, the average total recovery rate is 0.546, which is 40% higher than in the estimated model. Removing conflicts of interest therefore produces an additional 29% increase. These results imply that the observed bankruptcy process is quite inefficient. Asymmetric information between creditors generates a significant inefficiency, and conflicts of interest among creditors generate an even larger inefficiency.

To convert the inefficiency into aggregate dollars, we note that the combined liabilities across all large Chapter 11 filings is \$146 billion in a typical year.[15] Multiplying this total annual debt by the change in total recovery rate, $0.546 - 0.389$, yields \$23 billion per year. In other words, we find that the two frictions we study destroy an average of \$23 billion per year in the U.S., which is significant.

The remaining rows of Table 4 help explain where these results come from. We first focus on mechanics, then economic intuition. The average total recovery rate can be decomposed as (1) the fraction of firms liquidated times the average liquidation value, plus (2) the fraction of firms reorganized times the average reorganization value, minus (3) the average direct costs. All values are scaled by total debt, D. The average liquidation value is the average of L_i/D_i across the firms i that (endogenously) get liquidated. The average reorganization value equals $V_{i,h,t}\theta_{i,k,t}$ averaged across firms i that (endogenously) get reorganized at time t under either $k = S$ or J. Table 4 contains the terms in this decomposition.

We start with term (3), the average direct cost. This cost is the sum of the average fixed costs of going to court (from c_0) and the average total per-period costs of being in court (from c_1).[16] Average fixed costs decrease from 0.038 to 0.034 when we remove asymmetric information, and to 0.015 when we additionally remove conflicts of interest. This change occurs because the fraction

[15] See Altman et al. (2019). A Chapter 11 filing is large if the firm has at least \$100 million in liabilities. We average across 1996 to 2017.

[16] More precisely, the average fixed cost of going to court is c_0 times the fraction of cases going to court. The average total per-period costs of being in court equal the fraction of cases going to court times the average number of periods in court times c_1.

of cases resolved pre-court increases from 0.250 to 0.327 to 0.711 across these three models. The average total per-period costs start at an already low value, 0.011, decrease further to 0.004 and 0.001. This decrease occurs because fewer cases go to court, and the average duration of court cases decreases from 13.7 months to 7.1 (with symmetric information) and 4.9 (under the social planner). The total direct costs drop from 0.049 to 0.039 when we remove asymmetric information, and to 0.016 in the social-planner benchmark. These cost reductions explain a non-trivial 23% of the efficiency improvements in the symmetric-information benchmark, and 21% of the efficiency improvements in the social-planner benchmark.[17] These improvements come from resolving more cases before court and reducing the duration of cases that do go to court.

We now focus on term (1), the part of recovery rates coming from liquidations. We see that the fractions of firms liquidated versus reorganized are quite similar between the estimated model and the social-planner benchmark. The inefficiencies, therefore, do not result simply from "too many" firms being liquidated or reorganized. Furthermore, average liquidation values are almost identical between the estimated and social-planner models. Inefficiencies are not a result of low-value liquidations, in other words.

The symmetric-information benchmark differs slightly from the other two models. With symmetric information, slightly more firms are liquidated, and average liquidation values are higher. These improved liquidation outcomes are the largest source of efficiency improvements in the symmetric-information benchmark, explaining 43% of the total efficiency improvements.[18]

Comparing the estimated model and the social-planner benchmark, by far the largest efficiency improvements result from term (2), the part of recovery rates coming from reorganizations. While the fraction of firms reorganized is quite similar across the two models, the average value of firms upon reorganization increases from 0.446 to 0.587. This increase in term (2) explains 75% of the overall efficiency improvements in the social-planner benchmark.[19] Reorganization values are also considerably higher, 0.480, in the symmetric-information benchmark, which helps increase term (2) by 34%. In sum, we find that removing asymmetric information—and

[17] Note that 23% = (0.049-0.039)/(0.432-0.389), and 21%=(0.049-0.016)/(0.546-0.389).
[18] Note that 43% = (0.418*0.136-0.370*0.103)/(0.432-0.389).
[19] Note 75% = (0.882*0.587-0.897*0.446)/(0.546-0.389).

especially conflicts of interest—would produce bankruptcy reorganizations of significantly higher value.

Why does removing these frictions increase reorganization values? We consider three possible explanations. The first is that frictions result in the wrong firms being liquidated versus reorganized. To quantify this channel, we track each simulated case across the three models, and we tabulate the fraction of all cases that are liquidated in the estimated model but reorganized in the benchmark. These are firms that should have been reorganized, but the frictions led them to be liquidated instead. And vice-versa, we tabulate the fraction of cases that switch from reorganization in the estimated model to liquidation in the benchmark model. We find few cases of the wrong firms being liquidated versus reorganized. Comparing the estimated model to the social-planner benchmark, only 3.7% of cases switch to liquidation, and only 2.2% switch to reorganization. The fractions are slightly higher in the symmetric information benchmark, but they remain under 7%. These low rates imply that excess liquidation and excess continuation are quantitatively small problems.

The second explanation is that firms are being reorganized at the wrong time. Waiting too long can destroy going-concern value due to loss of customers, employees, and the other elements captured in our model by $\rho < 1$. Reorganizing too soon could also be a mistake, because creditors' reorganization plans can improve over time. To gauge these effects, Table 4 reports the average months elapsed before a reorganization. A pre-court reorganization is coded as a zero duration. We see that the average time until reorganization decreases from 10.3 months in the estimated model to 5.4 months with symmetric information and 2.4 months under the social planner. These results indicate that frictions do lead firms to be reorganized at the wrong time. Many firms would be better off reorganizing much earlier, preferably before going to court.

A third explanation is that the wrong creditor leads the reorganization. In the estimated model, a creditor can end up leading a reorganization even though the opposing creditor has higher ability. It can be optimal for the high-ability opposing creditor to accept such a deal if it knows that delay will destroy significant value, or if the high-ability creditor faces a low

probability of proposing in the future (e.g., due to the judge favoring the other creditor). To quantify this channel, Table 4 shows the fraction of reorganizations led by the creditor with the lower ability. This fraction is zero in the social-planner benchmark. In the estimated model, 29% of reorganizations are led by the low-ability creditor, which helps explain why average reorganization values are lower than in the social-planner benchmark (0.446 versus 0.587). Surprisingly, 44% of reorganizations in the symmetric-information benchmark are led by the low-ability creditor, which helps explain why average reorganization values do not improve much relative to the estimated model (0.480 versus 0.446).[20]

The economic intuition for these results is as follows. With symmetric information, there is less strategic delay. Creditors no longer have an incentive to make tough, low-ball offers that prevent the counterparty from signaling its own strong bargaining power. Even controlling for delay, it makes sense that reorganizations produce more value under symmetric information. If both creditors can see each other's true reorganization ability, then creditors can better judge whether liquidation or reorganization is the better outcome.

Conflicts of interest add another layer of inefficiency. Even if creditors perfectly know each other's ability, they may still disagree on the optimal resolution plan. This is because they each maximize their share of the pie, not the total size of the pie. A creditor will sometimes try to grab a larger share of the pie, even if doing so shrinks the whole pie. For example, suppose a low-ability creditor is deciding whether to accept a proposal, and the creditor knows he has a good chance of making a counter-proposal next period. Proposing confers more bargaining power, which would allow him to seize a larger share of the pie. It can therefore be privately (but not socially) optimal for the low-ability creditor to reject the current proposal, even if it is quite strong. This delay makes the reorganization value decay, and the low-ability creditor could even end up leading the reorganization. Both effects reduce the total size of the pie.

