



EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Olivier De Jonghe, Daniel Lewis

Identifying relationship-level effects
using covariance restrictions



No 3238

Challenges for Monetary Policy Transmission in a Changing World Network (ChaMP)

This paper contains research conducted within the network “Challenges for Monetary Policy Transmission in a Changing World Network” (ChaMP). It consists of economists from the European Central Bank (ECB) and the national central banks (NCBs) of the European System of Central Banks (ESCB).

ChaMP is coordinated by a team chaired by Philipp Hartmann (ECB), and consisting of Diana Bonfim (Banco de Portugal), Margherita Bottero (Banca d’Italia), Emmanuel Dhyne (Nationale Bank van België/Banque Nationale de Belgique) and Maria T. Valderrama (Oesterreichische Nationalbank), who are supported by Melina Papoutsi and Gonzalo Paz-Pardo (both ECB), 7 central bank advisers and 8 academic consultants.

ChaMP seeks to revisit our knowledge of monetary transmission channels in the euro area in the context of unprecedented shocks, multiple ongoing structural changes and the extension of the monetary policy toolkit over the last decade and a half as well as the recent steep inflation wave and its reversal. More information is provided on its [website](#).

Abstract

We propose a new model in which relationship-specific effects or shocks are identified in a bipartite network under mild covariance restrictions, generalizing the influential Abowd et al. (1999) framework. For example, separate demand shocks are identified for each bank from which a firm borrows. We show how previous approaches break down when confronted with such heterogeneity, while our novel identification strategy yields a simple estimator that is consistent and asymptotically normal, under weaker network density assumptions than previous approaches. The methodology performs well in empirically-calibrated simulations. We apply our approach to identify relationship-level credit demand and supply shocks for thousands of firms and banks across nine Euro-area countries and three distinct economic episodes. We formally reject the Abowd et al. (1999) assumptions in nearly every country-period and show that within-firm/bank shock variation is of comparable scale to between firm/bank variation. We document considerable bias in Abowd et al. (1999) style estimates and associated regressions, while finding significant deleterious effects of the post-2022 monetary contraction on exposed firms. We highlight novel heterogeneity in the transmission of monetary policy.

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

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

Non-Technical Summary

Many important questions in economics concern how outcomes are shaped by a specific *relationship* between two parties: the wage a worker earns at a particular firm, the price a consumer pays to a particular retailer, the volume of credit a firm obtains from a particular bank. In each case, the observed outcome reflects contributions from both sides of the relationship – the worker’s productivity and the firm’s pay policy, the consumer’s preferences and the retailer’s pricing, the firm’s creditworthiness and the bank’s lending conditions. Separating these two-sided effects is central to empirical work across labour economics, industrial organization, international trade, and finance.

The dominant tool for this separation is the two-way fixed effects estimator, introduced by Abowd et al. (1999) – henceforth AKM – and since applied to a vast range of settings involving bipartite networks, that is, networks in which two distinct types of agents form relationships with one another but not with agents of their own type. AKM decomposes the observed outcome into an effect attributed to one side and an effect attributed to the other. This approach assumes that each agent’s effect is *the same across all of its relationships*. A worker is assumed to bring identical productivity regardless of which firm employs her; a firm is assumed to have homogeneous credit demand across banks. This homogeneity assumption is both conceptually restrictive and, as we show, empirically untenable.

The contribution of this paper is twofold. First, we develop a new econometric framework applicable to any bipartite network setting in which two outcome variables are jointly observed for each relationship. The core goal is to decompose this pair of outcomes into two underlying structural shocks – one associated with each side of the market – at the level of the individual relationship, rather than at the level of the agent. We explain this methodology in the context of corporate credit markets, where the two outcomes are loan quantities and interest rates, and the two structural objects of interest are credit demand and supply. In this setting, rather than assigning a single demand shock to each firm (as if its demand were the same across all its lenders) or a single supply shock to each bank (as if its supply conditions were the same for all its borrowers), we recover a separate demand and supply shock for every firm-bank pair in every period. By exploiting the covariance structure of prices and quantities across an agent’s relationships, we can separately identify relationship-level demand and supply without imposing distributional assumptions or restricting the degree of heterogeneity. Importantly, the correlations in shocks that invalidate the standard two-way fixed effects approach actually facilitate our identification, rather than preventing it. We also establish the formal statistical properties of our estimator and propose a simple, implementable test of the two-way fixed effects assumptions that the literature has

so far lacked.

Second, we apply this framework to nine Euro-area countries over the period 2019 to 2023, making use of AnaCredit. Our sample spans three distinct economic periods: the Covid pandemic, the post-pandemic inflation surge, and the subsequent monetary policy tightening. Our test of the homogeneity assumption yields unambiguous results: we reject it in 25 of 27 country-period combinations studied, confirming that standard approaches are unlikely to produce reliable estimates in this setting. We further show that the heterogeneity being masked is economically large. Within-firm variation in credit demand across relationships is roughly one-third as large as variation across firms; within-bank variation in credit supply across borrowers is actually *larger* than the variation across banks.

The consequences of imposing homogeneity are not merely statistical. We document a stark illustration of the resulting bias: when demand is measured using the standard firm-time fixed effect —as in Khwaja and Mian (2008) or Amiti and Weinstein (2018)— a positive demand shock is associated with a *lower* interest rate – the opposite of what economic theory predicts. The reason is that the fixed effect conflates genuine demand with supply factors, because banks’ idiosyncratic supply conditions contaminate what is labelled “firm demand”. This contamination propagates into regressions for firm outcomes. Using the 2022 monetary policy tightening by the ECB – which raised rates by 450 basis points in fourteen months – as a quasi-natural experiment, we show that firms borrowing at a floating rate (or via short-term loans) experienced a significant contraction in credit supply relative to firms that had locked in fixed rates before the tightening. When demand is controlled for using the standard fixed effects approach, this estimated effect is attenuated to near-zero. Our identified demand measure restores a large and highly significant effect. The difference translates into firms’ real activity: floating-rate borrowers experienced significantly lower total assets, turnover, and credit-to-assets ratios over the following two years, effects that would be invisible or severely underestimated using conventional methods.

Finally, because our recovered shocks vary at the relationship level, we can now study how monetary policy shocks propagate heterogeneously across firm-bank relationships. For example, following a contractionary monetary policy surprise, firms with a higher share of fixed-rate loans reduce their demand for new credit while banks simultaneously *increase* their supply to these same firms, viewing them as safer counterparties.

These findings have broad implications. They suggest that a reassessment of a large body of empirical work on credit markets may be warranted, and that access to price data alongside credit quantities can substantially sharpen identification. The methodology is general – applicable to labour markets, trade networks, and industrial organization settings.

1 Introduction

Many econometric settings require the decomposition of the outcome of an observed match into two components – one associated with each side. For instance, in seminal work, Abowd et al. (1999) (hence AKM) decompose log wages into worker and employer characteristics using two-way fixed effects. However, this problem is present in any bipartite network: that is, a network with two types of agents who interact with agents of the opposite type, but not their own type (see, e.g., Bonhomme 2020). For instance, in influential papers, Khwaja and Mian (2008) and Amiti and Weinstein (2018) use fixed effects to decompose relationship-level observations between firms and banks into credit demand and supply shocks, and Abowd et al. (2004) consider relationships between importers and exporters. Other applications abound throughout finance, production networks, and international trade.

Indeed, the majority of existing “reduced-form” research has followed the AKM paradigm, whereby worker and firm fixed effects (or firm and bank fixed effects) are used to recover worker- and firm-level effects (demand and supply shocks) using two-way fixed effects. This framework imposes the often unrealistic assumption that these effects are homogeneous across a worker’s employers and a firm’s workers, or across a firm or bank’s credit relationships. These assumptions have been called into question in recent work. For instance, Borovičková and Shimer (2024) propose a model in which workers obtain a higher surplus from firms with similar productivity to their own. Paravisini et al. (2023) study the specialization of banks (across export markets), a well-documented phenomenon that is also inconsistent with the AKM assumption that each firm has a uniform demand shock across relationships, and each bank issues a uniform supply shock to all firms. Alternatively, banks may ration credit based on correlated firm riskiness, and firms value relationship lending (e.g., Anderson et al. 2022), plausibly generating heterogeneity in supply and demand shocks, respectively. Indeed, a model with homogeneous shocks precludes important policy questions, for instance how credit supply may change differentially across firms following changes to regulatory or monetary policy.

In this paper, we generalize the AKM model to accommodate and identify heterogeneous relationship-level effects, while nesting the AKM model with i.i.d. errors as a special case. To do so, we add a second observable. For instance, in the labour market setting, this could be the duration of a given match, while for corporate credit it is natural to study the interest rate on credit, along with its volume (capturing price and quantity). We assume that this bivariate relationship-level vector is a linear combination of two relationship-level effects, for instance, relationship-level demand and supply shocks, or match-specific worker and employer effects. Rather than assuming that effects are homogeneous across relationships, we assume that

they are *correlated*. Crucially, we allow two-way correlation – so demand and supply shocks can each be correlated across both firms and banks. We illustrate how this realistic feature leads to bias in AKM-type models. Identification exploits the covariance patterns of this observable vector across firms and across banks, very similarly to how heteroskedasticity can furnish identification in time series settings (following, for instance, Rigobon (2003); see Lewis (2025) for a review). In a simple example, we show that our estimator can be interpreted as an “internal instruments” procedure, leveraging information contained across relationships to generate an instrument.

