Estimating the Impact of Loan Supply Shocks *

Nittai K. Bergman[1], Alejandro Casado[2], Rajkamal Iyer[3], and Itay Saporta-Eksten[4]

[1]Tel Aviv University

[2]Bank of Spain

[3]Imperial College London, CEPR and ABFER

[4]University of Manchester, Tel Aviv University, CEPR and IZA

November 2025

Abstract

Using a simple model of firm borrowing with standard ingredients, we show that commonly used empirical approaches in the literature do not recover the impact of credit supply shocks on loan-level lending, on total firm-level borrowing or on real outcomes. We propose new estimators that recover these effects. We apply our methodology to the 2011 credit crisis in Spain and show that it implies significantly smaller effects of loan supply shocks than those generated by current empirical approaches.

1. Introduction

Disruptions to the financial sector are widely recognized as key drivers of macroeconomic fluctuations, shaping business cycles, amplifying financial crises, and mediating the transmission of monetary policy (see, e.g., Gertler and Gilchrist (1994); Bernanke and Gertler (1995); Kashyap and Stein (2000); Peek and Rosengren (2000)). Understanding the effects of loan supply shocks – both on credit allocation and on real economic outcomes – is therefore a central question in economics. At the heart of this inquiry lies a core empirical challenge: disentangling shifts in credit supply from movements in credit demand. In this paper, we revisit this central question. Using a simple theoretical framework, we show that commonly used empirical approaches do not recover the impact of credit supply shocks on loan-level lending, on total firm-level borrowing or on real outcomes. We then propose an alternative methodology to do so.

*We thank Dominic Cucic, Nir Jaimovich, Asim Khwaja, David Matsa, Atif Mian and Ramana Nanda for useful discussions and comments. We thank Kemal Emre Macar for outstanding research assistance.

As is well known, understanding the impact of loan supply shocks requires distinguishing between two fundamentally different effects – each tied to a distinct set of economic questions. The first is the effect on a firm's *total borrowing* across all its lenders (henceforth, total firm borrowing effect). This aggregate response determines how credit supply disruptions transmit to investment, employment, and output, making it the central object for macroeconomic analysis. The second is the effect of a shock to a specific bank on the firm's borrowing from that bank (henceforth, loan-level effect). This loan-level response, at times referred in the literature as the "bank lending channel", is critical for evaluating bank behavior, credit reallocation, and issues pertaining to financial stability (See, e.g., Allen and Gale (2000)). Crucially, these two effects are distinct as firms can substitute borrowing across lenders.

To formalize and estimate these two objects – the impact of loan supply shocks on total firm borrowing and on loan-level lending – we begin by developing a micro-founded framework with standard ingredients: firms borrow from multiple banks to finance investments, banks are subject to different credit supply shocks, and firms can substitute borrowing between banks in response to these supply shocks. The model gives rise to two distinct effects through which credit supply shocks impact borrowing. The first is a scale effect, which captures how a change in the average cost of credit impacts a firm's total borrowing. The second is a substitution effect, which reflects how a firm reallocates its borrowing across lenders in response to relative changes in banks' financing conditions. We use our framework to show how to map the scale and substitution effects into the two objects of interest. In particular, we show that the impact of a loan supply shock on total firm borrowing is given by the scale effect, while the impact of supply shocks on loan-level lending to the firm is a *combination* of both the scale and the substitution effects.

Using our framework, we then analyze two common approaches used to estimate the impact of loan supply shocks. The first is the well-known Khwaja and Mian (2008) estimator (henceforth KM), commonly used in much of the literature. This estimator uses firm fixed effects to control for unobserved firm-level demand shocks in loan-level regressions where firms borrow from multiple banks. We show that the KM estimator captures the elasticity of substitution of firm borrowing across banks experiencing different shocks, i.e. the substitution effect. Thus, while a non-zero KM estimate indicates that supply shocks to banks are transmitted to loans, it does not map to either the loan level effect or the total firm borrowing effect.

A second common approach in the literature attempts to capture the impact of *total* supply shocks – which reflect the full change in a bank's cost or ability to lend – by using regressions with bank and firm fixed effects (Amiti and Weinstein (2018)). Rather than focusing on an observed, specific supply shifter as the KM method does – e.g., the impact of banks' deposit shock on lending – this approach interprets the bank and firm fixed effects as capturing the total supply and total demand shocks, respectively. These fixed effects are then typically used in firm real outcome regressions. Using our framework, we show that the bank and firm fixed effects do not in fact correspond to demand and supply shocks; Part of the supply shock is captured by the firm fixed effect rather than the bank fixed effect. This, in turn, implies that regressions of real outcomes on bank and firm fixed effects misestimate the impact of loan supply shocks.

We continue by deriving novel estimators that consistently identify the scale and substitution effects of specific loan supply shocks, and then use them to recover both the total firm-level and loan-level impact of loan supply shocks. To do so, we introduce a "scale-substitution regression", which is derived from our theoretical framework; This regression analyzes changes in firm level borrowing from each bank and includes on the right hand side terms that capture both the *average* supply shock faced by a firm (the scale component) and the *difference* in supply shocks across its lenders (the substitution component). This regression allows for unobserved variation in firm-level credit demand, which enters in the error term and may be correlated with loan supply shocks. To obtain consistent estimates of the scale effect, we propose a strategy based on comparing two versions of the substitution effect: (i) a biased estimate of the substitution term, obtained from the scale-substitution regression, which embeds the covariance between supply and demand shocks, and (ii) a consistent estimate of the same substitution term, obtained via the standard KM regression. The difference between these two estimates thus isolates the bias term, allowing us to back out the covariance between loan supply and demand and, in turn, correct the estimate of the scale effect in the scale-substitution regression. Armed with the scale and substitution effect estimates, we can then recover the total firm-level credit response, as well as the loan-level response, to specific loan supply shifters.

Next, we use our framework to recover the *total* loan supply shock (as opposed to a specific supply shifter) experienced by each bank and the loan demand shocks experienced by firms. This allows us to estimate the impact of total supply shocks on total firm level and loan-level lending, as well as on firm-level real outcomes. We obtain total supply and demand shocks by appropriately combining the firm and bank fixed effects from lending regressions with estimates of the scale and substitution elasticities derived from the analysis using a specific supply shifter.

We apply our novel estimation framework to the 2011 Spanish debt crisis, a period when Spanish banks with real estate exposure faced negative supply shocks stemming from the collapse of the real estate market. Using banks' pre-crisis real estate exposure as a shifter of loan supply, we find that the standard KM method overestimates the loan-level effect by 50%. We further find that the impact on total firm-level borrowing is negligible, indicating that firms were able to mitigate the shock by reallocating borrowing. Finally, we demonstrate that methodologies using bank fixed effects as a proxy for total loan supply shocks overestimate the real effects of these supply shocks on firm-level outcomes by up to 100%.

Our paper revisits a foundational question in the credit supply literature: how to estimate the impact of loan supply shocks. In doing so, we re-examine two of the most influential empirical approaches used to estimate the effects of loan supply shocks: (1) the methodology introduced by Khwaja and Mian (2008), which provides an important contribution by using firm fixed effects to control for unobserved firm loan-demand and identify the existence of the bank lending channel; and (2) the methodology introduced by Amiti and Weinstein (2018), which provides an important insight that bank fixed effects are informative of the total supply shocks hitting banks.[1] While these studies have provided crucial insights into the impact of loan supply shocks, we show that in a framework where firms can substitute borrowing between lenders, neither method recovers the key objects of interest: the effect of loan supply shocks on total firm-level borrowing and on loan-level lending. We then propose an alternate methodology to estimate these objects of interest.[2]

We proceed as follows. Section 2 provides our theoretical framework introducing the scale and substitution effects, and linking to the objects of interest – the impact of loan supply shocks at the total firm-level and at the loan-level. Section 3 introduces the empirical counterpart to the model. Section 4 re-examines existing methodologies for estimating the impact of loan supply shocks through the lens of our model. Section 5 uses our framework to present a methodology to estimate the impact of loan supply shocks – both observable supply shifters as well as total supply shocks – on loan-level lending and total firm-level borrowing. Finally, Section 6 applies our methodology to the case of the 2011 Spanish debt crisis. Section 7 concludes.

2. Understanding the Impact of Loan Supply Shifters: A Simple Model

To fix ideas, in this section we introduce a simple model which analyzes how loan supply shocks affect lending. In considering supply shocks, we will distinguish between what we will call the "total supply shock", which measures the total change in the cost of borrowing from a given bank, and a "specific supply shifter", which is a particular variable that shifts loan supply, but does not necessarily capture the entire shift in loan supply. Examples of such loan supply shifters in the literature include banks' exposure to the interbank lending market and the degree of credit line drawdowns experienced by banks. Empirically, total

[1] See for example, Iyer et al. (2014), Chodorow-Reich (2014), Jiménez et al. (2017), and Greenwald, Krainer and Paul (forthcoming) for KM applications, and Amiti, McGuire and Weinstein (2017), and Alfaro, García-Santana and Moral-Benito (2021) for applications of regressions using both bank and firm fixed effects to estimate the impact of total loan supply shocks.

