Identifying relationship-level effects using covariance restrictions

Olivier De Jonghe
Daniel Lewis

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP12/26

Identifying Relationship-level Effects Using Covariance Restrictions *

Olivier De Jonghe† Daniel Lewis‡

July 14, 2026

Abstract

We propose a novel decomposition to identify relationship-specific effects or shocks in a bipartite network under covariance restrictions. We show existing two-way fixed effects decompositions are ill-suited to correlations consistent with realistic heterogeneity. Our strategy yields a simple consistent estimator. We estimate relationship-level credit demand and supply shocks across nine euro-area countries and three episodes. We find evidence demand and supply each have both firm and bank components, and within-firm/bank shock variation is of comparable scale to between-firm/bank variation. Regressions using firm fixed effects as demand controls are at odds with economic theory, while our method uncovers significant deleterious effects of the post-2022 monetary contraction on exposed firms.

Keywords: two-way fixed effects, supply shock, demand shock, corporate credit, identification, higher moments, networks
JEL codes: C33, C58, E44, G21, G30

*An earlier version of this paper circulated as "Identifying Heterogeneous Supply and Demand Shocks in European Credit Markets". 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. Funded by the European Union (ERC, CreMac, StG-101161596). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. We are grateful to Gabriel Chodorow-Reich, Ugo Albertazzi (discussant), Diana Bonfim, Max Bruche (discussant), Ferre De Graeve, Hans Degryse, Aureo de Paula, Martina Jasova (discussant), Peter Jiang, Edoardo Rainone (discussant), Enrico Sette, and Andrei Zeleneev for helpful discussions, as well as seminar participants at UCL, the Banque de France, the Bank of England, the European Central Bank, the Federal Reserve Bank of New York, Penn State, Paris School of Business, the University of Cyprus, and the University of East Anglia for helpful comments, conference participants at the Barcelona Summer Forum (Structural Shocks Identification), Encounters in Econometric Theory, the Leuven Summer Event (Macro), the NBER Summer Institute (EFFE), the 7th Summer Conference on Financial Intermediation and Corporate Finance, the ECB-MPC Taskforce on Banking Analysis for Monetary Policy, IBEFA (ASSA meetings 2026), and the UEA Time Series Workshop.

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

‡University College London, CEPR (corresponding author: daniel.j.lewis@ucl.ac.uk).

1 Introduction

Many econometric settings require the decomposition of the outcome of an observed match into two components, with one associated with each side of a market. In a seminal paper, Abowd et al. (1999) (henceforth AKM) decompose log wages into worker and employer characteristics using two-way fixed effects. This decomposition problem applies broadly across bipartite networks, where two distinct types of agents interact exclusively with agents of the opposite type (see Bonhomme 2020 for a review). For instance, Khwaja and Mian (2008) and Amiti and Weinstein (2018) use fixed effects to decompose firm-bank credit relationships into credit demand and supply shocks, while Abowd et al. (2004) consider relationships between importers and exporters. Similar empirical challenges and applications abound throughout corporate finance, production networks, and international trade.

The majority of subsequent reduced-form research has adopted the AKM framework, relying on two-way fixed effects to recover agent-level components. However, these decompositions are premised on homogeneous effects, which are uniform across a worker's employers or a bank's credit relationships. Recent work questions this assumption across multiple literatures. Borovičková and Shimer (2024) and Bonhomme et al. (2023) demonstrate that sorting, worker mobility, and assortative matching can violate uniform worker-firm effects. In credit markets, Paravisini et al. (2023) document bank specialization across export markets, while Chodorow-Reich (2014) shows the empirical importance of relationship lending. Furthermore, banks may ration credit based on correlated firm riskiness, and firms uniquely value specific lending relationships (e.g., Anderson et al. 2022), plausibly generating meaningful heterogeneity in supply and demand shocks. Moreover, a model with homogeneous shocks precludes crucial policy questions, such as how credit supply propagates differentially across heterogeneous firms following regulatory or monetary interventions.

We propose an alternative decomposition to identify heterogeneous relationship-level effects. To do so, we introduce a second observable variable. In a labour setting, this could be match duration studied alongside wages; in credit markets, it is natural to study the interest rate on credit alongside its volume. We model this bivariate, relationship-level vector as a linear combination of two relationship-level effects, such as match-specific worker and employer effects or relationship-level demand and supply shocks. Rather than imposing homogeneity across relationships, we assume that these effects are *correlated*. Crucially, we permit *two-way correlation*, meaning for instance that demand shocks can be correlated not just within a firm but also across firms, for the same bank. We illustrate how this realistic feature naturally leads AKM-style decompositions that rely on two-way fixed effects to conflate demand and supply effects. Our identification strategy exploits the cross-sectional covariance of the observable vector across both firms and banks, adapting the intuition of identification via heteroskedasticity from the time-series literature (Rigobon 2003; see Lewis 2025 for a review) to a cross-sectional network setting.

Our strategy gives rise to a simple, closed-form estimator of the matrix linking observables and relationship-level effects, and we establish its consistency and asymptotic normality. We also provide a test for the presence of two-way correlation, under which traditional fixed effects approaches may fail to separate demand and supply. Our estimator is consistent provided a non-vanishing share of agents on each side of the market have degree (number of relationships) increasing with the number of opposite-type agents. Under this assumption, we decompose all relationships with vanishing error, even for agents with only two relationships. Recently, Jochmans and Weidner (2019) established that the consistency of traditional two-way fixed effects estimates requires the number of connections per agent and their immediate neighbors to grow to infinity. Our requirement is more plausible in our empirical setting where most multi-bank firms borrow from only two or three lenders, while a few large firms maintain extensive banking relationships. Inference adapts standard variance estimators, and we provide Stata/Python routines for implementation. Empirically-calibrated simulations (based on the Italian credit market) demonstrate strong finite-sample performance.

We view our decomposition as an alternative to AKM; each requires different assumptions that will be more plausible in some settings than others. Our approach hinges on two key structural assumptions that reflect its limitations. First, we assume the same coefficient matrix maps observables to underlying shocks across all relationships, though the scale of the shocks themselves can vary. While potentially restrictive, the econometrician can weaken this assumption by conditioning on narrower sample slices, such as limiting attention to similar firm sizes within a given industry. Moreover, this restriction must be weighed against the similarly stark assumption of standard decompositions that all shocks are equal across relationships for a given agent. Second, our decomposition is defined by the assumption that effects of opposite type (demand and supply) are orthogonal across relationships, aligning with macroeconomic notions of structural shocks. While we argue that this condition is suitable for intensive-margin variation, it requires careful interpretation when network formation is involved and does not allow us to study aggregate time-series variation. Finally, like AKM, our approach is applicable when agents have multiple relationships within the sample to leverage cross-relationship covariance.

Throughout the paper, we interpret our methodology and assumptions through the lens of our empirical application, the supply and demand of corporate credit. This setting easily maps to a range of other financial market contexts. However, the methodology also extends to labor markets, product markets, and broader industrial organization settings where price and quantity data are available across multi-agent matches, such as Nielsen consumer panels.

We apply our identification strategy to recover quarterly relationship-level demand and supply shocks for thousands of firms and banks across nine euro-area credit markets from 2019 to 2023 using the AnaCredit dataset. We establish seven key sets of empirical facts. First, we document substantial cross-country heterogeneity in the slopes of credit demand and supply curves, though they remain stable within countries over time. Second, under our assumptions, we statistically reject the one-way correlation restriction required for AKM decompositions to disentangle demand and supply in the vast majority of country-periods. Third, we find pervasive heterogeneity across relationships, with within-firm and within-bank variation in demand and supply shocks being substantial, on the same order of magnitude as variation between firms and banks. Fourth, regressions using AKM-based measures yield results at odds with economic theory; demand effects appear contaminated by supply shocks, with a positive demand shock *lowering* interest rates. Fifth, as a result, controlling for AKM-style demand measures in regressions, following Khwaja and Mian (2008), attenuates the estimated impact of the post-2022 monetary contraction, whereas controlling for our demand measures uncovers a significant and substantial reduction in credit growth for exposed firms borrowing at floating or short-term rates. Sixth, this supply-driven borrowing contraction leads to significantly worse real firm-level outcomes in terms of total assets, turnover, and credit-to-asset ratios, which are similarly attenuated when using AKM controls. Finally, following identified monetary policy shocks, *relationships* with higher shares of fixed-rate loans exhibit lower relative credit demand, but the relative supply of credit to these same firms increases, consistent with banks assessing lower interest rate risk. These empirical results illustrate how our alternative decomposition can speak to important policy questions regarding the distributional and real effects of contractionary monetary policy actions — insights that may be lost under aggregate or homogeneous shock assumptions.

Our paper connects to several strands of literature. Existing approaches to incorporating heterogeneity into decompositions of this type have typically relied on interactions or other functions of fixed or random effects, which remain homogeneous, to do so (e.g., Chamberlain 1980; Graham et al. 2014; Amiti and Weinstein 2018; Crippa 2025). Our model maintains a tractable linear structure by embedding heterogeneous relationship-level shocks directly. Our identification strategy also relates to the granular instrumental variables approach of Gabaix and Koijen (2024), which similarly exploits idiosyncratic shock variation. However, while granular identification leverages size-weighted idiosyncratic demand to instrument for market-level aggregates under a uniform price, our approach exploits the variance within an individual firm's relationships and requires price variation at the relationship level.

We differ from several contemporaneous extensions of the AKM model. Crippa (2025) introduces complementarities via an interaction function but maintains homogeneous unit-level effects. There is a growing contemporaneous literature questioning the suitability of AKM decompositions specifically for credit markets. Gutierrez et al. (2025) apply a grouped fixed effects estimator (Bonhomme and Manresa 2015) to allow for latent group-specific supply shocks, but they retain homogeneous firm demand and require many relationships for each firm to consistently estimate group membership. Bruche et al. (2026) focus on firm preferences over specific lenders, measured using a multinomial logit framework based on observables, but do not generate relationship-level demand shocks or speak to supply heterogeneity. After constructing a structural banking model to isolate supply substitution, Bergman et al. (2025) make similar observations to ours about how AKM-type decompositions may not recover true relationship-level effects in the face of heterogeneity and propose estimators to recover the effects of both supply shifters and total supply shocks on firm outcomes. However, their framework restricts attention to two-bank firms and is silent on demand-side heterogeneity. Lastly, Pietrosanti and Rainone (2026) provide instrumental variables valid in the presence of network spillovers, which can be viewed as endogenous responses to structural shocks and can contaminate AKM-type decompositions; our framework can accommodate these interactions on the intensive margin.

Empirically, we build directly on the literature pioneered by Khwaja and Mian (2008), and pursued by Amiti and Weinstein (2018), Greenstone et al. (2020), Cingano et al. (2016), and many others, which principally uses fixed effects-based measures of demand and/or supply shocks as regressors. For instance, Khwaja and Mian (2008) exploit an exogenous event impacting bank liquidity to estimate heterogeneous credit supply shocks to firms and their effects, while controlling for homogeneous credit demand with firm-time fixed effects. By relaxing the common foundational assumption of relationship homogeneity, our methodology provides an alternative framework for identifying relationship-level effects or shocks in matched data environments.

The remainder of the paper is organized as follows. Section 2 provides a simple example showing how AKM-type decompositions can fail to recover true effects in the presence of heterogeneous shocks and details the intuition underlying our strategy. Section 3 presents the formal econometric framework, asymptotic properties, and Monte Carlo simulations. Section 4 applies our approach to euro-area credit markets and contrasts the results with those based on AKM-derived benchmarks. Section 5 concludes.

2 The challenge of heterogeneous shocks

In this section, we present a simple example to illustrate both a limitation of AKM-style decompositions and the intuition behind our identification strategy. Following our empirical application, we frame this illustration within a corporate credit market, though the underlying results generalize directly to any bipartite matching setting. Let $u_{fb} = (u^d_{fb}, u^s_{fb})'$ denote the relationship-specific demand and supply shocks for firm f and bank b. Consider an economy with two firms and two banks. Assume bank 1 is specialized – for instance, in finance for export – and that a change in trade policy has led both firms to expand their export operations. Suppose that

$$
\begin{aligned}
u^d_{f1} &= d_f + a, \quad && f \in \{1,2\}, \\
u^d_{f2} &= d_f - a, \quad && f \in \{1,2\}, \\
u^s_{fb} &= s_b + e_f \quad && f \in \{1,2\}, b \in \{1,2\}.
\end{aligned}
$$

Here, $a > 0$ reflects a common bank-specific demand component driven by export specialization, resulting in a substitution effect from bank 2 to bank 1.[1] The parameter d_f represents a firm demand effect, s_b represents a bank supply effect, and e_f captures a firm-specific component of supply, due to both banks similarly assessing firms' risk profiles, say. For expositional clarity, we omit idiosyncratic errors. If $a = e_f = 0$, this setup collapses exactly to a standard AKM decomposition where demand shocks are purely firm-specific (d_f) and supply shocks are purely bank-specific (s_b). Otherwise, demand includes a bank-specific component and supply includes a firm-specific component, with both representing realistic features of credit market heterogeneity, which we refer to as *two-way correlation*. We show below that two-way correlation causes traditional two-way fixed effects estimators to conflate demand and supply.

To demonstrate this fact, assume for simplicity the change in loan quantity between firm f and bank b is the sum of these shocks, $\Delta l_{fb} = u^d_{fb} + u^s_{fb}$, yielding the system

$$
\begin{aligned}
\Delta l_{f1} &= d_f + s_1 + e_f + a, \\
\Delta l_{f2} &= d_f + s_2 + e_f - a.
\end{aligned}
$$

Fitting firm and bank fixed effects to Δl_{fb} as an AKM decomposition (using firm 2 as the

[1] The symmetric impact of a on u^d_{f1} and u^d_{f2} is chosen for expositional simplicity and can be relaxed without altering our conclusions.

reference group) identifies the objects

$$\tilde{d}_1 = (d_1 - d_2) + (e_1 - e_2)$$
$$\tilde{s}_1 = s_1 + a + (d_2 + e_2)$$
$$\tilde{s}_2 = s_2 - a + (d_2 + e_2).$$

Evidently, the AKM decomposition fails to identify either the firm-specific common components of demand or the bank-specific common components of supply. Because firm 2 is the reference unit, the estimated firm effect $\tilde{d}_1$ is intended to capture firm 1's relative demand component, $d_1 - d_2$. Instead, it is contaminated by the relative firm-specific supply component, $(e_1 - e_2)$. Similarly, the estimated bank supply effects $\tilde{s}_b$ are contaminated by the bank-specific demand shock $\pm a$ and the reference unit-specific effects. Even if these reference firm components are normalized to zero ($d_2 = e_2 = 0$), the parameters are not identified: $\tilde{d}_1 = d_1 + e_1$, $\tilde{s}_1 = s_1 + a$, and $\tilde{s}_2 = s_2 - a$. Thus, an AKM-style decomposition may fail to recover the true shocks when there is correlation in demand for a specific bank across firms, or any correlation in supply to a specific firm across banks. Appendix C.1 demonstrates that this result holds symmetrically if a bank is selected as the reference unit. This finding complements similar arguments based on a structural credit model in Bergman et al. (2025), which highlight how supply substitution can distort AKM demand controls. Our reduced-form argument shows that this limitation extends to supply itself as well as demand and generalizes across network settings.

However, rather than impeding identification, these cross-relationship correlations can actually be leveraged to separate the underlying shocks if a second observable variable is introduced. Suppose we observe both the change in the interest rate (Δr_{fb}) and the loan quantity (Δl_{fb}), linked to the underlying shocks through the system:

$$\begin{pmatrix} \Delta r_{fb} \\ \Delta l_{fb} \end{pmatrix} = A u_{fb}, \;\; A = \begin{bmatrix} 1 & A_{12} \\ A_{21} & 1 \end{bmatrix}, \tag{1}$$

where the matrix A encodes the relative response of each observable to the relationship-level shocks, and the unit diagonal normalizes the scale.[2] Standard economic theory implies that both shocks positively impact loan volume ($A_{21} > 0$), while demand and supply exert opposite forces on the interest rate ($A_{12} < 0$). If the parameter matrix A can be identified, the shocks can be fully recovered by inverting (1).

Let a, d_f, s_b, and e_f be mutually orthogonal, mean-zero random variables with variances

[2] Without a scale normalization, it is impossible to separate the magnitude of the underlying shocks from the scale of their impacts.

