

Identifying Heterogeneous Supply and Demand Shocks in European Credit Markets*

Olivier De Jonghe[†] Daniel Lewis[‡]

July 1, 2025

Abstract

We propose a new model in which relationship-specific supply and demand shocks are non-parametrically identified in bipartite data under mild assumptions. For example, separate heterogeneous supply shocks are identified for each firm to which a bank lends. We show that a simple estimator is consistent, derive its limiting distribution, and illustrate its performance in simulations. Using these methods, we identify the heterogeneous distributions of supply and demand shocks for thousands of banks and firms in 11 European countries using the Anacredit dataset. Our estimates characterise how both quantity and price elasticities, and thus supply and demand curves, have changed in those 11 markets in recent years. The shock distributions exhibit within-firm/bank heterogeneity that is not well-explained by conventional fixed effects approaches, which only capture between-firm/bank heterogeneity. This unexplained heterogeneity correlates strongly with economically meaningful relationship-level characteristics and macroeconomic policy measures. These results have important implications for policy, identification assumptions in empirical work, and modeling exercises.

Keywords: supply shock, demand shock, corporate credit, identification, higher moments, networks, fixed effects

JEL codes: C33, C58, E44, G21, G30

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Bank of Belgium and the Eurosystem. We are grateful to seminar participants at UCL, the Banque de France and the European Central Bank for helpful comments and Ferre De Graeve, Aureo de Paula and Andrei Zeleneev for helpful discussions.

[†]European Central Bank, National Bank of Belgium, Ghent University, and Tilburg University.

[‡]University College London, daniel.j.lewis@ucl.ac.uk. Research supported by 2024 ERC Starting Grant CreMac.

1 Introduction

A fundamental challenge in empirical settings is to disentangle the effects of supply and demand. Approaches taken in the literature range from reduced form (for instance, unobserved heterogeneity models, e.g., Abowd et al. (1999)), to those blending instrumental variables and structural specifications (e.g., Berry et al. (1995), MacKay and Miller (*forthcoming*)) to highly structural (e.g., Olley and Pakes (1996)). While sometimes the elasticities of prices and quantities with respect to supply and demand are of interest, supply and demand “shocks” themselves – unpredictable innovations to supply or demand – can also offer important insights. This is often the case, for example, in studies of credit markets (e.g. Khwaja and Mian (2008), Greenstone et al. (2020), Chava and Purnanandam (2011), or Amiti and Weinstein (2018)) – researchers want to measure how much of an observed change in credit quantity may be associated with a change in demand from the firm, versus a change in credit supplied by the bank. This knowledge may allow policymakers, for instance, to understand whether credit fluctuations are supply- or demand-driven, and hence craft an appropriate policy response. For concreteness, we maintain the firm-bank credit market setting as a running example throughout this paper.

Credit markets are an example of a bipartite network – one in which there are two types of agents, who form relationships with agents of the other type, but not their own type. This paper studies the many-to-many bipartite network setting. Although specialised, this setting is actually relatively common in applied work and has received considerable methodological attention, building on the seminal contribution of Abowd et al. (1999). Other empirical environments include importers and exporters and many types of financial markets where some firms serve as dealers of a type of asset to collections of buyers; for a review of such settings, see Bonhomme (2020). While Abowd et al. (1999) originally intended to recover permanent characteristics of the workers and the employers as fixed effects, the same methodology has been widely used to identify supply and demand shocks, under a homogeneity assumption. That is, a firm experiences a single demand shock in each period, which uniformly impacts the quantity of credit it demands across all banks with which it interacts. Likewise, each bank adjusts its supply of credit to all firms uniformly. As a result, demand shocks are purely firm-specific objects, and supply shocks bank-specific objects, and the “shocks” are thus estimable as simple fixed (or random) effects in a linear regression (e.g., Abowd et al. (1999), Bonhomme (2020), Khwaja and Mian (2008), Greenstone et al. (2020), Amiti and Weinstein (2018)). These are very strong assumptions, at odds with models of credit shock transmission; they also rather limit the types of policy questions that a researcher can answer. Indeed, after a tightening of some regulatory requirement for banks,

a policymaker would be keenly interested to know if banks shift credit away from riskier firms towards safer borrowers: What are the distributional effects of regulatory changes? This is not possible in the standard framework since all shocks are necessarily homogeneous by assumption.

In this paper, we depart from the simple additive model studied by Abowd et al. (1999) and others to consider one in which each relationship has specific – generally unique – associated supply and demand shocks in each period. To do so, we crucially study a pair of innovations to observed outcomes – price and quantity, say – rather than a single outcome, which has typically been the change or growth in loan quantity in the credit markets literature.¹ We assume that the innovation to price and quantity in each firm-bank relationship is a linear combination of the firm-bank specific supply and demand shocks, with the associated matrix of elasticities constant across firm and bank pairs. Demand shocks from different firms to the same bank may be correlated, as may demand shocks from one firm to different banks – and vice versa for supply shocks from banks to firms. The key restriction is that the demand shock from one firm to a particular bank is uncorrelated with the supply shock from a different bank to that firm; symmetrically, the supply shock from one bank to a particular firm is uncorrelated with the demand shock from a different firm to that bank. This is a largely innocuous assumption, since it is much weaker than the theoretical properties usually stipulated for supply and demand shocks. Under these assumptions, we show that the matrix of elasticities linking price and quantity to supply and demand shocks is non-parametrically identified from two covariance matrices: one measuring the covariance of price and quantity innovations across firms, holding the lending bank fixed, and the other measuring the covariance of price and quantity innovations across banks, holding the borrowing firm fixed. Given these matrices, the elasticities are available in closed form and the shocks themselves are recoverable by simply inverting the relationship between price and quantity and the shocks. Crucially, the identification argument requires only a single period of data, meaning that while elasticities are assumed constant across relationships, they need not be unduly restricted over time.

Next, we establish the asymptotic properties of an estimator based on the sample counterparts of those two covariance matrices. As is typical in network settings, this is a non-standard problem, since, even as the number of firms and banks grows large, dependence of new observations with previous observations does not in general fade. Asymptotically, the final observations may be strongly correlated with the first – for instance if the final bank lends to a common firm with the first bank. Naturally, some further restrictions are needed

¹This is not strictly novel; Graeve and Karas (2014) also consider price and quantity in a SVAR-based study of bank runs.

on the shocks – in particular, we assume that supply and demand shocks are independent, and can be decomposed into two components: one that may be correlated across firms and one that may be correlated across banks. Then, under suitable assumptions on the growth rates of the numbers of banks and firms, the sample counterparts of these covariance matrices are consistent. The resulting estimator for the matrix of elasticities is also consistent and asymptotically normal. The asymptotic variance may be consistently estimated with a surprisingly simple estimator, very similar to a cluster-robust variance estimator. Crucially, the assumptions required for these asymptotic results do not imply that the network of firms and banks is unrealistically dense; there need only be some non-vanishing fraction of “large” firms/banks whose number of trading partners increases at the same rate as the total number of banks/firms. All of these results extend to estimation strategies where data is pooled across multiple periods in order to obtain a richer sample.

We examine the performance of our proposed estimator in simulations. The simulations clearly demonstrate the consistency of our approach for the matrix of elasticities, and exhibit relatively small biases even in samples considerably smaller than those present in the data. Predictably, as more periods of data are pooled together, there are efficiency gains and reduced bias, even when the shocks exhibit persistence over time. Using our asymptotic variance estimators, simple t -tests control size well at the nominal level for sample sizes in the range observed in the data. When the shocks are persistent, the autocorrelation-robust version of our variance estimator continues to control size at the nominal level. Next, we compare our estimates of the average demand/supply shock a firm/bank experiences to those obtained using the standard fixed effects approach. The simulations show that, even if a homogeneous “average shock” is the object of interest, taking averages of the firm-bank specific shocks we recover is much more accurate than estimating fixed effects. This is because our estimator uses information in both the firm and bank dimensions to separate supply or demand shocks, rather than just one or the other.

We use the new methodology to study European credit markets from 2019-2023 using the AnaCredit dataset. We recover considerable time variation in the elasticities of price and quantity to supply and demand shocks over the pandemic, inflation, and monetary policy tightening periods. With each time period, there is also considerable heterogeneity by country. Representing the elasticities as supply and demand curves demonstrates that markets are frequently characterised by inelastic supply or demand, but that this again varies by period and country.

Studying the shock distributions uncovers a wealth of findings. We fit a series of fixed effects models to the estimated shocks and find that fixed effects, like those the literature relies on, explain very little of the variation in the distributions we uncover. Notably, a

vanishingly small share of variation in demand shocks is explained by fixed effects. This is consistent with further findings that the within-firm and within-bank variation is comparable to or larger than the between-firm or between-bank variation in demand and supply shocks, respectively. Crucially, the variation in both demand and supply shocks not captured by standard fixed effects is economically meaningful. It is strongly correlated with relationship-specific variables measuring loan type and the importance of the relationship. The signs of these relationships vary with the period considered, yet they generally remain highly significant. We show that time variation in the obtained relationships are linked to changes in monetary and macroprudential policy. A bank’s subjective assessment of a firm’s probability of default predicts heterogeneous variation in supply shocks but has zero correlation with demand shocks; this strongly indicates that demand and supply shocks have been reliably identified. Removing fixed effects from the regression and including bank- or firm-level regressors reveals significant relationships with the shocks but even less explanatory power. Fixed effects still capture unobservable heterogeneity not explained by the covariates available in AnaCredit, even though they leave the majority of shock variation unexplained. In a final exercise, we also show how studying heterogeneity in the credit demand and supply shock may resolve ambiguities or hidden effects when analyzing heterogeneity in credit growth and interest rate changes.

This paper builds on several literatures. We are not aware of any direct antecedents for the model we propose, but it is closely related to both fixed and random effects approaches to modeling heterogeneity in bipartite networks. In particular, it can be viewed as a generalisation of both Abowd et al. (1999) as well as earlier work by Chamberlain (1980); the latter is perhaps closer since it views the underlying “shocks” as random effects whose moments are to be estimated, as opposed to objects to be directly estimated. However, these approaches maintain the same key homogeneity assumption discussed above – each firm has one fixed/random effect “demand shock” in a given period, and each bank a single fixed/random effect “supply shock”. Any heterogeneity is noise in the error term (typically i.i.d. within each period), as opposed to an object to be studied. In contrast, we model each *relationship* as experiencing two unique shocks in each period – one which has a structural interpretation as a demand shock, and the other as a supply shock. Some papers, like Graham et al. (2014), do consider extensions to nonlinear random effects models. However, in these models the corresponding effects are still assumed to be some function of the homogeneous effects from each side of the market. Rather than assuming that the demand shocks a firm issues to different banks are all *identical*, we assume merely that there exists some non-zero *correlation* between them, and likewise for a bank’s supply shocks. Indeed, the possibility of homogeneous shocks, like in the fixed effects model of Abowd et al. (1999), or random effects

models, is nested as a special case, as we discuss below. However, it is also possible to view our contribution as providing a structural extension of the Abowd et al. (1999) model. For instance, the Abowd et al. (1999) methodology could model the conditional mean of the innovations to price and quantity, with our model providing a structural interpretation to the *error terms* arising from that familiar specification. For these reasons, we view our model as highly versatile: it relaxes a key homogeneity assumption and nests existing models, but can alternatively be seen as a complement to those models.

These empirical results have important implications in several directions. First, they provide insights potentially useful for policy, by characterising how price and quantity are likely to respond to shocks within a given credit market, as well as illustrating how firms' credit demand and banks' credit supply respond to different types of events at a granular, distributional level. Second, by testing the correlations between supply and demand shocks and various observable characteristics, our results allow a laboratory within which identification assumptions may be tested for use in less data-rich credit market datasets. Finally, our results provide important targets for modeling exercises: we provide correlations for a bank's credit supply with firm characteristics, for instance, that modelers can seek to match.

While the setting of the problem is novel, the structure of the identification argument parallels strategies more commonly found in time series settings. In particular, it is connected to identification approaches based on heteroskedasticity, for example in SVARs, see for instance Rigobon (2003). In these methodologies, two or more covariance matrices are used to identify a common matrix of elasticities. However, to the best of our knowledge, such arguments have never been used for objects other than simple variance-covariance matrices; here, we apply the argument to covariance matrices between different firms and between different banks. Ostensibly, the identification strategy may seem reminiscent of the "granular IV" approach of Gabaix and Kojen (2024), since both exploit idiosyncratic variation in demand shocks for identification. However, our contribution uses the variation *within* each firm's demand shocks – across its relationships – for identification, and focuses on higher moments. Moreover, our identification rests on higher moments of the shocks – their covariance across relationships – while granular identification exploits the fact that a cleverly constructed weighted average of demand is a valid instrument for aggregate demand (in the price equation) or price (in the demand equation). A key feature of granular identification is also the use of a common price across entities, while we require price variation not only at the entity level, but the relationship level.

While asymptotic analysis is largely missing from the original Abowd et al. (1999) paradigm, it was recently provided by Jochmans and Weidner (2019), who note that in the real world data studied by papers like these, incidental parameters bias arising from

many firms having a small number of relationships means that these estimates may be unreliable. In contrast, the methodology we propose leverages information from all firms and all banks simultaneously to estimate the elasticity matrix decomposing price and quantity innovations into shocks. Note that a straightforward comparison of rates is not possible since this paper proposes to consistently estimate the matrix of elasticities mapping observed innovations to shocks, while Jochmans and Weidner (2019) provide conditions for the consistency of the fixed effects – and thus shock estimates – themselves. Their variance bounds are complicated functions not only of the degree of a given firm or bank, but also those of every entity it is connected to. However, we do discuss inference for moments of the recovered shocks below; the sampling uncertainty around an estimate of the average shock associated with any firm or bank is inversely proportional to its degree.

Empirically, this paper is related to the literature recovering supply and demand shocks in corporate credit markets – much of which owes its origins to Abowd et al. (1999). For instance, Greenstone et al. (2020) use a shift-shares approach to recover bank lending supply shocks at the county level. Khwaja and Mian (2008) exploit an exogenous event impacting bank liquidity to estimate heterogeneous credit supply shocks to firms and their effects. Amiti and Weinstein (2018) develop a methodology to separate firm credit demand shocks from bank supply shocks and quantify their effects on overall investment dynamics. However, all of these approaches maintain the assumption of homogeneous shocks, implicitly or explicitly making the assumptions underlying a fixed effects model, and ultimately estimating at least demand shocks as fixed effects.

The remainder of the paper is organised as follows. Section 2 describes the setting, existing methodologies, and the novel identification argument. Section 3 defines the estimators and establishes their asymptotic properties. Section 4 presents simulation results. Section 5 introduces the AnaCredit dataset and analyses our key credit market elasticities vary both across countries and across time. Section 6 estimates the distributions of supply and demand shocks in the same data and establishes which observable characteristics of firms and banks explain heterogeneous supply and demand dynamics. Section 7 concludes.

2 Identification

In this paper we consider the loan market between firms and banks, although our results apply equally to any many-to-many bipartite network. Loan quantities between a firm-bank pair are simultaneously determined by both supply and demand effects. We want to decompose changes in loan quantity into supply- and demand-driven components. Typically, this has been achieved by assuming a standard Abowd et al. (1999)-style model for changes

in loan volume with unobserved heterogeneity:

$$\Delta l_{fb}^* = d_f + s_b + \Gamma X_{fb} + \epsilon_{fb}, f = 1, \dots, F, b = 1, \dots, B, \quad (1)$$

where Δl_{fb}^* is a potential outcome that is realized if firm f has a relationship with bank b , and $\Delta l_{fb} = \Delta l_{fb}^* D_{fb}$ is observed, where D_{fb} is an indicator for whether a lending relationship exists between f and b . d_f is the firm-specific demand component and s_b the bank-specific supply component. Additionally, ϵ_{fb} is mean-zero and i.i.d. It is generally assumed that $E[\epsilon_{fb}|D, X, \mathbf{d}, \mathbf{s}] = 0$, where \mathbf{d}, \mathbf{s} stack the realised firm/bank effects; that is, X , and the relationships encoded by D , are strictly exogenous. This allows the use of standard reduced-form regression techniques. However, it does accommodate endogenous network formation, since D may only be exogenous after additionally conditioning on $X, \mathbf{d}, \mathbf{s}$. Often, models include a third time subscript on each observation. (1) can be cast as either a fixed effects or random effects model depending on the setting. The former is most common in classic treatments of bipartite networks (e.g., Abowd et al. (1999)), as well as leading empirical studies of credit markets, for instance Greenstone et al. (2020), Khwaja and Mian (2008), and Amiti and Weinstein (2018). As shown in Abowd et al. (1999), under the exogeneity assumption above, Γ can be consistently estimated, and we can focus instead on the simpler model

$$\Delta l_{fb} = d_f + s_b + \epsilon_{fb}, \quad (2)$$

where it is understood that a) we limit our attention to observed relationships (i.e. $D_{fb} = 1$) and b) any covariates, X , have been partialled out. We will henceforth focus on this simpler model, and return to the problem of including covariates in Section 3.3. The statistical properties of the model depend on the structure of the network, but Jochmans and Weidner (2019) show that the *degree* (i.e. the number of connections of a firm or bank) appears in the denominator of the variance (although the numerator also depends in a complicated manner on the same quantity). This means that – holding certain properties of the network fixed – if the number of banks a firm interacts with goes to infinity (or vice versa), the firm and bank effects can be consistently estimated. This is admittedly a strong assumption, but implicitly underpins the approaches of existing empirical work; otherwise, these estimates suffer from incidental parameters bias (e.g., Abowd et al. (2004), Andrews et al. (2008), Jochmans and Weidner (2019)).

A crucial uninterrogated assumption is that demand shocks, d_f , vary across firms, but, for a given firm, impact its borrowing from each bank identically. The same is true for supply shocks, s_b , with respect to each firm. Any heterogeneity across relationships is driven by random noise in ϵ_{fb} , which is i.i.d. across firms and banks, and cannot be given any

structural interpretation. However, it could be the case that, for example, faced with tighter regulatory requirements or financial conditions, the supply shocks from bank b to different firms are different — “riskier” firms might face a more negative supply shock as the bank attempts to rebalance its loan portfolio. Instead of assuming that shocks are *homogeneous* for each entity, we will instead make the weaker assumption that they are *correlated* for each entity; the associated variation facilitates identification of the relationship-specific shocks. Allowing for negatively correlated shocks across relationships also allows for the possibility of *substitution effects*, which are ruled out in the homogeneous setting. If a firm’s demand behaviour is better expressed as a distribution of different shocks across lending relationships (and similarly for banks’ supply shocks), this distribution may be helpful to characterise heterogeneity in credit dynamics across time, macrofinancial conditions, and policy actions.

2.1 A new model for price and quantity

In this paper, we propose a new model that nests (1) as a special case. A key distinction of our approach is that we jointly study two simultaneously determined outcome variables – price and quantity, say – in this case the change in the interest rate paid on loans and the growth in loan volume. More generally,

$$\eta_{fb}^i = A_{i1}u_{fb}^d + A_{i2}u_{fb}^s, \quad i \in \{1, 2\}, \quad (3)$$

where demand shocks u_{fb}^d and supply shocks u_{fb}^s are mean-zero and correlated across f and b .² Note that this model nests the structure in (2) by letting

$$u_{fb}^d = (d_f + a_{fb}^d)/A_{i1} \quad (4)$$

$$u_{fb}^s = (s_b + a_{fb}^s)/A_{i2} \quad (5)$$

$$\epsilon_{fb} = a_{fb}^d/A_{i1} + a_{fb}^s/A_{i2}, \quad (6)$$

provided that the sequence d_f is uncorrelated with s_b conditional on D_{fb} . The mapping between the two models is not unique; there are generally infinitely many ways to decompose the i.i.d. ϵ_{fb} into uncorrelated components a_{fb}^d, a_{fb}^s such that the structure in (3) holds. Thus, one way to view the present model is as a generalisation of the random effects version of (2), where ϵ_{fb} can be decomposed into two components, one of which is associated structurally with supply, and the other with demand, so there are *relationship-specific* random effects, $d_f + a_{fb}^d$ and $d_f + a_{fb}^s$, and no additional error term.

Stacking the equations in (3), we will study firm-bank observations η_{fb} as a linear com-

²Stronger assumptions are introduced in Assumption 3 in order to derive asymptotic properties.

bination of firm-bank specific demand and supply shocks:

$$\eta_{fb} \equiv \begin{pmatrix} \Delta r_{fb} \\ \Delta l_{fb} \end{pmatrix} = A \begin{pmatrix} u_{fb}^d \\ u_{fb}^s \end{pmatrix} = Au_{fb}. \quad (7)$$

Only mild assumptions are required for identification. Let \bar{D} in a conditional expectation denote the remainder of D not specified by the other conditions; for instance, in $E[\cdot | D_{fb} = 1, \bar{D}]$, $\bar{D} = D_{-fb}$.

Assumption 1. *The following hold:*

1. *A is invertible and constant across firm-bank pairs,*
2. $E[u_{fb}|D_{fb} = 1, \bar{D}] = 0$,
3. $E[u_{fb}^d u_{fb'}^s | D_{fb} = 1, D_{fb'} = 1, \bar{D}] = 0$, $E[u_{fb}^d u_{f'b}^s | D_{fb} = 1, D_{f'b} = 1, \bar{D}] = 0$, $b' \neq b$, $f' \neq f$.

A encodes the elasticity of both loan size and interest rates to demand and supply shocks. Second, the shocks are mean zero. The third point requires that demand shocks from firm f to one bank b are orthogonal to supply shocks to that firm f from a different bank b' ; demand shocks from one firm f to a bank b are orthogonal to supply shocks from that bank to a different firm f' . Crucially, we do not need to make any assumptions about network formation; the stipulated points must hold in the observed sample (that for which $D_{fb} = 1$), possibly after controlling for some covariates or demeaning (see Section 3.3), but we do not need to impose any form of exogeneity. This is because the central requirement for identification is the structure of the identifying equations (e.g., (12) below) with respect to A , rather than the precise values of the non-zero moments of the shocks. Demeaning may induce an intercept shift in all shocks for a given time period, but this is analogous to selecting the reference unit in a fixed effects model. In what follows, unless otherwise noted, all expectations and covariances should be understood to be taken conditional on D , as in Assumption 1; however, we suppress the conditioning to avoid notational clutter.

To identify (7), we use two sets of covariance equations. The first is

$$\text{cov}(\eta_{fb}, \eta_{f'b}) = \begin{bmatrix} A_{11}^2 E[u_{fb}^d u_{f'b}^d] + A_{12}^2 E[u_{fb}^s u_{f'b}^s] & - \\ A_{11} A_{21} E[u_{fb}^d u_{f'b}^d] + A_{12} A_{22} E[u_{fb}^s u_{f'b}^s] & A_{21}^2 E[u_{fb}^d u_{f'b}^d] + A_{22}^2 E[u_{fb}^s u_{f'b}^s] \end{bmatrix} \quad (8)$$

the covariance of η_{fb} across firms, holding b fixed. This has the structure

$$\Sigma_{FF} \equiv \text{cov}(\eta_{fb}, \eta_{f'b}) = A \Lambda_{FF} A', \quad (9)$$

where

$$\Lambda_{FF} = \begin{bmatrix} E[u_{fb}^d u_{f'b}^d] & 0 \\ 0 & E[u_{fb}^s u_{f'b}^s] \end{bmatrix}. \quad (10)$$

The second set of equations is

$$\Sigma_{BB} \equiv \text{cov}(\eta_{fb}, \eta_{f'b}) = A \Lambda_{BB} A'. \quad (11)$$

Let θ denote the six free parameters between the entries of A and $\Lambda_{FF}, \Lambda_{BB}$; while these three objects contain 8 parameters, two of them must be fixed to impose a scale normalisation. For instance, we set $\Lambda_{FF} = I_2$. Together these covariances provide six equations,

$$m(\theta) = \begin{bmatrix} \text{vech}(\Sigma_{FF} - A \Lambda_{FF} A') \\ \text{vech}(\Sigma_{BB} - A \Lambda_{BB} A') \end{bmatrix} = 0. \quad (12)$$

Finally, the identification condition is given in Assumption 2:

Assumption 2. $\Lambda_{FF} \neq c \Lambda_{BB}$ for any scalar c .

The following identification result is well-known for equations having the structure of (12) (e.g., Rigobon (2003)).

Proposition 1. *If Assumption 2 holds, then θ is the unique solution to (12) up to scale, sign, and column ordering.*

A is identified as long as the covariances between the two shocks across firms (holding bank fixed) are not exactly proportional to their covariances across banks (holding firm fixed). The result follows directly from linear algebra arguments; A is identified in closed form as the eigenvectors of $\text{cov}(\eta_{fb}, \eta_{f'b}) \text{cov}(\eta_{fb}, \eta_{f'b})^{-1} = A \Lambda_{FF} \Lambda_{BB}^{-1} A^{-1}$. Intuitively, separately exploiting heterogeneity within a firm but across banks and within a bank but across firms provides two sets of linearly independent moments in terms of supply and demand shocks, which are linked by A .

