Dissecting Bankruptcy Frictions

Winston Wei Dou

Lucian A. Taylor

Wei Wang

Wenyu Wang*

December 18, 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 4%, and eliminating conflicts of interest would increase them by an additional 18%. Without these frictions, an extra 14% of cases would be resolved before going to court, and the remaining court cases would be 73% 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

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 efficient is the U.S. bankruptcy system for corporate restructuring? The answer clearly matters for insolvent firms, but it also matters for healthy firms' capital structure choices and securities pricing. 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. Some cases should be resolved quickly and out of court, but instead they experience long court battles, which amplify these costs ("excess delay").

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 senior and junior creditors. 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 then 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 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,

---

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

quantifying the system's overall efficiency requires an economic model.

We overcome these challenges by structurally estimating a new 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. 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 creditors will split the proceeds. 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 learning and finding a better reorganization plan. Conflicts of interest can also lead creditors to 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. Relative to the literature, our sample has the most comprehensive coverage of 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 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 payouts, especially for cases resolved early, help identify the initial quality of creditors' reorganization plans. The way in which payouts are split between creditors helps identify their relative bargaining power. Data on the length of court cases and likelihood of reorganization help to disentangle how fast reorganization values decay and how fast creditors learn. Overall, the model does a good job fitting the observed distributions of creditors' recovery rates, the timing of outcomes, the negative relation between debt recovery and case duration, the frequency of outcomes (liq-

uidation 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 4% if we remove asymmetric information, and it increases an additional 18% if we also remove conflicts of interest. The frictions together destroy about $11.4 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 modest inefficiency, and conflicts of interest among creditors generate a significant inefficiency.

We find that excess delay is the primary source of inefficiency. 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 3 percentage points (from 70% to 67%) when we remove asymmetric information, and further decreases by 11 percentage points (to 56%) when we remove conflicts of interest. The average duration of the remaining cases decreases from 16.7 to 13.4 months without asymmetric information, and to 4.5 months without conflicts of interest. In other words, removing these frictions would reduce court cases' duration by 73%. Less delay results in lower legal, accounting, and other direct costs. More important, less delay results in higher reorganization values, because there is less decay in the going-concern value. Surprisingly, the inefficiencies from excess liquidation and excess continuation are quantitatively small, together making up just 6.5% of the total inefficiency we find.

How do the frictions cause excess delay? Asymmetric information increases uncertainty about the counterparty's type. This uncertainty leads creditors to make lowball offers out of a

precautionary motive to avoid overpaying. These lowball offers are often rejected, causing excess delay. Asymmetric information also results in inefficient screening, because rejecting a lowball offer reveals little about one's type, which makes information asymmetry persist. Conflicts of interest cause further delay. If the two creditors have similar reorganization skill, they have an incentive to reject each other's offers in hopes of making a counteroffer and extracting better deal terms in the future. More simply, by "playing tough" with each other, creditors delay the case, allowing the direct and indirect costs of bankruptcy to grow.

Related Literature. Discussions and theories of bankruptcy inefficiencies extend back at least to Jensen and Meckling (1976), Bulow and Shoven (1978), Baird (1986), Jackson (1986), Bebchuk (1988), Giammarino (1989), Gertner and Scharfstein (1991), Aghion et al. (1992), and Bradley and Rosenzweig (1992). We extend this line of theories by making the total bankruptcy payout creditor-specific, which helps create a model that can be taken to the data. We then quantify the inefficiencies and their economic sources.

Several papers provide reduced-form evidence that bankruptcy frictions exist. 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. An important exception is Stromberg (2000), who models and quantifies inefficiencies in Swedish cash auction

4

bankruptcies. These papers also highlight 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), Pulvino (1998), 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), Glover (2016), and Dou et al. (2019), 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 costs? Also, by studying counterfactuals without frictions, we provide a benchmark for judging whether the observed bankruptcy 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.

Three 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) and Antill (2019) estimate the losses from inefficient liquidations in bankruptcy. Similar to us, Jenkins and Smith (2014) find that excess liquidation produces a small inefficiency. Antill (2019) instead finds a large inefficiency from excess liquidation. Our papers differ in

several other ways. Unlike Jenkins and Smith (2014) and Antill (2019), we model the dynamics of bankruptcy cases and find that excess delay is a major source of inefficiency. Whereas we focus on how conflicts of interest and asymmetric information create inefficiencies, Jenkins and Smith (2014) only explore conflicts of interest, and Antill (2019) does not model the sources of any inefficiency. Antill (2019) instead estimates a Roy model with data on random judge assignment, which has the benefit of imposing minimal structure on the problem.

1   Baseline 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 whose 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 Rubinstein (1982), Bebchuk (1984), Chatterjee and Samuelson (1987), and Spier (1992), among others.

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 creditors are bargaining with each other.[2] 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 reach an agreement out of court in $t = 0$, then bargaining continues in court in each period $t = 1, 2, ...$ until creditors reach an agreement. The firm incurs a fixed cost from going to court and variable costs that depend on the length of the court case. Specifically, the firm incurs a one-time fixed cost of $c_0 D$ if bargaining is not resolved pre-court. Parameter c_0 captures both direct

---

[2]Supporting this assumption, recent empirical evidence shows that equity holders are typically wiped out and lose their bargaining power in cases filed since the turn of the century. For example, Bharath et al. (2014) and Kim (2018) document that shareholder recovery in bankruptcy has experienced a secular decline since the Bankruptcy Reform Act of 1978, and it declined to near zero by the early 2000s. This time-series pattern is consistent with the general trend of strengthening creditor control in bankruptcy (Ayotte and Morrison, 2009; Ayotte et al., 2013).

costs, such as legal and other professional fees from going to court, as well as indirect costs, such as the loss of customers or employees. The cost c_0 should be interpreted as the cost triggered by going to court minus the sum of (1) the cost of coordinating creditors and achieving a settlement out of court, and (2) any benefits of going to court (e.g., improved liquidity, as in Ayotte and Skeel, 2013). For example, $c_0 = 0$ implies that the cost of going to court exactly equals the cost of coordinating creditors and reaching an agreement pre-court. In addition, the firm incurs variable direct costs of $c_1 D$ during each period the case stays in court. Consistent with this assumption, LoPucki and Doherty (2004) find that professional fees increase with case duration. At the end of period t, accumulated costs from c_0 and c_1 are denoted $C_t = \mathbf{1}_{\{t>0\}} (c_0 + c_1 t) D$.

The outcome for the firm is either reorganization or piecemeal liquidation, with bargaining resolved either before court or in court. In a liquidation, the firm's assets are sold for a known amount L, 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 $O_{S,t} = \min(L - C_t, D_S)$, and the junior creditor collects the residual, $O_{J,t} = L - C_t - \min(L - C_t, D_S)$. APR thus creates an asymmetry between the two creditors.

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

$$V_t = \rho^{t-1} V_0, \quad t \geq 1. \tag{1}$$

A lower ρ indicates larger per-period indirect costs of bankruptcy.

Leading a reorganization requires skill. This skill reflects the quality of the reorganization plan and the creditor's managerial ability. We allow the senior and junior creditors to have

different levels of reorganization skill, which makes the total surplus creditor-specific. We allow these skill levels to change randomly over time. Specifically, the senior and junior creditors have skill $\theta_{S,t}$ and $\theta_{J,t}$, respectively, at time t. Both θ values lie in the interval $[0,1]$. If creditor $k \in \{S, J\}$ leads the reorganization in period t, then the firm's value, and hence the combined payout to both creditors upon emergence, is

$$U_t(\theta_{k,t}) \equiv \theta_{k,t}V_t - C_t = \theta_{k,t}V_0\rho^{t-1} - C_t, \quad \text{with } k \in \{S, J\}. \tag{2}$$

This assumption implies that higher skill produces a higher reorganization value.

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 (i.e., $\theta = 1$), consistent with the evidence of Kahl (2002). We also include learning to allow the possibility that some amount of delay is efficient. We capture these effects by assuming $\theta_{J,t}$ and $\theta_{S,t}$ follow independent, increasing Markov processes. Specifically, if a creditor's reorganization skill is θ_t at time t, then his reorganization skill 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 reorganization skill. We choose the generalized beta distribution for a few reasons. It guarantees that next period's skill level is between the current level and the maximum of 1. (We do not allow skill to decrease, because a creditor would discard any inferior plans it finds.) The speed of learning

---

[3]Skill could also increase if a high-skill investor buys the stake of a low-skill creditor. While such a transaction would be public knowledge, the new investor's skill at reorganizing this specific firm is arguably private information.

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 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 generalized beta distribution significantly improves tractability.[4]

Figure 1 illustrates our assumptions about learning and reorganization values, using estimated model parameters from Section 3.2. The top panel shows how creditors' median skill levels increase over time. With these parameter values, skill increases quite slowly. Not shown in the figure, shocks to reorganization skill 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_t, which this figure normalizes to 1 at $t = 1$. The lower lines plot the medians of $\theta_{k,t} V_t$, the firm's reorganization value under each creditor k's skill. 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.

Figure 2 illustrates how bargaining works each period, including the pre-court period. Each period is divided into two subperiods, "morning" and "afternoon." Proposals are made in the morning, responses in the afternoon. At the beginning of period t, the values of $\theta_{S,t}$ and $\theta_{J,t}$ are private information, and the counterparties' beliefs about them are $F_\beta(\theta_{S,t}|\ell_{S,t})$ and $F_\beta(\theta_{J,t}|\ell_{J,t})$, respectively. The lower bounds $\ell_{S,t}$ and $\ell_{J,t}$ that characterize the beliefs are publicly known. One creditor, say creditor $k \in \{S, J\}$, 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 $1 - \lambda_J$.[5] 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, creditor k (for example) proposes reorganizing the firm under her own

---

[4]The beta distribution guarantees the property of increasing hazard rates (e.g., Bebchuk, 1984; Nalebuff, 1987) and invariant truncated conditional distributions. More precisely, the conditional distribution of θ_{t+1} given that $\theta_{t+1} > \ell$ for any $\ell \in [\theta_t, 1)$ is $F_\beta(\theta_{t+1}|\ell)$.

[5]The random-proposal scheme is common in bankruptcy models (e.g., Posner and Kordana, 1999; Eraslan, 2008; and Antill and Grenadier, 2019) as well as in the game-theoretic literature on dynamic bargaining models (e.g., Merlo and Wilson, 1995, 1998). Our assumption could reflect randomness in which creditor is first to prepare a detailed proposal, which creditor's proposal the judge supports, or which creditor's plan the debtor firm supports during the exclusivity period.

plan and paying the counterparty $\xi_{k,t}$, with the remaining value going back to herself. The subscript t on $\xi_{k,t}$ means $\xi_{k,t}$ depends on information up to the beginning of period t. The reorganization proposal reveals creditor k's reorganization skill $\theta_{k,t}$, because, for example, the proposal includes a detailed business plan.[6] Based on this information, the responding creditor $\overline{k}$ updates his belief about $\theta_{k,t+1}$ to $F_\beta(\theta_{k,t+1}|\ell_{k,t+1})$ with $\ell_{k,t+1} = \theta_{k,t}$. Meanwhile, the proposing creditor k keeps the same belief about $\theta_{\overline{k},t+1}$, characterized by $F_\beta(\theta_{\overline{k},t+1}|\ell_{\overline{k},t})$. Creditor k can also propose liquidating assets piecemeal for the total payout of $L - C_t$, which is split according to APR. Liquidation proposals automatically end the bargaining, but they do not necessarily result in a liquidation. If the responding creditor $\overline{k}$ prefers not to liquidate the firm, he can instead reorganize the firm under his own reorganization plan as long as he pays the proposing creditor k what she would receive upon a liquidation, under APR. Finally, creditor k can propose waiting by making a reorganization offer that will be rejected for sure (i.e., by proposing a very low $\xi_{k,t}$), effectively moving the game to the next period.

When the afternoon begins, the reorganization skills change from $\theta_{S,t}$ and $\theta_{J,t}$ to $\theta_{S,t+1}$ and $\theta_{J,t+1}$, respectively.[7] Based on his updated reorganization skill $\theta_{\overline{k},t+1}$, the responding creditor $\overline{k}$ weighs how much he would gain by accepting the proposal (i.e., the payment $\xi_{k,t}$) against how much he would get by declining the proposal and waiting (i.e., the continuation value, denoted $W_{\overline{k},t+1}$). If $\xi_{k,t} \geq W_{\overline{k},t+1}$, the responding creditor accepts the offer and the game ends; otherwise, he rejects the offer and the game moves to the next period. When the responding creditor $\overline{k}$ calculates $W_{\overline{k},t+1}$, he takes into account that rejecting the offer would partially reveal his skill, $\theta_{\overline{k},t+1}$. In other words, $\overline{k}$ internalizes the "screening effect" of rejecting an offer. Specifically, if $\overline{k}$ rejects the offer, the lower bound that characterizes the proposing creditor k's belief about

---

[6]Supporting this assumption, the Online Appendix contains an extended model in which proposals perfectly and endogenously reveal skill levels. In the extension, there is a small probability that judges "cram down" a plan. Cram down serves as a commitment device that guarantees the separating equilibrium in which proposing creditors perfectly reveal their skill.

[7]We allow skill levels to change between proposal and response, because in practice there is significant delay between these events, for two specific reasons. First, the bankruptcy code and rules introduce a delay of at least 28 days between offer (the filing of a plan and disclosure statement) and response. After the filing of a Chapter 11 case, the parties of interest (e.g. creditors, the trustee) should be given at least 28 days of notice prior to the court holding a hearing on confirmation of a plan and approval of disclosure statement (see Rule 2002 (b), Rule 3017, Rule 3018 of United States Bankruptcy Code & Rules (2017 edition)). Second, bankruptcy reorganization/liquidation plans and disclosure statements are lengthy legal documents that typically contain hundreds of pages. It takes significant time for a counterparty to study and digest these documents.