$\sigma_a^2, \sigma_d^2, \sigma_s^2$, and σ_e^2. For this stylized example, we impose the additional restriction that $\sigma_s^2 = \sigma_e^2$, meaning the bank-specific and firm-specific components of supply share equal variance. Consider four cross-sectional moments:

$$\begin{aligned}
E\left[\Delta r_{11}\Delta l_{21}\right] &= A_{21}\sigma_a^2 + A_{12}\sigma_s^2, \\
E\left[\Delta r_{11}\Delta l_{12}\right] &= A_{21}(\sigma_d^2 - \sigma_a^2) + A_{12}\sigma_e^2, \\
E\left[\Delta r_{11}\Delta r_{21}\right] &= \sigma_a^2 + A_{12}^2\sigma_s^2, \\
E\left[\Delta r_{11}\Delta r_{12}\right] &= (\sigma_d^2 - \sigma_a^2) + A_{12}^2\sigma_e^2.
\end{aligned}$$

These moments capture the cross-firm and cross-bank covariances of interest rates and quantities, holding one side of the market fixed. Taking the ratio of the differences between the cross-firm and cross-bank covariances identifies A_{21}:

$$\begin{aligned}
\frac{E\left[\Delta r_{11}\Delta l_{21}\right] - E\left[\Delta r_{11}\Delta l_{12}\right]}{E\left[\Delta r_{11}\Delta r_{21}\right] - E\left[\Delta r_{11}\Delta r_{12}\right]} &= \frac{A_{21}\sigma_a^2 + A_{12}\sigma_s^2 - (A_{21}(\sigma_d^2 - \sigma_a^2) + A_{12}\sigma_e^2)}{\sigma_a^2 + A_{12}^2\sigma_s^2 - ((\sigma_d^2 - \sigma_a^2) + A_{12}^2\sigma_e^2)} \\
&= \frac{A_{21}(2\sigma_a^2 - \sigma_d^2)}{2\sigma_a^2 - \sigma_d^2} \\
&= A_{21}.
\end{aligned} \tag{2}$$

This closed-form expression maps directly to an instrumental variables (IV) strategy, where the relationship-specific interest rate Δr_{11} is an instrument in the regression of $(\Delta l_{21} - \Delta l_{12})$ on $(\Delta r_{21} - \Delta r_{12})$:

$$\frac{E\left[\Delta r_{11}(\Delta l_{21} - \Delta l_{12})\right]}{E\left[\Delta r_{11}(\Delta r_{21} - \Delta r_{12})\right]} = \frac{E\left[\Delta r_{11}\Delta l_{21}\right] - E\left[\Delta r_{11}\Delta l_{12}\right]}{E\left[\Delta r_{11}\Delta r_{21}\right] - E\left[\Delta r_{11}\Delta r_{12}\right]}.$$

The corresponding exogeneity condition requires

$$\begin{aligned}
&E\left[\Delta r_{11}\left((s_1 - s_2 + e_2 - e_1) - A_{21}A_{12}(s_1 - s_2 + e_2 - e_1)\right)\right] \\
&= E\left[\Delta r_{11}(1 - A_{21}A_{12})(s_1 - s_2 + e_2 - e_1)\right] \\
&= (1 - A_{21}A_{12})E\left[(a + d_1 - s_1 - e_1)(s_1 - s_2 + e_2 - e_1)\right] \\
&= (1 - A_{21}A_{12})(\sigma_s^2 - \sigma_e^2) \\
&= 0,
\end{aligned}$$

which holds exactly under the restriction $\sigma_s^2 = \sigma_e^2$. Intuitively, differencing – or comparing outcomes across relationships – systematically eliminates the endogeneity induced by the shared components of supply. The instrument is relevant provided $2\sigma_a^2 - \sigma_d^2 \neq 0$. Once A_{21}

is recovered, A_{12} is immediately identified from the covariance of $(\Delta r_{fb},\ \Delta l_{fb})'$, for instance, allowing the econometrician to invert A in (1) and recover the vector u_{fb}.

The assumption that $\sigma_s^2 = \sigma_e^2$ is restrictive; we adopt it here only to endow our strategy with an IV interpretation. Our general formulation below imposes no restrictions of this type, but, while identification is still achieved in closed form, it no longer boasts an IV interpretation. The problem above is also over-identified, even in this simple 2-firm, 2-bank example.[3] The estimator we propose exploits possible combinations of firms and banks connected in the network to increase precision.

3 A decomposition of heterogeneous effects

In this section, we present our general identification result and the corresponding estimator. We interpret the identification assumptions through the lens of our illustrative application. We next characterize the estimator's asymptotic distribution and inference. We compare the assumptions to those required for AKM to decompose the true shocks. Finally, we summarize the results of an empirically calibrated Monte Carlo study found in Appendix D.

3.1 The AKM decomposition

The classical AKM framework (see e.g., Bonhomme 2020 for a review) models an observable match outcome, such as the change in loan quantity Δl_{fb}, by decomposing it into homogeneous firm and bank components meant to capture demand and supply shocks:

$$\Delta l_{fb} = d_f + s_b + \Gamma X_{fb} + \varepsilon_{fb},\ \forall \{f, b\} : D_{fb} = 1, f = 1, \ldots, F,\ b = 1, \ldots, B,$$

where D_{fb} is an indicator for whether a relationship exists between f and b. To identify the firm-specific demand effect, d_f, and bank-specific supply effect, s_b, each agent must have at least two relationships. The idiosyncratic error term ε_{fb} is assumed to be mean-zero and independently distributed across relationships. Strict exogeneity is typically assumed, $E\left[\varepsilon_{fb} | D, \mathbf{X}, \mathbf{d}, \mathbf{s}\right] = 0$, where the vectors $\mathbf{d}$ and $\mathbf{s}$ stack the realized agent effects and $\mathbf{X}$ stacks the observables. Under this exogeneity assumption, the parameter vector Γ can be consistently estimated, and the covariates X_{fb} can be partialed out via a standard Frisch-Waugh-Lovell argument. This allows the researcher to focus on the simplified model:

$$\Delta l_{fb} = d_f + s_b + \varepsilon_{fb}. \tag{3}$$

[3] Note that we could permute either or both of the firm and bank indices to construct a new IV estimator.

The AKM decomposition defines firm demand, d_f, as invariant across all lenders from which a firm borrows, and bank supply, s_b, as invariant across all of its borrowers. As illustrated in the preceding section, if the true shocks are heterogeneous and exhibit two-way correlation, these fixed effects cease to align with the latent economic shocks, since the imposed structure is misaligned with the DGP. Put differently, for an AKM decomposition to recover demand and supply shocks, they must be homogeneous and all correlation in demand effects across relationships must be captured by d_f and all correlation in supply captured by s_b.

3.2 A new relationship-level decomposition

We consider a vector of relationship-level observables, $\eta_{fb} \equiv (\Delta r_{fb}, \Delta l_{fb})'$, for instance, the change in the interest rate and loan quantity between firm f and bank b. These observables are jointly determined by two relationship-level effects, $u_{fb} \equiv (u^d_{fb}, u^s_{fb})'$, (demand and supply shocks):

$$\eta_{fb} \equiv \begin{pmatrix} \Delta r_{fb} \\ \Delta l_{fb} \end{pmatrix} = A \begin{pmatrix} u^d_{fb} \\ u^s_{fb} \end{pmatrix} = A u_{fb}, \ \forall \{f, b\} : D_{fb} = 1, f = 1, \ldots, F, \ b = 1, \ldots, B, \tag{4}$$

where the matrix A is left unrestricted. While (4) models η_{fb} as an exact combination of the two shocks, additional relationship-level idiosyncratic factors can be accommodated, as explained in Section 3.5. We focus our development on intensive-margin variation on a static, pre-determined network structure, consistent with firms that may draw down lending facilities with certain banks, but rarely terminate relationships entirely. We thus abstract from network formation and destruction and condition our analysis on the observed network structure in the realized sample, $D = \{D_{fb}\}$. We discuss modifications to accommodate matching in the context of a labour market example below.

Assumption 1.

1. *The matrix A is invertible and constant across relationships,*
2. $\operatorname{cov}\left(u^d_{fb}, u^s_{fb'}\right) = 0, \ b' \neq b,$
3. $\operatorname{cov}\left(u^d_{fb}, u^s_{f'b}\right) = 0, \ f' \neq f.$

A encodes the response coefficients of both observables to the effects. The cross-relationship orthogonality conditions are definitional for the decomposition of the match outcome into distinct relationship-level effects. If u_{fb} are interpreted as shocks, it is natural to further assume that they are mean-zero. We expand on these assumptions below.

To identify A, and subsequently recover the latent relationship-level shocks via $u_{fb} = A^{-1}\eta_{fb}$, we exploit two distinct sets of cross-sectional covariance equations. The first captures the covariance of η_{fb} across firms while holding the bank fixed,

$$\text{cov}\left(\eta_{fb}, \eta'_{f'b}\right) \equiv \Sigma_{FF} = A\Lambda_{FF}A',$$

where $\Lambda_{FF} = \begin{bmatrix} \text{cov}(u^d_{fb}, u^d_{f'b}) & 0 \\ 0 & \text{cov}(u^s_{fb}, u^s_{f'b}) \end{bmatrix}$. The second captures the covariance across banks while holding the firm fixed:

$$\text{cov}\left(\eta_{fb}, \eta'_{fb'}\right) \equiv \Sigma_{BB} = A\Lambda_{BB}A',$$

where Λ_{BB} is defined symmetrically. After imposing a standard scale normalization, such as setting the diagonal elements of A to unity, or normalizing $\Lambda_{FF} = I_2$, there are six free parameters stacked in the vector θ. Together, these covariance structures yield six equations:

$$m(\theta) = \begin{bmatrix} \text{vech}\left(\Sigma_{FF} - A\Lambda_{FF}A'\right) \\ \text{vech}\left(\Sigma_{BB} - A\Lambda_{BB}A'\right) \end{bmatrix} = 0. \tag{5}$$

Finally, we require that the covariance of effects across firms is not proportional to the covariance of effects across banks:

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

The following result establishes identification:

Proposition 1. *If Assumptions 1 and 2 hold, then θ is the unique solution to (5) up to scale, sign, and column ordering of A.*

This identification strategy builds on simultaneous equation models identified via heteroskedasticity (e.g., Rigobon 2003). When the diagonals of Λ_{FF} and Λ_{BB} are non-zero, A is identified in closed form as the matrix of eigenvectors of $\Sigma_{FF}\Sigma_{BB}^{-1} = A\Lambda_{FF}\Lambda_{BB}^{-1}A^{-1}$. Intuitively, separately exploiting variations within a firm across banks, and within a bank across firms, provides two linearly independent sets of moments linked by A. While a single covariance matrix only determines A up to an arbitrary orthogonal rotation, a second linearly independent covariance matrix ensures uniqueness up to scale and column order. The matrix A directly summarizes the responses of interest rates and credit volumes to demand and supply shocks, tracing out the slopes of the curves. Once A is identified, the relationship-specific shocks u_{fb} are immediately recovered for each match by inverting (4).

We now interpret the identifying conditions in Assumption 1 in our leading corporate credit setting. Credit demand and supply shocks are likely heterogeneous due to lender specialization (e.g., Paravisini et al. 2023), substitution effects, or various portfolio adjustment incentives. Supply shocks may be correlated across firms for a given bank due to aggregate shifts in the bank's lending posture or credit reallocation (measured by $\text{cov}(u^s_{fb}, u^s_{f'b})$), and across banks for a given firm due to common assessments of the firm's underlying creditworthiness (measured by $\text{cov}(u^s_{fb}, u^s_{fb'})$). Symmetrically, demand shocks are correlated across banks for a given firm due to aggregate changes in the firm's financing needs or substitution behavior (measured by $\text{cov}(u^d_{fb}, u^d_{fb'})$), and across firms for a given bank due to bank specialization or product attractiveness (measured by $\text{cov}(u^d_{fb}, u^d_{f'b})$). These covariances populate the diagonals of the within-firm Λ_{FF} and within-bank Λ_{BB} shock covariance matrices, which, jointly with A, dictate the reduced-form price and quantity covariances, Σ_{FF} and Σ_{BB}.

Assumption 1.1 requires that the relative responses of the interest rate and loan quantity on the intensive margin to the underlying shocks are linear and invariant across the sample within a given time period. This condition need not hold intertemporally, since only a single time period is required for identification. If a researcher is concerned that the relative response elasticities vary across firm characteristics – such as industry, size, credit rating, collateral, or maturity – the estimation sample can be partitioned into narrower slices where this invariance is more appealing. For instance, in our application, we study the euro-area but allow for country-period-specific elasticities. In addition, in Appendix E.3, we demonstrate how, for a single country (Spain), attention can be limited further to subsamples by region, industry, or firm size. Crucially, this restriction does not require the *scale* of effects to be constant, only that the responses of price relative to those of quantity are fixed.

In macroeconomic settings (and random-effects models), demand and supply shocks like those in (4) are assumed to be orthogonal or mutually independent. Without this independence, they lack a structural interpretation, as they could either stem systematically from a common proximal cause or embody endogenous responses to one another. This interpretation implies that orthogonality holds even *within* a relationship; our cross-relationship orthogonality restrictions in Assumptions 1.2 and 1.3 are therefore weaker assumptions. For u_{fb} to be unpredictable structural shocks, X_{fb} will typically include lagged controls. While motivated by this construct, we interpret the conditions in a somewhat weaker manner: on average, there should be no *systematic* relationship between shocks of opposite types across relationships, so that these orthogonality conditions hold. There may be situations where demand and supply shocks positively comove, for example if a firm receives positive news about future productivity, raising both its demand for credit and banks' willingness to lend to it. But there may also be cases where, following negative news, a firm demands more credit, as banks' retract credit supply. Our assumptions do not rule either of these cases out, but rather state that, on average, such comovements are zero. These covariance restrictions are definitional to our decomposition, just as the homogeneity assumption, ruling out two-way correlation, is definitional to the AKM decomposition. Economic intuition clarifies the content of these restrictions. Assumption 1.2 ensures that Λ_{BB} is diagonal. This requires a reorientation or transmission delay, meaning bank b''s supply shock to firm f cannot instantaneously alter firm f's credit demand from bank b, and that no omitted common factor simultaneously drives both demand from and supply to a given firm. Symmetrically, Assumption 1.3 ensures that Λ_{FF} is diagonal. This condition is satisfied when firms are atomistic, so firm f's demand shock cannot directly alter bank b's supply decision toward firm f', and when cross-firm spillovers operate as endogenous equilibrium *responses* to the structural shocks rather than contemporaneous shock innovations themselves. This assumption generally rules out unobserved industry-wide shocks unless X_{fb} includes industry controls. Similar interpretations extend to other financial environments featuring bipartite matching or trade settings involving networks of importers and exporters.

Adapting this methodology to settings like the worker-firm environment of Abowd et al. (1999) requires accounting for outcomes of sequential matches rather than an intensive margin. There, the network structure, D, cannot be held fixed, so Assumptions 1.2 and 1.3 must be modified to depend explicitly on the realized network, yielding $\text{cov}(u^d_{fb}, u^s_{fb'}|D) = 0$ and $\text{cov}(u^d_{fb}, u^s_{f'b}|D) = 0$. Taking unconditionally orthogonal effects as definitional, these conditional covariance restrictions may be satisfied under several existing matching frameworks. Examples include unilateral selection models (a firm decides to hire a worker purely based on the firm's idiosyncratic demand for that worker's profile, and workers accept all offers, with worker effects determined *ex post*, see e.g., Bernard et al. 2019), independent mutual consent models (workers and firms consent to a match if the independent quality draw they observe surpasses a threshold, see, e.g., Comola and Fafchamps 2014), or selection strictly on observables included in X_{fb} (see, e.g., Graham 2017, Fafchamps and Gubert 2007).

3.3 Estimation

Estimation proceeds by replacing the population covariance matrices, Σ_{FF} and Σ_{BB}, with their sample analogues, S_{FF} and S_{BB}. Without loss of generality, we assume that the observable vector η_{fb} is mean-zero in-sample, which follows from either demeaning or projecting out relationship-level covariates, X_{fb}. The natural estimators are the sample averages across all unique pairwise combinations of firms sharing a common bank, and all banks sharing a common firm, respectively:

$$S_{FF} = \frac{1}{N_{FF}} \sum_{b=1}^{B} \sum_{f' \neq f} \eta_{fb} \eta'_{f'b}$$
$$S_{BB} = \frac{1}{N_{BB}} \sum_{f=1}^{F} \sum_{b' \neq b} \eta_{fb} \eta'_{fb'},$$

where $N_{FF} = \frac{1}{2} \sum_{b=1}^{B} F_b(F_b - 1)$ and $N_{BB} = \frac{1}{2} \sum_{f=1}^{F} B_f(B_f - 1)$. Here, F_b denotes the number of distinct firms connected to bank b, B_f represents the number of distinct banks connected to firm f, and the indices $f' \neq f$ and $b' \neq b$ indicate summations over all unique, unordered pairs of connected agents. Given S_{FF} and S_{BB}, the sample analogue of the moments $m(\theta)$ is given by

$$q(\boldsymbol{\eta}, \theta) = \begin{pmatrix} \operatorname{vech}\left(S_{FF} - A\Lambda_{FF}A'\right) \\ \operatorname{vech}\left(S_{BB} - A\Lambda_{BB}A'\right) \end{pmatrix},$$

where $\boldsymbol{\eta}$ stacks all observed relationship-level vectors η_{fb}. The minimum distance estimator $\hat{\theta}$ solves the just-identified sample system,

$$q(\boldsymbol{\eta}, \hat{\theta}) = 0. \tag{6}$$

An initial solution for the response matrix, $\tilde{A}$, is obtained in closed form as $\tilde{A} = \operatorname{evec}\left(S_{FF} S_{BB}^{-1}\right)$, where $\operatorname{evec}(\cdot)$ denotes the matrix of right eigenvectors of its argument. This spectral decomposition determines $\tilde{A}$ uniquely up to sign, scale, and column permutation.

To complete estimation, we must impose a normalization for A and establish a column ordering to label the latent shocks. We first normalize $\tilde{A}$ such that the covariance of both shocks across firms is unity, $\Lambda_{FF} = I_2$. We then select the column permutation of the-rescaled matrix that minimizes the distance to the reference matrix,

$$\begin{pmatrix} \operatorname{std}(\Delta r_{fb}) & -\operatorname{std}(\Delta r_{fb}) \\ \operatorname{std}(\Delta l_{fb}) & \operatorname{std}(\Delta l_{fb}) \end{pmatrix}.$$

This criterion signs and labels the columns to obtain the signed permutation of A that is most consistent with the first shock being a demand shock and the second a supply shock. The resulting transformed matrix defines our final estimator, $\hat{A}$. For the rescaled and labeled bank-level shock covariances, $\hat{\Lambda}_{BB} = \hat{A}^{-1} S_{BB} \hat{A}'^{-1}$. Appendix B provides further details.

3.4 Asymptotic properties

Deriving the asymptotic distribution of $\hat{A}$ requires regularity conditions on the network structure and error processes, which we formalize in Appendix A. To summarize, Assumption 3 decomposes each relation-specific shock into two mutually independent processes: one operating across firms and the other across banks. We require these components to possess finite eighth moments and non-degenerate covariance matrices. Assumption 4 governs the network structure, requiring a non-vanishing share of agents whose degrees grow proportionally with the number of agents on the opposite side of the market (B or F).[4] We permit the two sides of the market to grow at asymmetric rates, provide that neither side asymptotically dominates the other ($F/B^2 \to 0$ as $F, B \to \infty$ and vice versa).