[2] There are of course papers which analyze the impact of credit supply shocks using other approaches, see for example Peek and Rosengren (1997), Kashyap and Stein (2000), Paravisini (2008), Jiménez et al. (2012), Greenstone, Mas and Nguyen (2020), and Paravisini, Rappoport and Schnabl (2023).

supply shocks are not directly observed, while supply shifters, by definition, are.

The model will introduce two key parameters of interest governing the impact of loan supply shocks: the impact on total firm-level borrowing, as well as the impact on loan-level borrowing. i.e. on the firm's level of borrowing from each of its lenders. In doing so, we consider both the impact of a specific loan supply shifter as well as the impact of the total supply shock.

2.1. Model Setup

Consider a firm raising capital for a project. For simplicity, assume that the firm has no internal funds but has the option of raising external finance from two different sources of financing, which we call banks.[3] Let L_j denote the amount borrowed from bank $j \in \{1,2\}$, so that total borrowing (and hence total investment) is $L = L_1 + L_2$. Investing an amount L in the project generates cash flow with present value given by $R(L)$, which is assumed to be concave.[4]

External borrowing is costly, with the deadweight cost associated with raising an amount L_j from bank j given by $a_j c(L_j)$, with $c(L_j) = L_j^\rho$ and $\rho > 1$.[5] The parameters a_j shift the cost of external finance obtained from bank j, capturing bank-level credit supply shocks. Borrowing an amount L_1 and L_2 is thus associated with a deadweight loss of $\widetilde{C}(L_1, L_2) = \left(a_1 L_1^\rho + a_2 L_2^\rho\right)$. These deadweight cost functions capture financial frictions, such as those having to do with information or moral hazard frictions, in a reduced form way. In what follows, we analyze the elasticity of loan-level lending, L_j, and the elasticity of total firm borrowing, L, to loan supply shocks, as captured by variation in a_j.

In deciding its level of investment, the firm maximizes the second-best NPV, inclusive of the external financing costs:

$$\begin{gathered}\max_{L_1, L_2} \left\{R\left(L_1 + L_2\right) - \left(L_1 + L_2\right) - \widetilde{C}(L_1, L_2)\right\} \\ \text{with } \widetilde{C}(L_1, L_2) = \left(a_1 L_1^\rho + a_2 L_2^\rho\right).\end{gathered} \tag{1}$$

As is standard, this problem can be written in two steps. The firm chooses total borrowing, L:

$$\max_L \left\{R\left(L\right) - L - C(L)\right\}, \tag{2}$$

with $C(L)$, the cost minimization function associated with (1):

$$C(L) := \min_{L_1, L_2} \left\{a_1 L_1^\rho + a_2 L_2^\rho \,|L_1 + L_2 = L\right\}. \tag{3}$$

[3]These assumptions are easily generalizable, for the case where the firm can partially rely on internal funds to invest in the project, as well as the case where the firm can borrow from more than two banks.

[4]The results hold for the case where R is locally concave at the optimal level of investment.

[5]$c(L_j)$ is thus increasing and convex as in Stein (2003).

Solving (3), it is straightforward to show (see appendix) that the cost minimization function associated with borrowing a total amount L across both banks is given by:

$$C(L) = \kappa_C L^{\rho}, \text{ with } \kappa_C := \left(a_1^{\frac{1}{1-\rho}} + a_2^{\frac{1}{1-\rho}}\right)^{1-\rho}. \tag{4}$$

2.2. Substitution vs. Scale Effects

The parameter ρ pins down the elasticity of substitution between borrowing from the two banks, where this substitution elasticity is equal to $\frac{1}{\rho-1}$. To see this, note that the ratio of the first order conditions of the cost minimization problem implies that the optimal borrowing from each bank satisfies:

$$\log\left(\frac{L_1}{L_2}\right) = -\frac{1}{\rho-1}\log\left(\frac{a_1}{a_2}\right),$$

which, taking a_1/a_2 as the relative costs of bank 1 and bank 2 lending, yields $\frac{1}{\rho-1}$ as the definition of the elasticity of substitution.

The impact on total firm borrowing Next, we consider the impact of changes to the cost of external finance on total firm borrowing. Denoting log changes over time (between periods t and $t+1$) with the Δ operator and log-linearizing the first order conditions, we show in the appendix that to the first order the response of total firm lending across both banks to changes in the cost of external finance is given by:

$$\Delta \log L = \theta\left(s_1 \Delta \log a_1 + s_2 \Delta \log a_2\right), \tag{5}$$

where $s_j = \frac{L_j}{L}$ is the pre-shock lending share of bank j.[6]

Equation (5) shows that total lending is determined by the weighted average of the loan supply shocks across the two banks, with θ the elasticity that captures how change in total lending responds to the weighted average change in lending costs.[7] We call this effect (capturing how average loan supply shocks impact total firm lending) the 'scale effect'.

What determines the size of θ? We show in the appendix that:

$$\theta = \frac{1}{\eta_{G',L} - \eta_{C',L}} = \frac{1}{\left(\frac{1+C'}{C'}\right)\eta_{R',L} - \eta_{C',L}}, \tag{6}$$

where $G(L) = R(L) - L$ is the first best NPV of investment (exclusive of the external financing cost) when investing an amount L, $G' = \frac{dG(L^*)}{dL}$, $R' = \frac{dR(L^*)}{dL}$ and $C' = \frac{dC(L^*)}{dL}$ are the marginal NPV, the marginal PV and

[6] Note that by the nature of the approximation, the shares s_j are determined by the *pre-shock* levels of lending.

[7] Alternatively, it is the exact elasticity of lending to the weighted geometric mean of the cost.

the marginal external finance cost at the optimal level of borrowing, L^*, receptively, and we use η to denote the partial elasticity operator.[8] As would be expected, the elasticity of total borrowing to financing costs is negative, since marginal PV (R') declines with L while marginal external finance cost (C') rises with L.

The first equality in Equation (6) shows that θ depends on two forces: (1) the elasticity of the marginal financing cost, $\eta_{C',L}$, and (2) the elasticity of the project's marginal NPV, $\eta_{G',L}$. To gain intuition on these two forces, it is useful to examine the problem through the firm's FOC, which pins down total firm borrowing: $G'(L) = C'(L)$. Consider an upwards shift in the marginal cost curve $C'(L)$, which results from a change in the cost parameters a_j. As a result, the firm will reduce total borrowing, L, until the FOC is restored. Now, if the marginal NPV, $G'(L)$, declines quickly with L – i.e., $\eta_{G',L}$ is large in absolute value – then a relatively small change in borrowing generates a large increase in marginal present value, implying that the firm reduces borrowing only modestly. Put differently, under these circumstances, demand for credit will be relatively inelastic, and hence the upward shift in the marginal financing cost curve will not affect total firm borrowing, L, by much: the magnitude of θ will be low. Analogously, if the marginal financing cost $C'(L)$ rises steeply with L – i.e., $\eta_{C',L}$ is large – then restoring the FOC requires a smaller reduction in L, and hence the magnitude of θ will be small.

The second equality in (6) results directly from the fact that $\eta_{G',L} = \frac{1+C'}{C'}\eta_{R',L}$. It thus reveals a third force that influences θ: the ratio of total marginal costs $(1+C')$ to marginal external financing costs (C').[9] As is intuitive, θ, the elasticity of total firm borrowing and investment to external finance costs, declines in absolute value when external financing costs become relatively less important (i.e. when $\frac{C'}{1+C'}$ declines). Indeed, when the relative importance of marginal external financing cost declines to zero, the elasticity of investment to external financing cost goes to zero as well.

The loan-level effect Turning to the borrowing of the firm from a specific bank, we show in the appendix that to the first order the response of lending from a specific bank j to changes in the cost of external finance is given by:

$$\Delta \log L_j = \theta \left(s_1 \Delta \log a_1 + s_2 \Delta \log a_2\right) - \frac{1}{\rho - 1} s_{-j} \Delta \log \left(\frac{a_j}{a_{-j}}\right). \tag{7}$$

Equation (7), which we refer to as the scale-substitution equation, decomposes the impact of shocks to the cost of external finance on loan-level lending into two intuitive components, capturing scale and substitution effects.[10] The first term in the right hand side of the equation, $\theta\left(s_1 \Delta \log a_1 + s_2 \Delta \log a_2\right)$, is a scale effect

[8] Formally, given a function y and a variable x, we denote the partial elasticity of y with respect to x by $\eta_{y,x} := \frac{\partial y}{\partial x}\frac{x}{y}$.