[20]The fraction of reorganizations led by the low-ability creditor increases under symmetric information mainly because duration decreases. Recall that the junior's estimated initial ability is less than the senior's. With shorter duration, there is a lower chance that the junior's ability crosses above the senior's ability. Given the short duration and high estimated value of λ_J, the senior has a low chance of ever making a reorganization proposal. The end result is that any observed reorganizations occur early, and they are proposed by a junior creditor whose ability is typically lower than the senior's.

5 Conclusions

We find that corporate bankruptcy in the U.S. is quite inefficient, due to information asymmetries between creditors and especially to conflicts of interest among them. Eliminating these frictions would increase average total payouts by 40%, in part by making cases resolve faster, and in part by improving the value of firms that reorganize. These results come from structurally estimating a dynamic bargaining model with two-sided asymmetric information.

Of course, we recognize that the economic frictions we study are real and cannot be easily eliminated. Our results imply that reducing the frictions, if possible, would have large benefits. Finding contracting, policy, or other means of reducing these frictions is an interesting area for future work.

Our study focuses on bankruptcy frictions related to bargaining among creditors. There are other bankruptcy frictions and inefficiencies that could be interesting to study in future research. For example, to what degree is investment during bankruptcy suboptimal, as in Gertner and Scharfstein (1991)? How important are coordination costs among creditors, agency conflicts in the management team, and search frictions in the liquidation market? The literature provides reduced-form evidence that each of these frictions exists, but their quantitative effects on bankruptcy's efficiency remains unclear. Finally, this paper quantifies the ex post efficiency of bankruptcy. Understanding its ex ante efficiency is another important area for future work.

References

Admati, A. R. and M. Perry: 1987, 'Strategic Delay in Bargaining'. *The Review of Economic Studies* **54**(3), 345–364.

Aghion, P., O. Hart, and J. Moore: 1992, 'The economics of bankruptcy reform'. *Journal of Law, Economics & Organization* **8**, 523–546.

Ahern, K. R.: 2012, 'Bargaining power and industry dependence in mergers'. *Journal of Financial Economics* **103**(3), 530–550.

Altman, E., E. Hotchkiss, and W. Wang: 2018, *Corporate Financial Distress, Restructuring, and Bankruptcy*. Wiley.

Altman, E. I., E. Hotchkiss, and W. Wang: 2019, *Corporate Financial Distress, Restructuring, and Bankruptcy: Analyze Leveraged Finance, Distressed Debt, and Bankruptcy*, Wiley Finance Series. Wiley, fourth edition.

Andrade, G. and S. N. Kaplan: 1998, 'How costly is financial (not economic) distress? Evidence from highly leveraged transactions that became distressed'. *Journal of Finance* **53**(5), 1443–1493.

Antill, S. and S. R. Grenadier: 2019, 'Optimal capital structure and bankruptcy choice: Dynamic bargaining versus liquidation'. *Journal of Financial Economics*.

Ayotte, K. and E. Morrison: 2009, 'Creditor control and conflict in Chapter 11'. *Journal of Legal Studies* **1**, 511–551.

Baird, D. G.: 1986, 'The uneasy case for corporate reorganizations'. *The Journal of Legal Studies* **15**(1), 127–147.

Bebchuk, L.: 1984, 'Litigation and Settlement under Imperfect Information'. *RAND Journal of Economics* **15**(3), 404–415.

Bebchuk, L. A.: 1988, 'A new approach to corporate reorganizations'. *Harvard Law Review* **101**(4), 775–804.

Bernstein, S., E. Colonnelli, and B. Iverson: 2017, 'Asset Allocation in Bankruptcy'. *Journal of Finance*.

Bharath, S., V. Panchapagesan, and I. Werner: 2014, 'The changing nature of Chapter 11'.

Bris, A., I. Welch, and N. Zhu: 2006, 'The costs of bankruptcy: Chapter 7 liquidation versus Chapter 11 reorganization'. *The Journal of Finance* **61**(3), 1253–1303.

Brown, J. and D. A. Matsa: 2015, 'Boarding a sinking ship? An investigation of job applications to distressed firms'. *Journal of Finance* **71**, 507–550.

Chang, T. and A. Schoar: 2013, 'Judge specific differences in chapter 11 and firm outcomes'. Working paper.

Chatterjee, K. and L. Samuelson: 1987, 'Bargaining with Two-sided Incomplete Information: An Infinite Horizon Model with Alternating Offers'. *Review of Economic Studies* **54**(2), 175–192.

Davydenko, S. A., I. A. Strebulaev, and X. Zhao: 2012, 'A market-based study of the cost of default'. *Review of Financial Studies* **25**(10), 2959–2999.

Edmans, A., I. Goldstein, and W. Jiang: 2012, 'The real effects of financial markets: The impact of prices on takeovers'. *The Journal of Finance* **67**(3), 933–971.

Eraslan, H. K. K.: 2008, 'Corporate Bankruptcy Reorganizations: Estimates From A Bargaining Model'. *International Economic Review* **49**(2), 659–681.

Gertner, R. and D. Scharfstein: 1991, 'A theory of workouts and the effects of reorganization law'. *Journal of Finance* **46**, 1189–1221.

Giammarino, R. M.: 1989, 'The Resolution of Financial Distress'. *The Review of Financial Studies* **2**(1), 25–47.

Gilson, S.: 1990, 'Bankruptcy, boards, banks, and blockholders: Evidence on changes in corporate ownership and control when firms default'. *Journal of Financial Economics* **27**, 355–387.

Glover, B.: 2016, 'The expected cost of default'. *Journal of Financial Economics* **119**, 284–299.

Graham, J. R., H. Kim, S. Li, and J. Qiu: 2016, 'Employee costs of corporate bankruptcy'.

Grieser, W. and Z. Liu: 2018, 'Corporate Investment and Innovation in the Presence of Competitor Constraints'. *Review of Financial Studies* **Forthcoming**.

Gruber, M. J. and J. B. Warner: 1977, 'Bankruptcy costs: Some estimates'. *Journal of Finance* **32**.

Hartigan, J. A.: 1975, *Clustering algorithms*. John Wiley & Sons, New York-London-Sydney. Wiley Series in Probability and Mathematical Statistics.

Hastie, T., R. Tibshirani, and J. Friedman: 2009, *The elements of statistical learning*, Springer Series in Statistics. Springer, New York, second edition. Data mining, inference, and prediction.

Hortacsu, A., G. Matvos, C. Syverson, and S. Venkataraman: 2013, 'Indirect costs of financial distress in durable goods industries: The case of auto manufacuturers'. *Review of Financial Studies* **26**(5), 1248–1290.

Ivashina, V., B. Iverson, and D. C. Smith: 2016, 'The ownership and trading of debt claims in Chapter 11 restructurings'. *Journal of Financial Economics* **119**(2), 316–335.

Iverson, B. C., J. Madsen, W. Wang, and Q. Xu: 2018, 'Learning by Doing: Judge Experience and Bankruptcy Outcomes'.

Jackson, T. H.: 1986, *The Logit and Limits of Bankruptcy Law*. Harvard University Press, Cambridge.

Jenkins, M. and D. C. Smith: 2014, 'Creditor conflict and the efficiency of corporate reorganization'. *Working paper*.

