Judicial Enforcement and the Transmission of Bankruptcy through Firm Networks

Geraldo Cerqueiro[2] Gil Nogueira[1,2]

February 1, 2025

Abstract

We show that judicial enforcement affects the transmission of bankruptcy through firm networks. To this end, we use the 2014 reform of the Portuguese judicial map as a source of exogenous variation in court congestion. We combine this reform with a detailed dataset covering all bankrupt firms and their suppliers with defaulted trade credit. Our findings indicate that reducing court congestion to the minimum observed level lowers the probability of bankruptcy among connected suppliers from 6% to 3%. We show that this difference is driven by the amplification of credit losses in "busy courts", which disproportionately affect creditors facing greater financing frictions.

1 Banco de Portugal. 2 Católica-Lisbon CLSBE. We thank Statistics Portugal (INE) for providing access to the data. Any views expressed are those of the authors and do not necessarily represent the views of Banco de Portugal, the European Central Bank, or the Eurosystem.

1 Introduction

Firms borrow from their suppliers and lend to their customers, making trade credit a valuable source of financing for non-financial firms. Recent estimates indicate that the aggregate volume of trade credit is around 20 percent of world GDP, a figure comparable to the volume of outstanding corporate bonds (Boissay et al. (2020)).

However, reliance on trade credit makes firms vulnerable to credit risk, since their clients can be hit by liquidity shocks and default on outstanding debt. These shocks can ripple through supply chains, triggering a cascade of bankruptcies (Kiyotaki and Moore (1997), Battiston et al. (2007), Jacobson and Von Schedvin (2015)) and potentially having large macroeconomic impacts (Acemoglu et al. (2012)).

This paper shows that efficient judicial enforcement mitigates the propagation of bankruptcies along supply chains. We measure the quality of enforcement using plausibly exogenous variation in court congestion generated by a reform of the Portuguese judicial map.[1] We show that lowering court congestion to its minimum observed level halves the likelihood that a supplier with defaulted trade credit from a bankrupt debtor also becomes bankrupt.

The Portuguese judicial map underwent a major redesign in 2014 as part of the country's bailout agreement. The reform aggregated existing courts into fewer, larger judicial districts and introduced specialized commerce courts to handle bankruptcy cases. As part of this process, judges were reassigned to the newly created courts. However, their allocation was based on projected case inflows and did not account for existing backlogs, a fact we confirm empirically. This oversight, combined with the transfer of all pending cases from smaller, dispersed courts to larger ones, led to substantial heterogeneity in backlog volumes.

We combine the judicial reform with a dataset of all creditors of firms that filed for bankruptcy between January and August 2014. All these cases were pending as of September

[1] Several studies, including Ponticelli and Alencar (2016), Iverson (2018), and Müller (2022), use case backlogs as a measure of the quality of judicial enforcement.

1, when the new court map took effect, and were transferred on this date to the newly designated courts. Importantly, insolvency cases must be filed in the judicial district where the debtor is headquartered or conducts most of its economic activity. Using fiscal identifiers, we match debtors and creditors to a firm-level dataset containing financial information for all firms in Portugal, allowing us to track these firms from 2007 to 2020.

We perform our empirical analysis in four steps.

First, we investigate whether the level of court congestion is correlated with the economic characteristics of judicial districts. Court congestion is defined as the number of backlog cases per judge in the court where the bankruptcy case was transferred on September 1, 2014. We find that court congestion is unrelated to local labor market conditions, the importance of the local banking sector, the local economic structure, or the judicial characteristics of the local court in the old map. Thus, differences in court congestion appear to be plausibly exogenous to local economic conditions.

Second, we analyze how court congestion affects the transmission of corporate bankruptcy through firm networks. We find that a supplier with defaulted trade credit is more likely to go out of business if the debtor's bankruptcy procedure is adjudicated in a more congested court. This result holds in both a simple regression controlling only for year dummies and a more saturated model that includes region-by-year, industry-by-year, and creditor-by-debtor fixed effects, along with additional creditor characteristics interacted with year dummies. The stability of our point estimate confirms that our measure of congestion is unrelated to local and industry-specific economic shocks. We also show that bankruptcy propagation only occurs when the defaulting debtor is large relative to the supplier.

We perform additional tests to demonstrate the robustness of this finding. Analyzing year-by-year coefficients, we find no evidence of pre-trends, since courts exhibited similar rates of bankruptcy propagation prior to the reform. Following Ponticelli and Alencar (2016), we also implement an alternative empirical strategy that compares pairs of adjacent municipalities belonging to different judicial districts. These municipalities differ in court congestion levels but are otherwise comparable in observable characteristics, thereby accounting for potentially confounding economic shocks. Our main finding holds.

Third, we examine the mechanisms driving our finding that busier courts amplify the transmission of corporate bankruptcies through firm networks. Following Jacobson and von Schedvin (2015), we focus on two potential channels: lost sales from the debtor's exit as a client (operational) and credit losses (financial).

If the operational channel were at play, we would expect that if a debtor is assigned to a busier court, either affected suppliers lose this client more quickly or the loss of this client has a greater impact. However, several tests suggest this channel does not explain our results. First, we find that debtors continue their operations for longer when their bankruptcy cases are assigned to busier courts. Second, we show that the degree of bankruptcy propagation is unaffected by the availability of alternative clients or the specificity of the trade credit relationship. These findings are inconsistent with suppliers experiencing faster declines in demand in more congested courts.

Instead, our evidence points to the financial channel: court congestion amplifies credit losses by delaying the resolution of bankruptcy cases, causing creditors to be paid later and less. These delays exacerbate suppliers' financial distress, increasing their likelihood of insolvency. Using data on impaired receivables, we confirm that court congestion increases credit losses for suppliers. These losses should be particularly damaging for financially constrained firms with limited capacity to absorb liquidity shocks. Consistent with this view, we show that the bankruptcy propagation mechanism is entirely driven by firms in industries characterized by high financial frictions, measured by external finance and liquidity dependence.

Fourth, we conduct a counterfactual analysis to quantify the effect of court efficiency on aggregate bankruptcy propagation rates. In our sample, the observed average bankruptcy propagation rate is 5.9%. We simulate two scenarios: one where all courts operate with the efficiency of the least congested court, reducing the propagation rate to 3.2%, and another where all courts operate at the efficiency level of the most congested court, increasing the propagation rate to 7.4%. Thus, moving from the most to the least efficient courts more than doubles bankruptcy propagation to trade creditors, underscoring the critical role of judicial efficiency in amplifying financial contagion within trade credit chains.

A growing body of literature explores how shocks propagate throughout the economy, often through different types of networks. One set of studies examines intra-firm networks. For instance, Giroud and Mueller (2019) show that local shocks can spread across U.S. regions through firms' internal networks of establishments. Similarly, Gilje et al. (2016) find that local bank lending is influenced by shocks affecting the bank in distant regions.

Another set of studies focuses on inter-firm networks, where shocks can propagate along the supply chain. The seminal work by Acemoglu et al. (2012) demonstrates how sectoral shocks can have aggregate effects due to intersectoral supplier-customer linkages. Recent evidence further underscores the economic importance of this propagation mechanism. Barrot and Sauvagnat (2016), using supplier-customer links reported by U.S. publicly listed firms, find that customers of suppliers affected by natural disasters experience a 20 percent drop in sales following the event. Similarly, Jacobson and Von Schedvin (2015) highlight the significance of trade credit chains in the propagation of corporate failures, showing that a trade debtor's failure is associated with a 53 percent increase in the supplier's risk. Fazio et al. (2021) provide evidence that firms located in places with congested courts are more adversely affected by shocks to firms in their network because of difficulties in contracting with new firms.