Theorem 1. *Suppose* $\theta_0 \in \text{interior}(\Theta)$ *and* Θ *is compact. Under Assumptions 1 to 4, following appropriate column labeling and scale normalization,* $\hat{\theta}$ *satisfies*

1. **Consistency:** $\hat{\theta} \xrightarrow{p} \theta_0$;
2. **Asymptotic Normality:** $\tilde{\Phi}(\hat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, \mathbf{W})$;
3. **Variance Estimation:** $\hat{\mathbf{W}} \xrightarrow{p} \mathbf{W}$;

where the normalizing matrix is given by

$$\tilde{\Phi} = \left(\begin{bmatrix} \sqrt{B} & 0 \\ 0 & \sqrt{F} \end{bmatrix} \otimes I_3 \right) \hat{\Phi}, \quad \hat{\Phi} = \left. \frac{\partial q(\boldsymbol{\eta}, \theta)}{\partial \theta'} \right|_{\hat{\theta}},$$

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

$\hat{\mathbf{W}}_{BB}$ *is defined symmetrically, and* $\hat{\mathbf{W}} = \begin{bmatrix} \hat{\mathbf{W}}_{FF} & 0 \\ 0 & \hat{\mathbf{W}}_{BB} \end{bmatrix}$.

Theorem 1 establishes the consistency and asymptotic normality of $\hat{A}$, and an estimator for the asymptotic variance. In practice, the asymptotic covariance matrix of $\hat{\theta}$ is computed as $\Omega = \tilde{\Phi}^{-1}\mathbf{W}\tilde{\Phi}'^{-1}$. $\hat{\mathbf{W}}_{FF}$ and $\hat{\mathbf{W}}_{BB}$ are standard clustered variance matrices, at the bank and firm levels, respectively.

[4]This network restriction contrasts with the requirements for consistent estimation of individual fixed effects in traditional AKM settings. As shown by Jochmans and Weidner (2019), AKM requires the degree of all nodes to diverge to infinity. Our approach is also distinct from the variance decomposition methods in Kline et al. (2020), which do not speak to the effect estimates themselves and require at least two time periods with time-invariant agent effects to avoid more restrictive degree conditions.

Testing for two-way correlation. A consequence of Theorem 1 is that it provides a test for whether homogeneity or one-way correlation hold. Section 2 illustrated that two-way correlation can prevent AKM from recovering even the firm component of demand shocks or the bank component of supply shocks. Rather, the shocks must either be homogeneous or else demand consist of firm and purely idiosyncratic components and supply of bank and purely idiosyncratic components. In either case, the covariance of demand shocks across firms is zero, and the covariance of supply shocks across banks is zero. These two restrictions are over-identifying in our system (6), so under Assumptions 1-4 and the null hypothesis that these conditions hold,

$$\min_{\theta:\Lambda_{FF}=\text{diag}(0,1),\Lambda_{BB}=\text{diag}(1,0)} \tilde{q}(\boldsymbol{\eta},\theta)'\hat{\mathbf{W}}^{-1}\tilde{q}(\boldsymbol{\eta},\theta) \xrightarrow{d} \chi_2^2,$$

where $\tilde{q}(\boldsymbol{\eta},\theta) = \left(\begin{bmatrix}\sqrt{B} & 0 \\ 0 & \sqrt{F}\end{bmatrix} \otimes I_3\right) q(\boldsymbol{\eta},\theta)$.[5] This test provides a tool to evaluate whether empirical correlation patterns violate the structure underlying the AKM decomposition under our maintained assumptions.

3.5 Extensions

Including covariates. In models of unobserved heterogeneity, it is common to include covariates, as in (3), either to purge predictable variation so u_{fb} can be interpreted as shocks, or as regressors of interest in their own right. We can extend our baseline framework to incorporate a vector of covariates, X_{fb},

$$\eta_{fb} = Au_{fb} + \Gamma X_{fb}, \tag{7}$$

where Γ represents a matrix of coefficients. In our empirical application, X_{fb} includes one-period lagged relationship-level characteristics. In Appendix C.2, we show that, under the preceding assumptions, if the structural shocks are additionally mean independent of X_{fb}, Γ can be consistently estimated via OLS .

Including idiosyncratic errors. The baseline specification in (4) writes the observable vector η_{fb} as an exact linear combination of the latent demand and supply shocks, a standard convention in macroeconometric settings. However, realizations may feature unmodeled shocks unrelated to demand or supply, measurement error, or errors reflecting non-linearities

[5] Proposition 1 remains satisfied under these restrictions; for instance, the elements of Σ_{FF} simplify, $\Sigma_{FF,22} = A_{22}^2$ and $\Sigma_{FF,21} = A_{22}A_{12}$, ensuring that A remains identified from the remaining variance blocks.

not captured in a linear framework. We can accommodate such factors by introducing an additive error term:

$$\eta_{fb} = Au_{fb} + \varepsilon_{fb}.$$

If ε_{fb} is strictly idiosyncratic – uncorrelated across distinct relationships – its presence has no bearing on identification. Because these match-specific errors are cross-sectionally independent, they vanish from the cross-firm and cross-bank covariance matrices, leaving Σ_{FF} and Σ_{BB} unchanged. The asymptotic properties of our minimum distance estimator are likewise unaltered, as these errors contribute only higher-order terms. The only drawback is that relationship-level idiosyncratic errors will contaminate relationship-level shock estimates: $\hat{u}_{fb} = \hat{A}^{-1}(Au_{fb} + \varepsilon_{fb})$. However, if the ultimate objective involves aggregating such estimates to the agent level (i.e., constructing firm-level demand controls, as in Sections 4.3 and 4.4), this measurement error vanishes asymptotically as an agent's degree grows.

Exploiting multiple time periods. While our identification strategy and asymptotic theory require only a single cross-section of data, empirical settings often feature sparse network structures or lower-frequency adjustment, so cross-sectional variation may be limited. Researchers can leverage a multi-period panel to expand the effective sample size and improve precision. Pooling observations across multiple periods leaves our identification argument unchanged, subject to the obvious requirement that the response matrix A remain invariant over the pooled estimation window. Computing S_{FF} and S_{BB} as averages over both cross-sectional relationships and time dimensions yields equivalent asymptotic distributions. If the shocks are not serially correlated, the diagonal blocks of the asymptotic variance should now cluster at bank-time and firm-time levels; if the shocks are persistent, the variance should continue to cluster at bank and firm level.

3.6 Contrasting AKM and covariance restrictions

While our decomposition relaxes key assumptions relative to the AKM approach, it introduces important restrictions of its own. It is instructive to summarize the trade-offs between these two frameworks. The AKM decomposition suffers from two documented limitations: it does not accommodate or recover relationship-level heterogeneity, and it fails to recover even the common components of shocks when the underlying shocks exhibit two-way correlation across both firms and banks. In contrast, our approach directly models relationship-level shocks. The first key trade-off is that our framework requires the response matrix A to be invariant across relationships within a given cross-section. This is a restriction on the relative transmission of the shocks across the two outcomes, rather than on their absolute magnitudes.[6] Because classical AKM applications consider a single outcome equation, there are no analogous cross-outcome restrictions. A drawback of this flexibility is that decomposing two separate outcomes following AKM is not guaranteed to provide internally consistent results. While the of an invariant A under our approach is significant, its empirical relevance can be mitigated and evaluated by estimating the model on plausibly homogeneous subsamples, as we demonstrate in Appendix E.3.

The invariance of A also imposes a linear relationship between observables and shocks. While several extensions of the AKM framework incorporate non-linearities, these adaptations are often motivated by a desire to capture match-level heterogeneity through non-linear agent effects. Our framework achieves the same objective directly by modeling shocks at the relationship level. As noted in Section 3.5, the introduction of idiosyncratic match-level errors does not distort identification, ensuring that this exact linear structure is not a binding restriction. Furthermore, our linear structure offers asymptotic advantages; Theorem 1 establishes consistency of our parameters with only a fraction of agents well-connected. This contrasts with the AKM environment, where Jochmans and Weidner (2019) demonstrate that for all effects to be consistently estimated, the degree of all agents must diverge.

Second, our framework introduces orthogonality restrictions on opposite-type shocks across different relationships (and within a relationship to facilitate asymptotic results), which is definitional to our decomposition. This condition is motivated by standard macroeconometric definitions of structural shocks. While it rules out scenarios where demand and supply innovations *systematically* co-move, it can accommodate occasional patterns of co-movement, provided these correlations cancel out in expectation in population. The classical AKM decomposition, conversely, leverages two-way fixed effects, which impose no restrictions on cross-shock covariances within a relationship, thereby accommodating arbitrary common drivers of demand and supply. However, for the fixed effects to correspond to underlying shocks, an alternative strong assumption on covariances must hold: demand shocks can only be correlated across banks (holding firm fixed) and supply shocks can only be correlated across firms (holding bank fixed).

In sum, our covariance restriction framework and the AKM model offer alternative decompositions. Each recovers the true shocks under DGPs featuring different covariance patterns, and fails to recover the shocks when its respective core assumptions are violated. The choice between them ultimately rests on the specific economic environment and the researcher's objects of interest. To aid this decision, our methodology provides an over-identification test to reject the one-way correlation structure required by AKM, which is valid under our iden-

[6] Since A and the latent shocks are identified up to scale, the underlying variances of the shocks may vary arbitrarily across relationships.

tification assumptions. Furthermore, as we demonstrate in Section 4.3, estimating an AKM specification across multiple outcome variables can reveal whether the recovered fixed effects appear to isolate the intended structural objects. AKM simply cannot speak to heterogeneity, which is increasingly a topic of first-order interest for researchers, while our methodology allows the researcher to recover effect or shock estimates at the relationship level, which may offer rich economic possibilities.

3.7 Monte Carlo Study

Appendix D details a Monte Carlo study assessing our estimator's finite-sample performance. The design is empirically calibrated using our estimates for the Italian data from 2022 Q3 to 2023 Q4. In the simulations, we vary the dimensions of the market by setting $F = 1000B$ while holding the network density fixed at its empirical value.[7]

When the number of banks is small ($B = 10$), a setting significantly sparser than the 100–200 lenders observed in the Italian data and lower than any country in our sample, the percentage bias for the coefficients in A can be non-trivial (see Table D.1). However, this bias decays considerably for $B = 25$ and by $B = 100$, which remains below the empirical bank count, falls to 1-2%. For $B = 250$, the parameter estimates are virtually unbiased. Further, pooling four periods of cross-sectional data together mitigates small-sample distortions remarkably well, capping the maximum bias at 7% for $B = 10$ and at 5% for $B = 25$. This highlights the value of pooling quarterly observations in empirical setups with smaller cross-sections.

Next, we evaluate the empirical size of t-tests at a 5% nominal size for each coefficient in A (Table D.2). While notable size distortions emerge when $B = 10$ or $B = 25$, the empirical size improves once $B = 100$, ranging between 4.8% and 8.3%. Pooling again yields significant improvements, controlling size for $B = 25$ between 5.5% and 7.8%, though distortions persist at $B = 10$.

Finally, we evaluate our estimator's ability to recover the common agent-level components of these heterogeneous shocks. The simulated DGP incorporates a firm component in demand and a bank component in supply, but exhibits the two-way correlations we highlight as problematic for two-way fixed effects. We aggregate our estimated relationship-level shock estimates to the agent level, compute their correlation with the true components, and compare the results against standard two-way fixed effects estimates (D.3). Our aggregated firm-level demand shocks are 7–10 times more strongly correlated with the true firm demand

[7]While the Italian market features a large absolute number of multi-bank firms, its network density ranks as the third-lowest among the nine countries in our full study, and there is minimal variation in the average degree of firms across countries.

components than the fixed effects estimates across all specifications. Similarly, our aggregated bank shocks exhibit higher correlations with the true supply component for $B > 25$ in a single period, and for all $B \geq 25$ when pooling data. These findings highlight the challenge facing two-way fixed effects when two-way correlation exists and also illustrate our estimator's efficiency gains from leveraging data over the entire network rather than relying solely on an agent's immediate neighbors. As a result, our approach frequently outperforms conventional fixed effects, particularly for low-degree agents like most firms.

4 Demand and Supply in European Credit Markets

We apply our methodology to estimate credit demand and supply dynamics using the Eurosystem's AnaCredit dataset, a harmonized loan-level registry containing the near-universe of corporate credit exposures exceeding €25,000 in the euro area (see Kosekova et al. 2025, for an extensive overview). Technical documentation regarding the underlying data and collection is detailed in the AnaCredit Manual (European Central Bank 2019). Although the original records are collected at a monthly frequency and instrument level, we collapse the data to a firm-bank-quarter frequency. This quarterly aggregation is chosen because contract-level adjustments to loan quantities and interest rates occur relatively infrequently at a monthly horizon. In aggregating individual instruments, we focus exclusively on the three dominant categories of commercial credit used to finance corporate working capital and investments: term loans, credit lines, and revolving credit.

For each active firm-bank relationship in a given quarter, we compute the total committed credit. We define the change in loan quantity as the quarter-on-quarter change in this total committed amount rather than utilized credit. This choice reflects the fact that committed credit lines represent the outcome of both demand and supply forces, whereas variation in utilization is primarily driven by the unilateral demand choices of the firm given pre-existing commitments. We convert these quarterly variations into a symmetric growth rate:

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

The corresponding price measure, $\Delta r_{fb,t}$, tracks the quarter-on-quarter change in the volume-weighted average interest rate across all active instruments within the relationship, incorporating the prevailing values of reference rates for floating-rate contracts. The relationship-level observable vector is then

$$\eta_{fb,t} = \begin{pmatrix} \Delta r_{fb,t} \\ \Delta l_{fb,t} \end{pmatrix}.$$[8]

Our methodology is restricted to the intensive margin and we limit our estimation sample to multi-bank firms. Consequently, a firm is included in quarter t only if it maintains active borrowing relationships with at least two distinct banks in both t and $t-1$. To insulate our estimators from reporting anomalies and extreme tail observations, we winsorize both $\Delta l_{fb,t}$ and $\Delta r_{fb,t}$ at the 5% level. As a final preprocessing step, we project the observable vector onto quarter fixed effects and a vector of lagged relationship- and bank-sector-specific controls to ensure that $\eta_{fb,t}$ captures unpredictable, non-aggregate variation.[9]

Our empirical analysis covers 9 euro-area countries over the sample period from July 2019 to December 2023: Austria (AT), Belgium (BE), Germany (DE), Spain (ES), France (FR), Greece (GR), Italy (IT), the Netherlands (NL), and Portugal (PT). We partition this sample into three distinct 6-quarter estimation windows. This pooling framework balances several econometric considerations. First, while quarterly estimation across nine separate nations would yield an unmanageable set of coefficients for the purpose of this illustration, pooling the entire 2019–2023 period into a single sample would make the assumption of an invariant transmission matrix A implausible due to major structural shifts in credit markets. The 6-quarter window strikes a balance. Second, these three sub-periods isolate distinct macroeconomic and policy regimes. The first window (2019Q3–2020Q4) spans the pandemic onset and its immediate aftermath, characterized by substantial real-economy disruptions and large-scale public credit support interventions. The second window (2021Q1–2022Q2) covers the onset of global inflation prior to active monetary tightening. The final window (2022Q3–2023Q4) captures an aggressive monetary contraction and credit tightening. It is intuitive that the demand and supply elasticities embedded in A vary across these regimes. Third, while some core euro-area countries feature exceptionally large networks, smaller or highly consolidated banking sectors benefit significantly from the increased effective sample size resulting from pooling six quarters of data, as illustrated in the Monte Carlo Study (Appendix D).

Table E.4 in the Appendix presents descriptive statistics for the number of unique firms,

[8] As we note in the Conclusion, future work could consider augmenting the observation equation with loan maturity, for example, to generate over-identifying restrictions.

[9] We residualize using four variables at the firm-bank-time level and two at the bank-sector-time level to net out systematic variation related to pricing rigidity, collateralization, loan composition, relationship intensity, and the bank's sectoral positioning. Specifically, these are the lagged shares of a firm's borrowing from a bank that are fixed-rate, secured by collateral, or structured as credit lines/term loans, alongside the lender's share in the firm's total credit, the bank's market share in the firm's sector, and the bank's overall exposure to that sector.

banks, and active firm-bank relationships across each country and sub-period within our final estimation sample. Mirroring the stylized facts documented by Kosekova et al. (2025), these metrics highlight substantial variation in credit market structure across Europe. Multi-bank borrowing is pervasive in economies like Italy and Portugal, yet remains relatively rare in the Netherlands, even when comparing countries of comparable economic scale. Banking sector concentration similarly diverges; the number of active lenders ranges from roughly 20 institutions in Belgium, Greece, and the Netherlands to over 800 in Germany. As a result, our final estimation samples vary substantially in scale, ranging from approximately 8,000 unique relationships per window in the Netherlands to more than 700,000 in Italy, and does not map one-to-one with country size.

4.1 Credit demand and supply patterns over countries and time

We begin by reporting the point estimates of A, which measure the transmission of demand and supply shocks to interest rates and loan volumes. Figure 1 plots these responses across the three macroeconomic regimes for each country along with their 95% confidence intervals. To ensure comparability across separate banking systems and time periods, the coefficients are scaled to represent responses to a unit standard deviation shock, with price changes in interest rate basis points and quantity adjustments in symmetric growth rates.

First, these coefficient estimates provide a validation of our decomposition. Across all 27 country-period blocks, our framework uncovers one shock that moves price and quantity in the same direction – consistent with a demand shock – and a second shock that drives price and quantity in opposite directions, matching a supply shock, up to statistical significance.[10] Because our methodology leaves the raw estimate of A entirely unrestricted (using the sign pattern only to label columns), this pattern is an empirical result, and not by construction.