A itself is of interest, containing the elasticities determining how supply and demand shocks pass through to loan volume and interest rates. Indeed, A is sufficient to compute the slopes of the supply and demand curves in the loan market. From A , u_{fb} is additionally immediately recovered for each firm-bank pair.

3 Estimation and Asymptotic Properties

In this section, we introduce a simple estimator for A and establish its asymptotic properties.

3.1 Estimation

Estimation proceeds with the sample analogues of the covariances used for identification. While sample averages, the structure of these estimators is non-standard due to the type of covariances involved. In particular, the natural estimators for the two covariance matrices are

$$S_{FF} = \frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \eta_{fb} \eta'_{f'b} \quad (13)$$

$$S_{BB} = \frac{1}{N_{BB}} \sum_{f=1}^F \sum_{b' \neq b} \eta_{fb} \eta'_{fb'} \quad (14)$$

where $N_{FF} = \frac{1}{2} \sum_{b=1}^B F_b (F_b - 1)$, $N_{BB} = \frac{1}{2} \sum_{f=1}^F B_f (B_f - 1)$, and F_b is the number of firms connected to bank b and B_f the banks connected to firm f . Note that sums over the indices $f' \neq f$ and $b' \neq b$ indicate sums over all unique combinations of different indices, without duplication. We assume without loss of generality that the sample average of η_{fb} is zero; this can be achieved by de-meaning the data.

After estimating Σ_{FF} and Σ_{BB} using S_{FF} and S_{BB} , respectively, the sample analogue of $m(\theta)$ is

$$q(\boldsymbol{\eta}, \theta) = \begin{pmatrix} \text{vech}(S_{FF} - A\Lambda_{FF}A') \\ \text{vech}(S_{BB} - A\Lambda_{BB}A') \end{pmatrix}, \quad (15)$$

where $\boldsymbol{\eta}$ contains all observations η_{fb} , $b = 1, \dots, B$, $f = 1, \dots, F$. $\hat{\theta}$ is the minimum distance estimator solving

$$q(\boldsymbol{\eta}, \theta) = 0 \quad (16)$$

or equivalently minimising the just-identified minimum distance objective function

$$Q = q(\boldsymbol{\eta}, \theta)' q(\boldsymbol{\eta}, \theta). \quad (17)$$

$\hat{\theta}$ is available in close form. In particular, a solution for A in (16) that is unique up to scale, sign, and column order is

$$\tilde{A} = \text{evec}(S_{FF} S_{BB}^{-1}), \quad (18)$$

where $\text{evec}(\cdot)$ denotes the matrix of left eigenvectors of its argument.

The column order of this initial estimate must be determined – which shock is the demand shock, and which supply – and a normalisation imposed. We select the column order that

minimises the Frobenius norm between \tilde{A} and $\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$, so that the first shock is a demand shock and the second shock is a supply shock. We finally normalise the resulting matrix such that covariance of both shocks across firms is unity, $\Lambda_{FF} = I_2$. The matrix resulting from these transformations is \hat{A} . Appendix A formalises this procedure. The value of Λ_{BB} consistent with this normalisation can be immediately obtained, and hence the full vector $\hat{\theta}$ consistent with the labeling and normalisation.

3.2 Asymptotic properties

While the identification argument needed to recover the supply and demand shocks holds under mild non-parametric assumptions, slightly stronger assumptions on the shocks facilitate asymptotic results. In particular, we assume an additive component structure for the shocks:

Assumption 3. *Demand and supply shocks have the structure*

$$u_{fb}^d = e_{fb}^d + v_{fb}^d \quad (19)$$

$$u_{fb}^s = e_{fb}^s + v_{fb}^s. \quad (20)$$

where $e_{fb}^i, i \in \{d, s\}$ is mean zero and independent of all innovations except for $e_{fb'}^i$ and $v_{fb}^i, i \in \{d, s\}$ is mean zero and independent of all innovations except for $v_{f'b}^i$. All innovations have strictly positive variance and finite eighth moments, and

$$\lim_{F, B \rightarrow \infty} \frac{B}{N_{FF}^2} \sum_{b=1}^B \text{var} \left(\sum_{f' \neq f} \text{vech}(v_{fb} v_{f'b}') \right) \text{ and } \lim_{F, B \rightarrow \infty} \frac{F}{N_{BB}^2} \sum_{f=1}^F \text{var} \left(\sum_{b' \neq b} \text{vech}(e_{fb} e_{fb'}') \right) \quad (21)$$

are symmetric positive definite, where e_{fb} and v_{fb} stack the bank and firm demand and supply components, respectively.

This structure is flexible enough to capture rich correlation patterns in supply and demand shocks. In particular, e_{fb}^d and e_{fb}^s are firm-specific demand and supply components, respectively, which may be correlated across banks (for a given firm), but are mutually independent and independent of bank-specific components and other firms' supply and demand components. Likewise, v_{fb}^d and v_{fb}^s are bank-specific supply and demand components that may be correlated across firms (for a given bank), but are mutually independent and independent of firm-specific components and other banks' supply and demand components. This structure allows for (positive or negative) correlation of both demand and supply shocks across banks, holding firm fixed, or across firms, holding bank fixed. With this flexible

structure in hand, we can derive the asymptotic distributions of the sample counterparts of the covariances required for identification and thus that of $\hat{\theta}$. The assumptions on the variances ensure that asymptotic variances of S_{FF} and S_{BB} are valid variance matrices and the moment assumptions ensure that they are consistently estimable.

We next impose assumptions on the network structure.

Assumption 4. *The following limits hold:*

1.

$$\lim_{F,B \rightarrow \infty} \frac{N}{FB} = \kappa \in (0, 1], \quad N \equiv \sum_{b=1}^B F_b = \sum_{f=1}^F B_f; \quad (22)$$

2.

$$\frac{B}{F^2} \rightarrow 0 \text{ as } F, B \rightarrow \infty; \quad (23)$$

3.

$$\frac{F}{B^2} \rightarrow 0 \text{ as } F, B \rightarrow \infty. \quad (24)$$

The first point of Assumption 4 is relatively mild; it guarantees that the density of connections available to the econometrician grows proportionally to B and F , but only for a non-vanishing share of firms and banks. The latter two do not impose that B and F increase at the same rate, but rather that neither grows to dominate the other.

Theorem 1. *Suppose $\theta_0 \in \text{interior}(\Theta)$, which is compact. After suitable normalisation and labeling of the columns of A , under Assumptions 1 to 4, $\hat{\theta} \xrightarrow{p} \theta$; $\tilde{\Phi}(\hat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, \mathbf{W})$, where*

$$\tilde{\Phi} = \left(\begin{bmatrix} \sqrt{B} & 0 \\ 0 & \sqrt{F} \end{bmatrix} \otimes I_3 \right) \Phi, \quad \Phi = E \left[\frac{\partial q(\mathbf{X}, \theta)}{\partial \theta'} \right]. \quad (25)$$

More practically, the asymptotic variance of the parameter estimates, $\hat{\theta}$, is given by $\Omega = \tilde{\Phi}^{-1} \mathbf{W} \tilde{\Phi}'^{-1}$. Theorem 1 shows that $\hat{\theta}$ is a consistent estimator and that a suitable rescaling has an asymptotically normal distribution. The proof of Theorem 1 is non-standard, since the asymptotic framework of Assumption 4 is novel.

As shown in the proof in Appendix B, \mathbf{W} is block diagonal, with the diagonal blocks consisting of matrices \mathbf{W}_{FF} and \mathbf{W}_{BB} . We now consider the problem of estimating \mathbf{W} . \mathbf{W}_{FF} can be estimated using

$$\hat{\mathbf{W}}_{FF} = \frac{B^2}{N_{FF}^2} \frac{1}{B} \sum_{b=1}^B \left(\sum_{f' \neq f} \text{vech}(\eta_{fb} \eta'_{f'b}) - \text{vech}(S_{FF}) \right) \left(\sum_{f' \neq f} \text{vech}(\eta_{fb} \eta'_{f'b}) - \text{vech}(S_{FF}) \right)' \quad (26)$$

This is essentially an unbalanced clustered variance formula (over banks) for the mean of $\text{vech}(\eta_{fb}\eta'_{f'b})$, and simply a rescaling of the sample variance of $\sum_{f' \neq f} \text{vech}(\eta_{fb}\eta'_{f'b}) - \text{vech}(S_{FF})$ across $b = 1, \dots, B$ banks. Define $\hat{\mathbf{W}}_{BB}$ analogously and let $\hat{\mathbf{W}} = \begin{bmatrix} \hat{\mathbf{W}}_{FF} & 0 \\ 0 & \hat{\mathbf{W}}_{BB} \end{bmatrix}$.

Proposition 2. *Under Assumptions 1-4, $\hat{\mathbf{W}} \xrightarrow{p} \mathbf{W}$.*

Thus, \hat{W} is a consistent estimator for the asymptotic variance in Theorem 1, so we can conduct inference on parameters of interest in $\hat{\theta}$.

3.3 Including covariates

It is common in models of unobserved heterogeneity to include covariates, as in (1), whose coefficients, in this setting, can be viewed as nuisance parameters. A more general version of (7) is then

$$\eta_{fb} = \Gamma X_{fb} + A u_{fb}. \quad (27)$$

In Abowd et al. (1999) and other fixed and random effects settings, the coefficients on observed covariates are typically shown to be consistently estimable, so X_{fb} can be partialled out and analysis focus on the heterogeneous effects of interest. The same is true in this novel setting; Lemma 2 in the Appendix shows that Γ can be consistently estimated under Assumptions 1-4 and an additional assumption involving X_{fb} :

Assumption 5. *u_{fb} is mean independent of X_{fb} , given D ,*

$$E[u_{fb}|X_{fb}, D_{fb} = 1, \bar{D}] = E[u_{fb}|D_{fb} = 1, \bar{D}], \quad (28)$$

and X_{fb} is mean zero for observed relationships,

$$E[X_{fb}|D_{fb} = 1, \bar{D}] = 0. \quad (29)$$

This assumption ensures that X_{fb} can be treated as exogenous within the observed sample, while refraining from imposing any exogeneity assumptions on network formation with respect to u_{fb} . Note that the second part is without loss of generality since X_{fb} can always be demeaned within the observed sample. Since the coefficients on suitable covariates can be consistently estimated, the requirement under Assumption 1 that the shocks are mean zero is without loss of generality, since η_{fb} may be demeaned in a pre-processing step. Moreover, if there is a concern that the demand shocks issuing from a subset of firms in the same industry, for example, may be correlated, the data may be first demeaned by industry.

The inclusion of covariates also allows us to recast our model as a complement to procedures like Abowd et al. (1999), rather than an alternative. Indeed, consider the model

$$Y_{fb} = x_f^F + x_b^B + \eta_{fb}, \quad (30)$$

where Y_{fb} , x_f , and x_b (firm and bank fixed effects) are all 2×1 vectors, and η_{fb} is modeled as in (7). This is a standard Abowd et al. (1999) model, except we impose additional structure on the error term: that is, it is a linear combination of supply and demand shocks, that are not i.i.d., but rather correlated across firms and banks. This setup allows the researcher to first control for unobserved firm- and bank-specific characteristics *before* embarking on the structural decomposition this paper proposes. However, this strategy relies on there remaining non-negligible correlation across firms and banks in η_{fb} *after* extracting common variation through the fixed effects in order for the identification condition in Assumption 2 to hold. Moreover, in order for the \hat{A} to be consistent, the fixed effects $\{x_f^F\}$, $\{x_b^B\}$ must all be consistently estimated, which in general will require the degree of every firm and bank to diverge as $F, B \rightarrow \infty$, which is a significantly stronger assumption than Assumption 4. Nevertheless, it is an assumption made implicitly in much of the empirical literature, so may not be seen as prohibitive in certain settings.

3.4 Exploiting intertemporal variation

So far we have exploited just a single time period to identify and estimate A . However, in practice the variance of the shocks may be substantial and the number of observations in a single period may not be that large; for instance, the number of major banks in a particular country may be modest, or the number of multi-bank firms small. In these cases, pooling data over several time periods can provide a greater effective sample size and more precise estimates. Using multiple time periods – adding a third t subscript to $\eta_{fb,t}$ – does not impact the asymptotic results above: the average of T (possibly correlated) consistent and asymptotically normal estimators remains consistent and asymptotically normal. If the shocks are uncorrelated over time, as is typically assumed (in keeping with most theoretical notions of structural shocks) the original estimator of \hat{W} remains consistent for the asymptotic variance, treating units in each time period as additional banks/firms:

$$\hat{\mathbf{W}}_{FF,T}^{i.i.d.} = \frac{B^2}{N_{FF}^2 TB} \sum_{t=1}^T \sum_{b=1}^B \left(\sum_{f' \neq f} \text{vech}(\eta_{fbt} \eta'_{f'bt}) - \text{vech}(S_{FF,T}) \right) \left(\sum_{f' \neq f} \text{vech}(\eta_{fbt} \eta'_{f'bt}) - \text{vech}(S_{FF,T}) \right)' \quad (31)$$

However, if the shocks are serially correlated, an autocorrelation robust estimator is required. In particular,

$$\hat{\mathbf{W}}_{FF,T}^{rob} = \frac{B^2}{N_{FF}^2} \frac{1}{TB} \sum_{b=1}^B \left(\sum_{t=1}^T \sum_{f' \neq f} \text{vech}(\eta_{fbt} \eta'_{f'bt}) - \text{vech}(S_{FF,T}) \right) \left(\sum_{t=1}^T \sum_{f' \neq f} \text{vech}(\eta_{fbt} \eta'_{f'bt}) - \text{vech}(S_{FF,T}) \right)' \quad (32)$$

where the only difference is that the estimator now computes the (scaled) variance of B double sums over t and $f' \neq f$, instead of simply $f' \neq f$. Consistency of both follows by the same arguments as in Proposition 2.

3.5 Studying the shock distributions

When supply and demand shocks are estimated by fixed effects, as in, for instance, Khwaja and Mian (2008) and Amiti and Weinstein (2018), they are typically treated as consistently estimated objects to study, rather than generated regressors. Besides conceptual complications – since shocks are, after all, random variables – consistent estimation requires the number of connections or “degree” of each bank or firm to tend to infinity, holding network structure fixed, see for instance Jochmans and Weidner (2019). In our setting, in contrast, the *decomposition* between supply and demand shocks is consistently estimated under the asymptotic framework in Assumption 4, although relationship-specific shocks themselves cannot be consistently estimated since they are random variables, not parameters. However, since the shock estimates are a linear combination of observed innovations, they are unbiased and the estimates converge to the true random shocks as $F, B \rightarrow \infty$. As the degree of any firm or bank increases, the variation in its *average* estimated shocks also vanishes, in this sense recovering something comparable to the fixed effects estimates seen in the preceding literature. We document the performance of such “average heterogeneous shocks” in our simulation study.

Regressions involving the estimated shocks face a generated regressors problem. Proposition 4 and 5 in Appendix E jointly show that such feasible regression coefficients are consistent, derive their asymptotic distribution, and provide consistent estimators for their asymptotic variances. These variance estimators are linear combinations of standard (unbalanced) clustered variances. Thus, the estimated shocks may be used as dependent variables in regressions with only minor modification to standard practice. These results and a detailed discussion can be found in Appendix E.

4 Simulations

In this section, we demonstrate the performance of the proposed estimator. We also compare the estimator's ability to recover average supply and demand shocks to estimating (7) as a fixed effects model.

In these simulations, we set

$$A = \begin{bmatrix} 0.0761 & -0.0687 \\ 0.0124 & 0.0610 \end{bmatrix}. \quad (33)$$

which corresponds to empirical estimates for Italy in the monetary policy tightening subsample. We also match the moments of the supply and demand shocks to those estimated in that specification. We consider both i.i.d. shocks and shocks which may have some persistence in firm and bank components. For i.i.d. shocks, we generate the shocks according to

$$u_{fb}^i = z_f^i + z_b^i + z_{fb}^i, \quad i = \{d, s\}, \quad f = 1, \dots, F, \quad b = 1, \dots, B, \quad (34)$$

where each of the components is independent and normally distributed with mean zero and empirically calibrated variance. For persistent shocks, we use the same structure as before, but now allow z_f^i and z_b^i to be independent mean-zero AR(1) processes, with autoregressive parameters matching the data and resulting shock variances equal to those in the i.i.d. case.

We consider three different sample sizes: a small sample, with 10 banks, a small empirical sample, with 25 banks, a moderate empirical sample (comparable to that in most countries studied in Section 5) with 100 banks, and a large sample, with 500 banks. In each case, we set $F = 1000B$. We consider settings using either a single quarter ($T = 1$) or 4 quarters of data ($T = 4$). We draw relationship dummies uniformly to match the density of the sparse network structure seen in the data, subject to the constraint that each firm borrows from at least two banks. For each design, we draw 1000 Monte Carlo samples.

Table 1 reports the mean and standard deviation of the estimates of A . Starting with the $T = 1$ case, for $B = 10$, which is a smaller sample size than in almost all countries in our empirical application, the bias can be moderately large particularly for A_{21} . These results are potentially driven by the fact that, with sample sizes this small (with such a sparse network), complex values for \hat{A} are obtained in some samples, so that these statistics are computed across only those samples for which \hat{A} is real. However, as B increases to 25 and then 100, comparable to the order of magnitude observed in most countries, the bias decreases quickly for all parameters. The standard deviation falls steadily for all parameters. Finally, for $B = 500$, all biases are negligible and the standard deviations are vanishingly

Table 1: Bias and standard deviation of parameter estimates

	$B = 10$		$B = 25$		$B = 100$		$B = 500$	
$T = 1$, i.i.d.	bias	s.d.	bias	s.d.	bias	s.d.	bias	s.d.
A_{11}	-0.04	0.03	-0.02	0.02	-0.01	0.01	0.00	0.00
A_{21}	-0.32	0.02	-0.09	0.02	-0.02	0.01	-0.00	0.00
A_{12}	-0.13	0.03	-0.08	0.02	-0.02	0.01	-0.00	0.00
A_{22}	-0.10	0.01	-0.06	0.01	-0.01	0.00	-0.00	0.00
<hr/>								
$T = 4$, i.i.d.	bias	s.d.	bias	s.d.	bias	s.d.	bias	s.d.
A_{11}	-0.02	0.02	-0.01	0.01	-0.00	0.01	0.00	0.00
A_{21}	-0.05	0.02	0.03	0.01	-0.01	0.00	-0.01	0.00
A_{12}	-0.05	0.01	-0.01	0.01	-0.00	0.00	-0.00	0.00
A_{22}	-0.04	0.01	-0.01	0.00	-0.00	0.00	-0.00	0.00
<hr/>								
$T = 4$, dep.	bias	s.d.	bias	s.d.	bias	s.d.	bias	s.d.
A_{11}	-0.02	0.02	-0.01	0.01	-0.00	0.01	0.00	0.00
A_{21}	-0.05	0.02	0.03	0.01	-0.01	0.00	-0.01	0.00
A_{12}	-0.05	0.01	-0.01	0.01	-0.00	0.00	-0.00	0.00
A_{22}	-0.04	0.01	-0.01	0.00	-0.00	0.00	-0.00	0.00

Notes: The table reports the bias (in relative terms) and standard deviation of estimates of each entry in A across 1000 Monte Carlo samples generated according to the DGPs described in the text. The top panel considers three sample sizes for $T = 1$, with shocks independent over time: $B = 10, 25, 100, 500$, and $F = 1000B$ in each case. The second panel considers $T = 4$ time periods, where the shocks are independent in each, and the final considers $T = 4$, where the shocks are dependent over time.

small as well. As T increases to 4, for i.i.d. data the bias falls substantially, essentially vanishing for $B = 100$, as do the standard deviations. When the data is instead dependent, the results are identical to the independent case up to two digits, due to the low level of presence detected in the data. The results make a clear case for pooling data across multiple time periods due to the reduction in bias, as we opt to do in the simulations. Overall, the estimation strategy performs extremely well in realistic sample sizes. In empirical terms, this potentially casts doubt on the empirical results for Ireland, where the sample size is smallest, but suggests results are likely reliable for other countries studied.

Table 2 presents the empirical size of nominal 5% t -tests for each parameter in the same specifications. The size distortions are small for all parameters when $B = 100$ or larger, and moderate for smaller sample sizes. When $B \geq 25$, for the majority of parameters and specifications, empirical sizes range from 4-6%. Otherwise, there are few clear patterns in the results to point to. Note that there is little difference between the robust and non-robust tests here because the empirically-calibrated persistence is small and negative. When shocks instead have positive persistence, as in additional calibrations considered, the differences in size control can be much more pronounced.

Table 2: Empirical size of t -tests

	$B = 10$		$B = 25$		$B = 100$		$B = 500$	
$T = 1$, i.i.d.	SU	robust	SU	robust	SU	robust	SU	robust
A_{11}	10.7	10.7	8.1	8.1	6.0	6.0	5.7	5.7
A_{21}	10.7	10.7	9.1	9.1	6.5	6.5	5.4	5.4
A_{12}	15.4	15.4	8.9	8.9	4.9	4.9	5.3	5.3
A_{22}	18.5	18.5	14.2	14.2	5.9	5.9	5.0	5.0
$T = 4$, i.i.d.	SU	robust	SU	robust	SU	robust	SU	robust
A_{11}	5.1	7.5	4.7	6.0	5.8	6.1	6.1	6.2
A_{21}	5.6	7.9	4.4	5.2	5.2	5.3	6.6	6.8
A_{12}	7.6	8.3	5.9	6.1	5.6	6.0	5.4	5.0
A_{22}	11.3	12.7	6.0	7.1	5.3	6.1	4.9	5.1
$T = 4$, dep.	SU	robust	SU	robust	SU	robust	SU	robust
A_{11}	5.3	7.4	4.5	6.1	6.0	6.1	6.2	6.2
A_{21}	5.8	7.6	4.7	5.1	4.9	5.3	6.9	6.6
A_{12}	7.3	8.3	5.5	6.1	5.6	5.9	5.6	5.5
A_{22}	11.5	13.4	6.6	7.2	6.2	6.2	5.1	5.5

Notes: The table reports the empirical rejection rates of 5% t -tests for each entry in A across 1000 Monte Carlo samples generated according to the DGPs described in the text. For each specification, rejection rates are reported using the baseline variance estimator that assumes serially uncorrelated shocks (“SU”) as well as the robust variance estimator that allows for dependence (“robust”). The top panel considers three sample sizes for $T = 1$, with shocks independent over time: $B = 10, 25, 100, 500$, and $F = 1000B$ in each case. The second panel considers $T = 4$ time periods, where the shocks are independent in each, and the final considers $T = 4$, where the shocks are dependent over time.

Table 3 considers the problem of estimating the average demand shock for each firm and supply shock for each bank. Given the DGP described, a fixed effects estimator is well-specified for this object, but the average heterogeneous shock is as well. In particular, the “demand” firm fixed effects are z_f^d and the “supply” bank fixed effects are z_b^s , and we have $E\left[B^{-1} \sum_{b=1}^B u_{fb}^d | z_f^d\right] = z_f^d$ and $E\left[F^{-1} \sum_{f=1}^F u_{fb}^s | z_b^s\right] = z_b^s$. For each Monte Carlo sample, we compute the correlation of both the average heterogeneous shocks for each firm (bank) with z_f^d (z_b^s) as well as the correlation of each estimated firm (bank) fixed effect with z_f^d (z_b^s). We report the average correlation for each across samples. For all specifications and estimators, the correlations are much stronger for bank fixed effects than firm fixed effects, reflecting the large number of relationships for each bank relative to each firm. Across specifications, the firm average heterogeneous demand shock is substantially better correlated with the true firm fixed effect than the estimated fixed effect – by a factor of 2 to 3. As explained, this is because although the firm average can only make use of a firm’s limited relationships, the new estimator leverages information across the entire dataset to separate supply and

Table 3: Correlations of average shocks and fixed effects

	$B = 10$		$B = 25$		$B = 100$		$B = 500$	
$T = 1$, i.i.d.	avg. het.	FE						
Firm	0.11	0.05	0.11	0.04	0.11	0.04	0.21	0.07
Bank	0.90	0.97	0.94	0.97	0.98	0.97	1.00	0.98
$T = 4$, i.i.d.	avg. het.	FE						
Firm	0.11	0.04	0.11	0.04	0.11	0.04	0.21	0.07
Bank	0.96	0.97	0.98	0.97	0.99	0.97	1.00	0.98
$T = 4$, dep.	avg. het.	FE						
Firm	0.11	0.04	0.11	0.04	0.11	0.04	0.21	0.07
Bank	0.96	0.97	0.98	0.97	0.99	0.97	1.00	0.98

Notes: The table reports the average correlations (across 1000 Monte Carlo samples) of the average estimated heterogeneous demand and supply shocks (“avg. het.”) and the true “fixed effects” as well as the average correlations of the estimated fixed effects (“FE”) with the true fixed effects in the DGP described in the text. For each firm, we consider the average demand shock across banks and for each bank we consider the average bank supply shock across firms. The top panel considers three sample sizes for $T = 1$, with shocks independent over time: $B = 10, 25, 100, 500$, and $F = 1000B$ in each case. The second panel considers $T = 4$ time periods, where the shocks are independent in each, and the final considers $T = 4$, where the shocks are dependent over time.

demand factors. For all $B = 10$ specifications and $B = 25$ with $T = 1$, the estimated bank fixed effects have slightly higher correlations with the true bank fixed effects than the bank average supply shocks. However, as more observations become available (through more firms/banks and/or time periods) the average bank supply shocks become more accurate, as the advantage to exploiting information from the whole dataset grows more pronounced.