Our identification strategy gives rise to a simple, closed-form estimator of the matrix linking observables and relationship-level effects. We establish its consistency and asymptotic normality under some additional assumptions on the relationship-level effects and the network structure. We exploit these results to propose a simple formal statistical test of key AKM assumptions, under which our model is over-identified. In a recent paper, Jochmans and Weidner (2019) establish conditions for consistency of two-way fixed effects estimates, like AKM estimates. Loosely speaking, these require the number of connections of each agent (or its direct neighbours) to tend to infinity in order for the decomposition to be consistent. In contrast, we only require that a non-vanishing share of agents on each side of the market are similarly “well-connected” (their number of relationships increases with the number of agents of opposite type) in order to consistently decompose *all effects*, even those for agents with only a handful of relationships. This assumption is far more plausible in our application, where most multi-bank firms have two or three lenders, but the largest firms may have many more. It is also consistent with most workers having only a few employers over their careers. Inference adapts standard variance estimators, and we provide a Stata/Python routine to implement our procedure. Empirically-calibrated simulations (based on the Italian credit market) illustrate the accuracy of our estimator in realistic sample sizes, that our inference procedures achieve correct coverage, and that the estimated shocks compare favourably to those based on AKM.

Our approach hinges on two key assumptions, which reflect the limitations our methodology. First, we assume that the same matrix of coefficients relates observables and relationship-level effects for all observations. For applications where each agent has multiple relationships simultaneously (for instance, corporate credit markets), there is no requirement that this matrix is constant over time, however. The plausibility of this assumption is ultimately an empirical question, but the econometrician can restrict the sample to weaken it – for instance, limiting attention to firms of a similar size in a given industry. This assumption also means that we view our method as best-suited for analysing either the intensive margin (as in corporate credit markets, in our empirical application) or the extensive margin (as in

labour markets), but not both simultaneously, since different coefficients could hold for each margin. Second, the key assumption on the effects themselves that defines our decomposition is that effects of opposite type are orthogonal across relationships. This condition is consistent with macroeconomic notions of structural shocks. These structures may appear restrictive relative to those imposed by the fixed effects of AKM, but only because that framework does not consider relationship-level effects at all. This assumption may also be incompatible with some mechanisms for endogenous network formation. Additionally, our method is designed to capture cross-sectional variation in effects/shocks, not aggregate time series variation; indeed, in our application, the shocks are normalized to be mean zero in each period. Our approach is also only applicable when agents have multiple relationships (and some agents have many relationships), for instance banks lend to many firms and firms borrow from multiple banks, or several employment spells are observed for workers, at firms with multiple employees (however, this is also the case for AKM). The requirement of constant effects, restriction to the intensive margin and agents with multiple relationships, and the covariance restrictions that define the shocks are the main limitations of our methodology.

Throughout the paper, we interpret our methodology and assumptions through the lens of two examples: supply and demand for corporate credit and the classic AKM setting of worker-firm match quality. The former can easily map to a range of other financial market settings. However, the methodology also appears well-suited to product markets and applications in industrial organization more broadly, where identifying supply and demand is crucial. For instance, Nielsen-type consumer data contains quantity and price data for many consumers who buy comparable products from multiple sellers, satisfying the requirements for our approach.

We apply our identification strategy to recover relationship-level demand and supply shocks for thousands of firms and banks in nine Euro-area credit markets at quarterly frequency from 2019 to 2023, using the AnaCredit dataset. We present seven key sets of empirical facts. First, we document considerable heterogeneity in the slopes of the demand and supply curves across countries, although they remain relatively stable within a given country over time. Second, using our novel statistical test, we are able to formally reject the AKM assumptions for the vast majority of country-periods. Third, we recover considerable within-firm and bank variation in demand and supply shocks – on the same order of magnitude as that between firms and banks. Fourth, we present evidence of substantial bias in AKM-style demand and supply effects estimated in this data, following Khwaja and Mian (2008). In particular, demand effects appear contaminated by supply shocks, resulting in a positive demand shock *lowering* interest rates. Fifth, as a result, using AKM-style demand measures as controls in regressions following Khwaja and Mian (2008) and others

leads to severely attenuated estimates of the effect of the post-2022 monetary contraction on credit for firms exposed to rising interest rates (those who borrow at a floating rate or short-term); controlling for our demand measures delivers a significant and substantial reduction in credit growth for these firms. Sixth, this supply-driven reduction in borrowing leads to significantly worse firm-level outcomes in terms of total assets, turnover, and credit to total assets ratios, with AKM-style estimates once more attenuated. Finally, we examine how the incidence of relationship-level demand and supply shocks varies with relationship-level characteristics following identified monetary policy shocks. We find relatively lower demand for credit from firms with a higher share of fixed-rate loans following news of higher interest rates; on the other hand, the relative supply of credit to these firms increases, consistent with banks assessing lower interest rate risk.

These empirical results illuminate a bountiful research agenda. First, our methodology – by recovering relationship specific shocks – can be used to analyse how firms are differentially impacted by policy actions, like the contractionary monetary policy shocks we study in Section 4.5. This is valuable information for policymakers concerned with the distributional effects of policy, not just the aggregate credit supply. We believe that empirical results of this nature are novel, and thus provide a new set of empirical facts to discipline theoretical models of the transmission of various types of policies. Second, our results provide a laboratory to assess the plausibility of possible identifying assumptions for settings with less-rich data. For instance, the results in Section 4.1 highlight that, in certain countries, credit growth can be viewed as essentially supply-driven, with demand channels negligible. Finally, our results demonstrating the bias of AKM-style approaches – coupled with contemporary findings by Bergman et al. (2025), Bruche et al. (2026), Gutierrez et al. (2025), and Pietrosanti and Rainone (2026) – suggest that a reevaluation of such regressions, commonly following Khwaja and Mian (2008), may be necessary. With the possibility of more reliably identified controls for demand, researchers may find stronger effects of supply disturbances.

This paper connects to several literatures. We are not aware of any direct antecedents of the model we propose, but it nests as special cases both two-way fixed and random effects approaches to modeling heterogeneity in bipartite networks (e.g., AKM, Chamberlain 1980). Our model is more general; besides adding a second equation, it relaxes the homogeneity assumption and allows for arbitrary correlation of both types of effects across both associated agents' relationships, which can be interpreted as relaxing the standard assumption of i.i.d. errors. Some papers, like Graham et al. (2014), Amiti and Weinstein (2018), and Crippa (2025) consider non-linear random or fixed effects models, but the resulting relationship-level effects are still an arbitrary function of homogeneous effects.

Our identification argument is most closely related to identification via heteroskedastic-

ity. For instance, Rigobon (2003) shows how, in timeseries data, two different covariance matrices, taken from different regimes, can be used to identify simultaneous equations. Our strategy applies the same intuition but to two different types of cross-sectional covariance matrices. Our strategy may appear reminiscent of the “granular IV” approach of Gabaix and Koijen (2024), since both exploit idiosyncratic variation in demand shocks for identification. However, our contribution uses the variation *within* each firm’s demand shocks – across its relationships – for identification. Our approach leverages differences in the covariance of effects across relationships, while granular identification exploits the fact that a cleverly constructed weighted average of demand is a valid instrument for aggregate demand (in the price equation) or price (in the demand equation). A key feature of granular identification is also the use of a common price across entities, while we require price variation not only at the agent level, but the relationship level.

More recently, several contemporaneous papers have considered relaxing various assumptions made in the AKM framework. Crippa (2025) modifies the linear additive structure for the interaction function to allow for complementarities. The interaction function (the mapping between effects and outcomes) remains a function of homogeneous unit-level effects, with all heterogeneity coming through non-linearities in the econometrician’s chosen functional form, making his contribution quite distinct from ours. Gutierrez et al. (2025) relax the AKM homogeneity assumption in the context of credit markets by allowing for group-specific heterogeneity in the transmission of the bank’s supply shock, while still assuming homogeneous firm demand (captured by firm-time fixed effects). Their method is an application of the Bonhomme and Manresa’s (2015) grouped fixed effects estimator, where firms’ group membership is latent, so credit growth is decomposed into a firm-time component plus a bank \times firm-group \times time interaction term capturing differential propagation. However, this approach cannot account for the type of demand heterogeneity we describe in Section 2, which is supported by our empirical findings, and a considerable number of relationships is needed for grouped fixed effects estimators to reliably estimate group membership and thus fixed effects. Bruche et al. (2026) modify the Khwaja and Mian (2008) setting to allow for firm preferences over different banks. Such preferences are captured as a scalar index which is the estimated match probability from a multinomial logit model. Unlike our approach, their goal is not to measure heterogeneous demand shocks, and they do not consider heterogeneity in supply; they also must impose a functional form for firm preferences and mainly focus on demand heterogeneity that is correlated with observables. Bergman et al. (2025) make similar observations to ours about the inability of AKM-type estimates to recover true relationship-level effects, based on a structural model of credit demand and supply, in which supply substitution effects may be conflated with demand shocks. They use this model to

propose estimators that can recover the effects of both supply shifters and total supply shocks on firm outcomes, under the assumed structural model. Their contribution is thus limited to the banking setting, and more specifically firms who borrow from two banks, while ours applies to matched data settings more generally. Moreover, their model does not account for heterogeneity in firms' credit demand that may be driven by bank characteristics other than supply shifters, and any heterogeneity in credit supply is driven by the previous period's lending shares. Finally, also specific to the credit market setting, Pietrosanti and Rainone (2026) show how spillovers across relationships can bias AKM-type estimates and propose an "overlapping portfolio instrumental variables" strategy to identify such spillovers, based on a linear model of network interactions, restricted to homogeneous demand and supply innovations. Such spillovers are best viewed as endogenous responses to structural demand and supply shocks, and are accommodated within our framework, once attention is restricted to the intensive margin.