Second, although the coefficients evolve over time within countries, cross-country variation generally dominates these temporal shifts. There are some notable exceptions to this pattern. The price response to demand shocks ($\hat{A}_{11}$) increases significantly during the monetary contraction in Austria, Belgium, Germany, Greece, Italy, and Portugal. Concurrently, the sensitivity of pricing to supply shocks ($\hat{A}_{12}$) falls across Germany, France, Greece, Italy, and Portugal, implying that these credit markets became increasingly demand-driven. Quantity responsiveness to demand shocks ($\hat{A}_{21}$) declines over time in Belgium, France, Greece, and Italy, while quantity sensitivity to supply innovations ($\hat{A}_{22}$) drops in Spain, France, Italy, and Portugal.

To provide a more intuitive economic interpretation of these parameters, the coefficients

[10] While there are a small number of point estimates that may display the wrong sign pattern, the offending estimates are never statistically distinct from zero.

Figure 1: Responses of Price and Quantity to Demand and Supply

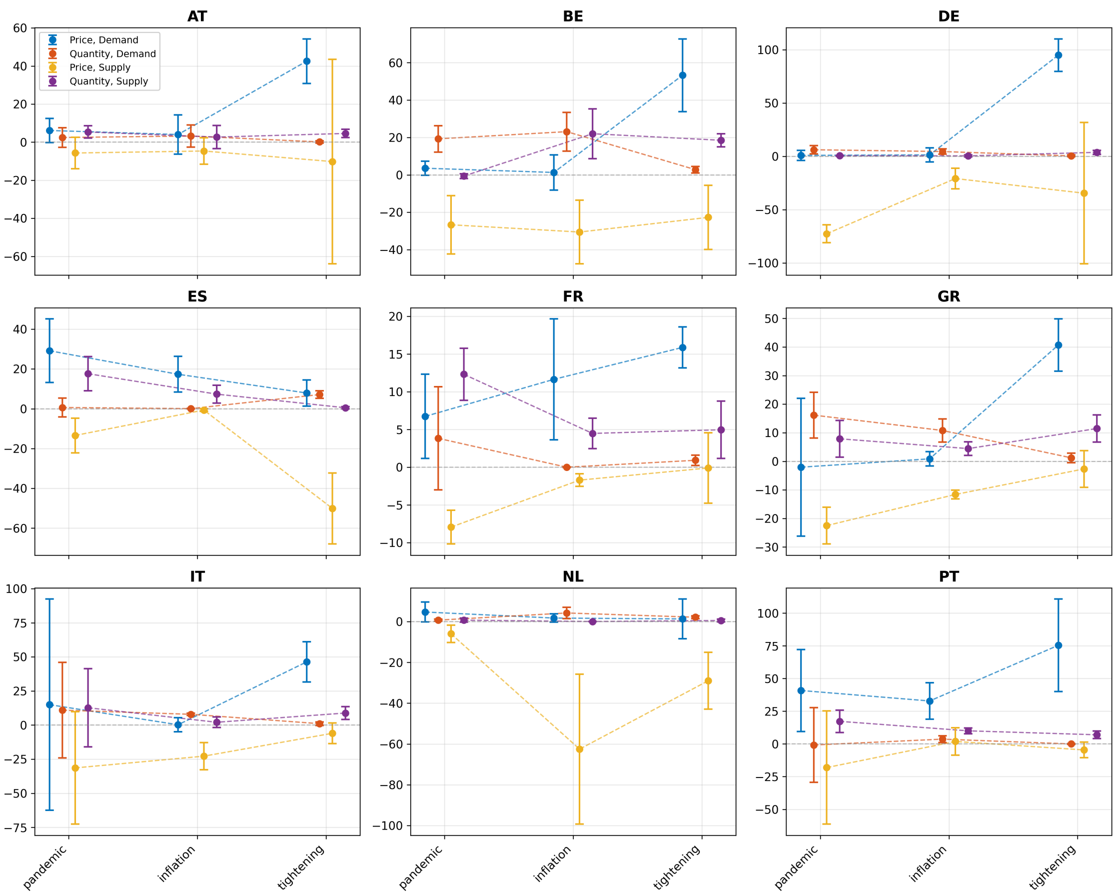

Notes: This figure displays estimates of the entries of $\hat{A}$ across nine European countries. For a given country, a subplot reports sequentially the response of price to demand shocks ($\hat{A}_{11}$); the response of quantity to demand shocks ($\hat{A}_{21}$); the response of price to supply shocks ($\hat{A}_{12}$); and the response of quantity to supply shocks ($\hat{A}_{22}$). The four coefficients are reported for the three estimation periods: the pandemic (2019Q3–2020Q4), the inflationary period (2021Q1–2022Q2), and the monetary tightening episode (2022Q3–2023Q4). Point estimates are scaled by the standard deviation of the corresponding structural shock. Error bars represent 95% confidence intervals constructed using autocorrelation robust standard errors.

in A can be translated into the implied slopes of credit demand and supply curves. Figure 2 summarizes how these curves vary by country and time by plotting the angle each curve forms with the horizontal axis, which avoids the scaling issues associated with reporting near-vertical slopes.[11] An angle of 0° indicates a perfectly elastic curve, 45° marks unit elasticity, and 90° signifies a perfectly inelastic relationship.

During the pandemic (blue diamonds), Austria, France, Greece, and Italy display standard downward-sloping demand and upward-sloping supply schedules. However, half of the

[11]Figure E.1 in the Appendix plots conventionally demand and supply curves for each country-period.

economies feature at least one near-vertical, inelastic curve. Credit demand is highly inelastic in Belgium, Germany, and the Netherlands, which aligns with firms aggressively seeking emergency liquidity to survive pandemic-related operational halts regardless of borrowing costs. Conversely, the credit supply curve is nearly vertical in Spain, Portugal, and the Netherlands.

During the subsequent inflationary episode (orange circles), slopes remain stable in Austria, Germany, and Greece, but shift substantially elsewhere. Inelastic regimes remain widespread. Credit demand becomes highly inelastic in Germany, the Netherlands, and Italy, indicating corporate borrowers prioritized raising funds to cover inflation-driven operating costs. In contrast, credit supply turns highly inelastic in Spain, France, and Portugal. This pattern is consistent with lenders restricting credit volumes to avoid locking in long-term loan yields that face real depreciation from persistent inflation. Note that while Portugal displays a marginally upward-sloping demand curve, this specific estimate is statistically insignificant and highly noisy.

During the monetary policy tightening regime (green squares), all countries exhibit at least one inelastic curve, predominantly an inelastic supply curve. This widespread steepening reflects constrained credit in response to tighter conditions fostered by monetary policy, with interest rates serving to ration credit. Spain and the Netherlands emerge as the only exceptions to this supply inelasticity with highly inelastic credit demand curves instead (Germany likewise has nearly inelastic demand). Elsewhere, demand ranges from highly elastic schedules (France and Greece) to a standard downward slope.

Overall, these estimates underscore that while monetary policy actions were uniform across the euro area, national credit markets behave with considerable heterogeneity. These divergent dynamics also offer valuable insights for designing or evaluating localized macroeconomic and macroprudential policies. Investigating the institutional drivers behind these cross-country differences in the responsiveness of price and quantity remains an open area for empirical research. An important caveat is that these results are representative only of multi-bank firms, which tend to be larger, safer, and better positioned to substitute across lenders than single-bank borrowers. Extending this linear framework to incorporate single-bank relationships represents another promising direction for future work.

4.2 Assessing the plausibility of an AKM decomposition

Before moving to regression analyses following leading empirical designs, we evaluate whether standard two-way fixed effects specifications are compatible with the data. We first execute the test derived in Section 3.4, which leverages the over-identifying restrictions of our decom-

Figure 2: Demand and Supply Elasticity Estimates Across Countries and Time Periods

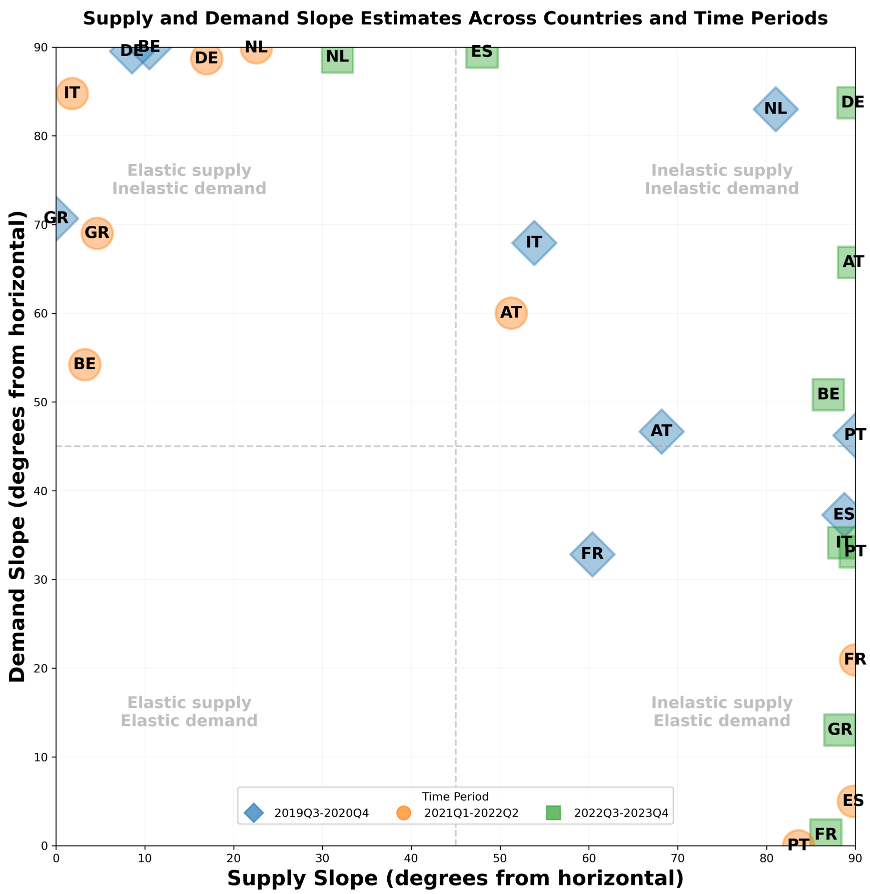

Notes: This figure presents demand and supply slope estimates across nine European countries over three time periods: 2019Q3–2020Q4 (blue diamonds), 2021Q1–2022Q2 (orange circles), and 2022Q3–2023Q4 (green squares). Slopes are derived from the supply-demand decomposition coefficients in A, where supply slope $= A_{11}/A_{21}$ and demand slope $= A_{12}/A_{22}$. Rather than raw slopes (which may diverge), we plot the angle of the curves relative to the $x-$axis. An angle of $0°$ represents a perfectly flat (elastic) curve, $45°$ corresponds to unit elasticity, and $90°$ represents a vertical (perfectly inelastic) curve. The reference lines at $45°$ delineate four quadrants representing combinations of elastic and inelastic demand and supply.

position to test for the one-way correlation patterns required by the AKM decomposition to separate demand and supply. Across the 27 distinct country-period blocks, the data reject the null hypothesis of one-way correlations at the 5% level in 25 cases.[12] While this over-identification test relies on the validity of our underlying covariance restrictions, these widespread rejections provide strong preliminary evidence that two-way fixed effects may struggle to separate clean demand and supply innovations in this setting.

Table 1 documents the extent of match-level heterogeneity that would be masked by a

[12]The two exceptions are the Netherlands in period 3 and Portugal in period 1. For many country-periods, the null hypothesis can be rejected with far smaller p-values.

Table 1: Between and within variation

	Collapse at the firm-time level						
	p10	p25	p50	p75	p90	IQR	StD
Average demand innovation	-0.677	-0.253	0.000	0.171	0.677	0.424	0.646
Std dev demand innovation	0.019	0.063	0.225	0.863	1.681	0.801	0.780

	Collapse at the bank-time level						
	p10	p25	p50	p75	p90	IQR	StD
Average supply innovation	-0.218	-0.088	0.009	0.095	0.231	0.184	0.399
Std dev supply innovation	0.267	0.485	0.712	0.952	1.266	0.467	0.511

Notes: Within and between variation in the (standardized) relationship level demand and supply innovations can be assessed at either the firm or bank level. In the upper panel, we start from the firm-bank-time level demand innovations and compute the average firm-time demand innovation and the within-firm standard deviation in the demand innovation. In the lower panel, we construct similar measures at the bank level focusing on the supply innovations. For each of these collapsed firm-time and bank-time measures, we provide summary statistics to shed light on the between-firm (bank) heterogeneity versus the within-firm (bank) heterogeneity in demand (supply) innovations.

decomposition making the conventional homogeneity restriction. For each firm-period, we calculate the average demand shock across its lending relationships alongside the corresponding within-firm standard deviation. The upper panel tracks the distribution of these metrics across firm-quarters, while the lower panel presents a symmetric breakdown for supply shocks at the bank level. The estimates reveal that within-unit variation across relationships is large, frequently matching or exceeding between-unit variation. Specifically, the median within-firm standard deviation of demand shocks is 0.225, approximately 35% of the cross-sectional between-firm standard deviation (0.646). At the 75th percentile of the distribution, the within-firm standard deviation exceeds the between-firm variation. This pattern is even more pronounced for credit supply: the median within-bank standard deviation (0.712) is 78% greater than the cross-sectional variation observed between banks (0.399).

These patterns call into question the premise that demand and supply shocks can be meaningfully approximated using a homogeneous decomposition like AKM. Our findings demonstrate instead that, for most firms and banks, the heterogeneity across their relationships exceeds the average differences across firms or banks. Consequently, strategies relying on standard two-way fixed effects may discard a primary component of economically meaningful variation. While one might object that this heterogeneity could be just noise, due to imposing that all variation is either demand or supply (rather than idiosyncratic errors as in AKM), the results in Section 4.5 demonstrate that this within-agent variation is not noise, but instead economically meaningful.

4.3 Assessing validity using price responses

To assess the empirical plausibility of previous identification strategies, we evaluate the recovered demand and supply measures against a basic tenet of economic theory with unambiguous predictions: demand shocks that increase quantities should increase prices, whereas supply shocks that increase quantities should decrease prices. In doing so, we show that AKM-style firm fixed effects appear to be confounded by supply shocks.

Table 2 reports the results of regressing relationship-level interest rate changes – residualized against our baseline controls – on demand and supply measures, pooled over countries and time periods, after controlling for country-time fixed effects. Standard errors are two-way clustered at the firm and bank levels. Column (1) displays the results using our relationship-level demand and supply innovations identified via covariance restrictions. The estimated coefficients match economic theory: a one-standard-deviation increase in the demand innovation raises lending rates by 21.9 basis points, while a symmetric supply innovation reduces them by 18.7 basis points. Both effects are highly significant and economically meaningful, explaining 48% of the total variation in match-level interest rate adjustments.[13]

Column (2) replaces our innovations with firm-time and bank-time fixed effects estimates from the quantity equation, alongside the corresponding residuals.[14] The AKM firm-time effect, typically deployed as a proxy for credit demand shifts, has a negative coefficient of -0.483, suggesting that this measure is actually dominated by *supply* factors. The bank-time effect has the expected negative sign, but the model's explanatory power collapses to an adjusted R^2 of 0.01. Additionally, the AKM residual exhibits a highly significant negative correlation with interest rate changes, demonstrating that residual idiosyncratic variation in the quantity equation remains systematically relevant for prices. To isolate the potential source of the sign reversal of the demand measure, Columns (3) and (4) implement sequential substitution tests. In Column (3), controlling for our identified supply innovation flips the coefficient on the AKM firm-time component to a positive 1.151. This reversal suggests that the firm-time fixed effect does contain demand variation, but that it is confounded by supply factors in Column (2) for which our supply innovation effectively controls. Symmetrically, Column (4) shows that controlling for our demand innovation steepens the loading of the AKM bank-time fixed effect from -0.751 to -1.260, which is consistent with the removal of attenuation due to the firm-time effect being contaminated by supply. Taken together, these patterns are consistent with two-way fixed effects estimates being mixtures of both structural

[13]The remaining variation reflects variation in the responses A across country-period blocks; within any single country-period estimation field, these innovations explain 100% of the price variance by construction.

[14]Following Khwaja and Mian (2008) and Amiti and Weinstein (2018), these components are extracted from the quarterly specification $\Delta l_{fb,t} = d_{f,t} + s_{b,t} + \Gamma X_{fb,t} + \varepsilon_{fb,t}$, estimated by country and six-quarter window after partialing out relationship controls. $X_{fb,t}$ is as in our baseline.

Table 2: Innovations versus fixed effects: Impact on Interest rate changes

	(1)	(2)	(3)	(4)
	Change in the relationship-level interest rate			
Demand innovation	0.219***			0.261***
	(0.008)			(0.012)
Supply innovation	-0.187***		-0.259***	
	(0.007)		(0.009)	
AKM Firm-Time FE		-0.483***	1.151***	
		(0.054)	(0.084)	
AKM Bank-Time FE		-0.751***		-1.260***
		(0.096)		(0.104)
AKM Residual		-0.470***	1.150***	-1.549***
		(0.054)	(0.082)	(0.111)
Adjusted R^2	0.48	0.01	0.25	0.36

Notes: The table reports coefficients from regressions of relationship-level interest rate changes on demand and supply measures. The dependent variable is the quarterly change in the relationship-level lending rate (expressed in percentage points) between bank b and firm f in period t, $\Delta P_{fb,t}$, after residualizing with respect to relationship-level controls, $X_{fb,t}$. In column (1), Demand innovation and Supply innovation are our relationship-level shocks (f, b, t) identified using covariance restrictions (standardized to have mean zero and unit standard deviation) . In column (2), the independent variables are AKM firm-time fixed effect (f, t), AKM bank-time fixed effect (b, t), and the AKM residual (f, b, t), obtained from the regression $\Delta l_{fb,t} = d_{f,t} + s_{b,t} + \Gamma X_{fb,t} + \varepsilon_{fb,t}$, where $\Delta l_{fb,t}$ is the quarterly change in credit from bank b to firm f in period t. Column (3) tests if firm-time FEs recover demand conditional on substituting our identified supply measure. Column (4) tests if bank-time FEs recover supply conditional on substituting our identified demand measure. Standard errors are two-way clustered by bank and firm. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively. All regressions are based on 30,569,791 observations and include country-time fixed effects.

forces rather than identified measures of demand and supply.

