Supply or Demand:
What Drives Fluctuations in the Bank Loan Market?

Carlo Altavilla — European Central Bank

Miguel Boucinha — European Central Bank

Paul Bouscasse* — Columbia University

October 2021

Abstract

We propose a new methodology to identify aggregate demand and supply shocks in the bank loan market. We present a model of sticky bank-firm relationships, estimate its structural parameters in euro area credit register data, and infer aggregate shocks based on those estimates. To achieve credible identification, we leverage banks' exposure to various sectors' heterogeneous liquidity needs during the COVID-19 Pandemic. We find that developments in lending volumes following the pandemic were largely explained by demand shocks. Fluctuations in lending rates were instead mostly determined by bank-driven supply shocks and borrower risk. A by-product of our analysis is a structural interpretation of two-way fixed effects regressions in loan-level data: according to our framework, firm- and bank-time fixed effects only separate demand from supply under certain parametric assumptions. In the data, the conditions are satisfied for supply but not for demand: bank-time fixed effects identify true supply shocks up to a time constant, while firm-time fixed effects are contaminated by supply forces. Our methodology overcomes this limitation: we identify supply and demand shocks at the aggregate and individual levels.

Keywords: credit demand, credit supply

JEL codes: E51, G21, G32

*We thank Olivier Darmouni, Matthieu Gomez, Veronica Guerrieri, Juan Herreño, Jennifer La'O, Cameron LaPoint, Emi Nakamura, Jón Steinsson, David E. Weinstein, Michael Woodford, and participants at the European Central Bank's (ECB) workshop on Banking Analysis for Monetary Policy. Paul worked on this paper while he was a PhD Trainee at the ECB. The views expressed in this paper are our own and do not necessarily reflect those of the ECB.

1 Introduction

Actual loan developments are the result of continuous interactions between demand and supply forces. The respective contributions of demand and supply, however, are not observable. At the same time, positive and normative statements on loan markets dynamics crucially depend on understanding their underlying source of fluctuations. For instance, credit supply restrictions where the bank credit rationing is driven by balance sheet constraints could be addressed by policies improving banks' intermediation capacity. Contractions in loan demand resulting from a drop in fixed investment should instead by countered with policies aiming at improving the return on investment. This illustrates why disentangling bank-supply shocks from firm-borrowing shocks is at the very core of the empirical banking literature.

The methodologies that have been used to separate demand and supply vary according to the econometric techniques and granularity of the data used in the empirical models. The studies broadly fall into three main categories: i) macro-econometric studies using aggregate data; ii) studies based on survey data; and iii) micro-econometric studies employing granular data.

Macro-econometric studies include both multivariate time-series models (Eckmeier and Ng, 2015, Moccero et al., 2014, Giannone et al., 2019, Altavilla et al., 2019) and DSGE models (Gerali et al., 2010, Christianio et al., 2010). The identification of demand and supply is typically achieved by imposing a set of restrictions (e.g. sign restrictions, in the case of vector autoregressive models) to the variance-covariance matrix of the multivariate model. These studies, while describing the aggregate lending dynamics, struggle to concretely capture lender and borrower heterogeneity.

Other studies use either aggregate or individual results obtained in regular lending surveys conducted by many central banks (including the FED, Bank of England, the ECB). The identification of demand vs supply in this case is facilitated by the fact that participants are asked to decompose their lending conditions into the contribution of demand and supply factors (Altavilla et al., 2021, Bassett et al., 2014, Del Giovane et al., 2011, Lown and Morgan, 2006). A drawback of these studies is that the self-reported dynamics they rely on can be distorted by noise and potential biases in banks' replies. Also, survey tend to report qualitative answers that in general cannot be accurately translated into quantitative statements.

Micro-econometric studies overcome many of the issues outlined above by using granular information at lender or transaction level collected in very large datasets, often confidential. Bank-level information is employed to analyse the impact of bank balance sheet characteristics on their ability to supply credit to firms and households (e.g., Kashyap and Stein, 2000). A widely used identification strategy employs credit registry data at the bank-firm level using firm-time fixed effects to control for changes in demand at the individual firm level, under the assumption that variation in lending to the same firm by different banks is driven by credit supply (Khwaja and Mian, 2008). One major drawback of studies relying on this approach is that they are unable to jointly identify demand and supply: they typically control for demand to obtain an identified effect on credit supply. An important contribution in the direction of simultaneously identifying and quantifying demand and supply is the one by Amiti and Weinstein (2018) and increasingly used in recent papers (Alfaro et al., 2021, Berg et al., 2021). The major drawback of these studies concerns their reliance on cross-sectional heterogeneity which does not allow for the identification of aggregate dynamics. In addition, studies relying on granular datasets tend to be focused on a specific country or loan segment (e.g. syndicated loans) which can limit their external validity.

Our approach combines the best of the macro and micro worlds: we use granular data for credible identification, while making statements about aggregate shocks. We do so based on structural assumptions.

Building on previous literature (Paravisini et al., 2020, Herreño, 2021), we lay down a model of sticky bank-firm relationships. Despite complex micro-foundations, the model boils down to simple demand and supply curves at the bank-firm level. Four parameters determine the slopes of those curves. On the demand side, those are the elasticity of substitution across banks, and the elasticity of substitution away from bank funding — these elasticities are defined with respect to the gross lending interest rate. On the supply side, the parameters are the elasticity of the bank's markup with respect to its market share, and the elasticity of the bank's total credit supply with respect to its lending rate. The latter is governed by the convexity of the bank's cost function — when it needs to fund more loans, does the marginal cost of funds increase? While we make structural assumptions on the market for bank loans, our framework is agnostic about the rest of the economy. Thus, one could embed it in a DSGE.

A by-product of our framework is a structural interpretation of the Amiti and Weinstein (henceforth AW) regression that we alluded to above. They propose distinguishing supply from demand through a regression of loan growth at the firm-bank level on time, bank-time, and firm-time fixed effects. Bank-time fixed effects, the reasoning goes, are specific to the bank, so they must be supply. Similarly, firm-time fixed effect must be demand. We show that, in our model, this is true only under certain conditions for the structural parameters. For the bank-time fixed effect to identify the true supply shocks, supply must be infinitely elastic. Otherwise, the fixed effects are contaminated by demand influences. For the firm-time fixed effect to identify the true demand shock, the two demand elasticities — across banks and away from bank funding — must be equal. When we take our model to the data, we find the first condition to be satisfied, but not the second one. Finally, as discussed above, even if these conditions are met, the AW framework does not allow for the identification of aggregate shocks to demand and supply, it does so only in the cross-section. A key contribution of our methodology is to overcome this limitation.

We then take our model to the data, focusing on the aftermath of the COVID-19 Pandemic. To kickstart identification, we propose a demand-driven instrument which leverages banks' heterogeneous exposure to each sector's liquidity needs during the pandemic. The crisis and the associated containment measures forced the temporary closure of a wide range of businesses, with consequences for their liquidity management and consequently, for credit demand. The asymmetry of the shock across firms allows us to build a demand-driven instrument which sharpens the identification. We find that banks do not raise their lending rate in reaction to demand shocks. In our model, this implies that if we regress the change in the lending rate on firm- and bank-time fixed effects, the obtained bank-time fixed effects are valid supply shifters — we can use them to identify the demand curve. Having identified the demand curve, we can recover the structural demand shocks, and in turn use these shocks as instruments to identify the supply curve. To sum up, our methodology allows us to identify two curves with a single external instrument.

We then show that, up to a first-order approximation, the model collapses to an aggregate demand and an aggregate supply curve. The slopes of these curves are determined by two parameters that we estimated in the granular data: the elasticity of a firm's total credit demand, and the elasticity of a bank's total credit supply. Therefore, we can recover aggregate demand and supply shocks by re-estimating the model on aggregate data. We emphasize that the two slopes are tightly disciplined by the granular data. In short, we used a micro estimation as an input to a macro estimation exercise. The output of the latter exercise is a decomposition of aggregate fluctuations in bank credit into demand, supply, and risk influences.

The granular analysis is based on AnaCredit, a new and unique a credit register covering all euro area countries. The heterogeneity of institutional frameworks, structural and cyclical economic performance across these countries strengthens the external validity of our results. Importantly, while our methodology leverages the granularity of the micro data for identification, it is parsimonious in the scope of information needed to implement it. It requires only 2 variables: lending volumes and lending rates. The macro estimation is also parsimonious in terms of data requirements: it is based on aggregate data on lending volumes, loan rates and borrower risk.

Our main empirical result is that developments in lending volumes were largely explained by demand shocks in the wake of the COVID-19 Pandemic. Fluctuations in lending rates were instead mostly determined by bank-driven supply shocks and borrower risk. This is because, in the micro data, we find a flat supply curve and a rather steep demand curve. Indeed, the inverse elasticity of supply, which determines the slope of the aggregate supply curve, is essentially 0: banks' loan supply was perfectly elastic — a result that might have to do with the significant support from different policy areas during the period. On the other hand, the elasticity of credit demand is between 1 and 2 — for a 100 basis point increase in their lending rates, firms contract borrowing by 1–2%.

We also estimate two parameters that have no bearing on our aggregate result, but that are interesting in themselves: the semi-elasticity of a bank's markup to its market share, and a firm's elasticity of substitution across banks. The interest rate mark-up is virtually unchanged following a demand-driven change in the market share. The elasticity of substitution across banks is between 7 and 23 depending on the methodology used for estimation. This implies that, when faced with a 100 basis points increase in the lending rate charged by a certain bank, a firm diverts a percentage in the range of 7–23% of its borrowing towards other banks.

The remainder of the paper is organized as follows. Section 2 outlines the model including firms' credit demand, banks' profit maximisation problem and resulting supply curve and the solution of the model. This section also presents a structural interpretation of the Amiti-Weinstein regression. Section 3 provides an overview of the data used in the application, including loan level and aggregate data. Section 4 presents the identification strategy and the results of the estimation of the elasticities of interest based on micro data. Section 5 includes the aggregation exercise, macro estimation and subsequent decomposition of the movements in lending volumes and rates into the contribution of demand, supply, and borrower risk. Section 6 concludes.

2 Environment

2.1 Firms

We assume that firm f's demand for credit obeys a nested constant elasticity of substitution (CES) demand system:

$$\text{bank-specific:} \qquad L^D_{fbt} = \xi^D_{fbt} \left(\frac{R_{fbt}}{R_{ft}} \right)^{-\gamma} L^D_{ft} \tag{1}$$

$$\text{firm-level:} \qquad L^D_{ft} = R^{-\varphi}_{ft} \Omega^D_{ft} \tag{2}$$

L^D_{fbt} is demand for credit from bank b: it is a function of a taste shock (ξ^D_{fbt}), the gross interest rate offered by bank b (R_{fbt}) relative to the interest rate index offered by all banks to firm f (R_{ft}), and total demand for credit (L^D_{ft}). That total credit demand is itself a function of the interest rate index offered to the firm (R_{ft}) and a demand shock (Ω^D_{ft}). The sum of the taste shocks is normalized to unity ($\sum_b \xi^D_{fbt} = 1$).

The parameters of interest are γ, the elasticity of substitution across banks, and φ, the elasticity of total credit demand.

We offer micro-foundations for these equations in the appendix (section A.1) based on Paravisini et al. (2020) and Herreño (2021). But they can also be seen as ad-hoc demand curves for bank credit that proxy for the stickiness of bank-firm relationships.

2.2 Banks

2.2.1 Simple Case: Risk-Free Loans and Constant Marginal Cost

Banks are monopolistic competitors: they maximize profits while taking into account the CES demand functions. For pedagogical purposes, we first consider a simple case without risk and with constant marginal cost of funding for the bank. Profits are then given by:

$$\sum_f R_{fbt} L^S_{fbt} - R^*_{bt} \left(\sum_f L^S_{fbt} \right)$$

Maximization subject to equations (1) and (2) implies:[1]

$$R_{fbt} = \mathcal{M}\left(s_{fbt}\right) R^*_{bt} \tag{3}$$

The markup, $\mathcal{M}(.)$, is defined as:

$$\mathcal{M}\left(s_{fbt}\right) \equiv \frac{\zeta_{fbt}}{\zeta_{fbt} - 1}$$

where ζ_{fbt} is a weighted average of the two elasticities:

$$\zeta_{fbt} \equiv \left(1 - s_{fbt}\right) \gamma + s_{fbt} \varphi$$

and s_{fbt} is bank b's market share in firm f's credit:

$$s_{fbt} = \frac{R_{fbt} L^D_{fbt}}{R_{ft} L^D_{ft}}$$

To better understand equation (3), consider two polar cases: $s_{fbt} = 0$ and $s_{fbt} = 1$. If $s_{fbt} = 0$, the bank represents an atomistic share of the firm's borrowing. In that case, the markup obeys the standard markup formula with CES demand: $\gamma/(\gamma - 1)$. Since the bank is small with respect to firm f's credit, the interest rate it charges does not influence the firm's total borrowing. So the bank does not take into account substitution away from bank lending; and φ, which is the elasticity of total borrowing, does not matter to the markup. On the other hand, with $s_{fbt} = 1$, the bank is a classic monopolist. It does not

[1] See the appendix for a detailed proof.

need to worry about substitution towards other banks, governed by γ, and only cares about substitution away from bank lending. Thus, its markup is $\varphi/(\varphi-1)$. In general, when s_{fbt} is between 0 and 1, the bank takes both margins into account by weighting the elasticities according to its market share. Note that s_{fbt} remains an endogenous object that is determined by demand shocks and interest rates. Thus, equation (3) is an equilibrium relationship that ties the interest rate to the bank's cost of funding, but it alone does not determine the interest rate.

Equation (3) is not new to our paper. Endogenous markups are in fact a classic result in international economics, where they are used to explain the behavior of the terms of trade and real exchange rate (Atkeson and Burstein, 2008). The concept was brought to the banking literature by Herreño (2021).

2.2.2 General Case

In general loans are risky and banks may face convex cost of funding. To introduce risk, we assume that firm f defaults with probability PD_{ft}, implying a probability of repayment of $1-PD_{ft}$. To account for banks' possibly heterogeneous perception of risk, we assume that bank b's subjective probability of repayment by firm f is $\xi^S_{fbt}(1-PD_{ft})$. To model convex costs of funding, we make the marginal cost of funds depend on total lending. That is, total cost of funds is given by: $R^*_{bt}\left(\sum_f L_{fbt}\right)^{1+\chi^{-1}}/(1+\chi^{-1})$. Thus, the bank's profit function becomes:

$$\sum_f \xi^S_{fbt}\left(1-PD_{ft}\right)R_{fbt}L^S_{fbt} - \frac{R^*_{bt}}{1+\chi^{-1}}\left(\sum_f L^S_{fbt}\right)^{1+\chi^{-1}}$$

Of course, this formulation collapses to the special case of section 2.2.1 when $\xi^S_{fbt}=1$, $PD_{ft}=0$ and $\chi^{-1}=0$.