$\theta_{\overline{k},t+1}$ is updated to $\ell_{\overline{k},t+1} = \ell(\xi_{k,t})$. In turn, the proposing creditor k internalizes the screening effect when choosing the proposal $\xi_{k,t}$ at the beginning of the period, and she understands the equilibrium belief updating function $\ell(\xi)$ for any proposal ξ she would make.[8]

Each creditor is risk neutral and maximizes its expected payout. 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. The creditors behave rationally at every node of the game given beliefs, and their beliefs are derived from the equilibrium strategies by Bayes' rule. Technical details on the equilibrium are described in Appendix A.

1.2 Discussion

Before describing the model's predictions, we address potential concerns about the setup and omitted factors.

We do not 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 instead respond to motions made by the debtor or creditors. Judges must approve a proposed plan before sending it out for a vote, but judges apply a fairly lenient feasibility standard.[9] 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 reorganization proposals require the other creditor's approval. Judges in practice can 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. Finally, we do not model judge-specific preferences. While

---

[8]The existence of $\ell(\cdot)$ relies upon regularity conditions including increasing hazard rates, which is common for screening models.

[9]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.

such preferences matter in small bankruptcy cases, there is little evidence that they matter in the large bankruptcies we study.[10]

Our data include 33 cases in which the firm sells all or substantially all of its assets through Section 363, meaning the firm sells its business without requiring the junior creditor's approval. Our model accommodates these outcomes. Recall that the senior creditor in our model can make a "liquidation proposal," which unilaterally ends the bargaining. (The junior creditor can also make a liquidation proposal, but we find this does not occur in equilibrium.) Liquidation proposals therefore capture several key features of Section 363 sales: They unilaterally end the case, they are generally proposed by the senior creditor (Ayotte and Morrison, 2009), and they cannot be vetoed by the junior creditor. Section 2.1 explains how we classify Section 363 sales in our data.

U.S. bankruptcy law encourages creditors to share their information through direct, private communication, which could reduce information asymmetry about reorganization skill. Our baseline model does not include private communication, because conflicts of interests may prevent the creditors from truthfully revealing their information. In Section 5, for robustness, we estimate an extended model in which private communication can reveal skill, and we show that our results do not change significantly.

We do not model the initial 120-day exclusivity period during which only the debtor firm can submit a plan. Although these plans are nominally submitted by the company's management, there is usually at least one creditor class that supports and sponsors the plan. We therefore view these plans as effectively being proposed by one creditor class.

We do not allow a creditor to combine his own financial plan with the counterparty's business plan. Supporting this assumption, *formal* plans submitted by the senior and junior typically look quite different. During *informal* conversations, however, a creditor will sometimes propose com-

---

[10]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. 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, and they show that judge fixed effects explain little to no variation in Chapter 11 outcomes.

bining the counterparty's business plan with his own financial plan. Since those conversations lack the commitment that comes with a formal court proposal, such "mimicking proposals" are simply cheap talk that need not affect how creditors actually behave in court. More important, a creditor cannot attain the counterparty's skill level by simply copying his written business plan, because skill also depends on how well the plan is executed, and execution skill arguably varies across creditors.

Gilson et al. (2000) and Ayotte and Morrison (2018) show disputes on corporate valuation in bankruptcy, arising from creditors strategically over- or under-valuing the firm. This behavior is consistent with our model, in which each creditor tries to grab as much of the surplus as possible. For the senior (junior) creditor, proposing a lower (higher) firm valuation is equivalent to requesting a larger share of the surplus.

Since our model begins with a firm that has already reached insolvency, this paper can only quantify the ex post efficiency of the bankruptcy process. Some degree of ex post inefficiency could be efficient ex ante, for example, by inducing firms to borrow less (e.g., Glover, 2016 and Kim, 2018). Quantifying the ex ante efficiency of Chapter 11 is an important area for future work.

Neither our model nor our data includes private workouts, in which equityholders and creditors reach an agreement without filing for bankruptcy. The bargaining in private workouts looks quite different from our model, because private workouts involve the shareholders, the APR rule need not hold, and voting thresholds are higher, which leads to greater coordination and holdout problems. Also, firms that reach a private workout tend to be less distressed, making them hard to compare to our sample's firms. One can think of our model as describing firms that fail to reach a private workout. Our results apply only to that selected group of firms.

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 the solution, including its derivation.

We solve the model numerically via dynamic programming. Solving the model entails finding

the two creditors' value functions and policy functions. For creditor k, the state variables are creditor k's true reorganization skill ($\theta_{k,t}$), creditor k's skill lower bound ($\ell_{k,t}$) as perceived by the counterparty $\bar{k}$, counterparty $\bar{k}$'s reorganization skill lower bound ($\ell_{\bar{k},t}$) as perceived by creditor k, and the period (t). The bankruptcy is guaranteed to be resolved by some period T defined by $\rho^T V_0 < L$. 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 describe the creditors' optimal offers, starting with their business plans (i.e., liquidation, waiting, or reorganization). Figure 3 plots the business plans for different combinations of reorganization skill. The horizontal axis denotes the proposing creditor's true skill, and the vertical axis denotes the proposing creditor's perception of the responding creditor's skill. The red areas represent the regions in which creditors make waiting offers, the black areas represent liquidation offers, and the light blue areas represent reorganization offers. The top and bottom subplots show the offers made at $t = 0$ and $t = 2$, respectively.

We see a few interesting patterns. If the senior creditor has very high reorganization skill, he always proposes reorganization, regardless of the junior's perceived reorganization skill. The value of waiting is low for a high-skill senior creditor, because waiting is costly and skill does not have much more room to grow. If the senior creditor has intermediate skill, the value of waiting increases, because his skill has more room to grow. The incentive to make a waiting offer is especially strong when the two creditors' skill levels are similar, because the senior creditor expects it to be difficult to convince the junior to compromise at this stage, and he hopes to gain higher skill through learning in the future. If the senior creditor's skill is very low, he prefers making a liquidation offer, especially when the junior's skill is high. The senior creditor in this case finds the protection provided by APR in liquidation to be more valuable than waiting longer, and further delay in court is likely to eat up its share in the firm.

For the junior creditor, we again see that high reorganization skill induces more reorganization offers and low reorganization skill induces more waiting offers. However, the junior creditor never proposes liquidation. This is because the liquidation value L is assumed to be lower than D_S in our example, so liquidation will leave the junior creditor with a zero payout. Therefore, the

14

junior creditor strongly prefers reorganization to liquidation.

Comparing the pre-court period (the top two panels) to the in-court period (the bottom two panels), we also find that the senior creditor is more likely to make a liquidation offer in the pre-court period. This happens because the liquidation values do not improve over time, and if the firm is liquidated pre-court, it saves the fixed cost of going to court ($c_0 D$).

Figure 4 illustrates creditors' optimal financial offers, which describe how the proposing creditor offers to split the total payout from reorganization between the two creditors. We define the value split as the fraction of the total payout offered to the responding creditor. The top two panels illustrate how the value split proposed by the senior creditor varies with the senior's own skill and the junior's perceived skill, and the bottom two panels illustrate how the value split proposed by the junior creditor varies with the junior's own skill and the senior's perceived skill.

We find that the value split decreases in the proposer's own skill and increases in the responder's perceived skill. This prediction is expected, because high relative skill increases the creditor's relative bargaining power. Even though the top and bottom panels show similar patterns, the fraction offered by the senior to the junior is overall lower than the fraction offered by the junior to the senior. This asymmetry occurs because the senior creditor is protected by APR in liquidation, and liquidation is an alternative (i.e., an outside option) to reorganization.

Finally, to illustrate the types of inefficiencies that arise in the model, Figure 5 shows two cases simulated from the model. The left panels show each creditor's reorganization skill over time, and the right panels show the type of offer made in each period t as well as the total recovery rate if the firm were reorganized by the proposing creditor in period t. Circles indicate waiting offers and squares represent reorganization offers. Red solid represents the junior creditor, blue dashed the senior.

Row 1 shows a simulation in which the senior starts out with a higher reorganization skill than the junior. The senior creditor gets to propose in the first few rounds. As shown in the right panel, the senior creditor initially makes waiting offers in hopes that his skill will increase. His skill does increase, leading him to make a reorganization offer at $t = 2$. The junior creditor,

seeing that his own skill (and hence bargaining position) is weak, accepts the senior's offer. This example illustrates that some delay can be efficient, because it allows learning.

The second simulation (row 2) illustrates how inefficiency can arise in the bargaining process. We start from the simulation above but make the junior's skill the same as the senior's, so the blue "×" markers and red "+" markers overlap. With higher skill, the junior now rejects the senior's reorganization proposal made at $t = 2$. Bargaining now ends at $t = 4$ with a significantly lower total recovery rate. This is an example of excess delay, meaning the firm would have been better off reorganizing earlier. Comparing the two simulations above, we see that when the proposer has high skill, it is not necessarily good to have a responder with high skill. Fierce competition between creditors with comparable skill may lead to significant delay, which entails high direct costs and decay in the going-concern value.

Excess delay in our model, including the simulation above, comes from both asymmetric information and conflicts of interest. Both frictions lead creditors to play tough with each other, delaying the case and potentially destroying part of the total surplus. Asymmetric information creates uncertainty about the counterparty's skill. This uncertainty, combined with the reluctance to overpay the counterparty, leads creditors to make precautionary low-ball offers. Such offers are typically rejected, causing delay. Making matters worse, almost any creditor, even one with low skill, would reject a low-ball offer, so low-ball offers are not very helpful in screening the counterparty's skill. Low-ball offers therefore make the asymmetric-information problem persist.[11]

Conflicts of interest lead to further delay. 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. Inefficient delay occurs when a creditor rejects a "good" proposal in hopes of making her own proposal next period, which would allow her to capture a bigger share of the surplus. Rejecting a "good" proposal can be privately optimal even if the creditor knows that delay will

---

[11]Several theory papers show that asymmetric information can lead to delay in bargaining. In Admati and Perry (1987), which features one-sided private information, agents delay making offers to signal the strength of their type and avoid overpaying. Cramton (1992) extends their analysis to allow two-sided private information. Our mechanism is closer to that in Chatterjee and Samuelson (1987), which also features two-sided private information. In both our models, players delay by making low-ball offers in order to avoid overpaying. Players in Chatterjee and Samuelson (1987) also make low-ball offers to signal their bargaining power, but this effect is negligible in our model, and in fact our players have a small, opposing screening motive.

destroy part of the total surplus.

Rubinstein (1982) shows that conflicts of interest do not necessarily cause excess delay in bargaining, because players will instead bargain immediately to achieve an efficient outcome. Conflicts of interest lead to inefficient delay in our model due to two features, both of which are designed to match the U.S. bankruptcy system. First, reorganization skill is creditor-specific, so players are not bargaining over a common surplus. This feature violates the "separation principle" of Merlo and Wilson (1998).[12] Second, skill and bargaining power are not perfectly linked. For example, suppose the creditor with high reorganization skill would like to settle immediately—and this would be socially efficient—but the low-skill creditor gets the opportunity to propose. This creditor cannot propose a plan that satisfies the high type, so the low type instead proposes waiting in order to maximize her own payout, thereby causing inefficient delay.[13]

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. Table 1 provides summary statistics. To construct our 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 that 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.[14] We cross-check with New Generation

---

[12]The separation principle is the following: The set of equilibrium agreement states depends exclusively on the total outcome process, independently of how the proposer is selected. Merlo and Wilson (1998) show that the separation principle ensures the efficiency of immediate agreement.

[13]The Online Appendix provides simple numerical examples illustrating why conflicts of interest produce excess delay in our model but not in previous models such as Rubinstein (1982), Bebchuk and Chang (1992), and Merlo and Wilson (1998).

[14]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).

Research's bankruptcydata.com as of March, 2016 on case status and outcome. After removing dismissed cases and pending cases, and filings by financial institutions (SIC 6000-6999), we have 629 bankruptcy filings, only 2 of which are Chapter 7 filings. From the LoPucki database, we collect each 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 sales of all or substantially all assets through Section 363 or the reorganization plan.[15]

From New Generation Research, Public Access to Court Electronic Records (PACER), National Archives at various locations, and U.S. Bankruptcy Courts for various districts, we are able to retrieve reorganization or liquidation plans and disclosure statements that are confirmed by the bankruptcy court for 520 cases. 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 (details in Appendix B). Second, they provide a detailed classification of claim/debt classes and the estimated amount owed or outstanding of each class of claims (including accrued interest, if any), 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 our final sample of 311 Chapter 11 filings.

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, 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

---

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

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.

The model's four possible outcomes are a pre-court reorganization, pre-court piecemeal liquidation, in-court reorganization, or in-court piecemeal liquidation. We classify sample cases into these outcomes by taking into account whether the filing is Chapter 7 or 11, whether the case converts to Chapter 7, whether the filing is prepackaged or pre-negotiated, whether the final outcome is a reorganization or sale of all or substantially all assets, and the case's duration. Details on our decision algorithm are in Appendix B, but we emphasize two features here. First, filings that are prepackaged or pre-negotiated are typically classified as a pre-court outcome. Second, we classify the sale of all or substantially all assets (either through a bankruptcy plan or Section 363) as a reorganization, not a piecemeal liquidation, because the business remains intact. Part of a reorganization involves finding the best management team for the business, regardless of whether that team is part of another firm or not. This classification also agrees with our model's assumption that a reorganization requires creditor skill; it is plausible that reaching a good M&A outcome requires skill.

The duration of each case is measured from the date of Chapter 11 filing to the plan confirmation date or the date of the Section 363 sale of all or substantially all assets.

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 interests" test for a 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 piecemeal liquidation of the firm's assets. This analysis is typically performed by independent financial advisors that are retained by the debtor firm. Since the majority of U.S. bankruptcy courts started to maintain electronic case dockets on PACER only after 2002, we search for independent liquida-

tion analyses for all sample cases filed from 2003 to 2014 in the LoPucki database. 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 372 firms in the LoPucki database. For our sample firms with missing liquidation analysis, 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_0, the firm's highest possible initial reorganization value. We estimate V_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. First, we estimate each sample firm's potential Tobin's Q as the median Q among all Compustat firms in the same industry and year. Second, we estimate V_0 by multiplying the potential Q by the firm's pre-filing book assets. Appendix B contains additional details.