These insights are robust to aggregating the data to the firm-time and bank-time levels. Our results reinforce the rejections of the one-way correlation restrictions documented in the previous section and match the contamination illustrated analytically in Section 2. They also accord with contemporary evidence by Bergman et al. (2025), who argue from a structural model that firm fixed effects reflect loan supply shocks in addition to demand shocks and document considerable bias in Khwaja and Mian (2008)-style estimates of the effect of banks' real estate exposure in Spain.

4.4 The impact of decompositions on firm outcome regressions

While the price regressions offer insights into the plausibility of demand and supply controls, they do not mirror typical empirical specifications, where firm-level real outcomes are the primary objects of interest (e.g., Khwaja and Mian 2008; Chodorow-Reich 2014; Jiménez et al. 2020). In this section, we demonstrate how the contamination documented above distorts parameters in real outcome regressions.

Following the literature, we first assess the impact of an exogenous supply shock on credit growth via a quasi-natural experiment inspired by Core et al. (2025). We exploit the large, sudden monetary policy tightening initiated by the ECB in July 2022, which raised policy rates by 450 basis points over the subsequent fourteen months. Because the timing and magnitude of this tightening cycle were largely unanticipated, firms' pre-determined exposure to interest rate fluctuations, determined by their *ex ante* choice of floating versus fixed-rate contracts and debt maturity structure, is exogenous relative to the shock. Our objective is to identify the differential change in credit growth between exposed (floating-rate) and unexposed (fixed-rate) firms while controlling for confounding credit demand. For each horizon h during the 2022Q3–2023Q4 window, we estimate the following relationship-level cross-sectional regression:

$$\ln(Q_{fb,2022Q2+h})-\ln(Q_{fb,2022Q2}) = \iota_b^h+z_{c,s}^h+\pi^h\mathbf{1}_{f,pre}[exposed]+\sum_{j=1}^{h}\vartheta_j^h\cdot demand_{fb,2022Q2+j}+\epsilon_{fb}^h, \tag{8}$$

where the coefficient of interest, π^h, measures the supply effect on loan quantities. The term ι_b^h represents bank fixed effects, which absorb bank-level differences in deposit franchises (Drechsler et al. 2017), interest rate risk exposure (Gomez et al. 2021), or valuation losses on securities (Greenwald et al. 2024). The vector $z_{c,s}^h$ denotes country-sector fixed effects that capture asymmetric real-economy shocks. Estimating (8) over an expanding horizon up to six quarters traces the cumulative impact of the monetary tightening episode.

Figure 3 plots the responses. A baseline specification, omitting demand controls altogether (black line), shows that exposed firms experienced significantly lower credit growth from 2022Q2 onward. When controlling for our relationship-level demand innovations (u_{fb}^d, blue line), the estimated supply contraction from 2022Q4 to 2023Q2 is more than twice as large.[15] This divergence reveals the presence of a countervailing positive demand channel with floating-rate firms demanding more credit. These findings are consistent with banks viewing firms with lower exposure to risks from rising interest rates as more attractive bor-

[15] While our demand innovations are relationship-specific, aggregating them to a firm-level average yields nearly identical results.

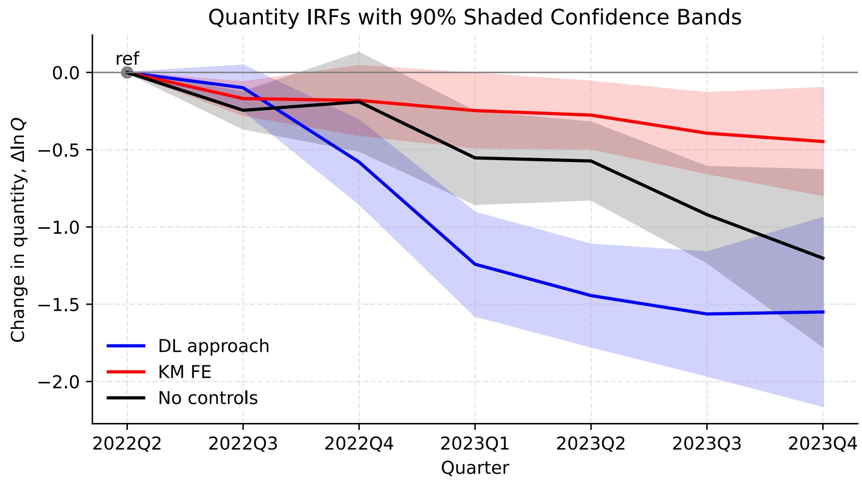

Figure 3: Dynamic Effect of Tightening on Quantities

Notes: This figure plots dynamic coefficients over the monetary policy tightening episode capturing the differential credit growth of floating-rate versus fixed-rate borrowing firms. The aim is to uncover the heterogeneous credit supply impact. All specifications include bank fixed effects and country-sector fixed effects to absorb heterogeneous bank and country-sector responses over the tightening episode. The regressions are at the bank-firm relationship level. The solid lines show the point estimates for three specifications: specification without controls (black), specification with credit demand captured by firm-time fixed effects in a first-stage Khwaja and Mian (2008) decomposition (red), and specification with our credit demand innovations (blue). Shaded areas represent 90% confidence bands based on two-way clustered standard errors at the bank and firm level. The reference quarter is 2022Q2.

rowers, while firms with greater exposure seek to secure credit before rates rise higher. In contrast, controlling for demand with firm-time fixed effects as in Khwaja and Mian (2008) (red line) substantially attenuates the response, rendering the supply effect barely significant at most horizons. This is consistent with our price-response analysis: the firm-time fixed effects appear to be dominated by supply factors, so controlling for it washes out the supply channel this strategy seeks to identify, more so than omitting demand controls entirely.

We next document how this credit supply contraction transmitted to firm-level real outcomes, as in the framework of Khwaja and Mian (2008), Chodorow-Reich (2014), Cingano et al. (2016), and Jiménez et al. (2020). We estimate annual regressions of the form

$$Y_{f,2021Q4+h} - Y_{f,2021Q4} = \tilde{\iota}^h_{b^*(f)} + \tilde{z}^h_{c,s} + \tilde{\pi}^h \mathbf{1}_{pre}[exposed] + \sum_{j=1}^{h} \tilde{\vartheta}^h_j \cdot demand_{f,2021Q4+j} + \tilde{\epsilon}^h_f, \quad h \in \{4, 8\} \tag{9}$$

where $Y_{f,t}$ represents a firm-level real outcome evaluated at the one-year ($h = 4$) or two-year ($h = 8$) horizon relative to the pre-tightening baseline (2021Q4). The term $b^*(f)$ isolates the firm's primary lender, allowing $\tilde{\iota}^h_{b^*(f)}$ to absorb main-bank characteristics, while $\tilde{z}^h_{c,s}$ controls for country-sector trends. demand$_{f,t}$ is a quarterly demand control.

Table 3 reports results for three key dependent variables: (log) total assets (panel 1), (log) turnover (panel 2), and the change in the credit to total assets ratio (panel 3). The left side reports results for 2022 and the right for 2023, with the first column on each side featuring no demand control, the second the firm fixed effect, following Khwaja and Mian (2008), and the third our demand innovations aggregated at firm level. Our object is the coefficient on firm exposure, $\tilde{\pi}^h$. For total assets, the responses in 2022 are negative and significant, reflecting a largely supply-driven effect. In 2023, however, the estimate controlling for our demand measure is nearly 1.5 times as large as that controlling for the firm fixed effect, and for both periods controlling for the fixed effect attenuates responses relative to including no control at all, again consistent with supply factors contaminating the firm fixed effect. For turnover, the ranking of point estimates is unchanged, with our demand measure delivering the largest responses and the fixed effect specification the smallest. For the credit to total assets ratio, in 2022 the point estimates are more than twice as large and significant (at the 10% level) with no control or our demand measure compared to when the firm fixed effect is included. For 2023, all coefficients are significant, but the coefficient using our measure is again more than 1.5 times that with the firm fixed effect.

In sum, we find that the credit supply contraction induced by the 2022 monetary tightening significantly curtailed real economic activity among interest rate-exposed firms, driving down total assets, turnover, and leverage ratios. Accounting for match-specific demand heterogeneity helps to uncover this channel, while relying on conventional firm fixed effects as a demand proxy misattributes supply retrenchment to the borrower, reflecting the supply contamination we documented in the previous section.

Taking stock, we have highlighted how using two-way fixed effects as demand and supply measures, following an AKM decomposition, can deliver results at odds with economic theory, since the one-way correlation structure it imposes on demand and supply is unlikely to hold – a fact we detected in 25 of 27 country-periods. The consequences of mixing of demand and supply in the fixed effects extend beyond price effects to real outcome regressions, where the resulting demand controls substantially attenuate the responses of both credit volume and several real outcomes amongst firms exposed to the post-2022 credit supply contraction, even relative to omitting demand controls altogether. In this setting, demand measures based on our novel decomposition appear more economically credible.

4.5 Heterogeneous Transmission of Monetary Policy

Finally, we exploit the *heterogeneity* in our demand and supply innovations to examine the heterogeneous transmission of monetary policy across relationships. Although quarterly

Table 3: Real Effects

	(1)	(2)	(3)	(4)	(5)	(6)
			ln(Total Assets)			
Floating or short-term borrower	-0.695***	-0.611***	-0.726***	-1.065***	-0.795***	-1.163***
	(0.182)	(0.171)	(0.168)	(0.245)	(0.225)	(0.220)
Observations	514493	509511	514493	415080	415080	415080
Adjusted R^2	0.02	0.03	0.02	0.03	0.06	0.04
Demand	-	KM	DL	-	KM	DL
Horizon	2022	2022	2022	2023	2023	2023
			ln(Turnover)			
Floating or short-term borrower	-0.749***	-0.664***	-0.758***	-1.272***	-1.087***	-1.359***
	(0.226)	(0.221)	(0.212)	(0.306)	(0.303)	(0.278)
Observations	486395	481690	486395	394249	394249	394249
Adjusted R^2	0.05	0.05	0.05	0.07	0.07	0.07
Demand	-	KM	DL	-	KM	DL
Horizon	2022	2022	2022	2023	2023	2023
			Credit-to-Total assets			
Floating or short-term borrower	-0.294*	-0.129	-0.322*	-1.258***	-0.723***	-1.289***
	(0.150)	(0.146)	(0.167)	(0.349)	(0.269)	(0.366)
Observations	514493	509511	514493	415080	415080	415080
Adjusted R^2	0.02	0.15	0.06	0.03	0.17	0.07
Demand	-	KM	DL	-	KM	DL
Horizon	2022	2022	2022	2023	2023	2023

Notes: For each panel, the dependent variable is the change in the indicated firm-level variable, from 2021 to the year reported as "Horizon". Specifically, the top panel considers (log) total assets, the middle (log) turnover, and the bottom the credit to total assets ratio. The regressions follow the form in (9) , where the demand measure included (if any) is indicated by "Demand". KM (Khwaja-Mian) refers to the AKM-type demand measure, while "DL" (De Jonghe-Lewis) refers to our relationship-level identified demand innovations, collapsed at the firm level. Quarterly KM and DL demand controls are included for each quarter of the estimation horizon (i.e. 4 (8) quarters in columns (1)-(3) ((4)-(6)). The specification also includes a main bank fixed effect and country-sector fixed effect. Standard errors are two-way clustered at main bank level and country-sector level.

demeaning mechanically purges aggregate level effects from our innovations, these measures can still capture how macroeconomic shocks transmit heterogeneously based on relationship-specific characteristics. We estimate specifications of the form:

$$u^i_{fb,t} = \varphi^i \cdot share_{fb,t-1} + \pi^i_{mp} \cdot mp_t \times share_{fb,t-1} + \pi^i_{cbi} \cdot cbi_t \times share_{fb,t-1} + \text{fixed effects} \quad \ldots \quad (10)$$

where $i \in \{\text{demand}, \text{supply}\}$. The interaction coefficients, π^i_{mp} and π^i_{cbi}, capture how demand and supply innovations covary with pre-determined, 1-quarter lagged relationship attributes ($share_{fb,t-1}$) following monetary policy (MP) and central bank information (CBI) shocks from Jarociński and Karadi's (2020) popular series.

Table 4 reports the results, where Columns (1)–(3) use our demand innovations as the dependent variable and Columns (4)–(6) employ supply innovations. We consider two key covariates for $share_{fb,t-1}$: the share of fixed-rate loans, to capture the floating-rate channel (Ippolito et al. 2018; Gürkaynak et al. 2022), and the share of collateralized loans, to capture the bank lending channel (Ivashina et al. 2022). Our estimates reveal contrasting demand and supply patterns across these contract structures. Following a contractionary monetary policy shock, a higher share of fixed-rate loans is associated with significantly lower credit demand. This pattern is consistent with borrowers who have secured credit at lower fixed rates avoiding new, costlier debt. Conversely, following a positive central bank information shock signaling expanding economic activity, these same firms who borrow at a fixed rate increase demand, anticipating higher interest rates in the near future. Turning to supply, following a contractionary MP shock, a larger fixed-rate share is associated with a significant increase in credit supply as banks perceive these counterparties as less exposed to interest rate risk. Following a positive CBI shock, however, supply to fixed-rate borrowers falls as banks retrench from fixed-rate exposures to prioritize floating-rate originations likely to be more profitable as rates rise. Meanwhile, loan collateralization does not meaningfully explain credit demand variations following either shock. On the supply side, however, a higher share of collateralized loans is associated with a significantly greater supply of credit following a contractionary MP shock, confirming banks actively prioritize asset-based lending over cash-flow-based alternatives during periods of contraction.

These *relationship-level* estimates augment three strands of the monetary transmission literature. On the demand side, our results complement the floating-rate channel documented by Ippolito et al. (2018), who show that firms with unhedged bank debt exhibit heightened sensitivity to monetary policy. We extend this by showing that this channel operates at the relationship level and manifests asymmetrically on each side of the market, with exposed firms reducing relative demand while banks concurrently reallocate relative supply toward insulated, fixed-rate borrowers. On the supply side, our collateral findings reinforce structural models of the collateral channel (Kiyotaki and Moore 1997; Chaney et al. 2012) by documenting that banks explicitly shift toward secured exposures as monetary conditions contract. Finally, while traditional frameworks emphasize firm-level characteristics to explain asymmetric transmission (Gertler and Gilchrist 1994; Ottonello and Winberry 2020), we document that substantial transmission heterogeneity exists within a single firm across its separate banking relationships – a granular dimension that current macroeconomic models omit.

From a policy perspective, these results highlight the distributional and procyclical consequences of monetary policy. The incidence of central bank tightening depends on the pre-existing credit contract structures. Firms reliant on floating-rate or unsecured loans face a double penalty: higher debt-servicing costs compounded by an endogenous contraction in relative loan supply as banks de-risk. This amplification mechanism –whereby the banking sector's risk management response reinforces rather than dampens the contractionary effect on vulnerable borrowers – raises concerns about procyclicality. Because the distribution of these contract structures is heavily shaped by macroprudential policy tools such as borrower-based limits or interest rate risk regulations, these findings reveal a crucial interaction where macroprudential choices enacted during benign conditions directly dictate the distributional costs of monetary tightening when it arrives.

Importantly, these relationship-level dynamics could not be uncovered using previous identification strategies, which generate demand and supply measures strictly at the aggregated firm or bank level. These results also address the potential concern that the demand and supply heterogeneity we recover is just "noise", treating variation appearing in the AKM decomposition's idiosyncratic errors as part of demand or supply. Instead, these findings show that within-firm and within-bank variation – precisely what appears in those idiosyncratic errors – is significantly associated with economically meaningful relationship characteristics following a monetary contraction.

5 Conclusion

This paper proposes a novel decomposition to recover relationship-level effects or shocks as an alternative to two-way fixed effects framework of Abowd et al. (1999). By exploiting the covariance structure of observed outcomes across an agent's relationships, we identify parameters and underlying innovations without distributional assumptions. The identification strategy leverages linearity and mutual orthogonality assumptions in exchange for the possibility of uncovering rich patterns of heterogeneity and simple estimation and inference.

Table 4: Heterogeneous transmission of monetary policy and central bank information shocks

	(1)	(2)	(3)	(4)	(5)	(6)
	Demand innovation			Supply innovation		
Share of fixed rate loans	0.014		0.014	-0.018***		-0.018***
	(0.012)		(0.012)	(0.007)		(0.007)
MP × Share of fixed rate loans	-0.537***		-0.538***	0.674***		0.692***
	(0.152)		(0.153)	(0.155)		(0.153)
CBI × Share of fixed rate loans	0.961***		0.960***	-0.385***		-0.392***
	(0.176)		(0.175)	(0.118)		(0.118)
Share of collateralized loans		0.015***	0.015***		-0.009	-0.009
		(0.006)	(0.005)		(0.008)	(0.009)
MP × Share of collateralized loans		-0.026	-0.064		0.395***	0.431***
		(0.071)	(0.075)		(0.067)	(0.068)
CBI × Share of collateralized loans		0.068	0.087		-0.012	-0.028
		(0.092)	(0.083)		(0.073)	(0.070)
Adjusted R^2	0.08	0.08	0.08	0.04	0.04	0.04

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 (MP) shocks and central bank information (CBI) shocks. The two firm-bank characteristics are the share of borrowing by firm f from bank b at time $t-1$ that is fixed rate and the share of borrowing by firm f from bank b at time $t-1$ that is collateralized. The monetary policy and central bank information shocks are obtained from Jarociński and Karadi (2020). All specifications include firm-quarter fixed effects as well as bank-industry-location-quarter fixed effects and include 26,210,226 observations across 9 euro-area countries over the period 2019Q3–2023Q4. Standard errors are two-way clustered by bank and firm.