Maximization implies:

$$R_{fbt}=\mathcal{M}(s_{fbt})\frac{R^*_{bt}\left(L^S_{bt}\right)^{\chi^{-1}}}{\xi^S_{fbt}(1-PD_{ft})}$$

Like equation (3), the latter relationship means that the interest rate is a variable markup over marginal cost. Unlike in equation (3), marginal cost is not constant as long as $\chi^{-1}>0$. Moreover, it is adjusted for risk ($\xi^S_{fbt}(1-PD_{ft})$).

χ, the elasticity of credit supply, is a parameter that we will seek to estimate in section 4.

2.3 Summary

Our model can be summarized by two equations, the demand and supply curves at the firm-bank level:

$$\text{demand:} \qquad L^D_{fbt} = \xi^D_{fbt} R^{-\gamma}_{fbt} R^{\gamma-\varphi}_{ft} \Omega^D_{ft} \tag{4}$$

$$\text{supply:} \qquad R_{fbt} = \mathcal{M}(s_{fbt}) \frac{R^*_{bt}\left(L^S_{bt}\right)^{\chi^{-1}}}{\xi^S_{fbt}(1-PD_{ft})} \tag{5}$$

For some theoretical results and estimation exercises, we'll use log.-changes from $t-1$ to t, which we denote with a tilde ($\tilde{.}$). While the demand curve is linear in logarithms, the supply curve requires a first-order approximation. We log.-linearize around a situation where $s_{fbt} = \bar{s}$. Since interest rates are close to 0 — in our sample, the mean is below 2% —, this approximation can only be reasonable. We obtain:

$$\text{demand:} \qquad \tilde{L}^D_{fbt} = -\gamma \tilde{R}_{fbt} + (\gamma - \varphi)\tilde{R}_{ft} + \tilde{\Omega}^D_{ft} + \tilde{\xi}^D_{fbt} \tag{6}$$

$$\text{supply:} \qquad \tilde{R}_{fbt} \approx \frac{\mu'' \tilde{R}_{ft} + \tilde{\pi}_{ft} - \tilde{\Omega}_{bt} + \chi^{-1}\tilde{L}^S_{bt} + \mu\tilde{\xi}^D_{fbt} - \tilde{\xi}^S_{fbt}}{1+\mu''} \tag{7}$$

where:

$$\mu \equiv \frac{\mathcal{M}'(\bar{s})}{\mathcal{M}(\bar{s})} \qquad \mu' = \bar{s}\mu \qquad \mu'' \equiv (\gamma-1)\mu' \qquad \tilde{\pi}_{ft} \equiv -\mathrm{d}\log(1-PD_{ft}) \qquad \Omega^S_{bt} \equiv \frac{1}{R^*_{bt}}$$

Notice the change of notation from R^*_{bt} to Ω^S_{bt}. We do this so that an increase in Ω^S_{bt} represents a positive supply shock.

While the approximation of the supply curve is bound to be good, there is a potential problem with equations (6–7): they do not allow for the possibility that ξ^D_{fbt} or $\xi^D_{fb,t-1}$ be equal to 0, as they feature the logarithms of both of these quantities. In practice, this means that we would have to ignore new or disappearing relationships. In section 4, we will answer this concern by estimating equations (4–5) in level as well as in log.-changes.

2.4 Structural Interpretation of the Amiti-Weinstein Regression

In order to distinguish demand from supply on the loan market, Amiti and Weinstein (2018), henceforth AW, run a regression of the form:

$$\tilde{L}_{fbt} = c^L_t + \alpha^L_{ft} + \beta^L_{bt} + \epsilon^L_{fbt}, \qquad \mathrm{E}\epsilon^L_{fbt} = 0 \tag{8}$$

where c_t^L, α_{ft}^L and β_{bt}^L are time, firm-time and bank-time fixed effects. The left-hand side is the log.- or percentage change in the stock of loans at the firm-bank level. The firm-time fixed effect catches what is specific to the firm, hence should capture demand influences according to AW; while the bank-time fixed effect should capture supply influences.[2] Given its simplicity, this identification scheme has been used in other papers (Amado and Nagengast, 2016, Amity et al., 2017, Alfaro et al., 2021). Other authors have run regressions of the same flavor, with region-time (Berton et al., 2018, Greenstone et al., 2020) or region-industry-size-time (Degryse et al., 2019) fixed effects, instead of firm-time fixed effects.

Before interpreting equation (8), note that it requires a normalization. Indeed, for any set of $(c_t^L, \alpha_{ft}^L, \beta_{bt}^L)$, equation (8) is observationally equivalent to:

$$\tilde{L}_{fbt} = c_t^{L'} + \alpha_{ft}^{L'} + \beta_{bt}^{L'} + \epsilon_{fbt}^L$$

where: $c_t^{L'} = 0$, $\alpha_{ft}^{L'} = c_t^L + \alpha_{ft}^{L'}$ and $\beta_{bt}^{L'} = \beta_{bt}^L$. There is in fact an infinity of potential normalizations.[3] Therefore, we follow AW and drop the first firm and bank dummies $(f, b = 1)$. So that α_{ft}^L and β_{bt}^L can be interpreted as relative to the first firm and bank. We denote with umlaut a deviation from the first firm or bank, for example: $\ddot{\tilde{R}}_{ft} = \tilde{R}_{ft} - \tilde{R}_{f=1,t}$, or $\ddot{\tilde{L}}_{bt} = \tilde{L}_{bt} - \tilde{L}_{b=1,t}$.

What does this regression identify within our model? Ideally, we would like the firm-time fixed effect to identify the true demand shock $(\tilde{\Omega}_{ft}^D)$, perhaps combined with the default probability $(\tilde{\pi}_{ft})$; while the bank-time fixed effect should identify the true supply shock $(\tilde{\Omega}_{bt}^S)$. To solve for $\tilde{L}_{fbt}$, we can substitute out $\tilde{R}_{fbt}$ in equation (4) thanks to equation (5):

$$\tilde{L}_{fbt} = \underbrace{\tilde{\Omega}_{ft}^D - \underline{\gamma}\tilde{\pi}_{ft} + (\underline{\gamma} - \varphi)\tilde{R}_{ft}}_{\text{firm specific}} + \underbrace{\underline{\gamma}\left(\tilde{\Omega}_{bt}^S - \chi^{-1}\tilde{L}_{bt}\right)}_{\text{bank specific}} + \underbrace{(1 - \underline{\gamma}\mu)\tilde{\xi}_{fbt}^D + \underline{\gamma}\tilde{\xi}_{fbt}^S}_{\text{error term}} \tag{9}$$

where: $\underline{\gamma} = \gamma/(1 + \mu'')$.

Equation (9) makes clear that the simple intuition behind the AW regression is not necessarily true. What is firm- or bank-specific does not have to be demand or supply. For instance, $\tilde{R}_{ft}$, the interest rate index charged to firm f is firm-specific, yet an equilibrium object. Similarly, the bank-specific part features the bank's total lending $(\tilde{L}_{bt})$.

[2]This short summary does not do full justice to AW's paper. Part of their contribution is to propose an estimator that builds on equation (8). One advantage of their estimator is that it can handle new and disappearing relationships: with the log.-change on the left-hand side, one can have neither; with the percentage change, one cannot have new relationships for the denominator is 0. When there are no new relationships, their estimator is equivalent to estimating equation (8) by weighted least squares with the percentage change on the left-hand side, weighting each firm-bank pair by lagged lending $(L_{fb,t-1})$. In this section, we are being conceptual, hence we abstract from these subtleties.

[3]For instance, we could also have: $c_t^{L'} = 0$, $\alpha_{ft}^{L'} = \alpha_{ft}^L$ and $\beta_{bt}^{L'} = c_t^L + \beta_{bt}^L$, or even: $c_t^{L'} = 3c_t^L$, $\alpha_{ft}^{L'} = \alpha_{ft}^L - c_t^L$ and $\beta_{bt}^{L'} = \beta_{bt}^L - c_t^L$.

We formalize this point in proposition 1 — the proof is in the appendix.[4] To do so, we introduce some notations: let X_t denote the matrix that contains the time, firm-time and bank-time dummies for period t where each row is a pair (f, b). We denote Q_{ft}/Q_{bt} the row corresponding to firm f/bank b of matrix $(X_t'X_t)^{-1}X_t'$. Finally, $\tilde{\Xi}_t^L$ is a vector whose rows are the error terms $(1-\underline{\gamma}\mu')\tilde{\xi}_{fbt}^D + \underline{\gamma}\tilde{\xi}_{fbt}^S$.

Proposition 1 *The Amiti-Weinstein regression identifies:*

$$\hat{\alpha}_{ft}^L = \ddot{\tilde{\Omega}}_{ft}^D - \underline{\gamma}\ddot{\tilde{\pi}}_{ft} + (\underline{\gamma} - \varphi)\ddot{\tilde{R}}_{ft} + Q_{ft}\tilde{\Xi}_t^L$$
$$\hat{\beta}_{bt}^L = \underline{\gamma}\left(\ddot{\tilde{\Omega}}_{bt}^S - \chi^{-1}\ddot{\tilde{L}}_{bt}\right) + Q_{bt}\tilde{\Xi}_t^L$$

If $\varphi = \gamma$ and $\chi^{-1} = 0$, the Amiti-Weinstein regression identifies the structural demand and supply shifters, with some measurement error:

$$\hat{\alpha}_{ft}^L = \ddot{\tilde{\Omega}}_{ft}^D - \varphi\ddot{\tilde{\pi}}_{ft} + Q_{ft}\tilde{\Xi}_t^L \qquad \hat{\beta}_{bt}^L = \varphi\ddot{\tilde{\Omega}}_{bt}^S + Q_{bt}\tilde{\Xi}_t^L$$

This proposition confirms what was already apparent in equation (9): in general, the AW regression does not identify structural objects. This is potentially highly problematic for identification: for instance, a bank that is exposed to positive demand shocks would see its fixed effect lowered, i.e. positive demand shocks would be partially interpreted as *negative* supply shocks. Indeed, to service demand, banks raise more expensive funds. They pass this cost to all borrowers, which chokes off some of the demand, thus looking like a negative supply shocks. Luckily, we shall find in the empirical section that $\chi^{-1} = 0$ is reasonable, which means that the AW regression does identify the true supply shock. We are less optimistic, on the other hand, about the demand side since we find: $\gamma > \varphi$. We come back to these issues in section 4.

The second lesson from proposition 1 is that the AW regression only identifies relative objects.[5] Indeed the umlaut $\ddot{}$ denotes deviation from another firm or bank, which we have chosen to be the first ones; for instance: $\ddot{\tilde{\Omega}}_{ft}^D = \tilde{\Omega}_{ft}^D - \tilde{\Omega}_{f=1,t}^D$. Any aggregate demand or supply shock — a shock that affects all firms or banks symmetrically — would not be captured by this estimator. For example, suppose that the demand shock is the sum of an aggregate component ($\tilde{\Omega}_t^{Da}$) and an idiosyncratic one ($\tilde{\Omega}_{ft}^{Di}$): $\tilde{\Omega}_{ft}^D = \tilde{\Omega}_t^{Da} + \tilde{\Omega}_{ft}^{Di}$. What appears in the expression for $\hat{\alpha}_{ft}^L$ is: $\ddot{\tilde{\Omega}}_{ft}^D = \tilde{\Omega}_{ft}^{Di} - \tilde{\Omega}_{f=1,t}^{Di}$, which is only a function of the idiosyncratic shocks. The aggregate shocks will be captured by the time fixed effect. To illustrate this, suppose for simplicity that we are in the ideal case where $\gamma = \varphi$ and $\chi^{-1} = 0$; and that $\tilde{\Omega}_{ft}^D$, $\tilde{\pi}_{ft}$ and $\tilde{\Omega}_{bt}^S$ obey the aggregate/idiosyncratic

[4] In their appendix, Chang et al. (2021) analyze a close cousin of the AW regression, the Khwaja and Mian (2008) one. The demand side of our proposition is reminiscent of one of their results.

[5] To be clear, we are not implying that AW say otherwise. The goal of this paragraph is to make a reader who wouldn't be familiar with their framework aware of this fact.

structure mentioned earlier. Then, the time fixed effect identifies:

$$\hat{c}_t^L = \underbrace{\tilde{\Omega}_t^{Da} - \varphi\tilde{\pi}_t^a + \varphi\tilde{\Omega}_t^{Sa}}_{\text{aggregate}} + \underbrace{\tilde{\Omega}_{f=1,t}^{Di} - \varphi\tilde{\pi}_{f=1,t}^i + \varphi\tilde{\Omega}_{b=1,t}^{Si}}_{\text{idiosyncratic}} + \underbrace{Q_{1t}\tilde{\Xi}_t^L}_{\text{error}}$$

The idiosyncratic term would be different if we had chosen to drop other dummies, but the aggregate part will be there as long as a time fixed effect remains in the regressors. It is a linear combination of aggregate demand and supply objects. Hence, even in the ideal parametric case ($\gamma = \varphi$ and $\chi^{-1} = 0$), the AW regression cannot separate aggregate demand from aggregate supply. It can only do so in the cross-section. A key part of our contribution is to overcome this limitation.

Finally, $Q_{ft}\tilde{\Xi}_t^L$ and $Q_{bt}\tilde{\Xi}_t^L$ appear in the formulas because the true fixed effects are estimated with error. In practice, this measurement error may imply spurious negative correlation between $\hat{\alpha}_t^L$ and $\hat{\beta}_t^L$. This problem is well-known in the labor literature where researchers have been running regressions such as equation (8) since Abowd et al. (1999).[6]

3 Data

3.1 Loan-level Data

Our dataset, AnaCredit, is the new credit register of the euro area.[7] Data collection is harmonized across all Member States (19 countries). A bank must report to its country's national central bank every corporate loan for which its exposure to the borrower is above 25,000€. As a result, there are more than 25 million loans, observed every month. The process was initiated in 2011, and actual data collection started in September 2018. Some papers have already used data collected from the preparatory phase that gathered the independently-collected credit registers of different countries (Altavilla et al., 2020a,b), but our paper is, as far as we know, the first to use the harmonized data.