Empirically, this paper is related to the literature recovering or controlling for demand and supply shocks in corporate credit markets that – with the exception of the recent contributions described above – largely leverages AKM. In the most influential paper, 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. Amiti and Weinstein (2018) develop a methodology that builds on AKM to separate firm credit demand shocks from bank supply shocks and quantify their effects on overall investment dynamics. Often, these (estimated) fixed effects are used as regressors. Greenstone et al. (2020) use estimated bank fixed effects from an AKM model in a shift-shares approach to recover bank lending supply shocks at the county level. Estimated firm-time fixed effects are often used as control variable in real effects regressions (see, e.g., Cingano et al. 2016). However, all of these approaches maintain the assumption of homogeneous shocks, implicitly or explicitly making the assumptions underlying a fixed effects model, and in general label them as demand (firm) or supply (bank).

The remainder of the paper is organized as follows. First, Section 2 shows how AKM estimates can be biased in the presence of heterogeneous shocks in a simple example, before illustrating our identification approach. Section 3 describes the general identification argument, provides asymptotic results, and summarizes a Monte Carlo study. Section 4 applies our methodology to Euro-area credit markets to document the bias in AKM-style estimates and provide novel empirical results. Section 5 concludes.

2 Confronting the challenge of heterogeneous shocks

In this section, we present a simple example to illustrate both the limitations of the AKM framework and the intuition behind our identification strategy. In line with our application, we consider the market for corporate credit as an illustration; the same results equally apply in other settings, like those discussed in Section 3.2, after relabelling. Let $u_{fb} = (u_{fb}^d, u_{fb}^s)'$ 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

$$u_{f1}^d = d_f + a, \quad f \in \{1, 2\}, \quad (1)$$

$$u_{f2}^d = d_f - a, \quad f \in \{1, 2\}, \quad (2)$$

$$u_{fb}^s = s_b + e_f \quad f \in \{1, 2\}, b \in \{1, 2\}. \quad (3)$$

There is a common component, $a > 0$, to the demand for credit from bank 1 due to increased export activity, with an equal and opposite decrease in demand from bank 2 (a substitution effect).¹ Each firm also has a firm-specific demand fixed effect across both relationships, d_f . For simplicity, we omit idiosyncratic error terms. On the supply side, each bank has a homogeneous bank-specific supply shock (s_b), but there is also a firm-specific component that is common across banks, due to both banks similarly assessing firms' risk profiles, say. If $a = e_f = 0$, this setting reduces to AKM, where demand shocks have a firm-specific common component, d_f , and supply shocks a bank-specific common component, s_b . Otherwise, the demand shocks also include a bank-specific component and the supply shocks a firm-specific component, due to realistic factors (bank specialization and heterogeneous firm risk). The AKM framework is not equipped to handle these empirically probable features.

To demonstrate this fact, assume for simplicity that the change in loan quantity between f and b is determined by

$$\Delta l_{fb} = u_{fb}^d + u_{fb}^s, \quad (4)$$

so

$$\Delta l_{f1} = d_f + s_1 + e_f + a, \quad (5)$$

$$\Delta l_{f2} = d_f + s_2 + e_f - a. \quad (6)$$

¹The equal and opposite impact of a on u_{f1}^d and u_{f2}^d is for expositional simplicity and can be relaxed without affecting the conclusions below.

Fitting firm and bank fixed effects to Δl_{fb} , following AKM (using firm 2 as the reference group), yields

$$\tilde{d}_1 = (d_1 - d_2) + (e_1 - e_2) \quad (7)$$

$$\tilde{s}_1 = s_1 + a + (d_2 + e_2) \quad (8)$$

$$\tilde{s}_2 = s_2 - a + (d_2 + e_2). \quad (9)$$

The AKM approach aims to identify the firm-specific common components of the demand shocks and the bank-specific common components of the supply shocks but fails to do either. Why? Since firm 2 is the reference unit, \tilde{d}_1 should recover firm 1's demand component, relative to that of firm 2, $d_1 - d_2$. While this term does appear in \tilde{d}_1 , it is contaminated by a second term, $(e_1 - e_2)$, which is firm 1's common component of supply, relative to that of firm 2 (due to risk factors, etc.). Turning to supply, \tilde{s}_1 recovers s_1 , plus two additional terms. The first, a , is bank 1's common component of demand (due to it specializing in finance for export) and is unrelated to supply. The second, $d_2 + e_2$, consists of both the demand and supply components specific to the reference firm, firm 2. \tilde{s}_2 has a parallel structure. Even if we set the components specific to the reference firm to zero, $d_2 = e_2 = 0$, both shocks are still mis-identified, $\tilde{d}_1 = d_1 + e_1$, due to contamination by the firm-specific component of supply, and $\tilde{s}_1 = s_1 + a$, $\tilde{s}_2 = s_2 - a$, due to contamination by the bank-specific component of demand. As this simple example demonstrates, the AKM approach fails whenever there is any correlation in demand from a particular bank across firms or any correlation in supply to a particular firm across banks. Appendix C.1 shows that analogous results hold if we instead use one of the banks as the reference unit. This illustration is closely related to similar points made by Bergman et al. (2025), who argue from a structural model for lending. Their model is specific to the banking setting and highlights how heterogeneity in supply shocks may be reflected in AKM-style demand measures. In contrast, our reduced-form argument easily maps to general applications and highlights potential bias in both demand and supply measures.

The AKM methodology does not account for the possibility that demand shocks may exhibit correlations across banks or that supply shocks may exhibit correlations across firms, in the empirically plausible manner described above. However, rather than complicating identification, we establish below that such correlations – not permitted in the AKM approach – actually facilitate identification. How does our strategy accomplish this? While we present the general case in Section 3, we illustrate the intuition in this stylized example. Let $a, d_1, d_2, s_1, s_2, e_1, e_2$ be mean-zero random variables, with respective variances $\sigma_a^2, \sigma_d^2, \sigma_s^2, \sigma_e^2$. For the purposes of this illustration, additionally assume that $\sigma_s^2 = \sigma_e^2$, so that the bank-

specific and firm-specific components of supply have equal variance. We introduce a second observable, Δr_{fb} , the change in the interest rate charged by bank b to firm f , and suppose that

$$\begin{pmatrix} \Delta r_{fb} \\ \Delta l_{fb} \end{pmatrix} = Au_{fb}, \quad A = \begin{bmatrix} 1 & A_{12} \\ A_{21} & 1 \end{bmatrix}. \quad (10)$$

A encodes the relative responses of the observables to the relationship-level shocks. The unit diagonal is simply a scale normalization.² In the lending example, we typically expect both firm demand and bank supply to increase lending, so $A_{21} > 0$, while demand and supply have opposite effects on the interest rate, so $A_{12} < 0$. Above, we implicitly assumed that $A_{21} = 1$. If we can recover A , then we can identify the shocks by inverting (10). It is possible to identify A_{21} using instrumental variables (IV) (and subsequently recover A_{12}).

To begin, consider four moments:

$$E[\Delta r_{11}\Delta l_{21}] = A_{21}\sigma_a^2 + A_{12}\sigma_s^2 \quad (11)$$

$$E[\Delta r_{11}\Delta l_{12}] = A_{21}(\sigma_d^2 - \sigma_a^2) + A_{12}\sigma_e^2 \quad (12)$$

$$E[\Delta r_{11}\Delta r_{21}] = \sigma_a^2 + A_{12}^2\sigma_s^2 \quad (13)$$

$$E[\Delta r_{11}\Delta r_{12}] = (\sigma_d^2 - \sigma_a^2) + A_{12}^2\sigma_e^2. \quad (14)$$

These moments respectively compute the covariance of interest rate and loan quantity across firms, holding bank fixed; the covariance of interest rate and loan quantity across banks, holding firm fixed; the covariance of interest rates across firms, holding bank fixed, and the covariance of interest rates across banks, holding firm fixed.