5 Supply and Demand in European Credit Markets

5.1 Data

We use our new methodology to estimate supply and demand dynamics in European credit markets using the ECB’s AnaCredit dataset. The AnaCredit dataset is a new harmonised euro area loan-level credit registry developed by the Eurosystem that contains the universe of loans to corporations in the Euro area over €25,000. Loans are reported at a monthly frequency from September 2018. The dataset contains many types of loans, including term loans, financial leases, trade receivables, credit lines, revolving credit, and syndicated loans. For a detailed discussion of the AnaCredit dataset, see the recent paper by Kosekova et al. (forthcoming). For more technical documentation of the data included, see the AnaCredit

Manual, European Central Bank (2019).

We focus our analysis on the 11 largest euro area countries, following Kosekova et al. (forthcoming). These are Austria (AT), Belgium (BE), Germany (DE), Spain (ES), Finland (FI), France (FR), Greece (GR), Ireland (IE), Italy (IT), the Netherlands (NL), and Portugal (PT). Our sample spans January 2019 to December 2023. We collapse the monthly data to quarterly frequency since we study changes in the loan quantities and interest rates, and there are few such changes at monthly frequency. We limit our focus to three types of credit: term loans, credit lines, and revolving credit. Like Kosekova et al. (forthcoming), we do not consider syndicated loans, since they are often not fully documented in AnaCredit. Given the nature of our identification approach, we limit our sample to firms that have borrowing relationships with multiple banks included in the dataset.

In each quarter, we compute the total committed credit for each firm and bank relationship over all instruments amongst the three types considered: term loans, credit lines, and revolving credit. We measure the the change in loan quantity as the quarter-on-quarter change in this total committed. We choose to focus on changes in committed quantities instead of utilised quantities since the former represent the outcome of interaction between both firm and bank, whereas the latter can be viewed as simply demand driven changes, given previously committed quantities. We convert the quarterly change to a quarterly growth rate using the formula

$$\Delta l_{fb,t} = \frac{l_{fb,t} - l_{fb,t-1}}{0.5l_{fb,t} + 0.5l_{fb,t-1}}. \quad (35)$$

We compute the prevailing interest rate for a relationship using the volume-weighted average across all instruments of the three types studied. Where the rate is indexed to a reference rate, we do so using the prevailing value of that reference rate. The quarterly change in the interest rate, $\Delta r_{fb,t}$ is then the quarter-on-quarter change in this weighted average interest rate. The object of interest for our analysis is then

$$\eta_{fb,t} = \begin{pmatrix} \Delta r_{fb,t} \\ \Delta l_{fb,t} \end{pmatrix}. \quad (36)$$

We drop observations for which the $\eta_{fb,t} = 0$; that is, observations for which a relationship exists between firm f and bank b , but for which there is no change at time t . We winsorise the data at the 1% level. We also address occasional data quality issues, for instance inconsistency in the units used to report interest rates. Taking this final sample for each country, as a final pre-processing step we demean $\eta_{fb,t}$ for each quarter. A full description of data cleaning and construction can be found in the Online Appendix.

We opt to estimate supply and demand dynamics for three 6-quarter windows for three reasons. First, given we study 11 different countries, presenting quarterly results would become unmanageable, while using a single pooled sample from 2019-2023 would make the assumption of a common A matrix non-credible due to dramatic changes in credit markets during this period. The 6-quarter approach strikes a balance between the two. Second, and relatedly, three 6-quarter windows cleanly capture three distinct economic periods for credit markets. The first sample, 2019Q3-2020Q4, spans the lead-up to and pandemic period, which featured severe disruptions and emergency support measures. The second sample, 2021Q1-2022Q2, covers the subsequent inflationary episode, up until the point of monetary policy intervention. Finally, 2022Q3-2023Q4 was a period of contractionary monetary policy and credit tightening. It is reasonable to believe that the elasticities in A , *inter alia*, the supply and demand curves of national credit markets, changed from each of these periods to the next along with macroeconomic and financial conditions. Third, while some of the countries we study feature very large numbers of banks and multi-bank firms, smaller countries and those with more consolidated lending relationships benefit from the increased sample size associated with pooling 6 quarters of data. Our results are quite stable with respect to the precise timing of the subsamples. We have experimented with both 4- and 8-quarter windows loosely aligning with the same major economic events as well as rolling windows, and the magnitudes of \hat{A} remain largely unchanged. These results are available upon request.

Table 9 in the Appendix reports summary statistics for each country, including the total number of firms and banks and total number of relationships in each of the three subsamples described above. Note that not every relationship may have a non-zero observation in a given quarter. The table illustrates that, as documented by Kosekova et al. ([forthcoming](#)), there are stark differences in the structure of credit markets across the countries studied. Indeed, particularly comparing to the total number of firms in the AnaCredit dataset in each country, it is clear multi-bank firms are quite common in some countries, like Italy and Spain (each with well over 100,000), but relatively rare in others, like Ireland and the Netherlands (200-650 and under 1,700, respectively). The number of banks also varies starkly, ranging from 9-16 for Greece and Ireland to around 800 for Germany. As a result, the effective sample size is quite variable across the countries considered, ranging from about 400 relationships per quarter in Ireland for two of the subsamples to about 5-600,000 per quarter in Italy.

5.2 Price and quantity elasticities over time

The first set of results consists of the elasticities contained in \hat{A} over the three subsamples considered. In particular, these are the elasticities of price with respect to demand shocks,

price with respect to supply shocks, quantity with respect to demand shocks, and quantity with respect to supply shocks. They are reported in Figures 1 and 2, respectively. The x -axis indicates the three different subsamples, and the y -axis indicates the elasticity for each of those samples. We report 95% confidence intervals using our robust variance estimator.

The first reassuring feature of the results is that the sign pattern of the elasticities is consistent with identifying both a demand shock and a supply shock. As with any statistical identification approach, there is no guarantee that this is the case. However, we see that for virtually every country-subsample pair, we recover one shock that positively impacts both price and quantity (demand), and one that negatively impacts price and positively impacts quantity (supply). While this is very occasionally violated – and only during the turbulent and likely noisy pandemic and inflationary periods – none of these violations are statistically significant. Thus, we can have confidence that our statistical identification strategy does indeed recover supply and demand shocks.

There is considerable variation in the scale of the elasticities both over countries and time. For the most part, the quantity elasticities with respect to both shocks are somewhat smaller in magnitude than those of price; in general, interest rates seem to be the main margin of adjustment in response to supply and demand shocks. The elasticity of price with respect to demand (Figure 1a) is statistically significant for 5 countries during the pandemic, 6 during inflation, and 8 during tightening. During the pandemic, it is largest for Greece, Ireland, Italy, and the Netherlands (1.34, 8.67, 1.24, 1.21) and smallest for Austria and Finland (0.05, 0.02). During inflation, it is largest for Italy, the Netherlands, and Portugal (1.19, 1.27, 2.45) and smallest for Austria and Belgium (0.23, -0.19). During tightening, it is largest for Belgium, Finland, and the Netherlands (1.09, 1.04, 1.04) and smallest for Spain and Portugal (0.13, 0.03).

The elasticity of price with respect to supply (Figure 1b) is statistically significant for 10 countries during the pandemic, 7 during inflation, and 6 during tightening. During the pandemic, it is largest for Belgium, Germany, the Netherlands, and Portugal (-1.22, -2.42, -1.48, -1.41) and smallest for Spain and France (-0.20, -0.20). During inflation, it is largest for Belgium, Germany, and Ireland (-1.50, -1.73, -2.50) and smallest for Spain and France (-0.04, -0.04). During tightening, it is largest for Germany and Portugal (-1.96, -1.12) and smallest for France (-0.17).

The elasticity of quantity with respect to demand (Figure 2a) is statistically significant for 6 countries during the pandemic, 4 during inflation, and 5 during tightening. During the pandemic, it is largest for Austria, Belgium, Germany, Finland, and Portugal (0.37, 0.51, 0.38, 0.35, 0.36) and smallest for Spain, France, and Greece (0.00, 0.01, -0.01). During inflation, it is largest for the Belgium and Ireland (0.80, 0.67) and smallest for Spain, France,

and the Netherlands (-0.00, 0.01, 0.02). During tightening, it is largest for Belgium, Germany, Spain, Ireland, and Portugal (0.25, 0.25, 0.22, 0.48, 0.24) and smallest for Austria, Finland, and the Netherlands (0.01, 0.00, 0.02).

The elasticity of quantity with respect to supply (Figure 2b) is statistically significant for 6 countries during the pandemic, 7 during inflation, and 7 during tightening. During the pandemic, it is largest for France and Greece (0.43, 0.48) and smallest for Belgium, Germany, and Finland (0.03, 0.09, 0.06). During inflation, it is largest for Belgium (0.69) and smallest for Germany and Finland (0.07, 0.02). During tightening, it is largest for Belgium, France, and Ireland (0.33, 0.29, 0.40) and smallest for Germany, Spain, and Portugal (0.02, 0.02, 0.01).

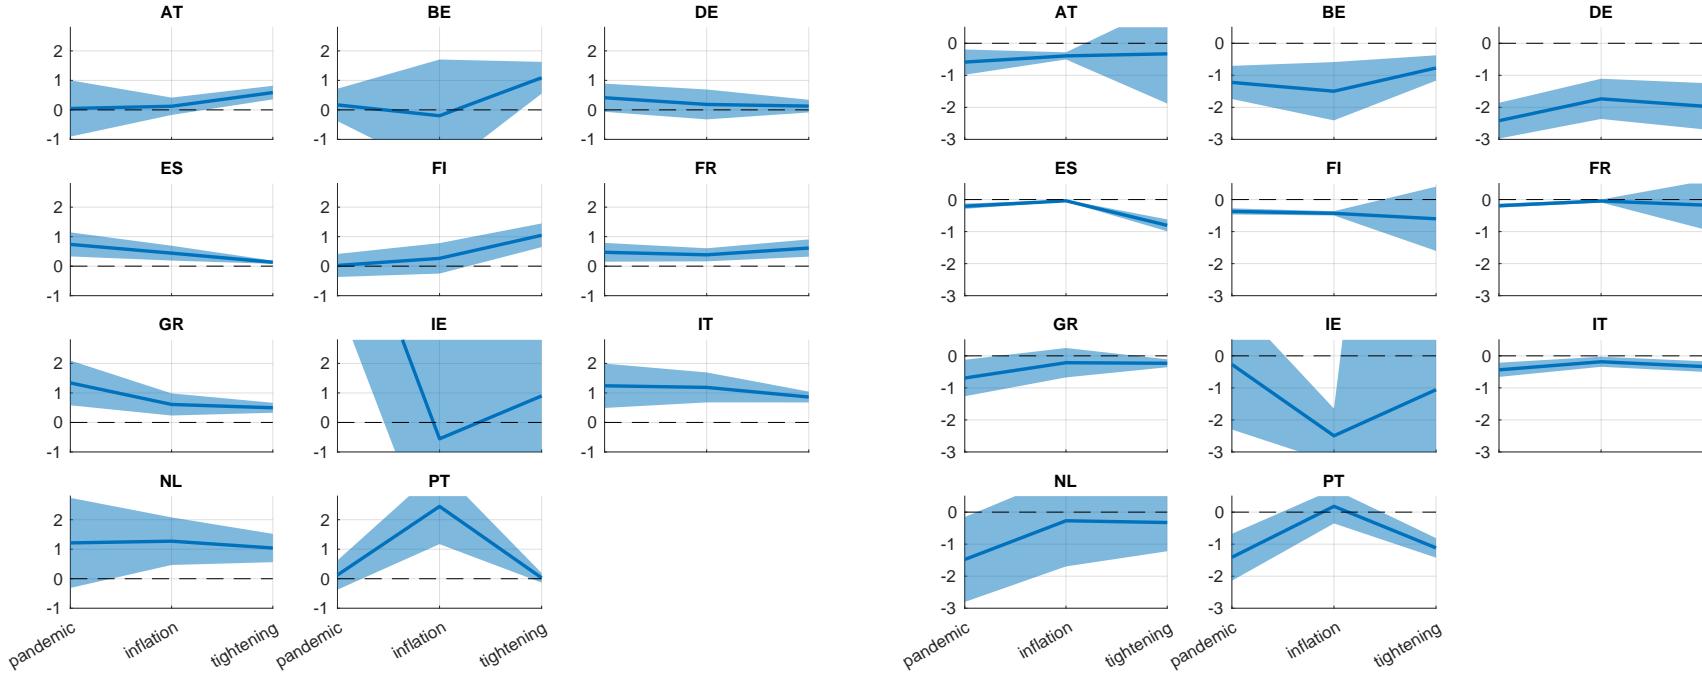
We also observe similar dynamics for certain elasticities amongst certain families of countries. The price elasticity with respect to demand is stable before rising significantly during tightening in Austria, Belgium, and Finland, while it peaks prominently during the inflationary episode in Portugal, and to a much lesser degree in the Netherlands. Germany, Spain, Greece, and Italy experience gradual declines in this elasticity. The elasticity of price with respect to supply displays a prominent peak during inflation in Italy and Portugal, and to a much lesser extent in Germany, France, Spain, Greece, and the Netherlands. Belgium and Ireland show a modest decline in this elasticity during tightening. The elasticity of quantity with respect to demand shows modest increases during the tightening episode in Spain, France, and Italy, with gradual declines in Austria, Belgium, Germany, Finland, and Portugal. The quantity elasticity with respect to supply peaks during tightening in Belgium the Netherlands, and Portugal, with the opposite pattern in Finland. There are gradual declines in Germany, Spain, and Greece.

Figure 1: Elasticity of Price with Respect to Demand and Supply

(a) Elasticity of price with respect to demand shocks, \hat{A}_{11} ,

(b) Elasticity of price with respect to supply shocks, \hat{A}_{12}

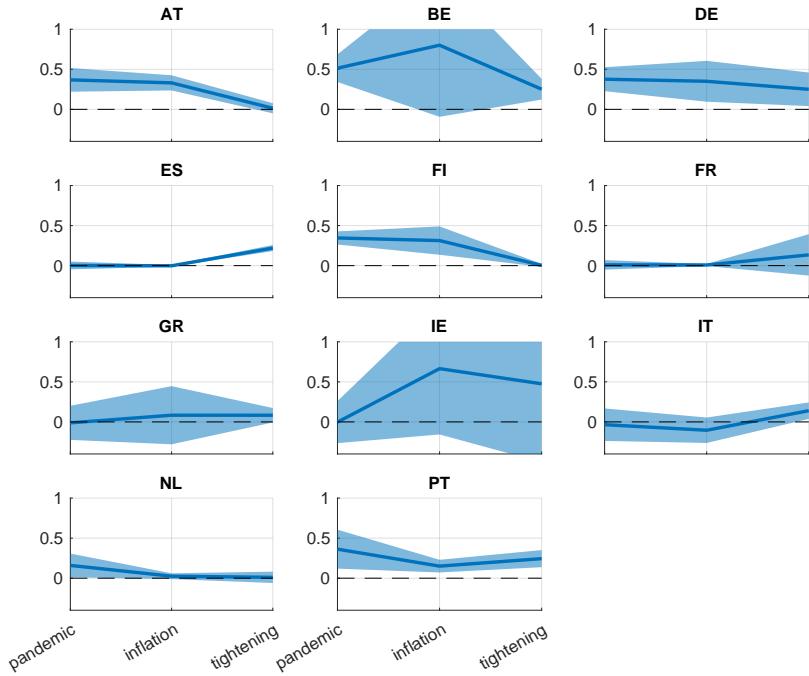
25



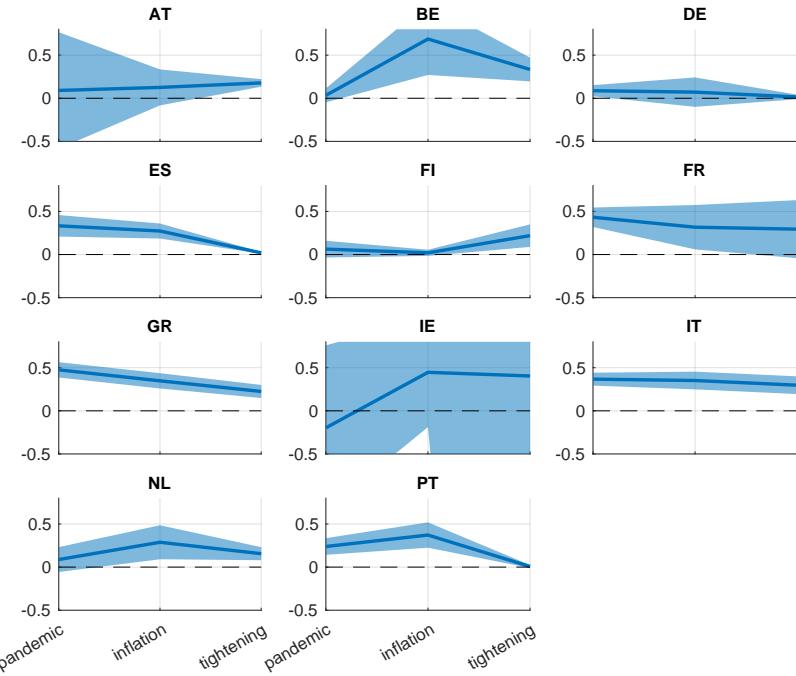
Notes: Panel (a) plots the elasticity of price with respect to credit demand shocks, \hat{A}_{11} . Panel (b) plots the elasticity of price with respect to credit supply shocks, \hat{A}_{12} . For each country, the figure plots the elasticities across the three 6-quarter sub-samples described in the text, which we label as “pandemic”, “inflation”, and “tightening”. The estimates are normalised so that both demand and supply shocks have unit variance in each country in each subsample. The confidence intervals are 95% intervals based on the robust variance estimator.

Figure 2: Elasticity of Quantity with Respect to Demand and Supply

(a) Elasticity of quantity with respect to demand shocks, \hat{A}_{21} ,



(b) Elasticity of quantity with respect to supply shocks, \hat{A}_{22}



Notes: Panel (a) plots the elasticity of quantity with respect to credit demand shocks, \hat{A}_{21} . Panel (b) plots the elasticity of quantity with respect to credit supply shocks, \hat{A}_{22} . For each country, the figure plots the elasticities across the three 6-quarter sub-samples described in the text, which we label as “pandemic”, “inflation”, and “tightening”. The estimates are normalised so that both demand and supply shocks have unit variance in each country in each subsample. The confidence intervals are 95% intervals based on the robust variance estimator.

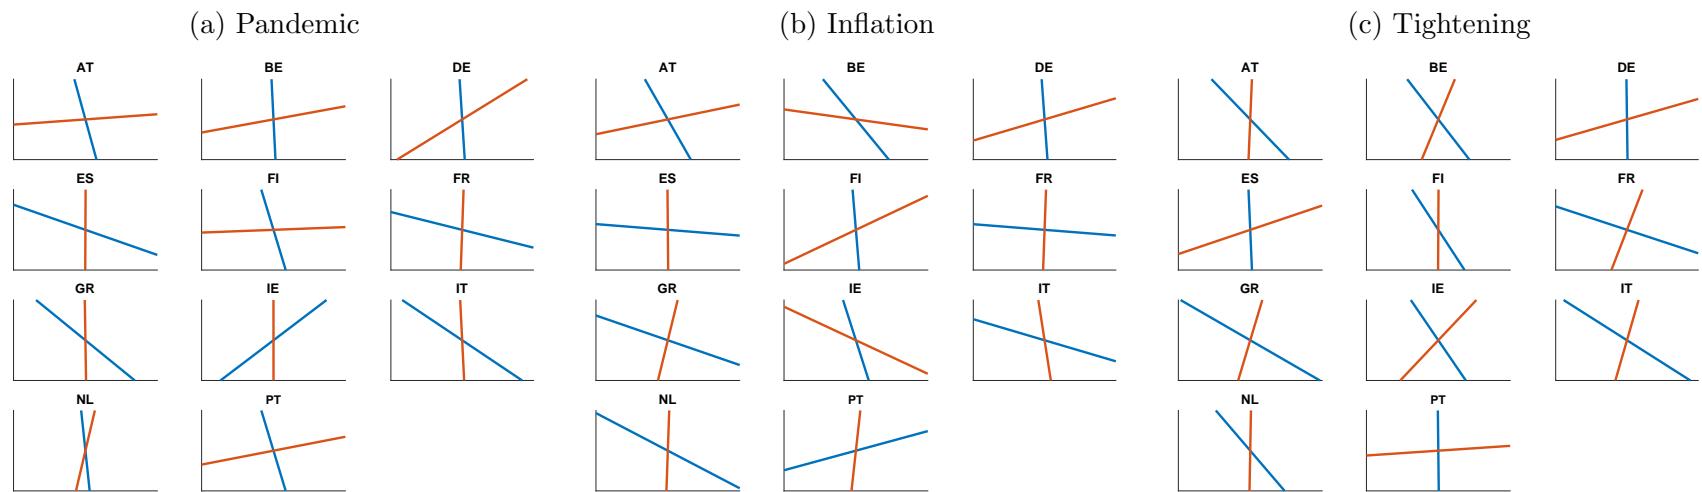
5.3 Supply and demand curves over time

The entries of the matrix A represent the elasticities of both price and quantity with respect to the shocks of demand and supply. As a result, they contain all the information needed to identify the slope of both the supply and demand curves. In particular, the slope of the supply curve is given by A_{11}/A_{21} – that is, it is identified by the relative movements of price and quantity in response to an identified demand shock. Symmetrically, the slope of the demand curve is identified by A_{12}/A_{22} . As a result, we can identify both supply and demand curves for each country and subsample, which may be a more intuitive way to represent the results in the previous section. Figure 3 plots these curves for each country and sub-sample.

The figures show how the dynamics of supply and demand have varied over countries and over time. We omit confidence intervals to avoid crowding the figures; note that they tend to be quite wide for many of the slopes, since they are the ratios of estimated parameters, some of which were quite noisily estimated in the previous section. During the pandemic, Figure 3a, most countries exhibit at least one curve that is essentially vertical (inelastic). For Belgium, Germany, the Netherlands, and to a lesser extent Austria, Finland, and Portugal, demand is inelastic. This is consistent with firms demanding the funding required to survive such a major economic shock at whatever price it was available. For Spain, France, Greece, Ireland, Italy and, to a lesser extent, the Netherlands, the supply is inelastic. This is consistent with banks supplying credit at essentially fixed prices following significant intervention into credit markets by the ECB, for instance, the Pandemic Emergency Purchase Program.

During the post-pandemic inflationary episode, Figure 3b, inelastic curves are again present in most countries. Germany, Finland, and to a lesser extent, Ireland, exhibit inelastic demand. This is consistent with firms demanding credit to cover rising costs due to inflation at whatever price it was available. On the other hand, Spain, France, Greece, Italy, the Netherlands, and Portugal exhibit inelastic supply. This is consistent with banks being wary to issue term loans whose repayment may be depreciated by inflation. For some countries, these results mark a change from the pandemic. Portugal goes from having largely inelastic demand to inelastic supply. Belgium and the Netherlands see a flattening of demand. It is worthwhile to note that while Belgium, Ireland, and Italy appear to have downward-sloping supply curves, these slopes are very noisily estimated and not statistically significant.

Figure 3: Supply and Demand Curves across Countries and Time



Notes: For each country, each panel plots the estimated supply and demand curves for the indicated period, constructed from \hat{A} using the method described in the text. Confidence intervals are omitted so as not to crowd the figures.