Our study contributes to this literature by showing that the propagation of corporate bankruptcy through supply chains decreases with the quality of judicial enforcement in bankruptcy resolution. This finding is especially relevant, given that business bankruptcy institutions perform poorly in many countries, including developed economies (Djankov et al. (2008)). A high rate of bankruptcy propagation implies that idiosyncratic shocks affecting specific industries are more likely to spread, potentially causing bankruptcy cascades and aggregate economic effects. Thus, our micro-evidence supports a growing body of research investigating the role of bankruptcy resolution for the business cycle (Jordà et al. (2022), Kornejew et al. (2024)). These studies argue that environments with poorly functioning bankruptcy systems face more severe economic contractions after a boom. Our findings provide a simple explanation for this pattern: inefficient bankruptcy resolution increases bankruptcy costs to firms in the network of bankrupt firms.

The paper proceeds as follows. Section 2 provides information about the bankruptcy regime in Portugal and explains in more detail the reform of the judicial map that took place in 2014, including how judges were allocated to the new courts. Section 3 discusses the data and variables, while section 4 presents the empirical methodology. Section 5 presents the main results and discusses the possible mechanisms. Section 6 assesses the aggregate relevance of our findings, and Section 7 concludes.

2 Institutional Background

2.1 Corporate Bankruptcy in Portugal

According to Portuguese Law, a firm is considered insolvent if at least one of the criteria is met: (1) the firm cannot fulfill their overdue obligations, or (2) its liabilities are greatly above its assets (the extent is not quantified). The insolvency procedure starts with a petition filed by the debtor or by one of its creditors. If the petition is filed by a creditor, the debtor can submit a defense. Insolvency cases must be filed in the judicial district where the debtor is headquartered or has the most of its economic activity. Therefore, creditors and debtors have limited discretion in choosing the court where they file their petitions.

The insolvency petition is assigned to a judge. If case the request is accepted, all collection efforts by creditors must stop and the bankruptcy process starts. The judge appoints a trustee to oversee the bankruptcy process. The trustee manages the debtor's assets, receives claims from creditors, and organizes a meeting with creditors where they can propose a plan for debt repayment or asset liquidation. This meeting takes place 45 to 60 days after the bankruptcy declaration. Any proposed repayment plan must be first approved by a two-thirds majority of creditors and then accepted by the judge. If creditors do not reach an agreement, the debtor's assets are liquidated and used to repay creditors. The trustee is responsible for the execution of the repayment plan or for the sale of the assets and subsequent repayment to creditors. Creditors often file additional complaints, arguing for example against debts in their favor that were not recognized by the court, or against potential acts of the debtor that devalued firm assets. Once the repayment plan is completed or the assets are fully liquidated, the judge must close the bankruptcy process.

2.2 Reform of the Judicial Map

The Portuguese judicial map underwent a major redesign in September 2014 under the Law for the Judicial System Organization.[2] This reform was part of the broader structural changes mandated by the financial assistance program that Portugal signed in 2011 with the European Union, European Central Bank, and International Monetary Fund. It targeted the inefficiencies of the judicial system, identified as a significant barrier to the country's economic growth. The reform had two main objectives.

The first objective was to merge the smaller courts into larger, more efficient units. Portugal is administratively divided into 20 administrative units, which are further subdivided into 308 municipalities. Before 2014, Portugal had 231 local jurisdictions, each serving either a single municipality or a small group of municipalities. After the reform, these were reduced to 23 larger judicial districts. Major urban centers like Lisbon and Porto, which concentrate a large fraction of economic activity, were further subdivided, with Lisbon divided into three

[2] The general guidelines were approved in August 2013 (Law 62/2013), but full implementation, including the new court structure, only occurred with Law 49/2014 in March 2014. The regulation governing judge allocation (Regulation 161/2014) was issued in August, during judicial holidays, with the new map becoming effective on September 1, 2014. In 2014, judicial holidays ran from July 16 to August 31.

districts and Porto into two. Each of the remaining eighteen districts has its main court in the district capital.

The second objective was to increase specialization. Each judicial district in Portugal is organized into a central court and local court departments. The central court departments serve the entire judicial district and may include sections dedicated to civil, criminal, and specialized matters, such as family, labor, or commercial law. "Commercial courts" handle business law matters, including bankruptcy cases. When available, bankruptcy cases are assigned to a commercial court. Judicial districts may have more than one commercial court, each with its own non-overlapping jurisdiction within the district. In districts without a commercial court, bankruptcy cases are typically assigned to the civil section of a local court.

The 2014 reform expanded the number of commercial courts from 4 to 20. We exclude cases processed in preexisting commercial courts before the reform to avoid potential differences in procedural history. By concentrating on cases transferred from generic courts to commercial courts, we ensure a more homogeneous sample, excluding both the most and least economically developed regions. Specifically, this approach excludes Lisbon and Porto, which already had commercial courts prior to 2014, as well as less economically developed interior regions with fewer insolvency cases and no specialized courts. In robustness tests we include all courts handling bankruptcy cases. Figure 1 shows the location of the newly created commercial courts.

2.3 Judge Assignment

The 2014 judicial reform required reassigning all judges from the old courts to the newly created ones. This involved two decisions: determining the number of judges needed for each court and selecting the assigned judges.

To determine the number of judges, policymakers employed a simple rule.[3] First, they calculated the average number of cases filed annually over the past three years in the old jurisdictions that were being merged into the new, larger jurisdictions. This average served as an estimate of how many cases the new court would receive each year. To determine how many judges were needed, this annual case inflow was divided by the maximum number of cases that each judge should handle in a year. For commercial courts, this limit was set at 200 cases per judge.

For example, if a newly created commercial court was expected to handle 500 cases annually, this number was divided by 200, resulting in 2.5 judges. Since fractional judges are not possible, they rounded this up to three judges for that court. This rule was applied in most cases. However, in a few jurisdictions, they added one or two additional judges. Removing these extra judges does not alter our results.

To determine individual judge allocations, policymakers invited all judges to submit until transfer requests by 31 May 2014, specifying their preferred placements. The allocation process prioritized seniority and specialization to ensure a balanced distribution of judges. Judges who did not submit requests were automatically assigned to fill the remaining vacancies.[4]

2.4 Case Migration and Court Congestion

The method of assigning judges based on expected case inflows led to significant disparities in congestion levels across the new jurisdictions (see Figure 1). When the reform took effect on September 1, 2014, all pending bankruptcy cases were transferred from the old courts to the newly created, larger courts. These new jurisdictions combined the backlogs of several smaller courts, each of which had varying levels of pending cases. However, the reform did not consider the differences in these pre-existing backlogs when assigning judges. As a result,

[3] *Linhas Estratégicas para a Reforma da Organização Judiciária* (in Portuguese).
[4] These rules are detailed in *Deliberação n.º 1084/2014* (in Portuguese).

some new courts inherited a disproportionately large number of pending cases compared to others, creating uneven levels of congestion across judicial districts.

This discrepancy was noticed by lawyers, prosecutors, and judges, all of whom attributed it to poor planning.[5] Others argue that the reform had mainly a cost-saving function, and the location for the new courts was based on the availability of vacant buildings. For example, Dias and Gomes (2018) state in page 182:

"The geographical distribution of the specialised sections of the new courts, taking advantage of existing court buildings by grouping those sections together, was carried out without any comprehensible criteria, apparently on the basis of making use of existing physical structures. The judicial actors themselves have reported many examples of odd locations and discontented citizens. No known studies justify these options and there have been instances of specialised sections being placed in geographical areas where the caseload is low, while in other areas in the same district there is no section to deal with the existing caseload."

2.5 Empirical Analysis of Judicial Assignments

Next, we evaluate whether the judge allocation decisions described in the official documentation are supported by the data. Specifically, in Table 1, we analyze data from the Ministry of Justice on caseload volume and judge assignments around the 2014 map reform to test two premises. The first is that the number of judges allocated to each court was unrelated to the inflow of case backlogs. The second is that there was limited effort to assign more experienced judges to courts with larger backlogs.