AnaCredit features hundreds of variables, but our approach is parsimonious and we will primarily use 2 of those: loans outstanding and interest rate. We restrict our attention to non-financial corporations, and collapse the data at the bank-firm-quarter level. In table 1, we report summary statistics for the full dataset (panel A, column (1)), the dataset where interest rates are not missing (panel B, columns (3–4)), the cleaned dataset that underlies the regressions in log.-changes (panel B, columns (1–2)) and the cleaned dataset that underlies the regressions in level (panel B, columns (3–4)). To go from columns (1–2) to columns (3–4) in panel A, we take out bank-firm-time observations for which the interest rate is missing on at least one of

[6] On this issue, known as "low-mobility bias", see Andrews et al. (2008) and Kline et al. (2020).
[7] See European Central Bank (2019) for more details.

the loans — there may be several loans outstanding within a relationship at a given time. In the appendix (section B), we describe in detail the cleaning choices that lead from panel A to panel B. There is significant attrition because we drop a firm-time set of observations if the interest rate is missing, or the reporting is suspicious, in any of its relationships. We prefer to err on the conservative side: otherwise, we would be constructing meaningless loan and interest rate indices when we estimate φ in section 4.2.1. The level dataset is smaller than the log.-change one because, in the log.-change case, we keep a firm-time if it survives the cleaning steps at time t and $t-1$; while in the level case, we keep a firm if it survives these steps at all times — in order to keep a balanced panel.

Table 1: Summary statistics

	(1) Loans (€)	(2) Rate (%)	(3) Loans (€)	(4) Rate (%)
Panel A: Full dataset				
	All		Rate non missing	
Mean	674,197	3.07	610,192	2.88
Median	103,510	1.61	100,837	1.68
St. dev.	9,447,703	17.84	7,934,400	15.98
IQR	249,754	1.50	242,681	1.51
# obs.	46,489,276	44,237,752	40,404,520	40,404,520
# banks	2,994	2,945	2,915	2,915
# firms	4,842,384	4,692,170	4,581,349	4,581,349
# quarters	10	10	10	10
Panel B: Clean dataset				
	Log.-change		Level	
Mean	548,965	1.91	535,631	1.96
Median	102,675	1.75	100,501	1.79
St. dev.	6,086,502	1.12	5,990,428	1.06
IQR	233,932	1.34	231,222	1.30
# obs.	27,415,200	27,415,200	21,437,578	21,437,578
# banks	2,721	2,721	2,161	2,161
# firms	3,972,785	3,972,785	2,705,339	2,705,339
# quarters	9	9	10	10

Note: summary statistics at the bank-firm-quarter level. See section 3.1 for details.

Overall, the cleaning steps work in the expected direction: since they have more loans within a relationship, bigger firms are more likely to have the interest rate missing on one of their loans. Therefore, the mean size of a relationship slightly declines. Finally, notice that the standard deviation of the interest rate is suspiciously high before cleaning. In the early part of the sample, some banks sometimes report interest rates in percent instead of decimal numbers: 0.05 (5%) becomes 5 for instance. To deal with this issue, we divide by 100 when the interest rate is above 0.3 (30%), and trim the bottom and top percentiles. These operations affect the mean and standard deviation, but the median stays similar, which is reassuring.

3.2 Aggregate Data

We use 3 time series on aggregate loans in the euro area: loans outstanding, lending rates on loans outstanding and average probability of default. The first two are standard series, available from the ECB's Statistical Data Warehouse. The third one is based on bank's confidential regulatory reporting. The economy-wide probability of default is computed by aggregating data reported by each bank. These series are plotted in figure 1 since 2004 for loans and rates, since 2014 for the probability of default. Since we focus on the pandemic in the aggregate exercise, we only use the data from the last quarter of 2019 to the last quarter of 2020.

4 Micro Estimation

4.1 γ: Elasticity of Substitution across Banks

4.1.1 Identification Problem

Consider the demand curve, equation (6), which we reproduce here for convenience:

$$\tilde{L}^D_{fbt} = -\gamma \tilde{R}_{fbt} + (\gamma - \varphi)\tilde{R}_{ft} + \tilde{\Omega}^D_{ft} + \tilde{\xi}^D_{fbt}$$

In order to estimate γ, one may be tempted to run the following OLS regression:

$$\tilde{L}_{fbt} = -\gamma \tilde{R}_{fbt} + \theta^D_{ft} + \tilde{\xi}^D_{fbt} \tag{10}$$

where θ^D_{ft} is a firm-time fixed effect that would soak up $(\gamma - \varphi)\tilde{R}_{ft} + \tilde{\Omega}^D_{ft}$. Of course, the issue with this approach is that $\tilde{\xi}^D_{fbt}$ may be correlated with $\tilde{R}_{fbt}$ in equilibrium, as shown by equation (7). Besides its direct effect on $\tilde{R}_{fbt}$ through the variable markup (μ'), $\tilde{\xi}^D_{fbt}$ will affect $\tilde{L}^S_{bt}$. Thus an OLS estimation of equation (10) is potentially biased.

What we need is a supply shock that is uncorrelated with $\tilde{\xi}^D_{fbt}$, in order to instrument $\tilde{R}_{fbt}$ with it. To find it, we regress the change in the interest rate at the bank-firm level on time, firm-time and bank-time fixed effects:

$$\tilde{R}_{fbt} = c^R_t + \alpha^R_{ft} + \beta^R_{bt} + \epsilon^R_{fbt}, \qquad \mathrm{E}\epsilon^R_{fbt} = 0 \tag{11}$$

This is equation (8), a.k.a. the AW regression, except that we put the interest rate on the left-hand side. Equation (7) shows that, like $\tilde{L}_{fbt}$, $\tilde{R}_{fbt}$ can be expressed as the sum of firm-specific, bank-specific and error

Figure 1: Aggregate data

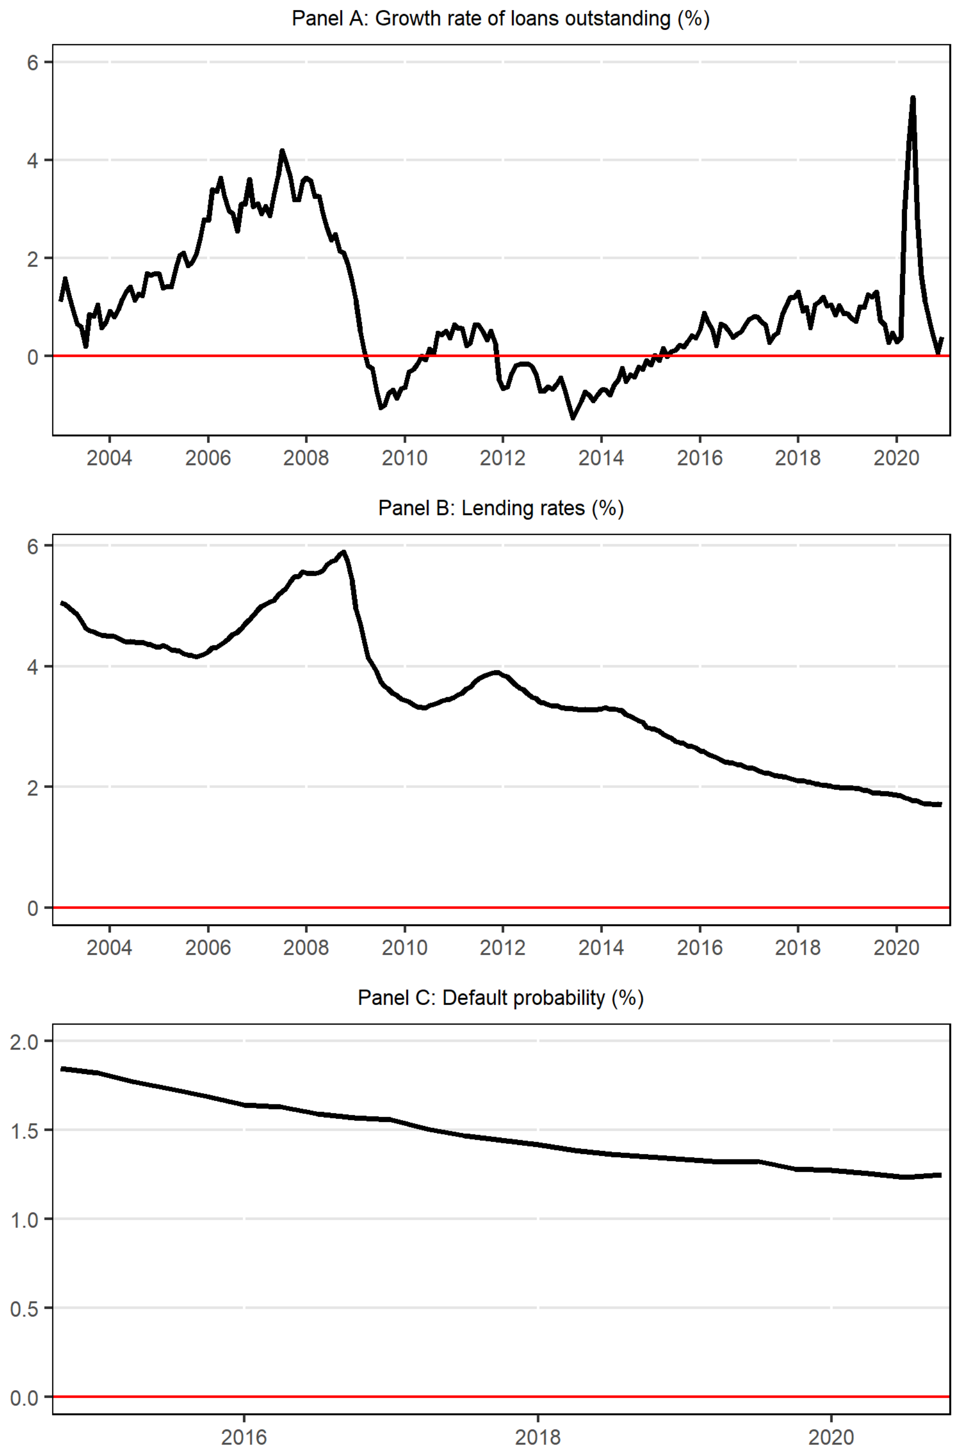

Note: quarter-on-quarter growth rate of loans outstanding (panel A), lending rates (panel B) and probability of default on loans to non-financial corporations in the euro area.

terms. Therefore the logic that underlies proposition 1 implies that once we run regression (11), we obtain the following bank-time fixed effects:

$$\hat{\beta}^R_{bt} = \underbrace{\frac{-\ddot{\tilde{\Omega}}_{bt} + \chi^{-1}\ddot{\tilde{L}}^S_{bt}}{1+\mu''}}_{\text{true fixed effect}} + \underbrace{Q_{bt}\tilde{\Xi}^R_{fbt}}_{\text{measurement error}} \tag{12}$$

where $\tilde{\Xi}^R_{bt}$ is the matrix whose rows contain the error terms that appear in equation (7): $\mu'\tilde{\xi}^D_{fbt} - \tilde{\xi}^S_{fbt}$.

Equation (12) raises two separate issues. The first one was already identified in proposition 1: as long as $\chi^{-1} > 0$, the true bank-time fixed effect is contaminated by demand through $\ddot{\tilde{L}}^S_{bt}$. The second one is that of measurement error: since Ξ^R_t contains ξ^D_{fbt}, $\hat{\beta}^R_{bt}$ might be correlated with the latter through the measurement error. The first problem is by far the most complicated to address, and we devote the next section to it.

4.1.2 A Demand-Driven Instrument

To deal with the first problem, we invoke an external demand-driven instrument. The COVID-19 crisis coincided with a sharp increase in lending during the first three quarters of 2020 (figure 2). Of course, this buildup might be due to demand as much as supply. But table 2 demonstrates that it was asymmetric across industries. There are large differences between industries that need emergency lending, such as hotels and restaurants, and those that were probably unaffected by the crisis such as agriculture. We will leverage banks' heterogeneous exposure to those industries to identify their supply curve.

To isolate the demand component of the lending buildup, we use its asymmetry across industries. That is, for each bank b, we construct:

$$XP_b = \sum_s w_{bst_0} \log\left(\frac{L_{st_1}}{L_{st_0}}\right), \qquad w_{sbt_0} = \frac{L_{sbt_0}}{L_{bt_0}}$$

where L_{st} is total lending to country-industry s at time t. The weights w_{sbt_0} are given by bank b's lending to country-industry s relative to its total lending before the pandemic. We also experiment with a version of this exposure measure where we leave out bank b's lending to construct the lending change $L_{s,-b,t_1}/L_{s,-b,t_0}$:

$$XP_b = \sum_s w_{bst_0} \log\left(\frac{L_{s,-b,t_1}}{L_{s,-b,t_0}}\right)$$

The dates t_0 and t_1 are the fourth quarter of 2019 (2019Q4), and the third one of 2020 (2020Q3). In the rest of the section, we omit time subscript as we run the regressions in changes, hence have a single time period.

Once we have constructed this measure of exposure, we follow a two-step procedure:

1. We run regression (11), with the change in the interest rate from 2019Q4 to 2020Q3 on the left-hand

Figure 2: Loans outstanding over the sample

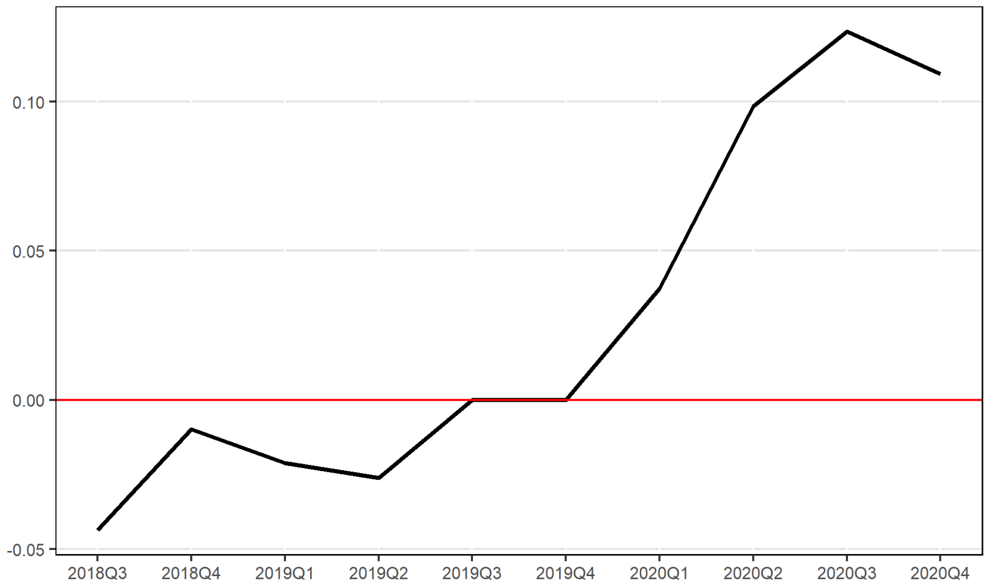

Note: Natural logarithm of loans outstanding in AnaCredit. 2019Q4 is normalized to 0.