Figure 3c presents results for the period of monetary policy tightening. Once more, inelastic curves remain common. Germany, Spain, and Portugal exhibit inelastic supply curves. This is consistent with firms demanding a relatively fixed amount of credit regardless of the increased funding costs. On the other hand, Austria, Finland, the Netherlands, and, to a lesser extent Belgium, France, Greece, and Italy exhibit inelastic supply curves. This is consistent with banks constraining credit in response to tighter conditions fostered by monetary policy, with interest rates serving to ration credit. Relative to the inflationary episode, Finland goes from having largely inelastic demand to inelastic supply, while the opposite is true for Spain and Portugal. Supply also dramatically steepens in Belgium.

Overall, these results demonstrate both the heterogeneity of credit markets across countries and time as well as remarkably similar patterns. While monetary policy was the same across the Euro-area, credit markets clearly behaved differently. Such results can be informative for evaluating the effects of potentially heterogeneous local policies as well as motivating such policies given the variation in credit dynamics.

6 What Explains Heterogeneity in Supply and Demand Shocks?

From the outset, one of the features of our identification strategy was relaxing the homogeneity assumption of supply and demand shocks, making an important advantage the ability to study the heterogeneous distributions of those shocks. In this section, we consider the distributions of supply and demand shocks over country, time, and observable relationship-specific, firm, and bank characteristics. This is a very rich source of empirical targets, so space constraints dictate that the present analysis only touches the surface of possible insights.

6.1 Within-firm/bank vs. between-firm/bank variation

We begin by recovering the identified shocks in each quarter by inverting the estimated \hat{A} matrix for the corresponding subsample and applying it to decompose the observed $\eta_{fb,t}$. Following the assumptions underlying many of the papers using credit registers, the within-firm dispersion in innovations should be small (homogeneous credit demand assumption). We thus explore to what extent fixed effects (capturing the “average shocks”) can explain the variation in supply and demand shocks. Figure 4 plots the adjusted R^2 over time after regressing the estimated supply and demand shocks on various combinations of fixed effects,

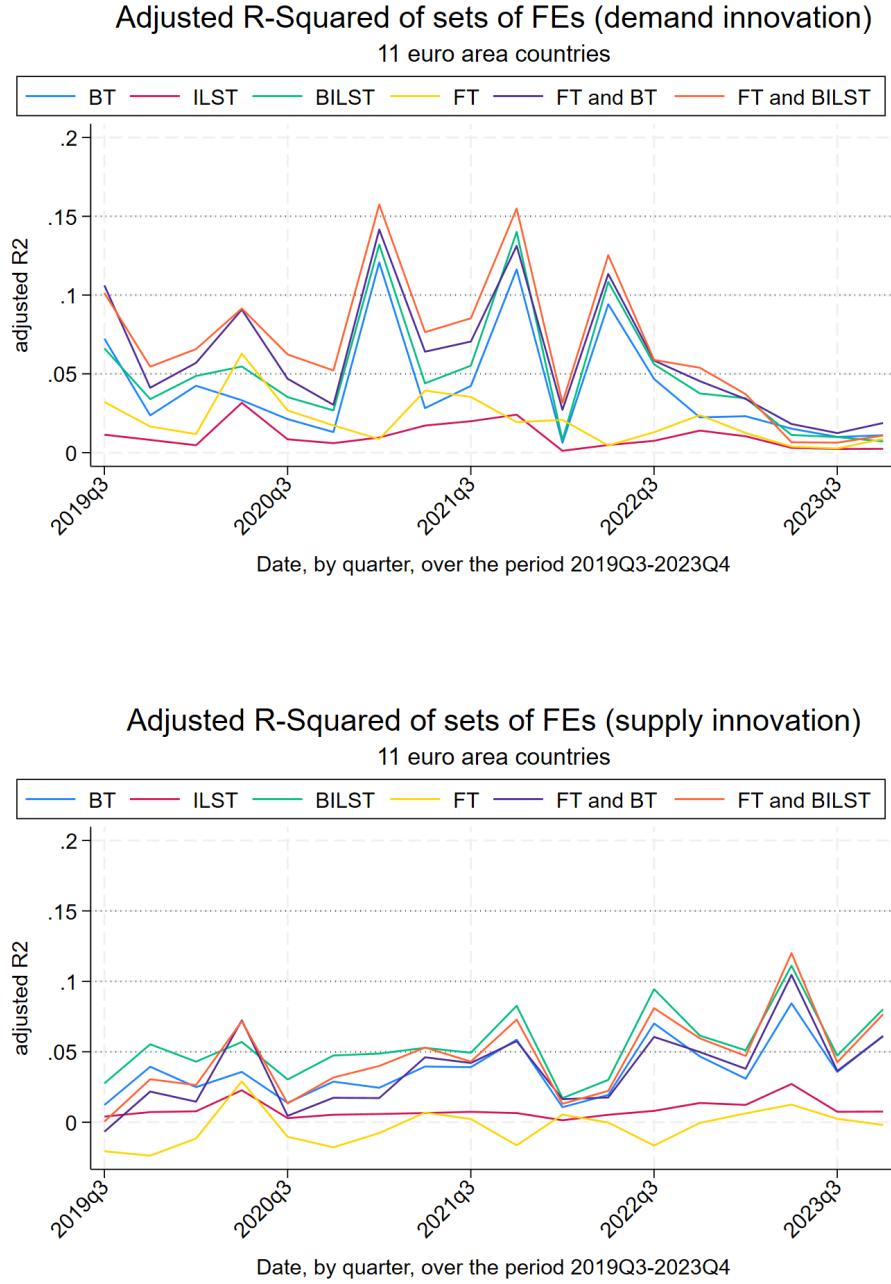
pooling all countries for ease of interpretation.³

We consider four types of fixed effects and combinations thereof: bank-time (BT), firm-time (FT), industry-location-sector-time (ILST), and bank-industry-location-size-time (BILST). The left panel starts with demand shocks. There are several notable features. First, the overall share of variation explained by even the richest combination of fixed effects (FT and BILST) generally explains less than 10% of variation in the estimated shocks. FT, which are typically used in the literature to capture demand shocks, almost always explain less than 5% of variation; the one exception is Q2 of 2020, the onset of the pandemic, which is the period where one would indeed expect a monolithic demand shock. The figures in the Appendix show that, of the four types of fixed effects, FT explain the most variation in unadjusted terms, but after adjusting for the very large number of fixed effects used, the adjusted value generally falls below that of BT and BILST. The right panel repeats the exercise for supply shocks. Once more, even the richest fixed effects generally fail to capture even 10% of variation in the supply shocks. Notably, FT are generally associated with a negative or essentially zero adjusted R^2 . On the other hand, BT explain more variation than either FT or ILST, consistent with bank-side factors being a more important determinant of supply than borrower characteristics.

Figure 5 conducts a similar exercise focusing instead on observables: the growth in loan quantity and the change in interest rates. In the typical empirical setting, the researcher only has access to the former, and estimates FT and BT fixed effects to measure supply and demand shocks using this data alone. Similar conclusions apply to this raw data, even if one was skeptical of our recovered shock distributions. Here, we find that FT explain essentially zero variation in loan quantity, after adjusting the R^2 . BT explain more variation, peaking around 5% during the inflationary episode, while interacting with ILST adds somewhat more information. Interestingly, fixed effects explain more variation in interest rates, with BT explaining as much as 15% in 2023, and FT and BILST together explaining as much as 25%. However, we reiterate that price data is not available in many credit datasets and is even more rarely studied. The trends in explanatory power are quite different for quantity compared to supply and demand. Studying the former would suggest that fixed effects peaked in efficacy during the inflationary episode, while the latter actually points to a gradual decline in explanatory power for demand shocks, despite some peaks during inflation. Crucially, studying loan quantity, which combines both supply and demand shocks, the only time that FT have positive explanatory power is Q2 of 2020, which captures the surge in demand for emergency credit at the onset of the pandemic; otherwise, the adjusted R^2 values

³Figure 6 in the Appendix additionally reports unadjusted R^2 values. Figure 7 in the Appendix collapses results into the three 6-quarter subsamples but breaks them down by country.

Figure 4: Variation in Demand and Supply shocks explained by Sets of Fixed Effects



Notes: This graph presents the adjusted R^2 values from regressing time-varying firm-bank specific demand and supply shocks on sets of fixed effects. Data is pooled over all countries. The fixed effects regressions are done separately for each quarter, over the period 2019Q3 to 2023Q4. The different lines depict adjusted R^2 values for various fixed effects combinations, including Bank-Time (BT), Industry-Location-Size-Time (ILST), Bank-Industry-Location-Size-Time (BILST), Firm-Time (FT), and two way fixed effects combining either FT and BT or FT and BILST in one set-up. The y-axis shows the adjusted R^2 values, with dotted reference lines at 0.2 intervals. The upper panel considers demand shocks and the lower supply.

are consistently negative.

Another way to compare within- and between- variation is to directly study the distribution of the average shocks collapsed at firm and bank levels. Table 4 computes summary statistics for the shocks, studying both between- and within- firm and bank variation. For demand, we consider the properties at firm-time level, and for supply at bank-time level, mirroring how each has been studied in the literature, using fixed effects. The first row shows that, after collapsing at firm-time level, demand shocks are essentially symmetrically distributed around zero. The final entry computes the between-firm-time standard deviation of demand shocks, 0.665. The next three rows compute various measures of the dispersion of the innovations. Of particular interest, the final row computes the distribution of the standard deviation of demand shocks within firm-time. The median such standard deviation is 0.291, over 40% of the within-firm standard deviation. This undermines the credibility of the homogeneity assumption; there is substantial dispersion of demand shocks within the firm at a given time. The second panel repeats the exercise for supply shocks at bank-time level. Here again, we see an essentially symmetric distribution, after collapsing at bank level, with a standard deviation of 0.453. However, the median within-bank-time standard deviation is higher at 0.803, almost 80% larger than the between bank-time variation. This suggests that the dispersion of supply shocks *within* a bank in a given period dominates the variation of the average supply shock across banks. Once more, this is evidence against the types of homogeneity assumptions commonly maintained in the literature. In contrast, it is clear support for the presence of important within-firm and within-bank heterogeneity.

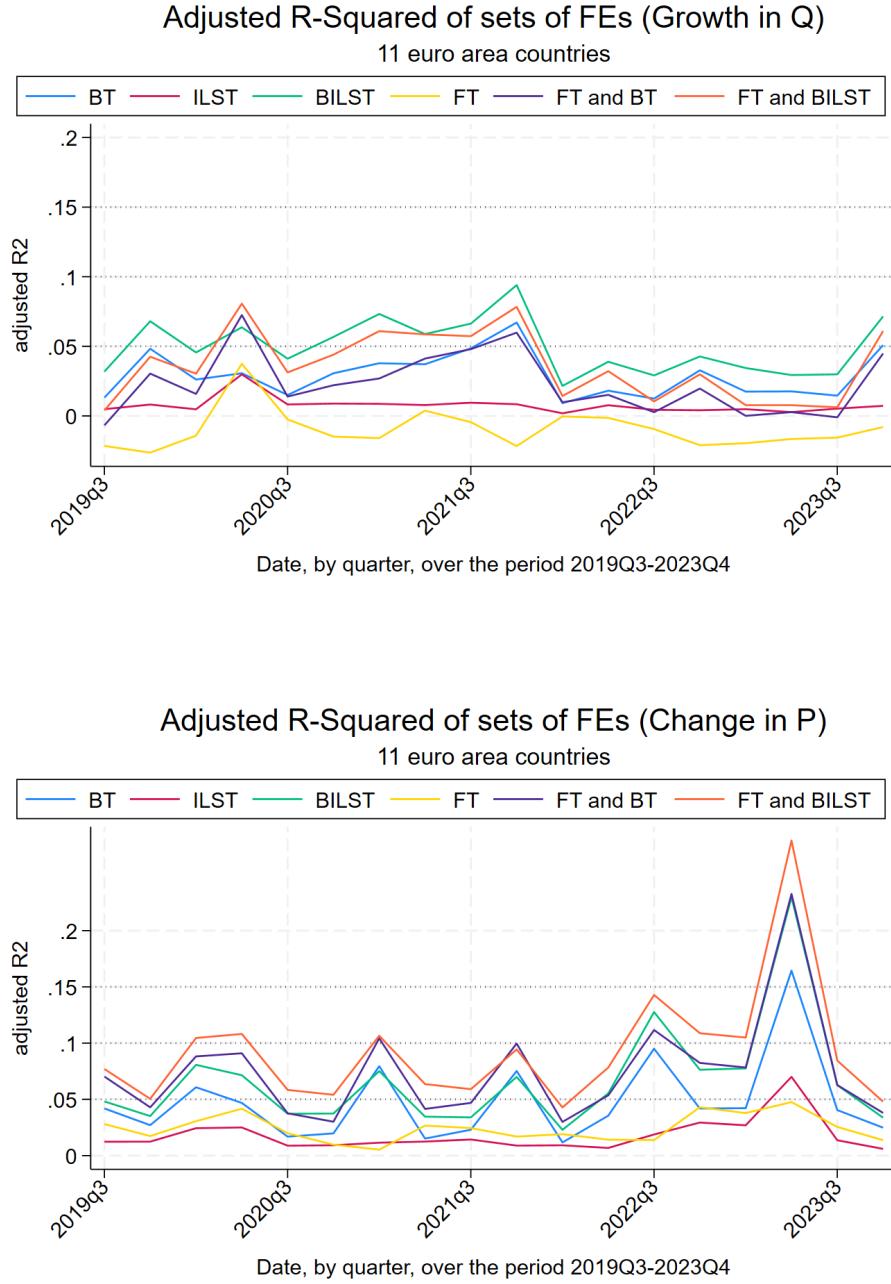
6.2 Observable characteristics predict shocks

Our methodology enables us to examine the drivers of heterogeneity in credit demand and credit supply. We can run analyses similar to a large body of empirical banking studies, but with estimates of time-varying firm-bank credit demand and supply as left-hand side variables. In this section, we illustrate this with a limited set of firm- and bank-specific characteristics as well as observable characteristics at the firm-bank relationship level. We will only consider characteristics that can be constructed using AnaCredit data.⁴ The goal here is not to be complete in the included characteristics, but to highlight the new insights that our approach can offer.

What we aim to show in these first analyses is that firm characteristics are also relevant for supply, and that bank characteristics matter for credit demand. Furthermore, we will illustrate that high-dimensional fixed effects may still have relevance in regressions of credit

⁴This approach ensures the analysis can be replicated without access to confidential regulatory bank data or commercial databases with firm balance sheet information.

Figure 5: Variation in Credit growth and Price changes explained by Sets of Fixed Effects



Notes: This graph presents the adjusted R^2 values from regressing Credit growth (upper) and Price changes (lower) on sets of fixed effects. The fixed effects regressions are done separately for each quarter, over the period 2019Q3 to 2023Q4. Data is pooled over 11 euro-area countries. The different lines depict adjusted R^2 values for various fixed effects combinations, including Bank-Time (BT), Industry-Location-Size-Time (ILST), Bank-Industry-Location-Size-Time (BILST), Firm-Time (FT), and two way fixed effects combining either FT and BT or FT and BILST in one set-up. The y-axis shows the adjusted R^2 values, with dotted reference lines at 0.05 intervals.

Table 4: Between and within variation

Collapse at the firm-time level							
	p10	p25	p50	p75	p90	IQR	StD
Number of banks	2.000	2.000	2.000	3.000	4.000	1.000	1.272
Average demand innovation	-0.636	-0.199	0.005	0.192	0.628	0.391	0.665
Range of demand innovation	0.021	0.110	0.472	1.370	2.899	1.260	1.324
IQR of demand innovation	0.019	0.097	0.393	1.134	2.450	1.037	1.180
Std dev demand innovation	0.014	0.071	0.291	0.811	1.699	0.739	0.792

Collapse at the bank-time level							
	p10	p25	p50	p75	p90	IQR	StD
Number of firms	3.000	9.000	30.000	107.000	440.000	98.000	2642.576
Average supply innovation	-0.360	-0.127	-0.002	0.135	0.377	0.262	0.453
Range of supply innovation	0.257	1.649	4.546	7.112	8.523	5.463	3.110
IQR of supply innovation	0.006	0.061	0.245	0.578	1.236	0.516	0.751
Std dev supply innovation	0.188	0.439	0.803	1.178	1.576	0.739	0.554

Notes: For the first panel, starting from the firm-bank-time level dataset, we collapse it to the firm-time level and compute four summary statistics for the time-varying innovations to firm-bank specific credit demand. We compute the average firm-time innovation, as well as three measures of within firm-time dispersion in the innovations (spread, interquartile range, and standard deviation). For each of these “collapsed” firm-time observations, we provide summary statistics to indicate the spread. The second panel repeats the exercise for supply innovations at bank level.

demand and credit supply. Finally, we introduce a number of bank-firm characteristics that received much less attention in the current theoretical and empirical banking literature (compared with bank and firm characteristics) and inspect how they relate to credit demand and credit supply.

We shed light on the above with six regressions with results reported in Table 5. Columns 1–3 of Table 5 explore demand shocks, while columns 4–6 focus on supply shocks. Essentially, these regressions examine whether the available characteristics within AnaCredit explain *between* variation in the demand and supply shocks, as well as within variation. The first block of variables are bank characteristics, the middle block are firm characteristics, while the bottom block are characteristics of the firm-bank credit relationship.

The first thing that stands out is that bank characteristics enter significantly in a regression framework explaining variation in credit demand. Similarly, firm characteristics are predictors of cross-sectional variation in credit supply. We find that the bank’s total assets and its credit-to-asset ratio are both significantly positive predictors of demand. Firm total assets, firm leverage, and firm age are all significantly negative predictors of credit demand. Younger and potentially rapidly expanding firms require more credit. A bank’s exposure to a specific sector or region affects credit supply significantly positively (suggesting a role for specialization). Banks supply significantly more credit to firms that are more dominant within their sector, but less to larger firms, more leveraged firms, and older firms. Our

structural decomposition thus enables researchers to address the concern that one otherwise has to remain largely agnostic about the economic origins of the shocks captured by high-dimensional fixed effects (Beaumont et al. 2019). ⁵

Across different columns, we use different sets of fixed effects. Columns 1 and 4 include time-fixed effects, while columns 2 and 5 substitute bank characteristics with bank-industry-location-quarter (BILT) fixed effects. Columns 3 and 6 replace firm characteristics with firm-quarter fixed effects. Bank characteristics still matter for credit demand, even after controlling for all unobserved time-varying firm heterogeneity (column 3). Similarly, firm characteristics matter for credit supply even when bank characteristics are replaced with granular fixed effects that are typically attributed to proxy for credit supply (column 6). Remarkably, there is an asymmetry in the stability of coefficients and significance levels in specifications with and without fixed effects. Bank characteristics are more stable in demand regressions than in supply regressions when fixed effects are added. Similarly, firm characteristics are more robust in supply regressions (columns 4-5) than in demand regressions (columns 1-2). This differential sensitivity suggests the critical role of unobserved heterogeneity in shaping demand and supply shocks. Also note that the (unadjusted) R^2 of the credit demand and credit supply regressions falls dramatically (from 0.40 to 0.00) as fixed effects are removed from the regressions, even as observable variables are added. This suggests that although they may not explain much variation in demand shocks (after adjustment), the fixed effects do still capture considerable unobserved heterogeneity that is not well-explained by the limited observable characteristics available within AnaCredit.

In the bottom panel, we introduce four characteristics that are specific to the firm-bank relationship. These four variables will be the focal point in the next subsection as they allow explaining the within firm and within bank heterogeneity in credit demand and credit supply. What we already surmise from the current regressions is that across specifications, the four relationship-specific characteristics remain consistent in sign and magnitude, underscoring their robustness. Across characteristics, we observe all possible combinations of signs and (in)significance on their impact on credit demand and credit supply. Some characteristics have a similar signed impact on both credit demand and supply, whereas others affect demand

⁵Beaumont et al. 2019 write the following: “the firm-time component may as well relate to demand (variations in input costs, productivity shocks) as to supply. In the event of a default on a trade bill, for instance, all banks might want to cut lending to the defaulting firm at the same time, leading to a low value of the firm-time fixed effect in a credit growth regressions. Similarly, a low bank-specific fixed effect may either indicate that the bank is cutting back lending due to funding constraints or that borrowers are all reducing their credit exposure to this particular bank.” We thus document that, on average during our sample period, bank characteristics indeed matter for credit demand, and firm characteristics also matter for credit supply. An interesting application for future research would be to look at typical shocks studied in the empirical banking literature (change in regulatory requirement, funding shocks, ...) and assess how these relationships behave during these specific episodes.

Table 5: Bank, firm and relationship characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	Demand innovation (f,b,t)			Supply innovation (f,b,t)		
ln(Total Assets - Bank) (b,t-1)	0.007*** (0.002)	0.008*** (0.002)	-0.002* (0.002)	-0.000 (0.002)		
Corporate lending to Total Assets - Bank (b,t-1)	0.037*** (0.010)	0.037*** (0.009)	-0.005 (0.011)	-0.006 (0.016)		
Bank Sectoral Market Share (f,b,t-1)	-0.021 (0.078)	-0.056 (0.058)	0.163** (0.067)	0.104 (0.120)		
Bank Sectoral Exposure (f,b,t-1)	0.037** (0.019)	0.124* (0.065)	0.088*** (0.032)	0.190*** (0.069)		
Bank Geographical Market Share (f,b,t-1)	-0.030 (0.030)	-0.052 (0.032)	-0.068* (0.038)	-0.049 (0.047)		
Bank Geographical Exposure (f,b,t-1)	-0.013 (0.013)	-0.014 (0.016)	0.046*** (0.014)	0.059*** (0.016)		
ln(Total Assets - Firm) (f,t-1)	-0.002 (0.002)	-0.004** (0.002)	-0.023*** (0.002)	-0.020*** (0.003)		
Bank credit to Total Assets - Firm (f,t-1)	-0.041*** (0.007)	-0.055*** (0.007)	-0.127*** (0.008)	-0.124*** (0.007)		
ln(Firm Age, in months) (f,t-1)	-0.018*** (0.003)	-0.014*** (0.003)	-0.023*** (0.004)	-0.031*** (0.003)		
Firm Geographical Market Share (f,t-1)	-0.148* (0.084)	0.033 (0.067)	0.619*** (0.074)	0.454*** (0.063)		
Firm Sectoral Market Share (f,t-1)	-0.101 (0.066)	-0.162** (0.068)	0.179** (0.084)	0.054 (0.068)		
Share of fixed rate loans (f,b,t-1)	-0.137*** (0.017)	-0.159*** (0.023)	-0.157*** (0.018)	0.139*** (0.014)	0.179*** (0.020)	0.118*** (0.015)
Share of collateralized loans (f,b,t-1)	-0.011 (0.014)	-0.011 (0.012)	0.021 (0.017)	0.037*** (0.012)	0.030*** (0.010)	0.039*** (0.013)
Share of Credit lines and Term Loans (f,b,t-1)	0.054*** (0.017)	0.059*** (0.020)	0.072*** (0.020)	-0.112*** (0.026)	-0.142*** (0.027)	-0.124*** (0.032)
Share of bank in a firm's overall borrowing (f,b,t-1)	-0.160*** (0.016)	-0.172*** (0.016)	-0.182*** (0.020)	-0.320*** (0.017)	-0.316*** (0.013)	-0.329*** (0.022)
Observations	9657889	9878671	10023898	9657889	9878671	10023898
R-squared	0.00	0.16	0.40	0.01	0.18	0.39
Adjusted R-squared	0.00	0.06	0.03	0.01	0.07	0.01
Time-FE	Yes	-	-	Yes	-	-
Firm×Time FE	-	-	Yes	-	-	Yes
Bank×Industry×Location×Time FE	-	Yes	-	-	Yes	-

Notes: This table shows the regression results obtained when regressing time-varying innovations to firm-bank specific **credit demand** (columns 1-3) and **credit supply** (columns 4-6) on six bank characteristics, five firm characteristics and four bank-firm relationship-specific characteristics. In columns 1 and 4, we include time fixed effects. In columns 2 and 5, the bank characteristics are absorbed by bank-industry-location-quarter (BILT) fixed effects. In columns 3 and 6, the firm characteristics are absorbed by firm-quarter (FT) fixed effects. Standard errors are obtained from a weighted sum of two standard cluster variances, one clustered at the firm level, and the other at the bank level (see Appendix E). The unbalanced sample spans firm-bank level observations over the period 2019Q3-2023Q4 for 11 euro area countries.

positively and supply negatively (or vice versa). Likewise, some have an impact on demand, supply, both, or none.