[5] Recall that the reform was developed and implemented under external pressure as part of the "bailout deal". The judicial map reorganization faced strong opposition from various entities, including local authorities, lawyers, and judges, all of whom perceived it as rushed. The main criticisms were the lack of planning and the absence of a transition period. Several testimonials by these stakeholders are provided in the following link (in Portuguese): https://portal.oa.pt/media/117987/nº130-set2015.pdf. There is ample anecdotal evidence that the abrupt change in the judicial map caused confusion and disruption in court operations. The massive migration of cases (around 3.5 million) to the new courts also caused the collapse of the IT platform (CITIUS) that stores all legal documents. The platform became fully operational by the end of October, almost two months after the new judicial map was implemented.

In Panel A, we analyze municipality-level data covering all jurisdictions in Portugal. The explanatory variable is the number of judges assigned to each new local court as of September 2014. Since each newly created court covers multiple municipalities, we cluster standard errors at the court level.

The dependent variable in column (1) is the number of incoming cases over the five years following the 2014 judicial map reform. The point estimate, approximately 200, aligns with the policymakers' premise that each judge should handle 200 cases per year. In column (2), the dependent variable is the number of backlog cases that migrated to the jurisdiction in September 2014. On average, each judge inherited 850 pending cases at the time of the reform—more than four times the threshold used to allocate judges—suggesting that backlogs were not a factor in judicial assignment decisions. Column (3) examines backlog cases per judge, with a negligible and statistically insignificant point estimate, confirming that the number of judges allocated was unrelated to the congestion experienced in September 2014.[6]

In Panel B, we analyze data at the judge level. The sample includes all judges transferred from old courts to the newly created courts handling insolvency cases in September 2014. The dependent variable is the number of backlog cases per judge. In column (1), the explanatory variable is the number of corporate insolvency cases handled in 2013 at the judge's previous court. In column (2), it is the average duration (in months) of insolvency processes in 2013 at the same court. Standard errors are clustered at the court level. While both point estimates are positive, they are statistically insignificant.

[6]As discussed in the next section, we use in our analysis the number of backlog cases per judge to measure court congestion.

3 Data and Variables

3.1 Data

The data for our analysis is obtained from CITIUS, the official IT platform used by the Portuguese judicial system and maintained by the Ministry of Justice. From this platform, we obtain records of all bankruptcy cases filed in Portugal in 2014 that remained pending as of September 1, 2014. This ensures that these bankruptcy cases were filed in one of the old courts and transferred on this date to a new court. We also obtain from the Ministry of Justice the number of judges and pending cases transferred to each newly created court. We use these variables to create our measure of court congestion.

For each bankruptcy case, we identify the insolvent firm and all firms that filed creditor claims. We also obtain the date of filing, any events pertaining to court rulings in the process, and court identifiers. Using their fiscal numbers, we match debtors and creditors to a firm-level dataset containing financial information for all firms in Portugal. This matching allows us to restrict our sample to non-financial creditor firms with debts due from the bankrupt firm. We follow these firms from 2007 to 2020.

3.2 Court Congestion

Following Ponticelli and Alencar (2016), Iverson (2018), and Müller (2022), among others, we measure court congestion as the ratio of pending cases to the number of judges in the court where the bankruptcy case is transferred to on September 1, 2014. Figure 1 shows large differences in congestion levels across jurisdictions. As explained in Section 2, when judges returned from judicial holidays to their newly assigned courts, they faced widely varying backlogs of bankruptcy cases to handle. Section 2 also provides evidence that these differences in backlog seem to be driven primarily by inefficiencies in judge allocation.

According to our definition, a congested court is one where judges have a high average case workload. Judges' intervention is required at several stages of the bankruptcy process, and a higher workload typically results in longer delays at these stages, thereby extending the overall bankruptcy process. Pereira and Wemans (2018) document significant variation in the duration of corporate bankruptcy cases in Portugal, with much of this variation driven by court location. Their study indicates that in 2015, the 25th percentile of case duration until closure was 10 months, while the 75th percentile was 70 months. They further show that median duration varied substantially across courts, ranging from 8 months in the fastest courts to 57 months in the slowest.

Given how the judicial map reform was planned and executed, our measure of congestion likely captures plausible exogenous variation. To verify this, we examine below how our congestion measure correlates with various local characteristics.

3.3 Dependent Variables

Our main dependent variable is *Bankrupt*, an indicator of whether a supplier with outstanding debt from a firm in bankruptcy also enters bankruptcy. In our robustness we use as alternative the variable *operating*, an indicator of whether the supplier has positive sales. Both variables capture the conditional probability that a supplier experiences financial distress after one of its clients with unpaid debts enters bankruptcy.

3.4 Other Variables

All explanatory variables are measured as of 2013, one year before the reform. *Creditor sales* is in million of euros. We measure the creditworthiness of the creditor with *Probability of default*, which we obtain from Antunes et al. (2016). The ratio *Debtor / Creditor assets* measures the relative size of the insolvent debtor relative to the creditor. *Receivables* is scaled by the firm's assets and it measures the importance of outstanding trade credit for the creditor. We winsorize all ratios at the 1 percent level to reduce the influence of outliers.

3.5 Descriptive Statistics

Table 2 provides summary statistics for our sample. All variables are measured as of 2013, except for court congestion, which is measured as of September 2014. As explained in Section 2, our main analysis focuses on jurisdictions with commercial courts, excluding cases from preexisting commercial courts prior to the judicial map reform. In robustness tests, we relax this sample restriction and include all courts in the analysis. Our sample includes 3,101 suppliers with outstanding debts from 370 firms with bankruptcy declared in 2014 but before the introduction of the new court map. The analysis focuses on creditors and debtors with financial statement data in 2013. The total number of debtor-creditor links is 54,418.

The average court congestion is 922.5, representing the number of pending cases assigned to each judge as a result of the judicial map reorganization in September 2014. The standard deviation highlights substantial variation in congestion levels across courts, as shown in Figure 1. Creditor size also varies significantly, as reflected in the high standard deviation of creditor sales. On average, debtors are nearly five times larger than creditors, a result driven by a few very large debtors, as the median debtor-to-creditor asset ratio is 0.34. Finally, accounts receivable make up 35 percent of assets, on average.

4 Empirical Methodology

Our data are at the debtor-creditor pair, where a given debtor i is a bankrupt firm and creditor j is a supplier with defaulted trade credit. Each bankrupt debtor in our sample owes money to multiple creditors, and a few specific creditors (such as internet providers) are owed money from multiple bankrupt debtors.

We start by examining whether characteristics of judicial districts can predict court congestion. We perform this analysis at the debtor-creditor level to maintain the relative importance of jurisdictions. We estimate the following cross-sectional regression:

$$Log(Congestion)_{i(c)} = \alpha + \gamma W_c + \epsilon_{ijc}. \quad (1)$$

The subscript c denotes the judicial district where the bankruptcy case was filed. The dependent variable is the log of the number of backlog cases per judge in the court where the bankruptcy case is transferred to on September 1, 2014. W_c is a vector of district-level characteristics in 2013.

Next, we examine the effect of court congestion on the probability that supplier j with defaulted credit from bankrupt client i also files for bankruptcy using the following equation:

$$Bankrupt_{i(c)jt} = \alpha_{ij} + \alpha_{rt} + \alpha_{st} + \sum_{t \neq -1} \beta_t I_t Log(Congestion)_{i(c)} + \sum_{t \neq -1} \Gamma_t I_t X_j + \epsilon_{i(c)jt}. \quad (2)$$

$Bankrupt_{i(c)jt}$ equals one if supplier j enters bankruptcy at time t, and equals zero otherwise. Once a supplier files for bankruptcy we leave the variable set at one, so that β_t measures the cumulative effective of court congestion. I_t is an indicator variable equal to 1 in year t. $Congestion_{i(c)}$ is the number of backlog cases per judge in court c, where the bankruptcy case is transferred to on September 1, 2014. The vector X_j denotes several observable characteristics of the supplier measured as of 2013.