Table 2: Loan growth in the first three quarters of 2020

Industry	Growth rate (%)
Accomodation and food services	28.6
Professional, scientific and technical activities	25.8
Administrative and support service activities	25.5
Arts, entertainment and recreation	23.9
Manufacturing	22.3
Information and communication	21.1
Transportation and storage	20.6
Education	20.2
Public administration	18.1
Wholesale and retail trade	12.3
Other services	11.2
Utilities	11.1
Construction	10.2
Water supply and waste management	9.0
Mining	8.2
Real estate	8.0
Agriculture	6.5
Human health and social work	1.4
Total	14.4

side. For each bank, this gives us a fixed effect $\hat{\beta}_b$.

2. We regress the estimated bank fixed effect on the bank's log.-change in total lending $\tilde{L}_b$, instrumenting with XP_b:

$$\hat{\beta}_b = c + d\tilde{L}_b + e_b^{\beta} \tag{13}$$

From equation (12), we can see that if XP_b is orthogonal to $\ddot{\tilde{\Omega}}_b$ and $Q_{bt}\tilde{\Xi}_{fbt}^R$, this procedure identifies $\chi^{-1}/(1+\mu'')$. We discuss the identification assumption after presenting the results.

There is, to say the least, limited evidence that $\chi^{-1}/(1+\mu'')$ is greater than 0 (table 3). This is good news for estimation as it implies that, at least in this context, estimating bank-time fixed effects is a good way to identify supply shocks within our model.

Table 3: Endogeneity of $\hat{\beta}_{bt}^R$

	(1) All	(2) LOO
$\chi^{-1}/(1+\mu'')$	-0.003*	-0.003
	(0.002)	(0.003)
# obs.	2,038	2,038
F-statistic	53	34

Note: Estimation of equation 13. In column (1), we use the exposure measure where industry-level lending growth includes bank b; in column (2), we use the version where we leave out bank b. Banks are weighted by pre-pandemic lending in step 2. See section 4.1.2 for details.

Could this result be the consequence of a failure of the identification assumption? As we have already mentioned, COVID exposure needs to be orthogonal to bank-specific supply shocks ($\ddot{\tilde{\Omega}}_b^D$) and the error term ($Q_{bt}\tilde{\Xi}_{fbt}^R$). Since the latter depends on firm-bank shocks, there is little reason to expect that it correlates with exposure, which loads on industry-level changes at the country level. Correlation with supply shocks, on the other hand, is more worrisome: it could be that banks that are exposed to sectors most affected by the pandemic cut lending in the face of a deteriorating balance sheet. That story, however, should bias us upward, not downward: those banks should raise their interest rate, not lower it. Hence it would reinforce our point: that χ^{-1} must be close to 0. For such bias to go downward, the correlation would need to be negative: that banks which are doing well are more exposed to troubled sectors. That seems implausible. One last possible objection is that there may be positive supply influences at the industry level. For instance, public loan guarantees might be targeted toward sectors that need emergency funding, thus making those loans less risky and lowering their interest rate. That kind of variation, however, would be soaked up by the firm-time fixed effects in the first step.

4.1.3 γ: Estimation and Results

We will run regression (10), using $\hat{\beta}^R_{bt}$ as an instrumental variable (IV). Indeed, since $\chi^{-1}/(1+\mu'')$ is indistinguishable from 0, we have, for all pratical purposes:

$$\hat{\beta}^R_{bt} = -\frac{\ddot{\tilde{\Omega}}_{bt}}{1+\mu''} + Q_{bt}\tilde{\Xi}^R_{fbt} \tag{14}$$

For $\hat{\beta}^R_{bt}$ to be a valid instrument, it must satisfy the exclusion restriction: $E\left(\hat{\beta}^R_{bt}\tilde{\xi}^D_{fbt}\right) = 0$. Sufficient conditions for this restriction to be true are:

$$E\left(\tilde{\Omega}_{bt}\tilde{\xi}^D_{fbt}\right) = 0 \tag{15}$$

$$E\left(\left(Q_{bt}\tilde{\Xi}^R_{fbt}\right)\tilde{\xi}^D_{fbt}\right) = 0 \tag{16}$$

To make sure that equation (16) holds, we (i) randomly divide firms into 10 groups indexed by j, (ii) estimate equation (11) while leaving out group j to obtain $\hat{\beta}^{R,-j}_{bt}$, (iii) use $\hat{\beta}^{R,-j}_{bt}$ as the instrument for firms of group j. Since ξ^D_{fbt} and ξ^S_{fbt} are independent across firms by assumption, equation (16) is verified. Therefore, the identification assumption boils down to equation (15): for a given interest rate, firms should not prefer banks that are doing better or worse.

We present the results in table 4. In columns (1) and (2), we run the regression in log.-changes. In column (1), we use $\hat{\beta}^R_{bt}$ as the IV, while in column (2) we use the $\hat{\beta}^{R,-j}_{bt}$ described in the previous paragraph. In both cases, we find a γ around 7. γ is the elasticity of substitution across banks, with respect to the gross interest rate, which is approximately the same as the semi-elasticity with respects to the net interest rate. In concrete terms, it means that if bank A increases its interest rate by 100 basis points relative to bank B, a firm that has a relationship with both of these banks decreases its borrowing by 7% from bank A relative to bank B. Finally, note that the F-statistic of the first stage is spectacular. Being a fixed effect, β^R_{bt} absorbs a substantial part of the variation in $\tilde{R}_{fbt}$, thus yielding a powerful first stage, and making weak instrument concerns irrelevant.

As we have alluded to in section 2.3, applying logarithms to equation (4) is problematic if there are a lot of new or disappearing relationships. To explicitly take those zeros into account, we estimate that equation in level. Specifically the econometric model is:

$$L_{fbt} = \xi^D_{fbt} \times \exp\left(-\gamma \log R_{fbt} + \theta^D_{ft}\right) \tag{17}$$

We estimate this with Poisson pseudo-maximum-likelihood (PPML). This method has been a common device

Table 4: Substitution across banks

	(1) Log.	(2) Log.	(3) PPML	(4) PPML
γ	7.096***	7.075***	12.356***	22.334***
	(1.817)	(1.872)	(3.854)	(4.873)
# obs.	6,750,296	6,748,463	7,765,469	781,563
# firms	649,904	649,783	374,426	37,737
# banks	2,426	2,343	1,976	1,805
# quarters	9	9	9	9
F-stat.	12,248	11,261	30,820	3,940
Split	N	Y	N	Y

Note: estimates for γ, the elasticity of substitution across banks. In columns (1–2), we estimate γ in log.-changes, following equation (10). In column (1), we use the simple fixed effect $\hat{\beta}^R_{bt}$ as IV. In column (2), we use $\beta^{R,-j}_{bt}$, which is estimated by splitting firms into groups. In columns (3–4), we estimate γ by Poisson pseudo-maximum-likelihood (PPML), following equation (17). Standard errors are two-way clustered at the firm and bank levels in columns (1–3). They are block-bootstrapped at the firm level in column (4). F-stat. is the Kleibergen-Paap rk statistic of the first stage in columns (1–2), and the F-statistic returned by the prediction step in columns (3–4). See section 4.1.3 for details.

in the trade literature since Santos Silva and Tenreyro (2006).[8] Our setting, however, presents two difficulties. First, the interest rate is potentially correlated with ξ^D_{fbt}. Second, we generally do not observe the interest rate when $L_{fbt} = 0$. Since the whole point of this exercise is to take the zeros into account, not having those observations would defeat the purpose. To work around these problems, we first estimate the level-version of equation (11):

$$\log R_{fbt} = c^R_t + \alpha^R_{ft} + \beta^R_{bt} + \epsilon^R_{fbt}, \qquad E\epsilon^R_{fbt} = 0 \tag{18}$$

We then predict $\log R_{fbt}$ based on this regression and estimate equation (17) with the prediction on the right-hand side:

$$L_{fbt} = \xi^D_{fbt} \times \exp\left(-\gamma \widehat{\log R_{fbt}} + \theta^D_{ft}\right), \qquad E\left(\xi^D_{fbt} \mid \widehat{\log R_{fbt}}\right) = 1 \tag{19}$$

Of course, if the second step were linear — and if we observed all R_{fbt} —. this procedure would be equivalent to two-stage least squares where $\hat{\beta}^R_{bt}$ is the IV. The results are reported in column (3) of table 4. $\hat{\gamma}$ is higher than with the log.-change specification, but in the same ballpark.

The procedure that we have just described raises one last issue: the standard errors of column (3) do not account for the fact that the right-hand side variable, $\widehat{\log R_{fbt}}$, is estimated. To compute accurate standard errors, we block-bootstrap at the firm level. Since estimating equation (19) is computationally costly, estimating on the whole set of observations of each bootstrap sample is infeasible. So, for each

[8] We implement this with the `ppmlhdfe` command in Stata (Correia et al., 2020). For the linear regressions, we use `reghdfe` (Correia, 2016).

bootstrap sample, we estimate equation (18) on 90% of the firms, predict $\log R_{fbt}$ for the remaining 1%, and estimate equation (19) on those. Of course, this procedure is not very efficient, but it is the best we can do within computational constraints. The results, column (4), are similar to those of column (3).

4.2 Other Parameters

4.2.1 φ: Elasticity of Total Credit

The identification of φ follows the same idea as that of γ. Formally, we run a regression that just follows the logarithmic version of equation (2):

$$\tilde{L}^D_{ft} = -\varphi \tilde{R}_{ft} + \tilde{\Omega}^D_{ft} \tag{20}$$

$\tilde{R}_{ft}$ is instrumented with a weighted sum of $\hat{\beta}^R_{bt}$:

$$\sum_b \frac{L_{fb,t-1}}{\sum_{b'} L_{fb',t-1}} \hat{\beta}^R_{bt}$$

Like before, the exclusion restriction boils down to two equations:

$$E\left(\tilde{\Omega}_{bt}\tilde{\Omega}_{ft}\right) = 0 \tag{21}$$

$$E\left(\left(Q_{bt}\tilde{\Xi}^R_{fbt}\right)\tilde{\Omega}^D_{ft}\right) = 0 \tag{22}$$

Equation (22) can be satisfied by using the leave out fixed effects, $\hat{\beta}^{R,-j}_{bt}$. Equation (21) is the real assumption: firm-specific demand shocks and bank-specific supply shocks must be uncorrelated.

Since L^D_{ft} and R_{ft} are indices, they need to be constructed with a value for γ. We try the different estimates of table 4, and also experiment with the linear case ($\gamma \to \infty$).[9]

We present the results in panel A of table 5: φ is between 1 and 2. Columns (1), (3) and (5) feature country-industry-time fixed effects; columns (2), (4) and (6) country-industry-size-time ones. All columns also feature firm fixed effects. Along columns (1–2), (3–4) and (5–6), we vary the value of $\hat{\gamma}$. That value matters little to $\hat{\varphi}$. Indeed, one can show that, up to a first-order approximation, the value of γ does not matter to the values of changes in L^D_{ft}.

In panel B, we estimate equation (2) in level, adapting the procedure described earlier for γ. We first estimate equation (18), predict the firm-level interest rate index, $\log R_{ft}$, by regressing it on a weighted sum

[9] With $\gamma \to \infty$, one can show that $L^D_{ft} \to \sum_b L^D_{fbt}$ and $R_{ft} \to \sum_b (L^D_{fbt}/L^D_{ft}) R_{fbt}$.

Table 5: Elasticity of total credit demand

	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Log. changes			
φ	1.838***	1.621***	1.838***	1.622***	1.837***	1.621***
	(0.064)	(0.064)	(0.064)	(0.064)	(0.064)	(0.064)
# obs.	20,598,996	20,598,215	20,598,996	20,598,215	20,598,996	20,598,215
# firms	3,296,185	3,296,102	3,296,185	3,296,102	3,296,185	3,296,102
# quarters	9	9	9	9	9	9
F-stat.	818	649	818	649	819	650
			Panel B: PPML			
φ	2.475**	2.321**	2.458**	2.325**	2.479**	2.346**
	(1.093)	(0.979)	(1.095)	(0.978)	(1.097)	(0.978)
# obs.	15,961,253	15,960,445	15,961,253	15,960,445	15,961,253	15,960,445
# firms	2,316,257	2,316,175	2,316,257	2,316,175	2,316,257	2,316,175
# quarters	9	9	9	9	9	9
			Parameters			
$\hat{\gamma}$	7.075	7.075	12.356	12.356	∞	∞
			Fixed effects			
Firm	Y	Y	Y	Y	Y	Y
Cntry-indstry-tm	Y	N	Y	N	Y	N
Cntry-indstry-sz-tm	N	Y	N	Y	N	Y

Note: estimates for φ, the elasticity of total credit demand. In panel A, we estimate φ in log.-changes, following equation (20). In panel B, we estimate it by Poisson pseudo maximum likelihood (PPML), following equation (23). All regressions feature firm fixed effects, and country-industry-time or country-industry-size-time fixed effects. Standard errors are clustered at the firm level. F-stat. is the Kleibergen-Paap rk statistic of the first stage in panel A, and the F-statistic of the prediction step in panel B. See section 4.2.1 for details.

of the estimated bank-time fixed effects, and finally estimate:

$$L_{ft} = \Omega^D_{ft} \times \exp\left(-\gamma \widehat{\log R_{ft}} + \theta^D_f + \iota^D_{st}\right), \qquad E\left(\Omega^D_{ft} \mid \widehat{\log R_{ft}}\right) = 1 \tag{23}$$

where θ^D_f and ι^D_{st} are firm and country-industry-time fixed effects. The estimated φ tends to be higher than when estimated with log.-changes (2–2.5).

4.2.2 Supply curve: μ and χ^{-1}

Once we've estimated γ, we can use equation (1) to recover estimates of the demand taste shocks, which we denote $\hat{\xi}^D_{fbt}$. We will use these structural shocks to identify the demand curve.