[9] Concavity of $R(L)$ at L^* ensures that $\eta_{R',L} < 0$, ruling out knife-edge cases such as linear $R(L)$ where the project has infinite NPV.

[10] The case of N banks is analogous, where, similar to equation (7), the impact of shocks to external finance on the lending of

which captures the impact of the weighted average of the external finance shocks experienced by the banks lending to a given firm, on individual bank lending. The second term in the right hand side of Equation (7), $\frac{1}{\rho-1}s_2\Delta\log\left(\frac{a_1}{a_2}\right)$, reflects the substitution effect in lending, in that when the relative cost of lending from bank 1 increases, borrowing from bank 1 declines according to the elasticity of substitution $\frac{1}{\rho-1}$.

Another instructive way to write the scale-substitution equation in (7) is

$$\Delta\log L_j = \left(\theta s_j - \frac{1}{\rho-1}s_{-j}\right)\Delta\log a_j + \left(\theta s_{-j} + \frac{1}{\rho-1}s_{-j}\right)\Delta\log a_{-j}. \tag{8}$$

This formulation shows the two effects of a bank-specific loan supply shock on loan-level lending. In particular, a proportional change in a_j (holding a_{-j} constant) affects lending via two channels: a scale-effect, given by θs_j, reflecting the change in the cost of external finance, and a substitution effect, $-\frac{1}{\rho-1}s_{-j}$, reflecting the firm's tilting of borrowing from the affected bank (j) to the unaffected bank ($-j$). Note that as the share of lending from bank j rises, the impact of loan supply shocks to bank j become dominated by the scale effect.

Comparing magnitudes It is useful to compare the magnitudes of the three elasticities discussed above: (1) the elasticity of substitution, $\frac{1}{\rho-1}$, (2) the scale elasticity, θ, and (3) the loan-level lending elasticity, $\theta s_j - \frac{1}{\rho-1}s_{-j}$. To do this, note that Equation (4) implies that the elasticity of the marginal cost of external finance, $\eta_{C',L}$, is simply equal to $\rho-1$, which, together with (6), implies that

$$\theta = \frac{1}{\left(\frac{1+C'}{C'}\right)\eta_{R',L} - (\rho-1)}. \tag{9}$$

Given that $\eta_{R',L} < 0$ and $C' > 0$, it is easy to show that:

$$-\frac{1}{\rho-1} < \theta s_j - \frac{1}{\rho-1}s_{-j} < \theta < 0. \tag{10}$$

That is, in absolute value, the elasticity of substitution is larger than the loan-level lending elasticity, which is larger than the scale elasticity. Indeed, theoretically, these three values can be quite different from each other.[11] For example, as ρ approaches one – i.e. the cost function is close to linear – the elasticity of

bank j is given by:

$$\Delta\log L_j = \theta\left(\sum_{k=1}^{N} s_k\Delta\log a_k\right) - \frac{1}{\rho-1}\sum_{k=1}^{N} s_k\Delta\log\left(\frac{a_j}{a_k}\right),$$

Once again, the impact of shocks to the cost of external finance on loan-level lending can be decomposed into scale and substitution effects; scale effects driven by the impact of uniform supply shocks (across banks) on lending while substitution effects are driven by how differential changes to the cost of external finance between lenders cause the firm to reallocate borrowing from one bank to the other.

[11] The only case where the three values coincide is when the PV function, R, is linear.

substitution $\frac{1}{\rho-1}$ goes to infinity, while θ will depend on the PV function and on the relative importance of marginal cost (and can even tend to zero with small values of C').

2.3. Specific Supply Shifters versus Total Loan Supply Changes

Up to this point we have considered the impact of bank-level loan supply shocks (a_j), which reflect changes in the cost of financing banks' provide to their borrowing firms. These shocks, however, are not observable in empirical analysis. Consider, therefore, an observable variable w_j, which shifts loan supply in bank j. Examples of such an observable shifters analyzed in the literature include such variables as bank-level deposit flows (Khwaja and Mian (2008)), exposure to the interbank lending market on the eve of a financial crisis (Iyer et al. (2014)), credit line drawdowns (Greenwald, Krainer and Paul (forthcoming)), etc. Using a linear projection, we can relate the total change in cost of external finance provide the bank, $\Delta \log a_j$, to the bank-level shifter, w_j:

$$\Delta \log a_j = b_0 + b_1 w_j + \chi_j, \tag{11}$$

where without loss of generality, we assume that the transmission coefficient, b_1 is positive (so that a higher supply shifter is associated with a higher cost of financing). In what follows, we refer to the bank-level loan-supply shifter, w_j, as a *specific* loan supply shifter to emphasize the fact that many such potential shifters exist. Indeed, these other shifters are captured by χ_j in equation (11).[12] This is in contrast to variation in the bank-level cost of external finance, a_j, which captures the *total* change in loan supply at the bank (and is hence unique over a given time period for any given bank). Note also that to more easily relate the analysis to prior empirical work, we assume without loss of generality that it is the level of w, rather than changes in w, that is correlated with the loan supply shock, $\Delta \log a_j$.

Analogously to Equation (5), it is then easy to show that to the first order the response of total firm borrowing (i.e., across both banks) to the loan supply shifter satisfies:

$$\Delta \log L = \text{constant} + b_1 \theta \left(s_1 w_1 + s_2 w_2\right) + \tilde{\chi}, \quad \text{with } \tilde{\chi} = \theta \left(s_1 \chi_1 + s_2 \chi_2\right). \tag{12}$$

That is, the change in total lending to the firm responds to the weighted average of the loan-supply shifter across the two banks lending to the firm ($s_1 w_1 + s_2 w_2$), according to the scale effect elasticity θ, and the transmission coefficient b_1.

Similarly, analogously to Equation (8), it is also straightforward to show that the response of loan-level lending of bank j to the specific loan supply shifter is given by:

$$\Delta \log L_j = \text{constant} + \left(b_1 \theta s_j - \frac{b_1}{\rho - 1} s_{-j}\right) w_j + \left(b_1 \theta s_{-j} + \frac{b_1}{\rho - 1} s_{-j}\right) w_{-j} + \tilde{\chi}_j; \tag{13}$$

[12] By the properties of linear projection, χ_j is uncorrelated with w_j, and captures the other supply shifters orthogonalized to w_j.

with $\tilde{\chi}_j = \tilde{\chi} + \frac{1}{(\rho-1)} s_{-j} (\chi_{-j} - \chi_j)$.

2.4. Taking Stock

Taken together, the four equations – (5), (8),(12), and (13) – describe the four objects of interest in understanding the impact of loan supply shocks on lending.

Equation (5) describes the impact of total supply shocks on total firm-level borrowing across both banks. This effect is given by $\theta (s_1 \Delta \log a_1 + s_2 \Delta \log a_2)$.

Equation (8) describes the impact of total supply shocks on loan-level lending: $\left(\theta s_j - \frac{1}{\rho-1} s_{-j}\right) \Delta \log a_j + \left(\theta s_{-j} + \frac{1}{\rho-1} s_{-j}\right) \Delta \log a_{-j}$, which includes both the direct effect, $\left(\theta s_j - \frac{1}{\rho-1} s_{-j}\right) \Delta \log a_j$, and the cross effect $\left(\theta s_{-j} + \frac{1}{\rho-1} s_{-j}\right) \Delta \log a_{-j}$.

Equation (12) describes the impact of a specific supply shifter, w, on total firm-level borrowing. This is given by $b_1 \theta (s_1 w_1 + s_2 w_2)$.

Finally, Equation (13) describes the impact of a specific supply shifter w on loan-level lending: $\left(b_1 \theta s_j - \frac{b_1}{\rho-1} s_{-j}\right) w_j + \left(b_1 \theta s_{-j} + \frac{b_1}{\rho-1} s_{-j}\right) w_{-j}$. Similar to Equation (8), this includes both the direct effect of w_j on bank j lending to the firm, as well as the cross effect of w_{-j} on bank j lending.

3. Connecting the model to data

Our immediate goal is to relate the model to empirical analysis commonly performed in the literature. To this end, we begin by introducing random components and explicit demand shifters to the model described above (while making the time dimension explicit). We assume that the cost function for firm i in borrowing an amount L_{ijt} from bank j at time t can be written as:

$$c_j (L_{ijt}) = a_{jt} u_{ijt} L_{ijt}^{\rho}, \tag{14}$$

where a_{jt} is the bank-level cost shifter, and u_{ijt} captures a random time t bank-firm level component of borrowing costs.