Our coverage of bankrupt firms is much larger than previous studies that rely on detailed information about debt structure, debt recovery, and restructuring outcomes. Our sample, however, is potentially subject to selection biases, because we only study firms that file for bankruptcy, we apply several data filters, and we require non-missing values for our variables. Appendix B discusses these biases in detail and compares our sample to a broader universe of distressed firms. An important caveat is that our results apply only to firms in our sample and not necessarily to, for example, small firms, private firms, firms with a single creditor class, or firms that reach a private workout.

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. We estimate seven model parameters: $\theta_{S,0}$, the senior creditor's initial reorganization skill; $\theta_{J,0}$, the junior creditor's initial reorganization skill; 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_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, μ, defined as the number of calendar

months per period. Parameter μ captures how long it takes for a creditor to make a proposal, for lawyers and the court to review it, and for the opposing creditor to respond to it. This parameter summarizes a variety of institutional and informational frictions.

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_0, the initial maximum reorganization value; and L, the firm's liquidation value.[16] 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. We use nine 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 2 shows the local sensitivity of our nine simulated moments to our seven parameters.

The first moment is the average log number of months between observed proposals, for in-court cases. Table 2 shows that this moment is helpful for identifying μ, the number of months per period. In our model, one period consists of one proposal by a creditor. By measuring the

---

[16]Since the model normalizes total debt, D, to one, we scale D_J, V_0, and L by the value of D before taking these parameter values to the model. Doing so makes D_S redundant.

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 proposals, both in the actual and simulated data.

Moment two is the fraction of cases that result in a reorganization, conditional on the case being resolved in court. This moment is most informative about β, which controls the speed of learning. In Table 2, 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_t sooner, which tends to make reorganization more attractive than liquidation. Conversely, in the limit where learning is infinitely slow (i.e., very high β), reorganization values $\theta_{k,t}V_t$ will always be low, so reorganization will typically remain unattractive compared to liquidation. Table 2 shows that this moment also helps pin down ρ, because slower value decay (i.e., a high ρ) increases V_t, again making reorganization more attractive than liquidation.

Moment three is the average log duration of court cases, in months. Table 2 shows that, once μ is pinned down, this third moment mainly helps identify ρ. Specifically, longer court cases indicate a higher value of ρ. A high value of ρ means reorganization value decays slowly, so there is a low cost of waiting another period, hence court cases tend to last longer.

Arguably the toughest identification challenge is disentangling ρ and β, because both influence the costs and benefits of waiting. Table 2 confirms that moments $1-3$ move in different directions with ρ and β, as we require for identification. A higher ρ produces more months between observable plans, longer court cases, and *more* reorganizations in court. A higher β, however, has little effect on months between plans or case duration, and it produces *fewer* reorganizations.

The fourth moment is the fraction of cases that are resolved in court. As expected, once the previous parameters are pinned down, this moment is highly 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 fifth (sixth) moments is the senior (junior) creditor's average recovery rate among cases that result in a pre-court reorganization. These moments are highly informative about the creditor's initial reorganization skill, $\theta_{S,0}$ ($\theta_{J,0}$), as we see in Table 2. This result is also expected.

If a reorganization occurs in period 0 in the model, then the reorganization value is $\theta_{k,0}V_0$ for whichever creditor k leads the reorganization. If the senior has higher initial skill, 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 skill for the senior creditor (for example) leads to lower recovery rates for the junior. This result also makes sense. If the senior creditor has higher skill, 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 fraction of the pie.

The seventh moment is the junior creditor's average fraction gain, conditional on an in-court reorganization. We define the junior's fraction gain as the junior's dollar amount recovered divided by the total dollar amount recovered by both creditors. This moment is designed to be informative about λ_J, the junior's probability of proposing. Table 2 confirms that once the previous parameters are pinned down, this moment helps identify λ_J. The fraction gain measures how the junior and senior creditor split up the bankruptcy proceeds, so 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 fraction 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.

Two additional moments help identify several parameters. Moment eight 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 face value, D. This moment captures the total surplus among these cases and is informative about quite a few parameters. For example, this surplus increases in ρ and decreases in β, indicating that slower value decay and faster creditor learning both improve bankruptcy payouts. The surplus also increases in the senior and junior creditor's initial skill, because these parameters set the starting point of the learning curve. Lastly, this surplus increases in the fixed cost of going to court, because a higher fixed cost leads only high-surplus cases to select into going to court.

The last moment is the slope coefficient from a regression of the log total recovery rate on the duration of the case. We estimate this regression using cases that are resolved in court and result in a reorganization, because the moment then has a clear link to parameter ρ. From equation (2), the log total reorganization payout, gross of direct costs C_t, equals

$$\log[U_t(\theta_{k,t}) + C_t] = (t - 1) \log \rho + \log \theta_{k,t} + \log V_0.$$

We see that the slope coefficient of log total payouts on duration $(t-1)$ increases in ρ. Intuitively, if value decays more slowly (i.e. higher ρ), then the total recovery rate should have a higher, less-negative slope on duration. Table 2 confirms that this predicted slope is indeed positively related to ρ. The slope also has a mechanically large relation to months per period (μ), because the regression uses duration measured in calendar time.

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 3 shows how the model fits the nine moments targeted in the SMM estimation. The t-statistics test whether each data moment matches its model counterpart. Overall, the differences between the model-implied moments and data moments are economically small and statistically insignificant, indicating a good fit.

The average log months between plans is 1.71 in the model and 1.77 in the data. The average log duration (months) of in-court cases is 2.61 in the model, 2.57 in the data.[17] The model fits these features of the data quite well. The model also does a decent job matching the percent of cases resolved in court: 70.1% in the model, 73.3% in the data. For cases resolved in court, the fraction that results in a reorganization is 0.90 in the model, 0.88 in the actual data.

Next, we see that the model, by taking into account APR, is able to capture that senior

---

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

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 85.7% in the model, 87.8% in the data, and the junior creditor's recovery rate is 19.2% in the model and 22.1% in the data. 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 fraction of gain) and the pie's total size (i.e., total recovery rate). Aggregating the creditors, the total recovery rate is around 37%, both in the model and the data.

The slope coefficient of total recovery rate on case duration is significantly negative in the actual data. This descriptive, reduced-form result is new to the literature, and it foreshadows a main result from our counterfactual analysis: Delay destroys value and represents an important source of inefficiency. The model produces a negative coefficient similar to the empirical one. Figure 6 illustrates this relation by plotting the average total recovery rate for cases grouped into six-month duration bins. In both the model and the data, cases that take longer to resolve yield lower average payouts to creditors, and the simulated values closely match the data. For example, for cases that settle within 6 months, the average total recovery rate is about 40%, and the recovery rate drops sharply to 32% if the cases last longer than 2 years.

As an out-of-sample test, we check how well the model can match the full distribution of key variables. Figure 7 shows that 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 low. The junior's recovery distribution also matches reasonably well (Panel B). Both distributions are bi-modal, both in the model and data. Panel C shows that the model does a fairly 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), showing that most plans are proposed within a five-month interval since the last observed plan. Overall, Figure 7 suggests that the distributions and functional forms assumed in the model are reasonable.

3.2 Parameter Estimates

Table 4 contains parameter estimates from SMM. The estimate of μ implies that each model period is roughly 4.6 months long. This estimate summarizes how fast creditors can make and respond to a proposal, and how fast lawyers and the court can process a proposal. According to our model, one reason for bankruptcy cases' long duration is that each round of bargaining takes significant time.

The estimated initial reorganization skill of the senior and junior creditors ($\theta_{S,0}$ and $\theta_{J,0}$) are 0.28 and 0.36, respectively. These estimates imply that the senior creditor would initially produce a reorganization value that is 28% of the firm's maximum potential value, and the junior would produce a value that is 36%. It is plausible that junior creditors are more skilled on average, because junior 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 9.84. Figure 1 shows how to interpret this value. Panel A of the figure simulates creditor skill over time using the estimated values of β, $\theta_{S,0}$, and $\theta_{J,0}$. With these values, it takes 3 periods (roughly 14 months) for the median junior creditor's reorganization skill to increase from its initial value of 0.36 to 0.5. Even after 8 periods (roughly 36 months), creditors' skill levels are still bounded away from their maximum value of one. Learning does occur, in other words, but it is slow.

The estimate of ρ is 0.884. This value implies that the indirect costs of bankruptcy cause 11.6% of the firm's maximum reorganization value to decay each period (4.6 months) during court. Panel B of Figure 1 illustrates the implications. The solid black line shows how the maximum reorganization value, V_t, decays if $\rho = 0.884$. We see that 22% of the maximum reorganization value decays after 2 periods (roughly 9 months), and 53% decays after 8 periods (roughly 36 months). The remaining lines show that learning and value decay combine to make creditors' median reorganization value roughly constant at first, and then it decreases. There is randomness around these medians, however, and creditors hope to receive positive shocks to their skill and (hence) reorganization value.

The fixed cost of going to court, c_0, is estimated to be roughly 4.4% of the firm's total debt

value. Going to court entails a large cost, which the model requires to explain why less than 80% of cases are resolved in court. The high estimate of c_0 reflects not just the direct costs of going to court, but also the indirect costs triggered by a bankruptcy filing. It makes sense that filing for bankruptcy without a prepackaged or pre-negotiated deal hurts the firm's reputation and imposes real operating costs.

The junior's probability of proposing, λ_J, is estimated at 34.6%. This value implies that junior creditors have relatively low bargaining power, which the model needs to fit the low fraction of gain captured by the junior creditor. A value of λ_J below 50% is plausible, because debtor-in-possesion (DIP) financing is quite common during our sample period, and contractual provisions in DIP loans grant significant control to senior creditors during the bankruptcy process (Eckbo and Wang, 2019).

We can connect our parameter estimates to other papers that measure the costs of bankruptcy. Making the connection is a challenge, though. Many existing papers estimate the cost of financial distress, which is conceptually different from the cost of bankruptcy. Also, many papers estimate just one specific component of the costs,[18] whereas we estimate total costs. Arguably the closest comparison is to Davydenko et al. (2012), who estimate the average total cost of bankruptcy to be 30.5% of the market value of assets. Our estimates imply unconditional average direct costs of bankruptcy equal to at most 3.3% of the face value of debt.[19] It is harder to extract a single number for indirect costs in our model, in part because it requires additional assumptions on the value of the firm's assets absent bankruptcy. For a back-of-the-envelope calculation, note that reorganizations in our model occur on average in period $\bar{t} = 4.66$, so indirect costs destroy roughly $(V_0 - V_{\bar{t}})/V_0 = 1 - \hat{\rho}^{4.66-1} = 36\%$ of those firms' maximum potential value. This value is surprisingly close to the estimate of Davydenko et al. (2012).[20]

---

[18]Examples include Maksimovic and Phillips (1998), Hortacsu et al. (2013), Brown and Matsa (2016), and Graham et al. (2016).

[19]See Table 5 for this calculation, which requires the estimated values of c_0 and c_1, the frequency of in-court resolutions, and the duration of court cases.

[20]To map our model's value to the estimate of Davydenko et al. (2012), we make the implicit assumption that bankruptcy causes the firm's maximum potential value (V_t) to decay, but bankruptcy does not affect the path of reorganization skill ($\theta_{k,t}$). The logic is that a firm's owners will try to make the firm as valuable as possible, which corresponds to increasing $\theta_{k,t}$, regardless of whether the firm goes bankrupt.

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 reorganization skill at all times. Creditors still face uncertainty about future skill. The opposing creditor's true skill replaces its perceived skill as a state variable, and a creditor's own skill 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 skill but not future skill, as in the previous counterfactual model. Each period, the social planner chooses whether to wait, liquidate, reorganize under the senior's reorganization skill, or reorganize under the junior's reorganization skill. 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, indirect costs captured by $\rho < 1$, and slow creditor learning captured by $\beta \gg 1$.

The columns of Table 5 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 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. 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 5, we see that removing the asymmetric-information friction increases the average total recovery rate from 0.351 to 0.365, an 4% increase. In the social-planner benchmark, the average total recovery rate is 0.429, which is 22% higher than in the estimated model. Removing conflicts of interest therefore produces an additional 18% increase. These results imply that the observed bankruptcy process is quite inefficient. Asymmetric information between creditors generates a modest inefficiency, and conflicts of interest among creditors generate a significant inefficiency.

To convert the inefficiency into aggregate dollars, we note that the average year sees $146 billion in combined liabilities across Chapter 11 filings of firms that have at least $100 million in liabilities (Altman et al., 2019). Multiplying this total annual debt by the change in total recovery rate, $0.429 - 0.351$, yields $11.4 billion per year. In other words, we find that the two frictions we study destroy an average of $11.4 billion per year in the U.S.

The remaining rows of Table 5 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 fixed and 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 $V_t \theta_{k,t}$ averaged across firms that (endogenously) get reorganized at time t under either $k = S$ or J. Table 5 contains the terms in this decomposition.

We start with term (3), the average fixed and direct costs. These costs are the sum of the average fixed cost of going to court (from c_0) and the average variable direct costs of being

in court (from c_1).[21] The average fixed cost decreases from 0.029 to 0.028 when we remove asymmetric information, and to 0.023 when we additionally remove conflicts of interest. This change occurs because the fraction of cases resolved pre-court increases from 0.299 to 0.333 to 0.436 across these three models. The average direct cost of being in court starts at an already low value, 0.004, and decreases further to 0.003 and 0.001. This decrease occurs because fewer cases go to court, and the average duration of court cases decreases from 16.7 months to 13.4 (with symmetric information) and 4.5 (under the social planner). Adding these together, the average direct and fixed costs drop from 0.033 to 0.031 when we remove asymmetric information, and to 0.024 in the social-planner benchmark. These cost reductions explain a non-trivial 14.3% of the efficiency improvements in the symmetric-information benchmark, and 11.5% of the efficiency improvements in the social-planner benchmark.[22] 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 rate that stems from liquidations. We see that the fractions of firms liquidated versus reorganized are similar between the estimated model and the social-planner benchmark (0.209 v.s. 0.181). The inefficiencies, therefore, do not result simply from "too many" or "too few" firms being liquidated or reorganized. Furthermore, average liquidation values are only slightly lower in the estimated model than in the social-planner model (0.263 v.s. 0.272). Inefficiencies are not a result of low-value liquidations, in other words. Not surprisingly, the statistics for the symmetric-information benchmark fall between the estimated model and the social planner's model.