Jiang, W., K. Li, and W. Wang: 2012, 'Hedge funds and Chapter 11'. *The Journal of Finance* **67**(2), 513–560.

Kim, D. S.: 2018, 'Shareholder Recovery and Leverage'. Working paper, University of Pennsylvania.

Lemmon, M. L., M. R. Roberts, and J. F. Zender: 2008, 'Back to the Beginning: Persistence and the Cross-Section of Corporate Capital Structure'. *Journal of Finance* **63**(4), 1575–1608.

LoPucki, L. and J. Doherty: 2004, 'The determinants of professional fees in large bankruptcy reorganization cases'. *Journal of Empirical Legal Studies* **1**, 111–141.

Lopucki, L. M. and J. W. Doherty: 2008, 'Professional overcharging in large bankruptcy reorganization cases'. *Journal of Empirical Legal Studies* **5**.

Ma, S., J. T. Tong, and W. Wang: 2018, 'Selling Innovation in Bankruptcy'. Working paper, Queen's University.

MacQueen, J.: 1967, 'Some methods for classification and analysis of multivariate observations'. In: *Proc. Fifth Berkeley Sympos. Math. Statist. and Probability (Berkeley, Calif., 1965/66)*. Univ. California Press, Berkeley, Calif., pp. Vol. I: Statistics, pp. 281–297.

Maksimovic, V. and G. Phillips: 1998, 'Asset efficiency and reallocation decisions of bankrupt firms'. *The Journal of Finance* **53**(5), 1495–1532.

Michaelides, A. and S. Ng: 2000, 'Estimating the rational expectations model of speculative storage: A Monte Carlo comparison of three simulation estimators'. *Journal of econometrics* **96**, 231–266.

Opler, T. C. and S. Titman: 1994, 'Financial distress and corporate performance'. *The Journal of Finance* **49**(3), 1015–1040.

Posner, E. A. and K. A. Kordana: 1999, 'A Positive Theory of Chapter 11'. *New York University Law Review* **74**(161), 161 – 233.

Spier, K. E.: 1992, 'The Dynamics of Pretrial Negotiation'. *The Review of Economic Studies* **59**(1), 93–108.

Stromberg, P.: 2000, 'Conflicts of interests and market liquidity in bankruptcy auctions'. *Journal of Finance* **55**, 2641–2692.

Figure 1: Dynamics of Ability and Reorganization Values

This figure shows simulated dynamics of creditor ability and reorganization value, using estimated parameter values from Table 3. To create the top panel, we initialize the creditors' abilities $\theta_{S,0}$ and $\theta_{J,0}$ at 0.46 and 0.34, respectively. We then randomly draw future values of ability from the generalized beta distribution, as in equation (3), using the value $\beta = 12.43$. The top panel plots the median simulated values of $\theta_{S,t}$ and $\theta_{J,t,}$. In the bottom panel, the solid line equals the maximum reorganization value, $V_{h,t}$, computed as in equation (1) with $\rho = 0.934$. This figure normalizes $V_{h,0}$ to 1. The lower lines show the product of $V_{h,t}$ and the medians of $\theta_{S,t}$ and $\theta_{J,t}$. These products equal the reorganization values for the senior and junior creditors, respectively.

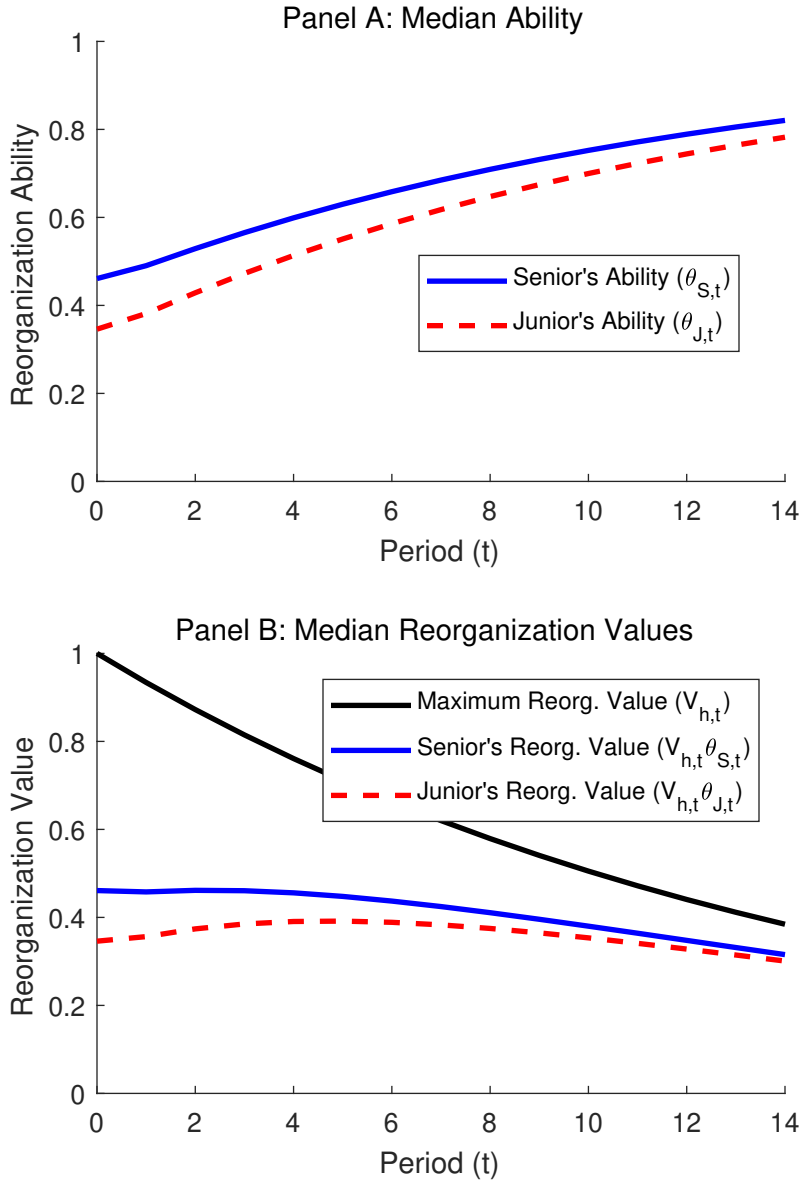


Figure 2: Optimal Offer Types

This figure shows creditors' optimal offers in our model. The horizontal axis denotes true ability, and the vertical axis denotes perceived ability. The red areas represent the regions in which creditors make waiting offers, the gray areas represent the regions of quitting offers, and the blue areas represent the region of reorganization offers. The top two subplots show the offers made by the senior and junior creditor in the pre-court period (t=0), and the bottom two subplots show offers made during an in-court period (t=2).