We saturate the model with various fixed effects. α_{ij} are creditor-by-debtor fixed effects, which force comparison over time of the same creditor-debtor pair. α_{rt} are region-by-year fixed effects, which control for time-varying regional shocks. α_{st} are industry-by-year fixed effects, which control for time-varying shocks affecting industries.[7] We double-cluster standard errors at both the creditor and judicial district levels.

[7]Regions correspond to the first two digits of the municipality codes available in the balance sheet data and correspond to 18 districts (*distritos*) in mainland Portugal and 4 distinct divisions in the Azores and Madeira islands). We use CAE codes to classify industries into 21 groups based on firms' primary economic activities. The CAE code system is very similar to NACE, the standard system used for statistical purposes within EU countries.

The main identification assumption in equation 2 is that differences across jurisdictions in court congestion are unrelated to local economic conditions. The discussion in Section 2 supports this assumption, first noting and then empirically confirming that judges were assigned to the newly created courts based on projected case inflows rather than existing case backlogs. Policymakers appear to have overlooked the implications of merging smaller judicial districts into larger ones and transferring pending cases accordingly, which led to significantly higher backlog volumes in some courts.

However, we cannot not fully rule out a possible reverse causality concern. If for example firms located in a particular jurisdiction are hit by a negative shock, one should expect an increase in litigation activity (including bankruptcy filings) that may raise congestion levels in the affected courts. We address these concerns about the exogeneity of our congestion measure in several ways. First, we include industry-year and region-year fixed effects in equation 2 to control for any time-varying shocks affecting specific industries or regions. Second, we analyze whether cross-sectional differences in court congestion can be predicted by various economic characteristics of the judicial districts. Third, we examine the dynamic behavior of bankruptcy filings around the judicial map reorganization. Specifically, finding no evidence of pre-trends speaks against the reverse causality argument and thus reinforces the validity of our empirical strategy. Fourth, we demonstrate that our results hold when using an alternative empirical strategy based on comparing adjacent municipalities along judicial district borders.

5 Results

5.1 What Drives Differences in Court Congestion?

In Table 3, we investigate whether the level of court congestion is correlated with the economic characteristics of judicial districts. Finding no such correlations would strengthen the plausibility of court congestion being exogenous. We perform this test using the regression specified in equation 1. We keep the analysis at the debtor-creditor level to maintain the relative importance of jurisdictions. The dependent variable is the log of the number of backlog cases per judge in the court where the bankruptcy case is transferred to on September 1, 2014. All explanatory variables are either measured as of 2013 or expressed as 2012-2013 growth rates. We double-cluster standard errors at both the creditor and court levels.

Column 1 shows that court congestion is unrelated to labor market conditions, as measured by earnings per worker, earnings growth, and the unemployment rate. Column 2 employs two measures of the size of the banking sector: the number bank branches per capita and mortgage credit per capita. Both measures are statistically insignificant. Columns 3 and 4 employ measures related to the local economic structure. Court congestion in Column 3 is unrelated to per capita Gross value added (GVA) and to GVA growth, while in Column 4 it is unrelated to the share of manufacturing GVA and to the share of SMEs. All estimated coefficients are statistically insignificant and economically small.

In Column 5 we show that court congestion in 2014 is also unrelated to judicial characteristics of courts in 2013, prior to the reform of the judicial map. Disposition time is the average monthly duration of completed bankruptcy cases, a standard measure of court efficiency. Change backlog is the percent change in the number of pending bankruptcy cases in all pre-reform courts that were aggregated into the new jurisdictions. Both coefficients are economically small and statistically insignificant.

5.2 Bankruptcy Propagation

In this section, we study how court congestion affects the transmission of corporate bankruptcy through firm networks. In particular, we assess how congestion affects the probability that a supplier with defaulted trade credit from a bankrupt debtor also becomes bankrupt. Each observation is a debtor-creditor pair and the estimation includes only the commerce courts created in the 2014 reform. The empirical specification is presented in equation (2) and the estimation results are presented in Table 4. The dependent variable is an indicator of whether the supplier files for bankruptcy. Once this indicator switches from zero to one, it remains at one. For simplicity, the table presents the estimated cumulative effect of court congestion on the probability of bankruptcy six years following the judicial map reform. We double-cluster standard errors at both the creditor and court levels.

In column 1 we estimate a specification that includes only year fixed effects to account for aggregate economic conditions. In columns 2 and 3 we saturate the regression with region-by-year, industry-by-year, and creditor-by-debtor fixed effects. It also includes as control variables the supplier's probability of default and log sales, both measured in 2013 and interacted with year dummies. The region-by-year fixed effects absorb any time-varying regional shocks. The industry-by-year absorb any time-varying industry-specific shocks. The creditor-by-debtor fixed effects forces comparison along each creditor-debtor pair. The remaining control variables are intended to absorb differences among creditors in their risk profile.

The estimates in columns 1 and 2 convey the same message: a supplier with defaulted trade credit is more likely to go out of business if the debtor's bankruptcy procedure is adjudicated in a more congested court. The estimates are statistically significant and economically relevant. For example, take the point estimate in column 2 and suppose we increase court congestion by one standard deviation from its sample median. The estimate indicates that the probability the creditor becomes bankrupt increases by 0.75 percentage points. This represents 13 percent of the unconditional bankruptcy probability, which in our sample equals 5.9 percent.

We also note that the inclusion of the full set of control variables and fixed effects in column 2 does not alter significantly the estimate in column 1. This is important because it shows that our measure of congestion is indeed unrelated to local- and to industry-specific economic shocks.

In column 3 we interact the congestion variable with *Large debtor*, an indicator equal to 1 if the the ratio of debtor assets to creditor assets in 2013 is at or above the sample median. This variable has been used in the literature as an implicit measure of debtor importance (Jacobson and Von Schedvin (2015)). If the debtor is relatively large, its bankruptcy is more likely to push the supplier toward insolvency as well. The results in column 3 confirm this prediction. Bankruptcy propagation in congested courts only occurs when the debtor is relatively large, and the estimated effect is twice as large as that in column 2. When the debtor is relatively small, we find no effect of court congestion on the supplier's bankruptcy probability.

Table 4 provides the cumulative effect of court congestion six years after the reform. Figure 2 plots year-by-year coefficients to examine the dynamic behavior of bankruptcy filings surrounding the judicial map reorganization. The figure shows no evidence of pre-trends, which supports a causal interpretation of the results. Furthermore, the increase in bankruptcy filings appears gradual and persistent over time, rather than abrupt. This is a plausible adjustment, since bankruptcy contagion is likely to evolve progressively. The point estimate for the year 2020 matches the one in column 1 of Table 4, which captures the average cumulative effect across all debtor-creditor pairs.

As previously shown in column 3 of Table 4, bankruptcy propagation occurs when the debtor is large relative to the supplier. Figure 3 extends the year-by-year analysis by separating the sample into subsamples based on whether the debtor-to-creditor size ratio is above or below the median. The blue line remains indistinguishable from zero across all years, indicating no evidence of bankruptcy propagation when the debtor is relatively small. In contrast, the red line reveals that propagation is exclusively driven by cases where the debtor is large, making its bankruptcy more likely to have financial repercussions for the supplier.

5.3 Alternative Specifications and Robustness Tests

Table 5 demonstrates the robustness of our main result (Table 4, column 2) to alternative specifications. Our baseline sample includes the 16 commerce courts established in the 2014 judicial map reform, excluding cases transferred from existing commercial courts, as well as less developed areas without commerce courts.