Now, construct the ratio

$$\frac{E[\Delta r_{11}\Delta l_{21}] - E[\Delta r_{11}\Delta l_{12}]}{E[\Delta r_{11}\Delta r_{21}] - E[\Delta r_{11}\Delta r_{12}]} = \frac{A_{21}\sigma_a^2 + A_{12}\sigma_s^2 - (A_{21}(\sigma_d^2 - \sigma_a^2) + A_{12}\sigma_e^2)}{\sigma_a^2 + A_{12}^2\sigma_s^2 - ((\sigma_d^2 - \sigma_a^2) + A_{12}^2\sigma_e^2)} \quad (15)$$

$$= \frac{A_{21}(2\sigma_a^2 - \sigma_d^2)}{2\sigma_a^2 - \sigma_d^2} \quad (16)$$

$$= A_{21}. \quad (17)$$

With A identified, it is straightforward to invert (10) and recover the vector u_{fb} . Inspection of the above moments shows that the closed form expression in (15) is equivalent to using

²Without some normalization, it is impossible to disentangle the magnitude of shocks from the scale of their effects.

Δr_{11} as an instrument in the regression of $(\Delta l_{21} - \Delta l_{12})$ on $(\Delta r_{21} - \Delta r_{12})$:

$$\frac{E[\Delta r_{11}(\Delta l_{21} - \Delta l_{12})]}{E[\Delta r_{11}(\Delta r_{21} - \Delta r_{12})]} = \frac{E[\Delta r_{11}\Delta l_{21}] - E[\Delta r_{11}\Delta l_{12}]}{E[\Delta r_{11}\Delta r_{21}] - E[\Delta r_{11}\Delta r_{12}]} \quad (18)$$

Thus, identification in this simple example can be achieved using an IV strategy. The exogeneity condition is

$$E[\Delta r_{11}((s_1 - s_2 + e_2 - e_1) - A_{21}A_{12}(s_1 - s_2 + e_2 - e_1))] \quad (19)$$

$$= E[\Delta r_{11}(1 - A_{21}A_{12})(s_1 - s_2 + e_2 - e_1)] \quad (20)$$

$$= (1 - A_{21}A_{12})E[(a + d_1 - s_1 - e_1)(s_1 - s_2 + e_2 - e_1)] \quad (21)$$

$$= (1 - A_{21}A_{12})(\sigma_s^2 - \sigma_e^2) \quad (22)$$

$$= 0, \quad (23)$$

where the last equality follows from the maintained assumption that $\sigma_s^2 = \sigma_e^2$. Intuitively, differencing – or comparing outcomes across relationships – precisely eliminates any correlation between Δr_{11} and the error term when $\sigma_s^2 = \sigma_e^2$. Simple algebra verifies that the instrument is relevant whenever $2\sigma_a^2 - \sigma_d^2 \neq 0$.

Of course, the assumption that $\sigma_s^2 = \sigma_e^2$ is arbitrary and restrictive; we maintain it here only because it endows our strategy with a simple IV interpretation. Our general identification result 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 general estimator we construct below will consider all possible combinations of firms and banks present in the network to increase precision.

3 Identification and inference

In this section, we present our general identification result and the associated estimator. We interpret the identification assumptions through the lens of two illustrative applications. We next outline results characterizing the estimator's asymptotic distribution and inference. Finally, we summarize the results of an empirically calibrated Monte Carlo study found in Appendix D.

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

3.1 The AKM model

The standard AKM approach takes some observable variable, like the change in loan quantity, Δl_{fb} , and seeks to decompose it into firm fixed effects measuring “demand shocks” and bank fixed effects measuring “supply shocks”. These shocks are assumed to be homogeneous across relationships,

$$\Delta l_{fb}^* = d_f + s_b + \Gamma X_{fb} + \varepsilon_{fb}, \quad f = 1, \dots, F, \quad b = 1, \dots, B, \quad (24)$$

where Δl_{fb}^* is a potential outcome that is realized if firm f has a relationship with bank b , and $\Delta l_{fb} = \Delta l_{fb}^* D_{fb}$ is observed, where D_{fb} is an indicator for whether a lending relationship exists between f and b . d_f is the firm-specific demand component and s_b the bank-specific supply component. Additionally, ε_{fb} is mean-zero and i.i.d. It is generally assumed that $E[\varepsilon_{fb} | D, \mathbf{X}, \mathbf{d}, \mathbf{s}] = 0$, where $D \equiv \{D_{fb}\}$ encodes the entire observed network structure and $\mathbf{d}, \mathbf{s}, \mathbf{X}$ stack the realized firm/bank effects and observables; that is, X , and the relationships encoded by D , are strictly exogenous, allowing the use of standard reduced-form regression techniques. As shown in Abowd et al. (1999), under this exogeneity assumption, Γ can be consistently estimated, and we can focus instead on the simpler model

$$\Delta l_{fb} = d_f + s_b + \varepsilon_{fb}, \quad (25)$$

where it is understood that a) we limit our attention to observed relationships (i.e. $D_{fb} = 1$) and b) any covariates, X , have been partialled out following a Frisch-Waugh argument. We focus on this simpler model. We showed above how the AKM approach can fail to recover the effects of interest under realistic assumptions. The AKM model assumes that d_f , firm f 's demand, is homogeneous across the banks that it borrows from, and s_b , firm b 's supply, is homogeneous across the firms that it lends to. Any correlation of demand across banks or supply across firms causes identification to fail. Put differently, the i.i.d. assumption for ε_{fb} implies that any and all correlation in demand and supply shocks is captured by the fixed effects.

Instead, we relax the homogeneity assumption, replacing it with the alternative assumption that demand and supply shocks are each merely *correlated* across both firms and banks. We exploit the associated covariances for identification.

3.2 A new model for relationship-level effects

We consider a vector of observables, η_{fb} , for instance, the change in the interest rate and loan quantity between firm f and bank b , that are jointly determined by two relationship-level

effects, for instance demand and supply shocks:

$$\eta_{i,fb} = A_{i1}u_{fb}^d + A_{i2}u_{fb}^s, \quad i \in \{1, 2\}.^4 \quad (26)$$

Stacking the equations in (3.2), we consider relationship-level observations, η_{fb} , (the change in interest rate and loan quantity) as a linear combination of relationship-level effects (demand and supply shocks),

$$\eta_{fb} \equiv \begin{pmatrix} \Delta r_{fb} \\ \Delta l_{fb} \end{pmatrix} = A \begin{pmatrix} u_{fb}^d \\ u_{fb}^s \end{pmatrix} = Au_{fb}, \quad (27)$$

where A is left unrestricted.

Assumption 1 gives the key conditions required for identification.

Assumption 1.

1. A is invertible and constant across relationships,
2. $E[u_{fb}|D] = 0$,
3. $E[u_{fb}^d u_{fb'}^s | D] = 0, b' = b$,
4. $E[u_{fb}^d u_{fb}^s | D] = 0, f' = f$.

A encodes the response coefficients of both observables to both effects, or shocks. Second, the effects or shocks are mean zero for observed relationships; this is without loss of generality following demeaning. We take the unconditional orthogonality,

$$E[u_{fb}^d u_{fb'}^s] = 0, b' = b, \quad (28)$$

$$E[u_{fb}^d u_{fb}^s] = 0, f' = f. \quad (29)$$

as definitional for the decomposition of the match outcome into the relationship-level effects. We interpret the meaning and plausibility of *conditional* orthogonality in Assumption 1.3-4, as well as the rest of the assumptions for several examples below, after stating the identification result. Henceforth, we drop the conditioning on D , and all expectations and covariances should be understood to be taken with respect to observed relationships.

⁴Note that this model nests the structure in (25) by letting

$$\begin{aligned} u_{fb}^d &= (d_f + a_{fb}^d)/A_{i1} \\ u_{fb}^s &= (s_b + a_{fb}^s)/A_{i2} \\ \varepsilon_{fb} &= a_{fb}^d/A_{i1} + a_{fb}^s/A_{i2}, \end{aligned}$$

for some uncorrelated a_{fb}^d, a_{fb}^s . The mapping between the two models is not unique, since there are generally infinitely many ways to decompose the i.i.d. ε_{fb} such that the structure in (3.2) holds.

The goal of this exercise is to identify parameters of A , since from (27) knowledge of A , and hence A^{-1} , allows the recovery of u_{fb} from observed data. To that end, we use two sets of covariance equations. The first,

$$\text{cov}(\eta_{fb}, \eta'_{f'b}) \equiv \Sigma_{FF} = A\Lambda_{FF}A', \quad (30)$$

is the covariance of η_{fb} across firms, holding b fixed, where

$$\Lambda_{FF} = \begin{bmatrix} E[u_{fb}^d u_{f'b}^d] & 0 \\ 0 & E[u_{fb}^s u_{f'b}^s] \end{bmatrix}. \quad (31)$$

The second set of equations is

$$\text{cov}(\eta_{fb}, \eta'_{f'b'}) \equiv \Sigma_{BB} = A\Lambda_{BB}A', \quad (32)$$

with Λ_{BB} similarly defined. After imposing a scale normalization for each of the two shocks, (i.e. A has a unit diagonal, $\Lambda_{FF} = I_2$, or $\Lambda_{BB} = I_2$), there are six free parameters among A and $\Lambda_{FF}, \Lambda_{BB}$; let θ vectorize them. Together, these covariances provide six equations,

$$m(\theta) = \begin{bmatrix} \text{vech}(\Sigma_{FF} - A\Lambda_{FF}A') \\ \text{vech}(\Sigma_{BB} - A\Lambda_{BB}A') \end{bmatrix} = 0. \quad (33)$$

Finally, suppose that the covariance of effects across firms (holding bank fixed) is not proportional to that across banks (holding firm fixed):

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

The following result establishes identification:

Proposition 1. *If Assumptions 1 and 2 hold, then θ is the unique solution to (33) up to scale, sign, and column ordering of A .*

This is a well-known result for equations having the structure in (33) (e.g., Rigobon 2003). When the diagonals of Λ_{FF} and Λ_{BB} are non-zero, it follows directly from linear algebra arguments: A is identified in closed form as the eigenvectors of $\Sigma_{FF}\Sigma_{BB}^{-1} = A\Lambda_{FF}\Lambda_{BB}^{-1}A^{-1}$. Intuitively, separately exploiting heterogeneity within a firm but across banks and within a bank but across firms provides two sets of linearly independent moments in terms of demand and supply shocks, which are linked by A . With one covariance matrix alone, A is uniquely determined only up to orthogonal rotations, but with a second, linearly independent matrix, it is unique up to scale and column order. A itself is of interest, containing the responses

of interest rates and loan volumes to demand and supply, for instance, which can be used to construct implied demand and supply curves. From A , u_{fb} is additionally immediately recovered for each relationship by inverting (27).