Let us go to a linear version of the supply curve:[10]

$$\tilde{R}_{fbt} = \mu \Delta s_{fbt} + \tilde{\pi}_{ft} - \tilde{\Omega}^S_{bt} + \chi^{-1} \tilde{L}^S_{bt} - \tilde{\xi}^S_{fbt} \tag{24}$$

Like with γ and φ, we face the identification problem that Δs_{fbt} and $\tilde{L}^S_{bt}$ are correlated with $\tilde{\Omega}^S_{bt}$ and $\tilde{\xi}^S_{fbt}$

[10]Remember that a capital delta, Δ, denotes time differentiation. We use the version of the supply curve where the market share is linearized in level so that we can keep the zeros.

in equilibrium. Notice that $\hat{\xi}^D_{fbt}$ in fact provides two instruments. The first one, $\Delta\hat{\xi}^D_{fbt}$, which moves at the firm-bank-time level can be the IV for the market share; the second one, $\sum_f w_{ft}\Delta\hat{\xi}^D_{fbt}$, which moves at the bank-time level can be the IV for $\tilde{L}_{fbt}$. Finally, we estimate:

$$\tilde{R}_{fbt} = \mu\Delta\tilde{s}_{fbt} + \chi^{-1}\tilde{L}_{fbt} + \theta^S_{ft} + e^S_{fbt} \tag{25}$$

where θ^S_{ft} is a firm-time fixed effect and e^S_{fbt} is the model equivalent of $-(\tilde{\Omega}^S_{bt} + \tilde{\xi}^S_{fbt})$. The identification assumption boils down to:

$$E\left((\Delta\xi^D_{fbt})\tilde{\Omega}^S_{bt}\right) = 0 \tag{26}$$

$$E\left((\Delta\xi^D_{f'bt})\tilde{\xi}^S_{bt}\right) = 0 \tag{27}$$

Note that the assumption in equation (26) has already been made in section 4.1.3.[11] The new assumption is equation (27): that demand and supply shocks must be orthogonal.

We find μ and χ^{-1} to be very close to 0. Neither result is surprising. μ is the semi-elasticity of the markup: $\mu = 1$, for instance, would imply that the interest rate increases by 100 basis points when the market share increases by 1%. In reality, μ seems to be very close to 0. As for χ^{-1}, table 6 is not so much an estimate as a sanity check. Indeed, in section 4.1.3, we estimated γ, hence $\hat{\xi}^D_{fbt}$, under the assumption that $\chi^{-1}/(1+\mu'')$ was approximately 0 — that assumption was the result of the estimation conducted in section 4.1.2. Hence, it should be that we find $\chi^{-1} \approx 0$ once we use shocks constructed under that assumption.

Table 6: Supply curve

	(1)	(2)	(3)
μ	-0.001*	0.001	-0.003***
	(0.001)	(0.001)	(0.000)
χ^{-1}	0.001	0.002	0.001
	(0.001)	(0.001)	(0.000)
# obs.	6,751,756	6,751,756	6,751,756
# firms	649,985	649,985	649,985
# banks	2,553	2,553	2,553
# quarters	9	9	9
F-stat.	124	125	206
$\hat{\gamma}$	7.075	12.356	∞

Note: estimates for μ, the semi-elasticity of the markup, and χ^{-1}, the inverse elasticity of supply, following equation (25). Standard errors are two-way clustered at the firm and bank levels. F-stat. is the Kleibergen-Paap rk statistic of the first stage. See section 25 for details.

[11] To be precise, in equation (26), it is the level change of ξ^D_{fbt} that appears while it is its log.-change in equation (15).

5 Macro Estimation

5.1 Aggregation

In principle, we should be able to recover the firm- and bank-specific shocks by applying the formulas:

$$\tilde{\Omega}^D_{ft} = \tilde{L}^D_{ft} + \varphi \tilde{R}_{ft}$$
$$\tilde{\Omega}^S_{bt} = \chi^{-1} \tilde{L}^S_{bt} + \tilde{\pi}_{bt} - \tilde{R}_{bt}$$

We could then define the aggregate demand and supply shocks as some average over firms and banks:

$$\bar{\Omega}^D_t = \sum_f w_{f,t-1} \tilde{\Omega}_{ft}$$
$$\bar{\Omega}^S_t = \sum_b w_{b,t-1} \tilde{\Omega}_{bt}$$

Unfortunately, incomplete coverage prevents us from doing so: we do not observe the interest rate and the probability of default of every firm, and we do not observe every loan. So that the decomposition of aggregate fluctuations wouldn't be exact.

Fortunately, we show in the appendix (section A.7) that, up to a first-order approximation:

$$\bar{L}^D_t = -\varphi \bar{R}_t + \bar{\Omega}^D_t \tag{28}$$
$$\bar{R}_t = \chi^{-1} \tilde{L}^S_t + \bar{PD}_t - \bar{\Omega}^S_t \tag{29}$$

where $\bar{L}_t$ is the growth rate of loans outstanding, $\bar{R}_t$ is the change in the interest rate, $\bar{\Omega}^D_t$ and $\bar{\Omega}^S_t$ are aggregate demand and supply shocks. It may seem surprising that neither γ nor μ appears in these equations. γ does not have a first-order effect on the firm's total borrowing. In that respect, the fact that $\hat{\gamma}$ didn't matter to our estimate of φ (section 4.2.1) is reassuring about the quality of the approximation. μ, the semi-elasticity of the markup, embodies banks' markup power with respect to a firm. Within a firm, a bank's gain in market share is another bank's loss. Hence, those wash out in the aggregate.

To obtain time series of $\bar{\Omega}^D_t$ and $\bar{\Omega}^S_t$, we could take a rearranged version of equations (28) and (29), with our estimates of φ and χ^{-1} as slopes of the demand and supply curve:

$$\bar{\Omega}^D_t = \bar{L}^D_t + \hat{\varphi} \bar{R}_t$$
$$\bar{\Omega}^S_t = \hat{\chi}^{-1} \tilde{L}^S_t + \bar{PD}_t - \bar{R}_t$$

Plugging aggregate data on loans, interest rates and probability of default on the right-hand side, we would recover aggregate demand and supply shocks. We could then compute first- and second-order moments of the sample distribution of those shocks. We do a slightly more complicated version of this simple exercise by estimating the model in a Bayesian fashion. Compared to the naive approach just described, a formal Bayesian estimation allows us to construct credible intervals on the first- and second-order moments of the distribution of the aggregate shocks.

5.2 Estimation

We assume that the vector $(\bar{\Omega}_t^D, \bar{\Omega}_t^S, \bar{PD}_t)$ is normally distributed:

$$\begin{pmatrix} \bar{\Omega}_t^D \\ \bar{\Omega}_t^S \\ \bar{PD}_{ft} \end{pmatrix} \sim \mathcal{N}\left(\bar{\Omega}, \Sigma\right) \tag{30}$$

where: $\bar{\Omega} = (\bar{\Omega}^D, \bar{\Omega}^S, 0)$.

Prior choices are summarized in table A.1. The most important choices are those for the two slopes, φ and χ^{-1}. We take normal distribution with mean and standard deviation equal to the point estimates and standard deviations of section 4.[12] We truncate those distributions below 0. The other prior distributions are flat. We describe them in details in the appendix (section C).

We conduct the Bayesian computations in Stan, through its R interface, RStan (Stan Development Team, 2020, 2021). Stan relies on Hamiltonian Monte Carlo sampling (Betancourt, 2017). We report some moments of the posterior distribution in the appendix (table A.2).

5.3 Decomposition

Before we jump to the results, let us clarify how the identification works. Thanks to the estimation of section 4, which is based on microeconomic data, we are able to place informative prior distributions on the slope of the aggregate demand and supply curves, φ and χ^{-1}. Due to the simplicity of equations (28–29), the model can be represented into a simple aggregate supply and demand diagram (figure 3). Since we found $\chi^{-1} \approx 0$, we draw a flat supply curve. On the other hand, since φ is low, the demand curve is steep. A positive supply shock corresponds to a downward move of the supply curve, the interest rate goes down. As the economy moves along the demand curve, loans go up, but they do so only slightly as the demand curve is steep. A decline in risk would act in the same direction by moving the supply curve down. On the

[12] We take those of panel A, column (1) in table 5, and column (1) in table 6.

other hand, a positive demand shock will not affect the interest rate, but increases loans one-for-one. Thus, figure 3 announces our quantitative results: fluctuations in the interest rate should be dominated by supply and risk, and fluctuations in loans should be dominated by demand with perhaps minor supply influences.

Figure 3: AD-AS diagram

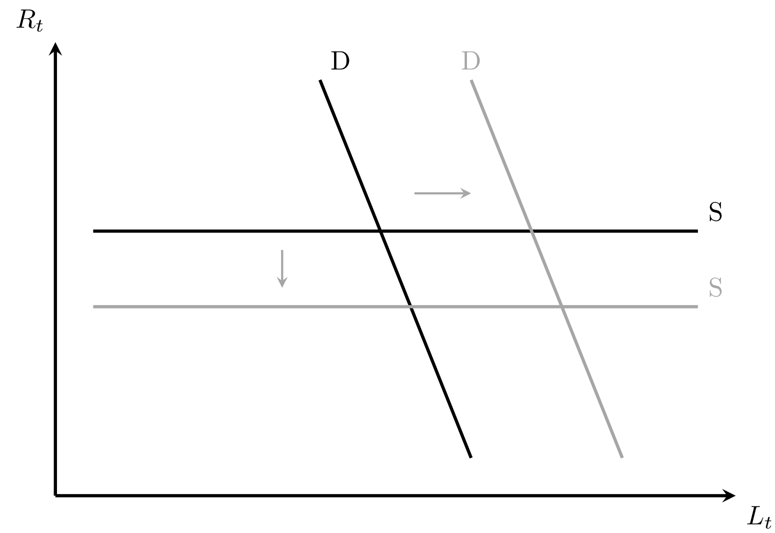

We can use equations (28) and (29) to solve for the equillibrium values of $\bar{L}_t$ and $\bar{R}_t$:

$$\bar{L}_t = \frac{1}{1+\varphi\chi^{-1}}\bar{\Omega}_t^D + \frac{\varphi}{1+\varphi\chi^{-1}}\bar{\Omega}_t^S - \frac{\varphi}{1+\varphi\chi^{-1}}\bar{PD}_t \tag{31}$$

$$\bar{R}_t = \underbrace{\frac{\chi^{-1}}{1+\varphi\chi^{-1}}\bar{\Omega}_t^D}_{\text{demand}} - \underbrace{\frac{1}{1+\varphi\chi^{-1}}\bar{\Omega}_t^S}_{\text{supply}} + \underbrace{\frac{1}{1+\varphi\chi^{-1}}\bar{PD}_t}_{\text{risk}} \tag{32}$$

Equations (31) and (32) imply that aggregate fluctuations in loans and interest rates can be fully decomposed into three terms: demand, supply and risk.

In figure 4, we show the decomposition in practice. On panel A, we decompose the growth rate of loans outstanding. As expected from the discussion of figure 3, demand dominates. In particular, there is a surge in demand in the first two quarters of 2020. The early part of the sample shows some supply influences particularly in 2014, where they partly counteract demand headwinds. The influence of supply is much more visible in the changes in lending rates (panel B). Until 2018, there were persistently positive supply shocks, which translate into a fall in lending rates. The fall in the probability of default also contributed significantly to the fall in lending rates. One, however, should keep in mind that due to the proximity of the zero lower bound, the movements in lending rates are small. The y-axis is expressed in percentage points, so that the biggest movement (2014Q4) was a drop of barely more than 10 basis points.

Just like we were able to decompose the growth rate of loans, and the change in lending rates, we can

Figure 4: Decomposition

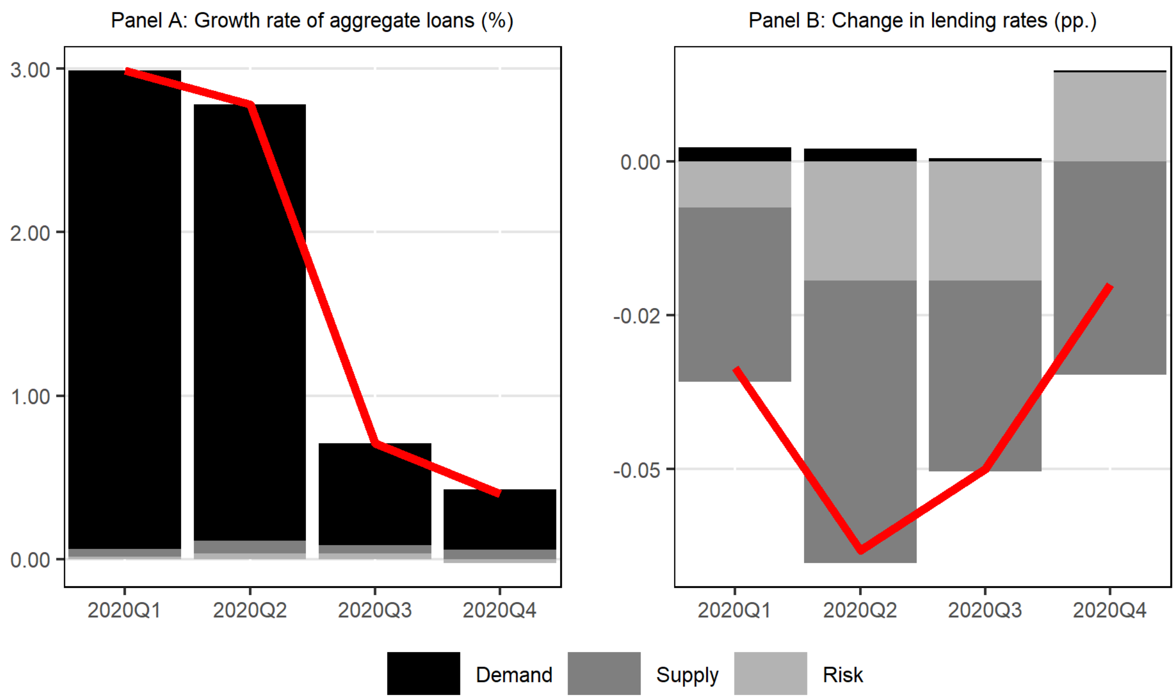

Note: decomposition of the growth rate of aggregate loans (panel A) and change in lending rates (panel B) in the euro area. The red line is the data. The black, dark and light gray bars are the posterior mean for the contributions of demand, supply and risk. They sum to the red line by construction. See equations (31–32) for the formula.

decompose their variances by applying the variance operator to both sides of equations (31) and (32). We show the results of this decomposition in table 7. It formalizes what was already apparent on figure 4. Demand accounts for most of the variance in loan growth. Supply and risk respectively account for 70 and 29% (at the posterior mean) of the variance of changes in lending rates. The covariance terms are small and statistically indistinguishable from 0.