In column 1 of Table 5 we re-estimate the model including all courts. The point estimate remains statistically significant. In column 2, we return to the baseline sample and modify the dependent variable. We replace *Bankrupt* with an indicator equal to one if the creditor submits financial statements, which we use as a proxy for going concern. The point estimate remains statistically significant.

In column 3, we implement an alternative empirical strategy to address any lingering concerns about potential confounding economic conditions driving our results. Specifically, we follow Ponticelli and Alencar (2016) and restrict our sample to pairs of adjacent municipalities that belong to different judicial districts. These municipalities have different levels of court congestion but are otherwise comparable in observable characteristics, thereby controlling for potentially confounding economic shocks. The point-estimate in column 3 is economically large and statistically significant, corroborating our main result.

5.4 Mechanisms

Our findings show that when a firm undergoes bankruptcy in a congested court, its suppliers with defaulted trade credit are more likely to also go bankrupt. This suggests that busier courts amplify the transmission of corporate bankruptcy through firm networks. To explain this effect, we build on Jacobson and von Schedvin (2015), who identify two main channels through which corporate failures propagate: lost demand and credit losses. Next, we discuss how court congestion may affect each of these channels.

5.4.1 Lost Demand

The first channel is operational. When a firm goes out of business, its suppliers lose a client. If this client is important and difficult to replace, the firm's bankruptcy reduces the suppliers' future sales. The resulting revenue loss may make these suppliers economically unviable, increasing their own risk of bankruptcy.

However, it is not obvious how court congestion would amplify this operational channel. For this channel to explain our findings, we would need to observe affected suppliers facing more rapidly shrinking demand when the debtor's bankruptcy case is assigned to a busier court. However, one can argue that court congestion could plausibly have the opposite effect. A slower bankruptcy process may keep the insolvent firm operating longer, allowing its suppliers to maintain a business relationship while seeking alternative clients. If this is the case, more congested courts would actually dampen—rather than amplify—bankruptcy propagation.

In Table 6 we perform two types of tests to shed light on this empirical question. First, we examine how court congestion affects for how long a bankrupt firm maintains its operations. We estimate the regression specified in equation 2, where the dependent variable is an indicator of whether the debtor reports sales. We find that court congestion does not significantly affect the duration of the bankrupt firm's operations. The point estimate in column 1 of Table 6 is positive but statistically insignificant at conventional levels. Therefore, court congestion does not seem to lead to a quicker demand contraction.

Second, we examine whether the degree of bankruptcy propagation is affected by the availability of alternative clients or by the specificity of the trade credit relationship. In column 2, we define the degree of market thickness as the logarithm of the number of firms in the same industry and region where the debtor operates. If the bankrupt debtor is one of many firms operating in the same region and industry, then it should be easier for the supplier to find new clients. If this is the case, court congestion should have a smaller probability to push the supplier to bankruptcy. In column 2, we add an interaction between the court congestion variable and an indicator of whether market thickness is above the sample median. The estimated coefficient is negative but statistically insignificant.

In column 3, we use the ratio of total R&D expenses to total assets in the creditor's industry as of 2013 to measure the specificity of a trade relationship. Jacobson and Von Schedvin (2015) argue that trade relationships are more exclusive and difficult to replace in more innovative industries. In column 3, we add an interaction between court congestion and an indicator of whether the industry's R&D intensity is above the sample median. If the increase in bankruptcy propagation in congested courts was driven by demand shrinkage, we should find a positive coefficient for the interaction term. However, the estimated interaction in column 3 is negative and insignificant.

Taken together, these results indicate that our main finding is not driven by a contraction in demand.

5.4.2 Credit Losses

The second channel is financial. When a debtor firm declares bankruptcy, its suppliers incur significant losses on defaulted trade credit. Court congestion amplifies these losses by delaying the resolution of bankruptcy cases. A slower resolution means that creditors are paid later and receive less, as the firm's assets may depreciate over time and judicial costs accumulate (Bris et al, 2006; Thorburn, 2000). As credit losses grow, the supplier's financial distress intensifies, raising the likelihood of insolvency.

We evaluate the importance of this channel in two steps. First, we examine in Table 7 whether court congestion affects suppliers' credit losses.[8] We use the same sample and empirical specification as in column 2 of Table 4. The dependent variable is cumulative im-

[8]There is very limited empirical evidence on how recovery rates for creditors depend on court congestion. Iverson (2018) shows that commercial banks report lower charge-offs on C&I loans in less congested courts. We are not aware of any study that provides comparable evidence for suppliers.

paired receivables over net account receivables. In column 1, the estimation sample includes only firms that submitted financial statements. To address potential attrition bias, column 2 also includes firms without financial information, for which we assign the last observed value. The regressions include region-by-year, industry-by-year, and creditor-by-debtor fixed effects. They also include as control variables the supplier's probability of default and log sales, both measured in 2013 and interacted with year dummies. We double-cluster standard errors at both the creditor and court levels.

The results in Table 7 confirm that court congestion amplifies credit losses for suppliers. The point estimate is consistent across the two specifications and captures the cumulative effect of court congestion on the fraction of trade credit losses six years after the judicial map reform. This estimate is both statistically significant and economically meaningful. Specifically, it indicates that a one-standard-deviation increase in court congestion from the sample median raises the average ratio of impaired to outstanding receivables by 1.8 percentage points.

Court congestion increases suppliers' credit losses. Following Jacobson and von Schedvin (2015), we examine whether financially constrained suppliers are disproportionately exposed to trade credit failures when courts are congested. To do so, we add an interaction term in equation 2 between the congestion variable and an indicator for firms in industries with high financing needs. Identification in this regression comes from comparing the effect of court congestion on bankruptcy for suppliers with stronger versus weaker financing frictions.

We present the results in Table 8. In column 1, we replicate our baseline result to facilitate the comparison of coefficients. In column 2, we interact the congestion measure with an indicator for suppliers in industries with high external finance dependence, and in column 3, with an indicator for high liquidity dependence.[9] The point estimates for the standalone congestion measure in columns 2 and 3 indicate no bankruptcy propagation in

[9] We calculate these two measures at the industry level using firm financial statements for Portuguese firms.

industries with low financing frictions. Suppliers in these industries have greater capacity to absorb liquidity shocks, making them more resilient to credit losses. In contrast, both interaction terms are statistically significant and economically meaningful, indicating that the bankruptcy propagation mechanism is primarily driven by financially constrained firms that are more vulnerable to credit losses.

6 Aggregate relevance

The results so far underscore the impact of court efficiency on the propagation of bankruptcies within trade credit networks. To explore the broader implications of our results, we conduct a simple counterfactual analysis to assess how different levels of court efficiency affect bankruptcy propagation rates. In this analysis we use the same baseline sample as Table 3.

Table 9 presents the number of bankruptcies under three scenarios. For each scenario, we measure bankruptcy propagation as the ratio of the number of bankruptcies within the debtor network to the total number of creditor-debtor pairs in the sample.

The base scenario contains the observed number of supplier bankruptcies in the network of bankrupt debtors. The average bankruptcy propagation rate in our sample is 5.9%. We then simulate two alternative counterfactual scenarios.

In the first scenario, we estimate the number of supplier bankruptcies that would occur if all courts were as efficient as the least congested court in our sample. Under this scenario, the bankruptcy propagation rate drops to 3.2%, nearly halving the transmission of corporate bankruptcy to trade creditors.

In the second scenario, we estimate the number of supplier bankruptcies that would occur if all courts operated at the efficiency level of the most congested court. In this case, the rate of bankruptcy propagation rises to 7.4%, a 26% increase relative to the base case.