We now interpret the identifying conditions in Assumption 1 in two settings.

Example 1: Corporate credit. Credit demand and supply shocks are likely heterogeneous due to specialization (e.g., Paravisini et al. 2023), substitution effects, or various incentives to adjust portfolios. Supply shocks are correlated across firms for a given bank due to either an overall change in the bank’s lending or substitution (measured by $E[u_{fb}^s u_{f'b}^s]$), and across banks, for a given firm, due to a common evaluation of the firm’s creditworthiness (measured by $E[u_{fb}^s u_{f'b'}^s]$).⁵ Demand shocks are correlated across banks for a given firm due to either an overall change in a firm’s demand for credit or substitution (measured by $E[u_{fb}^d u_{f'b'}^d]$), and across firms, for a given bank (measured by $E[u_{fb}^d u_{f'b}^d]$), due to bank specialization or the attractiveness of its products, for example. These covariances appear on the diagonals of the within-firm Λ_{FF} and within-bank Λ_{BB} shock covariance matrices, and – in conjunction with A – dictate the reduced-form covariances of prices and quantities, Σ_{FF} and Σ_{BB} .

Assumption 1.1 imposes that the responses of both the interest rate and loan quantity are linear and constant across a chosen sample, within a time period. Note that there is no requirement that this assumption holds intertemporally. It is also possible to focus on a subset of firms – within a particular industry, region, rating, or size class such that the assumption is deemed more plausible. In our application, we focus on the intensive margin, since different coefficients may apply to new relationships and ending relationships compared to adjustments within a relationship.

In most macroeconomic settings, demand and supply shocks are assumed to be orthogonal or independent. Otherwise, they cannot be considered structural shocks, as they would have a proximal cause – the factor driving both – or else still exhibit endogenous responses to each other. This notion of shocks implies orthogonality holds even *within* a relationship, and the unconditional orthogonality across relationships in (29) is a weaker assumption. In order to be structural shocks, u_{fb} should be unpredictable, so X_{fb} will typically contain lagged controls. Appendix C.2 shows that the conditional orthogonality in Assumptions 1.3-4 is implied if u_{fb}^d , $u_{f'b'}^s$ and u_{fb}^d , $u_{f'b}^s$ are mean independent and the information contained in the network structure, $D = \{D_{fb}\}$, satisfies some reduced-form conditions that place conditions on selection into relationships. 1.3 ensures that Λ_{BB} is diagonal and, economically, can be interpreted as requiring a reorientation delay (bank b ’s supply to firm f cannot impact

⁵As a reminder, all expectations are understood to be over observed relationships, or conditional on D .

firm f 's demand from bank b) and that no common factor drives both demand from and supply to firm f . Assumption 1.4 ensures that Λ_{FF} is diagonal and can be interpreted as requiring that firms are atomistic (firm f 's demand cannot impact bank b 's supply to firm f') and that any spillovers are endogenous *responses* to the shocks (since bank b 's supply shock to f' cannot contemporaneously impact firm f 's demand shock). It also typically rules out industry shocks, unless industry indicators or shocks are included in X_{fb} . The above discussion can easily be adapted for alternative settings common in finance with two mutually exclusive types of counterparties, or to international trade, with importers and exporters.

Some researchers may nevertheless view demand and supply shocks within a relationship as potentially correlated, for instance if demand is correlated with firm riskiness (if, say, a firm obtains a new patent, which could drive both demand for credit and make its debt appear less risky to lenders). Across relationships, Assumptions 1.2-4 also mean that our methodology cannot capture aggregate fluctuations (which are removed by demeaning), for instance where firm demand is driven by the evolution of total factor productivity or bank supply is driven by funding conditions. Both of these points highlight possible conceptual limitations of a statistical decomposition like the one we propose, although it retains significant flexibility compared to the existing AKM approach, as demonstrated in Section 2.

Example 2: Labour match quality. In the seminal AKM setting, log-wages are decomposed into worker- and employer-specific effects, reflecting quality or productivity. In our setting, a second observable is required, for instance the match duration, so η_{ij} is a vector containing the log-wage of worker i and employer j and the length of the employment spell, for instance. u_{ij} consists of two unobserved components: the worker match-specific quality, u_{ij}^w , and the employer match-specific quality, u_{ij}^e . Worker match-specific quality is likely heterogeneous: the same worker might be far more productive as a tech executive than as an intern, but in the AKM framework, quality is fixed over the life-cycle (even as workers change industries). Employers may offer on-the-job training, for instance, which will benefit some workers more than others, resulting in heterogeneous employer effects across workers. However, worker effects should be correlated across employers, holding worker fixed, due to innate ability, and across workers, holding employer fixed, since certain features of a company may enable all workers to be more productive. Employer effects should be correlated across workers, holding employer fixed, since firms differ in their productivity, impacting all workers, and across employers, holding worker fixed, due to assortative matching.

Assumption 1.1 requires that the effects of worker and employer match-specific quality on log-wages and employment duration are both linear and constant. This needs to hold

over time (a worker's employment spells occur sequentially over time) and across workers/employers. It is possible to focus on a single sector or workers of a fixed age, education, or experience level, if this assumption is deemed to be too strong in general.

The match-specific effects are definitionally unconditionally orthogonal across relationships, (29). However, only successful matches are observed, so the conditional orthogonality in Assumption 1.3-4 restricts possible matching functions or the stochastic process generating u_{ij} . We provide simple reduced-form sufficient conditions for these to hold in Appendix C.2 and discuss hiring behaviour with which they could be consistent. There are likely many possible combinations of assumptions on the stochastic and matching processes that could satisfy these conditions in practice.

3.3 Estimation

Estimation proceeds with the sample analogues of Σ_{FF}, Σ_{BB} . The natural estimators are simply sample averages across all pairwise combinations of firms borrowing from each bank, and all banks lending to each firm, respectively:

$$S_{FF} = \frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \eta_{fb} \eta'_{f'b} \quad (34)$$

$$S_{BB} = \frac{1}{N_{BB}} \sum_{f=1}^F \sum_{b' \neq b} \eta_{fb} \eta'_{fb'}, \quad (35)$$

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

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

$$q(\boldsymbol{\eta}, \theta) = \begin{pmatrix} \text{vech}(S_{FF} - A\Lambda_{FF}A') \\ \text{vech}(S_{BB} - A\Lambda_{BB}A') \end{pmatrix}, \quad (36)$$

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

$$q(\boldsymbol{\eta}, \hat{\theta}) = 0. \quad (37)$$

$\hat{\theta}$ is available in closed form. In particular, an initial estimate $\tilde{A} = \text{evec}(S_{FF}S_{BB}^{-1})$ is unique

up to scale, sign, and column order, where $\text{vec}(\cdot)$ denotes the matrix of right eigenvectors of its argument.

The column order of this initial estimate must be determined, equivalently the shocks/effects must be labeled (i.e., as demand and supply), and the scale normalized. In our application, we select the column order that minimizes the distance between \tilde{A} and

$$\begin{pmatrix} \text{std}(\Delta r_{fb}) & -\text{std}(\Delta r_{fb}) \\ \text{std}(\Delta l_{fb}) & \text{std}(\Delta l_{fb}) \end{pmatrix} \quad (38)$$

so that the first shock is a demand shock and the second shock is a supply shock. We finally normalize the resulting matrix such that the covariance of both shocks across firms is unity, $\Lambda_{FF} = I_2$. Our estimate, \hat{A} , is the matrix resulting from these transformations; Appendix B provides details. $\hat{\Lambda}_{BB}$ can be immediately recovered, completing $\hat{\theta}$.

3.4 Asymptotic properties

Stronger assumptions on the shocks and network are required to derive the asymptotic distribution of \hat{A} . We present these formally in Appendix A. To summarize, Assumption 3 imposes additional assumptions on the shocks or effects: each shock consists of two independent components, one of which is correlated across firms, and the other of which is correlated across banks. It also requires that the shocks have finite eighth moments and that the covariance matrices of the firm and bank components of the shocks are non-degenerate. Assumption 4 on the network structure guarantees that there is a non-vanishing share of firms (banks) whose number of connections increases proportionally with B (F).⁶ It further requires that neither B nor F dominates the other asymptotically (i.e., $F/B^2 \rightarrow 0$ as $F, B \rightarrow \infty$), but realistically allows them to grow at different rates.