Empirically, we document extensive insights into euro-area corporate credit markets. While the slopes of credit demand and supply curves are relatively stable within countries over time, they exhibit substantial cross-country heterogeneity. A specification test detects two-way correlations of shocks across agents, making AKM decompositions likely to conflate demand and supply. As a consequence, AKM-style "demand" effects exhibit a counter-theoretical negative impact on interest rates, appearing to be dominated by supply-side variation. Consequently, using such proxies to control for credit demand attenuates responses in regressions studying the credit supply contraction following the 2022 monetary tightening cycle. In contrast, our covariance-identified demand measures help to uncover a large, persistent contraction in credit supply to exposed firms over the subsequent six quarters. This supply reduction transmitted directly to real outcomes, lowering total assets, turnover, and leverage for exposed firms, effects that are systematically underestimated using a firm fixed effect as a demand control. Finally, the relationship-level heterogeneity in our shocks reflects differential transmission of monetary policy and central bank information via contract structures, informing modeling exercises and macroprudential policy debates.

On the theoretical front, fruitful avenues for future research include extensions to accommodate single-relationship agents or formalizing the extensive margin of relationship formation. Incorporating additional relationship-level observables, such as loan maturity, could provide valuable over-identifying restrictions to test underlying assumptions or place structure on temporal variation in the response parameters. While extending the framework to absorb aggregate time-series variation (e.g., Amiti and Weinstein 2018) is conceptually appealing, it remains a challenge, as aggregate macro-shocks inherently act as proximal causes driving both sides of the market simultaneously.

Finally, while our empirical focus is corporate credit markets, our decomposition is broadly applicable across several fields. It translates directly to other financial settings, international trade, and production networks. Furthermore, application to the classic AKM setting of matched employer-employee data could unveil match-specific wage-setting dynamics, with rich consumer data another exciting avenue.

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

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

Anderson, G., Bahaj, S., Chavaz, M., Foulis, A., & Pinter, G. (2022). Lending relationships and the collateral channel. *Review of Finance*, *27*(3), 851–887. https://doi.org/10.1093/rof/rfac041

Bergman, N., Casado, A., Iyer, R., & Saporta-Eksten, I. (2025). *Estimating the impact of loan supply shocks* (tech. rep.). CEPR Discussion Paper No. 20865. https://cepr.org/publications/dp20865

Bernard, A. B., Moxnes, A., & Saito, Y. U. (2019). Production networks, geography, and firm performance. *Journal of Political Economy*, *127*(2), 639–688. https://doi.org/10.1086/700764

Bonhomme, S., Holzheu, K., Lamadon, T., Manresa, E., Mogstad, M., & Setzler, B. (2023). How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics*, *41*(2), 291–322. https://doi.org/10.1086/720009

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.

Bonhomme, S., & Manresa, E. (2015). Grouped patterns of heterogeneity in panel data. *Econometrica*, *83*(3), 1147–1184. http://www.jstor.org/stable/43616962

Borovičková, K., & Shimer, R. (2024). *Assortative matching and wages: The role of selection* (Working Paper No. 33184). National Bureau of Economic Research. https://doi.org/10.3386/w33184

Bruche, M., Farinha, L., Kokas, S., Sette, E., & Tsoukas, S. (2026). *Tracing the impact of bank shocks under bank-specific credit demand* (tech. rep.). CEPR Discussion Paper No. 20865. https://cepr.org/publications/dp21335

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

Chaney, T., Sraer, D., & Thesmar, D. (2012). The collateral channel: How real estate shocks affect corporate investment. *American Economic Review*, *102*(6), 2381–2409. https://doi.org/10.1257/aer.102.6.2381

Chodorow-Reich, G. (2014). The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis *. *The Quarterly Journal of Economics*, *129*(1), 1–59. https://doi.org/10.1093/qje/qjt031

Cingano, F., Manaresi, F., & Sette, E. (2016). Does credit crunch investment down? New evidence on the real effects of the bank-lending channel. *The Review of Financial Studies*, *29*(10), 2737–2773. https://doi.org/10.1093/rfs/hhw040

Comola, M., & Fafchamps, M. (2014). Testing unilateral and bilateral link formation. *The Economic Journal*, *124*(579), 954–976. https://doi.org/10.1111/ecoj.12071

Core, F., De Marco, F., Eisert, T., & Schepens, G. (2025). *Inflation and floating-rate loans: Evidence from the euro-area* (tech. rep.). CEPR Discussion Paper No. 20354. https://cepr.org/publications/dp20354

Crippa, F. (2025). *Identification, estimation, and inference in two-sided interaction models* (tech. rep.).

Drechsler, I., Savov, A., & Schnabl, P. (2017). The deposits channel of monetary policy. *The Quarterly Journal of Economics*, *132*(4), 1819–1876. https://doi.org/10.1093/qje/qjx019

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

Fafchamps, M., & Gubert, F. (2007). The formation of risk sharing networks [Papers from a Symposium: The Social Dimensions of Microeconomic Behaviour in Low-Income Communities]. *Journal of Development Economics*, *83*(2), 326–350. https://doi.org/https://doi.org/10.1016/j.jdeveco.2006.05.005

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

Gertler, M., & Gilchrist, S. (1994). Monetary policy, business cycles, and the behavior of small manufacturing firms. *The Quarterly Journal of Economics*, *109*(2), 309–340. https://doi.org/10.2307/2118465

Gomez, M., Landier, A., Sraer, D., & Thesmar, D. (2021). Banks' exposure to interest rate risk and the transmission of monetary policy. *Journal of Monetary Economics*, *117*, 543–570. https://doi.org/https://doi.org/10.1016/j.jmoneco.2020.03.011

Graham, B. S. (2017). An econometric model of network formation with degree heterogeneity. *Econometrica*, *85*(4), 1033–1063. https://doi.org/https://doi.org/10.3982/ECTA12679

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

Greenwald, D., Krainer, J., & Paul, P. (2024). *Monetary transmission through bank securities portfolios* (Working Paper No. 32449). National Bureau of Economic Research. https://doi.org/10.3386/w32449

Gutierrez, B., Villacorta, A., & Villacorta, L. (2025). *(unobserved) Heterogeneity in the bank lending channel: Accounting for bank-firm interactions and specialization* (tech. rep.). https://dx.doi.org/10.2139/ssrn.4608656

Gürkaynak, R., Karasoy-Can, H. G., & Lee, S. S. (2022). Stock market's assessment of monetary policy transmission: The cash flow effect. *The Journal of Finance*, *77*(4), 2375–2421. https://doi.org/https://doi.org/10.1111/jofi.13163

Hoeffding, W. (1948). A Class of Statistics with Asymptotically Normal Distribution. *The Annals of Mathematical Statistics*, *19*(3), 293 –325. https://doi.org/10.1214/aoms/1177730196

Ippolito, F., Ozdagli, A. K., & Perez-Orive, A. (2018). The transmission of monetary policy through bank lending: The floating rate channel. *Journal of Monetary Economics*, *95*, 49–71. https://doi.org/https://doi.org/10.1016/j.jmoneco.2018.02.001

Ivashina, V., Laeven, L., & Moral-Benito, E. (2022). Loan types and the bank lending channel. *Journal of Monetary Economics*, *126*, 171–187. https://doi.org/https://doi.org/10.1016/j.jmoneco.2021.11.006

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

Jiménez, G., Mian, A., Peydró, J.-L., & Saurina, J. (2020). The real effects of the bank lending channel. *Journal of Monetary Economics*, *115*, 162–179. https://doi.org/https://doi.org/10.1016/j.jmoneco.2019.06.002

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

Kiyotaki, N., & Moore, J. (1997). Credit cycles. *Journal of Political Economy*, *105*(2), 211–248. Retrieved January 29, 2026, from http://www.jstor.org/stable/10.1086/262072

Kline, P., Saggio, R., & Sølvsten, M. (2020). Leave-out estimation of variance components. *Econometrica*, *88*(5), 1859–1898. https://doi.org/https://doi.org/10.3982/ECTA16410

Kosekova, K., Maddaloni, A., Papoutsi, M., & Schivardi, F. (2025). Firm-bank relationships: A cross-country comparison. *The Review of Corporate Finance Studies*, cfaf004. https://doi.org/10.1093/rcfs/cfaf004

Lewis, D. J. (2025). Identification based on higher moments in macroeconometrics. *Annual Review of Economics*. https://doi.org/https://doi.org/10.1146/annurev-economics-070124-051419

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.

Ottonello, P., & Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, *88*(6), 2473–2502. https://doi.org/https://doi.org/10.3982/ECTA15949

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

Pietrosanti, S., & Rainone, E. (2026). *Cross-elasticities in credit markets* (tech. rep.). Banca d'Italia Temi Di Discussione No. 1436.

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

A Proof of Theorem 1

We first formally state the assumptions described in the text.

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

$$u^d_{fb} = e^d_{fb} + v^d_{fb}$$
$$u^s_{fb} = e^s_{fb} + v^s_{fb}.$$

where e^i_{fb}, $i \in \{d, s\}$ *is mean zero and independent of all innovations except for* $e^i_{fb'}$ *and* $v^i_{fb}, i \in \{d, s\}$ *is mean zero and independent of all innovations except for* $v^i_{f'b}$. *All innovations are identically distributed, have strictly positive variance and finite eighth moments, and*

$$\lim_{F,B\to\infty} \frac{B}{N^2_{FF}} \sum_{b=1}^{B} \text{var}\left(\sum_{f'\neq f} \text{vech}(v_{fb}v'_{f'b})\right) \text{ and } \lim_{F,B\to\infty} \frac{F}{N^2_{BB}} \sum_{f=1}^{F} \text{var}\left(\sum_{b'\neq b} \text{vech}(e_{fb}e'_{fb'})\right)$$

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

This structure allows either positive or negative correlation of each shock across both firms and banks; e^d_{fb} and e^s_{fb} are firm-specific demand and supply components (correlated over banks, and v^d_{fb} and v^s_{fb} are bank-specific demand and supply components (correlated over firms). The moment existence assumptions allow consistent estimation of the asymptotic variance and the variance assumptions ensure non-degenerate limiting distributions.

We next impose assumptions on the network structure.

Assumption 4. *The following limits hold:*

1.

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

2.

$$\frac{B}{F^2} \to 0 \textit{ as } F, B \to \infty;$$

3.

$$\frac{F}{B^2} \to 0 \textit{ as } F, B \to \infty.$$

While the first point of Assumption 4 does require that the average degree of a firm grows linearly with B (and vice versa for banks), this only requires that there is a non-vanishing share of firms (banks) whose number of connections increases proportionally with B (F), in contrast to the stronger firm-by-firm conditions in Jochmans and Weidner 2019. The latter points require that neither B nor F dominates the other asymptotically, but realistically allows them to grow at different rates.

We begin by proving two preliminary lemmas. We first prove that the variance of $\text{vech}(S_{FF})$ vanishes as $F, B \to \infty$.

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

Proof.

$$
\begin{aligned}
\text{var}(\text{vech}(S_{FF})) &= \text{var}\left(\frac{1}{N_{FF}}\sum_{b=1}^{B}\sum_{f'\neq f}\text{vech}\left(\eta_{fb}\eta'_{f'b}\right)\right) \\
&= N_{FF}^{-2}\text{var}\left(\text{vech}\left(\sum_{b=1}^{B}\sum_{f'\neq f}\eta_{fb}\eta'_{f'b}\right)\right) \\
&= N_{FF}^{-2}\sum_{b=1}^{B}\sum_{b'=1}^{B}\text{cov}\left(\text{vech}\left(\sum_{f'\neq f}\eta_{fb}\eta_{f'b}\right), \text{vech}\left(\sum_{\tilde{f}\neq\bar{f}}\eta_{\tilde{f}b'}\eta_{\bar{f}b'}\right)'\right) \\
&= N_{FF}^{-2}\sum_{b=1}^{B}\sum_{f'\neq f}\sum_{\tilde{f}\neq\bar{f}}\text{cov}\left(\text{vech}\left(\eta_{fb}\eta_{f'b}\right), \text{vech}\left(\eta_{\tilde{f}b}\eta_{\bar{f}b}\right)'\right) \qquad (11)\\
&+ 2N_{FF}^{-2}\sum_{b'\neq b}\sum_{f'\neq f}\text{cov}\left(\text{vech}\left(\eta_{fb}\eta_{f'b}\right), \text{vech}\left(\eta_{fb'}\eta_{f'b'}\right)'\right) \qquad (12)\\
&\equiv \omega_{FF},
\end{aligned}
$$

where in (11)-(12) terms are dropped for which $\left(\bar{f}, f\right) \neq (f, f')$ 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

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

terms, while the second summation contains

$$
\sum_{b'\neq b}\sum_{f'\neq f}\mathbf{1}_{fb}\mathbf{1}_{fb'}\mathbf{1}_{f'b}\mathbf{1}_{f'b'}
$$

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, $\mu_{F,B} \to \tilde{\mu} \in (0, 1]$. It follows that

$$
\lim_{F,B\to\infty}\frac{1}{F^2 B}\sum_{b=1}^{B}\frac{F_b(F_b-1)}{2} = \tilde{\delta} \in (0, 1/2] \qquad (13)
$$

and thus

$$\lim_{F,B\to\infty}\left(\frac{1}{F^2B}\sum_{b=1}^{B}\frac{F_b(F_b-1)}{2}\right)^2=\tilde{\delta}^2\in(0,1/4]\,,$$

and additionally

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

Then, since in the denominator $N_{FF}^2=\left(\frac{1}{2}\sum_{b=1}^{B}F_b(F_b-1)\right)^2=O\left(B^2F^4\right)$, the first term of ω_{FF} is $O(B^{-1})$, since the covariances are assumed to be finite, and thus converges to zero as $F,B\to\infty$. 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 $F,B\to\infty$. 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. The same is true for S_{BB} by symmetry, whose variance we denote as ω_{BB}. □

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

Lemma 2. *Under Assumptions 3 and 4, $S_{FF}\xrightarrow{p}\Sigma_{FF}$ and $S_{BB}\xrightarrow{p}\Sigma_{BB}$; the convergence of $q(\boldsymbol{\eta},\theta)$ 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

$$q_{FF}(\boldsymbol{\eta},\theta)=\operatorname{vech}\left(S_{FF}-A\Lambda_{FF}A'\right)$$

to $E\left[q_{FF}(\boldsymbol{\eta},\theta)\right]$ is also uniform in θ since it is entirely separable in θ and $\boldsymbol{\eta}$, where $\boldsymbol{\eta}$ denotes the sample of η_{fb}, and thus

$$\|q_{FF}(\boldsymbol{\eta},\theta)-E\left[q_{FF}(\boldsymbol{\eta},\theta)\right]\|=\|S_{FF}-\Sigma_{FF}\|,$$

independent of θ. By symmetry, the same arguments apply to $q_{BB}(\boldsymbol{\eta},\theta)$. □

Note that this result closely resembles those found in the literature on U-statistics, see for instance Hoeffding (1948). However, we opt to develop asymptotic properties directly to ensure the results are self-contained and make clear how the various assumptions contribute to the asymptotic results. We now prove Theorem 1.

Proof. Following the uniform consistency result of Proposition 2, $\hat{\theta} \xrightarrow{p} \theta$ by standard minimum distance results, for instance Theorem 2.1 of Newey and McFadden (1994), establishing the first part of the theorem.

We derive the limiting distribution of q_{FF}, q_{BB} separately. The proof of Lemma 1 contains the variances of each block; their covariance, ω_{FB} is:

$$\begin{aligned}\omega_{FB} &= \operatorname{cov}\left(\frac{1}{N_{FF}}\sum_{b=1}^{B}\sum_{f'\neq f}\operatorname{vech}\left(\eta_{fb}\eta'_{f'b}\right), \frac{1}{N_{BB}}\sum_{\tilde{f}=1}^{F}\sum_{\tilde{b}\neq b'}\operatorname{vech}\left(\eta_{\tilde{f}b'}\eta'_{\tilde{f}\tilde{b}}\right)'\right)\\ &= \frac{4}{N_{FF}N_{BB}}\sum_{f'\neq f}\sum_{b'\neq b}\operatorname{cov}\left(\operatorname{vech}\left(\eta_{fb}\eta'_{f'b}\right), \operatorname{vech}\left(\eta_{fb}\eta'_{fb'}\right)'\right).\end{aligned}$$

Thus, $\omega_{FB} = O(F^{-1}B^{-1})$ and vanishes at a faster rate than either of the variances.

Under Assumption 3, $\operatorname{vech}\left(\eta_{fb}\eta'_{f'b}\right)$ can be decomposed

$$\eta_{fb,i}\eta_{f'b,j} \equiv \alpha_{b,ff',ij} + \beta_{b,ff',ij} + \zeta_{b,ff',ij} + \xi_{b,ff',ij},$$

where

$$\begin{aligned}\alpha_{b,ff',ij} &= \left(A_{i1}e^{d}_{fb} + A_{i2}e^{s}_{fb}\right)\left(A_{j1}e^{d}_{f'b} + A_{j2}e^{s}_{f'b}\right)\\ \beta_{b,ff',ij} &= \left(A_{i1}v^{d}_{fb} + A_{i2}v^{s}_{fb}\right)\left(A_{j1}v^{d}_{f'b} + A_{j2}v^{s}_{f'b}\right)\\ \zeta_{b,ff',ij} &= \left(A_{i1}e^{d}_{fb} + A_{i2}e^{s}_{fb}\right)\left(A_{j1}v^{d}_{f'b} + A_{j2}v^{s}_{f'b}\right)\\ \xi_{b,ff',ij} &= \left(A_{i1}v^{d}_{fb} + A_{i2}v^{s}_{fb}\right)\left(A_{j1}e^{d}_{f'b} + A_{j2}e^{s}_{f'b}\right).\end{aligned}$$

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

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

since the inner summation is 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

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

jointly in F, B, where the limit follows since $O\left(\frac{BF^2}{N_{FF}}\right) \to O(1)$ by (13) and $B^{\frac{1}{2}}/F \to 0$ by Assumption 4.