Next, we allow the profitability of investment to change over time and across firms. In particular, we assume that the present value function is of the form

$$R_{it}(L_{it}) = B_{it} \widetilde{R}(L_{it}),$$

where L_{it} is firm i total borrowing (and investment) across both its banks, B_{it} is a time-varying, firm-specific parameter that shifts investment opportunities, and $\widetilde{R}$ is a time-constant investment function. This formulation thus allows for cross-sectional and time-series variation in firm level demand for loans, and in particular correlation between investment opportunities, B_{it}, and borrowing costs, a_{jt}.

Log-linearizing the first order conditions, we show in the appendix that the empirical counterpart of Equation (5), showing the response of total firm lending across both banks to the changes in the cost of external finance, is given by:

$$\Delta \log L_i = \text{constant} + \underbrace{\Delta \log \tilde{B}_i}_{x^*_{d,i}} + \underbrace{\theta\left(s_{i1}\Delta \log a_1 + s_{i2}\Delta \log a_2\right)}_{x^*_{s,i}} + \tilde{\nu}_i,\ \tilde{\nu}_i = s_{i1}\nu_{i1} + s_{i2}\nu_{i2} \tag{15}$$

where in this regression, $\Delta \log \tilde{B}_i$ is a manipulation of the shock to investment opportunities, B, and ν_{ij} is an error term which is linear in the $\log u_{ij}$ terms.[13] Equation (15) decomposes the change in firm borrowing into shifts due to demand and shifts due to supply effects. The demand component is given by $x^*_{d,i} := \Delta \log \tilde{B}_i$, whereas the supply component is given by $x^*_{s,i} := \theta\left(s_{i1}\Delta \log a_1 + s_{i2}\Delta \log a_2\right)$, i.e., the weighted average of the loan supply shifts multiplied by the scale elasticity, θ. Note that given that θ is negative, a negative cost shock amounts to a positive supply shock, $x^*_{s,i}$.

Similarly, the empirical counterpart of the scale-substitution equation in (8), which describes how lending by bank j to firm i is affected by the total supply shocks $\Delta \log a_j$ is given by:

$$\Delta \log L_{ij} = \text{constant} + \underbrace{\Delta \log \tilde{B}_i}_{x^*_{d,i}} + \underbrace{\left(s_{i,j}\theta - s_{i,-j}\frac{1}{\rho-1}\right)\Delta \log a_j}_{x^*_{s,i,j}} + \underbrace{\left(s_{i,-j}\theta + s_{i,-j}\frac{1}{\rho-1}\right)\Delta \log a_{-j}}_{x^*_{s,i,-j}} + \nu_{ij}. \tag{16}$$

Equation (16) decomposes the change in bank j lending to the firm into four components: (1) the demand component, $x^*_{d,i}$; (2) a component capturing how bank j lending to the firm is influenced by the total loan supply shock in bank j, given by $x^*_{s,i,j} := \left(s_{i,j}\theta - s_{i,-j}\frac{1}{\rho-1}\right)\Delta \log a_j$; (3) a component capturing how bank j lending to the firm is influenced by the total loan supply shock in bank $-j$, given by $x^*_{s,i,-j} := \left(s_{i,-j}\theta + s_{i,-j}\frac{1}{\rho-1}\right)\Delta \log a_{-j}$; and (4) a component capturing changes in borrowing costs at the bank-firm level, $\nu_{i,j}$.

Finally, in analyzing the lending impact of the specific loan supply shifter w (as opposed to the total supply shocks, a), the empirical counterparts of equations (12) and (13), are:

$$\Delta \log L_i = \text{constant} + \Delta \log \tilde{B}_i + b_1\theta\left(s_{i1}w_1 + s_{i2}w_2\right) + \tilde{\nu}_i + \tilde{\chi}_i, \tag{17}$$

and

$$\Delta \log L_{i,j} = \text{constant} + \Delta \log \tilde{B}_i + \left(b_1\theta s_{i,j} - \frac{b_1}{\rho-1}s_{i,-j}\right)w_j + \left(b_1\theta s_{i,-j} + \frac{b_1}{\rho-1}s_{i,-j}\right)w_{-j} + \nu_{i,j} + \tilde{\chi}_{i,j}, \tag{18}$$

[13] Specifically, without loss of generality, assume that $E\left(\log u_{ijt}\right) = 0$, so that $\Delta \log \tilde{B}_i$ is the deviation from the cross-sectional mean of $-\frac{R_i'(L_i)}{G_i'(L_i)}\theta\Delta \log B_i$. Further, $\nu_{i,j} = \frac{1}{(\rho-1)}\left(s_{i,j}\left(1+\theta(\rho-1)\right)-1\right)\Delta \log u_{i,j} + s_{i,-j}\frac{1}{(\rho-1)}\left(1+\theta(\rho-1)\right)\Delta \log u_{i,-j}$.

with $\tilde{\chi}_{i,j}$ defined as in (13).

4. Common Estimation Methods through the Lens of the Model

In this section, we use the model and its empirical counterpart to analyze two methods commonly used in the literature to estimate the impact of loan supply shocks – the KM estimator and lending regressions with bank and firm fixed effects.

4.1. Loan Regressions with Firm Fixed Effects through the Lens of the Model

Consider the following canonical KM specification in a multi-firm, multi-bank environment. Assume that between two periods, t and $t+1$, banks are affected by loan supply shocks that are potentially correlated with loan demand shocks to their portfolio firms.[14] Following KM, a common method of estimating the impact of loan supply shocks on bank lending is to run a regression of the form

$$\Delta \log L_{ij} = \beta_{KM} w_j + \delta_i + \varepsilon_{ij}, \tag{19}$$

where w_j is an observable, bank-level shifter of loan supply in bank j and δ_i is a firm-level fixed effect. The identifying assumption to obtain a consistent estimate of β_{KM} is that w_j is uncorrelated with ε_{ij} – i.e., conditional on the firm fixed effects, bank-level credit supply shocks, w_j, are uncorrelated with unobserved loan demand or supply shifters at the bank-firm level. This assumption allows for firm and bank matching – generating correlation between bank supply shocks and firm demand shocks – so long as the demand shocks are at the firm level.[15]

To connect this KM methedology to the model, we return to the maximization problem in (1), and write the first order condition with respect to the level of borrowing from bank j:

$$B_{it}\widetilde{R}'(L_{it}) = 1 + \rho a_{jt} u_{ijt} L_{ijt}^{\rho-1} \tag{20}$$

Taking logs of this first order condition, rearranging, using Equation (11), and taking first differences yields that for a specific supply shifter, w_j:

[14] The extension from two to multiple periods is trivial.

[15] As discussed in Paravisini, Rappoport and Schnabl (2023), if firm production can be broken down into various activities, and if banks specialize in financing different activities, firm fixed effects might not adequately address a correlation between demand shocks at the activity level and loan supply shocks. However, as Paravisini, Rappoport and Schnabl (2023) shows, even in the presence of specialization, if conditional on the firm fixed effects, the activity-level demand shocks are uncorrelated with the loan supply shifter then the assumption that ε_{ij} is uncorrelated with w_j will still hold.

$$\Delta \log L_{ij} = \text{constant} + \frac{1}{(\rho-1)}\Delta \log(B_i \widetilde{R}'(L_i) - 1) - \frac{b_1}{(\rho-1)} w_j - \frac{1}{(\rho-1)}(\Delta \log u_{ij} + \chi_j). \tag{21}$$

There is a natural correspondence between Equation (21) and the KM regression in Equation (19). Similar to the KM assumption, we assume that the error term in (21), $\Delta \log u_{ij}$, is uncorrelated with the credit supply shock, w_j. The error term χ_j is uncorrelated with w_j by construction given that Equation (11) is a linear projection. Further, the first term in the right hand side of (21), i.e. $\Delta \log(B_i \widetilde{R}'(L_i) - 1)$, varies at the firm level and so will be absorbed by the firm fixed effects in the KM regression in (19). Taking all this into account, running the KM regression in (19) (and maintaining the usual assumption that w_j is uncorrelated with ε_{ij}) implies that:

$$\beta_{KM} = -\frac{b_1}{\rho - 1}. \tag{22}$$

Put differently, the KM coefficient measures the *elasticity of substitution in borrowing between banks*, which is then scaled by b_1, the transmission coefficient between the credit supply shifter (w_j) and the bank-level cost of external finance (a_j). Another way of seeing that β_{KM} captures the elasticity of substitution is running the analogue of the KM regression (19):

$$\Delta \log \left(\frac{L_{i2}}{L_{i1}} \right) = \beta_{KM}(w_2 - w_1) + (\varepsilon_{i2} - \varepsilon_{i1}), \tag{23}$$

i.e., taking differences across loans within the same firm generates a standard elasticity of substitution equation.