In sum, the shocks we recover strongly correlate with the observable characteristics found in the AnaCredit dataset in ways that are economically meaningful. This is true at both the firm and the bank level – capturing “between” effects, and at the relationship-specific level – capturing “within effects”, or heterogeneity. The latter results are intrinsically unavailable from previous methodologies and thus provide novel insights. They also cast doubt on the suitability of the types of homogeneity assumptions maintained in the literature. In the following subsections, we dive deeper into these relationship-specific level or “within” effects.

6.3 The impact of relationship characteristics: subsample analyses

Table 4 in Subsection 6.1 documented significant variability in demand and supply shocks between firms and banks. In this subsection, we analyze whether this variability can be attributed to observable *relationship-specific* characteristics as captured in the AnaCredit data set. We integrate all observed and unobserved time-varying heterogeneity in firms and banks into granular fixed effects, and thus focus exclusively on relationship-level variables. The idea of this analysis is that, if our hypothesis is correct and supply and demand shocks exhibit meaningful heterogeneity within firms and banks, then, even after including fixed effects, the remaining variation should be correlated with observable relationship-specific characteristics.

Column 1 of Table 6 begins by examining how characteristics of the firm-bank credit relationship predict demand shocks (upper panel). For a given relationship, we find that the share of fixed-rate loans and the importance of a bank to a firm are both negative and highly significant predictors of demand shocks, conditional on fixed effects. Firms will demand more credit from banks, where firms’ past borrowing was mainly in the form of credit lines and term loans. These insights are obtained after controlling for a firm’s overall demand and a bank’s overall supply (through the included fixed effects).

In parallel, the lower panel of Table 6 considers relationship-specific predictors of supply shocks. Column (1) presents the full-sample regression. In contrast to the effect on credit demand, we find that the share of past lending that is fixed rate positively and significantly predicts supply shocks. Credit supply is also higher if a larger fraction of past borrowing is secured by collateral. On the other hand, the share of credit lines and term loans, and the share of a bank in a firm’s overall borrowing both significantly negatively predict supply shocks. Even after accounting for a bank’s overall supply – and overall supply to a particular

firm – supply will be smaller for a firm-bank pair with larger values of these predictors.

Table 10 in the appendix shows that the results of the first column are remarkably quantitatively robust to alternative clustering of the standard errors or the inclusion of bank-quarter fixed effects (rather than the more granular bank-industry-location-quarter fixed effects).

In column 2, we additionally include the probability of default assessed by the bank. The probability of default itself has zero effect on demand. In stark contrast to the demand results, the bank-specific assessment of firm credit risk is a highly significant negative predictor of credit supply. This finding makes sense: This implies that banks supply relatively less credit to firms that they deem to be riskier in a given period. Although it may seem like an obvious result, this finding – in contrast with the null result for demand – strongly suggests that we have credibly separately identified supply and demand shocks. Importantly, the other coefficients are little changed in this column. This first indicates that the effect of the probability of default is not related to other contractual characteristics. Moreover, since the inclusion of the PD variable effectively halves the sample (as it is not available for all relationships), it provides further robustness on the previously obtained results.

The results so far already clearly demonstrate that the heterogeneity present in the distributions of demand and supply shocks is economically meaningful and a potentially rich source of empirical insights leading to a better understanding of credit market dynamics. Subsequently, we take a first look at the stability of these coefficients in various sub-samples.

Table 6: Baseline and subsample analyses

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Demand innovation (f,b,t)						
Share of fixed rate loans (f,b,t-1)	-0.162*** (0.019)	-0.148*** (0.021)	-0.066*** (0.013)	-0.168*** (0.046)	-0.210*** (0.040)	-0.114*** (0.031)	-0.179*** (0.020)
Share of collateralized loans (f,b,t-1)	0.016 (0.016)	0.008 (0.024)	-0.002 (0.010)	0.022 (0.029)	0.034 (0.022)	0.064*** (0.015)	-0.028*** (0.005)
Share of Credit lines and Term Loans (f,b,t-1)	0.086*** (0.017)	0.085*** (0.020)	0.131*** (0.030)	0.127* (0.071)	0.042** (0.019)	0.040 (0.030)	0.102*** (0.020)
Share of bank in a firm's overall borrowing (f,b,t-1)	-0.206*** (0.023)	-0.214*** (0.031)	-0.117*** (0.027)	-0.085*** (0.012)	-0.320*** (0.045)	-0.319*** (0.015)	-0.123*** (0.013)
Probability of Default (f,b,t)		-0.023 (0.027)					
R-squared	0.52	0.52	0.51	0.56	0.50	0.56	0.48
Adjusted R-squared	0.04	0.03	0.05	0.09	0.02	0.06	0.03
	Supply innovation (f,b,t)						
Share of fixed rate loans (f,b,t-1)	0.156*** (0.021)	0.150*** (0.030)	0.008 (0.007)	0.027* (0.015)	0.311*** (0.042)	0.239*** (0.022)	0.089*** (0.009)
Share of collateralized loans (f,b,t-1)	0.034*** (0.013)	0.042** (0.018)	0.008 (0.013)	0.086*** (0.028)	0.012 (0.015)	0.048*** (0.013)	0.038* (0.020)
Share of Credit lines and Term Loans (f,b,t-1)	-0.150*** (0.030)	-0.204*** (0.043)	-0.235*** (0.036)	-0.129*** (0.027)	-0.126*** (0.035)	-0.063** (0.025)	-0.216*** (0.030)
Share of bank in a firm's overall borrowing (f,b,t-1)	-0.337*** (0.019)	-0.386*** (0.027)	-0.518*** (0.053)	-0.505*** (0.039)	-0.151*** (0.040)	-0.301*** (0.030)	-0.355*** (0.027)
Probability of Default (f,b,t)		-0.249*** (0.055)					
R-squared	0.52	0.52	0.51	0.53	0.52	0.57	0.47
Adjusted R-squared	0.04	0.02	0.04	0.03	0.06	0.07	0.02
Observations	10477109	5899787	3229758	2734506	4512845	4827038	5650071
Sample Period	201909-202312	201909-202312	201909-202012	202103-202206	202209-202312	201909-202312	201909-202312
Coverage	11 countries	11 countries	11 countries	11 countries	11 countries	Excl. Italy	Italy

Notes: Time-varying innovations to firm-bank specific **credit demand** (upper panel) or **credit supply** (lower panel) are regressed on four time-varying bank-firm characteristics. These are the share of borrowing by firm f from bank b at time $t - 1$ that is (1) fixed rate, (2) secured by collateral, or (3) credit line or term loan borrowing (with the remaining part mainly revolving credit). The fourth characteristic is the share of firm f 's borrowing from bank b at time $t - 1$ in firm f 's overall borrowing at time $t - 1$. The baseline sample (column 1) is an unbalanced sample of firm-bank level observations over the period 2019Q3-2023Q4 for 11 euro area countries. In column 2, we also include a bank-specific assessment of firm credit risk (resulting in a much smaller sample size). In subsequent columns, we estimate the specification of column 1 in various other subsamples. In columns 3 to 5, we split the sample into three periods of six quarters: the pandemic period 2019Q3-2020Q4, the inflationary period 2021Q1-2022Q2 and the monetary tightening period 2022Q3-2023Q4. In the last two columns, we return to the entire time period, but split the sample geographically. In column 6, we exclude Italy. In column 7, we only include Italian banks and firms. All specifications include firm-quarter fixed effects as well as bank-industry-location-quarter fixed effects. Standard errors are obtained from a weighted sum of two standard cluster variances, one clustered at the firm level, and the other at the bank level (see Appendix E).

In columns (3) to (5) of Table 6, we divide the sample into three periods of six quarters: the pandemic period 2019Q3-2020Q4, the inflationary period 2021Q1-2022Q2, and the monetary tightening period 2022Q3-2023Q4. In general, the signs of the significant coefficients remain similar in the three columns. For some variables, we observe sizable variation in the coefficients and thus in the economic magnitude. In particular, whether past borrowing is at a fixed rate appears to have a much larger effect on credit supply innovations in the last subsample covering the monetary tightening. During an episode where policy rates increase, banks have a more generous credit supply to firms who borrowed mainly fixed rate in the past. This can be linked to bank risk management, where banks have a more positive attitude towards firms whose interest expenses are predetermined. The share of a bank in the firm's total borrowing seems to gain importance for explaining heterogeneity in credit demand during the monetary tightening episode, whilst its impact on credit supply becomes economically smaller in the same subsample.

Columns (6) and (7) consider again the full sample period, but differ in the country coverage. Specifically, we assess how the results differ for Italy relative to all other countries. The upper and lower panels depict a slightly different story. The impact of the relationship characteristics is much more stable for credit supply than credit demand innovations. In the lower panel, the signs are similar in the two columns, and for two variables, the point estimates are similar across both columns. In the upper panel, we find that country-specificities play a larger role. For each of the four characteristics, we find sizable differences in economic magnitude and for two of them also in statistical significance or sign.

6.4 The distributional impact of monetary and macroprudential policy

The estimated coefficients in columns (3) to (5) of Table 6 indicate that the relationships between firm-bank characteristics and the credit demand and supply shocks exhibit time variation. We will now focus on two of these characteristics, namely the probability of default and the share of fixed-rate loans, to inspect whether the time variation is associated with policy changes, and may thus be given a structural interpretation. In particular, we will analyze whether monetary policy and macroprudential policy shape these relationships.

We interact the two firm-bank characteristics with the Monetary Policy (MP) and Central Bank Information (CBI) shocks from Jarociński and Karadi (2020). Changes in macroprudential policy are measured as the quarterly change in the index of 17 policy-action indicators from the IMF's integrated Macroprudential Policy (iMaPP) Database, originally constructed by Alam et al. (2019). The results are presented in 7.

Across specifications, we still find that firms' PD is not statistically significant in the credit demand regressions, though its interaction with CBI shocks is negative and statistically significant. This should be interpreted as follows: credit demand of borrowers with different risk profiles will be more dispersed the larger the economic news shock is. In addition, whether the news shock contains positive or negative news determines whether safer or riskier firms will have a higher credit demand (relative to the other group). When there is a positive (negative) CBI shock, which indicates an improvement (a deterioration) in the economic outlook, the safer (riskier) firms will demand more credit than the riskier (safer) firms. The relationship between credit supply and firms' PD seems not to be affected by the MP and CBI shocks. Adding changes in the macroprudential policy stance (columns 3 and 6) does not alter these conclusions. When macroprudential policy is tightened, the supply of credit to riskier firms decreases, and their demand for credit increases. However, the economic effects are smaller compared to the CBI shock on credit demand. Importantly, controlling for the interaction between PD and changes in macroprudential policy has an impact on the interaction effect of PD and MP surprises. The effect is also opposite in the credit demand and supply regressions. In the credit supply (demand) regression, the interaction term of the PD and MP surprises becomes larger (smaller) in absolute value and is almost significant. The coefficient suggests a reduced supply of credit to riskier borrowers during a tightening of monetary policy.

Table 7: The impact of monetary policy, central bank information and macroprudential policy

	(1)	(2)	(3)	(4)	(5)	(6)
	Demand innovation (f,b,t)			Supply innovation (f,b,t)		
Probability of Default (f,b,t)	0.016 (0.024)	0.011 (0.024)	-0.006 (0.026)	-0.212*** (0.050)	-0.207*** (0.049)	-0.185*** (0.050)
Monetary Policy (t) \times Probability of Default (f,b,t)	-0.397 (0.265)	-0.416 (0.255)	-0.044 (0.203)	-0.012 (0.175)	0.034 (0.183)	-0.446 (0.307)
Central Bank Information (t) \times Probability of Default (f,b,t)	-2.087*** (0.534)	-2.127*** (0.516)	-2.271*** (0.574)	-0.466 (0.302)	-0.349 (0.307)	-0.162 (0.329)
Share of fixed rate loans (f,b,t-1)			-0.127*** (0.022)	-0.125*** (0.021)	0.141*** (0.018)	0.141*** (0.018)
Monetary Policy (t) \times Share of fixed rate loans (f,b,t-1)			-0.511*** (0.127)	-0.533*** (0.127)	0.721*** (0.238)	0.725*** (0.237)
Central Bank Information (t) \times Share of fixed rate loans (f,b,t-1)		0.021 (0.300)	0.033 (0.313)		0.700* (0.366)	0.699* (0.365)
Quarterly Change in Macro-Prudential index (t) \times Probability of Default (f,b,t)				0.037*** (0.010)		-0.047* (0.025)
Quarterly Change in Macro-Prudential index (t) \times Share of fixed rate loans (f,b,t-1)				-0.017* (0.010)		0.003 (0.007)
Observations	5899787	5899787	5899787	5899787	5899787	5899787
R-squared	0.52	0.52	0.52	0.51	0.51	0.51
Adjusted R-squared	0.02	0.03	0.03	0.00	0.01	0.01
Firm \times Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Bank \times Industry \times Location \times Time FE	Yes	Yes	Yes	Yes	Yes	Yes
SE-cluster1	Bank	Bank	Bank	Bank	Bank	Bank
SE-cluster2	Firm	Firm	Firm	Firm	Firm	Firm

Notes: Time-varying innovations to firm-bank specific **credit demand** (columns 1-3) or **credit supply** (columns 4-6) are regressed on two firm-bank characteristics as well as their interactions with monetary policy shocks, central bank information shocks, and changes in macroprudential policy. The two firm-bank characteristics are a bank-specific assessment of firm credit risk and the share of borrowing by firm f from bank b at time t-1 that is fixed rate. The Monetary Policy and Central Bank Information shocks are obtained from Jarociński and Karadi (2020). Changes in macro-prudential policy are measured as the quarterly change in the index of 17 policy-action indicators from the IMF's integrated Macroprudential Policy (iMaPP) Database, originally constructed by Alam et al. (2019). All specifications include firm-quarter fixed effects as well as bank-industry-location-quarter fixed effects. Standard errors are obtained from a weighted sum of two standard cluster variances, one clustered at the firm level, and the other at the bank level (see Appendix E). The unbalanced sample spans firm-bank level observations over the period 2019Q3-2023Q4 for 11 euro area countries.

In addition, the interaction effects between MP and CBI shocks and the share of fixed-rate loans reveal interesting patterns. During monetary tightening episodes, the heterogeneity in credit supply and demand induced by past differences in fixed versus variable rate borrowing is much more sizable. This is in line with the results in column 5 of Table 6, where the sample period coincides with the monetary tightening period (2022Q3-2023Q4), where we also find larger coefficients compared to column 3 (the pandemic period). During a surprise tightening, banks' credit supply is more generous if borrowers borrowed predominantly fixed rate (from that bank) in the past. In bank-firm relationships with more historical variable rate borrowing, there will be more cautious supply as those firms' debt-bearing capacity is negatively affected by a monetary policy tightening. Positive CBI shocks increase the wedge in credit supply for fixed-rate borrowers to a similar extent.

In sum, the application to monetary and macroprudential policy changes and the established relationships therein also lend further credibility to our decomposition of credit growth and interest rate changes in their underlying demand and supply innovations. Moreover, these results illustrate another interesting application of our methodology. It illustrates how these heterogeneous shocks may be used to assess the distributional effects of financial or economic policies. We have showed an application with changes in monetary policy surprises and changes in macro-prudential policy. Future research could investigate whether the policy stance also matters or whether other structural differences between countries can explain cross-country heterogeneity in the relationships. Such an analysis would be helpful in interpreting and extrapolating results from single-country applications.

6.5 Credit growth, interest rate changes and credit demand and supply

When credit registers record both quantities and prices, implementing the Amiti and Weinstein (2018) or Khwaja and Mian (2008) approach to both dimensions could be considered as a useful framework for labeling relationships as demand or supply effects. While the signs of coefficients may be suggestive, we will now show that it is not necessarily conclusive and at best would only indicate which effect is the dominant one. Our approach, on the other hand, allows us to unveil the full picture of how a bank or firm characteristic may impact demand and/or supply and how this translates into quantity and price changes. We will illustrate these claims based on the results presented in Table 8 where we run a similar regression on four outcome variables: (i) Credit growth, (ii) Interest Rate changes, (iii) Credit demand innovations and (iv) Credit supply innovations.⁶

⁶In Table 11, we provide a table similar to Table 8, but based on the analysis performed in Table 7.

Table 8: Credit Growth and Interest rate changes vs. Demand and Supply innovations

	(1) Credit growth (f,b,t)	(2) Change in Interest Rate (f,b,t)	(3) Demand innovation (f,b,t)	(4) Supply innovation (f,b,t)
Share of fixed rate loans (f,b,t-1)	0.045*** (0.006)	-0.229*** (0.016)	-0.157*** (0.018)	0.158*** (0.016)
Share of collateralized loans (f,b,t-1)	0.047*** (0.011)	-0.021*** (0.008)	0.011 (0.013)	0.039*** (0.012)
Share of Credit lines and Term Loans (f,b,t-1)	-0.137*** (0.025)	0.118*** (0.017)	0.088*** (0.016)	-0.130*** (0.028)
Share of bank in a firm's overall borrowing (f,b,t-1)	-0.482*** (0.025)	-0.043*** (0.008)	-0.223*** (0.022)	-0.301*** (0.019)
Bank Sectoral Market Share (f,b,t-1)	-0.030 (0.040)	-0.010 (0.020)	-0.023 (0.023)	-0.001 (0.041)
Bank Sectoral Exposure (f,b,t-1)	0.225*** (0.042)	-0.069*** (0.019)	0.060 (0.038)	0.154*** (0.031)
Observations	12711274	12711274	12711274	12711274
R-squared	0.42	0.45	0.43	0.43
Adjusted R-squared	0.05	0.11	0.08	0.07
Firm \times Time FE	Yes	Yes	Yes	Yes
Bank \times Time FE	Yes	Yes	Yes	Yes
SE-cluster1	Bank	Bank	Bank	Bank
SE-cluster2	Firm	Firm	Firm	Firm
Sample	201909-202312	201909-202312	201909-202312	201909-202312
Coverage	11 countries	11 countries	11 countries	11 countries

Notes: This table shows the regression results obtained when regressing (i) Credit growth (column 1), (ii) Interest Rate changes (column 2), (iii) Credit demand innovations (column 3) and (iv) Credit supply innovations (column 4) on six variables that vary at the firm-bank-time level. These are the share of borrowing by firm f from bank b at time $t-1$ that is (1) fixed rate, (2) secured by collateral, or (3) credit line or term loan borrowing (with the remaining part being mainly revolving credit). The fourth characteristic is the share of firm f 's borrowing from bank b at time $t-1$ in firm f 's overall borrowing at time $t-1$. The fifth characteristic is the share of bank b 's lending to the sector of firm f in total borrowing of all firms in that sector (the market share of bank b in the sector of firm f). The sixth characteristic is the share of bank b 's lending to the sector of firm f in total lending by banking b (the sectoral exposure of bank b to the sector of firm f). We include firm-quarter fixed effects as well as bank-quarter (BILT) fixed effects, which subsume all firm and bank characteristics. Standard errors are clustered at the bank and firm level in columns 1 and 2. In columns 3 and 4, Standard errors are obtained from a weighted sum of two standard cluster variances, one clustered at the firm level, and the other at the bank level (see Appendix E). The unbalanced sample spans firm-bank level observations over the period 2019Q3-2023Q4 for 11 euro area countries.

The first two variables of Table 8 provide an interesting illustration. The share of past borrowing that was fixed rate or collateralized both positively impact credit growth while negatively affecting interest rate changes. These signs align with a supply-driven explanation. However, our decomposition shows that fixed-rate borrowing influences quantities and prices through a combination of contractionary demand and expansionary supply effects. The demand contraction offsets the supply expansion, leading to a relatively small net effect on credit growth but a significant negative effect on pricing.

In contrast, the impact of secured (collateralized) borrowing is purely a supply-side phenomenon, as indicated by the absence of demand-side effects. The result is small, but it produces a positive effect on credit growth and a negative one on pricing. In particular, relying solely on the magnitude of the coefficients in the credit growth regression could incorrectly suggest that the supply effects are equally important to those for fixed rate loans, which underscores the importance of incorporating price data and implementing our identification strategy to arrive at heterogeneous credit demand and supply effects.

The impact of the share of credit lines and term loans in past borrowing provides additional insight. Its effect on credit growth and interest rate changes is opposite in sign compared to the impact of the share of fixed rate borrowing. But this is not just due solely to a contractionary supply effect. This characteristic also has an expansionary effect on credit demand. The supply effect is approximately 40% stronger than the demand effect, resulting in a larger (absolute) impact on credit growth and a smaller impact on pricing compared to the share of fixed-rate loans. This asymmetry highlights the importance of considering the relative magnitudes of supply and demand effects when interpreting these relationships.

The effect of the main bank share on credit growth and price changes indicates that at least a contractionary demand effect should be at play. This is also effectively what we find in column 3. Our approach also reveals a contractionary supply effect. Together, these forces result in a significant contraction in credit with only a minimal effect on interest rates. This highlights the dual role of demand and supply channels in shaping credit outcomes in concentrated banking relationships.

Compared to the baseline approach, where we mainly focused on these four bank-firm relationship characteristics, we now also look at the impact of banks' market share in a sector and banks' exposure to a sector. The impact of banks' sectoral exposure aligns with previous findings in the literature, such as De Jonghe et al. (2024), who show that bank specialization in a sector enhances credit growth for healthy firms. However, as noted by Paravisini et al. (2023), firm-time fixed effects alone may not fully isolate supply-side effects, as bank specialization can also influence heterogeneous credit demand. In column 2, we find

that bank sectoral exposure also leads to more favorable pricing, which combined with the evidence in column 1 supports at least the idea that the supply effect dominates. Our credit demand and supply decomposition allows to conclude that in our setting the effect of bank sectoral exposure on credit market dynamics is mainly supply driven. The coefficient on credit demand is insignificant (though borderline) and much smaller in economic magnitude. But, if anything, it would imply that firms would direct their credit demand more to banks with more expertise in their sector.

7 Conclusion

This paper proposes a novel framework to study heterogeneous distributions of supply and demand shocks, and their impacts on price and quantity, in bipartite network data. The elasticities of price and quantity, or, alternatively, the shocks themselves, are non-parametrically identified under mild assumptions, and the elasticities can be consistently estimated. We illustrate the advantages of this new framework empirically. We characterize how the supply and demand dynamics of 11 European countries have evolved through the pandemic, subsequent inflation, and monetary policy tightening. The slopes of the supply and demand curves across different countries accord with contemporaneous macrofinancial conditions. Previous methodologies assuming homogeneous shocks generally capture less than 10% of variation in the shocks we identify. Moreover, we are able to study how – within a firm or bank – shocks covary with observable characteristics. The types of loans that characterise a relationship, as well as the importance of the relationship to the firm, robustly predict supply and demand shocks, although the sign may vary with financial conditions. The bank’s assessment of a firm’s probability of default strongly predicts supply shocks, but not demand shocks, validating our approach. We show that some time variation in the relationships between observable characteristics and shocks is associated with policy changes, demonstrating this methodology is a potentially powerful tool for understanding the distributional effects of policy.

The empirical analysis only begins to touch the surface of what is possible using modern credit register data; exhaustive study is left for dedicated future work. These types of results – demonstrating what identified supply and demand shocks do and do not correlate with – can be instructive for other empirical exercises. They may be useful to motivate alternative identifying assumptions. On the theoretical side, they may help to discipline models of credit supply and demand: they provide targets for how a bank’s credit supply should vary by firm, for example.

The methodology is also suited to a wide range of other empirical settings which remain to be explored, including labour market and trade applications. Econometrically, there is

scope to extend the dimensions of the analysis beyond price and quantity; indeed, the same approach could be extended to an arbitrary n -dimensional vector of observables that depends linearly on n or fewer underlying shocks in a bipartite network.