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 (0.791 v.s. 0.819), the average value of firms upon reorganization increases significantly, from 0.411 to 0.493. The resulting increase in term (2) explains 83% of the overall efficiency improvements in the social-planner benchmark.[23] In sum, we find that removing asymmetric information and

---

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

[22]Note that 14.3% = (0.033-0.031)/(0.365-0.351), and 11.5%=(0.033-0.024)/(0.429-0.351).

[23]Note 83% = 0.791(0.493-0.411)/(0.429-0.351).

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 excess liquidations, meaning firms that should have been reorganized, but the frictions led them to be liquidated instead. We then track excess continuations, meaning cases the social planner reclassifies from reorganization to liquidation. We find modest rates of excess liquidation and excess continuation. The social planner would reclassify 7.7% of cases from liquidation to reorganization, and 4.9% cases from reorganization to liquidation. The value gained from these reclassifications is also modest. The conditional average improvement in the recovery rate, moving from the baseline model to the social planner benchmark, is 0.062 for the excess liquidations and 0.006 for the excess continuations. Combining these extensive and intensive margins, we find that excess liquidation and continuation reduce the unconditional average total recovery rate by just 48 and 3 basis points, respectively. They together account for just 6.5% the total inefficiencies in the estimated model.[24] Therefore, our results suggest that excess liquidation and excess continuation are quantitatively small problems.

The second explanation is that firms are reorganized too late. Excess delay can destroy going-concern value due to loss of customers, employees, and the other indirect costs captured in our model by $\rho < 1$. Table 5 indicates that excess delay is a primary culprit for low reorganization values. The social planner would resolve an additional 14% of cases before going to court, and would reduce the average duration of remaining court cases by 12 months, or 73%. These results suggest that many firms would be better off reorganizing much earlier.

A third potential explanation is that the "wrong" creditor sometimes leads a reorganization. In the estimated model, a creditor can end up leading a reorganization even though the opposing

---

[24]Note that 48 basis points equals 0.077×0.062, and 3 basis points equals 0.049×0.006. The percent of the total inefficiency is (0.0048+0.0003)/(0.429-0.351)=6.5%.

creditor has higher reorganization skill. The opposing creditor may optimally allow such a deal if delay will destroy significant value, or if the opposing 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 5 shows the fraction of all cases in which a reorganization is led by the creditor with the lower reorganization skill. This fraction is zero in the social-planner benchmark, by construction. In the estimated model, 4.7% of cases result in a low-skill reorganization. Conditional on such an outcome, the average loss in total recovery rate is 20%. This loss equals V_t/D times the gap in skill ($\theta_{k,t}$) between the two creditors. The unconditional average loss from this type of inefficiency is therefore 0.94% of total debt value, which is 12% of the total inefficiency we find.[25] Low-skill reorganizations, in other words, are not a major source of inefficiency.

5 Robustness

In 2005 the U.S. Congress passed the Bankruptcy Abuse Prevention and Consumer Protection Act, which is viewed as making corporate bankruptcy more friendly to creditors in a variety of ways.[26] To check whether our results differ around the Act, we estimate our model independently in two time subperiods, 1996–2005 and 2006–2014. In our data moments, we find that the later subsample features fewer months between observed proposals, shorter court cases, and a slope coefficient of total recovery rate on court case duration that is more than double in magnitude (Table 6, Panel A). These differences in data moments produce differences in parameter estimates across subsamples (Panel B). We find that periods become shorter after 2005, perhaps because the new Act improved the speed of the court system. Creditors also learn faster post-2005 (i.e. β is lower), which also helps the model explain why cases become shorter. These forces tend to increase efficiency, but they are opposed by another force: the indirect costs of bankruptcy increase after 2005. To see this, note that ρ is slightly lower in the later subperiod, indicating larger indirect costs per model period, and model periods become shorter, so indirect costs per

---

[25]Note that 0.94%=0.047×0.200, and 12%=0.0094/(0.429-0.351).

[26]See Altman et al. (2019) for a review. Among other things, the Act changed the debtor's exclusivity period, composition of creditors' committees, treatment of vendors, treatment of commercial leases, use of key employee retention plans, and treatment of employee retirement benefits.

unit of calendar time become larger. The model needs these larger indirect costs to fit the recent subsample's steeper negative relation between total recovery and calendar-time duration (Panel A).[27] We find that these opposing forces roughly offset each other, producing a similar level of inefficiency in the two subperiods. Panel C shows that removing frictions increases the total recovery rate by 8.4% of debt value in the early subperiod and by 9.3% in the recent subperiod.

Next, we estimate in subsamples based on firm size. We expect larger bankruptcies to be more complex and hence slower. Panel A confirms that larger bankruptcies are more likely to be resolved in court and spend longer in court. Also consistent with greater complexity, we find that larger bankruptcies feature lower initial levels of skill and a slower speed of learning—it takes longer to find the best plan for a larger firm (Panel B). Waiting therefore has greater benefits for a larger firm, which helps the model explain why larger cases are slower. Longer cases, however, allow greater scope for excess delay and its associated costs. Panel C indicates that the social planner wishes to shorten large firms' court cases by 17 months, compared to 9 months for small firms. The social planner thereby increases reorganization values by more in large firms than small firms (0.083 versus 0.042). We therefore find a somewhat higher degree of inefficiency in larger firms (0.098 versus 0.063).

Next, we show that our results are robust to modeling asymmetric information more flexibly. Outside our baseline model, private communication among creditors could reveal their reorganization skill levels. To capture these effects, we estimate an extended version of our model in which creditors perfectly reveal their skill levels each period with probability p. Our baseline model assumes $p = 0$, which imposes a relatively high degree of information asymmetry. We estimate p using our original nine moments. Results are in Table 7. We estimate p to be 0.14, fairly close to its assumed value of zero in the baseline model. In the extended model, removing both frictions increases the average total recovery rate by 0.079, almost identical to the 0.078 increase from the baseline results.

---

[27]One potential explanation for the higher indirect costs after 2005 could be that firms increasingly hold intangible assets (e.g., Peters and Taylor, 2017), and such assets face higher indirect costs. Another explanation is that certain changes made by BAPCPA are seen as detrimental to bankrupt firms. For example, as explained by Gilson (2009), the lengthening of vendors' reclamation period imposes real costs on distressed retailers. Furthermore, restrictions on approving key employee retention plans likely result in flights of key employees from bankrupt firms (Goyal and Wang, 2017).

Industry conditions can vary over time, which leads to fluctuations in firms' potential reorganization value (V_t). These fluctuations can change the option value of waiting relative to our baseline model. For robustness, we extend the baseline model to make V_t stochastic. We replace Equation (1) with $V_t = \rho^{t-1}\hat{V}_t$, where $\hat{V}_t$ follows a two-state Markov process. Recall that our baseline estimation uses median industry Tobin's Q to measure firms' initial maximum reorganization value, V_0. To be consistent, we calibrate the Markov process for $\hat{V}_t$ to match the persistence and conditional volatility of median industry Tobin's Q. We then estimate our model's other parameters as before. Additional details are in the Online Appendix, and estimation results are in Table 7. Relative to our baseline results, the social planner resolves even more cases pre-court, but he shortens the remaining court cases' duration by less. Overall, the estimated inefficiency decreases somewhat, from 0.078 to 0.055, but it is still quite large.

Whereas creditors in the baseline model make proposals sequentially and randomly, the social planner can always choose either creditor to lead a reorganization. One potential concern is that this social planner is too powerful, which inflates the degree of inefficiency we find. Put differently, sequential and random proposing may be a feature of reality that a social planner or regulator cannot change. We therefore consider an alternative social planner who faces the same random-proposing scheme as in the baseline model. Specifically, the social planner is the same as before, except with probability λ_J each period the social planner can choose the junior creditor lead a reorganization, and with probability $1 - \lambda_J$ the social planner can choose the senior. We change the social-planner benchark but not the estimated model. As expected, the inefficiency is smaller when judged against this more-constrained social planner. Table 7 shows that the estimated inefficiency decreases somewhat, from 0.078 to 0.050, but it remains quite large.

6   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 22%, mainly by removing excess delay and

thereby 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 remain unclear.

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, and Organization* **8**(3), 523–46.

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.: 2019, 'Do the Right Firms Survive Bankruptcy?'. *Working Paper*.

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.

Ayotte, K. and E. R. Morrison: 2018, 'Valuation Disputes in Corporate Bankruptcy'. *University of Pennsylvania Law Review* **166**, 1819–1851.

Ayotte, K. and D. A. Skeel: 2013, 'Bankruptcy Law as a Liquidity Provider'. *The University of Chicago Law Review* pp. 1557–1624.

Ayotte, K. M., E. S. Hotchkiss, and K. S. Thorburn: 2013, 'Governance in Financial Distress and Bankruptcy'. In: K. K. Douglas Michael Wright, Donald S. Siegel and I. Filatotchev (eds.): *The Oxford Handbook of Corporate Governance*. Oxford University Press.

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.: 1988, 'A New Approach to Corporate Reorganization'. *Harvard Law Review* **101**, 775–804.

Bebchuk, L. A. and H. F. Chang: 1992, 'Bargaining and the Division of Value in Corporate Reorganization'. *Journal of Law, Economics, & Organization* **8**(2), 253–279.

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'.

Bradley, M. and M. Rosenzweig: 1992, 'The untenable case for Chapter 11'. *Yale Law Journal* **101**(5), 1043–1095.

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.

Brown, J. and D. A. Matsa: 2016, 'Boarding a Sinking Ship? An Investigation of Job Applications to Distressed Firms'. *The Journal of Finance* **71**(2), 507–550.

Bulow, J. I. and J. B. Shoven: 1978, 'The Bankruptcy Decision'. *The Bell Journal of Economics* **9**(2), 437–456.

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.

Chen, L., W. Dou, and Z. Qiao: 2013, 'Ensemble Subsampling for Imbalanced Multivariate Two-Sample Tests'. *Journal of the American Statistical Association* **108**(504), 1308–1323.

Cramton, P. C.: 1992, 'Strategic Delay in Bargaining with Two-Sided Uncertainty'. *The Review of Economic Studies* **59**(1), 205–225.

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.

Demiroglu, C. and C. James: 2015, 'Bank loans and troubled debt restructuring'. *Journal of Financial Economics* **118**(1), 192–210.

Dou, W. W., Y. Ji, D. Reibstein, and W. Wu: 2019, 'Inalienable Customer Capital, Corporate Liquidity, and Stock Returns'. *The Journal of Finance* **forthcoming**.

Eckbo, B. Espen, K. L. and W. Wang: 2019, 'Rent Extraction by Super-Priority Lenders'. *Working Paper.*

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.

Erickson, T. and T. M. Whited: 2002, 'Two-step GMM estimation of the errors-in-variables model using high-order moments'. *Econometric Theory* **18**, 776–799.

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.

Gilson, S., E. Hotchkiss, and R. Ruback: 2000, 'Valuation of Bankrupt Firms'. *Review of Financial Studies* **13**, 43–74.

Gilson, S. C.: 2009, 'Note on the Bankruptcy Abuse Prevention and Consumer Protection Act of 2005 (BAPCAP)'. *Harvard Business School Case.*

Gilson, S. C., K. John, and L. H. Lang: 1990, 'Troubled debt restructurings: An empirical study of private reorganization of firms in default'. *Journal of Financial Economics* **27**(2), 315–353.

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

Goyal, V. K. and W. Wang: 2017, 'Provision of management incentives in bankrupt firms'. *Journal of Law, Finance, and Accounting* **2**(1), 87–123.

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.*

Jensen, M. C. and W. H. Meckling: 1976, 'Theory of the firm: Managerial behavior, agency costs and ownership structure'. *Journal of Financial Economics* **3**(4), 305–360.

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

Kahl, M.: 2002, 'Economic Distress, Financial Distress, and Dynamic Liquidation'. *The Journal of Finance* **57**(1), 135–168.

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.

Merlo, A. and C. Wilson: 1995, 'A Stochastic Model of Sequential Bargaining with Complete Information'. *Econometrica* **63**(2), 371–399.

Merlo, A. and C. A. Wilson: 1998, 'Efficient delays in a stochastic model of bargaining'. *Economic Theory* **11**, 39–55.

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.

Nalebuff, B.: 1987, 'Credible Pretrial Negotiation'. *The RAND Journal of Economics* **18**(2), 198–210.

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

Peters, R. H. and L. A. Taylor: 2017, 'Intangible capital and the investment-q relation'. *Journal of Financial Economics* **123**, 251–272.

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

Pulvino, T. C.: 1998, 'Do asset fire sales exist? An empirical investigation of commercial aircraft transactions'. *The Journal of Finance* **53**(3), 939–978.

Rubinstein, A.: 1982, 'Perfect Equilibrium in a Bargaining Model'. *Econometrica* **50**(1), 97–109.

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

Stahl, I.: 1972, *Bargaining Theory*. Economic Research Institute, Stockholm School of Economics.

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

Figure 1: Dynamics of Skill and Reorganization Values

This figure shows simulated dynamics of creditor reorganization skill and reorganization value, using estimated parameter values from Table 4. To create the top panel, we initialize the creditors' abilities $\theta_{S,0}$ and $\theta_{J,0}$ at 0.28 and 0.36, respectively. We then randomly draw future values of reorganization skill from the generalized beta distribution, as in equation (3), using the value $\beta = 9.84$. 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_t, computed as in equation (1) with $\rho = 0.884$. This figure normalizes V_0 to 1. The lower lines show the product of V_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.