The scaled elasticity of substitution, $\beta_{KM} = -\frac{b_1}{\rho-1}$, is different from the lending impact of loan supply shocks. In particular, as shown in Equation (12), the impact of the specific supply shifter, w_j, on total firm borrowing is given by the elasticity $b_1\theta$, whereas Equation (13) shows that the impact of the specific supply shifter on loan-level lending to the firm is given by $b_1\theta s_j - \frac{b_1}{\rho-1}s_{-j}$. Equation (10) allows us to rank the relative magnitudes of these three elasticities. Given that $b_1 > 0$:

$$-\frac{b_1}{\rho-1} < b_1\theta s_j - \frac{b_1}{\rho-1}s_{-j} < b_1\theta < 0 \tag{24}$$

Thus, while informative of the existence of a loan-level effect, $\beta_{KM} = -\frac{b_1}{\rho-1}$, overestimates the elasticity of loan-level lending to the loan-supply shifter ($b_1\theta s_j - \frac{b_1}{\rho-1}s_{-j}$) as well as the response of total firm-level lending to the loan-supply shifter ($b_1\theta$).[16]

[16] This result is different from that obtained in the theoretical model in KM. This is because the KM theoretical model does not allow for substitution across banks by borrowers.

4.2. Loan Regressions with Bank and Firm Fixed Effects through the Lens of the Model

In its essence, the bank and firm fixed effect strategy employs a regression of the following nature:

$$\Delta \log L_{ij} = \phi_i + \zeta_j + \upsilon_{ij}, \tag{25}$$

where $\Delta \log L_{ij}$ is the log change in lending from bank j to firm i between t and $t+1$, ζ_j is a vector of bank fixed effects, ϕ_i is a vector of firm fixed effects, and υ_{ij} is a loan-specific error term.[17] The fixed effects are identified using banks that lend to more than one firm and using firms that borrow from more than one bank. In this specification, the firm fixed effects are commonly thought of in the literature as capturing firm-level demand, and as a consequence, the bank fixed effects are interpreted by the literature as measuring the change in lending due to the overall loan supply shock experienced by each bank (see, for example, Amiti, McGuire and Weinstein (2017), Amiti and Weinstein (2018), Alfaro, García-Santana and Moral-Benito (2021)).

To interpret this regression we return to the first order condition in (20). Log-linearizing once again, we obtain a variant of equation (21) for total supply shocks ($\Delta \log a_j$) rather than supply shifters (w_j):

$$\Delta \log L_{ij} = \frac{1}{(\rho-1)} \Delta \log (B_i \widetilde{R}'(L_i) - 1) - \frac{1}{(\rho-1)} \Delta \log a_j - \frac{1}{(\rho-1)} \Delta \log u_{ij}. \tag{26}$$

Comparing this equation to the bank and firm fixed effects regression in (25), it is readily seen that the vector of bank fixed effects correspond to the set of scaled bank-level cost shocks, $-\frac{1}{(\rho-1)}\Delta \log a_j$, where the scaling parameter is the elasticity of substitution $\frac{1}{(\rho-1)}$.[18] By the same argument, the firm fixed effect corresponds to $\frac{1}{(\rho-1)}\Delta \log (B_i \widetilde{R}'(L_i) - 1)$. Importantly, in addition to the demand shock, this term also incorporates the firm-level response to the loan supply shocks through the optimal choice of L_i. As such, somewhat surprisingly, the firm fixed-effect incorporates not just loan demand shocks, but also some of the response to loan supply shocks.

Next, to decompose the change in loans, $\Delta \log L_{ij}$, to demand and supply effects, taking into account the

[17] There are many variants of this regression in the literature. Amiti and Weinstein (2018) conduct a comprehensive analysis that formalizes how different variants of this regression – using percent change instead of log difference, using WLS instead of OLS, and careful treatment of entry and exit of bank-firm lending relationships – have different aggregation properties. Furthermore, running this regression in practice requires taking a stand on the normalization of the fixed effect. As this normalization is immaterial for the conceptual point we make here, we relegate the discussion about normalizations to the appendix.

[18] The elasticity of substitution scaling factor is explained by the fact that the bank fixed effects are estimated in the presence of firm fixed effects, and so they are identified only from within firm substitution between banks.

supply effect which shows up in $\frac{1}{(\rho-1)}\Delta\log(B_i\widetilde{R}'(L_i)-1)$, we rearrange Equation (16) to get:

$$\begin{aligned}\Delta\log L_{ij} &= \text{constant} + \underbrace{\Delta\log\tilde{B}_i}_{x^*_{d,i}} + \underbrace{\left(s_{i,j}\theta - s_{i,-j}\frac{1}{\rho-1}\right)\Delta\log a_j}_{x^*_{s,i,j}} + \underbrace{\left(s_{i,-j}\theta + s_{i,-j}\frac{1}{\rho-1}\right)\Delta\log a_{-j}}_{x^*_{s,i,-j}} + \nu_{ij} \\ &= \underbrace{\text{constant} + \Delta\log\tilde{B}_i + \overbrace{\left(\theta+\frac{1}{\rho-1}\right)(s_{i,j}\Delta\log a_j + s_{i,-j}\Delta\log a_{-j})}^{\tau_i}}_{\phi_i} + \underbrace{\left(-\frac{1}{\rho-1}\Delta\log a_j\right)}_{\zeta_j} + \nu_{ij}.\end{aligned} \tag{27}$$

This equation introduces a new object, τ_i, which is only a function of the loan supply shocks, $\Delta\log a_j$, and hence is a supply-driven component. This component is symmetric across the two banks – i.e, it changes at the firm level. A direct result of Equation (27) is that the bank and firm fixed effect can be written as:

$$\zeta_j = x^*_{s,i,j} + x^*_{s,i,-j} - \tau_i \tag{28}$$

$$\phi_i = x^*_{d,i} + \tau_i. \tag{29}$$

Thus, the bank fixed effect ζ_j is simply the sum of $x^*_{s,i,j}$, the supply component capturing how bank j's lending to the firm is influenced by the supply shock to bank j, and $x^*_{s,i,-j}$, the supply component capturing how bank j's lending to the firm is influenced by the supply shock to bank $-j$, net of the firm-level term τ_i. By the same token, the firm fixed effect, is a sum of two components: the demand component, $x^*_{d,i}$, and the firm-level term, τ_i.

Typical applications in the literature that use the bank and firm fixed effect model proceed by using the bank fixed effects to calculate a firm-level loan-supply shock measure. They do so by calculating for each firm the share weighted average of the bank fixed effects, with shares calculated over the banks lending to the firm. Thus, the literature calculates for each firm i: $\bar{\zeta}_i := s_{1,i}\zeta_1 + s_{2,i}\zeta_2$. This measure, however, does not map to the loan supply shock at the firm level. Indeed, Equation (15) shows that the supply shock at the firm level amounts to $x^*_{s,i} = \theta\,(s_{1,i}\Delta\log a_1 + s_{2,i}\Delta\log a_2)$, which is easy to show is different than $\bar{\zeta}_i$. In fact, similar to the loan-level case, $\bar{\zeta}_i$ is missing the firm-level term, τ_i:

$$\bar{\zeta}_i = x^*_{s,i} - \tau_i.$$

Furthermore, given that the bank fixed effects correspond to $\zeta_j = -\frac{1}{\rho-1}\Delta\log a_j$, we have that at the firm level:

$$\bar{\zeta}_i = -\frac{1}{\rho-1}\,(s_{1,i}\Delta\log a_1 + s_{2,i}\Delta\log a_2)\,.$$

Note that since $-\frac{1}{\rho-1} < \theta < 0$ (Equation (10)), the actual supply driven change in lending, $x^*_{s,i}$, is smaller in magnitude (i.e. less negative) than the erroneously calculated supply driven change in lending, $\bar{\zeta}_i$.

4.2.1. Applications of Regressions with Both Bank and Firm Fixed Effects

The fact that the firm fixed effects from regressions with both bank and firm fixed effects incorporate supply effects, and that the bank fixed effects are off by a scaling factor imply that the two most common applications of the fixed effect specifications commonly employed in the literature yield biased results. First, the aggregation (either to the bank level or to the economy level) of lending changes into supply versus demand driven changes misattributes firm-level supply changes to the demand channel rather than the supply channel.

Second, the literature estimates real-effect regressions, relating firm-level outcomes (such as employment changes) to the firm fixed effects, ϕ_i, and the share-weighted average of bank fixed effects, $\bar{\zeta}_i$, interpreting the coefficient on the latter as capturing the impact of the supply driven change in lending on the real effect being examined. Given that, as shown above, ϕ_i and $\bar{\zeta}_i$ do not correspond to supply and demand driven changes in lending, these regressions will generally yield biased estimates for the impact of loan supply shocks on real outcomes. To formally see this, note that given (15), the correctly specified real effect regression is of the form:

$$y_i = \gamma_0^* + \gamma_s^* x_{s,i}^* + \gamma_d^* x_{d,i}^* + \varphi_i^* \tag{30}$$

where y_i is the firm-level real outcome (e.g. investment or change in employment), while $x_{s,i}^*$ and $x_{d,i}^*$ are, respectively, the supply- and demand-driven changes in lending as estimated above.