These counterfactual findings underscore the substantial effect of court congestion on bankruptcy outcomes across the network. Our simulations show that moving from the least to the most efficient courts more than doubles bankruptcy propagation to trade creditors, highlighting the role of court efficiency in amplifying financial contagion within trade credit chains.

7 Conclusion

In this paper we analyze the effect of changing the quality of judicial enforcement on the transmission of bankruptcy through firm networks. As a natural experiment, we use the introduction of a new court map in Portugal and the transfer of bankruptcy cases that were previosuly assigned to generic courts.

Commerce courts had differing levels of congestion, but such congestion was unrelated with the characteristics of the places served by the courts. Using congestion in these courts as a source of exogenous variation in enforcement quality, we show that higher enforcement quality reduces the transmission of bankruptcy from bankrupt firms to firms in their network. In a counterfactual analysis, we show that decreasing congestion from observed levels to the level observed at the least congested court in the sample would reduce the transmission of bankruptcy from 6% to 3%.

Our findings have important policy implications. There is evidence that business bankruptcy institutions perform poorly in many countries, including developed economies (Djankov et al. (2008)). Our results suggest that strengthening these institutions could mitigate the cascading effects of financial distress, enhancing resilience to industry-specific shocks and reducing the risk of broader economic instability.

Figures

Figure 1: Court congestion

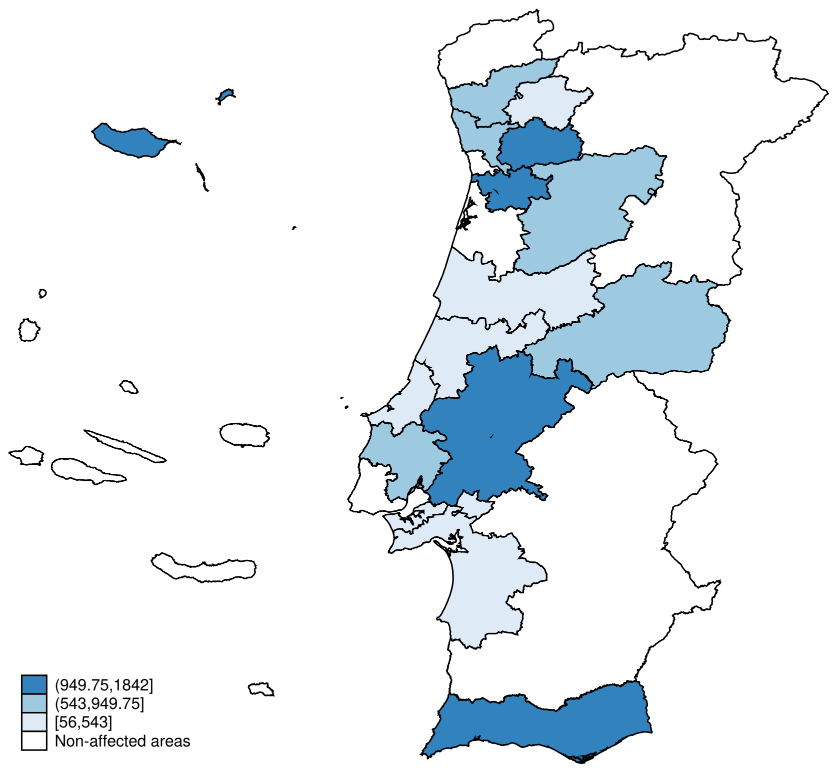

Notes: The figure depicts the average number of court cases per judge seat for newly created commerce courts.

Figure 2: Dynamic treatment effects: bankruptcy

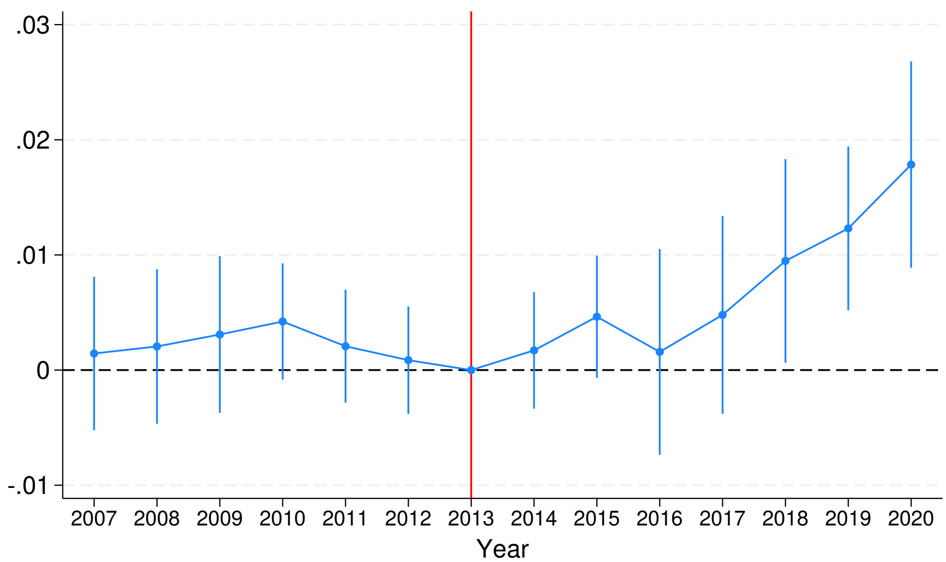

Notes: The figure depicts point estimates from equation (2). The dependent variable is an indicator equal to 1 if the firm has a bankruptcy record. Error bars denote 95% confidence intervals. We double-cluster standard errors at the court level and at the creditor level.

Figure 3: Dynamic treatment effects: bankruptcy, by debtor size

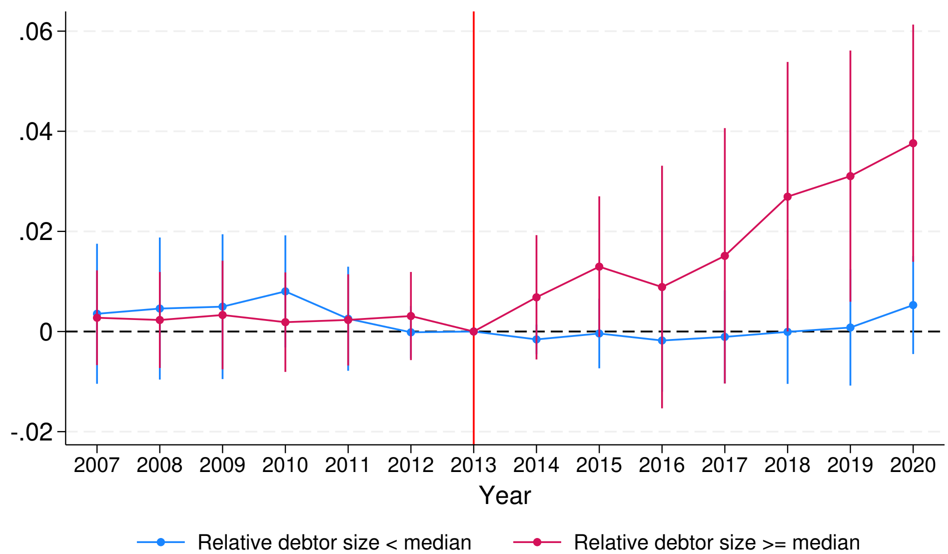

Notes: The figure depicts point estimates from equation (2). The dependent variable is an indicator equal to 1 if the firm has a bankruptcy record. Error bars denote 95% confidence intervals. We double-cluster standard errors at the court level and at the creditor level.