Theorem 1. *Suppose $\theta_0 \in \text{interior}(\Theta)$, which is compact. After suitable normalization and labeling of the columns of A , under Assumptions 1 to 4,*

1. $\hat{\theta} \xrightarrow{p} \theta_0$;
2. $\tilde{\Phi}(\hat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}(0, \mathbf{W})$;
3. $\hat{\mathbf{W}} \xrightarrow{p} \mathbf{W}$;

where

$$\tilde{\Phi} = \left(\begin{bmatrix} \sqrt{B} & 0 \\ 0 & \sqrt{F} \end{bmatrix} \otimes I_3 \right) \hat{\Phi}, \quad \hat{\Phi} = \left. \frac{\partial q(\boldsymbol{\eta}, \theta)}{\partial \theta'} \right|_{\hat{\theta}}, \quad (39)$$

⁶This contrasts with the requirement for consistency of AKM, which Jochmans and Weidner (2019) show necessitates the number of connections to tend to infinity for all firms and banks.

$$\hat{\mathbf{W}}_{FF} = \frac{B^2}{N_{FF}^2} \frac{1}{B} \sum_{b=1}^B \left(\sum_{f' \neq f} \text{vech}(\eta_{fb} \eta_{f'b}') - \text{vech}(S_{FF}) \right) \left(\sum_{f' \neq f} \text{vech}(\eta_{fb} \eta_{f'b}') - \text{vech}(S_{FF}) \right)', \quad (40)$$

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

The first point of Theorem 1 establishes consistency of $\hat{\theta}$, the second asymptotic normality, and the third shows that the asymptotic variances of S_{FF} and S_{BB} can be consistently estimated. More practically, the asymptotic variance of the parameter estimates, $\hat{\theta}$, is given by $\Omega = \tilde{\Phi}^{-1} \mathbf{W} \tilde{\Phi}'^{-1}$. It is noteworthy $\hat{\mathbf{W}}_{FF}$ and $\hat{\mathbf{W}}_{BB}$ are standard clustered variance estimators (where the clusters are banks and firms, respectively).

Testing AKM identification assumptions. A consequence of Theorem 1 is that we can test whether a necessary condition for AKM to be unbiased holds. As illustrated in Section 2, AKM will be biased unless the covariance of demand shocks across firms, holding bank fixed, is zero, and the covariance of supply shocks across banks, holding firm fixed, is zero. These two restrictions are over-identifying in (37), so, under the null hypothesis that the AKM conditions hold,

$$\min_{\theta: \Lambda_{FF} = \text{diag}(0,1), \Lambda_{BB} = \text{diag}(1,0)} \tilde{q}(\boldsymbol{\eta}, \theta)' \hat{\mathbf{W}}^{-1} \tilde{q}(\boldsymbol{\eta}, \theta) \xrightarrow{d} \chi^2_2, \quad (41)$$

where $\tilde{q}(\boldsymbol{\eta}, \theta) = \left(\begin{bmatrix} \sqrt{B} & 0 \\ 0 & \sqrt{F} \end{bmatrix} \otimes I_3 \right) q(\boldsymbol{\eta}, \theta)$. Comparing this statistic to the associated critical values provides a simple test to formally assess key AKM assumptions in various empirical settings.

3.5 Extensions

Including covariates. It is common in models of unobserved heterogeneity to include covariates, as in (24), whose coefficients can either be viewed as nuisance parameters, as is the case in our application, or objects of interest. Moreover, u_{fb} should be orthogonal to any observed lagged regressors in X_{fb} if they are to be interpreted as structural shocks.

A more general version of (27) is then

$$\eta_{fb} = Au_{fb} + \Gamma X_{fb}. \quad (42)$$

In our application, X_{fb} contains a variety of 1-period lagged relationship-level characteristics.

In Appendix C.3, we show that, if u_{fb} is additionally mean independent of X_{fb} , given D , then Γ can be consistently estimated via OLS, maintaining all previous assumptions.

Exploiting multiple time periods. Our identification strategy and asymptotic results all require only a single period of data for applications where agents engage in several relationships simultaneously (when relationships are sequential, as in the labour market example, a timeseries is required). However, it can be helpful to exploit a timeseries of observations to obtain a larger effective sample size, particularly in sparser networks or when changes in outcomes are rare. We demonstrate this in our simulation study, and ultimately pool several quarters of data in our application.

Replacing the identifying moments with expectations taken additionally over time does not impact identification, at the cost of assuming that A remains fixed over the periods considered. Likewise, estimation using the average values of S_{FF} and S_{BB} over time does not impact the asymptotic properties of the estimator. If the shocks/effects are not serially correlated, W_{FF} can be estimated using the clustered variance at bank-time level and W_{BB} at firm-time level. Allowing for serial correlation, variances must be clustered at the bank level and firm level, respectively.

3.6 Monte Carlo Study

Appendix D describes a Monte Carlo study assessing our estimator's performance in detail, empirically calibrated using our estimates for the Italian data from 2022 Q3 to 2023 Q4. We vary the number of firms and banks, setting $F = 1000B$, while holding the network density fixed at the value observed in the data. We summarize the key results here.

For $B = 10$ – far fewer than the 100-200 in the Italian data (and fewer than in any country we study) – the percentage bias for the entries of A can be quite large (see Table D.1). However, for $B = 25$, the bias falls considerably, and by $B = 100$ – still smaller than the number of banks in the data – it is only 1-2%. For $B = 250$, the estimates are unbiased. If, however, we pool four periods of data together, the bias is at most 7% even for $B = 10$, and at most 5% for $B = 25$, illustrating the potential value of pooling quarterly data in our application.

We next compute the empirical size of t -tests with nominal size of 5% for each of the coefficients in A (see Table D.2). For $B = 10, 25$, there are notable distortions for some coefficients, but for $B = 100$, empirical size ranges from 4.8-8.3%. Pooling data again proves helpful, with empirical size for $B = 25$ now controlled between 5.5% and 7.8%; distortions for $B = 10$ remain slightly larger.

Finally, we compare the ability of our methodology to recover the common component of heterogeneous shocks. Our DGP admits a firm component within the simulated demand shock and a bank component within the simulated supply shock. We estimate our relationship-level shocks, and then collapse them at the firm and bank level to obtain proxies for the firm and bank common components of demand and supply, respectively. We compute the average correlations between these and the true firm and bank components of demand and supply, respectively, across samples. We do the same for standard fixed effects estimates. For all specifications, our collapsed firm-level shocks are 7-10 times more strongly correlated with the true homogeneous firm component of the demand shocks than actual fixed effect estimates (see Table D.3). Our collapsed bank-level shocks are better correlated with the true homogeneous bank component of the supply shocks for $B > 25$ using a single time period (and $B \geq 25$ using multiple time periods). These results illustrate both the bias present in fixed effects estimates (since there is also a firm component of supply and bank component of demand in the DGP) and how our estimator uses data from the entire network to estimate the value of A and decompose observables into shocks for each agent. Taken together, these factors mean that – even when the target is a homogeneous agent-level shock – our estimation approach significantly outperforms fixed effect estimates, especially for agents with relatively few connections, like most firms, since those estimates use only data from agents’ own neighbours.

4 Demand and Supply in European Credit Markets

We use our new methodology to estimate demand and supply dynamics in European credit markets using the ECB’s AnaCredit dataset, which contains quantities and interest rates (and many other contract characteristics) on the near-universe of corporate loans in the Euro area.⁷ The data are recorded monthly and at the contract level, but we use quarterly data collapsed at the firm-bank level. We resort to a quarterly frequency since we study changes in loan quantities and interest rates, and there are relatively few such changes at the monthly frequency. In the aggregation of contracts, we focus on the three dominant types of credit used to finance working capital and investments (term loans, credit lines, and revolving credit).

⁷AnaCredit is a harmonized Euro-area loan-level credit registry developed by the Eurosystem that contains the universe of loans to corporations in the Euro area for all firm-bank relationships with an exposure of at least €25,000. Data collection started in September 2018 and records loans/instruments at a monthly frequency. The dataset contains many types of loans, including term loans, financial leases, trade receivables, credit lines, and revolving credit. For a detailed discussion of the AnaCredit dataset, see the recent paper by Kosekova et al. (2025). For more technical documentation of the data included, see the AnaCredit Manual, European Central Bank (2019).

In each quarter, we compute the total committed credit for each firm and bank relationship. We measure the change in loan quantity as the quarter-on-quarter change in this total committed amount. We choose to focus on changes in committed quantities instead of utilized quantities, since the former represent the outcome of interaction between both the firm and the bank, whereas the latter can be viewed as simply demand driven changes, given previously committed quantities. We convert the quarterly change to a quarterly growth rate using the formula

$$\Delta l_{fb,t} = \frac{l_{fb,t} - l_{fb,t-1}}{0.5l_{fb,t} + 0.5l_{fb,t-1}}. \quad (43)$$

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

$$\eta_{fb,t} = \begin{pmatrix} \Delta r_{fb,t} \\ \Delta l_{fb,t} \end{pmatrix}.^8 \quad (44)$$

Given the nature of our identification approach, we focus on the intensive margin (to reduce the role of relationship formation/destruction) and limit our sample to firms that have borrowing relationships with multiple banks. Hence, for a firm f to be included in time t , at least two firm-bank lending relationships must be observed in t and $t - 1$. Although we only retain observations on the intensive margin, the dataset remains unbalanced in nature due to new and terminating lending relationships.