However, instead of the independent variables in (30), the literature uses on the right hand side the share weighted average bank fixed effect, $\bar{\zeta}_i$ and the firm fixed effect ϕ_i. What is the result of running this regression? We have shown above that $\bar{\zeta}_i = x_{s,i}^* - \tau_i$ and $\phi_i = x_{d,i}^* + \tau_i$, which when plugged into (30) yields:

$$y_i = \gamma_0^* + \gamma_s^* \bar{\zeta}_i + \gamma_d^* \phi_i + \varphi_i^* + \tau_i \left(\gamma_s^* - \gamma_d^*\right).$$

In this equation, the composite error $\varphi_i^* + \tau_i \left(\gamma_s^* - \gamma_d^*\right)$ is correlated with the explanatory variables $\bar{\zeta}_i$ and ϕ_i through the object τ_i. As such, a regression of y_i on $\bar{\zeta}_i$ and ϕ_i would generally yield biased results for γ_s^* and for γ_d^*.

5. Estimating the Impact of Loan Supply Shocks

In this section we describe a process for estimating the four objects of interest in analyzing loan supply shocks. In particular, in the next subsection we show how to estimate the impact of a specific supply shifter on total firm-level lending and loan-level lending. In the following subsection we then show how to estimate the impact of total supply shocks (i.e., not driven by a specific supply shifter) on total firm-level lending and loan-level lending.

5.1. The Impact of Specific Supply Shifters

We are after the following two objects: first, $b_1\theta$, which given Equation (12), is the elasticity of total firm level lending to the weighted average of the loan supply shifter experienced by the banks lending to the firm (i.e. the scale effect); and second, $\left(b_1\theta s_j - \frac{b_1}{\rho-1}s_{-j}\right)$, which given equation (13), is the elasticity of bank j's lending with respect to the specific loan supply shifter, w_j. As shown above, the KM regression consistently estimates the scaled elasticity of substitution, $\frac{b_1}{1-\rho}$. Thus, because bank lending shares are observable, estimating the two desired elasticities boils down to estimating the scale elasticity $b_1\theta$.

To make progress, consider the empirical counterpart of the scale-substitution equation by re-arranging Equation (18) in the following manner:

$$\Delta\log L_{i,j} = \text{constant} + b_1\theta\underbrace{(s_{i,1}w_1 + s_{i,2}w_2)}_{x_{i,1}} - \frac{b_1}{(\rho-1)}\underbrace{s_{i,-j}\left(w_j - w_{-j}\right)}_{x_{i,j,2}} + \underbrace{\Delta\log\tilde{B}_i + \nu_{i,j} + \tilde{\chi}_{i,j}}_{e_{i,j}} \tag{31}$$

Note that if we could consistently estimate equation (31), we could recover the desired scale elasticity $b_1\theta$. Of course, the model in equation (31) cannot be estimated using OLS: the firm level demand shocks, $\tilde{B}_i$, which are potentially correlated with the supply shocks, w_j, are not observable and are part of the error term, $e_{i,j}$. Instead, what is estimatable is the following *scale-substitution regression*:

$$\Delta\log L_{i,j} = d_0 + d_1x_{i,1} + d_2x_{i,j,2} + r_{i,j}. \tag{32}$$

To the extent that the error term $e_{i,j}$, which includes the demand shock, is uncorrelated with the scale and substitution variables, $x_{i,1}$ and $x_{i,j,2}$, the OLS estimate for d_1 in (32) is consistent for the elasticity of interest, $b_1\theta$. However, when the demand shock *is* correlated with the loan supply shocks – the standard concern in the literature – equation (32) suffers from classical omitted variable bias. For example, as is intuitive, when increased credit cost is associated with decreased loan demand ($cov\left(w_{j(i)}, \Delta\log\tilde{B}_i\right) < 0$), then $\hat{d}_1$ from the empirical scale-substitution regression (32) would be downward biased – i.e. more negative than the true loan supply effect, $b_1\theta$. In what follows we provide a method to estimate $b_1\theta$.

5.1.1. Estimating the Lending Elasticities of a Specific Loan Supply Shifter: A New Estimator

We show in the appendix that if $E\left[x_{i,j,2}\left(\nu_{i,j} + \tilde{\chi}_{i,j}\right)\right] = 0$, i.e., the non-demand driven loan-level idiosyncratic shocks are uncorrelated with the substitution term, we can recover a consistent estimate for the loan supply effect $b_1\theta$. Specifically, combining estimates from the scale substitution regression (32) and the KM regression, yields the following consistent estimator for $b_1\theta$:

$$\widehat{b_1\theta} = \widehat{d}_1 - \frac{1}{\widehat{\delta}_{x_1,x_2}}(\widehat{\beta_{KM}} - \widehat{d}_2), \tag{33}$$

where $\widehat{\delta}_{x_1,x_2}$ is the regression coefficient from the univariate regression of $x_{i,j,2}$ on $x_{i,1}$.

We sketch here the idea for the proof of Equation (33), providing the formal derivation in the appendix. Using a symmetry argument, we show that $x_{i,j,2}$ is not correlated with the demand shock $\Delta \log \tilde{B}_i$, and hence the error terms $e_{i,j}$ is not correlated with $x_{i,j,2}$. The bias in d_1 therefore stems only from the correlation between $e_{i,j}$ and $x_{i,1}$. At the same time, the bias in d_2 is determined by the correlation between $e_{i,j}$ and $x_{i,1}$ and the correlation between $x_{i,1}$ and $x_{i,j,2}$. We can therefore use the difference between the unbiased β_{KM} estimator and the biased d_2 estimator to recover the covariance between $x_{i,1}$ and $e_{i,j}$, and use it to debias d_1 to obtain an estimate of $b_1\theta$. As such, the obtained estimator in Equation (33) relies on the difference between β_{KM} and d_2 and the regression coefficient $\widehat{\delta}_{x_1,x_2}$, which measures the correlation between $x_{i,j,2}$ and $x_{i,1}$.[19]

Before we turn to the empirical implementation of (33), it is useful to discuss the identification assumption $E\left[x_{i,j,2}\left(\nu_{i,j} + \tilde{\chi}_{i,j}\right)\right] = 0$. This assumption is in the same spirit as the typical assumption made: while bank supply shocks, w_j, are correlated with the demand component, they are uncorrelated with the non-demand driven idiosyncratic shocks. In this particular context, the assumption is that the idiosyncratic shocks are uncorrelated with the share weighted difference in the supply shifter $w_j - w_{-j}$.

To empirically implement the new estimators and to obtain standard errors, it is useful to run equations (23) and (32), along with the regression of x_1 on x_2, as a linear system. The estimators for $b_1\theta$ is then given by a combination of coefficients as in (33). Given that an observation in the scale-substitution regression (32) is a bank-firm, while the demand shocks are at the firm-level, it is important to cluster at the firm-level.

In sum, this process results in (1) a consistent estimator of the scale elasticity $b_1\theta$; (2) by using the KM regression to obtain $-\frac{b_1}{\rho-1}$, a consistent estimate of $\left(b_1\theta s_{i,j} - \frac{b_1}{\rho-1}s_{i,-j}\right)$, the elasticity of firm i's borrowing from bank j with respect to bank j's specific loan supply shifter, w_j; (3) a consistent estimate of the *average* elasticity of firm borrowing from a given bank with respect to the bank's loan supply shifter given by $0.5\left(b_1\theta - \frac{b_1}{\rho-1}\right)$, which is the average of the scale elasticity and β_{KM}; and (4) standard errors for all estimators.

[19] The idea that recovering or bounding the covariance is useful in the estimation of the total borrowing effect goes back to KM and is formalized in Jiménez et al. (2020). However, using our framework, we show that the estimator derived in the latter does not capture the impact of loan supply shocks on total firm lending (see appendix for further details).

5.2. The Impact of Total Supply Shocks

In this section we estimate the impact of total loan supply shocks – as opposed to the impact of a specific supply shifter – on each bank's lending to the firm and on total firm lending. Equation (5) shows that the impact of total supply shocks on total firm lending across the two banks is given by $x^*_{s,i} = \theta s_1 \Delta \log a_1 + \theta s_2 \Delta \log a_2$. Further, as shown in Equation (8), the impact of total supply shock to bank j on borrowing from bank j is given by $x^*_{s,i,j} = \left(\theta s_j - \frac{1}{\rho-1} s_{-j}\right) \Delta \log a_j$. Note that these two objects are not *elasticities* of lending to supply shocks. It is natural to look at the total impact of the supply shock, rather than the elasticity, given that the total supply shocks, $\log a_j$, themselves are unobservable.