We reemphasize that the macro results were already baked into the micro estimation. That demand would matter little to lending rates was a forgone conclusion as soon as we estimated $\chi^{-1} \approx 0$. We see this as a good thing. The identification of the aggregate exercise is tightly disciplined by parameters estimated in micro data. Those are identified with more plausible assumptions than we could dream about in aggregate data.

6 Conclusion

This paper presents a new methodology to identify aggregate demand and supply shocks in the bank loan market. We build a model of the bank loan market identifying structural elasticities that can be estimated

Table 7: Variance decomposition

	Mean	St. dev.	5%	95%
Loans outstanding				
Demand	0.993	0.045	0.941	1.041
Supply	0.000	0.003	0.000	0.002
Risk	0.004	0.035	0.000	0.008
Cov. D/S	0.001	0.016	-0.019	0.023
Cov. D/R	0.001	0.056	-0.043	0.047
Cov. S/R	-0.000	0.001	-0.001	0.001
Lending rates				
Demand	0.013	0.049	0.000	0.059
Supply	0.328	0.373	0.011	1.056
Risk	0.786	0.420	0.140	1.417
Cov. D/S	-0.003	0.063	-0.069	0.058
Cov. D/R	-0.021	0.107	-0.161	0.066
Cov. S/R	-0.103	0.536	-0.970	0.335

Note: posterior distribution of the variance decomposition of the growth rate of loans outstanding and the change in lending rates.

based on granular data. These include the elasticity of credit demand and supply, as well as firms' elasticity of substitution across banks and the elasticity of the mark-up applied by banks to their market share and the elasticity of supply. We then show that these can be used to inform a macroeconomic exercise that allows for a decomposition of aggregate fluctuations bank lending into the contribution of demand and supply. Importantly, while the methodology requires both granular and aggregate data, it is parsimonious in the scope of variables used, which are limited to lending volumes, lending rates and firm credit risk.

We apply this novel methodology to the developments in the euro area following the COVID-19 Pandemic. This exercise is based on the AnaCredit dataset, a new credit register covering all euro area countries, and on aggregate time series for the euro area. The main result is that fluctuations in lending volumes in the wake of the Pandemic were largely explained by demand shocks, whereas developments in lending rates were mostly driven by supply shocks and borrower risk.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Alfaro, L., García-Santana, M., and Moral-Benito, E. (2021). On the direct and indirect real effects of credit supply shocks. *Journal of Financial Economics*, 139(3):895–921.

Altavilla, C., Boucinha, M., Holton, S., and Ongena, S. (2021). Credit supply and demand in unconventional times. *Journal of Money, Credit and Banking*. Forthcoming.

Altavilla, C., Boucinha, M., Peydró, J.-L., and Smets, F. (2020a). Banking supervision, monetary policy and risk-taking: big data evidence from 15 credit registers. ECB Working Paper No. 2349.

Altavilla, C., Darracq Pariès, M., and Nicoletti, G. (2019). Loan supply, credit markets and the euro area financial crisis. *Journal of Banking and Finance*, 109:1–19.

Altavilla, C., Laeven, L., and Peydró, J.-L. (2020b). Monetary and macroprudential policy complementarities: evidence from European credit registers. CEPR Discussion Paper No. 15539.

Amado, J. and Nagengast, A. J. (2016). The effect of bank shocks on firm-level and aggregate investment. ECB Working Paper No. 1914.

Amiti, M. and Weinstein, D. E. (2018). How much do idiosyncratic bank shocks affect investment? Evidence from matched bank-firm loan data. *Journal of Political Economy*, 126(2):525–585.

Amity, M., McGuire, P., and Weinstein, D. E. (2017). Supply- and demand-side factors in global banking. NBER Working Paper No. 23536.

Andrews, M. J., Gill, L., Schank, T., and Upward, R. (2008). High wage workers and low wage firms: Negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 171(3):673–697.

Atkeson, A. and Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, 98(5):1998–2031.

Bassett, W. F., Chosak, M. B., Driscoll, J. C., and Zakrajšek, E. (2014). Changes in bank lending standards and the macroeconomy. *Journal of Monetary Economics*, 62:23–40.

Berg, T., Saunders, A., Schäfer, L., and Steffen, S. (2021). Brexit and the contraction of syndicated lending. *Journal of Financial Economics*, 141(1):66–82.

Berton, F., Mocetti, S., Presbitori, A. F., and Richiardi, M. (2018). Banks, firms, and jobs. *Review of Financial Studies*, 31(6):2113–2156.

Betancourt, M. (2017). A conceptual introduction to Hamiltonian Monte Carlo. arXiv e-print.

Chang, B., Gomez, M., and Hong, H. (2021). Sorting out the real effects of credit supply. NBER Working Paper No. 28842.

Christianio, L., Motto, R., and Rostagno, M. (2010). Financial factors in economic fluctuations. ECB Working Paper No. 1192.

Correia, S. (2016). A feasible estimator for linear models with multi-way fixed effects. Working paper.

Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *Stata Journal*, 20(1):95–115.

Degryse, H., De Jonghe, O., Jakovljevic, S., Mulier, K., and Schepens, G. (2019). Identifying credit supply shocks with bank-firm data: Methods and applications. *Journal of Financial Intermediation*, 40:200–225.

Del Giovane, P., Eramo, G., and Nobili, A. (2011). Disentangling demand and supply in credit developments: A survey-based analysis for italy. *Journal of Banking and Finance*, 35(10):2719–2732.

Eaton, J. and Kortum, S. (2002). Technology, geography, and trade. *Econometrica*, 70(5):1741–1779.

Eckmeier, S. and Ng, T. (2015). How do US credit supply shocks propagate internationally? A GVAR approach. *European Economic Review*, 74:128–145.

European Central Bank (2019). *AnaCredit Reporting Manual*. `https://www.ecb.europa.eu/stats/money_credit_banking/anacredit/html/index.en.html`.

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis*. Texts in Statistical Science. CRC Press, Boca Raton, third edition.

Gerali, A., Neri, S., Sessa, L., and Signoretti, F. M. (2010). Credit and banking in a DSGE model of the euro area. *Journal of Money, Credit and Banking*, 42:107–141.

Giannone, D., Lenza, M., and Reichlin, L. (2019). Money, credit, monetary policy, and the business cycle in the euro area: What has changed since the crisis? *International Journal of Central Banking*, 15:137–173.

Greenstone, M., Mas, A., and Nguyen, H.-L. (2020). Do credit market shocks affect the real economy? Quasi-experimental evidence from the great recession and "normal" economic times. *American Economic Journal: Economic Policy*, 12(1):200–225.

Herreño, J. (2021). The aggregate effects of bank lending cuts. Working paper.

Kashyap, A. K. and Stein, J. C. (2000). What do a million observations on banks say about the transmission of monetary policy? *American Economic Review*, 90(3):407–428.

Khwaja, A. I. and Mian, A. (2008). Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review*, 98(4):1413–1442.

Kline, P., Saggio, R., and Sølvsten, M. (2020). Leave-out estimation of variance components. *Econometrica*, 88(5):1859–1898.

Lown, C. and Morgan, D. P. (2006). The credit cycle and the business cycle: New findings using the loan officer opinion survey. *Journal of Money, Credit and Banking*, 38:1575–1597.

Moccero, D. N., Darracq Pariès, M., and Maurin, L. (2014). Financial conditions index and identification of credit supply shocks for the euro area. *International Finance*, 17(3):297–321.

Paravisini, D., Rappoport, V., and Schnabl, P. (2020). Specialization in bank lending: Evidence from exporting firms. Working paper.

Santos Silva, J. M. C. and Tenreyro, S. (2006). The log of gravity. *Review of Economics and Statistics*, 88(4):641–658.

Stan Development Team (2020). RStan: the R interface to Stan. R package version 2.21.2.

Stan Development Team (2021). *Stan Modeling Language Users Guide and Reference Manual*. Version 2.27.

APPENDIX

A Theoretical Derivations

A.1 Micro-Foundations of the Demand Equations

A.1.1 Equation (1)

These derivations follow Herreño (2021). Paravisini et al. (2020) offer a similar setup. They have been standard in the trade literature since Eaton and Kortum (2002).

The firm's production function is:

$$Y_{ft} = A_{ft} N_{ft}^{1-\psi}$$

with:

$$N_{ft} = \left(\int_0^1 \left(N_{ft}(\eta)\right)^{\frac{\sigma-1}{\sigma}} d\eta \right)^{\frac{\sigma}{\sigma-1}}$$

η indexes a continuum of tasks, performed by labor $N_{ft}(\eta)$. σ is the elasticity of substitution between those tasks. ψ is the curvature of the production function.

The total cost of financing task η is:

$$TC_{ft}(\eta) = R_{ft}(\eta) W_t N_{ft}(\eta)$$

where $R_{ft}(\eta)$ is the effective cost of funds for task η, and W_t is the wage. Banks finance task η with comparative advantage: the effective cost of financing task η is $R_{fbt}/\nu_{fbt}(\eta)$ where R_{fbt} is the actual interest rate offered by bank b and $\nu_{fbt}(\eta)$ is a productivity shock that determines the comparative advantage. As a result, the firm chooses the bank that minimizes the total cost of financing task η:

$$b = \underset{b'}{\operatorname{argmin}} \frac{R_{b'ft}}{\nu_{fb't}(\eta)}$$

Productivity shocks $\nu_{fbt}(\eta)$ (henceforth, we omit indexation by η) are drawn from a Fréchet distribution with cumulative distribution function (c.d.f.):

$$F_{fbt}(\nu_{fbt}) = \exp\left(-\xi_{fbt}^{D} \nu_{fbt}^{-\gamma'} \right)$$

The probability that R_{fbt}/ν_{fbt} is smaller than some $x \in (0, +\infty)$ is:

$$P\left(\frac{R_{fbt}}{\nu_{fbt}} < x\right) = P\left(\frac{R_{fbt}}{x} < \nu_{fbt}\right) = 1 - F_{fbt}\left(\frac{R_{fbt}}{x}\right) = 1 - \exp\left(-\xi^D_{fbt} R_{fbt}^{-\gamma'} x^{\gamma'}\right)$$

So the probability that $\min_b R_{fbt}/\nu_{fbt}$ is smaller than some x is:

$$\begin{aligned}
F_{ft}(x) \equiv P\left(\min_b \frac{R_{fbt}}{\nu_{fbt}} < x\right) &= 1 - P\left(\min_b \frac{R_{fbt}}{\nu_{fbt}} \geq x\right) \\
&= 1 - \Pi_b P\left(\frac{R_{fbt}}{\nu_{fbt}} \geq x\right) = 1 - \Pi_b\left(1 - P\left(\frac{R_{fbt}}{\nu_{fbt}} < x\right)\right) \\
&= 1 - \Pi_b \exp\left(-\xi^D_{fbt} R_{fbt}^{-\gamma'} x^{\gamma'}\right) = 1 - \exp\left(-\left(\sum_b \xi^D_{fbt} R_{fbt}^{-\gamma'}\right) x^{\gamma'}\right)
\end{aligned}$$

Now, given R_{fbt}/ν_{fbt}, the probability that bank b is chosen for a given task is given by:

$$P\left(\frac{R_{fbt}}{\nu_{fbt}} < \min_{b' \neq b} \frac{R_{fb't}}{\nu_{fb't}}\right) = 1 - P\left(\frac{R_{fbt}}{\nu_{fbt}} > \min_{b' \neq b} \frac{R_{fb't}}{\nu_{fb't}}\right) = \exp\left(-\left(\sum_{b' \neq b} \xi^D_{fbt} R_{fbt}^{-\gamma'}\right)\left(\frac{R_{fbt}}{\nu_{fbt}}\right)^{\gamma'}\right)$$

We can obtain the share of tasks for which the bank is chosen by integrating over $x = R_{fbt}/\ nu_{fbt}$:

$$\begin{aligned}
s_{fbt} &\equiv \int_0^{+\infty} P\left(x < \min_{b' \neq b} \frac{R_{fb't}}{\nu_{fb't}}\right) dF_{fbt}(x) \\
&= \int_0^{+\infty} \exp\left(-\left(\sum_{b' \neq b} \xi^D_{fbt} R_{fbt}^{-\gamma'}\right) x^{\gamma'}\right) \times \gamma' \xi^D_{fbt} R_{fbt}^{-\gamma'} x^{\gamma'-1} \exp\left(-\xi^D_{fbt} R_{fbt}^{-\gamma'} x^{\gamma'}\right) dx \\
&= \frac{\xi^D_{fbt} R_{fbt}^{-\gamma'}}{\sum_b \xi^D_{fbt} R_{fbt}^{-\gamma'}} \int_0^{+\infty} \gamma' \left(\sum_b \xi^D_{fbt} R_{fbt}^{-\gamma'}\right) x^{\gamma'-1} \exp\left(-\left(\sum_b \xi^D_{fbt} R_{fbt}^{-\gamma'}\right) x^{\gamma'}\right) dx \\
&= \frac{\xi^D_{fbt} R_{fbt}^{-\gamma'}}{\sum_b \xi^D_{fbt} R_{fbt}^{-\gamma'}} \left[-\exp\left(-\left(\sum_b \xi^D_{fbt} R_{fbt}^{-\gamma'}\right) x^{\gamma'}\right)\right]_{x=0}^{+\infty} = \frac{\xi^D_{fbt} R_{fbt}^{-\gamma'}}{\sum_b \xi^D_{fbt} R_{fbt}^{-\gamma'}} \\
&= \xi^D_{fbt}\left(\frac{R_{fbt}}{R_{ft}}\right)^{-\gamma'} \qquad \text{(A.1)}
\end{aligned}$$

where:

$$R_{ft} \equiv \left(\sum_b \xi^D_{fbt} R_{fbt}^{-\gamma'}\right)^{-\frac{1}{\gamma'}}$$

By the same logic, banks charge the same average effective interest rate in equilibrium. Indeed, the density of charging x and being chosen is:

$$P\left(x < \min_{b' \neq b} \frac{R_{fb't}}{\nu_{fb't}}\right) dF_{fbt}(x)$$

So the c.d.f. of charging r and being chosen is:

$$\int_0^r P\left(x < \min_{b' \neq b} \frac{R_{fb't}}{\nu_{fb't}}\right) dF_{fbt}(x) \; = s_{fbt} F_{ft}(r)$$

To get to the conditional — conditional on being chosen — distribution of the effective interest rate charged by b, we divide by the probability of being chosen, s_{fbt}. So the conditional distribution of the effective interest rate is the same for all banks: $F_{ft}(r)$.