We focus our analysis on 9 Euro-area countries over the sample period, July 2019 to December 2023. These are Austria (AT), Belgium (BE), Germany (DE), Spain (ES), France (FR), Greece (GR), Italy (IT), the Netherlands (NL), and Portugal (PT).⁹ The sample period is determined by the choice to estimate demand and supply dynamics for three 6-quarter windows. We do this for three reasons. First, given we study 9 different countries, presenting quarterly results would become unmanageable, while using a single pooled sample from 2019–2023 would make the assumption of a common A matrix not credible due to dramatic changes in credit markets during this period. The 6-quarter approach strikes a balance between the two. Second, and relatedly, three 6-quarter windows cleanly capture three distinct economic periods for credit markets. The first sample, 2019Q3–2020Q4, spans the lead-up to and pandemic period, which featured severe disruptions and emergency support

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

⁹The set of countries is guided by Kosekova et al. (2025) and insights of the Monte Carlo analysis.

measures. The second sample, 2021Q1–2022Q2, covers the subsequent inflationary episode, up until the point of monetary policy intervention. Finally, 2022Q3–2023Q4 was a period of contractionary monetary policy and credit tightening. It is reasonable to believe that the elasticities in A , *inter alia*, the demand and supply curves of national credit markets, changed from each of these periods to the next along with macroeconomic and financial conditions. Third, while some of the countries we study feature very large numbers of banks and multi-bank firms, smaller countries and those with more consolidated lending relationships benefit from the increased sample size associated with pooling 6 quarters of data.¹⁰ We address occasional data quality issues, for instance inconsistency in the units used to report interest rates, and mitigate the impact of outliers by winsorizing $\Delta l_{fb,t}$ and $\Delta r_{fb,t}$ at the 5% level. Taking this final sample for each country, as a final pre-processing step we residualize with respect to a set of lagged relationship-level characteristics and quarter fixed effects, so that $\eta_{fb,t}$ contains unpredictable non-aggregate variation.¹¹

Table E.4 in the Appendix reports information on the number of firms, bank, and firm-bank matches in each country, in each of the three subsamples described above. The statistics pertain to the samples used in the estimation and thus only consider firms who simultaneously borrow from at least two banks. The table illustrates that, as documented by Kosekova et al. (2025), there are stark differences in the structure of credit markets across the countries studied. Indeed, particularly when comparing to the total number of firms in the AnaCredit dataset in each country, it is clear that multi-bank firms are quite common in some countries, but relatively rare in others, even when comparing countries of similar size (Belgium or Portugal versus the Netherlands). The number of banks also varies starkly, ranging from about 20 in Belgium, Greece, and the Netherlands, to around 800 in Germany. As a result, the effective sample size is quite variable across the countries considered, ranging from about 8,000 relationships per period in the Netherlands to more than 700,000 in Italy. The number of firm-bank pairs does not rank one-to-one with country size.

¹⁰Our results are quite stable with respect to the precise timing of the subsamples. We have experimented with both 4- and 8-quarter windows loosely aligning with the same major economic events as well as rolling windows, and the magnitudes of \hat{A} remain largely unchanged.

¹¹We residualize using four variables at the firm-bank-time level and two at the bank-sector time level. These six variables net out systematic variation in loan quantities and prices related to contract pricing rigidity, risk mitigation, loan type composition, relationship importance and the bank’s sectoral positioning vis-à-vis the firm. Specifically, they are: the share of borrowing by firm f from bank b at time $t - 1$ that is (1) fixed rate, (2) secured by collateral, or (3) credit line or term loan borrowing (with the remaining part being mainly revolving credit). The fourth characteristic is the share of firm f ’s borrowing from bank b at time $t - 1$ in firm f ’s overall borrowing at time $t - 1$. The bank-sector-time level variables are a bank’s market share in the sector of firm f (the share of bank b ’s lending to the sector of firm f in total lending by all banks to that sector), and a bank’s exposure to the sector of firm f (the share of bank b ’s lending to the sector of firm f in total lending by bank b).

4.1 Credit demand and supply patterns over countries and time

We begin by reporting estimates of A , which encode the responses of price and quantity to demand and supply shocks over time. Figure 1 plots the price and quantity responses to demand and supply shocks across the three periods for each country, with 95% confidence intervals. The coefficients are scaled to represent the response to a unit standard deviation shock to render them comparable across countries and time; price responses are in interest rate basis points, and quantity responses are growth rates.

First, these results provide an important validation of our methodology. For every one of the 27 country-periods, there is one shock that moves price and quantity in the same direction and is thus consistent with a demand shock, and one shock that moves price and quantity in opposite directions and is thus consistent with a supply shock, up to statistical significance.¹² To emphasize, this is an empirical finding and is not guaranteed by the methodology.

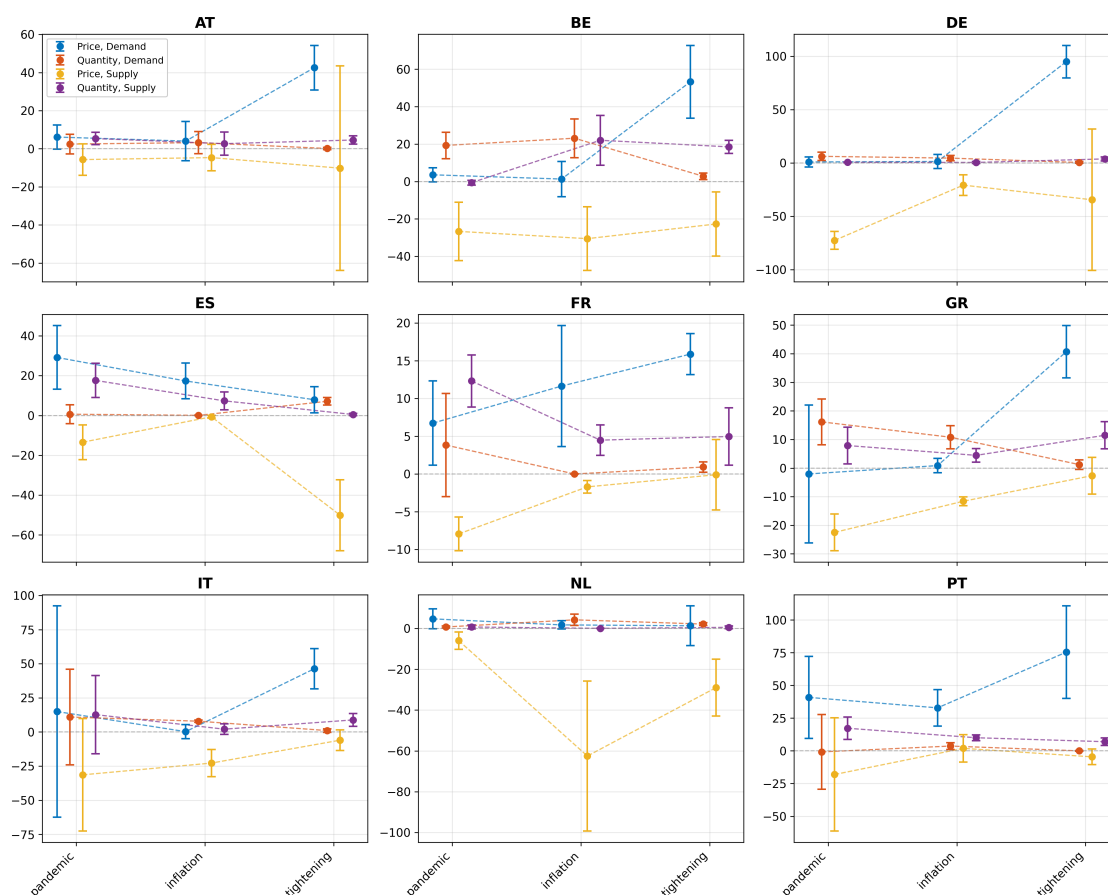
Second, while the coefficients certainly evolve over time for many countries, these changes are generally smaller in magnitude than the differences that exist between countries. There are some exceptions; for instance, the price response to demand shocks increases substantially in Austria, Belgium, Germany, Greece, Italy, and Portugal during the period of monetary contraction. Over the three periods considered, there has also been a commensurate reduction in the responsiveness of prices with respect to supply in many countries (Germany, France, Greece, Italy and Portugal), suggesting that they have become more demand-driven. The responsiveness of quantities to demand has fallen for several countries (Belgium, France, Greece, and Italy), while for others the responsiveness to supply has fallen (Spain, France, Italy and Portugal).

The coefficients in A directly imply the slopes of credit demand and supply curves for each country-period; this may be a more intuitive way to view the results. Figure 2 summarizes how these curves vary over country and time. For convenience, instead of plotting the slopes of the curves (since near-vertical curves have very large slopes), this figure depicts the angle that each curve makes with the x -axis; Figure E.1 in the Appendix plots the curves themselves separately for each country-period. The pandemic period is plotted with blue diamonds, the inflationary period with orange circles, and the monetary contraction with green squares.