We estimate these two objects by combining estimators from the bank and firm fixed effects regression in (25) together with estimators of the lending elasticity of specific supply shocks discussed in Section 5.1 in the following manner. As shown in equation (27), . Since the loan shares $s_{i,j}$ are observed, the only thing remaining to recover in order to estimate $x^*_{s,i}$ and $x^*_{s,i,j}$ is $\theta \Delta \log a_j$. This latter object is proportional to the bank fixed effect, ζ_j, but is off by a scaling factor of $-\theta(\rho - 1)$. We thus need to rescale the bank fixed effects. This is easily done by using the estimates obtained from a specific loan supply shock, $b_1 \theta$ and $\beta_{KM} = -\frac{b_1}{\rho-1}$, discussed in Section 5.1, and dividing one by the other. In particular, we have that

$$\theta \Delta \log a_j = \underbrace{-\frac{1}{(\rho-1)} \Delta \log a_j}_{\zeta_j} \frac{b_1 \theta}{\beta_{KM}}. \tag{34}$$

Having estimated $\theta \Delta \log a_j$ and $-\frac{1}{\rho-1} \Delta \log a_j$, we can then easily estimate the two objects of interest $\left(\theta s_j - \frac{1}{\rho-1} s_{-j}\right) \Delta \log a_j$ and $\theta s_1 \Delta \log a_1 + \theta s_2 \Delta \log a_2$, which measure the impact of total loan supply shocks on total firm-level and loan-level lending, respectively.

In sum, by combining results from the regression with bank and firm fixed effects together with information from a specific, bank-level shock to loan supply, w_j, we can estimate the change in lending to firms as a result of the sum total of loan supply shocks experienced by banks (and not just due to the specific loan supply shock captured by w_j).

5.2.1. Applications Using Total Supply Shocks

In the prior section we have shown how our framework can be used to recover the change in total lending to firm i due to all loan supply shocks experienced by the banks that lend to it. This was captured by the variable $x^*_{s,i} = \theta s_1 \Delta \log a_1 + \theta s_2 \Delta \log a_2$. It is then straightforward to recover the change in total lending to the firm driven by changes in demand. Indeed, as shown in (27), the change in lending due to demand factors is given by $x^*_{d,i} = \phi_i - \tau_i$, and from the discussion above $\tau_i = x^*_{s,i} - \bar{\zeta}_i$. We thus have both demand and supply side effects.

Real Effect Regressions As discussed in section 4.2.1, a typical use of regressions with both bank and firm fixed effects is to analyze how changes in loan supply affect real outcomes (such as employment changes, investment, etc). The standard approach in the literature is to proxy for supply and demand shocks with the bank and firm fixed effects, respectively. As we have shown in section 4.2.1, the firm fixed effect captures both loan demand and loan supply side effects, which implies that the impact of loan supply shocks is mismeasured. However, with the decomposition of firm level lending into loan demand and loan supply effects provided above, it is straightforward to analyze how total loan supply shocks affect various firm-level real outcomes, by running regression (30) using estimates for $x^*_{d,i}$ and $x^*_{s,i}$.

Decomposing Aggregate Lending to Demand and Supply Next, we use our framework to decompose the bank level changes in lending into demand and supply effects. Following Amiti and Weinstein (2018), we have that the firm and bank fixed effects can be used to decompose the change in bank level in the following manner:[20]

$$\overline{D_j} = \zeta_j + \sum_{i \in I_j} q_{ij} \phi_i, \tag{35}$$

where $\overline{D_j}$ is the average change in bank j's loans to all the firms it lends to; ζ_j and ϕ_i are the bank and firm fixed effects, respectively, obtained from the firm and bank fixed effect regression; $q_{ij} = \frac{L_{ij,t-1}}{\sum_{i \in I_j} L_{ij,t-1}}$ is the share of lending to firm i out of total lending bank j, and I_j is the set of firms borrowing from bank j.

To obtain a decomposition of the change in lending into demand and supply effects, the following calcu-

[20]To obtain the exact aggregation results from Amiti and Weinstein (2018), note that the set of ζ_j and ϕ_i should be obtained from an OLS regression where the dependent variable is the percent change in loans rather than their log difference. Furthermore, to obtain aggregate growth in bank j's lending (rather than the average), one needs to run the regression with WLS. Because we assume two banks per firm, with no entry or exit, the Amiti and Weinstein (2018) corrections for such events are not relevant.

lation needs to be applied to the bank and firm fixed effects:

$$
\begin{aligned}
\bar{D}_j &= \zeta_j + \sum_{i \in I_j} q_{ij}\phi_i \\
&= \zeta_j + \sum_{i \in I_j} q_{ij}\left(x^*_{d,i} + \tau_i\right) \\
&= \sum_{i \in I_j} q_{ij}(\tau_i + \zeta_j) + \sum_{i \in I_j} q_{ij}x^*_{d,i} \\
&= \underbrace{\sum_{i \in I_j} q_{ij}x^*_{s,i,j}}_{\text{Own Supply}} + \underbrace{\sum_{i \in I_j} q_{ij}x^*_{s,i,-j}}_{\text{Peer Supply}} + \underbrace{\sum_{i \in I_j} q_{ij}x^*_{d,i}}_{\text{Demand}},
\end{aligned} \tag{36}
$$

The first line is the Amiti and Weinstein (2018) decomposition into firm and bank fixed effects, while the subsequent lines are algebraic manipulations directly from Equation (28). Equation (36) provides the true decomposition of bank level lending into supply and demand elements, where the fact that the firm fixed effects also include loan supply effects (as captured by τ_i; see equation (27)) have been accounted for. The loan supply component is divided into two: (1) The "own supply" element which captures how bank j's total lending is influenced by supply shocks experienced by bank j itself, and (2) The "peer supply" element which captures how bank j's total lending is influenced by supply shocks experienced by bank j's "co-lenders" – i.e., the banks which lends to the firms to which bank j lends to.[21]

Finally, one can also decompose the average loan growth in the economy into demand and supply effects relying on $\bar{D} = \sum_j z_j \bar{D}_j$, where $z_j = \frac{\sum_{i \in I_j} L_{ij,t-1}}{\sum_j \sum_{i \in I_j} L_{ij,t-1}}$.

6. Application: The 2011 Spanish Debt Crisis

6.1. Background and Data

In this section, we apply our framework to the 2011 Spanish sovereign debt crisis. The bursting of the housing bubble in the late 2000s in Spain burdened Spanish banks with distressed real estate assets, triggering financial instability (see for example Baudino, Herrera and Restoy (2023)). We apply our methodology to this setting to measure the impact of loan supply shocks and contrast our results against those that would be generated with common empirical strategies used in the existing literature.

We employ a detailed dataset from Bank of Spain's confidential credit register (CIR) which includes loan-level credit registry data in Spain, combined with firm-level financial information available from the

[21] Note, that as is clear from (36) that the corrected decomposition inherits the aggregation properties shown in Amiti and Weinstein (2018).

Spanish Mercantile Register. We use the degree of banks' exposure as of the end of 2011 to real estate loans as a 'specific supply shifter' to proxy for the severity of lending constraints. Our analysis proceeds in two steps: first, we estimate the impact of loan supply shocks resulting from this specific real-estate-exposure supply shifter, and then, we estimate the impact of the total loan supply shock affecting banks. In doing so, we analyze both lending and real effects.

To match the analysis to our model and empirical framework, we consider only firms with two lenders in 2011 and 2012.We further limit our analysis to firms that have the same two lenders in 2011 and 2012 and also that have data on firm-level characteristics. Table 1 presents the descriptive statistics. There are a total of 29,964 firms. The average decline in lending between the period 2011 and 2012 is approximately 15%. In our analysis, we use real estate exposure of banks as the loan supply shifter. Real estate exposure is measured as the ratio of a bank's total real estate lending to its total lending. The sample is comprised of 115 banks, which on average have a real estate exposure of 57%.

6.2. The Impact of Real-Estate Exposure

Table 2 presents the estimates for the impact of real estate exposure obtained using our empirical framework and compares them to results obtained using the standard empirical techniques in the literature. The first row in the table reports the result of a standard KM lending regression, which uses the log change in lending at the bank-firm level between 2011 and 2012 as the dependent variable, the real estate exposure as the loan supply shifter, and which includes firm fixed effects to control for unobservable loan demand. As our framework shows, the KM coefficient on the real estate exposure (-0.183), captures only the substitution effect $-\frac{1}{\rho-1}$, scaled by b_1.