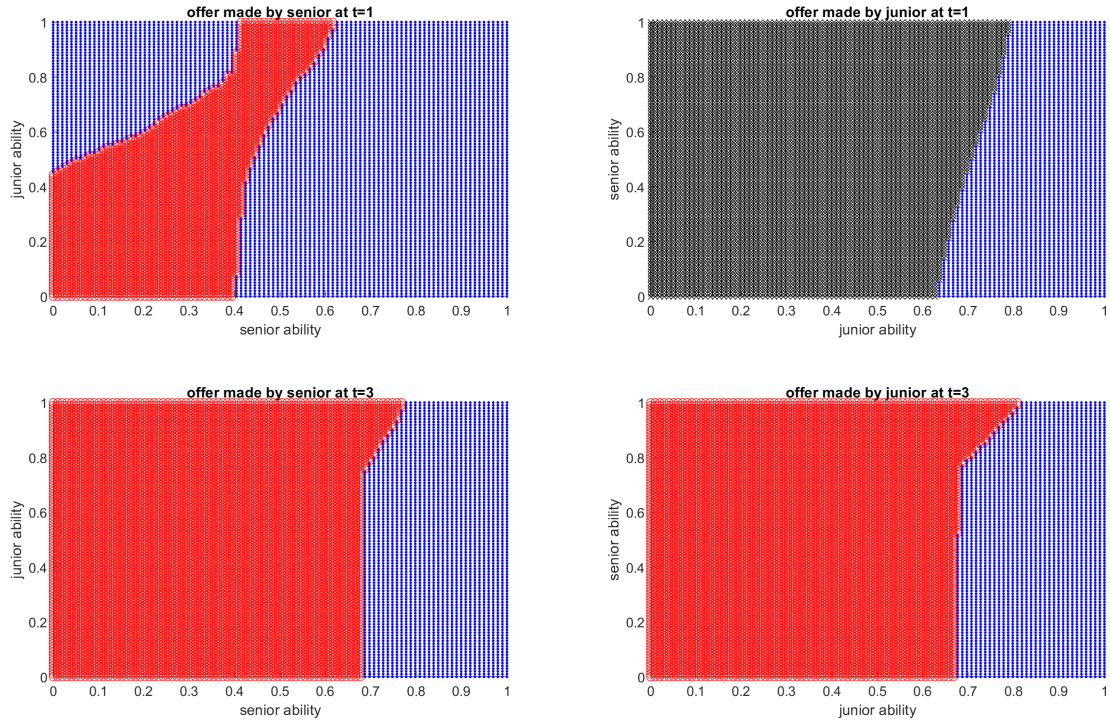


Figure 3: Optimal Financial Offers and Responses

The left panel plots $\xi_{J,0}$, the payout offered by the senior to the junior creditor, in the pre-court period, as a function of the senior's ability. The right panel plots the junior creditor's reservation value as a function of its ability. The blue (red) line illustrates the scenario where the senior's skill is 0.95 (0.50). If the senior's offered payout exceeds the junior's ability, then the junior accepts the offer, otherwise the junior rejects it.

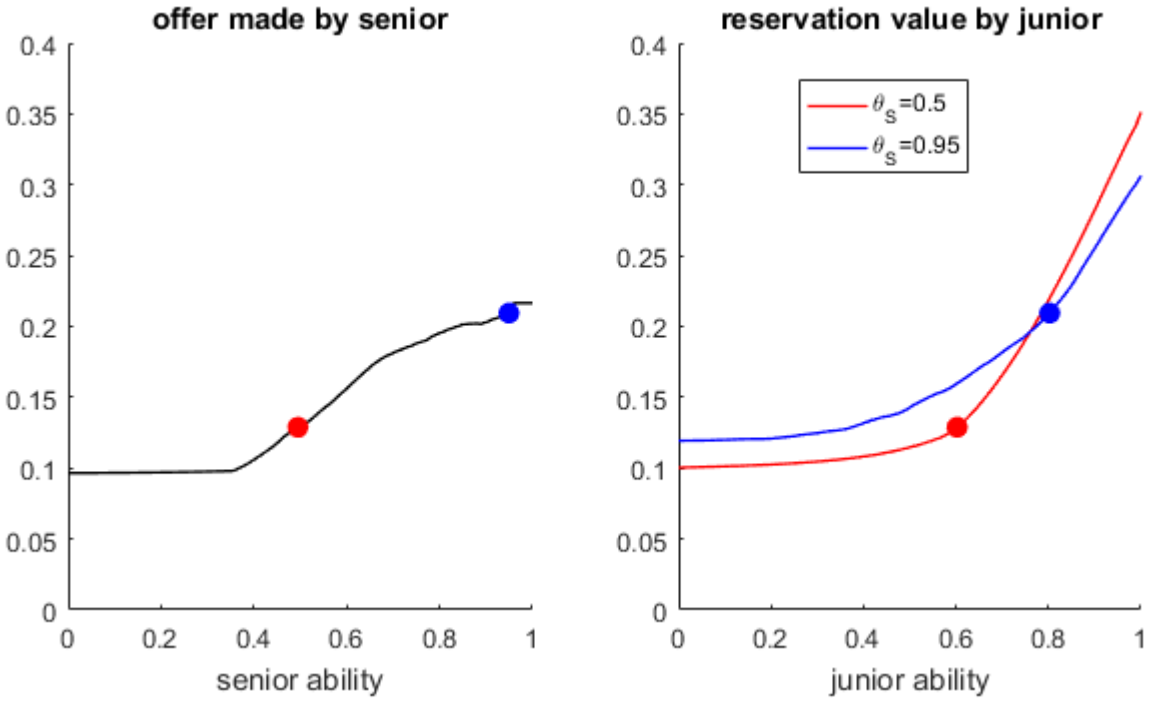


Figure 4: Examples of Simulated Bankruptcy Cases

This figure plots four simulations of the model. Each row corresponds to one simulation. The simulations differ in the realized paths of creditors' ability and the choice of proposer. The panels on the left column present the realized paths of the senior (red "x") and junior (blue '+') creditor's ability. The panels on the right column show the offer and the associated proposers. These plots contain three pieces of information: (1) who proposes (red indicates senior proposes, blue indicates junior proposes), (2) the offer type (circle means waiting offer, square means reorganization offer, triangle means quitting offer), and (3) the offer value. Waiting offers have value of zero. For the purpose of this figure, we scale offer values by the responder's face value of debt, so the offer values can be interpreted as the responder's recovery rate.