Tables

Table 1: Judge Assignment

Panel A: Court caseload

	Incoming cases	Backlog cases	Backlog per judge
	(1)	(2)	(3)
# Judges assigned	202.481***	850.583***	-3.080
	(31.594)	(199.069)	(41.240)
Observations	308	308	308

Notes: The unit of analysis is the municipality, and the sample includes all municipalities in Portugal. Each court created in the 2014 reform serves multiple municipalities. *Judges assigned* refers to the number of judges assigned to the new local court in September 2014. The dependent variable in column (1) is the number of incoming cases during the five years following the 2014 court map reform. In column (2), it is the number of backlog bankruptcy cases that migrated to the local court in September 2014. In column (3), the dependent variable equals the number of backlog cases from column (2) divided by the number of judges assigned to that jurisdiction. Standard errors are clustered at the court level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel B: Judge experience

	Backlog per judge	
	(1)	(2)
# Cases in previous court	0.269	
	(1.278)	
Average case duration in previous court		5.786
		(5.086)
Observations	129	129

Notes: The unit of analysis is the judge. The sample includes all judges transferred from the old courts to the newly created courts in September 2014 that handle insolvency cases. The dependent variable is the number of backlog cases that migrated to the new court in September 2014, divided by the number of judges assigned to that court. In column (1), the explanatory variable is the 2013 number of corporate insolvency cases handled in the old court where the judge was assigned. In column (2), it is the average duration (in months) of insolvency processes in 2013 in the old court where the judge was assigned. Standard errors are clustered at the court level. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 2: Descriptive statistics

	Mean	Median	Std. dev.
Congestion	922.500	818.500	(430.379)
Creditor sales	149.535	1.761	(683.981)
Probability of default	0.045	0.021	(0.066)
Debtor/creditor assets	4.759	0.338	(16.245)
Receivables	0.350	0.328	(0.223)
Observations		54,418	
Number of creditors		3,101	
Number of debtors		370	

Notes: The table depicts descriptive statistics for creditor-debtor bankruptcy linkages in the sample.

Table 3: Court congestion and county characteristics

	Log(Congestion)				
	(1)	(2)	(3)	(4)	(5)
Earnings/worker	-0.146				
	(0.547)				
Earnings growth	0.071				
	(0.052)				
Unemployment	0.022				
	(0.034)				
Banks per capita		0.082			
		(0.074)			
Mortgages per capita		-0.002			
		(0.002)			
GVA per capita			-0.001		
			(0.003)		
GVA growth			0.011		
			(0.009)		
Share manufacturing GVA				-0.001	
				(0.005)	
Share SMEs				-0.011	
				(0.015)	
Disposition time					-0.020
					(0.036)
Change backlog					0.314
					(0.391)
Observations	54,418	54,418	54,418	54,418	53,914

Notes: The table depicts estimates from equation 1 using county-level variables as independent variables. In column (1) the independent variables are the average earnings per worker, earnings per worker growth, and the unemployment rate. In column (2) the independent variables are the number of banks and mortgages per capita. In column (3) the independent variables are the gross value added per capita and the growth rate of the gross value added. In column (4) the independent variables are the gross value added of manufacturing firms as a share of total gross value added and the share of firms that are SMEs. In column (5) the independent variables are the average disposition time of insolvency cases and the change in the backlog of insolvency cases. Errors are double-clustered at the creditor and at the court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 4: Court congestion and bankruptcy propagation

	Bankruptcy		
	(1)	(2)	(3)
Log(Congestion)×(t=6)	0.017**	0.018***	0.005
	(0.007)	(0.005)	(0.006)
Log(Congestion)×Large debtor×(t=6)			0.031**
			(0.013)
Observations	52,542	52,542	52,542
Year FE	Yes	Yes	Yes
Controls	No	Yes	Yes
Region*Industry*Year FE	No	Yes	Yes
Creditor*Debtor FE	No	Yes	Yes

Notes: The table depicts estimates from equation 2. The dependent variable is an indicator equal to 1 if the creditor has a bankruptcy record. In columns (2) and (3) we use the creditor's 2013 probability of default and log sales, interacted with year fixed effects, as controls. Additionally, we control for region*industry*year fixed effects and debtor-creditor pair fixed effects. *Large debtor* is an indicator variable equal to one if the ratio of debtor assets to creditor assets in 2013 is at or above the median. Errors are double-clustered at the creditor and at the court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 5: Robustness tests

Dependent variable:	Bankruptcy	Operating	Bankruptcy
Courts in sample:	All	New commerce	Bordering commerce
	(1)	(2)	(3)
Log(Congestion)×(t=6)	0.009**	-0.017**	0.024**
	(0.004)	(0.006)	(0.010)
Observations	99,344	52,542	75,698
Controls	Yes	Yes	Yes
Region-Year FE	Yes	Yes	No
Industry-Year FE	Yes	Yes	Yes
Creditor-Debtor FE	Yes	Yes	Yes
County-pair-Year FE	No	No	Yes

Notes: This table shows estimates from alternative specifications to the baseline model (equation 2, and results in column 2 of Table 4). In column 1, we expand the estimation sample from newly created commerce courts to include all courts handling bankruptcy procedures, including both all commerce courts and non-specialized courts. In column 2, we change the dependent variable to a a dummy that equals one if the creditor has positive sales, and zero otherwise. In column 3, the dataset contains only bankruptcy filings from debtors headquartered in counties that are contiguous to other countries that are under the jurisdiction of different courts. Errors are double-clustered at the creditor and court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 6: The demand channel

	Operating debtor	Bankruptcy	
	(1)	(2)	(3)
Log(Congestion)×(t=6)	0.042	0.022**	0.035**
	(0.028)	(0.009)	(0.012)
Log(Congestion)×(t=6)×(Thickness>=median)		-0.009	
		(0.010)	
Log(Congestion)×(t=6)×(R&D>=median)			-0.018
			(0.013)
Observations	52,542	52,542	49,504
Controls	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes
Creditor-Debtor FE	Yes	Yes	Yes

Notes: The table depicts estimates from equation 2. In column (1) the dependent variable is equal to 1 if the debtor of the firm has positive sales in that year. In columns (2) and (3), the dependent variable is an indicator equal to 1 if the firm has a bankruptcy record. *Thickness* is the number of firms operating in 2013 in the same region and industry as the bankrupt debtor. *R&D/sales* is total industry sales over assets, measured for the industry of the creditor in 2013. We use the creditor's 2013 probability of default and log sales, interacted with year fixed effects, as controls. Errors are double-clustered at the creditor and at the court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 7: Impaired Receivables

	Impairment/Receivables (1)	(2)
Log(Congestion)×(t=6)	0.042***	0.042***
	(0.010)	(0.010)
Observations	49,151	51,189
Controls	Yes	Yes
Region-Year FE	Yes	Yes
Industry-Year FE	Yes	Yes
Creditor-Debtor FE	Yes	Yes
Closed firms	No	Yes

Notes: The table depicts estimates from equation 2. The dependent variable is the ratio of cumulative impaired receivables to net account receivables. In column (1) the sample contains only firms that report financial statements. In column (2) the sample contains all firms, in which we replace the impaired receivables ratio by the last observed value. We use the creditor's 2013 probability of default and log sales, interacted with year fixed effects, as controls. Errors are double-clustered at the creditor and at the court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 8: Bankruptcy and financing frictions

	Bankruptcy		
	(1)	(2)	(3)
Log(Congestion)×(t=6)	0.0178***	0.003	0.008
	(0.00457)	(0.005)	(0.008)
Log(Congestion)×High EFD×(t=6)		0.029**	
		(0.013)	
Log(Congestion)×High LD×(t=6)			0.018*
			(0.009)
Observations	52,542	52,542	52,542
Controls	Yes	Yes	Yes
Region-Year FE	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes
Creditor-Debtor FE	Yes	Yes	Yes

Notes: The table depicts estimates from equation 2. The dependent variable is an indicator equal to 1 if the firm has a bankruptcy record. *EFD* is an indicator variable equal to 1 if the external finance dependence of the industry of the firm is at or above the median. *LD* is an indicator variable equal to 1 if the liquidity dependence of the industry of the firm is at or above the median. We use the creditor's 2013 probability of default and log of sales, interacted with year fixed effects, as controls. Errors are double-clustered at the creditor and at the court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 9: Counterfactual number of bankruptcies

	Count	Var(%)	Propagation
Base	229	0	.059
All courts=most efficient	125	-45.415	.032
All courts=least efficient	289	26.201	.074

Notes: The table depicts the counterfactual number of bankruptcies under different scenarios. The base scenario depicts the number of observed liquidation bankruptcies in the network of bankrupt debtors included in the sample. The counterfactual scenario *all courts=most efficient* depicts the number of bankruptcies one would observe if all courts were as efficient as the most efficient court in the sample. The counterfactual scenario *all courts=least efficient* depicts the number of bankruptcies one would observe if all courts were as efficient as the least efficient court in the sample. *Propagation* is the ratio between the number of bankruptcies of firms in the network of the debtors included in the sample and the number of creditor-debtor pairs in the sample.