Rows 2 and 3 of Table 2 show the results of the scale-substitution regression, equation (32). The coefficients of interest d_1 and d_2 are -0.039 and -0.178, respectively. These coefficients are potentially biased for the scale and substitution effects since the demand shocks are not accounted for in the regression. We then apply equation (33) and obtain the estimate of $b_1\theta$, shown in the fourth row of Table 2. Because the difference between the biased (d_2) and unbiased (β_{KM}) estimates of the substitution effect is small, the estimated covariance term, $cov(x_{i,j,1}, e_{i,j})$, which captures the correlation between demand and supply shocks, is close to zero. This implies that $b_1\theta$ is close to the biased scale effect d_1.

Finally, the last row of Table 2 reports the estimate of the average loan level elasticity to be -0.116. As shown by the theoretical discussion, indeed this value is smaller (in absolute value) compared to β_{KM}, the latter of which is typically treated as the loan-level effect. The difference between the coefficients is 0.067 – i.e., β_{KM} overestimates the loan-level effect by over 50%. Estimated at -0.052, the firm level effect, $b_1\theta$, is even closer to zero and is not statistically significant at conventional levels.

6.3. The Impact of Total Supply Shocks on Real Outcomes

We turn now to estimate the impact of total supply shocks on lending and real outcomes. We start with the estimation of a lending regression with bank and firm fixed effects (equation (25)). Using our estimates of $b_1\theta$ and β_{KM}, along with equation (34) we recover the total supply shocks experienced by different banks, $x^*_{s,j}$. Figure 1 compares the distribution of supply shocks at the bank level ($x^*_{s,j}$) to the distribution of bank fixed effects (ζ_j). Given our theoretical framework, which implies that true supply shocks are a scaled version of the bank fixed effects (equation (34)), the distribution of supply shocks in the figure is a compressed version of the distribution of bank fixed effects. This highlights that the bank fixed effects overstate the extent of total supply shocks experienced by banks.

Next, following the discussion in 5.2.1 we recover estimates for the firm level supply and demand shocks, $x^*_{s,i}$ and $x^*_{d,i}$. This allows us to examine the impact of supply shocks on real outcomes including investment, employment, and value added. As discussed in section 4.2.1, the regression of real outcomes on the estimated bank and firm fixed effects from the lending regression deliver biased estimates for the impact of supply shocks. The results in Table 3 demonstrate that in our sample, the difference between the biased and unbiased estimates is large. In particular, from column (1), a one standard deviation change in averaged bank FE ($\bar{\zeta}_i$) is associated with 0.86 percent increase in investment. Correcting for the bias, column (2) shows that a one standard deviation increase in total supply ($x^*_{s_i}$) is associated with only a 0.41 percent increase in investment – less than half the magnitude of the biased coefficients.[22] Columns (3) to (6) reveal similar patterns for employment and for value added growth.

7. Conclusion

This paper revisits a foundational question in the credit supply literature: how to measure the impact of loan supply shocks on firm borrowing. Using a theoretical framework that endogenizes firm borrowing decisions while allowing for substitution in lending across banks, we show that two widely used empirical strategies – the Khwaja and Mian (2008) approach and the bank and firm fixed effects regression approach as in Amiti and Weinstein (2018) – do not recover the total firm- or loan-level impact of supply shocks.

We then provide a methodology to estimate total firm-level and loan-level effects of loan supply shocks. We do so for both specific loan supplier shifters as well as for total supply shocks. This methodology is

[22]Recall that a more positive $x^*_{s_i}$ reflects a larger positive supply shock.

particularly relevant for understanding how shocks to the financial system propagate to investment, employment, and output, and thus for informing the design of policies aimed at promoting financial stability and mitigating the amplification of crises. Applying our methodology to the case of the Spanish 2011 debt crisis, we show that our new estimators yield substantially smaller real effects of loan supply shocks.

References

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

Allen, Franklin, and Douglas Gale. 2000. "Financial contagion." *Journal of political economy*, 108(1): 1–33.

Amiti, Mary, and David E Weinstein. 2018. "How much do idiosyncratic bank shocks affect investment? Evidence from matched bank-firm loan data." *Journal of Political Economy*, 126(2): 525–587.

Amiti, Mary, Patrick McGuire, and David E Weinstein. 2017. "Supply-and demand-side factors in global banking." National Bureau of Economic Research.

Baudino, Patrizia, Mariano Herrera, and Fernando Restoy. 2023. "The 2008-14 banking crisis in Spain."

Bernanke, Ben S, and Mark Gertler. 1995. "Inside the black box: the credit channel of monetary policy transmission." *Journal of Economic perspectives*, 9(4): 27–48.

Chodorow-Reich, Gabriel. 2014. "The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis." *The Quarterly Journal of Economics*, 129(1): 1–59.

Gertler, Mark, and Simon Gilchrist. 1994. "Monetary policy, business cycles, and the behavior of small manufacturing firms." *The quarterly journal of economics*, 109(2): 309–340.

Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen. 2020. "Do credit market shocks affect the real economy? Quasi-experimental evidence from the great recession and ânormalâ economic times." *American Economic Journal: Economic Policy*, 12(1): 200–225.

Greenwald, Daniel L, John Krainer, and Pascal Paul. forthcoming. "The Credit Line Channel." *Journal of Finance*.

Iyer, Rajkamal, José-Luis Peydró, Samuel da Rocha-Lopes, and Antoinette Schoar. 2014. "Interbank liquidity crunch and the firm credit crunch: Evidence from the 2007–2009 crisis." *The Review of Financial Studies*, 27(1): 347–372.

Jiménez, Gabriel, Atif Mian, José-Luis Peydró, and Jesús Saurina. 2020. "The real effects of the bank lending channel." *Journal of Monetary Economics*, 115: 162–179.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina. 2012. "Credit supply and monetary policy: Identifying the bank balance-sheet channel with loan applications." *American Economic Review*, 102(5): 2301–2326.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina. 2017. "Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the Spanish dynamic provisioning experiments." *Journal of Political Economy*, 125(6): 2126–2177.

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

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

Paravisini, Daniel. 2008. "Local bank financial constraints and firm access to external finance." *The Journal of Finance*, 63(5): 2161–2193.

Paravisini, Daniel, Veronica Rappoport, and Philipp Schnabl. 2023. "Specialization in bank lending: Evidence from exporting firms." *The Journal of Finance*, 78(4): 2049–2085.

Peek, Joe, and Eric S Rosengren. 1997. "The international transmission of financial shocks: The case of Japan." *The American Economic Review*, 495–505.

Peek, Joe, and Eric S Rosengren. 2000. "Collateral damage: Effects of the Japanese bank crisis on real activity in the United States." *American economic review*, 91(1): 30–45.

Stein, Jeremy C. 2003. "Agency, information and corporate investment." *Handbook of the Economics of Finance*, 1: 111–165.

Table 1: Descriptive Statistics

	mean	s.d.
$\Delta \log L$	-0.153	0.413
Real Estate Exposure (w_j)	0.568	0.115
Number of firms	29,964	
Number of banks	115	

Table 2: Impact of Loan Supply Shock: Spanish Debt Crisis

	Interpretation	Source	Estimate (s.e)
β_{KM}	Substitution effect ($-\frac{b_1}{\rho-1}$)	Standard KM reg	-0.183 (0.020)
d_1	Biased scale effect	SUR (scale-substitution)	-0.039 (0.019)
d_2	Biased substitution effect	SUR (scale-substitution)	-0.178 (0.020)
Implied elasticities			
$b_1\theta$	firm-level effect on total lending	SUR applying section 5.1.1	-0.052 (0.035)
$0.5b_1\theta - 0.5\frac{b_1}{\rho-1}$	average loan-level effect on total lending	SUR applying section 5.1.1	-0.116 (0.020)

Table 3: Real Effect Regressions

	Long Term Assets Growth		Employment Growth		Value Added Growth	
	(1)	(2)	(3)	(4)	(5)	(6)
	Bank FE ($\bar{\zeta}_i$)	Supply Shock ($x^*_{s,i}$)	Bank FE ($\bar{\zeta}_i$)	Supply Shock ($x^*_{s,i}$)	Bank FE ($\bar{\zeta}_i$)	Supply Shock ($x^*_{s,i}$)
Total Supply Shock $\times 100$	0.858	0.413	0.149	-0.095	1.024	0.786
(s.e.)	(0.162)	(0.162)	(0.230)	(0.231)	(0.341)	(0.340)
Observations	29,939		26,472		26,837	

Notes: The table reports regressions of firm level outcomes on supply shocks. In columns (1) and (2) the dependent variable log change in long term assets. In columns (3) and (4) the dependent variable is employment growth rate. In columns (5) and (6) the dependent variable is value added growth, defined as log change in value added. In Columns (1), (3) and (5) the explanatory variable is the Bank FE ($\bar{\zeta}_i$), controlling also for the firm FE (ϕ_i). In Columns (2), (4) and (6) the explanatory variable is $x^*_{s_i}$, controlling also for $x^*_{d_i}$.

Figure 1: Bank Level Total Supply Shock Distribution