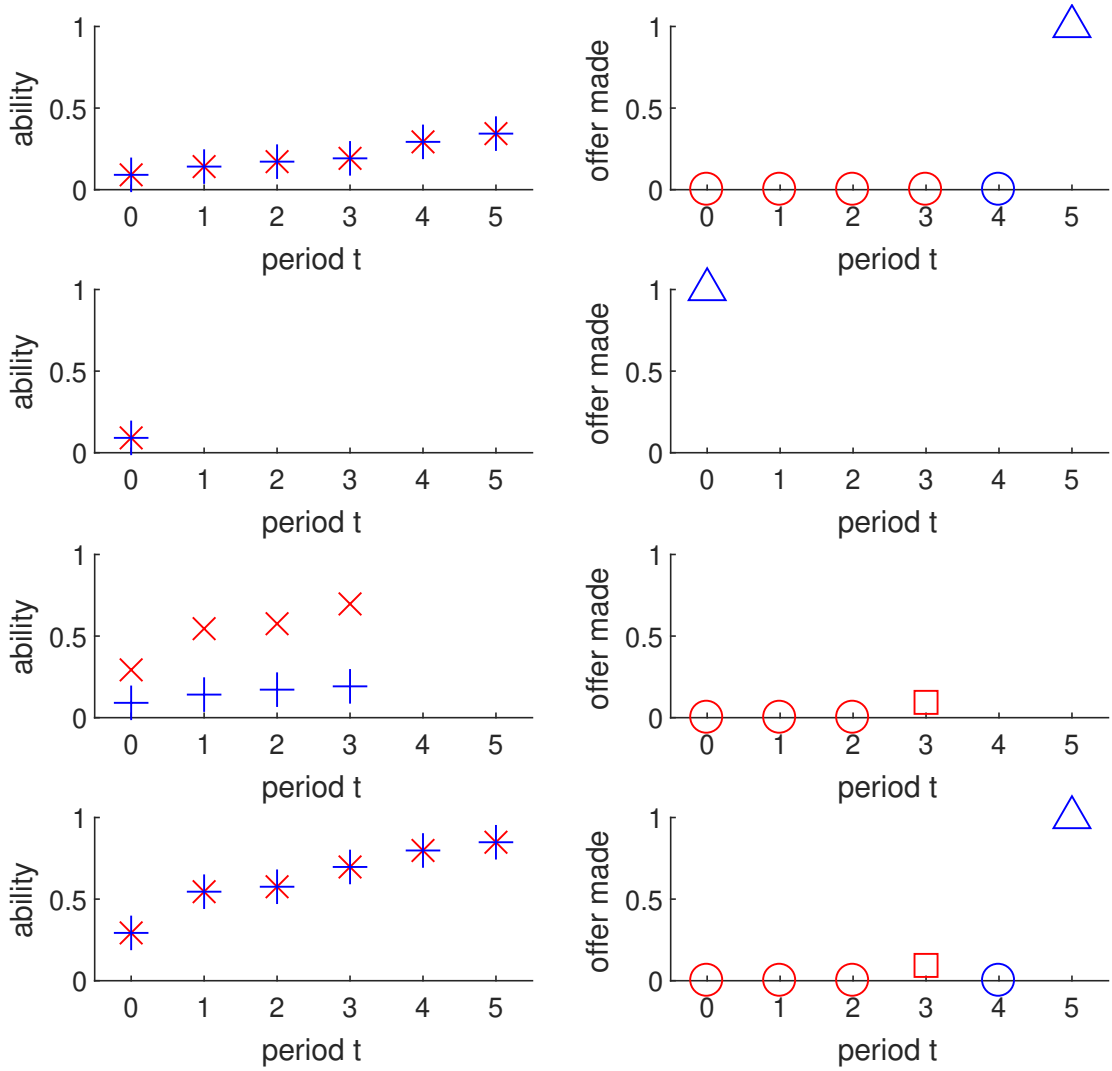


Figure 5: Comparing Simulated and Empirical Distributions

This figure plots the distributions of recovery rates (senior and junior), duration of court cases, and months between observed creditor proposals. Dark blue columns show results from data simulated off the estimated model. Grey bars show the empirical distributions. These histograms pool all cases (e.g. pre-court and in-court).

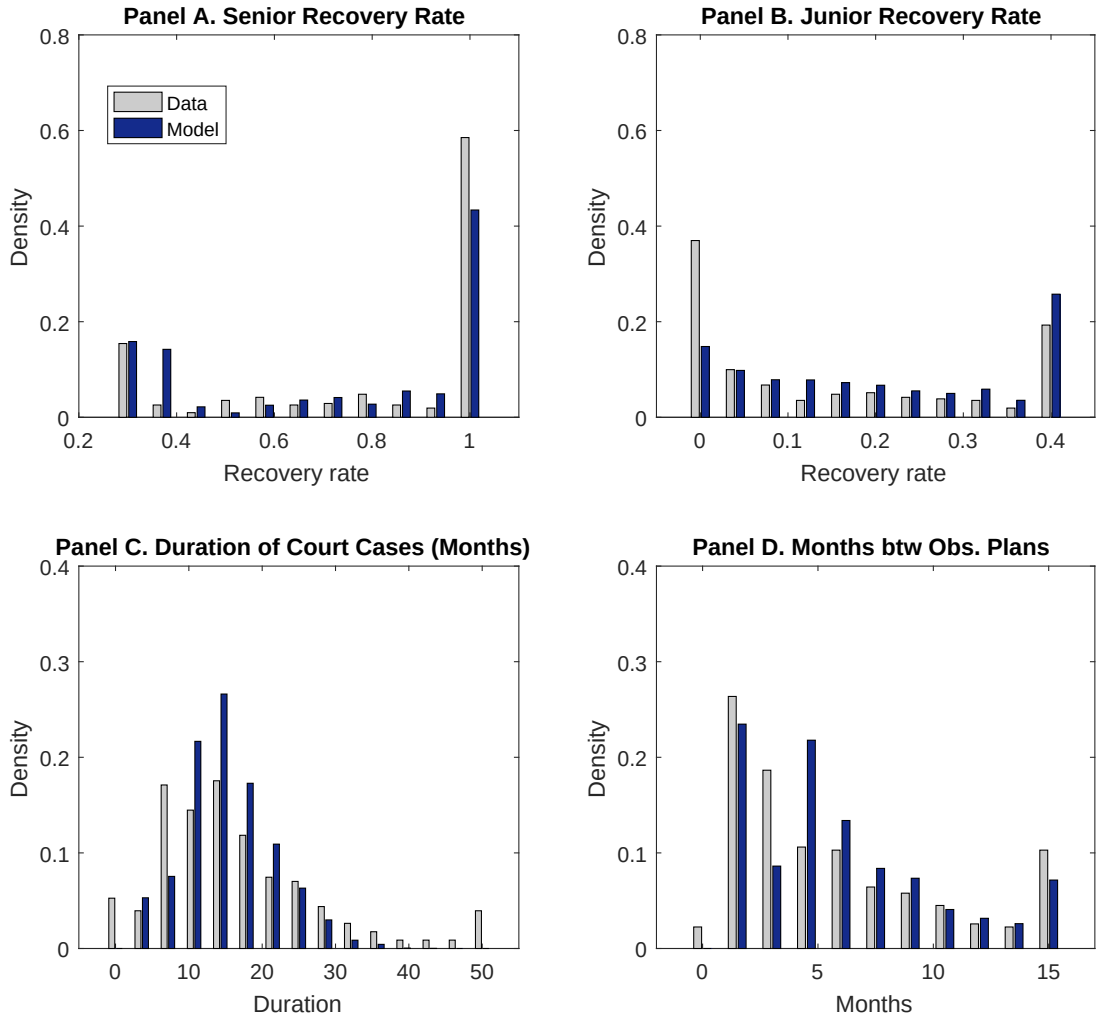


Figure 6: Recovery Rates Versus Case Duration

This figure plots the average total recovery rate versus the bankruptcy case's duration. The total recovery rate equals the total payout to both creditors scaled by their total debt. The red dashed line shows values simulated from the model. The black line shows values from the actual data . The grey shaded region is the 95% confidence interval from the actual data. The first bin contains cases resolved pre-court. The remaining bins contain cases of various lengths that are resolved in court.