Turning to $\beta_{b,ff',ij}$, $E[\beta_{b,ff',ij}] = E\left[\eta_{fb,i}\eta_{f'b,j}\right] = \Sigma_{FF,ij}$, since v^d_{fb} and v^s_{fb} are the only source of correlation across firms. Then

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

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, Lyapunov's condition holds. Then, by the Lyapunov Central Limit Theorem,

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

where

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

Joint asymptotic normality holds for $\sqrt{B}\left(\frac{1}{N_{FF}}\sum_{b=1}^{B}\sum_{f'\neq f}\beta_{b,ff'} - \text{vech}(\Sigma_{FF})\right)$, for $\beta_{b,ff'}$, the vector stacking unique entries over i, j, by the Cramer-Wold device, so

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

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

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

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 $\lim_{F,B\to\infty}\sqrt{B}O_p\left(N_{FF}^{-\frac{1}{2}}\right) = 0$ by Assumption 4. By symmetry, the same argument applies to $\xi_{b,ff',ij}$. Thus, combining the limits,

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

by Slutsky's Theorem. Stacking entries and using the joint normality established above for sums of $\beta_{b,ff'}$ yields $\sqrt{B}q_{FF}(\boldsymbol{\eta}, \theta) \xrightarrow{d} \mathcal{N}(0, \mathbf{W}_{FF})$, where $\mathbf{W}_{FF} \equiv V$. By symmetry, $\sqrt{F}q_{BB}(\boldsymbol{\eta}, \theta) \xrightarrow{d} \mathcal{N}(0, \mathbf{W}_{BB})$.

Joint asymptotic normality (as $F, B \to \infty$) of $\sqrt{B}q_{FF}(\boldsymbol{\eta}, \theta)$ and $\sqrt{F}q_{BB}(\boldsymbol{\eta}, \theta)$ follows from $\sqrt{F}\sqrt{B}\omega_{F,B} \to 0$, which ensures they are asymptotically independent. Thus

$$\begin{pmatrix} \sqrt{B}q_{FF}(\boldsymbol{\eta}, \theta) \\ \sqrt{F}q_{BB}(\boldsymbol{\eta}, \theta) \end{pmatrix} \xrightarrow{d} \begin{pmatrix} \Psi_b \\ \Psi_f \end{pmatrix} \sim \mathcal{N}(0, \mathbf{W}),$$

where

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

Regularity conditions for continuous differentiability and the Jacobian of q, denoted Φ, required for minimum distance estimation, are trivially satisfied due to the structure of q. Therefore, following standard arguments (e.g., Theorem 3.2, Newey and McFadden 1994),

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

where

$$\tilde{\Phi} = \left(\begin{bmatrix} \sqrt{B} & 0 \\ 0 & \sqrt{F} \end{bmatrix} \otimes I_3 \right) \hat{\Phi}, \quad \hat{\Phi} = \left.\frac{\partial q(\boldsymbol{\eta}, \theta)}{\partial \theta'}\right|_{\hat{\theta}}.$$

The final part of the proof establishes the consistency of estimators $\hat{\mathbf{W}}_{FF}, \hat{\mathbf{W}}_{BB}$. Note

$$\begin{aligned} \mathbf{W}_{FF} &= \lim_{F,B\to\infty} \frac{B}{N_{FF}^2} \sum_{b=1}^{B} \text{var}\left(\sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right) \\ &= \lim_{F,B\to\infty} \sum_{b=1}^{B} \frac{B(F_b(F_b-1)/2)^2}{N_{FF}^2} \text{var}\left(\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right) \qquad (14) \\ &= \frac{\tilde{\gamma}}{\tilde{\delta}^2}\, \tilde{\mu} \lim \text{var}\left(\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right) \\ &= \frac{\tilde{\gamma}\tilde{\mu}}{\tilde{\delta}^2} \Omega_{FF}, \end{aligned}$$

where $\Omega_{FF} \equiv \lim_{F\to\infty} \text{var}\left(\frac{1}{F_b(F_b-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right)$ for banks with $\lim F_b/F \in (0, 1]$. 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 $\frac{B(F_b(F_b-1)/2)^2}{N_{FF}^2} \to 0$, so their contribution to the variance in (14) is asymptotically negligible. For $\tilde{\mu}B$, the scaled summations are asymptotically independently and identically distributed.

Define

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

Adopting the decomposition in (??)-(??), $\hat{\mathbf{W}}_{FF}$ is a function of four types of terms. The argument above shows that sample averages of three of them converge in probability to zero, i.e. $\frac{1}{F^2}\sum_{f'\neq f}\alpha_{b,ff'} = O_p(F^{-1}) \to 0$, $\frac{1}{F^2}\sum_{f'\neq f}\zeta_{b,ff'} = O_p(F^{-1}) \to 0$, $\frac{1}{F^2}\sum_{f'\neq f}\xi_{b,ff'} = O_p(F^{-1}) \to 0$. Therefore, second powers involving those terms are asymptotically negligible. The only non-vanishing terms are those depending exclusively on $\beta_{b,ff'}$.

First, denote

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

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

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

for a fraction $\tilde{\mu}$ of banks as $F, B \to \infty$, and is asymptotically negligible for all others.

Since S_{FF} is consistent for Σ_{FF} by Lemma 2,

$$\begin{aligned}
&\frac{B}{N_{FF}^2}\sum_{b=1}^{B}\left(\sum_{f'\neq f}\beta_{b,ff'} - \text{vech}(S_{FF})\right)\left(\sum_{f'\neq f}\beta_{b,ff'} - \text{vech}(S_{FF})\right)' \\
&= \frac{1}{B}\sum_{b=1}^{B}\frac{B^2(F_b(F_b-1)/2)^2}{N_{FF}^2}\left(\frac{1}{F_b(F_b-1)/2}\sum_{f'\neq f}\beta_{b,ff'} - \text{vech}(S_{FF})\right) \\
&\times\left(\frac{1}{F_b(F_b-1)/2}\sum_{f'\neq f}\beta_{b,ff'} - \text{vech}(S_{FF})\right)' \\
&\xrightarrow{p} \lim_{F,B\to\infty}\frac{1}{B}\sum_{b=1}^{B}\frac{B^2(F_b(F_b-1)/2)^2}{N_{FF}^2}\nu_b\nu_b' = \tilde{\mu}\frac{\tilde{\gamma}}{\tilde{\delta}^2}\Omega_{FF} + (1-\tilde{\mu})\times 0 = \frac{\tilde{\gamma}\tilde{\mu}}{\tilde{\delta}^2}\Omega_{FF},
\end{aligned}$$

as the average of $\tilde{\mu}B$ i.i.d. random variables. Symmetric results apply to $\hat{\mathbf{W}}_{BB}$, defined analogously, completing the proof of Theorem 1.

□

Identifying Relationship-level Effects Using Covariance Restrictions

Olivier De Jonghe Daniel Lewis

ONLINE APPENDIX

July 14, 2026

B Labeling and normalization of $\hat{A}$

As noted in the text, the preliminary estimate $\tilde{A}$ requires labeling and normalization to obtain the supply and demand elasticities. In particular, let

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

where $P(\cdot)$ is a signed permutation matrix and $\tilde{\Upsilon}(\cdot)$ is a diagonal matrix imposing a scale normalization. In practice, we set $\tilde{\Upsilon} = |\tilde{\Lambda}_{FF}|^{\frac{1}{2}}$ (where $|\cdot|$ denotes the element-wise absolute value of a matrix), where $\tilde{\Lambda}_{FF} = \tilde{A}^{-1}S_{FF}\tilde{A}'^{-1}$, which normalizes the covariances of demand and supply shocks across firms (within a given bank) to be $\pm$ unity.[16] 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 minimizes the Frobenius norm of $P\tilde{A}\tilde{\Upsilon} - \begin{pmatrix} \text{std}(\Delta r_{fb}) & -\text{std}(\Delta r_{fb}) \\ \text{std}(\Delta l_{fb}) & \text{std}(\Delta l_{fb}) \end{pmatrix}$. That is, it is the signed permutation that minimizes the elementwise distance between $\hat{A}$ and $\begin{pmatrix} \text{std}(\Delta r_{fb}) & -\text{std}(\Delta r_{fb}) \\ \text{std}(\Delta l_{fb}) & \text{std}(\Delta l_{fb}) \end{pmatrix}$. Formally,

$$\hat{P}(\tilde{A}\tilde{\Upsilon}) = \text{argmin}_{P\in\mathcal{P}} \sqrt{\sum_{i=1}^{2}\sum_{j=1}^{2}\left(P\tilde{A}\tilde{\Upsilon} - \begin{pmatrix} \text{std}(\Delta r_{fb}) & -\text{std}(\Delta r_{fb}) \\ \text{std}(\Delta l_{fb}) & \text{std}(\Delta l_{fb}) \end{pmatrix}\right)_{ij}^{2}},$$

[16]Using the absolute value of $\tilde{\Lambda}_{FF}$ means that after normalization the entries of Λ_{FF} may be $+1$ or -1, depending on the covariance pattern in the data.

where $(M)_{ij}$ denotes the ij entry of matrix M. As a result, 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, $\hat{\theta}$.

C Additional theoretical results

C.1 Additional details on AKM with two-way correlations

In this appendix, we show that the results in the main text on the bias of AKM in the simple example (Section 2) are robust to alternative implementations.

Suppose instead of treating firm 2 as the reference group, we treat bank 2 as the reference group. Then we obtain

$$\begin{aligned}
\tilde{d}_1 &= d_1 + e_1 + (s_2 - a) \\
\tilde{d}_2 &= d_2 + e_2 + (s_2 - a) \\
\tilde{s}_1 &= s_1 - s_2 + (a - (-a))s_1 - s_2 + 2a.
\end{aligned}$$

The identified demand shocks are contaminated by two terms: the firm's common supply component, e_f, and the reference bank's common components (both supply and demand). $\tilde{s}_1$ should recover bank 1's supply relative to the reference bank, but it is contaminated by an additional term, the difference between the two banks' demand components. If we set the reference bank-specific components to zero to abstract from normalization, $s_2 = 0, u^d_{f2} = d_f$, we still obtain $\tilde{d}_f = d_f + e_f$ and $\tilde{s}_1 = s_1 + a$, so both supply and demand shocks remain unidentified.

Our identification strategy employs changes in prices in addition to the usual changes in quantity used in the banking literature. Suppose that Δr_{fb} is generated as

$$\Delta r_{fb} = u^d_{fb} - u^s_{fb},$$

and that we apply the AKM approach to Δr_{fb} instead of Δl_{fb}. Similar results hold, with just a flip in sign of supply-related terms. For instance, with firm 2 as the reference unit,

$$
\begin{aligned}
\tilde{d}_1 &= (d_1 - d_2) - (e_1 - e_2) \\
\tilde{s}_1 &= -s_1 + a + (d_2 - e_2) \\
\tilde{s}_2 &= -s_2 - a + (d_2 - e_2).
\end{aligned}
$$

The commentary is qualitatively unchanged from that in the text.

C.2 Partialing out covariates

In this section, we show that covariates, X_{fb}, can be partialed out of the regression model in (7), under some additional assumptions.

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

$$
E\left[u_{fb} | X_{fb}\right] = E\left[u_{fb}\right].
$$

Lemma 3. *In the model $\eta_{fb} = Au_{fb} + \Gamma X_{fb}$, Γ can be consistently estimated by OLS under Assumptions 1-5 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 var$\left(Au_{fb}X'_{fb}\right) < \infty$.*

Proof. By the usual algebra, the OLS estimator of Γ is

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

The expectation of the numerator of the error term is zero by Assumption 1: by the law of iterated expectations,

$$
\begin{aligned}
E\left[\frac{1}{N}\sum_{f=1}^{F}\sum_{b=1}^{B} Au_{fb}X'_{fb}\right] &= \frac{1}{N}\sum_{f=1}^{F}\sum_{b=1}^{B} AE\left[E\left[u_{fb}|X\right]X'_{fb}\right] \\
&= \frac{1}{N}\sum_{f=1}^{F}\sum_{b=1}^{B} AE\left[E\left[u_{fb}\right]X'_{fb}\right] \\
&= 0,
\end{aligned}
$$

using the conditional mean independence of Assumption 5. The variance of the numerator is

$$
\begin{aligned}
&\text{var}\left(\frac{1}{N}\sum_{f=1}^{F}\sum_{b=1}^{B} Au_{fb}X'_{fb}\right) \\
&= \frac{1}{N^2}\sum_{f=1}^{F}\sum_{b=1}^{B}\sum_{f'=1}^{F}\sum_{b'=1}^{B}\text{cov}\left(Au_{fb}X'_{fb}, X_{f'b'}Au'_{f'b'}\right) \\
&\leq \frac{F^2B^2}{N^2}\frac{1}{F^2B^2}\sum_{f=1}^{F}\sum_{b=1}^{B}(F+B-1)\text{var}\left(Au_{fb}X'_{fb}\right) \\
&= \frac{F^2B^2}{N^2}\frac{(F+B-1)}{FB}\text{var}\left(Au_{fb}X'_{fb}\right) \to 0,
\end{aligned}
$$

where the inequality uses the independence of X_{fb} for observations with no common indices. Applying Chebyshev's Inequality shows that the numerator converges in probability to 0. Since the correlation pattern in the denominator is identical, the variance of the denominator also vanishes as $F, B \to \infty$, so the denominator converges in probability to $E\left[X_{fb}X'_{fb}\right]$. Finally, applying Slutsky's Theorem yields $\hat{\Gamma} \xrightarrow{p} \Gamma$. □

D Monte Carlo Study

In this Appendix, we describe our Monte Carlo study in detail. In our simulations, we set

$$
A = \begin{bmatrix} 7.5101 & -1.5346 \\ 0.1670 & 2.2644 \end{bmatrix}.
$$

which corresponds to empirical estimates for Italy in the monetary policy tightening subsample. We also match the second moments of the supply and demand shocks to those estimated in that specification. We generate the shocks according to:

$$
u^i_{fb} = z^i_f + z^i_b + z^i_{fb}, \quad i \in \{d, s\}, \; f = 1, \dots, F, \; b = 1, \dots, B,
$$

where each of the components is independent and normally distributed with a mean of zero and empirically calibrated variance.

We consider four 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 4) with 100 banks, and a large sample with 250 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

Table D.1: Bias and standard deviation of parameter estimates

	$B = 10$		$B = 25$		$B = 100$		$B = 250$	
$T = 1$, i.i.d.	bias	std	bias	std	bias	std	bias	std
A_{11}	-0.12	0.21	-0.06	0.14	-0.02	0.07	-0.00	0.04
A_{21}	-0.35	4.64	-0.22	2.64	0.02	1.03	0.01	0.37
A_{12}	-0.59	1.66	-0.13	1.43	-0.01	0.75	-0.00	0.43
A_{22}	-0.26	0.59	-0.10	0.39	-0.00	0.07	-0.00	0.05
$T = 4$, i.i.d.	bias	std	bias	std	bias	std	bias	std
A_{11}	-0.04	0.12	-0.01	0.07	-0.01	0.04	-0.00	0.02
A_{21}	-0.07	1.97	-0.05	1.08	-0.00	0.51	-0.01	0.20
A_{12}	-0.03	1.19	0.04	0.76	-0.00	0.37	0.01	0.22
A_{22}	-0.04	0.23	-0.00	0.07	-0.00	0.04	-0.00	0.02

Notes: The table reports the bias and standard deviation (both relative to the corresponding entry of A) of estimates of each entry in A across 1,000 Monte Carlo samples generated according to the DGPs described in the text. The top panel considers four sample sizes for $T = 1$, with shocks independent over time: $B = 10, 25, 100, 250$, and $F = 1000B$ in each case. The second panel considers $T = 4$ time periods, where the shocks are independent in each period.

network structure observed 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 D.1 reports the relative bias and relative standard deviation of each entry of $\hat{A}$ over 1,000 samples. Starting with the $T = 1$ case, for $B = 10$, which is a smaller sample size than in all countries in our empirical application, the bias can be fairly large. These results are potentially driven by the fact that, with sample sizes this small (and with a relatively sparse network), complex values for $\hat{A}$ are obtained in some samples, so that these statistics are computed only across 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 when B increases, though it remains sizable for two coefficients. Finally, for $B = 250$, all biases are negligible, and standard deviations are further reduced. As T increases to 4, the bias falls substantially, essentially vanishing for $B > 25$, and the standard deviation also decreases, as expected. 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 our empirical analysis. Overall, the estimation strategy performs well with realistic sample sizes and suggests that our results are likely to be reliable for most countries studied.

Table D.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 in the case of $T = 1$, and moderate in the case of $B = 25$. When $B \geq 25$ and $T = 4$, size is better controlled, as expected.

Table D.2: Empirical size of t-tests

	$B = 10$	$B = 25$	$B = 100$	$B = 250$
$T = 1$				
A_{11}	18.3	13.2	8.3	5.4
A_{21}	15.8	10.2	4.8	3.4
A_{12}	11.1	10.7	7.3	4.1
A_{22}	25.0	12.8	6.1	6.2
$T = 4$				
A_{11}	10.3	7.8	6.7	4.8
A_{21}	9.3	6.8	6.1	6.1
A_{12}	7.4	5.5	5.4	4.9
A_{22}	9.7	5.6	5.5	5.5

Notes: The table reports the e empirical rejection rates of 5% t-tests for each entry in A across 1,000 Monte Carlo samples generated according to the DGP described in the text. For each specification, rejection rates are reported using the baseline variance estimator that assumes serially uncorrelated shocks. The top panel considers four sample sizes for $T = 1$, with shocks independent over time: $B = 10, 25, 100, 250$, and $F = 1000B$ in each case. The second panel considers $T = 4$ time periods, where the shocks are independent in each period.

Table D.3 considers the problem of estimating the average demand shock for each firm and average supply shock for each bank. Given the DGP described, there are indeed firm and bank common components corresponding to each of these objects, to which the average heterogeneous shocks should converge. In particular, the "demand" firm effects are z_f^d and the "supply" bank effects are z_b^s, and we have $E\left[B_f^{-1}\sum_{b=1}^{B_f} u_{fb}^d | z_f^d\right] = z_f^d$ and $E\left[F_b^{-1}\sum_{f=1}^{F_b} u_{fb}^s | z_b^s\right] = z_b^s$. 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). Finally, for comparison, we also compute the correlation of each estimated firm and bank fixed effect with the true firm and bank components of quantity (which mix both demand and supply effects)[17].