Panel A: Median Reorganization Skill

Panel B: Median Reorganization Value
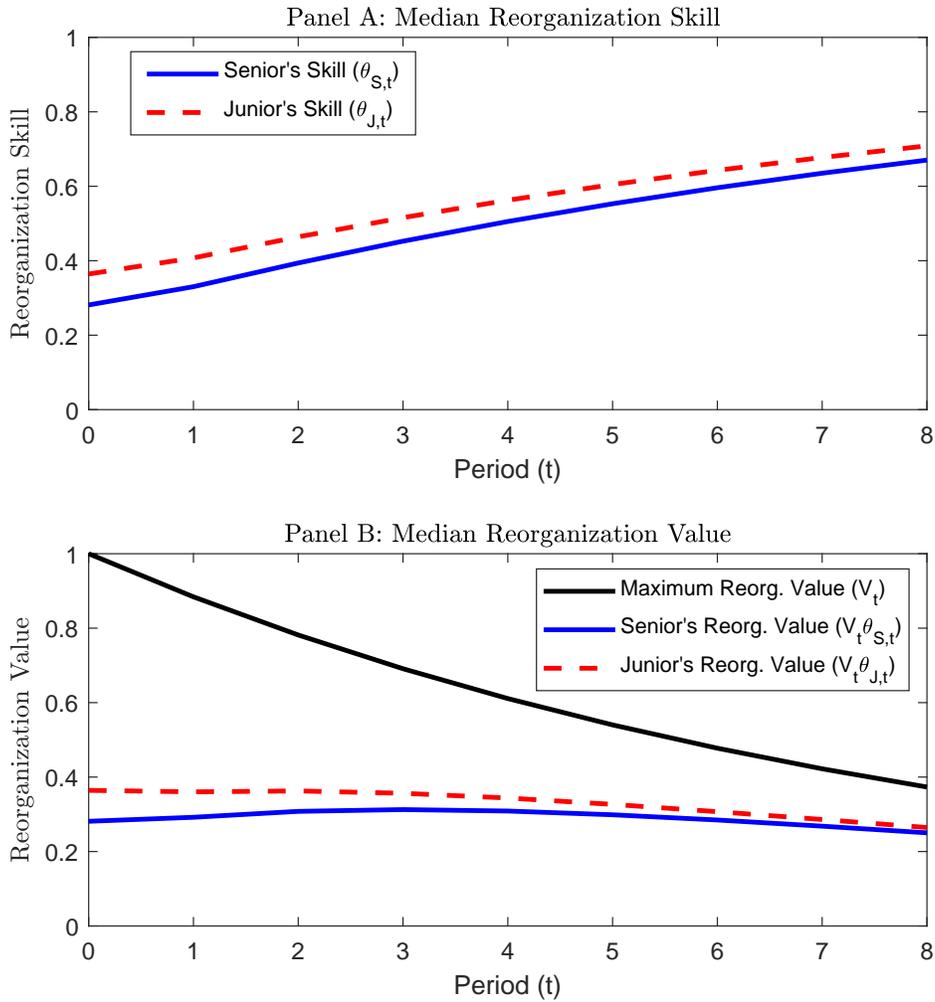
Figure 2: Timeline of Bargaining in Period t

Morning   **Afternoon**

Senior $\theta_{S,t}$ $\theta_{S,t}$ $\boxed{\theta_{S,t+1}}$ $\theta_{S,t+1}$

Public $\ell_{\bar{k},t},,\ell_{k,t}$ $\ell_{\bar{k},t},\boxed{\ell_{k,t+1}=\theta_{k,t}}$ $\ell_{\bar{k},t},\ell_{k,t+1}=\theta_{k,t}$ $\boxed{\ell_{\bar{k},t+1}=\ell(\xi_{k,t})},\ell_{k,t+1}=\theta_{k,t}$

Junior $\theta_{J,t}$ $\theta_{J,t}$ $\boxed{\theta_{J,t+1}}$ $\theta_{J,t+1}$

Actions   **Creditor k makes an offer $\xi_{k,t}$ to the counterparty**   **Creditor $\bar{k}$ responds to $\xi_{k,t}$ and period t ends**

Shocks   **Creditor k is chosen randomly to propose**   **Creditors randomly receive higher abilities**

42
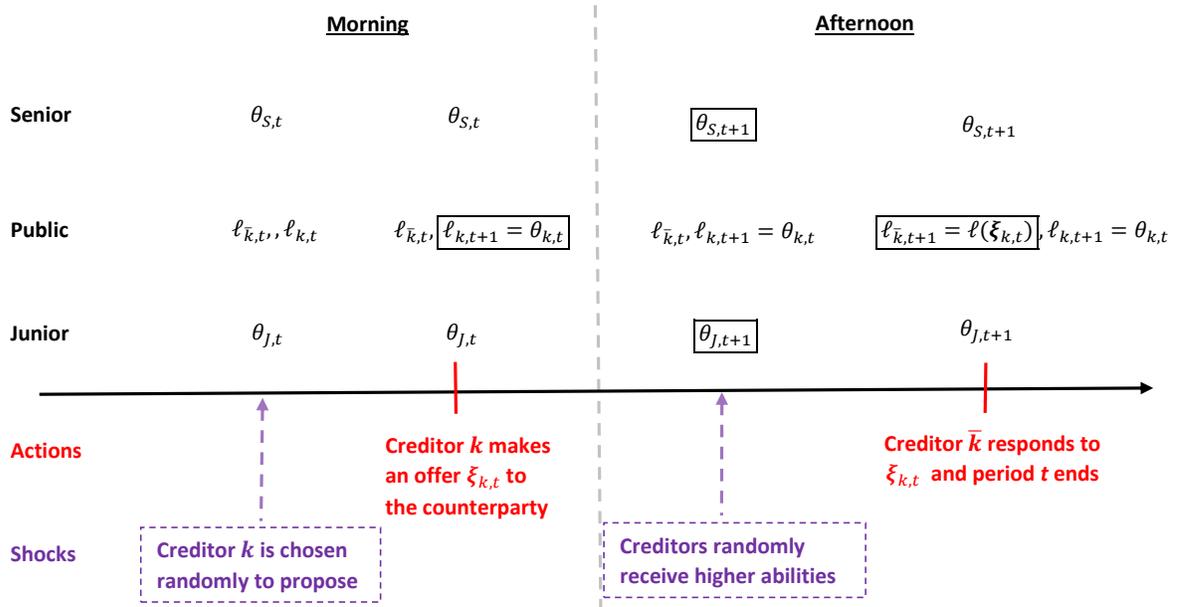
This figure shows creditors' optimal offers in our model. The horizontal axis denotes the proposer's true reorganization skill, and the vertical axis denotes the perceived responder's reorganization skill. The red areas represent the regions in which creditors make waiting offers, the black areas represent liquidation offers, and the light blue areas represent 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 ($t=2$). Parameters used for generating this figure are in Table 4.
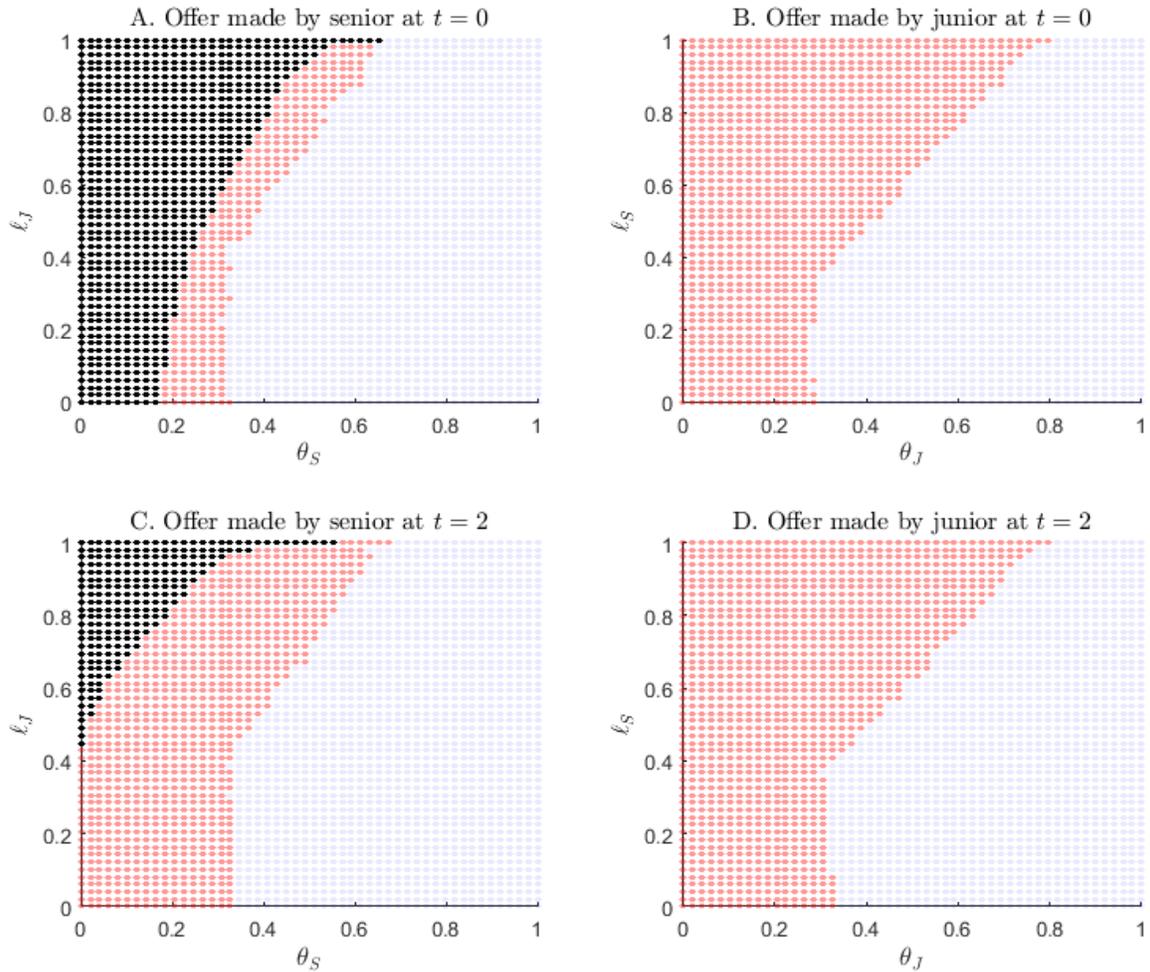
Figure 4: Optimal Financial Plan

This figure shows how creditors propose to split the firm's reorganization value in the pre-court period ($t=0$). The top two panels illustrate how the fraction of value offered by the senior to the junior change with the senior and junior's reorganization skill, and the bottom two panels illustrate how the fraction of value offered by the junior to the senior change with the junior and senior's reorganization skill. When we vary one creditor's reorganization skill, we fix the other creditor's reorganization skill at 0.5. Parameters used for generating this figure are in Table 4.
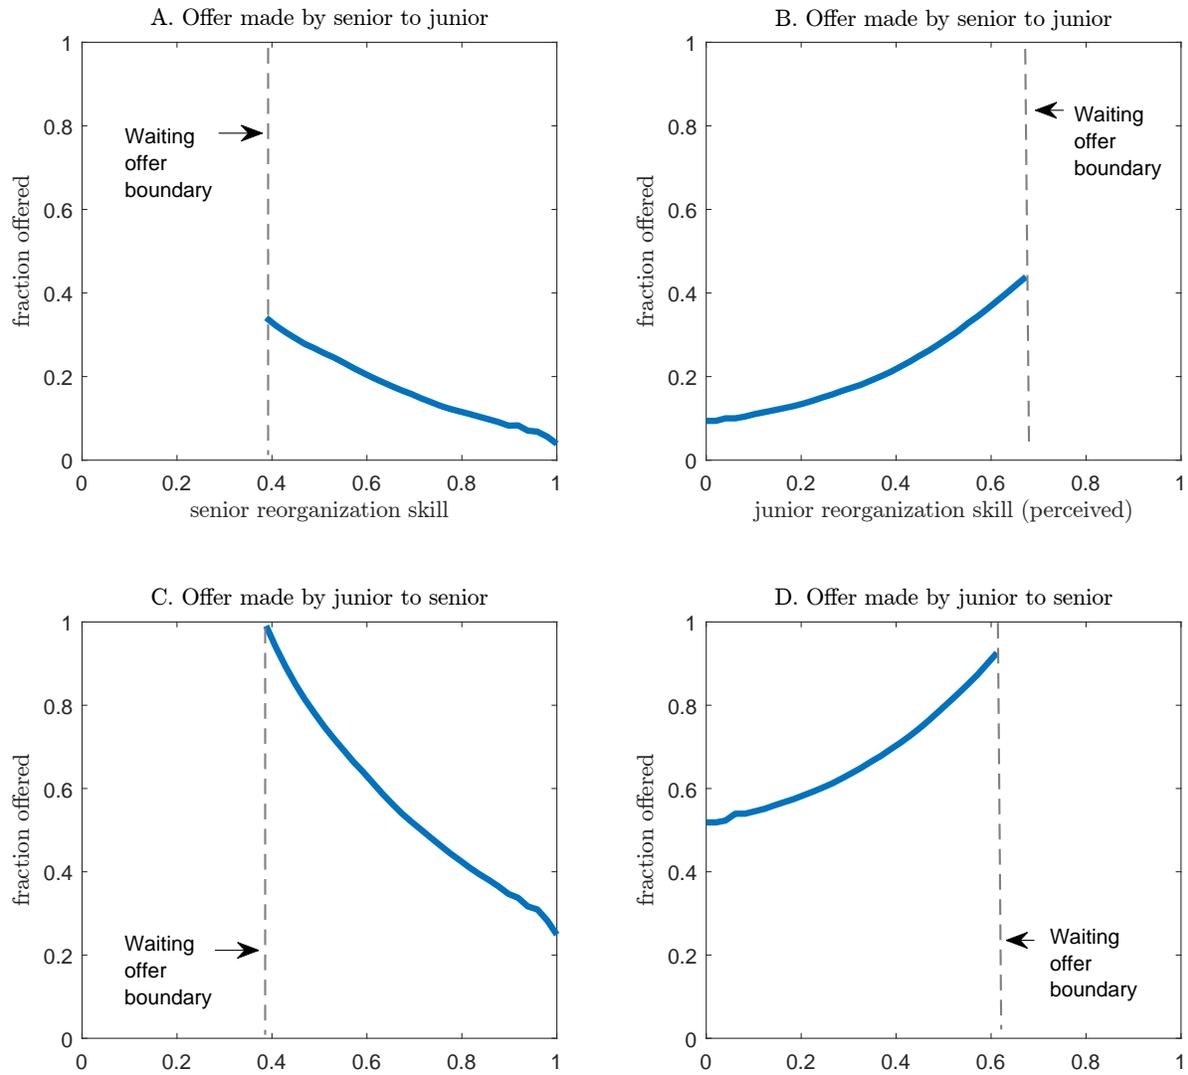
Figure 5: Examples of Simulated Bankruptcy Cases

This figure plots two simulations of the model. Each row corresponds to one simulation. The panels on the left column present the realized paths of the senior (blue "x") and junior (red "+") creditor's reorganization skill. The panels on the right column show the hypothetical total recovery rate if the case was settled at that point of time. The right panels contain three pieces of information: (1) who proposes (blue indicates senior proposes, red indicates junior proposes), (2) the offer type (circle means waiting offer, square means reorganization offer), and (3) the total recovery rate if the case was settled at time t by the proposer. The solid markers in the figure indicate case settlement.