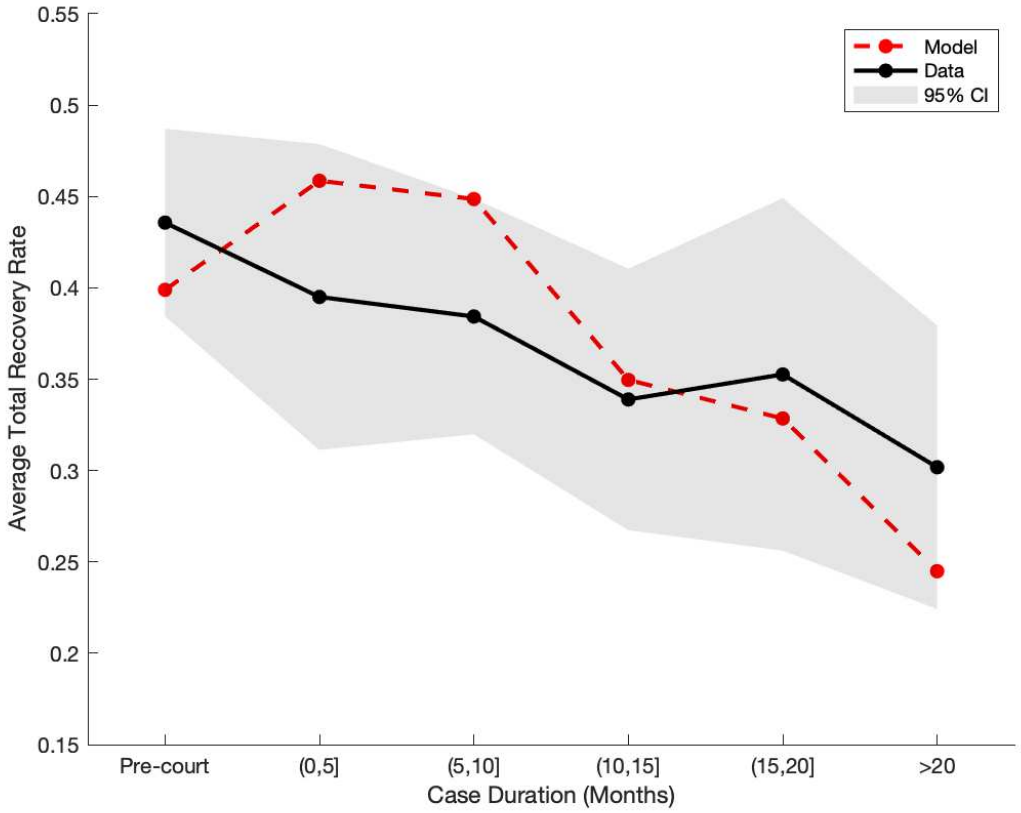


Table 1: Sensitivity of Moments to Parameters

This table shows the sensitivity of model-implied moments (in columns) with respect to model parameters (in rows). To make the sensitivities comparable across parameters and moments, we scale the sensitivities by a ratio of standard errors. The table contains the values of $\frac{dm}{dp}\frac{Stderr(p)}{Stderr(m)}$, where $\frac{dm}{dp}$ is the derivative of simulated moment m with respect to parameter p (evaluated at estimated parameter values from Table 3), $Stderr(p)$ is the estimated standard error for parameter p (also from Table 3), and $Stderr(m)$ is the estimated standard error for the empirical moment m (from Table 2). Moments are defined in detail in Section 2.3. Parameter ρ is the persistence of reorganization value, c_0 is the fixed cost of going to court, $\theta_{S,0}$ and $\theta_{J,0}$ are the initial abilities of the senior and junior creditor, β is the (inverse) speed of creditor learning, and λ_J is the probability that the junior proposes in a given period.

	Months	Ln Duration	Fraction	Recovery, Pre-Court Reorgs.		Frac. Reorg.	Avgs. Across In-Court Reorgs.	
Parameter	Btn. Plans	In Court	In Court	Senior	Junior	Among In-Court	J's Relative Gain	Total Recovery
Months/Period	**6.391**	6.507	0.000	0.000	0.000	0.000	0.000	0.000
ρ	2.321	**5.682**	3.225	-5.250	-0.945	0.318	0.110	-2.067
c_0	1.880	0.403	**-5.314**	-12.448	-0.602	-1.316	1.040	1.093
$\theta_{S,0}$	0.793	0.133	-0.149	**12.106**	-0.141	0.032	-1.722	0.607
$\theta_{J,0}$	2.117	-0.249	-3.397	-6.492	**0.758**	0.490	2.731	1.872
β	1.678	1.175	-3.030	-7.245	-0.638	**-0.528**	-0.951	-0.437
λ_J	-0.806	-1.297	-2.538	-5.671	-0.252	0.174	**1.278**	1.748

Table 2: Model Fit

This table shows how well the model fits the data moments that are targeted in SMM estimation. The $t-$statistics test whether the model moment equals the data moment. Moments are defined in detail in Section 2.3.

Moment	Model	Data	Std. Err.	t-stat.
Averages Across In-Court Cases:				
Ln Months Between Plans	1.861	1.766	0.059	-1.61
Ln Duration (Months)	2.620	2.568	0.058	-0.91
Fraction Reorganized	0.945	0.877	0.022	-3.10
Average Recovery Rates for Pre-Court Reorganizations:				
Senior	0.808	0.878	0.033	2.12
Junior	0.206	0.221	0.027	0.54
Fraction In Court	0.782	0.733	0.025	-1.97
Averages Across In-Court Reorganizations:				
Junior's Relative Gain	-0.065	-0.084	0.016	-1.19
Total Recovery Rate	0.375	0.365	0.020	-0.47

Table 3: Parameter Estimates

This table contains parameter estimates from the SMM estimation.

Parameter	Notation	Estimate	Std. Err.
Months per period		1.310	0.493
Senior's initial ability	$\theta_{S,0}$	0.460	0.102
Junior's initial ability	$\theta_{J,0}$	0.344	0.087
(Inverse) speed of creditor learning	β	12.43	1.26
Persistence of reorganization value	ρ	0.934	0.014
Fixed cost of going to court (%)	c_0	5.08	0.951
Junior's probability of proposing	λ_J	0.898	0.018

Table 4: Quantifying Bankruptcy Inefficiencies and Their Causes

This table compares implications from the estimated model and two counterfactual models. Parameter values shared by all three models are in Table 3. The first counterfactual model assumes symmetric information, meaning each creditor can perfectly observe the other creditor's skill at all points of time. The second counterfactual model assumes a social planner can see both creditors' current skill chooses offer types and offer acceptance/rejection so as to maximize the total expected payout to both creditors combined. All implications are computed from simulated data. Total recovery rate is the total payout to both creditors scaled by their total debt. Average liquidation value is the average of L/D across all firms that are liquidated. Average reorganization value is the average of $\theta_{k,t}V_{h,t}/D$ across firms that are reorganized, where k is the creditor who leads the reorganization and t is the period when the reorganization plan is agreed upon. The average fixed cost of going to court equals c_0 times the fraction of cases going to court. The average cost in court equals the fraction of cases going to court times the average number of periods in court times c_1. The table scales both direct costs by D. Frac. switching from liq. to reorg. is the fraction of all simulated cases that both (1) end in liquidation in the estimated model and (2) end in reorganization in the counterfactual model. Frac. switching from reorg. to liq. is defined similarly. Frac. or reorgs. led by low-skill creditor is the fraction of reorganizations in which the creditor leading the reorganization has lower skill than the other creditor, where both creditors' skill is measured at the time that reorganization is agreed upon.