We report the average correlations across 1000 samples. For all specifications and estimators, the correlations are much stronger for bank fixed effects than for firm fixed effects, reflecting the much larger number of relationships for each bank relative to each firm. Across specifications, the firm average heterogeneous demand shock is substantially better correlated with the firm component of demand than the estimated firm fixed effect – by a factor of 7 to 10. This is due to two factors. First, the fixed effects estimates are biased for the firm component of demand, since there is also a firm-specific component of supply in the DGP. Second, as explained in the text, although the firm fixed effect can only make use of a firm's very limited number of relationships, our new estimator leverages information across the

[17] In particular, these are $A_{21}z_f^d + A_{22}z_f^s$ and $A_{21}z_b^d + A_{22}z_b^s$, respectively

Table D.3: Correlations of average shocks and fixed effects

	$B = 10$		$B = 25$		$B = 100$		$B = 250$	
$T = 1$, i.i.d.	firm	bank	firm	bank	firm	bank	firm	bank
corr(True, Avg. Het.)	0.22	0.73	0.23	0.90	0.24	0.99	0.30	1.00
corr(True, FE)	0.03	0.99	0.03	0.99	0.03	0.99	0.04	0.99
corr(True Q, FE)	0.20	0.99	0.20	0.99	0.20	0.99	0.25	1.00
$T = 4$, i.i.d.	firm	bank	firm	bank	firm	bank	firm	bank
corr(True, Avg. Het.)	0.23	0.97	0.24	0.99	0.24	1.00	0.30	1.00
corr(True, FE)	0.03	0.95	0.03	0.98	0.03	0.99	0.04	0.99
corr(True Q, FE)	0.20	0.95	0.20	0.98	0.20	0.99	0.25	1.00

Notes: The table reports the average correlations across 1,000 Monte Carlo samples of the average estimated heterogeneous demand and supply shocks and the true firm components of demand and bank components of supply (corr(True, Avg. Het.)), the average correlations of the estimated fixed effects with the true firm components of demand and bank components of supply (corr(True, FE)), as well as the average correlations of the estimated fixed effects with the true fixed effects in the quantity equation (corr(True Q, FE)) in the DGP described in the text. The top panel considers four sample sizes for $T = 1$: $B = 10, 25, 100, 250$, and $F = 1000B$ in each case. The second panel considers $T = 4$ time periods.

entire dataset to separate supply and demand factors. For $B = 10$ and $B = 25$ with $T = 1$, the estimated bank fixed effects have slightly higher correlations with the true bank supply effects than the bank average heterogeneous supply shocks. However, as more observations become available (through more firms/banks and/or time periods) the bank average supply shocks become more accurate, as the advantage of exploiting information from the whole dataset grows more pronounced. The third set of correlations in each panel help to quantity the relative importance of the misspecification bias of fixed effects for the shock components versus network structure exploited in estimator performance. In particular, the correlations of the estimated firm fixed effects with the true firm components of quantity are considerably higher – although still lower than those for the average heterogeneous shocks – suggesting that the misspecification bias channel is more important than the network structure exploited in explaining the poor correlation of fixed effects estimates with the firm components of demand.

E Additional empirical results

E.1 Sample composition

We first summarize the number of firms, banks, and relationships by country and time period in Table E.4. This table characterizes the network structure present in the data.

Table E.4: Summary Statistics by Country and Time Period

	Pandemic			Inflation			Tightening		
Country	Firms	Banks	Pairs	Firms	Banks	Pairs	Firms	Banks	Pairs
AT	16,799	338	45,248	18,705	456	50,592	19,256	417	51,551
BE	28,195	20	61,387	28,478	21	62,266	30,121	21	65,771
DE	137,938	886	373,311	140,317	827	370,753	148,497	794	396,387
ES	142,902	100	426,001	152,289	101	469,943	149,558	99	434,075
FR	169,106	134	395,590	189,810	139	449,290	190,769	138	449,919
GR	7,458	16	20,983	8,802	15	24,709	9,714	14	23,759
IT	250,545	214	757,802	253,576	203	758,457	242,958	195	719,489
NL	4,201	19	8,732	3,729	20	7,714	3,930	21	8,477
PT	33,610	111	92,805	36,520	104	98,368	37,451	99	100,668

Notes: This table reports the sample sizes by country for each of the three sub-samples studied. The three time periods are: Pandemic (2019Q3-2020Q4), Inflation (2021Q1-2022Q2), and Tightening (2022Q3-2023Q4). For each country and period, we report the number of multi-bank firms (firms borrowing from at least two banks), the number of banks lending to these firms, and the total number of firm-bank relationships (N).

E.2 Supply and demand curves over time

Figure E.1 presents the results found in Figure 2 in the main text as supply and demand curves. This is a more familiar depiction of the slopes that are represented as angles in that figure. Naturally, the conclusions and patterns discussed in the main text apply to this figure as well.

E.3 The level of aggregation and the constant-A assumption

Treating A as constant within a country-period reflects a choice about the level of aggregation, not a restriction imposed by our identification strategy. Estimating a separate A for each country in a currency area already makes this choice explicit: applied to U.S. data, an analogous decision would be whether to let A vary by state rather than treat the country as a whole. The covariance restrictions we exploit can be estimated at any level of (dis)aggregation, provided enough firms borrow from multiple banks and enough banks lend to multiple firms within a cell to identify the relevant second moments. Pooling more coarsely reduces sampling variance at the risk of averaging over genuine heterogeneity; splitting further reverses this trade-off, shrinking the firms and, especially, the banks entering the two-way clustered variance estimator and widening the resulting confidence intervals. Whether a finer partition is informative or merely noisier is thus an empirical question, governed by the same bias-variance trade-off that arises whenever a homogeneity restriction is imposed to gain precision.

Figure E.2 illustrates this trade-off, re-estimating $\hat{A}$ for Spain in 2022Q3–2023Q4 after partitioning the sample by sector, region, and firm-size quartile. We choose this country-period purely for illustration: Spain is large and heterogeneous enough across regions, sectors, and firm sizes that each partition still retains sufficient firms and banks to estimate $\hat{A}$ with confidence intervals of comparable magnitude to the pooled estimate. We make no claim that this example is representative of the other 26 country-periods, nor that A would remain estimable under arbitrarily fine partitions of a smaller or more homogeneous sample. The exercise demonstrates only that the constant-A assumption is a choice available to the researcher, weighed against the same bias-variance considerations as any homogeneity restriction, and one that can be tested directly whenever the data permit.

Beyond serving as a robustness check on the constant-A assumption, the ability to estimate A at a finer level, when the data permit, speaks to questions in which cross-sectional heterogeneity in demand and supply elasticities is itself the object of interest. Quantifying how strongly credit supply and its pricing respond to policy shocks across sectors, regions, or firm sizes is a key input for heterogeneous-firm models of monetary transmission, and for calibrating targeted macroprudential tools, such as sectoral capital buffers or size-dependent credit guarantee schemes, whose rationale rests on such elasticities differing enough across groups to justify differentiated treatment.

E.4 Assessing validity using price responses

In this Appendix, we repeat the analysis of Table 2, except that we now collapse the dependent variable, the change in interest rate at the firm or bank level. Table E.5 reports the results; qualitatively, the findings are essentially identical to those in the main text, for regressions conducted at the relationship level.

Figure E.1: Supply and Demand Curves across Countries and Time

(a) Pandemic

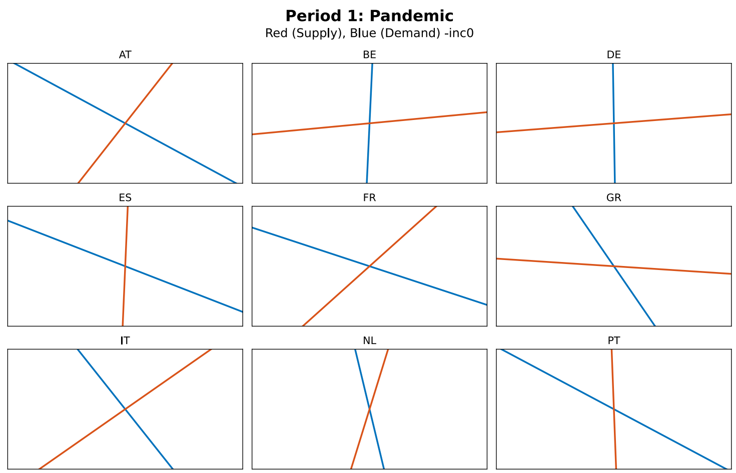

(b) Inflationary

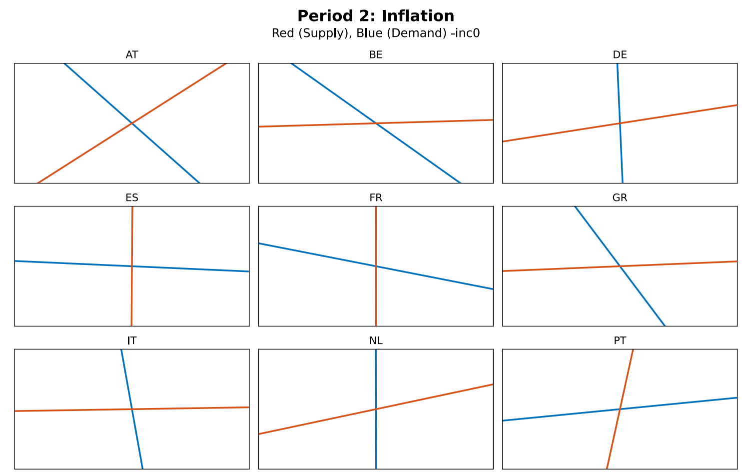

(c) Tightening

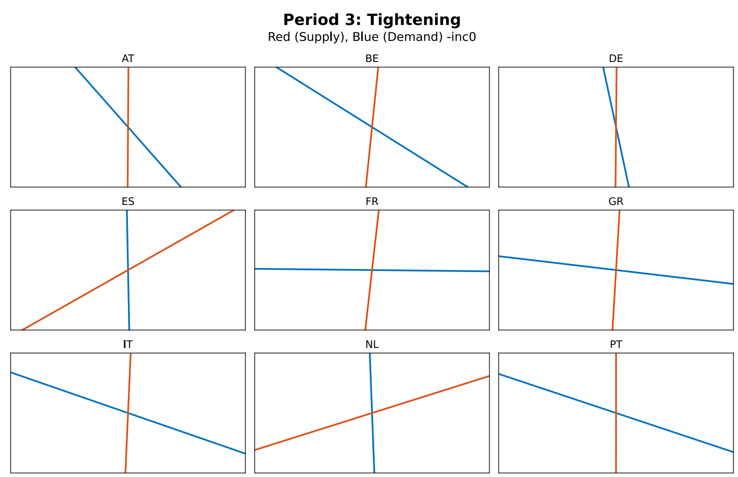

Notes: For each country, each panel plots the estimated supply and demand curves for the indicated period, constructed from $\hat{A}$ using the method described in the text. Confidence intervals are omitted so as not to crowd the figures.

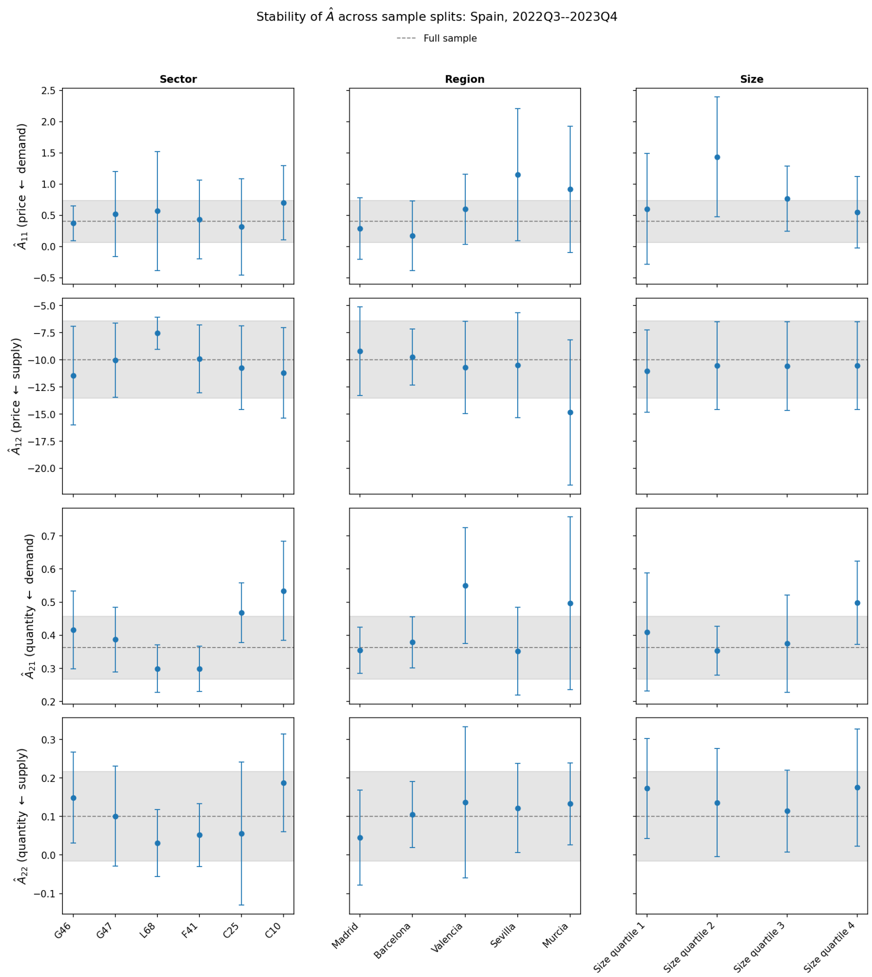

Figure E.2: Within-country heterogeneity of Supply and Demand curves

Notes: This figure assesses the stability of the estimated response matrix $\hat{A}$ within a single country-period, Spain, 2022Q3–2023Q4, when the sample is further partitioned along three dimensions: two-digit NACE sector (left column), province of the firm's headquarters (middle column), and quartile of firm size (right column). Rows report, in turn, the response of price to demand shocks ($\hat{A}_{11}$), the response of price to supply shocks ($\hat{A}_{12}$), the response of quantity to demand shocks ($\hat{A}_{21}$), and the response of quantity to supply shocks ($\hat{A}_{22}$). Size quartiles are based on (log) total assets, with firms ranked and re-assigned to a quartile separately in each quarter; a firm's quartile can therefore change from one quarter to the next within the 6-quarter estimation window, so "Q1" denotes the bottom quartile of the size distribution in a given quarter rather than a fixed panel of firms. Sectors and regions are not selected by an explicit statistical rule; but the five regions shown are in general Spain's largest by population, employment, and GDP, and the six sectors are chosen for their economic relevance. Blue points are the sub-sample point estimates, with error bars denoting 95% confidence intervals constructed using autocorrelation robust standard errors. The dashed horizontal line and shaded band in each panel report the point estimate and 95% confidence interval for the full Spain, 2022Q3–2023Q4 sample, i.e., one of the 27 country-period estimates underlying Figure 1. Within the sector and region columns, sub-samples are ordered from the largest to the smallest by number of firms; size quartiles are ordered Q1 (smallest) to Q4 (largest). Sector codes are: C10 (Manufacture of food products), C25 (Manufacture of fabricated metal products, except machinery and equipment), F41 (Construction of buildings), G46 (Wholesale trade, except of motor vehicles and motorcycles), G47 (Retail trade, except of motor vehicles and motorcycles), and L68 (Real estate activities).

Table E.5: Innovations versus fixed effects: Impact on Interest rate changes

	(1)	(2)	(3)	(4)
	Firm-Time average of dP (residualised)			
Average Demand innovation	0.229***			0.234***
	(0.000)			(0.000)
Average Supply innovation	-0.189***		-0.263***	
	(0.000)		(0.000)	
AKM Firm-Time FE		-0.483***	1.179***	
		(0.002)	(0.003)	
Firm-time average of AKM Bank-Time FE		-0.741***		-1.102***
		(0.007)		(0.006)
Observations	11601979	11539414	11539414	11539414
Adjusted R^2	0.47	0.01	0.25	0.29
	Bank-Time average of dP (residualised)			
Average Demand innovation (b,t)	0.376***			0.499***
	(0.017)			(0.017)
Average Supply innovation (b,t)	-0.362***		-0.406***	
	(0.008)		(0.009)	
Bank-Time average of AKM Firm-Time FE		-0.401***	0.223*	
		(0.129)	(0.117)	
AKM Bank-Time FE		-0.439***		-2.256***
		(0.077)		(0.190)
Observations	31823	31148	31148	31148
Adjusted R^2	0.62	0.08	0.40	0.42

Notes: The table reports coefficients from regressions of interest rate changes, collapsed at either firm- or bank-level, on supply and demand measures. In the top panel, the dependent variable is the quarterly value-weighted average change in the interest rate for firm f in period t, after residualizing relationship-level controls $X_{fb,t}$. In the bottom panel, the dependent variable is the quarterly value-weighted average change in the interest rate for bank b in period t, after residualizing relationship-level controls $X_{fb,t}$. In column 1, Demand innovation and Supply innovation are our firm- or bank-level aggregated relationship-level shocks identified using covariance restrictions. In column 2, the independent variables are AKM firm-time fixed effects (f, t) and AKM bank-time fixed effects (b, t) obtained from the regression $\Delta l_{fb,t} = d_{f,t} + s_{b,t} + \Gamma X_{fb,t} + \varepsilon_{fb,t}$, where $\Delta l_{fb,t}$ is the quarterly change in credit from bank b to firm f in period t. Column 3 tests if firm-time FEs recover demand conditional on substituting our identified supply measure. Column 4 tests if bank-time FEs recover supply conditional on substituting our identified demand measure. All columns include country×time fixed effects. Standard errors are clustered at the firm level in panel 1 and bank level in panel 2. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.