45
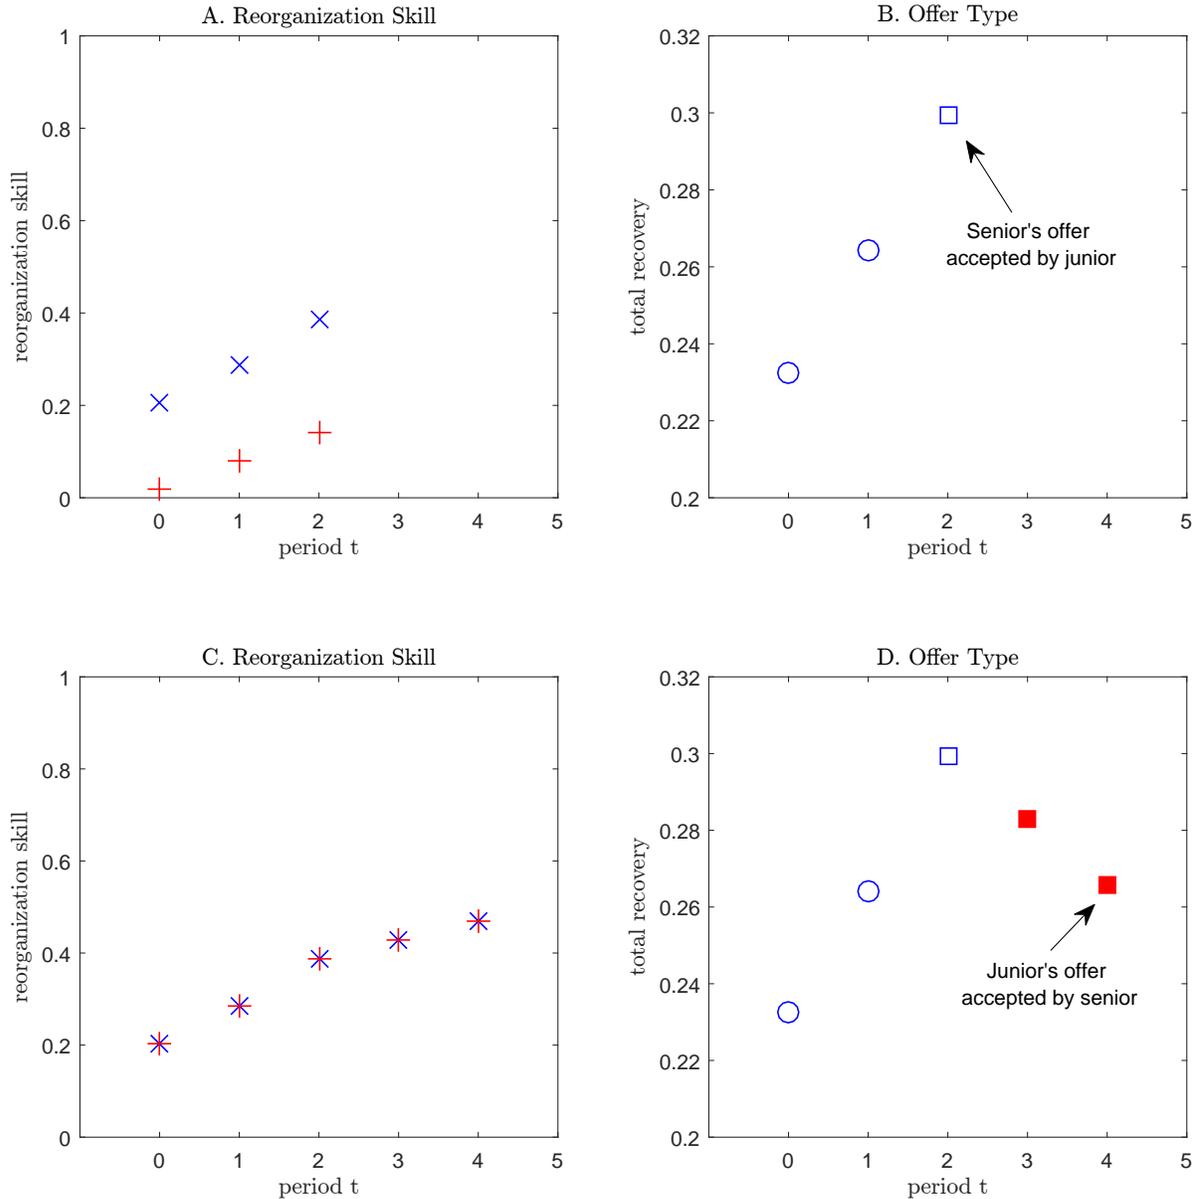
Figure 6: Recovery Rates versus Case Duration

This figure plots the average total recovery rate versus the bankruptcy case's duration, for cases resolved in court. 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.
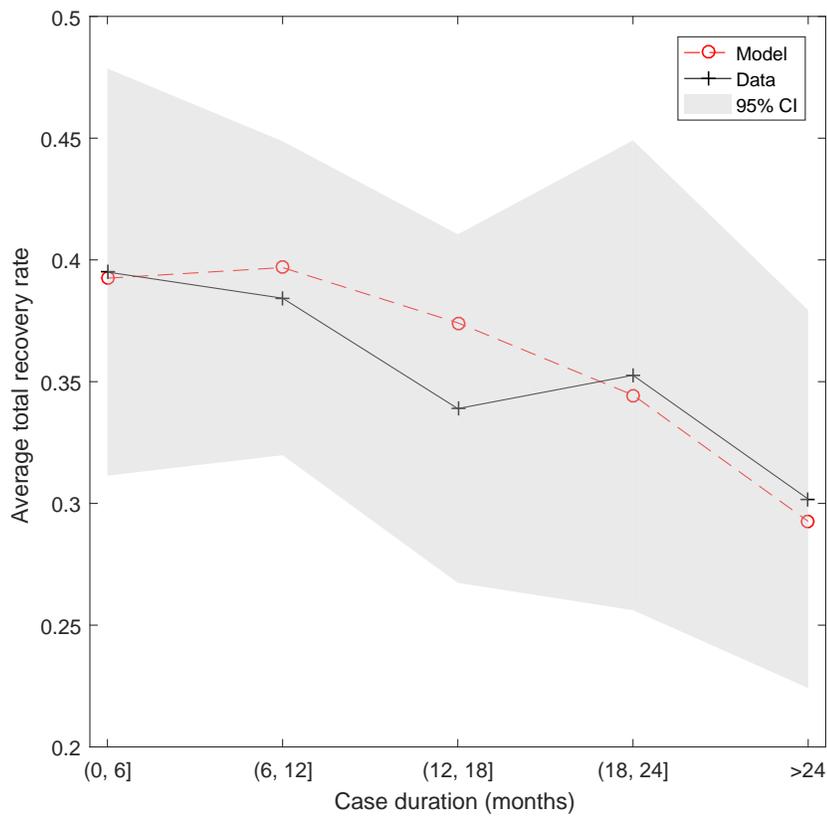
Figure 7: Comparing Simulated and Empirical Distributions

This figure plots the distributions of recovery rates (senior and junior), duration of court cases, and months between observed creditor proposals. Dark blue bars show results from simulation off the estimated model. Grey bars show the empirical distributions. These histograms pool all cases (e.g. pre-court and in-court).
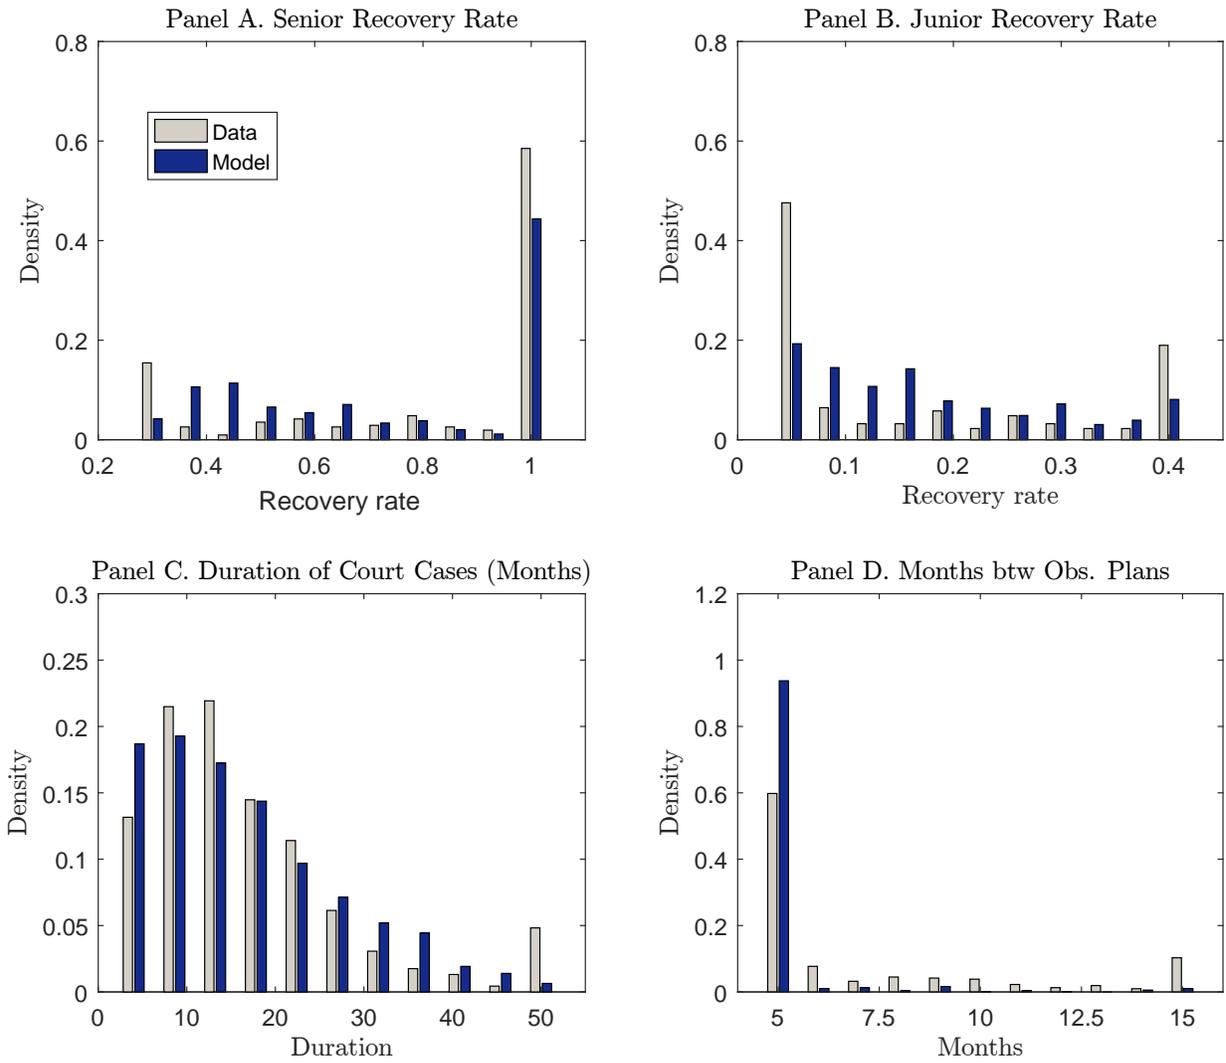
Table 1: Sample Overview

This table presents the summary statistics of our sample firms. Panel A reports the number of bankruptcies by year of filing and by Fama-French 12 industries. Panel B reports the mean and the median values of key characteristics of bankrupt firms. All financial variables are taken from the last fiscal year reported immediately prior to bankruptcy filing, retrieved from Compustat. Assets, Liabilities, and Sales are book assets, book liabilities and net revenue measured in millions of dollars, respectively. Employees are total headcount. ROA is EBITDA scaled by book assets. Leveraged is the ratio of book liabilities to book assets. Delaware and NYSD are indicators for cases filed in the District of Delaware and the New York Southern District, respectively. Reorg. Pre Court indicates pre-court reorganizations. Liq. Pre Court indicates pre-court liquidations; Reorg. In Court indicates in-court reorganizations. Liq. In Court indicates in-court liquidations. Months in bankruptcy measures the number of months from bankruptcy filing date to the plan confirmation date. Our sample consists of 311 large nonfinancial U.S firms that filed for Chapter 11 bankruptcies between 1996 and 2014.