		Counterfactual Models	
	Estimated	Symmetric	Social
Simulated Statistic	Model	Information	Planner
Average Total Recovery Rate	0.389	0.432	0.546
Average Direct Costs			
Fixed Cost of Going to Court (from c_0)	0.038	0.034	0.015
Costs in Court (from c_1)	0.011	0.004	0.001
Total Direct Costs	0.049	0.039	0.016
Fraction Liquidated	0.103	0.136	0.118
Average Liquidation Value	0.370	0.418	0.371
Fraction Reorganized	0.897	0.864	0.882
Average Reorganization Value	0.446	0.480	0.587
Fraction Resolved Pre-Court	0.250	0.327	0.711
Avg. Duration of Court Cases (Months)	13.7	7.1	4.9
Frac. Switching from Liq. To Reorg.	0.000	0.030	0.022
Frac. Switching from Reorg. To Liq.	0.000	0.063	0.037
Average Months Until Reorganization	10.3	5.4	2.4
Frac. of Reorgs. Led by Low-Skill Creditor	0.288	0.438	0.000

A Model

A.1 Equilibrium

The equilibrium in this game is entirely described by a pair of increasing sequences $\{\ell_{J,t}\}$ and $\{\ell_{S,t}\}$ characterizing the lowest perceived abilities and optimal proposed payments ξ_J and ξ_S.

We focus on equilibria that satisfy the skimming regularity condition (refinement). This is a standard intuitive assumption in the literature of dynamic bargaining (see, e.g. Spier, 1992).

Assumption 1 (Skimming) *The creditors' strategies are such that if type θ' accepts the counterparty's restructuring proposal R with positive probability, then all types $\theta'' < \theta'$ accepts the proposal R with probability* 1*.*

This assumption guarantees that the distribution of types that remain in each period is a truncation of the original distribution. This assumption is quite intuitive: a creditor who faces greater ability to restructure the firm is more likely to decline the counterparty's proposal and lead the restructuring by himself.

A.2 Solution

This is a standard (Markovian) stochastic game with double-sided asymmetric information. We solve the game recursively using the dynamic programming approach.

The end period. The equilibrium is solved recursively by backward induction. The "end period" is the first time t such that $\rho^{t-1}V_h \leq L$. In equilibrium, there is certain probability that the bargaining ends before the scenario $\rho^{t-1}V_h \leq L$ occurs. In that period, the creditors choose to quit the bargaining by liquidating the firm. The APR applies to split whatever is left.

Bellman Equations. Let's consider period t for any $t \geq 0$. The key is to establish the recursive Bellman equations for the "afternoon" continuation values $W_{S,t}(\theta_{S,t}, \ell_{S,t}, \ell_{J,t})$ and $W_{J,t}(\theta_{J,t}, \ell_{S,t}, \ell_{J,t})$ with the endogenous state $(\ell_{S,t}, \ell_{J,t})$ and private state "afternoon" ability $\theta_{J,t}$ or $\theta_{S,t}$. The information about $\theta_{S,t}$ and $\theta_{J,t}$ are revealed in the "afternoon" of period t.

The continuation value of the senior creditor at the end of period t follows the Bellman equation:

$$W_{S,t}(\theta_{S,t},\ell_{S,t},\ell_{J,t}) = \underbrace{(1-\lambda)\times\max\left\{O_{S,t+1},\ \max_{\xi_{S,t}}\mathbb{E}_t^S\left[\widetilde{M}_{S,t+1}(\xi_{S,t})\right]\right\}}_{\text{if } S \text{ proposes in the ``morning''}}$$

$$+\underbrace{\lambda\times\mathbb{E}_t^S\left[\max_{\zeta_{S,t+1}\in\{0,1\}}\widetilde{A}_{S,t+1}(\zeta_{S,t+1})\ \middle|\ \theta_{J,t}\geq\phi_{J,t}\right]\times\mathbb{P}_t^S\{\theta_{J,t}\geq\phi_{J,t}\}}_{\text{if } J \text{ proposes restructuring in the ``morning''}} \quad (4)$$

$$+\underbrace{\lambda\times\mathbb{E}_t^S\left[\max\{O_{S,t+1},U_{t+1}(\theta_{S,t+1})-O_{J,t+1}\}\right]\times\mathbb{P}_t^S\{\theta_{J,t}<\phi_{J,t}\}}_{\text{if } J \text{ decides to liquid in the ``morning''}}, \quad (5)$$

where $\mathbb{E}_t^S$ is the expectation of the senior creditor over $(\theta_{J,t},\theta_{J,t+1})$, i.e. the junior creditor's restructuring abilities in the "afternoon" of periods t and $t+1$, conditional on $\theta_{S,t}$ and $\ell_t = (\ell_{J,t},\ell_{S,t})$. Here, $\zeta_{S,t+1}=1$ means that the senior creditor accepts the offer proposed by the junior in the "morning" of period $t+1$. Here, $\phi_{J,t}$ is the threshold for the junior creditor to choose restructuring or liquidation.

The continuation value of the junior creditor follows the Bellman equation:

$$W_{J,t}(\theta_{J,t},\ell_{S,t},\ell_{J,t}) = \underbrace{\lambda\times\max\left\{O_{J,t+1},\ \max_{\xi_{J,t}}\mathbb{E}_t^J\left[\widetilde{M}_{J,t+1}(\xi_{J,t})\right]\right\}}_{\text{if } J \text{ proposes in the ``morning''}}$$

$$+\underbrace{(1-\lambda)\times\mathbb{E}_t^J\left[\max_{\zeta_{J,t+1}\in\{0,1\}}\widetilde{A}_{J,t+1}(\zeta_{J,t+1})\ \middle|\ \theta_{S,t}\geq\phi_{S,t}\right]\times\mathbb{P}_t^J\{\theta_{S,t}\geq\phi_{S,t}\}}_{\text{if } S \text{ proposes restructuring in the ``morning''}} \quad (6)$$

$$+\underbrace{(1-\lambda)\times\mathbb{E}_t^J\left[\max\{O_{J,t+1},U_{t+1}(\theta_{J,t+1})-O_{S,t+1}\}\right]\times\mathbb{P}_t^J\{\theta_{S,t}<\phi_{S,t}\}}_{\text{if } S \text{ chooses to liquid in the ``morning''}}, \quad (7)$$

where $\mathbb{E}_t^J$ is the expectation of the junior creditor over $(\theta_{S,t},\theta_{S,t+1})$, i.e. the senior creditor's restructuring abilities in the "afternoon" of periods t and $t+1$, conditional on $\theta_{J,t}$ and $\ell_{S,t}$. Here $\phi_{S,t}$ is the threshold for the senior creditor to choose restructuring or liquidation, and it is known to agents at the end of period t.

Senior Creditor's Payoffs in Period $t+1$. The payoffs of the senior creditor are both realized in the "afternoon" of period $t+1$.

If the senior creditor proposes in the "morning" of period $t+1$, the payoff to the senior creditor in the "afternoon" of period $t+1$, conditional on the choice $\xi_{S,t}$, is described as follows:

$$\widetilde{M}_{S,t+1}(\xi_{S,t}) = \underbrace{[U_{t+1}(\theta_{S,t+1}) - \xi_{S,t}]\,\mathbf{1}\{W_{J,t+1}(\theta_{J,t+1}, \ell_{S,t+1}, \ell_{J,t+1}) \leq \xi_{S,t}\}}_{\text{if } J \text{ accepts the offer}} \tag{8}$$

$$+ \underbrace{W_{S,t+1}(\theta_{S,t+1}, \ell_{S,t+1}, \ell_{J,t+1})\mathbf{1}\{W_{J,t+1}(\theta_{J,t+1}, \ell_{S,t+1}, \ell_{J,t+1}) > \xi_{S,t}\}}_{\text{if } J \text{ does not accept the offer}}, \tag{9}$$

where the decision variable $\xi_{S,t}$ depends only on the senior creditor's information up to the end of period t. In the "afternoon" of period $t+1$, the junior creditor observes $\xi_{S,t}$ and $\theta_{J,t+1}$ and chooses to accept the offer with $\xi_{S,t}$ (i.e. $\zeta_{J,t+1} = 1$) if and only if $W_{J,t+1}(\theta_{J,t+1}) \leq \xi_{S,t}$.