References

Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The network origins of aggregate fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Antunes, António, Homero Gonçalves, Pedro Prego et al., "Firm default probabilities revisited," *Economic Bulletin and Financial Stability Report Articles*, 2016, pp. 21–45.

Barrot, Jean-Noël and Julien Sauvagnat, "Input specificity and the propagation of idiosyncratic shocks in production networks," *The Quarterly Journal of Economics*, 2016, *131* (3), 1543–1592.

Djankov, Simeon, Oliver Hart, Caralee McLiesh, and Andrei Shleifer, "Debt enforcement around the world," *Journal of Political Economy*, 2008, *116* (6), 1105–1149.

Fazio, Dimas, Thiago Silva, and Janis Skrastins, "Court quality and economic resilience," *Available at SSRN 3517257*, 2021.

Gilje, Erik P, Elena Loutskina, and Philip E Strahan, "Exporting liquidity: Branch banking and financial integration," *The Journal of Finance*, 2016, *71* (3), 1159–1184.

Giroud, Xavier and Holger M Mueller, "Firms' internal networks and local economic shocks," *American Economic Review*, 2019, *109* (10), 3617–3649.

Iverson, Benjamin, "Get in line: Chapter 11 restructuring in crowded bankruptcy courts," *Management Science*, 2018, *64* (11), 5370–5394.

Jacobson, Tor and Erik Von Schedvin, "Trade credit and the propagation of corporate failure: An empirical analysis," *Econometrica*, 2015, *83* (4), 1315–1371.

Jordà, Òscar, Martin Kornejew, Moritz Schularick, and Alan M Taylor, "Zombies at large? Corporate debt overhang and the macroeconomy," *The Review of Financial Studies*, 2022, *35* (10), 4561–4586.

Kornejew, Martin, Chen Lian, Yueran Ma, Pablo Ottonello, and Diego J Perez, "Bankruptcy Resolution and Credit Cycles," Technical Report, National Bureau of Economic Research 2024.

Müller, Karsten, "Busy bankruptcy courts and the cost of credit," *Journal of Financial Economics*, 2022, *143* (2), 824–845.

Pereira, Manuel Coutinho and Lara Wemans, "How long does it take to enforce a debt in the Portuguese judicial system?," *Banco de Portugal Economic Studies*, 2018, *4*, 2.

Ponticelli, Jacopo and Leonardo S Alencar, "Court enforcement, bank loans, and firm investment: evidence from a bankruptcy reform in Brazil," *The Quarterly Journal of Economics*, 2016, *131* (3), 1365–1413.

NEW ANALYSIS

Table 10: Correlation with county characteristics

	Log congestion				
	(1)	(2)	(3)	(4)	(5)
Earnings/worker	-0.046				
	(0.608)				
Unemployment	0.022				
	(0.036)				
Banks per capita		10.587			
		(7.172)			
Mortgages per capita		-0.002			
		(0.002)			
GVA per capita			-0.002		
			(0.004)		
GVA growth			0.015*		
			(0.008)		
Share manufacturing GVA				-0.003	
				(0.004)	
Share SMEs				-2.154	
				(1.500)	
Disposition time					-0.003
					(0.040)
Change backlog					-0.005
					(0.444)
Constant	6.473***	6.575***	6.801***	222.107	6.717***
	(0.511)	(0.364)	(0.267)	(149.893)	(0.193)
N	54,418	54,418	54,418	54,418	52,948

The table depicts estimates from equation 1 using county-level variables as independent variables. In column (1) the independent variables are the average earnings per worker, earnings per worker growth, and the unemployment rate. In column (2) the independent variables are the number of banks and mortgages per capita. In column (3) the independent variables are the gross value added per capita and the growth rate of the gross value added. In column (4) the independent variables are the gross value added of manufacturing firms as a share of total gross value added and the share of firms that are SMEs. In column (5) the independent variables are the average disposition time of insolvency cases and the change in the backlog of insolvency cases. Errors are double-clustered at the creditor and at the court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 11: Correlation with county characteristics (SPAI counties)

	Log congestion					
	(1)	(2)	(3)	(4)	(5)	(6)
Earnings/worker	-0.068					
	(0.552)					
Earnings growth	0.059					
	(0.053)					
Unemployment	0.026					
	(0.032)					
Banks per capita		10.587				
		(7.172)				
Mortgages per capita		-0.002				
		(0.002)				
GVA per capita			-0.002			
			(0.004)			
GVA growth			0.015*			
			(0.008)			
Share manufacturing GVA				-0.003		
				(0.004)		
Share SMEs				-2.154		
				(1.500)		
Disposition time					-0.015	
					(0.028)	
Change backlog					0.011	
					(0.227)	
Firm switched headquarters county					-0.130	
					(0.105)	
Share governing coalition vote						-0.262
						(0.422)
N	54,418	54,418	54,418	54,418	52,332	54,418

The table depicts estimates from equation 1 using county-level variables as independent variables. In this table, we use data from SPAI to obtain the county of debtors and creditors. In column (1) the independent variables are the average earnings per worker, earnings per worker growth, and the unemployment rate. In column (2) the independent variables are the number of banks and mortgages per capita. In column (3) the independent variables are the gross value added per capita and the growth rate of the gross value added. In column (4) the independent variables are the gross value added of manufacturing firms as a share of total gross value added and the share of firms that are SMEs. In column (5) the independent variables are the average disposition time of insolvency cases, the change in the backlog of insolvency cases, and an indicator variable equal to 1 if the firm moved headquarters location until 2013. In column (6) the dependent variable is the fraction of voters who voted for the governing coalition in the 2013 local elections. Errors are double-clustered at the creditor and at the court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 12: Court congestion and bankruptcy propagation

	(1) Bankruptcy	(2) Bankruptcy	(3) Bankruptcy
Exposure*(Year=t+6)	0.017**	0.018***	0.004
	(0.007)	(0.005)	(0.006)
Exposure*Large debtor*(Year=t+6)			0.031**
			(0.013)
Observations	52,542	52,542	52,542
R-squared	0.020	0.544	0.544
Year FE	Yes	Yes	Yes
Controls	No	Yes	Yes
Region*Industry*Year FE	No	Yes	Yes
Creditor*Debtor FE	No	Yes	Yes

The table depicts estimates from equation 2. The dependent variable is an indicator equal to 1 if the creditor has a bankruptcy record. In columns (2) and (3) we use the creditor's 2013 probability of default and log sales, interacted with year fixed effects, as controls. Additionally, we control for region*industry*year fixed effects and debtor-creditor pair fixed effects. *Large debtor* is an indicator variable equal to one if the ratio of debtor assets to creditor assets in 2013 is at or above the median. Errors are double-clustered at the creditor and at the court levels. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.