Panel A: Distribution by Year of Filing and Fama-French 12 Industries					
Year	Number	% of Sample	Industry	Number	% of Sample
1996	8	2.6	Nondurables	35	11.3
1997	10	3.2	Durables	23	7.4
1998	13	4.2	Manufacturing	44	14.2
1999	24	7.7	Oil and Gas	13	4.2
2000	40	12.9	Chemicals	11	3.5
2001	47	15.1	Business Equipment	15	4.8
2002	35	11.3	Telecom	49	15.8
2003	24	7.7	Utility	5	1.6
2004	14	4.5	Wholesale and Retail	43	13.8
2005	8	2.6	Healthcare	14	4.5
2006	5	1.6	Other	59	19.0
2007	4	1.3	All	311	100
2008	12	3.9			
2009	33	10.6			
2010	8	2.6			
2011	7	2.3			
2012	5	1.6			
2013	10	3.2			
2014	4	1.3			
All	311	100			

Panel B: Firm and Bankruptcy Characteristics					
	Mean	Std. Dev.	25% Percentile	Median	75% Percentile
Assets ($m)	2,399	7,643	406	808	1,818
Liabilities ($m)	2,251	5,179	451	822	1,834
Sales ($m)	2,226	7,194	364	742	1,753
Employees	9,473	19,650	1,625	3,646	8,350
ROA	0.038	0.147	-0.003	0.062	0.104
Leverage	1.129	0.535	0.839	0.992	1.266
Delaware	0.486	0.501	0	0	1
NYSD	0.206	0.405	0	0	0
Reorg. Pre Court	0.270	0.445	0	0	1
Liq. Pre Court	0	0	0	0	0
Reorg. In Court	0.640	0.481	0	1	1
Liq. In Court	0.090	0.287	0	0	0
Months in bankruptcy	13.57	12.68	4.53	10.30	18.47

Table 2: 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 4), $Stderr(p)$ is the estimated standard error for parameter p (also from Table 4), and $Stderr(m)$ is the estimated standard error for the empirical moment m (from Table 3). Moments are defined in detail in Section 2.3. Parameter μ is the months per model period, β is the (inverse) speed of creditor learning, 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 skill levels of the senior and junior creditor, respectively, and λ_J is the probability that the junior proposes in a given period.

Panel A. Sensitivity of Moments to Parameters

					Moments				
Parameter	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
μ	**2.17**	0.00	2.21	0.00	0.00	0.00	0.00	0.00	0.41
β	-0.22	**-0.64**	-0.40	-1.44	-2.17	0.03	0.39	-0.33	-0.09
ρ	0.49	0.47	**1.30**	1.44	-2.94	-0.86	0.12	0.17	0.13
c_0	0.09	-0.43	-0.10	**-2.27**	-1.37	-0.51	-0.10	0.17	-0.05
$\theta_{S,0}$	-1.40	0.44	-0.31	1.06	**1.92**	-0.59	-0.95	0.30	0.04
$\theta_{J,0}$	-0.09	0.30	-0.24	-1.77	0.02	**0.40**	0.75	0.36	-0.03
λ_J	-0.51	0.49	-0.86	-1.24	1.28	0.29	**0.57**	-0.32	0.16

Panel B: Description of Moments

(1) Average log number of months between observed proposals for in-court cases.

(2) Fraction of cases that result in a reorganization, conditional on resolving in court.

(3) Average log duration of in-court cases, in months.

(4) Fraction of cases resolved in court.

(5) Senior creditor's average recovery rate in pre-court reorganizations.

(6) Junior creditor's average recovery rate in pre-court reorganizations.

(7) Junior creditor's average fraction of gain, conditional on an in-court reorganization.

(8) Total recovery rate averaged across all in-court reorganizations.

(9) Regression slope coefficient of log total recovery rate on case duration.

Table 3: 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. Error	t-stat.
Averages Across In-Court Cases:				
Ln Months Between Plans	1.711	1.769	0.060	-0.97
Fraction Reorganized	0.902	0.881	0.021	0.99
Ln Duration (Months)	2.608	2.571	0.058	0.64
Averages Across All Cases:				
Fraction Resolved In Court	0.701	0.731	0.025	-1.21
Average Recovery Rates for Pre-Court Reorganizations:				
Junior	0.192	0.221	0.027	-1.06
Senior	0.857	0.878	0.033	-0.63
Averages Across In-Court Reorganizations:				
Junior's Fraction of Gain	0.298	0.270	0.018	1.53
Slope of Ln Recovery on Duration	-0.017	-0.014	0.005	-0.59
Total Recovery Rate	0.375	0.370	0.019	0.25

Table 4: Parameter Estimates

This table contains parameter estimates from the SMM estimation. Parameter Mos./Period is the months per model period, β is the (inverse) speed of creditor learning, 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, respectively, and λ_J is the probability that the junior proposes in a given period.

Parameter	Notation	Estimate	Std. Error
Months per Period	μ	4.566	0.609
Senior's Initial Reorganization Skill	$\theta_{S,0}$	0.281	0.036
Junior's Initial Reorganiztion Skill	$\theta_{J,0}$	0.364	0.016
(Inverse) Speed of Creditor Learning	β	9.835	1.046
Persistence of Reorganization Value	ρ	0.884	0.006
Fixed Cost of Going to Court (%)	c_0	4.400	0.867
Junior's Probability of Proposing	λ_J	0.346	0.088

Table 5: 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 4. The first counterfactual model assumes symmetric information. The second counterfactual model assumes symmetric information and maximization of 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. The average fixed cost of going to court equals c_0 times the fraction of cases going to court, scaled by D. The average cost in court equals the fraction of cases going to court times the average number of periods in court times c_1, scaled by D. Average liquidation value is the average of L/D across all firms that are liquidated. Average reorganization value is the average of $\theta_{k,t}V_t/D$ across reorganized firms. Avg. gain due to eliminating excess liq. and reorg. is the average increase in total recovery rate from reassigning firms' outcomes between the estimated model relative to the two counterfactual models; it equals the sum of (1) fraction of all cases reassigned from liq. in the estimated model to reorg. in the counterfactual, times the average gain among those reassigned cases; and (2) the fraction of all cases reassigned from reorg. in the estimated model to liq. in the counterfactual, times the average gain among those reassigned cases. Avg. loss due to low-skill reorganization is the decrease in average total recovery rate coming from reorganizations being led by the creditor with lower skill ($\theta_{k,t}$); it equals the product of (1) the fraction of all cases in which there is a reorganization led by the creditor with lower skill; and (2) the average loss in recovery rate (relative to the reorganization being led by the high-skill creditor) among those low-skill reorganizations.

| | | Counterfactual Models | |
Simulated Statistic	Estimated Model	Symmetric Information	Social Planner
Avg. Total Recovery Rate	0.351	0.365	0.429
Avg. Fixed Cost of Going to Court (from c_0)	0.029	0.028	0.024
Avg. Direct Costs in Court (from c_1)	0.004	0.003	0.001
Avg. Fixed and Direct Costs	0.033	0.031	0.024
Fraction Liquidated	0.209	0.198	0.181
Avg. Liquidation Value	0.263	0.265	0.272
Fraction Reorganized	0.791	0.802	0.819
Avg. Reorganization Value	0.411	0.425	0.493
Fraction Resolved Pre-Court	0.299	0.333	0.436
Avg. Duration of Court Cases (Months)	16.7	13.4	4.5
Avg. Gain From Eliminating Excess Liq. and Reorg.	0.000	0.0048	0.0051
Frac. of Cases Switching from Liq. To Reorg.	0.000	0.012	0.077
Avg. Gain | Switching from Liq. To Reorg.	0.000	0.040	0.062
Frac. of Cases Switching from Reorg. To Liq.	0.000	0.001	0.049
Avg. Gain | Switching from Reorg. To Liq.	0.000	-0.014	0.006
Avg. Loss Due to Low-Skill Reorganizations	0.0094	0.0089	0.000
Frac. of Cases Low-Skill Reorgs.	0.047	0.062	0.000
Avg. Loss | Low-Skill Reorg.	0.200	0.144	0.000

Table 6: Comparing Results Across Year and Firm-Size Subsamples

This table presents results from independently estimating the model in four subsamples. The first two subsamples are formed based on the bankruptcy filing's year (1996-2005 versus 2006-2014). The latter subsamples are formed by comparing the firm's pre-filing book assets to the sample median. Panel A contains empirical values of selected moments targeted in SMM estimation. Moments are defined in detail in Section 2.3. Panel B reports SMM parameter estimates. Panel C reports selected model implications for variables defined in Table 5. Specifically, it reports the variables' average simulated values in the social-planner benchmark minus the estimated model.

	1996-2005	2006-2014	Large Firms	Small Firms
Panel A: Selected Data Moments				
Avg. Ln Months Between Plans	1.932	1.327	1.825	1.702
Avg. Ln Duration of Court Cases	2.754	2.076	2.712	2.403
Fraction Resolved in Court	0.743	0.701	0.789	0.673
Slope of Ln Recovery on Court Duration	-0.012	-0.030	-0.012	-0.011
Panel B: Parameter Estimates				
Months per Period (μ)	5.358	3.421	5.049	5.302
Senior's Initial Skill ($\theta_{S,0}$)	0.278	0.289	0.257	0.330
Junior's Initial Skill ($\theta_{J,0}$)	0.370	0.361	0.360	0.399
Inverse Speed of Creditor Learning (β)	9.95	5.46	10.19	9.67
Persistence of Reorganization Value (ρ)	0.881	0.856	0.876	0.85
Fixed Cost of Going to Court ($c_0, \%$)	4.65	4.24	4.32	3.00
Junior's Probability of Proposing (λ_J)	0.335	0.337	0.122	0.344
Panel C: Model Implications (Social Planner Model Minus Estimated Model)				
Avg. Total Recovery Rate	0.084	0.093	0.098	0.063
Avg. Reorganization Value	0.066	0.117	0.083	0.042
Fraction Resolved Pre-Court	0.107	0.132	0.319	0.424
Avg. Duration of Court Cases (Months)	-14.19	-6.19	-17.12	-9.10

Table 7: Robustness to Alternative Specifications

This table compares parameter estimates and counterfactual results between our baseline model and two alternative model specifications. Panel A reports SMM parameter estimates. Panel B reports selected model implications for variables defined in Table 5. Specifically, it reports the variables' average values in the social planner benchmark minus the estimated model. The first column shows results from our baseline estimation; these values match those in Tables 4 and 5. The specification with private communication adds a probability p that creditors' skill levels are perfectly revealed to each other each period. The specification with stochastic V_t allows the best-case reorganization value to fluctuate over time; specifically, it replaces Equation (1) with $V_t = \rho^{t-1}\hat{V}_t$, where $\hat{V}_t$ follows a two-state Markov process. In the last column, we compare the baseline estimated model to an alternative social planner that cannot choose the creditor with highest skill to lead a reorganization, but instead faces the same random-proposer scheme as in the baseline model.

	Baseline Model	Private Communication	Stochastic V_t	Alternative Social Planner
Panel A: Parameter Estimates				
Months per Period (μ)	4.566	5.524	4.388	4.566
Senior's Initial Skill ($\theta_{S,0}$)	0.281	0.323	0.244	0.281
Junior's Initial Skill ($\theta_{J,0}$)	0.364	0.349	0.355	0.364
Inverse Speed of Creditor Learning (β)	9.835	9.209	8.828	9.835
Persistence of Reorganization Value (ρ)	0.884	0.861	0.871	0.884
Fixed Cost of Going to Court (c_0, %)	4.40	4.62	3.70	4.40
Junior's Probability of Proposing (λ_J)	0.346	0.338	0.320	0.346
Probability of Revealing Skill (p)		0.141		
Panel B: Model Implications (Social Planner Model Minus Estimated Model)				
Avg. Total Recovery Rate	0.078	0.079	0.055	0.050
Avg. Reorganization Value	0.082	0.074	0.048	0.074
Fraction Resolved Pre-Court	0.137	0.271	0.192	0.267
Avg. Duration of Court Cases (Months)	-12.2	-10.2	-7.53	-11.2

A Technical Details on the Model

A.1 Perfect Bayesian Nash Equilibrium

The equilibrium in this game is entirely described by a pair of increasing sequences $\{\ell_{J,t}\}$ and $\{\ell_{S,t}\}$, which characterize the lower bounds for the perceived reorganization skills and optimally proposed payments $\{\xi_{J,t}\}$ and $\{\xi_S\}$.

We focus on equilibria that satisfy the skimming regularity (refinement) condition, an intuitive assumption widely adopted in the literature of dynamic bargaining with asymmetric information (e.g., Spier, 1992).

Assumption 1 (Skimming) *The creditors' strategies are such that if the responding creditor with reorganization skill θ' accepts the proposing creditor's restructuring proposal ξ with positive probability, then all responding creditors with reorganization skill $\theta'' < \theta'$ accept the proposal ξ 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 reorganization skill to restructure the firm is more likely to decline the counterparty's proposal and lead the restructuring by himself.

A.2 Solution

Stahl (1972) and Rubinstein (1982) were among the first to investigate formally the dynamic and strategic aspects of bargaining situations. We develop a modified version of Rubinstein (1982) bargaining model, incorporating separate growth options and asymmetric information. The screening feature of our model is similar to Bebchuk (1984). We solve the game recursively using the dynamic programming approach.

First, we describe the initial point of the dynamic programming procedure. The equilibrium is solved recursively by backward induction. The "end period" is the first time t such that $\rho^{t-1}V_0 \leq L$. In equilibrium, there is certain probability that the bargaining ends before the scenario $\rho^{t-1}V_0 \leq L$ occurs. In that period, the creditors choose to quit the bargaining by liquidating the firm. The APR applies to split whatever is left.

Next, we describe the Bellman equation for the senior creditor. Let us consider period t for any $t \geq 0$. The key is to establish the recursive Bellman equations for the continuation values at the beginning of the morning of period t, i.e., $W_{S,t}(\theta_{S,t}, \ell_{S,t}, \ell_{J,t})$ and $W_{J,t}(\theta_{J,t}, \ell_{S,t}, \ell_{J,t})$ with the endogenous state variables $(\ell_{S,t}, \ell_{J,t})$ and the exogenous (private) state variable $\theta_{J,t}$ or $\theta_{S,t}$. The private information about $\theta_{S,t}$ and $\theta_{J,t}$ is learned by the senior and junior, respectively, at the very beginning of the afternoon of period $t-1$.