The continuation values $W_{S,t+1}$ and $W_{J,t+1}$ are equilibrium value functions, and the common information expected by the creditors. The key is to track the endogenous state variables ℓ_t's evolution. The evolution of ℓ_t depends on the realization of proposing opportunity and the endogenous choice variables ξ and ζ. But importantly, $\ell_{S,t+1}$ does not depend on the junior creditor's choice $\zeta_{J,t+1}$, and $\ell_{J,t+1}$ does not depend on the senior creditor's choice $\zeta_{S,t+1}$.

When the senior creditor receives the proposal opportunity in the "morning" of period $t+1$, it holds that $\ell_{S,t+1} = \theta_{S,t}$ and $\ell_{J,t+1} = W^{-1}_{J,t+1}(\xi_{S,t}; \theta_{S,t})$; that is, $\xi_{S,t} = W_{J,t+1}(\ell_{J,t+1}, \theta_{S,t}, \ell_{J,t+1})$.

When is the updating information about $\ell_{J,t+1}$ and $\ell_{S,t+1}$ realized to two agents? The following is the description if the senior creditor receives the proposing opportunity:

- The update of $\ell_{S,t+1}$ is realized to the junior creditor right after he sees the proposal $\xi_{S,t}$. The update is perfectly perceived by the senior creditor at the very beginning of period $t+1$ right after he receives the proposing opportunity.

- The update of $\ell_{J,t+1}$ is realized to the junior creditor right after he sees the proposal $\xi_{S,t}$. The update is perfectly perceived by the senior creditor at the very beginning of period $t+1$ right after he receives the proposing opportunity.

Therefore, in the equilibrium, the updated belief $\ell_{t+1} = (\ell_{S,t+1}, \ell_{J,t+1})$ only depends on $\theta_{S,t}$ and $\ell_t = (\ell_{S,t}, \ell_{J,t})$, and thus known to the senior creditor at the very beginning of period $t+1$, if the senior creditor receives the opportunity to propose in the "morning".

If the junior creditor proposes in the "morning" of period $t+1$, the payoff to the senior creditor in the "afternoon" of period $t+1$, conditional on the choice $\xi^*_{J,t}$ and thus $\ell^*_{S,t+1}$, is described as follows:

$$\max_{\zeta_{S,t+1}\in\{0,1\}} \widetilde{A}_{S,t+1}(\zeta_{S,t}) = \underbrace{\xi^*_{J,t}\mathbf{1}\{W_{S,t+1}(\theta_{S,t+1},\ell_{S,t+1},\ell_{J,t+1}) \leq \xi^*_{J,t}\}}_{\text{if } S \text{ accepts the offer: } \zeta_{S,t+1}=1} \tag{10}$$

$$+ \underbrace{W_{S,t+1}(\theta_{S,t+1},\ell_{S,t+1},\ell_{J,t+1})\mathbf{1}\{W_{S,t+1}(\theta_{S,t+1},\ell_{S,t+1},\ell_{J,t+1}) > \xi^*_{J,t}\}}_{\text{if } J \text{ does not accept the offer: } \zeta_{S,t+1}=0}. \tag{11}$$

Junior Creditor's Payoffs in Period $t+1$**.** The setup is similar to the senior creditor's payoffs.

B Data

B.1 Estimation of missing liquidation values

This section describes how we estimate the missing values of L, the firm's liquidation ratio.

B.2 Maximum restructuring value ($V_{h,0}$)

We estimate each firm's potential Tobin's Q to be the 50th percentile Q in the same industry and year. To compute this measure, we first combine all observations from a given three-digit SIC industry across all years, subtract each year's median from Q, compute the 50th percentile value of these median-adjusted values, and finally add back the median from each industry×year. The rationale behind pooling and adjusting for yearly medians is to more accurately estimate the 50th percentiles by avoiding tiny subsamples. A firm's Q ratio is defined as market equity plus total debt plus preferred stock liquidating value minus deferred taxes and investment credit, all divided by total assets (as in Lemmon et al. 2008).

We adjust for medians, whereas Edmans et al. (2012) adjust for means. We find that means within industries years are highly sensitive to outliers, even if we were to winsorize our measures. We also depart from Edmans et al. (2012) by using the 50th rather than 80th percentile. We use the 50th percentile because it is unrealistic that a highly impaired, bankrupt firm would quickly reach a high valuation. Our results are not sensitive to the choice of percentile. With a different percentile, the creditors' estimated initial ability level and speed of learning would adjust to continue fitting the data, and the paths of reorganization value $\theta_{k,t}V_{h,t}$ are largely unchanged.

C Details on SMM estimation

We use SMM to estimate the vector parameters $\Theta = \{\rho, \beta, \theta_{S,0}, \theta_{J,0}, c_0, \lambda, \text{months per period}\}$. The SMM estimator $\hat{\Theta}$ searches for the parameter values that minimize the distance between the data moments and the model-implied moments:

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \left(\hat{m} - \frac{1}{S} \sum_{s=1}^{S} \hat{m}^s(\Theta) \right)' W \left(\hat{m} - \frac{1}{S} \sum_{l=1}^{S} \hat{m}^s(\Theta) \right).$$

Vector $\hat{m}$ contains the moments estimated from data, and $\hat{m}^s(\Theta)$ is the corresponding vector of moments estimated from the sth sample simulated using parameter vector Θ. W is the efficient weighting matrix, equal to the inverse of the estimated covariance of moments m. The efficient weighting matrix W is constructed using influence functions, following . We cluster by year interacted with industry. Specifically, we allow two cases' error terms to be correlated if the cases are from the same two-digit SIC industry and their years of filing differ by less than two years. Michaelides and Ng (2000) find that using a simulated sample 10 times as large as the empirical sample generates good small-sample performance. We choose $S = 40$ simulated samples to be conservative.

When simulating data, we feed observed values of the parameters D_J, $V_{h,0}$, and L into the model. One challenge is that these three parameters vary across our sample cases. (Note that since we normalize $D = 1$, the fraction of debt held by the senior is just $1 - D_J$.) Ideally, we would solve the model for each sample case's specific values of $\{D_J, V_{h,0}, L\}$, simulate data from

each of those model cases, then combine simulated cases into a single simulated data set. That approach is infeasible, however, because there are more than 300 cases in our sample, and solving the model even once takes considerable time. We therefore take an intermediate approach that captures a large part of the heterogeneity in our sample. We use a K-means algorithm to assign each sample case to one of ten clusters, where each cluster contains cases that share similar values of $\{D_J, V_{h,0}, L\}$. K-means is one of the simplest and most commonly used unsupervised learning algorithms for clustering problems. The method goes back to MacQueen (1967) and Hartigan (1975), and today it is quite standard (see, e.g., Chapters 13 and 14 of Hastie et al., 2009). The K-means algorithm has been used recently in the finance literature by, for example, Grieser and Liu (2018). The choice of ten is arbitrary, and this number can be increased with the help of more computing power. We record the mean values of $\{D_J, V_{h,0}, L\}$ for each clusters. When simulating data off the model, we solve the model for each of these ten median values of $\{D_J, V_{h,0}, L\}$, we simulate data off each of the ten model solutions, and we sample the ten simulations in proportion to the empirical frequency of each cluster.