References

Abowd, J. J., Haltiwanger, J., & Lane, J. (2004). Integrated longitudinal employer-employee data for the united states. *American Economic Review*, 94(2), 224–229. <https://doi.org/10.1257/0002828041301812>

Abowd, J. M., Kramarz, F., & Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2), 251–333. <https://doi.org/https://doi.org/10.1111/1468-0262.00020>

Alam, Z., Alter, A., Eiseman, J., Gelos, G., Kang, H., Narita, M., Nier, E., & Wang, N. (2019). Digging Deeper—Evidence on the Effects of Macroprudential Policies from a New Database. *IMF Working Papers*, (2019/066). <https://ideas.repec.org/p/imf/imfwpa/2019-066.html>

Amiti, M., & 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), 525–587. <https://doi.org/10.1086/696272>

Andrews, M. J., Gill, L., Schank, T., & 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. <https://doi.org/https://doi.org/10.1111/j.1467-985X.2007.00533.x>

Beaumont, P., Libert, T., & Hurlin, C. (2019). Granular borrowers. *Université Paris-Dauphine Research Paper*, (3391768). <https://doi.org/https://dx.doi.org/10.2139/ssrn.3391768>

Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4), 841–890. Retrieved September 26, 2024, from <http://www.jstor.org/stable/2171802>

Bonhomme, S. Econometric analysis of bipartite networks (B. Graham & Áureo de Paula, Eds.). In: *The econometric analysis of network data* (B. Graham & Áureo de Paula, Eds.). Ed. by Graham, B., & Áureo de Paula. Academic Press, 2020, pp. 83–121. ISBN: 978-0-12-811771-2. <https://doi.org/https://doi.org/10.1016/B978-0-12-811771-2.00011-0>

Chamberlain, G. (1980). Analysis of Covariance with Qualitative Data. *The Review of Economic Studies*, 47(1), 225–238. <https://doi.org/10.2307/2297110>

Chava, S., & Purnanandam, A. (2011). The effect of banking crisis on bank-dependent borrowers. *Journal of Financial Economics*, 99(1), 116–135. <https://doi.org/https://doi.org/10.1016/j.jfineco.2010.08.006>

De Jonghe, O., Mulier, K., & Samarin, I. (2024). Bank specialization and zombie lending. *Management Science*, 0(0), 0–0. <https://doi.org/https://doi.org/10.1287/mnsc.2023.01437>

European Central Bank. (2019). *Anacredit reporting manual: Part ii - datasets and data attributes* (tech. rep.). European Central Bank. https://www.ecb.europa.eu/pub/pdf/other/AnaCredit_Manual_Part_II_Datasets_and_data_attributes_201905~cc9f4ded23.en.pdf

Gabaix, X., & Kojen, R. S. J. (2024). Granular instrumental variables. *Journal of Political Economy*, 132(7), 2274–2303. <https://doi.org/10.1086/728743>

Graeve, F. D., & Karas, A. (2014). Evaluating theories of bank runs with heterogeneity restrictions. *Journal of the European Economic Association*, 12(4), 969–996. Retrieved May 29, 2025, from <https://www.jstor.org/stable/90023403>

Graham, B. S., Imbens, G. W., & Ridder, G. (2014). Complementarity and aggregate implications of assortative matching: A nonparametric analysis. *Quantitative Economics*, 5(1), 29–66. <https://doi.org/https://doi.org/10.3982/QE45>

Greenstone, M., Mas, A., & 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. <https://doi.org/10.1257/pol.20160005>

Jarociński, M., & Karadi, P. (2020). Deconstructing monetary policy surprises—the role of information shocks. *American Economic Journal: Macroeconomics*, 12(2), 1–43. <https://doi.org/10.1257/mac.20180090>

Jochmans, K., & Weidner, M. (2019). Fixed-effect regressions on network data. *Econometrica*, 87(5), 1543–1560. <https://doi.org/https://doi.org/10.3982/ECTA14605>

Khwaja, A. I., & Mian, A. (2008). Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review*, 98(4), 1413–42. <https://doi.org/10.1257/aer.98.4.1413>

Kosekova, K., Maddaloni, A., Papoutsis, M., & Schivardi, F. (forthcoming). Firm-bank relationships: A cross-country comparison. *Review of Corporate Finance Studies*.

MacKay, A., & Miller, N. H. (forthcoming). Estimating models of supply and demand: Instruments and covariance restrictions. *American Economic Journal: Microeconomics*.

Newey, W. K., & McFadden, D. Chapter 36 large sample estimation and hypothesis testing. In: Vol. 4. *Handbook of Econometrics*. Elsevier, 1994, pp. 2111–2245. [https://doi.org/https://doi.org/10.1016/S1573-4412\(05\)80005-4](https://doi.org/https://doi.org/10.1016/S1573-4412(05)80005-4).

Olley, G. S., & Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6), 1263–1297. Retrieved September 26, 2024, from <http://www.jstor.org/stable/2171831>

Paravisini, D., Rappoport, V., & Schnabl, P. (2023). Specialization in bank lending: Evidence from exporting firms. *The Journal of Finance*, 78(4), 2049–2085. <https://doi.org/https://doi.org/10.1111/jofi.13254>

Rigobon, R. (2003). Identification through heteroskedasticity. *The Review of Economics and Statistics*, 85(4), 777–792. Retrieved October 3, 2023, from <http://www.jstor.org/stable/3211805>

A Labeling and normalisation of \hat{A}

The preliminary estimate \tilde{A} requires labeling and normalisation to obtain the supply and demand elasticities. In particular, let

$$\hat{A} = \hat{P}(\tilde{A}\tilde{\Upsilon})\tilde{A}\tilde{\Upsilon}, \quad (37)$$

where $P(\cdot)$ is a signed permutation matrix and $\tilde{\Upsilon}(\cdot)$ is a diagonal matrix imposing a scale normalisation. In practice, we set $\tilde{\Upsilon} = \tilde{\Lambda}_{FF}^{\frac{1}{2}}$, where $\tilde{\Lambda}_{FF} = \tilde{A}^{-1}S_{FF}\tilde{A}'^{-1}$, which normalises the covariances of demand and supply shocks across firms (within a given bank) to be unity. A signed permutation matrix is one in which a subset of columns may be negative. We choose $\hat{P}(\tilde{A}\tilde{\Upsilon})$ to be the signed permutation matrix $P \in \mathcal{P}$, where \mathcal{P} is the set of all such matrices, that minimises the Frobenius norm of $P\tilde{A}\tilde{\Upsilon} - \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$. That is, it is the signed permutation that minimises the elementwise distance between \hat{A} and $\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}$. Formally,

$$\hat{P}(\tilde{A}\tilde{\Upsilon}) = \operatorname{argmin}_{P \in \mathcal{P}} \sqrt{\sum_{i=1}^2 \sum_{j=1}^2 \left(P\tilde{A}\tilde{\Upsilon} - \begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix} \right)_{ij}^2}, \quad (38)$$

where $(M)_{ij}$ denotes the ij entry of matrix M . This means that the first column of \hat{A} corresponds to the shock whose properties are most consistent with those of a demand shock, while the second corresponds to a supply shock, and the effects of the shocks are scaled such that the shocks' covariances across firms (within a given bank) are unity. Once \hat{A} is obtained, it is trivial to obtain $\hat{\Lambda}_{BB}$, and thus the full vector of estimated parameters, $\hat{\theta}$. In reporting the elasticities, for instance in Figures 1a-2b, we rescale them such that both supply and demand shocks have unit standard deviations, to render the elasticities more interpretable across countries and time periods.

B Proof of Theorem 1

We begin by proving a preliminary lemma and proposition. We first prove that the variance of $\operatorname{vech}(S_{FF})$ vanishes as $F, B \rightarrow \infty$.

Lemma 1. *Under Assumptions 3 and 4, $\operatorname{var}(S_{FF}) \rightarrow 0$ as $F, B \rightarrow \infty$.*

Proof.

$$\text{var}(\text{vech}(S_{FF})) = \text{var} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \text{vech}(\eta_{fb} \eta'_{f'b}) \right) \quad (39)$$

$$= N_{FF}^{-2} \text{var} \left(\text{vech} \left(\sum_{b=1}^B \sum_{f' \neq f} \eta_{fb} \eta'_{f'b} \right) \right) \quad (40)$$

$$= N_{FF}^{-2} \sum_{b=1}^B \sum_{b'=1}^B \text{cov} \left(\text{vech} \left(\sum_{f' \neq f} \eta_{fb} \eta'_{f'b} \right), \text{vech} \left(\sum_{\tilde{f} \neq \tilde{f}'} \eta_{\tilde{f}b'} \eta'_{\tilde{f}b'} \right)' \right) \quad (41)$$

$$= N_{FF}^{-2} \sum_{b=1}^B \sum_{f' \neq f} \sum_{\tilde{f} \neq \tilde{f}'} \text{cov} \left(\text{vech}(\eta_{fb} \eta'_{f'b}), \text{vech}(\eta_{\tilde{f}b} \eta'_{\tilde{f}b})' \right) \quad (42)$$

$$+ 2N_{FF}^{-2} \sum_{b' \neq b} \sum_{f' \neq f} \text{cov} \left(\text{vech}(\eta_{fb} \eta'_{f'b}), \text{vech}(\eta_{fb'} \eta'_{f'b'})' \right) \quad (43)$$

$$\equiv \omega_{FF}, \quad (44)$$

where in (43) terms are dropped for which $(\tilde{f}, f) \neq (f, f')$ since these are zero by the structure and independence conditions in Assumption 3. We now derive the asymptotic behaviour of this variance. Note that the first summation contains

$$\frac{1}{4} \sum_{b=1}^B (F_b(F_b - 1))^2 \quad (45)$$

terms, while the second summation contains

$$\sum_{b' \neq b} \sum_{f' \neq f} \mathbf{1}_{fb} \mathbf{1}_{fb'} \mathbf{1}_{f'b} \mathbf{1}_{f'b'} \quad (46)$$

terms, where $\mathbf{1}_{fb}$ is an indicator for whether there is a connection between f and b . Since the number of connections of any single bank, F_b , cannot be larger than F , Assumption 4 guarantees that F_b grows proportionally to F for a non-vanishing share of banks, $\mu_{F,B} \rightarrow \tilde{\mu} \in (0, 1]$. It follows that

$$\lim_{F,B \rightarrow \infty} \frac{1}{F^2 B} \sum_{b=1}^B \frac{F_b(F_b - 1)}{2} = \tilde{\delta} \in (0, 1] \quad (47)$$

and thus

$$\lim_{F,B \rightarrow \infty} \left(\frac{1}{F^2 B} \sum_{b=1}^B \frac{F_b(F_b - 1)}{2} \right)^2 = \tilde{\delta}^2 \in (0, 1], \quad (48)$$

and additionally

$$\lim_{F,B \rightarrow \infty} \frac{1}{F^4 B} \sum_{b=1}^B \frac{F_b^2 (F_b - 1)^2}{4} = \tilde{\gamma} \in (0, 1]. \quad (49)$$

Then, since in the denominator we have $N_{FF}^2 = \left(\frac{1}{2} \sum_{b=1}^B F_b (F_b - 1) \right)^2 = O(B^2 F^4)$, the first term of ω_{FF} is $O(B^{-1})$, since the covariances are assumed to be finite, and thus converges to zero as $F, B \rightarrow \infty$. The number of non-zero terms entering the second summation is bounded above by that for the fully dense network where all firms are connected to all banks, in which case there are $\frac{1}{4} B(B-1)F(F-1)$ terms. Then, given the denominator N_{FF}^2 , the second term is $O(F^{-2})$, and converges to zero as $F, B \rightarrow \infty$. Thus, the total variance $\omega_{FF} = O(B^{-1}) + O(F^{-2}) = O(B^{-1})$ by the third point of Assumption 4 and converges to zero asymptotically. \square

The consistency of S_{FF} for Σ_{FF} then follows from Lemma 1:

Proposition 3. *Under Assumptions 3 and 4, $S_{FF} \xrightarrow{p} \Sigma_{FF}$; the convergence of $\text{vech}(S_{FF} - A\Lambda_{FF}A')$ is uniform in θ .*

Proof. Consistency of S_{FF} for Σ_{FF} follows directly from Lemma 1 by applying Chebyshev's Inequality. The convergence of

$$f(\theta) = S_{FF} - A\Lambda_{FF}A' \quad (50)$$

to $E[f(\theta)]$ is also uniform in θ since it is entirely separable in θ and \mathbf{X} , where \mathbf{X} denotes the sample of η_{fb} , and thus

$$\|f(\theta) - E[f(\theta)]\| = \|S_{FF} - \Sigma_{FF}\|, \quad (51)$$

independent of θ . \square

By symmetry, the same arguments apply to the estimator S_{BB} for Σ_{BB} , whose variance I denote as ω_{BB} . Denote the two blocks of $q(\mathbf{X}, \theta)$ as

$$q(\mathbf{X}, \theta) \equiv \begin{pmatrix} q_{FF}(\mathbf{X}, \theta) \\ q_{BB}(\mathbf{X}, \theta) \end{pmatrix} \quad (52)$$

We can now prove Theorem 1.

Proof. Following the uniform consistency result of Proposition 3, $\hat{\theta} \xrightarrow{p} \theta$ by standard minimum distance results, see, e.g., Theorem 2.1 of Newey and McFadden (1994), establishing the first part of the theorem.

To prove asymptotic normality, we first derive the limiting distribution of the two blocks of $q(\cdot)$, q_{FF} , q_{BB} . While the proof of Lemma 1 contains expressions for the variances of each block, it remains to derive their covariance, ω_{FB} . We have

$$\text{cov}(\text{vech}(S_{FF}), \text{vech}(S_{BB})') = \text{cov} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \text{vech}(\eta_{fb} \eta'_{f'b}), \frac{1}{N_{BB}} \sum_{\tilde{f}=1}^F \sum_{\tilde{b} \neq b} \text{vech}(\eta_{\tilde{f}b} \eta'_{\tilde{f}\tilde{b}}) \right) \quad (53)$$

$$= \frac{4}{N_{FF} N_{BB}} \sum_{f' \neq f} \sum_{b' \neq b} \text{cov} \left(\text{vech}(\eta_{fb} \eta'_{f'b}), \text{vech}(\eta_{fb} \eta'_{f'b})' \right), \quad (54)$$

where we have leveraged the structure of the error terms in Assumption 3. Thus, the covariance is $O(F^{-1}B^{-1})$ and vanishes asymptotically at a faster rate than either of the variance blocks.

Under Assumption 3, $\text{vech}(\eta_{fb} \eta'_{f'b})$ can be decomposed into several components with different covariance properties; in particular, the ij element of $\eta_{fb} \eta'_{f'b}$ is

$$\eta_{fb,i} \eta_{f'b,j} = \left(A_{i \cdot} \begin{pmatrix} u_{fb}^d \\ u_{fb}^s \end{pmatrix} \right) \left(A_{j \cdot} \begin{pmatrix} u_{f'b}^d \\ u_{f'b}^s \end{pmatrix} \right) \quad (55)$$

$$= \left(A_{i \cdot} \begin{pmatrix} e_{fb}^d + v_{fb}^d \\ e_{fb}^s + v_{fb}^s \end{pmatrix} \right) \left(A_{j \cdot} \begin{pmatrix} e_{f'b}^d + v_{f'b}^d \\ e_{f'b}^s + v_{f'b}^s \end{pmatrix} \right) \quad (56)$$

$$\equiv \alpha_{b,ff',ij} + \beta_{b,ff',ij} + \zeta_{b,ff',ij} + \xi_{b,ff',ij}, \quad (57)$$

where

$$\alpha_{b,ff',ij} = (A_{i1} e_{fb}^d + A_{i2} e_{fb}^s) (A_{j1} e_{f'b}^d + A_{j2} e_{f'b}^s) \quad (58)$$

$$\beta_{b,ff',ij} = (A_{i1} v_{fb}^d + A_{i2} v_{fb}^s) (A_{j1} v_{f'b}^d + A_{j2} v_{f'b}^s) \quad (59)$$

$$\zeta_{b,ff',ij} = (A_{i1} e_{fb}^d + A_{i2} e_{fb}^s) (A_{j1} v_{f'b}^d + A_{j2} v_{f'b}^s) \quad (60)$$

$$\xi_{b,ff',ij} = (A_{i1} v_{fb}^d + A_{i2} v_{fb}^s) (A_{j1} e_{f'b}^d + A_{j2} e_{f'b}^s). \quad (61)$$

Working term by term, for $\alpha_{b,ff',ij}$, $E[\alpha_{b,ff',ij}] = 0$ as it is the product of independent firm-specific components. First, note that

$$\frac{\sqrt{B}}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \alpha_{b,ff',ij} = \frac{\sqrt{B} F^2}{N_{FF}} \sum_{b=1}^B \frac{1}{F^2} \sum_{f' \neq f} \alpha_{b,ff',ij} = \frac{\sqrt{B} F^2}{N_{FF}} \sum_{b=1}^B O_p(F^{-1}), \quad (62)$$

since the inner summation is the sample average of at most $F(F-1)/2$ uncorrelated mean-zero terms by the independence conditions in Assumption 3, and is thus $\sqrt{F^2}$ consistent by the weak LLN. Then

$$\frac{\sqrt{B}F^2}{N_{FF}} \sum_{b=1}^B O_p(F^{-1}) = O\left(\frac{BF^2}{N_{FF}}\right) O_p\left(B^{\frac{1}{2}}F^{-1}\right) \rightarrow 0 \quad (63)$$

jointly in F, B , where the limit follows since $O\left(\frac{BF^2}{N_{FF}}\right) \rightarrow O(1)$ by (47) and $B^{\frac{1}{2}}/F \rightarrow 0$ by Assumption 4.

Turning to $\beta_{b,ff',ij}$, first note that

$$E[\beta_{b,ff',ij}] = E[\eta_{i,fb,i}\eta_{f'b,j}] = \Sigma_{FF,ij}, \quad (64)$$

since v_{fb}^d and v_{fb}^s are the only source of correlation across firms. Then

$$\sqrt{B} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \right) \quad (65)$$

is a (scaled) sum of B mean-zero independent random variables $\sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij}$. Note that under Assumption 3, in particular the i.i.d. and finite moments assumptions on the shock components, it follows that, for some $\epsilon > 0$,

$$\left(\sum_{b=1}^B \text{var} \left(\sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \right) \right)^{-(1+\epsilon/2)} \sum_{b=1}^B E \left[\left| \sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \right|^{2+\epsilon} \right] = O(B^{-\frac{\epsilon}{2}}) \xrightarrow[F,B \rightarrow \infty]{} 0, \quad (66)$$

Lyapunov's condition.

It follows from Lyapunov's Central Limit Theorem that

$$\left(\sum_{b=1}^B \text{var} \left(\sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \right) \right)^{-\frac{1}{2}} \sum_{b=1}^B \sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \xrightarrow{d} \mathcal{N}(0, 1). \quad (67)$$

Then, rearranging we obtain

$$\frac{\sqrt{B}}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \xrightarrow{d} \psi_{b,ij} \sim \mathcal{N}(0, V_{ij}), \quad (68)$$

where

$$V_{ij} = \lim_{F,B \rightarrow \infty} \frac{B}{N_{FF}^2} \sum_{b=1}^B \text{var} \left(\sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \right). \quad (69)$$

Joint asymptotic normality then holds across $\sqrt{B} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right)$, for $\beta_{b,ff'}$ the vector stacking unique entries over i, j . Note that for any linear combination given by $t \in \mathbb{R}^3$, the resulting random variable $\frac{\sqrt{B}}{N_{FF}} \sum_{b=1}^B t' \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF})$ will satisfy Lyapunov's Condition by the same argument given above, since the asymptotic weight given to each $\beta_{b,ff',ij}$ is constant across entries, so the linear combination is asymptotically normal with mean zero and variance $\frac{\tilde{\gamma}}{\delta^2} t' V t$, where $V = \lim_{F,B \rightarrow \infty} \frac{B}{N_{FF}^2} \sum_{b=1}^B \text{var} \left(\sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right)$ which is equal to $\text{var}(t' \Psi_b)$, where Ψ_b stacks $\psi_{b,ij}$ across ij . Note that V is full rank by Assumption 3. Then, by the Cramer-Wold device, the original random vector is jointly asymptotically normal:

$$\sqrt{B} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right) \xrightarrow{d} \mathcal{N}(0, V). \quad (70)$$

For $\zeta_{b,ff',ij}$, which is mean zero,

$$\sqrt{B} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \zeta_{b,ff',ij} - E[\zeta_{b,ff',ij}] \right) = \sqrt{B} O_p \left(N_{FF}^{-\frac{1}{2}} \right), \quad (71)$$

since the expression in parentheses is the sample average of N_{FF} mean-zero finite variance uncorrelated random variables, converging in probability to zero (e.g., Chebyshev's Weak LLN). Then $\lim_{F,B \rightarrow \infty} \sqrt{B} O_p \left(N_{FF}^{-\frac{1}{2}} \right) = 0$ by Assumption 4. By symmetry, the same argument applies to $\xi_{b,ff',ij}$, since it has the same properties. Thus, combining the limits of each component,

$$\sqrt{B} \left(S_{FF,ij} - (A \Lambda_{FF} A')_{ij} \right) \xrightarrow{d} \mathcal{N}(0, V_{ij}) \quad (72)$$

by Slutsky's Theorem. Stacking entries and using the joint normality established above for β_b , the only non-vanishing term asymptotically, gives a multivariate normal limiting distribution for $\sqrt{B} q_{FF}(\mathbf{X}, \theta)$, with variance denoted $\mathbf{W}_{FF} \equiv V$. By symmetry, a parallel argument applies to $\sqrt{F} q_{BB}(\mathbf{X}, \theta)$, with variance \mathbf{W}_{BB} .

Joint asymptotic normality (as $F, B \rightarrow \infty$) of $\sqrt{B} q_{FF}(\mathbf{X}, \theta)$ and $\sqrt{F} q_{BB}(\mathbf{X}, \theta)$ follows because they are asymptotically independent. In particular, asymptotically, the distribution of each sum depends exclusively on sums of v_{fb} and e_{fb} , respectively, which are independent by Assumption 3; all terms inducing dependence are asymptotically negligible. Due to this independence, Ψ_b and Ψ_f are jointly normal, so we have

$$\begin{pmatrix} \sqrt{B}q_{FF}(\mathbf{X}, \theta) \\ \sqrt{F}q_{BB}(\mathbf{X}, \theta) \end{pmatrix} \xrightarrow{d} \mathcal{N}(0, \mathbf{W}), \quad (73)$$

where

$$\mathbf{W} = \begin{bmatrix} \mathbf{W}_{FF} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_{BB} \end{bmatrix}. \quad (74)$$

Given the structure of q and the identification condition, regularity conditions for continuous differentiability and the Jacobian of q , denoted Φ , required for minimum distance estimation, are trivially satisfied. Therefore, with this limiting distribution, following the standard minimum distance argument (e.g., Theorem 3.2, Newey and McFadden (1994)),

$$\tilde{\Phi}(\hat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, \mathbf{W}), \quad (75)$$

where

$$\tilde{\Phi} = \left(\begin{bmatrix} \sqrt{B} & 0 \\ 0 & \sqrt{F} \end{bmatrix} \otimes I_3 \right) \Phi, \quad \Phi = E \left[\frac{\partial q(\mathbf{X}, \theta)}{\partial \theta'} \right]. \quad (76)$$

□

C Proof of Proposition 2

Proof. The proof of Theorem 1 defined \mathbf{W}_{FF} ,

$$\mathbf{W}_{FF} = \lim_{F, B \rightarrow \infty} \frac{B}{N_{FF}^2} \sum_{b=1}^B \text{var} \left(\sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right) \quad (77)$$

$$= \lim_{F, B \rightarrow \infty} \sum_{b=1}^B \frac{B(F_b(F_b-1)/2)^2}{N_{FF}^2} \text{var} \left(\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right) \quad (78)$$

$$= \frac{\tilde{\gamma}}{\tilde{\delta}^2} \tilde{\mu} \lim \text{var} \left(\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right) \quad (79)$$

$$= \frac{\tilde{\gamma}\tilde{\mu}}{\tilde{\delta}^2} \Omega_{FF}, \quad (80)$$

where $\Omega_{FF} = \lim_{F_b \rightarrow \infty} \text{var} \left(\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right)$. In particular, the third equality follows from the fact that F_b increases proportionally to F for a non-vanishing share, $\tilde{\mu}$ asymptotically, of banks as a consequence of Assumption 4, and for the remaining banks $\frac{B^2(F_b(F_b-1)/2)^2}{N_{FF}^2} \rightarrow 0$, so their variance is asymptotically negligible in computing (79). For a fraction $\tilde{\mu}$ of banks the scaled summations are independently and asymptotically identically

distributed.