The cost index of one unit of labor is given by:

$$\begin{aligned}
C_{ft}^{1-\sigma} &= \int_0^\infty (R_{ft}(\eta) W_t)^{1-\sigma} \, d\eta \; = W_t^{1-\sigma} \int_0^1 x^{1-\sigma} \, dF_{ft}(x) \\
&= W_t^{1-\sigma} \int_0^\infty x^{1-\sigma} \left(\sum_b \xi_{fbt}^D R_{fbt}^{-\gamma'}\right) \gamma' x^{\gamma'-1} \exp\left(-\left(\sum_b \xi_{fbt}^D R_{fbt}^{-\gamma'}\right) x^{\gamma'}\right) dx \\
&= W_t^{1-\sigma} \left(\sum_b \xi_{fbt}^D R_{fbt}^{-\gamma'}\right)^{-\frac{1-\sigma}{\gamma'}} \int_0^\infty r^{\frac{1-\sigma}{\gamma'}} \exp(-r) \, dr \\
&= \Theta^{1-\sigma} W_t^{1-\sigma} \left(\sum_b \xi_{fbt}^D R_{fbt}^{-\gamma'}\right)^{-\frac{1-\sigma}{\gamma'}}
\end{aligned}$$

where:

$$\Theta \equiv \left(\gamma' \left(\frac{1-\sigma}{\gamma'} + 1\right)\right)^{\frac{1}{1-\sigma}}$$

To go from the second to the third line, we made the change of integration variable: $r = \left(\sum_b \xi_{fbt}^D R_{fbt}^{-\gamma'}\right) x^{\gamma'}$. $\Gamma(.)$ denotes the gamma function. So the cost index is:

$$C_{ft} = \Theta R_{ft} W_t \tag{A.2}$$

Now, to finance task η, the firms needs to borrow $L_{ft}(\eta) = W_t N_{ft}(\eta)/\nu_{fbt}(\eta)$, and repays:

$$R_{fbt} L_{ft}(\eta) = \frac{R_{fbt}}{\nu_{fbt}(\eta)} W_t N_{ft}(\eta) = R_{ft}(\eta) W_t N_{ft}(\eta)$$

So the demand for loans from bank b is:

$$
\begin{aligned}
R_{fbt}L_{fbt} &= \int_0^1 R_{ft}(\eta)W_tN_{ft}(\eta) \times \mathbb{1}\left[b = \underset{b'}{\operatorname{argmin}} \frac{R_{fb't}}{\nu_{fb't}(\eta)}\right] d\eta \\
&= \left(C_{ft}\right)^{\sigma} N_{ft} \int_0^1 \left(R_{ft}(\eta)W_t\right)^{1-\sigma} \times \mathbb{1}\left[b = \underset{b'}{\operatorname{argmin}} \frac{R_{fb't}}{\nu_{fb't}(\eta)}\right] d\eta \\
&= \left(C_{ft}\right)^{\sigma} N_{ft} s_{fbt} C_{ft}^{1-\sigma} = \Theta s_{fbt} R_{ft} W_t N_{ft}
\end{aligned}
$$

To go from the first to the second line, we have used the standard formula for CES demand: $N_{ft}(\eta) = (R_{ft}(\eta)W_t/C_{ft})^{-\sigma}N_{ft}$. To go from the second to the third line, we have used the previous results on the conditional distribution of $R_{fbt}/\nu_{fbt}(\eta)$. Therefore, we have:

$$
L_{fbt} = \Theta \xi_{fbt}^{D} \left(\frac{R_{fbt}}{R_{bt}}\right)^{-\gamma'-1} W_t N_{ft}
$$

Defining $\gamma \equiv \gamma' + 1$, we have:

$$
L_{fbt} = \xi_{fbt}^{D} \left(\frac{R_{fbt}}{R_{bt}}\right)^{-\gamma} L_{ft} \tag{A.3}
$$

$$
L_{ft} \equiv \Theta W_t N_{ft} \tag{A.4}
$$

We can easily check that:

$$
R_{ft} = \left(\sum_b \xi_{fbt}^{D} R_{fbt}^{1-\gamma}\right)^{\frac{1}{1-\gamma}} \tag{A.5}
$$

$$
L_{ft} = \left(\sum_b \left(\xi_{fbt}^{D}\right)^{\frac{1}{\gamma}} L_{fbt}^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}} \tag{A.6}
$$

This proves equation (1).

A.2 Equation (2)

The firm maximizes profits:

$$
P_t A_{ft} N_{ft}^{1-\psi} - C_{ft} N_{ft}
$$

It takes P_t and C_{ft} as given.[13] The first order condition is:

$$(1-\psi)P_t A_{ft} N_{ft}^{-\psi} = C_{ft}$$

Now, using equations (A.2) and (A.4), and rearranging:

$$L_{ft} = R^{-\frac{1}{\psi}} \left(\frac{(1-\psi)P_t A_{ft}}{\Theta^{1-\psi} W_t^{1-\psi}} \right)^{\frac{1}{\psi}}$$

This proves equation (2) with $\varphi \equiv 1/\psi$ and:

$$\Omega_{ft}^{D} \equiv \left(\frac{(1-\psi)P_t A_{ft}}{\Theta^{1-\psi} W_t^{1-\psi}} \right)^{\frac{1}{\psi}} \tag{A.7}$$

Thus, Ω_{ft}^{D} embeds idiosyncratic (A_{ft}) and aggregate ($P_t/W_t^{1-\psi}$) forces.

A.3 Demand with Default

To model default on the firm side, we can assume that productivity is 0 with some probability:

$$A_{ft} = \begin{cases} \bar{A}_{ft}, & \text{with probability } 1 - PD_{ft} \\ 0, & \text{with probability } PD_{ft} \end{cases}$$

The firm makes the decision to hire and borrow before it knows the realization of A_{ft}. If $A_{ft} = 0$ happens, the firm is unable to pay back its debt, defaults and makes 0 profits. The demand equations carry through with:

$$\Omega_{ft}^{D} = \left(\frac{(1-\psi)P_t \bar{A}_{ft}}{\Theta^{1-\psi} W_t^{1-\psi}} \right)^{\frac{1}{\psi}}$$

A.4 Derivation of Equations (3) and (5)

The bank maximizes:

$$\sum_f \xi_{fbt}^{S} \left(1 - PD_{ft}\right) R_{fbt} L_{fbt}^{S} - \frac{R_{bt}^{*}}{1+\chi^{-1}} \left(\sum_f L_{fbt}^{S} \right)^{1+\chi^{-1}}$$

[13] That P_t is taken as given could be relaxed. See Herreño (2021).

subject to $L^S_{fbt} = L^D_{fbt}$, where L^D_{fbt} is defined in equation (4). The first-order condition with respect to R_{fbt} is:

$$0 = \xi^S_{fbt}\left(1 - PD_{ft}\right)\left(L_{fbt} + R_{fbt}\frac{\partial L_{fbt}}{\partial R_{fbt}}\right) - R^*_{bt}\left(L^S_{bt}\right)^{\chi^{-1}}\frac{\partial L_{fbt}}{\partial R_{fbt}} \tag{A.8}$$

Now, using equation (4) and the definition of the interest rate index R_{ft}:

$$L_{fbt} = \xi^D_{fbt} R^{-\gamma}_{fbt} R^{\gamma-\varphi}_{ft} \Omega^D_{ft} = \xi^D_{fbt} R^{-\gamma}_{fbt}\left(\sum_b \xi^D_{fbt} R^{1-\gamma}_{ft}\right)^{\frac{\gamma-\varphi}{1-\gamma}} \Omega^D_{ft}$$

Therefore:

$$\begin{aligned}
\frac{\partial L_{fbt}}{\partial R_{fbt}} &= \left(-\gamma R^{-\gamma-1}_{fbt}\left(\sum_b \xi^D_{fbt} R^{1-\gamma}_{ft}\right)^{\frac{\gamma-\varphi}{1-\gamma}} + (\varphi-\gamma)\xi^D_{fbt} R^{-2\gamma}_{fbt}\left(\sum_b \xi^D_{fbt} R^{1-\gamma}_{ft}\right)^{\frac{\gamma-\varphi}{1-\gamma}-1}\right)\xi^D_{fbt}\Omega^D_{ft} \\
&= \left(-\gamma R^{-1}_{fbt} + (\varphi-\gamma) R^{-\gamma}_{fbt}\left(\sum_b \xi^D_{fbt} R^{-\gamma}_{ft}\right)^{-1}\right)\xi^D_{fbt} R^{-\gamma}_{fbt}\left(\sum_b \xi^D_{fbt} R^{1-\gamma}_{ft}\right)^{\frac{\gamma-\varphi}{1-\gamma}}\Omega^D_{ft} \\
&= -\left(\gamma + (\varphi-\gamma)\xi^D_{fbt}\frac{R^{1-\gamma}_{fbt}}{R^{1-\gamma}_{ft}}\right)\frac{L_{fbt}}{R_{fbt}} = -\left(\gamma(1-s_{fbt}) + \varphi s_{fbt}\right)\frac{L_{fbt}}{R_{fbt}} = -\zeta_{fbt}\frac{L_{fbt}}{R_{fbt}}
\end{aligned}$$

To go from the second to the third line, we used the definition of the interest rate index R_{ft} and factored out R^{-1}_{fbt}. Substituting the latter into equation (A.8) and dividing by L_{fbt}/R_{fbt}:

$$\xi^S_{fbt}\left(1 - PD_{ft}\right)\left(1 - \zeta_{fbt}\right)R_{fbt} + \zeta_{fbt} R^*_{bt}\left(\sum_f L^S_{fbt}\right)^{\chi^{-1}} = 0$$

Rearranging:

$$R_{fbt} = \frac{\tilde{\zeta}_{fbt}}{\tilde{\zeta}_{fbt} - 1}\frac{R^*_{bt}\left(L^S_{bt}\right)^{\chi^{-1}}}{\xi^S_{fbt}\left(1 - PD_{ft}\right)}$$

This is equation (5). Equation (3) follows with $\xi^S_{fbt} = 1$, $PD_{ft} = 0$ and $\chi^{-1} = 0$.

A.5 Linearization of the Supply Curve

We start from equation (5):

$$R_{fbt} = \mathcal{M}(s_{fbt})\frac{R^*_{bt}\left(L^S_{bt}\right)^{\chi^{-1}}}{\xi^S_{fbt}(1 - PD_{ft})} \tag{5}$$

Taking the log.-change from $t-1$ to t:

$$
\begin{aligned}
\tilde{R}_{fbt} &= \Delta \log\left(\mathcal{M}(s_{fbt})\right) + \tilde{R}^*_{bt} + \chi^{-1}\tilde{L}^S_{bt} - \mathrm{d}\log(1-PD_{ft}) - \xi^S_{fbt} \\
&\approx \mu\Delta\tilde{s}_{fbt} + \tilde{R}^*_{bt} + \chi^{-1}\tilde{L}^S_{bt} - \Delta\log(1-PD_{ft}) - \xi^S_{fbt} \\
&\approx \mu'\tilde{s}_{fbt} + \tilde{R}^*_{bt} + \chi^{-1}\tilde{L}^S_{bt} - \Delta\log(1-PD_{ft}) - \xi^S_{fbt}
\end{aligned}
$$

with: $\mu = \mathcal{M}'(\bar{s})/\mathcal{M}(\bar{s})$ and $\mu' = \bar{s}\mu$. A capital delta, Δ, denotes time differentiation. We obtain the second line by taking an approximation of $\mathcal{M}(s_{fbt})$ around a situation where $s_{fbt} = \bar{s}$. To go to the third line, we use: $\Delta s_{fbt} \approx \bar{s}\tilde{s}_{fbt}$. Now we use the fact that $s_{fbt} = \xi^D_{fbt}\left(R_{fbt}/R_{ft}\right)^{1-\gamma}$:

$$
\tilde{R}_{fbt} \approx \mu''\tilde{R}_{fbt} + \tilde{R}^*_{bt} + \chi^{-1}\tilde{L}^S_{bt} + \mu''\tilde{R}_{ft} - \mathrm{d}\log(1-PD_{ft}) + \mu'\tilde{\xi}^D_{fbt} - \xi^S_{fbt}
$$

with: $\mu'' = (\gamma-1)\mu'$. Rearranging:

$$
\tilde{R}_{fbt} \approx \frac{\mu''\tilde{R}_{ft} + \tilde{\pi}_{ft} - \tilde{\Omega}_{bt} + \chi^{-1}\tilde{L}^S_{bt} + \mu\tilde{\xi}^D_{fbt} - \xi^S_{fbt}}{1+\mu''} \tag{7}
$$

with: $\tilde{\pi}_{ft} = -\Delta\log(1-PD_{ft})$ and $\Omega_{bt} = 1/R^*_{bt}$.