The continuation value of the senior creditor at the beginning of period t (i.e., the very end of period $t-1$) follows the Bellman equation:

$$W_{S,t}(\theta_{S,t}, \ell_{S,t}, \ell_{J,t}) = \underbrace{(1-\lambda)\max\left\{O_{S,t}, \max_{\xi_{S,t}} \mathbb{E}_t^S\left[\widetilde{M}_{S,t+1}(\xi_{S,t})\right]\right\}}_{\text{if } S \text{ proposes in the "morning"}}$$

$$+ \underbrace{\lambda\mathbb{E}_t^S\left[\max_{\zeta_{S,t+1}\in\{0,1\}} \widetilde{A}_{S,t+1}(\zeta_{S,t+1}) \;\middle|\; \theta_{J,t} \geq \phi_{J,t}\right]\mathbb{P}_t^S\{\theta_{J,t} \geq \phi_{J,t}\}}_{\text{if } J \text{ proposes restructuring in the "morning"}} \qquad (4)$$

$$+ \underbrace{\lambda\mathbb{E}_t^S\left[\max\{O_{S,t}, U_{t+1}(\theta_{S,t+1}) - O_{J,t}\}\right]\mathbb{P}_t^S\{\theta_{J,t} < \phi_{J,t}\}}_{\text{if } J \text{ decides to liquid in the "morning"}}, \qquad (5)$$

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

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

$$\widetilde{M}_{S,t+1}(\xi_{S,t}) = \underbrace{[U_{t+1}(\theta_{S,t+1}) - \xi_{S,t}]\mathbf{1}\{W_{J,t+1}(\theta_{J,t+1}, \ell_{S,t+1}, \ell_{J,t+1}) \leq \xi_{S,t}\}}_{\text{if } J \text{ accepts the offer}} \qquad (6)$$

$$+ \underbrace{W_{S,t+1}(\theta_{S,t+1}, \ell_{S,t+1}, \ell_{J,t+1})\mathbf{1}\{W_{J,t+1}(\theta_{J,t+1}, \ell_{S,t+1}, \ell_{J,t+1}) > \xi_{S,t}\}}_{\text{if } J \text{ does not accept the offer}}, \qquad (7)$$

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

How do the endogenous state variables $\ell_{S,t}$ and $\ell_{J,t}$ evolve endogenously in this case? If the senior creditor receives the proposal opportunity in the morning of period t, $\ell_{S,t+1} = \theta_{S,t}$ and $\ell_{J,t+1} = \max(\theta_{J,t}^*, \ell_{J,t})$ with $\theta_{J,t}^*$ being determined by $\xi_{S,t} = W_{J,t+1}(\theta_{J,t}^*, \theta_{S,t}, \theta_{J,t}^*)$. The update of $\ell_{S,t+1}$ takes place right after the junior creditor sees the proposal $\xi_{S,t}$. The update is perfectly perceived by the senior creditor at the very beginning of period t right after he receives the proposing opportunity.

If the junior creditor proposes in the morning of period t, the payout to the senior creditor in the afternoon of period t, conditional on the junior's optimal choice $\xi_{J,t}$ that further depends

on $(\theta_{J,t}, \ell_{S,t}, \ell_{J,t})$, is described as follows:

$$\max_{\zeta_{S,t+1}\in\{0,1\}} \widetilde{A}_{S,t+1}(\zeta_{S,t}) = \underbrace{\xi_{J,t}\mathbf{1}\{W_{S,t+1}(\theta_{S,t+1}, \ell_{S,t+1}, \ell_{J,t+1}) \leq \xi_{J,t}\}}_{\text{if } S \text{ accepts the offer: } \zeta_{S,t+1} = 1} \tag{8}$$

$$+ \underbrace{W_{S,t+1}(\theta_{S,t+1}, \ell_{S,t+1}, \ell_{J,t+1})\mathbf{1}\{W_{S,t+1}(\theta_{S,t+1}, \ell_{S,t+1}, \ell_{J,t+1}) > \xi_{J,t}\}}_{\text{if } S \text{ does not accept the offer: } \zeta_{S,t+1} = 0}. \tag{9}$$

Finally, we describe the Bellman equation for the junior creditor. The continuation value of the junior creditor follows the Bellman equation:

$$W_{J,t}(\theta_{J,t}, \ell_{S,t}, \ell_{J,t}) = \underbrace{\lambda \max\left\{O_{J,t}, \ \max_{\xi_{J,t}}\mathbb{E}_t^J\left[\widetilde{M}_{J,t+1}(\xi_{J,t})\right]\right\}}_{\text{if } J \text{ proposes in the "morning"}}$$

$$+ \underbrace{(1-\lambda)\mathbb{E}_t^J\left[\max_{\zeta_{J,t+1}\in\{0,1\}}\widetilde{A}_{J,t+1}(\zeta_{J,t+1}) \ \Big| \ \theta_{S,t} \geq \phi_{S,t}\right]\mathbb{P}_t^J\{\theta_{S,t} \geq \phi_{S,t}\}}_{\text{if } S \text{ proposes restructuring in the "morning"}} \tag{10}$$

$$+ \underbrace{(1-\lambda)\mathbb{E}_t^J\left[\max\{O_{J,t}, U_{t+1}(\theta_{J,t+1}) - O_{S,t}\}\right]\mathbb{P}_t^J\{\theta_{S,t} < \phi_{S,t}\}}_{\text{if } S \text{ chooses to liquid in the "morning"}}, \tag{11}$$

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

If the junior creditor proposes in the morning of period t, the payout to the junior creditor in the afternoon of period t, conditional on the choice $\xi_{J,t}$, is described as follows:

$$\widetilde{M}_{J,t+1}(\xi_{J,t}) = \underbrace{[U_{t+1}(\theta_{J,t+1}) - \xi_{J,t}]\mathbf{1}\{W_{S,t+1}(\theta_{S,t+1}, \ell_{S,t+1}, \ell_{J,t+1}) \leq \xi_{J,t}\}}_{\text{if } J \text{ accepts the offer}} \tag{12}$$

$$+ \underbrace{W_{S,t+1}(\theta_{S,t+1}, \ell_{S,t+1}, \ell_{J,t+1})\mathbf{1}\{W_{S,t+1}(\theta_{S,t+1}, \ell_{S,t+1}, \ell_{J,t+1}) > \xi_{J,t}\}}_{\text{if } J \text{ does not accept the offer}}, \tag{13}$$

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

How do the endogenous state variables $\ell_{S,t}$ and $\ell_{J,t}$ evolve endogenously in this case? If the

junior creditor receives the proposal opportunity in the morning of period t, $\ell_{J,t+1} = \theta_{J,t}$ and $\ell_{S,t+1} = \max(\theta_{S,t}^*, \ell_{S,t})$ with $\theta_{S,t}^*$ being determined by $\xi_{J,t} = W_{S,t+1}(\theta_{S,t}^*, \theta_{S,t}^*, \theta_{J,t})$. The update of $\ell_{J,t+1}$ takes place right after the senior creditor sees the proposal $\xi_{J,t}$. The update is perfectly perceived by the junior creditor at the very beginning of period t right after she receives the proposing opportunity.

If the senior creditor proposes in the morning of period t, the payout to the junior creditor in the afternoon of period t, conditional on the senior's optimal choice $\xi_{S,t}$ that further depends on $(\theta_{S,t}, \ell_{S,t}, \ell_{J,t})$, is described as follows:

$$\max_{\zeta_{J,t+1}\in\{0,1\}} \widetilde{A}_{J,t+1}(\zeta_{J,t}) = \underbrace{\xi_{S,t}\mathbf{1}\{W_{J,t+1}(\theta_{J,t+1},\ell_{S,t+1},\ell_{J,t+1}) \le \xi_{S,t}\}}_{\text{if } J \text{ accepts the offer: } \zeta_{J,t+1} = 1} \tag{14}$$

$$+ \underbrace{W_{J,t+1}(\theta_{J,t+1},\ell_{S,t+1},\ell_{J,t+1})\mathbf{1}\{W_{J,t+1}(\theta_{J,t+1},\ell_{S,t+1},\ell_{J,t+1}) > \xi_{S,t}\}}_{\text{if } J \text{ does not accept the offer: } \zeta_{J,t+1} = 0}. \tag{15}$$

B Data

B.1 Sample selection

This subsection discusses the extent to which our sample is representative of the larger universe of distressed firms.

First, our sample does not include financially distressed firms whose creditors negotiate out of court and reach a deal without filing for bankruptcy (i.e. private workouts). We do not include these firms in our sample for two reasons. First, firms that successfully restructure out of court are likely to be those that are in less severe financial distress and face fewer obstacles impeding the negotiations (e.g., Gilson et al., 1990; Demiroglu and James, 2015), and are thus not comparable to bankrupt firms. Second, unfortunately, there are no standard definitions of out-of-court restructuring and databases that track the universe of these restructurings and debt recovery. Prior studies tend to capture only a fraction of this type of restructurings through exhaustive manual efforts. Nonetheless, since our sample includes prepackaged cases, our sample captures a sizable fraction of cases where creditors negotiate a plan out of court. In fact, firms that choose to file for bankruptcy often need the bankruptcy system to deal with specific issues such as tax and labor issues. Firms that achieve successful negotiations out of court but still file for bankruptcy are more comparable to those that file for bankruptcy and negotiate in court than those that avoid bankruptcy filing altogether. Overall, our sample could be biased towards firms that are more distressed and have a strong need to use the bankruptcy process to resolve distress.

Second, the requirement of non-missing information on debt classes, debt amount, and debt recovery for our estimation reduces the original sample of 626 bankruptcy filings (624 Chapter 11 and 2 Chapter 7) by US non-financial firms from 1996-2014 by 187, resulting in a sample of

439 cases. To illustrate the potential sample-selection bias, the Online Appendix compares key variables' means across the two subsamples. On average, the cases we exclude are smaller in firm size and have higher occurrence of liquidation relative to reorganization. This is largely due to the fact that the bankruptcy liquidation plans and disclosure statements tend not to provide precise estimated recovery of each debt class as reorganization plans and there are no plans filed for cases converted to Chapter 7 liquidation.

Moreover, comparing the 311 sample cases that have more than one class of debt with the 128 excluded cases with one class, we find that our sample cases have higher ROA, leverage, and tangibility, suggesting these firms have a stronger asset base and are relatively more likely to suffer financial distress than economic distress (e.g. Andrade and Kaplan, 1998). It is therefore not surprising that our sample cases have a higher likelihood of pre-court reorganization and lower likelihood of in-court liquidation. To conclude, compared to the universe of large-firm bankruptcies, our sample could be biased towards reorganization and firms that suffer financial distress rather than economic distress. This type of potential bias is often present in studies that require detailed information on debt structure and debt recovery in bankruptcy. For example, using Default & Recovery Database from Moody's to study the cost of default, Davydenko et al. (2012) cover a sample of 77 bankrupt firms, about 85% of which emerged from bankruptcy, a statistic that is comparable to the fraction of emerged firms in our sample.

B.2   Recovery rates

From the reorganization/liquidation plans and disclosure statements confirmed by the bankruptcy court, we obtain 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 calculate the actual recovery using the total value of cash, new debt, and new equity (using the closing price on the first trading day after emergence if available or estimate value given in the plan) given to each class of creditors. In a small subsample of debt with missing recovery, we obtain an estimate of recovery using the fitted value from an OLS regression of recovery on several observable characteristics: natural logarithm of debt amount, indicators on debt seniority (e.g. senior secured, senior unsecured, etc.), firm size, return on assets, leverage, asset tangibility, number of debt classes, an indicator for industry distress (measured by whether median sales growth within the firm's 2-digit SIC industry is less than -5%), and fixed effects for the Fama-French 12 industries.

B.3 Classification of bankruptcy outcomes

We classify sample cases into the four model outcomes (pre-court reorganization, pre-court liquidation, in-court reorganization, in-court liquidation) as follows. Pre-court reorganization occurs if the firm files a prepackaged or pre-negotiated plan at Chapter 11 filing, the firm successfully reorganizes or sells all or substantially 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. Pre-court liquidation occurs if the firm either (1) files Chapter 7, or (2) files a prepackaged or pre-negotiated Chapter 11 with an intent to liquidate, and it is liquidated in Chapter 11 or converted to Chapter 7. After applying this definition, we find no pre-court liquidations in our final sample. In-court reorganization occurs if either (1) the case is non-prepackaged or not pre-negotiated and the firm is reorganized, acquired, or sold as a whole through either Section 363 or a plan; or (2) the case is prepackaged or pre-negotiated, the firm is 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 or not pre-negotiated and the firm is liquidated piecemeal or converted to Chapter 7; or (2) the firm files a prepackaged or pre-negotiated plan with an intent to reorganize at Chapter 11 filing, yet the firm is liquidated piecemeal in Chapter 11 or converted to Chapter 7. Our classification is not affected by firms' selling only a small fraction of assets through Section 363.

B.4 Missing liquidation values (L)

This subsection describes how we estimate the missing values of L/D, the firm's liquidation value scaled by face value of debt. For firms with missing L, we first obtain an estimate of L scaled by book assets. This estimate is the fitted value from an OLS regression of $L/Assets$ on several observable characteristics: firm size, return on assets, leverage, asset tangibility, a dummy for whether there is fraud, the industry market-to-book ratio, fixed effects for the Fama-French 12 industries, and fixed effects for the company's state of incorporation. We estimate this regression using data from 204 bankrupt firms that filed for Chapter 11 from 2003-2014 in the UCLA LoPucki Database and have non-missing L and the other regression variables. When possible, we measure book assets value at the time of the liquidation analysis, otherwise we measure book assets at the last fiscal year before the bankruptcy. In the regression, leverage and asset tangibility enter positively at the 5% confidence level. The remaining regressors, excluding fixed effects, do not enter significantly at the 5% level. For 28 firms that indicate liquidation analysis is "Not Necessary" we assume their liquidation value must be low, and we set $L/Assets$ to the 10th percentile across firms with directly observable $L/Assets$. We multiply the estimated value of $L/Assets$ by $Assets/D$ to obtain an estimate of L/D.

B.5 Maximum reorganization value (V_0)

We estimate each firm's potential Tobin's Q to be the 50th percentile Q in the same SIC3 industry and year. 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 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 estimated speed of learning and creditors' initial skill levels would adjust to continue fitting the data, leaving the paths of reorganization value, $\theta_{k,t}V_t$, largely unchanged.

C Details on SMM estimation

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

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

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 Erickson and Whited (2002). 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_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 $\{D_J, V_0, L\}$, 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 $\{D_J, V_0, L\}$. 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., Chen et al., 2013, and 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 $\{D_J, V_0, L\}$ for each clusters. When simulating data off the model, we solve the model for each of these ten median values of $\{D_J, V_0, L\}$, 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.