Note that, following the decomposition in the proof of Theorem 1, $\hat{\mathbf{W}}_{FF}$ can be expanded in terms of four types of terms. That proof showed that sample averages of three of them over $f' \neq f$, holding b fixed and normalised by F^2 converged in probability to zero, for instance $\frac{1}{F^2} \sum_{f' \neq f} \alpha_{b,ff'} = O_p(F^{-1}) \rightarrow 0$. Since \hat{W}_{FF} features second powers of such sample averages, it is easy to show that any contributions owing to those terms are asymptotically negligible. The only non-vanishing terms in $\hat{\mathbf{W}}_{FF}$ are those depending exclusively on $\beta_{b,ff'}$. We thus study the behaviour of

$$\frac{B}{N_{FF}^2} \sum_{b=1}^B \left(\sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(S_{FF}) \right) \left(\sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(S_{FF}) \right)' . \quad (81)$$

First, denote

$$\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \xrightarrow{d} \nu_b, \quad (82)$$

where $\text{var}(\nu_b)$ is bounded for all b by Assumption 3. It follows that

$$\nu_b \sim \nu, \quad E[\nu] = 0, \quad \text{var}(\nu) = \Omega_{FF}, \quad (83)$$

for a fraction $\tilde{\mu}$ of banks as $F, B \rightarrow \infty$. Since S_{FF} is consistent for Σ_{FF} by Proposition 3,

$$\frac{1}{B} \sum_{b=1}^B \frac{B^2(F_b(F_b-1)/2)^2}{N_{FF}^2} \left(\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(S_{FF}) \right) \left(\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(S_{FF}) \right)' \quad (84)$$

$$\xrightarrow{p} \lim_{F,B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B \frac{B^2(F_b(F_b-1)/2)^2}{N_{FF}^2} \nu_b \nu_b' = \frac{\tilde{\gamma} \tilde{\mu}}{\tilde{\delta}^2} \Omega_{FF} + (1 - \tilde{\mu}) \times 0 = \frac{\tilde{\gamma} \tilde{\mu}}{\tilde{\delta}^2} \Omega_{FF}, \quad (85)$$

as the average of κB i.i.d. random variables consistent by Chebyshev's Inequality. □

D Lemma 2

Lemma 2. *In the model $\eta_{fb} = \Gamma X_{fb} + A u_{fb}$, Γ can be consistently estimated by OLS under Assumptions 1-4 provided that X_{fb} is independent of $X_{f'b'}$, for all $f' \neq f, b' \neq b$, X_{fb} has full rank and finite variance, and $\text{var}(A u_{fb} X'_{fb}) < \infty$.*

Proof. By the usual algebra, making the dependence on network structure D explicit, the

OLS estimator of Γ is

$$\hat{\Gamma} = \Gamma + \frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B A u_{fb} X'_{fb} D_{fb} \left(\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B X_{fb} X'_{fb} D_{fb} \right)^{-1}. \quad (86)$$

The expectation of the numerator of the error term is zero by Assumption 1. To see this, note that by the law of iterated expectations,

$$E \left[\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B A u_{fb} X'_{fb} D_{fb} \right] = \frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B A E \left[E [u_{fb} | X, D] X'_{fb} D_{fb} | D \right] \quad (87)$$

$$= \frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B A E \left[E [u_{fb} | D] X'_{fb} D_{fb} | D \right] \quad (88)$$

$$= E [D_{fb}] A E \left[E [u_{fb} | D_{fb} = 1, D_{-fb}] X'_{fb} | D_{fb} = 1, D_{-fb} \right] \quad (89)$$

$$= E [D_{fb}] A E [u_{fb} | D_{fb} = 1, D_{-fb}] E \left[X'_{fb} | D_{fb} = 1, D_{-fb} \right] \quad (90)$$

$$= 0. \quad (91)$$

The variance of the numerator is

$$\text{var} \left(\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B A u_{fb} X'_{fb} \right) \quad (92)$$

$$= \frac{1}{N^2} \sum_{f=1}^F \sum_{b=1}^B \sum_{f'=1}^F \sum_{b'=1}^B \text{cov} (A u_{fb} X'_{fb}, X'_{f'b'} A u'_{f'b'}) \quad (93)$$

$$\leq \frac{F^2 B^2}{N^2} \frac{1}{F^2 B^2} \sum_{f=1}^F \sum_{b=1}^B (F + B - 1) \text{var} (A u_{fb} X'_{fb}) \quad (94)$$

$$= \frac{F^2 B^2}{N^2} \frac{(F + B - 1)}{FB} \text{var} (A u_{fb} X'_{fb}) \rightarrow 0. \quad (95)$$

Since the correlation pattern is identical in the denominator, the variance of the denominator also vanishes as $F, B \rightarrow \infty$. Applying Chebyshev's Inequality separately to both the numerator and the denominator followed by Slutsky shows that $\hat{\Gamma} \xrightarrow{p} \Gamma$. \square

E Inference based on estimated shocks

Fixed-effects based approaches to shock estimation in the corporate credit literature – for instance, Khwaja and Mian (2008) and Amiti and Weinstein (2018) – aim to consistently

estimate supply and/or demand shocks. In our setting, the shocks specific to any firm-bank pair can be obtained by simply inverting \hat{A} and computing $\hat{u}_{fb} = \hat{A}^{-1}\eta_{fb}$. Due to the linearity assumption in (7), our methodology exploits the information in all F firms and B banks to estimate the matrix A , and thus identify the shocks specific to any firm-bank pair. In contrast, as shown by Jochmans and Weidner (2019), in Abowd et al. (1999)-type estimators, information is exploited in effect from only the direct neighbours of a given firm or bank. While our estimates of the relationship-specific shocks, \hat{u}_{fb} , are unbiased for the true values, asymptotic inference for an average firm demand shock or bank supply shock is conceptually problematic from a frequentist perspective. The reason is that, akin to the random effects literature, shocks are better viewed as random variables. It would be more appropriate to instead consider prediction intervals. This is a key conceptual distinction between the original Abowd et al. (1999) setting – where unobserved heterogeneity captures fixed worker and employer characteristics – and many subsequent empirical applications, where the unobserved heterogeneity takes the form of shocks or other random variables. Inference on a random variable is not a well-defined frequentist problem. In previous frameworks, it is perhaps more accurate to think about sampling uncertainty of fixed effects estimates of shocks vanishing asymptotically, with the sampling distribution collapsing around some realisation of the underlying random variable.

Although we do not impose that $B_f \rightarrow \infty$ for all firms, if we did it would be immediate to show (under Assumptions 1 – 4) that the sample average of the relationship-specific shock estimates for any firm would converge in probability to e_{fb}^d and e_{fb}^s , and similarly for banks to v_{fb}^d and v_{fb}^s . Such averages provide an analogous estimate to the supply shocks obtained by Amiti and Weinstein (2018), for instance. We explore the relative accuracy of these approaches in simulations below and find that such sample averages outperform well-specified fixed effects estimates in terms of their correlation with a true “average” supply or demand shock.

While inference for the shocks themselves is not a well-defined frequentist problem, it is natural to treat such random variables as generated regressors in regressions of interest. In particular, we use them as a dependent variable in a variety of specifications. We now establish the asymptotic properties of OLS estimators based on these generated regressors. First, we show that moments in terms of the estimated shocks and some other observable, Z_{fb} are consistent estimates of their population counterparts based on the true, infeasible shocks.

Proposition 4. *For some covariate, Z_{fb} , which is potentially correlated across both firms (holding b fixed) and banks (holding f fixed) and $E\left[Z_{fb}Z'_{fb} (u_{fb}^i)^2\right] < \infty$, under Assumptions*

1-4,

$$\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B Z_{fb} \hat{u}_{fb}^i \xrightarrow{p} E[Z_{fb} u_{fb}^i], \quad i \in \{s, d\}. \quad (96)$$

Proof. Note that

$$\hat{u}_{fb} = A^{-1} \eta_{fb} = u_{fb} + (\hat{A}^{-1} - A) \eta_{fb}, \quad (97)$$

where $(\hat{A}^{-1} - A) \xrightarrow{p} 0$ and is asymptotically independent of all η_{fb} . Then

$$\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B \hat{u}_{fb}^i Z_{fb} \rightarrow \text{plim} \frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B u_{fb}^i Z_{fb}. \quad (98)$$

Under the stated assumptions,

$$\text{var} \left(\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B u_{fb}^i Z_{fb} \right) \quad (99)$$

$$= \frac{1}{N^2} \sum_{f=1}^F \sum_{b=1}^B \sum_{f'=1}^F \sum_{b'=1}^B \text{cov}(u_{fb}^i Z_{fb}, u_{f'b'}^i Z_{f'b'}) \quad (100)$$

$$\leq \frac{F^2 B^2}{N^2} \frac{1}{F^2 B^2} \sum_{f=1}^F \sum_{b=1}^B (F+B-1) E[u_{fb}^{i2} Z_{fb} Z'_{fb}] \quad (101)$$

$$= \frac{F^2 B^2}{N^2} \frac{(F+B-1)}{FB} E[u_{fb}^{i2} Z_{fb} Z'_{fb}] \rightarrow 0. \quad (102)$$

Then the estimate is consistent for the expectation by Chebyshev's Inequality. \square

Now, consider regressions of the form

$$u_{fb}^i = \varsigma^i + \phi^i Z_{fb} + v_{fb}^i, \quad i \in \{s, d\}, \quad (103)$$

where $E[v_{fb}^i | Z_{fb}] = 0$. We make an additional assumption on the dependence structure of the regressors and regression errors across firms and banks, appropriate moments, and the asymptotic behaviour of F and B .

Assumption 6. Z_{fb} can be decomposed as $Z_{fb} = z_{fb,f} + z_{fb,b}$, where z_{fb}^f, z_{fb}^b are independent, but $z_{fb,f}$ may be dependent across banks, and $z_{fb,b}$ may be dependent across firms. For $i \in \{s, d\}$, v_{fb}^i can be decomposed as $v_{fb}^i = v_{fb,f}^i + v_{fb,b}^i$, where $v_{fb,f}^i, v_{fb,b}^i$ are independent, but $v_{fb,f}^i$ may be dependent across b , and $v_{fb,b}^i$ may be dependent across f . Additionally, $E[Z_{fb} Z'_{fb} (v_{fb}^i)^2]$ is positive semidefinite and $E[Z_{fb,l}^2 Z_{fb,k}^2 (v_{fb}^i)^4] < \infty$ for all regressors l, k . $\lim F/B$ exists and is in $[0, \infty]$.

With this assumption in hand, we can characterise the asymptotic distribution; details on the variance structure and estimators are given in the proof.

Proposition 5. *Under Assumptions 1-6, the OLS estimator of ϕ in (103) follows $\sqrt{\min(F, B)} (\hat{\phi} - \phi) \xrightarrow{d} \mathcal{N}(0, V_\phi)$; V_ϕ can be consistently estimated using clustered variance estimators.*

Proof. Proposition 4 shows that the numerator of the OLS estimator is consistent for its population counterpart (using the infeasible, true shocks). By a similar argument,

$$\frac{1}{N} \sum_{b=1}^B \sum_{f=1}^F Z_{fb} Z'_{fb} \equiv \hat{Q}_{ZZ} \xrightarrow{p} E [Z_{fb} Z'_{fb}] \equiv Q_{ZZ}, \quad (104)$$

so $\hat{\phi}^i \xrightarrow{p} \phi^i$.

Note that we can write $\hat{u}_{fb}^i = u_{fb}^i + (\hat{A}^{-1} A - I_2)_i u_{fb}$. Following Theorem 1,

$$\hat{u}_{fb}^i = \left(I_2 + O_p \left(\min(F, B)^{-\frac{1}{2}} \right) \right)_i u_{fb}. \quad (105)$$

We thus study the asymptotic behaviour in terms of the infeasible true shocks, before showing that this suffices to characterise that involving the estimated shocks.

Let $\tilde{\phi}^i$ be the infeasible OLS estimator based on the true shocks. As usual, we note that

$$\tilde{\phi}^i - \phi = \left(\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B Z_{fb} Z'_{fb} \right)^{-1} \frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B Z_{fb} v_{fb}^i, \quad (106)$$

and thus study $\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B Z_{fb} v_{fb}^i$. First, consider the asymptotic behaviour of

$$\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B Z_{fb} v_{fb}^i. \quad (107)$$

We can construct a similar decomposition of $Z_{fb} v_{fb}^i$ as in the proof of Theorem 1. In particular,

$$Z_{fb} v_{fb}^i = \tau_{fb,f}^i + \tau_{fb,b}^i + \lambda_{fb,fb}^i + \lambda_{fb,bf}^i \quad (108)$$

$$\tau_{fb,f}^i \equiv z_{fb,f} v_{fb,f}^i \quad (109)$$

$$\tau_{fb,b}^i \equiv z_{fb,b} v_{fb,b}^i \quad (110)$$

$$\lambda_{fb,fb}^i \equiv z_{fb,f} v_{fb,b}^i \quad (111)$$

$$\lambda_{fb,bf}^i \equiv z_{fb,b} v_{fb,b}^i, \quad (112)$$

and similarly decompose the summation into four corresponding summations. Consider first $\frac{1}{N} \sum_{b=1}^B \sum_{f=1}^F \tau_{fb,b}^i$, and assume for now that Z_{fb} is a scalar. By Assumption 6, this is a normalised sum of B independent mean zero random variables, $\sum_{f=1}^F \tau_{fb,b}^i$. By the moment assumptions made in Assumption 6, and definition of N , $\frac{\sqrt{B}}{B} \sum_{b=1}^B \left((B/N) \sum_{f=1}^F \tau_{fb,b}^i \right)$ satisfies Lyapunov's condition, so

$$\frac{\sqrt{B}}{B} \sum_{b=1}^B \left((B/N) \sum_{f=1}^F \tau_{fb,b}^i \right) \xrightarrow{d} \mathcal{N}(0, V_{ZB}^i). \quad (113)$$

If Z_{fb} is vector valued, joint asymptotic normality follows from the Cramer-Wold device, similarly to the proof of Theorem 1. By symmetry,

$$\frac{\sqrt{F}}{F} \sum_{f=1}^F \left((F/N) \sum_{b=1}^B \tau_{fb,f}^i \right) \xrightarrow{d} \mathcal{N}(0, V_{ZF}^i), \quad (114)$$

and joint convergence likewise holds when Z_{fb} is vector-valued.

Turning to $\frac{1}{N} \sum_{b=1}^B \sum_{f=1}^F \lambda_{fb,fb}^i$, we have

$$\sqrt{B} \left(\frac{1}{N} \sum_{b=1}^B \sum_{f=1}^F \lambda_{fb,fb}^i \right) = \sqrt{B} O_p \left(N^{-\frac{1}{2}} \right), \quad \sqrt{F} \left(\frac{1}{N} \sum_{b=1}^B \sum_{f=1}^F \lambda_{fb,fb}^i \right) = \sqrt{F} O_p \left(N^{-\frac{1}{2}} \right), \quad (115)$$

since the expression in parentheses is the sample average of N mean-zero finite variance uncorrelated random variables, converging in probability to zero (e.g., Chebyshev's Weak LLN). Then in either case, the expressions converge to zero by Assumption 4.

Thus,

$$\frac{\sqrt{\min(F, B)}}{N} \sum_{b=1}^B \sum_{f=1}^F Z_{fb} v_{fb}^i \xrightarrow{d} \varrho_B + \varrho_F, \quad (116)$$

where $\frac{\sqrt{\min(F, B)}}{N} \sum_{b=1}^B \left(\sum_{f=1}^F \tau_{fb,f}^i \right) \xrightarrow{d} \varrho_B$ and $\frac{\sqrt{\min(F, B)}}{N} \sum_{b=1}^B \left(\sum_{f=1}^F \tau_{fb,b}^i \right) \xrightarrow{d} \varrho_F$. The asymptotic behaviour depends on the relative rates of F, B . If $B/F \rightarrow 0$, then $\varrho_F = 0$, since $\min(F, B)$ is not of large enough asymptotic order to normalise the underlying summation (whose variance is $O(F^{-1})$). On the other hand, $\varrho_B \sim \mathcal{N}(0, V_{ZB}^i)$, so the entire summation is distributed $\mathcal{N}(0, V_{ZB}^i)$. By symmetry, the opposite is true if $B/F \rightarrow 0$. If $F/B \rightarrow c \in [1, \infty)$, then $\varrho_F \sim \mathcal{N}(0, c^{-1} V_{ZF}^i)$; if $c \in (0, 1)$, then $\varrho_B \sim \mathcal{N}(0, c V_{ZB}^i)$. ϱ_B and ϱ_F are independent since τ_{fb}^f, τ_{fb}^b are independent by Assumption 6. Then the sum converges to

$\mathcal{N}(0, V_{ZB}^i + c^{-1}V_{ZF}^i)$. Defining V_{Zv}^i as

$$V_{Zv}^i = \begin{cases} V_{ZB}^i, & B/F \rightarrow 0 \\ V_{ZF}^i, & F/B \rightarrow 0 \\ V_{ZB}^i + c^{-1}V_{ZF}^i, & F/B \rightarrow c \in [1, \infty) \\ cV_{ZB}^i + V_{ZF}^i, & F/B \rightarrow c \in (0, 1) \end{cases} \quad (117)$$

yields

$$\sqrt{\min(F, B)} (\tilde{\phi}^i - \phi^i) \xrightarrow{d} \mathcal{N}(0, V_\phi^i), \quad (118)$$

where $V_\phi^i = Q_{ZZ}^{-1} V_{Zv}^i Q_{ZZ}^{-1}$.

It remains to relate the limiting distribution of $\hat{\phi}^i$ constructed using the estimated shocks, \hat{u}_{fb}^i , to this one. Following (105),

$$\hat{\phi}^i = \left(1 + O_p\left(\min(F, B)^{-\frac{1}{2}}\right)\right) \tilde{\phi}^i \xrightarrow{d} \mathcal{N}(0, V_\phi^i), \quad (119)$$

since the first term converges to unity.

V_{ZB}^i and V_{ZF}^i are clustered variances at firm and bank level, respectively, for unbalanced panel data. Under the moment assumptions in Assumption 6, the standard clustered estimators are consistent. To illustrate this, consider the infeasible estimator

$$\tilde{V}_{ZB}^i = \frac{B}{N^2} \sum_{b=1}^B \left(\sum_{f=1}^F Z_{fb} v_{fb}^i \right) \left(\sum_{f=1}^F Z_{fb} v_{fb}^i \right)' - \left(\frac{1}{N} \sum_{b=1}^B \sum_{f=1}^F Z_{fb} v_{fb}^i \right) \left(\frac{1}{N} \sum_{b=1}^B \sum_{f=1}^F Z_{fb} v_{fb}^i \right)' \quad (120)$$

The second term is consistent for $\lim E[Z_{fb} v_{fb}^i] E[Z_{fb} v_{fb}^i]'$ following Proposition 4 and the continuous mapping theorem.

The variance of this object is the sum of at most $O(BF^4 + B^2F^2) = O(BF^4)$ non-zero covariances, since in order to have non-zero covariance terms must either have common b index or two common f indices, but the object is normalised by B^2N^{-4} , and $B^2N^{-4}O(BF^4) = O(B^{-1}) \rightarrow 0$. Note that

$$\frac{B}{N^2} \sum_{b=1}^B \left(\sum_{f=1}^F Z_{fb} v_{fb}^i \right) \left(\sum_{f=1}^F Z_{fb} v_{fb}^i \right)' = \frac{1}{B} \sum_{b=1}^B \left(B/N \sum_{f=1}^F Z_{fb} v_{fb}^i \right) \left(B/N \sum_{f=1}^F Z_{fb} v_{fb}^i \right)' \quad (121)$$

Therefore, by Chebyshev's inequality,

$$\frac{B}{N^2} \sum_{b=1}^B \left(\sum_{f=1}^F Z_{fb} v_{fb}^i \right) \left(\sum_{f=1}^F Z_{fb} v_{fb}^i \right)' \xrightarrow{p} \lim E \left[\left(B/N \sum_{f=1}^F Z_{fb} v_{fb}^i \right) \left(B/N \sum_{f=1}^F Z_{fb} v_{fb}^i \right)' \right]. \quad (122)$$

Combining the two terms, we have

$$\tilde{V}_{ZB}^i \xrightarrow{p} \lim E \left[\left(B/N \sum_{f=1}^F Z_{fb} v_{fb}^i \right) \left(B/N \sum_{f=1}^F Z_{fb} v_{fb}^i \right)' \right] - E [Z_{fb} v_{fb}^i] E [Z_{fb} v_{fb}^i]' \quad (123)$$

$$= \lim \text{var} \left(B/N \sum_{f=1}^F Z_{fb} v_{fb}^i \right) = V_{ZB}^i. \quad (124)$$

As above, $\hat{v}_{fb}^i = \left(I_2 + O_p \left(\min(F, B)^{-\frac{1}{2}} \right) \right)_i v_{fb}$, so $\hat{V}_{ZB}^i = \left(1 + O_p \left(\min(F, B)^{-\frac{1}{2}} \right) \right) \tilde{V}_{ZB}^i \xrightarrow{p} V_{ZB}^i$. By symmetry, the same is true for the analogous estimator \hat{V}_{ZF}^i . Finally, \hat{V}_{Zv}^i can be consistently estimated by combining $\hat{V}_{ZB}^i, \hat{V}_{ZF}^i, F, B$ according to (117), and thus $\hat{V}_\phi^i = \hat{Q}_{ZZ}^{-1} \hat{V}_{Zv}^i Q_{ZZ}^{-1} \xrightarrow{p} V_\phi^i$. \square

Three remarks are in order. First, the structure of the asymptotic variance is dependent on the relative rates of F, B . Thus, the researcher is required to take a stand on the appropriate asymptotic framework to compute the estimator. Second, \hat{V}_ϕ^i is a weighted sum of two standard cluster variances, one clustered at the firm level, and the other at the bank level. Note that these variance estimators must respect the fact that the panel data is unbalanced, however, due to the varying degrees of each firm and bank. Third, in the case that either $\text{var}(z_{fb,f}) = 0$ or $\text{var}(z_{fb,b}) = 0$, for instance in the cases of regressions using only firm or bank fixed effects, respectively, the variance estimators will still be asymptotically correct. Either \hat{V}_{ZB}^i or \hat{V}_{ZF}^i will go to zero, in the presence of irrelevant clusters. However, in the case where $B/F \rightarrow 0$ and only bank-invariant regressors are used (e.g, firm fixed effects), and vice versa, the limiting distribution may be degenerate. For instance, in the case where $B/F \rightarrow 0$, but the only regressors included are bank invariant (i.e. firm fixed effects), the asymptotic variance is normalised by $\sqrt{B} = \min(F, B)^{\frac{1}{2}}$, so $V_{Zv}^i = V_{ZB}^i$, but $V_{ZB}^i \rightarrow 0$. In this case, it is necessary to study the distribution of $\sqrt{F} (\hat{\phi}^i - \phi^-)$ in order to obtain a non-degenerate limiting distribution.

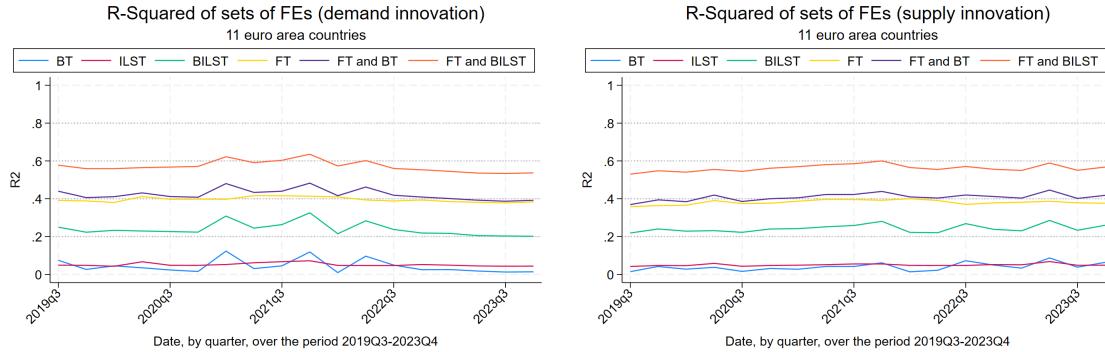
F Additional empirical results

Table 9: Summary statistics

	Pandemic			Inflation			Tightening		
	<i>F</i>	<i>B</i>	<i>N</i>	<i>F</i>	<i>B</i>	<i>N</i>	<i>F</i>	<i>B</i>	<i>N</i>
Austria	6,324	334	17,371	7,222	446	19,493	17,234	416	45,824
Belgium	12,511	19	27,129	13,398	20	29,297	16,891	21	37,107
Germany	59,059	848	151,185	60,468	808	155,567	95,451	774	242,030
Spain	108,521	99	323,796	100,198	101	302,326	114,485	96	328,883
Finland	7,649	172	16,324	7,019	158	15,026	13,749	144	30,155
France	60,156	129	142,101	74,498	132	176,373	57,476	131	135,142
Greece	3,536	16	9,645	4,042	15	10,074	8,165	14	20,072
Ireland	200	10	409	217	9	439	650	10	1,334
Italy	192,523	214	582,294	168,079	202	497,973	196,463	195	583,328
Netherlands	1,092	19	2,267	1,692	19	3,585	1,519	20	3,282
Portugal	22,700	110	62,724	25,288	103	68,216	29,881	99	80,965

Notes: This table reports the sample sizes by country for each of the subsamples studied. *F* indicates the number of multi-bank firms, *B* the number of banks and *N* the number of relationships between banks and such firms. The steps used to select this sample are described in the text.