A.6 Proof of Proposition 1

We start from equation (9):

$$
\tilde{L}_{fbt} = \underbrace{\tilde{\Omega}^D_{ft} - \underline{\gamma}\tilde{\pi}_{ft} + (\underline{\gamma}-\varphi)\tilde{R}_{ft}}_{\text{firm specific}} + \underbrace{\underline{\gamma}\left(\tilde{\Omega}^S_{bt} - \chi^{-1}\tilde{L}_{bt}\right)}_{\text{bank specific}} + \underbrace{(1-\underline{\gamma}\mu)\tilde{\xi}^D_{fbt} + \underline{\gamma}\tilde{\xi}^S_{fbt}}_{\text{error term}} \tag{9}
$$

We denote with superscript T the true fixed effects and error term:

$$
\begin{aligned}
\alpha^{LT}_{ft} &= \tilde{\Omega}^D_{ft} - \underline{\gamma}\tilde{\pi}_{ft} + (\underline{\gamma}-\varphi)\tilde{R}_{ft} \\
\beta^{LT}_{bt} &= \underline{\gamma}\left(\tilde{\Omega}^S_{bt} - \chi^{-1}\tilde{L}_{bt}\right) \\
\xi^L_{fbt} &= (1-\underline{\gamma}\mu)\tilde{\xi}^D_{fbt} + \underline{\gamma}\tilde{\xi}^S_{fbt}
\end{aligned}
$$

We consider each time period separately. For some t, let x_{fbt} be the vector with the constant, and the firm and bank dummies for firm-bank pair (f,b). We order the time dummy first, the firm dummies second, and the bank dummies third. Let X_t be the matrix where each row is an (f,b) pair. For $X_t'X_t$ to be invertible, we need to drop two dummies. As mentioned in the text, we drop the dummies for the first firm and bank. Finally, let Δ_t^{LT} be the vector: $(\alpha_{1t}^{LT}+\beta_{1t}^{LT}, \ddot{\alpha}_{1t}^{LT}, ..., \ddot{\alpha}_{F_t t}^{LT}, \ddot{\beta}_{2t}^{LT}, ..., \ddot{\beta}_{B_t,t}^{LT})$. We can rewrite equation (9) in vector form:

$$\mathcal{L}_t = X_t \Delta_t^{LT} + \Xi_t^{LT}$$

where $\mathcal{L}_t$ and Ξ_t^L are vectors where each row contains $\tilde{L}_{fbt}$ and ξ_{fbt}^L. Applying the standard linear projection algebra:

$$(X_t' X_t)^{-1} X_t' \mathcal{L}_t = \Delta_t^{LT} + (X_t' X_t)^{-1} X_t' \Xi_t^L$$

Or, equivalently:

$$\begin{pmatrix} \hat{c}_t^L \\ \hat{\alpha}_{2t}^L \\ \vdots \\ \hat{\alpha}_{F_t t}^L \\ \hat{\beta}_{2t}^L \\ \vdots \\ \hat{\beta}_{B_t t}^L \end{pmatrix} = \begin{pmatrix} \alpha_{1t}^{LT}+\beta_{1t}^{LT} \\ \ddot{\alpha}_{1t}^{LT} \\ \vdots \\ \ddot{\alpha}_{F_t t}^{LT} \\ \ddot{\beta}_{2t}^{LT} \\ \vdots \\ \ddot{\beta}_{B_t,t}^{LT} \end{pmatrix} + (X_t' X_t)^{-1} X_t' \Xi_t^L$$

A.7 Derivations of Aggregate Equations

We start from equations (A.5), (A.6), (2) and (5):

$$R_{ft} = \left(\sum_b \xi_{fbt}^D R_{fbt}^{1-\gamma}\right)^{\frac{1}{1-\gamma}} \tag{A.5}$$

$$L_{ft} = \left(\sum_b \left(\xi_{fbt}^D\right)^{\frac{1}{\gamma}} L_{fbt}^{\frac{\gamma-1}{\gamma}}\right)^{\frac{\gamma}{\gamma-1}} \tag{A.6}$$

$$L_{ft}^D = R_{ft}^{-\varphi} \Omega_{ft}^D \tag{2}$$

$$R_{fbt} = \mathcal{M}(s_{fbt}) \frac{R_{bt}^* \left(L_{bt}^S\right)^{\chi^{-1}}}{\xi_{fbt}^S (1 - PD_{ft})} \tag{5}$$

We denote $\mathcal{R}_{bt}^f$ the set of banks with which firm f has a relationship at time t or $t-1$, and $\mathcal{R}_{ft}^b$ the set of firms with which bank b has a relationship at time t or $t-1$.

We linearize in levels around an equilibrium where:

$$\underline{R}_{fbt} = \underline{R}_{fb,t-1} = 1 \qquad \underline{PD}_{ft} = \underline{PD}_{f,t-1} = 0 \qquad \underline{s}^D_{fbt} = \underline{s}^D_{fb,t-1} = n_{ft}^{-1}$$
$$\underline{\Omega}^D_{ft} = \underline{\Omega}^D_{f,t-1} = \Omega^D_{f,t-1} \qquad \underline{R}^*_{bt} = \underline{R}^*_{b,t-1} = R^*_{b,t-1}$$

where n_{ft} is the number of banks in $\mathcal{R}^f_{bt}$. Underlines $\underline{\cdot}$ denote that equilibrium, by opposition to actual realizations. Note that, for this situation to satisfy equations (1) and (5), we must have:

$$\underline{\xi}_{fbt} = n_{ft}^{-1} \qquad \underline{\xi}^S_{fbt} = \mathcal{M}(n_{ft}^{-1}) R^*_{b,t-1} \left(\underline{L}^S_{bt}\right)^{\chi^{-1}}$$

Also:

$$\underline{L}^S_{bt} = \sum_{f \in \mathcal{R}^b_{ft}} \underline{L}^S_{fbt} = \sum_{f \in \mathcal{R}^b_{ft}} \frac{\Omega^D_{f,t-1}}{n_{ft}} = \sum_{f \in \mathcal{R}^b_{ft}} \frac{\underline{L}_{ft}}{n_{ft}}$$

Linearization of equations (A.5), (A.6), (2) and (5) implies:

$$dR_{ft} \approx n_{ft}^{-1} \sum_{b \in \mathcal{R}^f_{bt}} dR_{fbt} \tag{A.9}$$

$$dL_{ft} \approx \sum_{b \in \mathcal{R}^f_{bt}} dL_{fbt} \tag{A.10}$$

$$dL_{ft} \approx -\varphi \Omega^D_{f,t-1}\, dR_{ft} + d\Omega^D_{ft} \tag{A.11}$$

$$dR_{fbt} \approx \mu_{ft}\, ds_{fbt} + \frac{dR^*_{bt}}{\underline{R}^*_{bt}} + \chi^{-1} \frac{dL^S_{bt}}{\underline{L}^S_{bt}} + dPD_{ft} - \frac{d\xi^S_{fbt}}{\underline{\xi}^S_{fbt}} \tag{A.12}$$

where: $\mu_{ft} = \mathcal{M}'(n_{ft}^{-1})/\mathcal{M}(n_{ft}^{-1})$. Plugging equation (A.12) into equation (A.9):

$$dR_{ft} = n_{ft}^{-1} \sum_{b \in \mathcal{R}^f_{bt}} \frac{dR^*_{bt}}{\underline{R}^*_{bt}} + \frac{\chi^{-1}}{n_{ft}} \sum_{b \in \mathcal{R}^f_{bt}} \frac{dL^S_{bt}}{\underline{L}^S_{bt}} + dPD_{ft} - n_{ft}^{-1} \sum_{b \in \mathcal{R}^f_{bt}} \frac{d\xi^S_{fbt}}{\underline{\xi}^S_{fbt}}$$

The market share drops out because they sum to 1 at the firm level. Weighting by $w_{ft} = \underline{L}_{ft}/\underline{L}_t$ and summing over f:

$$\begin{aligned}
\bar{R}_{ft} &\equiv \sum_f w_{ft}\, dR_{ft} \\
&= \sum_f \frac{w_{ft}}{n_{ft}} \sum_{b \in \mathcal{R}^f_{bt}} \frac{dR^*_{bt}}{\underline{R}^*_{bt}} + \chi^{-1} \sum_f \frac{w_{ft}}{n_{ft}} \sum_{b \in \mathcal{R}^f_{bt}} \frac{dL^S_{bt}}{\underline{L}^S_{bt}}
\end{aligned}$$

$$+\sum_f w_{ft}\, dPD_{ft} \;\; - \sum_f \frac{w_{ft}}{n_{ft}} \sum_{b \in \mathcal{R}^f_{bt}} \frac{d\xi^S_{fbt}}{\underline{\xi}^S_{fbt}}$$

Notice that we can invert the sums: $\sum_f \sum_{b \in \mathcal{R}^f_{bt}} = \sum_b \sum_{f \in \mathcal{R}^b_{ft}}$. We obtain:

$$\bar{R}_{ft} = \sum_b w_{ft} \frac{dR^*_{bt}}{\underline{R}^*_{bt}} + \frac{\chi^{-1}}{\underline{L}_t} \sum_b dL^S_{bt} \; + \sum_f w_{ft}\, dPD_{ft} \;\; - \sum_b \sum_{f \in \mathcal{R}_{fbt}} \frac{w_{ft}}{n_{ft}} \frac{d\xi^S_{fbt}}{\underline{\xi}^S_{fbt}}$$

Which we can rewrite as our aggregate supply relationship:

$$\bar{R}_{ft} = \bar{\Omega}^S_{ft} + \chi^{-1} \bar{L}^S_t + \bar{PD}_t$$

where:

$$\begin{aligned}
\bar{R}_{ft} &\equiv \sum_f w_{ft}\, dR_{ft} \\
\bar{L}^S_t &\equiv \frac{dL^S_t}{\underline{L}^S_t} \\
\bar{\Omega}^S_{ft} &\equiv \sum_b \left(w_{ft} \frac{dR^*_{bt}}{\underline{R}^*_{bt}} - \sum_{f \in \mathcal{R}^b_{ft}} \frac{w_{ft}}{n_{ft}} \frac{d\xi^S_{fbt}}{\underline{\xi}^S_{fbt}} \right) \\
\bar{PD}_{ft} &\equiv \sum_f w_{ft}\, dPD_{ft}
\end{aligned}$$

Notice that the first part of $\bar{\Omega}^S_{ft}$ i, up to a first order, a weighted average of $\tilde{\Omega}^S_{bt}$:

$$\sum_b w_{ft} \frac{dR^*_{bt}}{\underline{R}^*_{bt}} \approx \sum_b w_{ft} \tilde{\Omega}^S_{bt}$$

The second part is a weighted average of $d\xi^S_{fbt}$. By the law of large numbers, it converges to: $E\, d\xi^S_{fbt} = dE\xi^S_{fbt} = 0$, since $E\xi^S_{fbt} = 1$ by assumption.

Finally, we can easily check that equation (A.11) sums to the aggregate demand relationship:

$$\bar{L}^D_{ft} = -\varphi \bar{R}_{ft} + \bar{\Omega}^D_{ft}$$

where:

$$\bar{L}^D_t \equiv \sum_f \frac{dL^D_{ft}}{\underline{L}^D_t} = \frac{dL^D_{ft}}{\underline{L}^D_t}$$

$$\bar{\Omega}_t^D \equiv \sum_f \frac{d\Omega_{ft}^D}{\underline{L}_t^D} = \frac{d\Omega_{ft}^D}{\underline{\Omega}_t^D}$$

B Data Cleaning

We consider only non-financial firms (ESA S11). We also exclude firms that belong to the following industries: financial and insurance activities (NACE K), activities of households as employers (NACE T), activities of extraterritorial bodies (NACE U).

After collapsing the dataset at the firm-bank-quarter level, we apply the following cleaning steps:

1. Divide interest rates above 0.3 (30%) by 100;
2. Delete a bank-time if it reports:
   (a) A fall in lending of more than 30% in one quarter, and an increase of more 30% in the next;
   (b) Or an increase in lending of more than 30% one quarter, and a fall of more 30% in the next;
3. Delete a firm-time if:
   (a) It borrowed from one of the banks deleted in step 2;
   (b) The interest rate is missing on one of its loans;
   (c) The amount outstanding is negative on one of its loans;
   (d) The interest rate that it pays to one of its banks is in the first or last percentile of its country-time.

Step 1 is motivated by the fact in the early part of the sample (2018, first quarters of 2019), some banks seem to report interest rates in percent instead of decimal numbers. For instance, .05 (5%) becomes 5. We implement step 2 because some banks sometimes report a one-period drop of their total loans. This is suggestive of temporarily incomplete reporting. We've also seen an example where some of the loans of one bank were shifted to another bank of the same group for one period. In both cases, the threshold of 30% is arbitrary, but there is no perfect way around these issues. In step 3, we delete a firm-pair if at least one of its relationships is contaminated by problematic reporting: suspect bank (a), missing interest rate (b), negative amount (c), interest rate outlier (d). This step leads to a drastic drop in the number of observations (see the difference between panels A and B in table 1), but we see it as necessary to ensure that we consistently observe the firm, at a given time, in all of its relationships. While we could do away with that step in the estimation of γ — in principle, a missing relationship would be captured by the firm-time fixed effect —, we wouldn't be able to correctly construct the loan and interest rate indices when we estimate φ, let alone construct the demand shocks, ξ_{ft}^D and Ω_{ft}^D.

Finally, we keep a firm-time pair, if it survives these steps at time t and $t-1$ for the dataset that which underlies the regressions in log.-changes; and we keep a firm if it survives these steps for all time periods for the dataset in level — in order to keep a balanced sample.

C Prior and Posterior Distributions

We pick a normal prior on $\bar{\Omega}^D$ and $\bar{\Omega}^S$. For the covariance matrix, we adopt the reparametrization recommended by Gelman et al. (2013, section 15.4) into a scale vector and a correlation matrix:

$$\Sigma = \text{diag}(\mathcal{T}) \times \Sigma^0 \times \text{diag}(\mathcal{T})$$

The scale, $\mathcal{T}$, is a vector whose elements are the square roots of the diagonal elements of the covariance matrix, i.e. the standard deviations: $\mathcal{T}_j = \sqrt{\Sigma_{jj}}$. The elements of Σ_0, the correlation matrix, are the coefficients of correlation: $\Sigma^0_{ij} = \Sigma_{ij}/(T_i T_j)$. diag(.) is the operator which transforms a vector into a diagonal matrix whose elements are the rows of said vector. The advantage of this reparametrization is that we can then take a flat prior on the covariance matrix by taking a Cauchy distribution on the scale vector and a Lewandowski-Kurowicka-Joe (LKJ) distribution on the correlation matrix.[14] The LKJ distribution is a distribution over correlation matrices. It has a single parameter. If it equals 1, it is uniform over correlation matrices. As it grows, the prior favors less correlation and a stronger diagonal.

[14] See *Stan Users guide*, pp. 34–35 (Stan Development Team, 2021).

Table A.1: Prior distributions

Parameter	Prior
Elasticities	
φ	$\mathcal{N}(\hat{\varphi}, \hat{\sigma}_{\varphi})$
χ^{-1}	$\mathcal{N}(\hat{\chi}^{-1}, \hat{\sigma}_{\chi^{-1}})$
Shock distribution	
$\bar{\Omega}^{D}$	$\mathcal{N}(0, 0.05)$
$\bar{\Omega}^{S}$	$\mathcal{N}(0, 0.05)$
$\mathcal{T}$	Cauchy$(0, 2.5)$
Σ^{0}	LKJCorr(2)

Table A.2: Posterior distributions

	Mean	St. dev.	5%	95%
Elasticities				
φ	1.692	0.065	1.591	1.798
χ^{-1}	0.001	0.001	0.000	0.002
Shock distribution				
$\bar{\omega}^{D}$	0.015	0.015	-0.010	0.037
$\bar{\omega}^{S}$	0.000	0.000	0.000	0.001
$\bar{\omega}^{PD}$	-0.000	0.000	-0.000	0.000
σ^{D}	0.030	0.028	0.011	0.077
σ^{S}	0.000	0.000	0.000	0.000
σ^{PD}	0.000	0.000	0.000	0.001
ρ^{DS}	0.052	0.391	-0.583	0.702
$\rho^{D,PD}$	-0.107	0.382	-0.695	0.533
$\rho^{S,PD}$	-0.047	0.375	-0.665	0.586