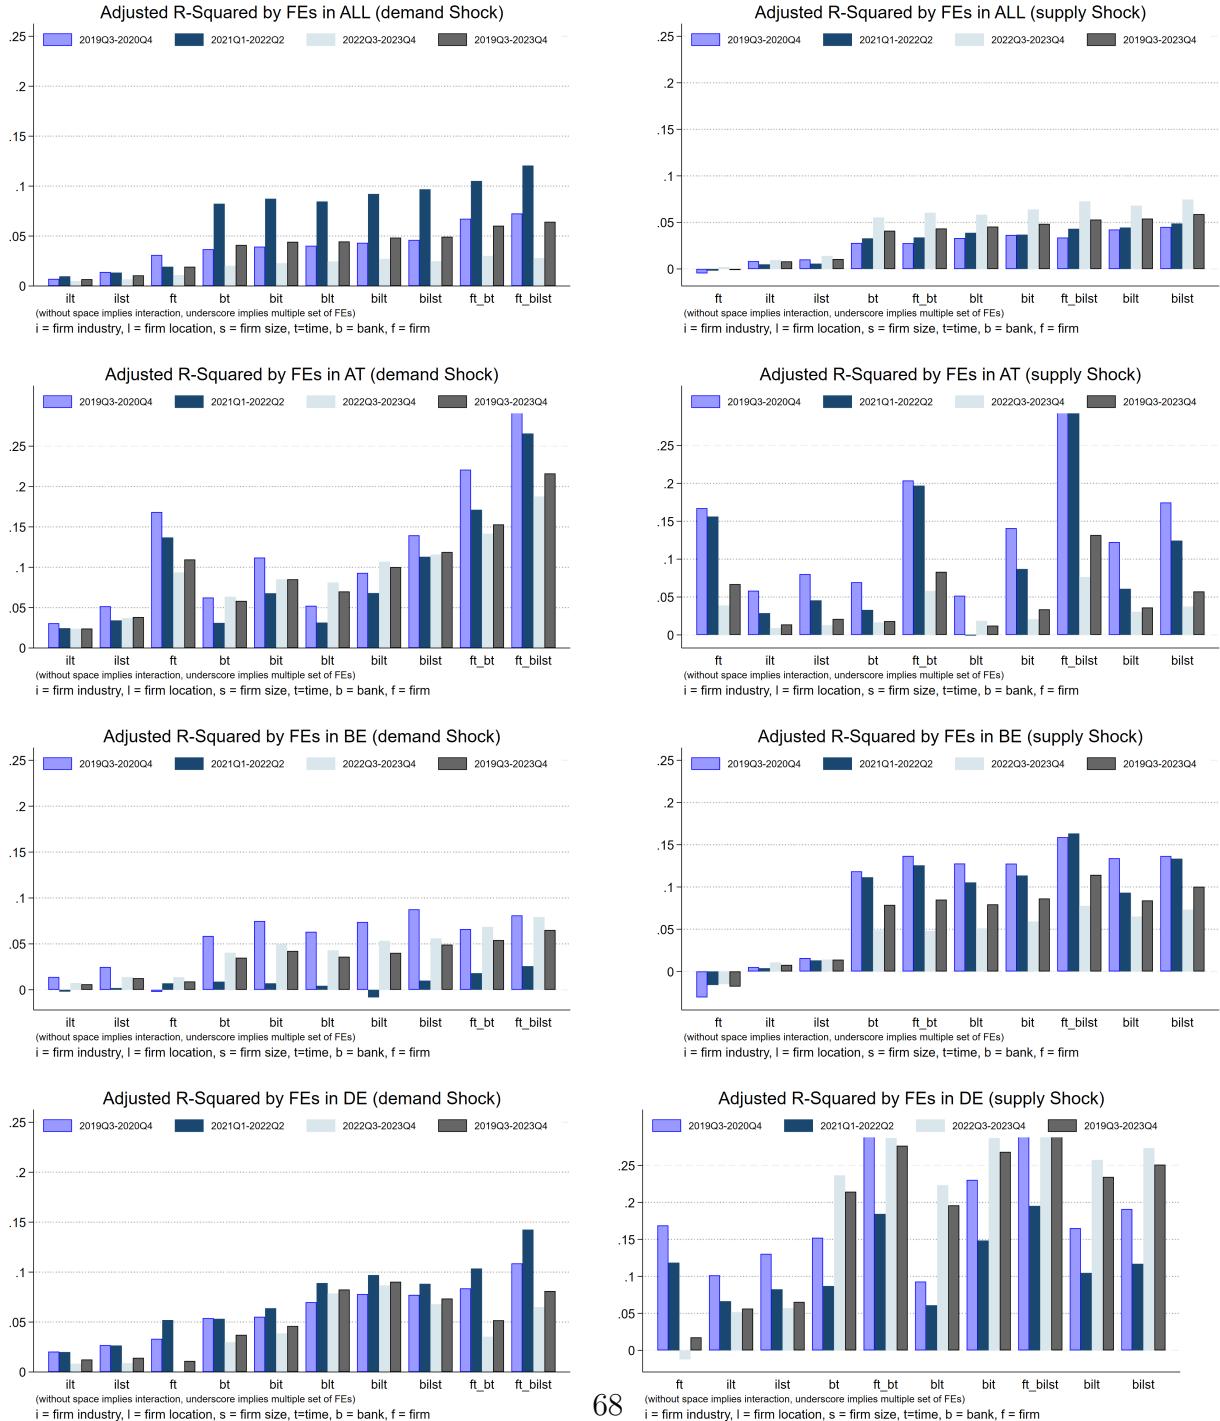
Figure 6: Variation in Demand and Supply shocks explained by Sets of Fixed Effects

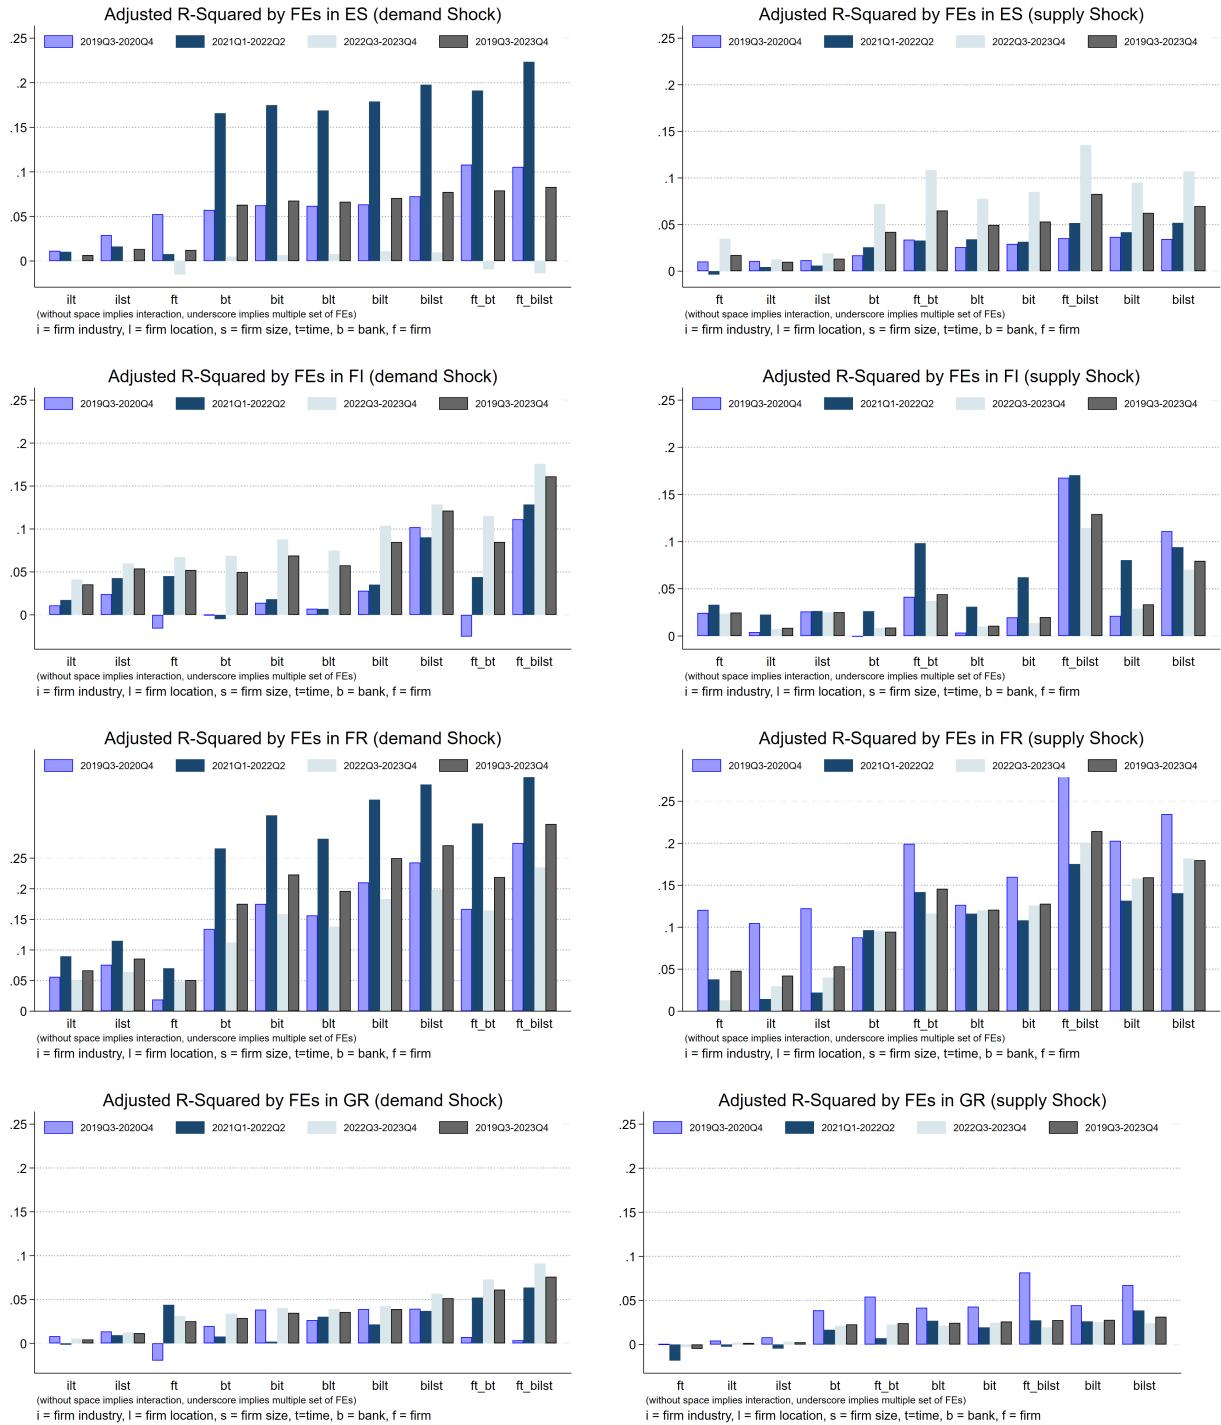


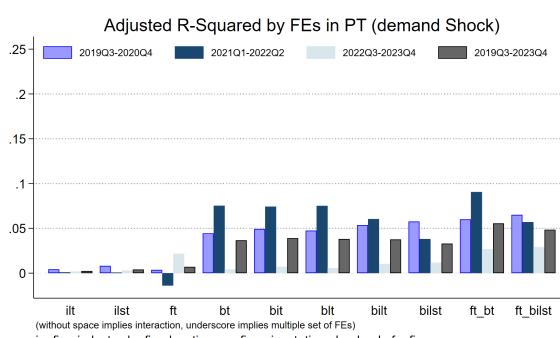
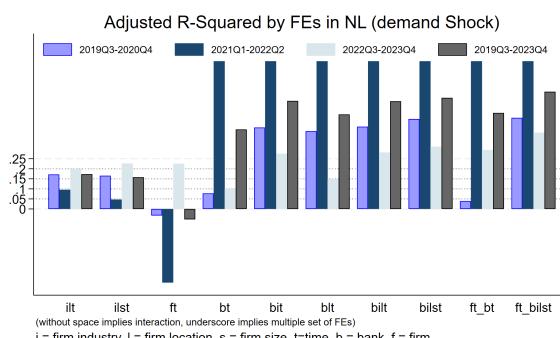
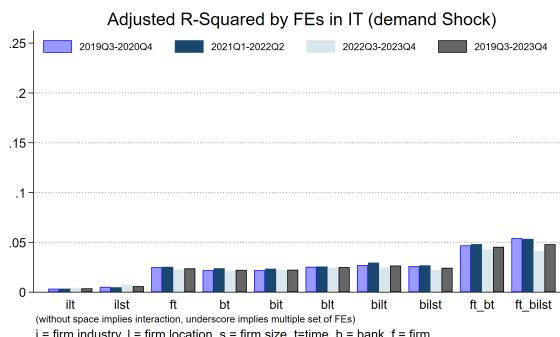
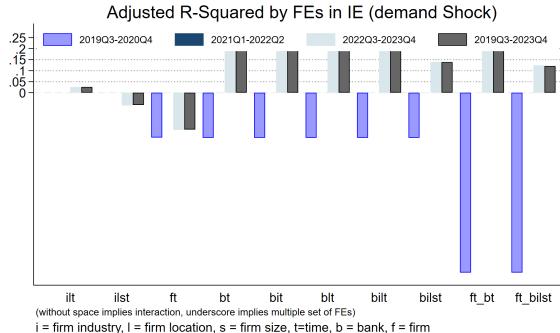
Notes: This graph presents the unadjusted R^2 values from regressing time-varying firm-bank specific demand and supply shocks on sets of fixed effects. Data is pooled over all countries. The fixed effects regressions are done separately for each quarter, over the period 2019Q3 to 2023Q4. The different lines depict R^2 values for various fixed effects combinations, including Bank-Time (BT), Industry-Location-Size-Time (ILST), Bank-Industry-Location-Size-Time (BILST), Firm-Time (FT), and two way fixed effects combining either FT and BT or FT and BILST in one set-up. The y-axis shows the R^2 values, with dotted reference lines at 0.2 intervals. The left panel considers demand shocks and the right supply.

Figure 7: Variation in Demand and Supply shocks explained by Fixed Effects: three 6 quarter periods and full period

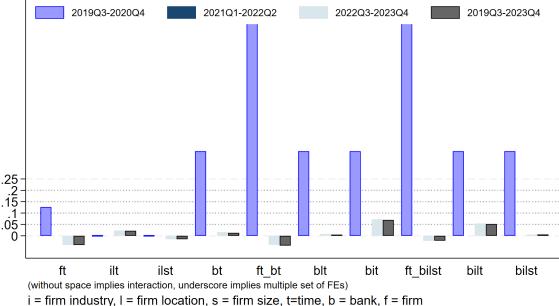
Notes: These graphs display adjusted R^2 values from regressing time-varying firm-bank specific demand (left) or supply (right) shocks on various sets of fixed effects (FEs). The first line of plots corresponds with pooling all countries. Each subsequent row corresponds with one country. For each subplot, we report the adjusted R^2 for 10 different sets of (combinations of fixed effects) across three six-quarter periods and the full sample period (2019Q3-2023Q4). The three six quarter periods are: 2019Q3-2020Q4, 2021Q1-2022Q2 and 2022Q3-2023Q4. Colors represent different time periods, with the entire period (2019Q3-2023Q4) shown in dark gray. The FEs are respectively (1) industry-location-time (ILT), (2) bank-time (BT), (3) industry-location-size time (ILST), (4) bank-location-time (BLT), (5) bank-industry-time (BIT), (6) bank-industry-location-time (BILT), (7) bank-industry-location-size-time (BILST), (8) firm-time (FT), (9) FT and BT, (10) FT and BILST. Time (T) is the quarterly frequency, where size is quintiles of the firm size distribution (in a BILT bin). The y-axis shows the R^2 values, with dotted reference lines at 0.2 intervals.



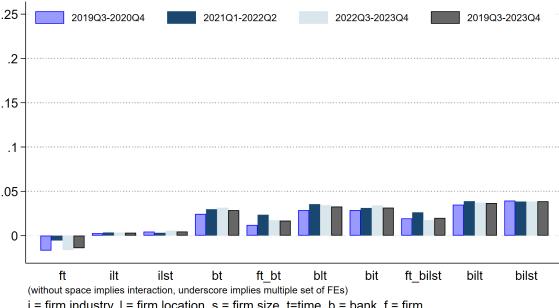




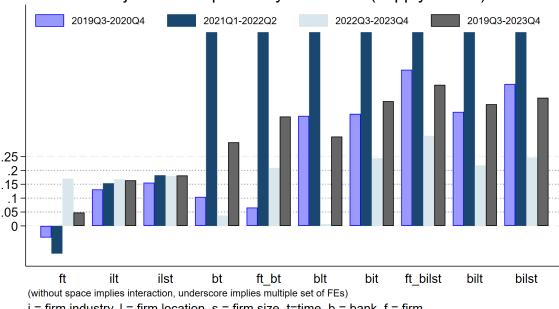
Adjusted R-Squared by FEs in IE (supply Shock)



Adjusted R-Squared by FEs in IT (supply Shock)



Adjusted R-Squared by FEs in NL (supply Shock)



Adjusted R-Squared by FEs in PT (supply Shock)

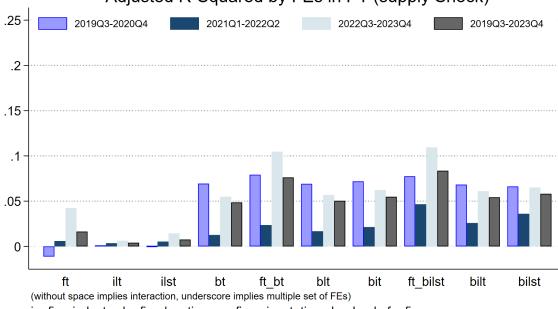


Table 10: Baseline: robustness to standard error clustering and fixed effects

	(1)	(2)	(3)	(4)	(5)
	Demand innovation (f,b,t)				
Share of fixed rate loans (f,b,t-1)	-0.162*** (0.019)	-0.162*** (0.017)	-0.162*** (0.001)	-0.162*** (0.019)	-0.157*** (0.018)
Share of collateralized loans (f,b,t-1)	0.016 (0.016)	0.016** (0.008)	0.016*** (0.001)	0.016 (0.016)	0.012 (0.013)
Share of Credit lines and Term Loans (f,b,t-1)	0.086*** (0.017)	0.086*** (0.016)	0.086*** (0.001)	0.086*** (0.017)	0.088*** (0.016)
Share of bank in a firm's overall borrowing (f,b,t-1)	-0.206*** (0.023)	-0.206*** (0.012)	-0.206*** (0.002)	-0.206*** (0.023)	-0.223*** (0.022)
R-squared	0.52	0.52	0.52	0.52	0.43
Adjusted R-squared	0.04	0.04	0.04	0.04	0.08
	Supply innovation (f,b,t)				
Share of fixed rate loans (f,b,t-1)	0.156*** (0.021)	0.156*** (0.018)	0.156*** (0.001)	0.156*** (0.021)	0.158*** (0.016)
Share of collateralized loans (f,b,t-1)	0.034*** (0.013)	0.034*** (0.008)	0.034*** (0.001)	0.034*** (0.013)	0.040*** (0.012)
Share of Credit lines and Term Loans (f,b,t-1)	-0.150*** (0.030)	-0.150*** (0.015)	-0.150*** (0.002)	-0.150*** (0.030)	-0.131*** (0.028)
Share of bank in a firm's overall borrowing (f,b,t-1)	-0.337*** (0.019)	-0.337*** (0.017)	-0.337*** (0.002)	-0.337*** (0.019)	-0.298*** (0.019)
R-squared	0.52	0.52	0.52	0.52	0.43
Adjusted R-squared	0.04	0.04	0.04	0.04	0.07
Observations	10477109	10477109	10477109	10477109	12752210
Firm \times Time FE	Yes	Yes	Yes	Yes	Yes
Bank \times Time FE	-	-	-	-	Yes
Bank \times Industry \times Location \times Time FE	Yes	Yes	Yes	Yes	-
SE-cluster1	Bank	Bankx-Time	Firm	Bank	Bank
SE-cluster2	-	-	-	Firm	-

Notes: Time-varying innovations to firm-bank specific **credit demand** (upper panel) or **credit supply** (lower panel) are regressed on four time-varying bank-firm characteristics. These are the share of borrowing by firm f from bank b at time $t-1$ that is (1) fixed rate, (2) secured by collateral, or (3) credit line or term loan borrowing (with the remaining part being mainly revolving credit). The fourth characteristic is the share of firm f 's borrowing from bank b at time $t-1$ in firm f 's overall borrowing at time $t-1$. We use an unbalanced sample of firm-bank level observations over the period 2019Q3-2023Q4 for 11 euro area countries. All specifications include firm-quarter fixed effects. In columns 1 to 4, we include bank-industry-location-quarter fixed effects. Columns 1 to 4 differ in the extent of clustering the standard errors. We cluster the standard errors at the bank level (column 1), bank-quarter level (column 2), firm level (column 3). In column 4, standard errors are obtained from a weighted sum of two standard cluster variances, one clustered at the firm level, and the other at the bank level (see Appendix E). In column 5, we replace the bank-industry-location-quarter fixed effects with the less granular bank-quarter fixed effects.

Table 11: The impact of monetary policy, central bank information and macroprudential policy

	(1) Credit growth (f,b,t)	(2) Change in Interest Rate (f,b,t)	(3) Demand innovation(f,b,t)	(4) Supply innovation(f,b,t)
Probability of Default (f,b,t)	-0.260*** (0.039)	0.079*** (0.024)	-0.006 (0.026)	-0.185*** (0.050)
Monetary Policy (t) \times Probability of Default (f,b,t)	0.064 (0.261)	0.789*** (0.195)	-0.044 (0.203)	-0.446 (0.307)
Central Bank Information (t) \times Probability of Default (f,b,t)	-1.120*** (0.349)	-0.255 (0.244)	-2.271*** (0.574)	-0.162 (0.329)
Share of fixed rate loans (f,b,t-1)	0.064*** (0.021)	-0.181*** (0.017)	-0.125*** (0.021)	0.141*** (0.018)
Monetary Policy (t) \times Share of fixed rate loans (f,b,t-1)	0.260** (0.114)	-0.922*** (0.122)	-0.533*** (0.127)	0.725*** (0.237)
Central Bank Information (t) \times Share of fixed rate loans (f,b,t-1)	-0.196** (0.079)	-0.588*** (0.192)	0.033 (0.313)	0.699* (0.365)
Quarterly Change in Macro-Prudential index (t) \times Probability of Default (f,b,t)	-0.016 (0.016)	0.043*** (0.010)	0.037*** (0.010)	-0.047* (0.025)
Quarterly Change in Macro-Prudential index (t) \times Share of fixed rate loans (f,b,t-1)	-0.008*** (0.003)	-0.015 (0.009)	-0.017* (0.010)	0.003 (0.007)
Observations	5899787	5899787	5899787	5899787
R-squared	0.51	0.53	0.52	0.51
Adjusted R-squared	-0.01	0.04	0.03	0.01
Firm \times Time FE	Yes	Yes	Yes	Yes
Bank \times Industry \times Location \times Time FE	Yes	Yes	Yes	Yes
SE-cluster1	Bank	Bank	Bank	Bank
SE-cluster2	Firm	Firm	Firm	Firm
Sample	201909-202312	201909-202312	201909-202312	201909-202312
Coverage	11 countries	11 countries	11 countries	11 countries

2

Notes: This table is similar in design to Table 8, but uses the specification as in columns 3 and 6 of Table 7. More specifically, the table shows the regression results obtained when regressing (i) Credit growth (column 1), (ii) Interest Rate changes (column 2), (iii) Credit demand innovations (column 3) and (iv) Credit supply innovations (column 4) on two firm-bank characteristics as well as their interactions with monetary policy shocks, central bank information shocks, and changes in macroprudential policy. The two firm-bank characteristics are a bank-specific assessment of firm credit risk and the share of borrowing by firm f from bank b at time t-1 that is fixed rate. The Monetary Policy and Central Bank Information shocks are obtained from Jarociński and Karadi (2020). Changes in macro-prudential policy are measured as the quarterly change in the index of 17 policy-action indicators from the IMF's integrated Macroprudential Policy (iMaPP) Database, originally constructed by Alam et al. (2019). All specifications include firm-quarter fixed effects as well as bank-industry-location-quarter fixed effects. Standard errors are clustered at the bank and firm level in columns 1 and 2. In columns 3 and 4, Standard errors are obtained from a weighted sum of two standard cluster variances, one clustered at the firm level and the other at the bank level (see Appendix E). The model only includes loans of fixed rate, short-term, and the