During the pandemic (blue diamonds), four countries (Austria, France, Greece, Italy) exhibit downward (upward) sloping demand (supply curves), but for half of the countries, at least one curve is essentially vertical (inelastic). For Belgium, Germany, and, to a lesser extent, the Netherlands, demand is inelastic. This is consistent with firms demanding the

¹²That is, while there are 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.

funding required to survive such a major economic shock at whatever price it was available. In Spain and Portugal, and again, to a lesser extent, the Netherlands, the supply curve is inelastic.

During the post-pandemic inflationary episode (orange circles), we observe stable slopes (compared to the pandemic period) in Austria, Germany and Greece. In most other countries, at least one of the slopes changes substantially and, as during the pandemic, inelastic curves are again present in most countries. Germany, the Netherlands, and Italy exhibit inelastic demand. This is consistent with firms demanding credit to cover rising costs due to inflation at whatever price is available. On the other hand, Spain, France, and to a lesser

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 structural 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.

extent Portugal exhibit inelastic supply. This is consistent with banks being wary of issuing term loans whose repayment may be depreciated by inflation. It is worthwhile to note that while Portugal appears to have an (marginally) upward-sloping demand curve, this slope is very noisily estimated and not statistically significant.

During the monetary policy tightening (green squares), all countries exhibit at least one inelastic curve, predominantly an inelastic supply. This is consistent with banks constraining credit in response to tighter conditions fostered by monetary policy, with interest rates serving to ration credit. Only in Spain and the Netherlands is the supply curve not nearly

perfectly inelastic. In these countries, the demand curve is inelastic (and nearly inelastic in Germany). In all other countries, the demand curve ranges from perfectly elastic (France and Greece) to a standard downward slope.

Overall, these results demonstrate both the heterogeneity of credit markets across countries and time, as well as remarkably similar patterns. While monetary policy was the same across the Euro area, credit markets behaved differently. Such results can be informative for evaluating the effects of potentially heterogeneous local policies, as well as motivating such policies, given the variation in credit dynamics. Of course, understanding the determinants of the differences in the responsiveness of price and quantity to demand and supply across countries and time is an important empirical direction, but one that is outside the scope of the present paper.

4.2 Assessing the AKM assumptions

Before employing our recovered demand and supply innovations in regression analysis, we assess whether previous strategies are compatible with the data. In Section 3.4, we propose a simple test of the AKM assumptions, exploiting over-identifying restrictions. We apply this test in each country-period. In 25 out of 27 country-periods, the data reject the null hypothesis of zero correlation between supply shocks across banks and demand shocks across firms at the 5% level, which, as explained in Section 2, is required to hold for AKM to identify demand and supply innovations.¹³ This is our first evidence that an AKM-based strategy is unlikely to provide reliable estimates in our empirical setting.

Table 1 reveals the extent of the heterogeneity masked by the AKM homogeneity assumption. For each firm-period, we compute the average demand innovation across borrowing relationships and the standard deviation of its demand innovations. The upper panel reports the distribution of these collapsed measures across firm-times. Similarly, the lower panel computes averages and standard deviations of supply innovations for each bank-period. The within-unit variation is at least comparable to the between-unit variation. The median within-firm standard deviation of demand shocks is 0.225, representing 35% of the between-firm standard deviation (0.646). At the 75th percentile, within-firm variation (0.863) exceeds between-firm variation. For supply shocks, the pattern is even more pronounced: the median within-bank standard deviation (0.712) is 78% larger than the between-bank standard deviation (0.399).

These patterns directly contradict the AKM homogeneity assumption, which posits that all meaningful variation occurs between firms (for demand) or between banks (for supply).

¹³The 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.

Our results show the opposite: for most firms and banks, the heterogeneity across their relationships exceeds the average differences across firms or banks. This implies that AKM fixed effects – which by construction capture only between-unit variation – discard rather than isolate the economically relevant signal.

4.3 Evidence of bias in AKM measures

In this section, we present evidence that AKM-style estimates of demand and supply shocks exhibit substantial bias. We do so in a simple setting where economic theory offers unambiguous predictions: demand shocks should raise prices (positive coefficient) while supply shocks should lower prices (negative coefficient). In particular, we show that firm-level “demand” estimates from an AKM model appear to capture supply effects.

Table 2 regresses the relationship-level interest rate change (specifically, the residual after regressing on relationship-level controls) on various demand and supply measures, controlling for country-time fixed effects. Note that these regressions pool observations across all countries and time periods. Unless otherwise noted, all standard errors are two-way clustered at firm and bank level, reflecting standard practice in the literature. Column (1) presents results using our demand and supply innovations identified via covariance restrictions. The signs align perfectly with economic theory: a one standard deviation increase in demand innovations raises interest rates by 21.9 basis points, while a one standard deviation supply

innovation reduces rates by 18.7 basis points. Both effects are highly statistically significant and economically meaningful. Moreover, these innovations explain 48% of the variation in relationship-level interest rate changes, indicating that they capture genuine price-relevant variation in credit market conditions.¹⁴

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*** (0.008)			0.261*** (0.012)
Supply innovation	-0.187*** (0.007)		-0.259*** (0.009)	
AKM Firm-Time FE		-0.483*** (0.054)	1.151*** (0.084)	
AKM Bank-Time FE		-0.751*** (0.096)		-1.260*** (0.104)
AKM Residual		-0.470*** (0.054)	1.150*** (0.082)	-1.549*** (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.

Column (2) reports results using AKM estimates (firm-time and bank-time fixed effects) as well as the residual in the AKM regression as regressors.¹⁵ The firm-time effect, often

¹⁴The only reason this number is not 100% is variation in A across country-periods. If we were to conduct the regression separately for each country-period, it would be 100% by construction.

¹⁵Following Khwaja and Mian (2008) and Amiti and Weinstein (2018), we obtain AKM estimates from the following specification: $\Delta l_{fb,t} = d_{f,t} + s_{b,t} + \Gamma \cdot X_{fb,t} + \varepsilon_{fb,t}$, where $\Delta l_{fb,t}$ is the quarterly change in credit from bank b to firm f in period t , $d_{f,t}$ are firm-time fixed effects that absorb all firm-level credit demand shocks, $s_{b,t}$ are bank-time fixed effects that capture bank-specific credit supply shocks, $X_{fb,t}$ are relationship-level controls, and $\varepsilon_{fb,t}$ is the error term. To mimic our setup as closely as possible, we partial out the relationship-level controls, $X_{fb,t}$, in a first stage regression estimated by country-six quarter period.

interpreted as a demand shifter in the lending literature, negatively impacts interest rate changes (-0.483), inconsistent with the expected sign for an exogenous demand shift. The bank-time effect has the expected negative sign, but the overall fit is low (Adj. $R^2 = 0.01$). Interestingly, the residual from the AKM regression also has a highly significant negative association with the interest rate, suggesting that considerable heterogeneity relevant to credit pricing decisions is left on the table by this methodology. To isolate whether the bias originates in the firm-time or bank-time fixed effects (or both), we conduct a sequential substitution test in columns (3) and (4). In column (3), conditioning on our identified supply innovation shifts the loading of the AKM firm-time component to a positive sign (1.151), consistent with it containing demand-related variation that is obscured by supply-related variation for which our identified supply innovation is an effective control. In column (4), conditioning on our demand innovation steepens the loading of the AKM bank-time component (-1.260 versus -0.751 in column (2)), consistent with attenuation from demand contamination. Together, these patterns suggest that AKM firm- and bank-time objects are mixtures of both forces in our setting, rather than clean demand/supply shifters, with implications not only for pricing but also likely for more sophisticated regressions of interest to empirical researchers.

Table E.5 in Appendix E repeats the analysis after collapsing at the firm-time level and bank-time level, which are the units of observation more typically used for AKM-style regressions. The findings are essentially unchanged. These results are entirely consistent with the findings of the previous section, which reject the AKM assumption of no correlation of the demand shock to a given bank across firms, as well as the illustrative two firms-two banks example presented in Section 2.

An alternative approach to assess and quantify the scope for bias is regressing the firm-time (bank-time) fixed effects on our demand and supply innovations (Table E.6 in Appendix E). Supply innovations have a 60% larger association with the firm-time fixed effects than do demand innovations. Supply innovations also explain a much larger fraction of the variation in the AKM firm-time components (38% of the variance in firm-time effects is explained by our supply innovations versus only 13% by demand innovations), directly contradicting the interpretation of firm-time effects as demand measures. Similarly, bank-time effects load 20% on demand versus 22% on supply, with comparable coefficients. These results provide direct empirical confirmation of the theoretical bias illustrated in Section 2, where we show how specialization and firm-level heterogeneity cause firm fixed effects to capture supply variation (through bank specialization) and bank fixed effects to be contaminated by firm-specific risk factors.

Our findings also accord with those in contemporary work 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 Effects of AKM bias in Firm Outcome Regressions

While the simple price regressions above offer insights into the contamination of AKM-style estimates, they are not aligned with the type of regressions typically estimated in empirical practice, where firm-level real outcomes are generally the objects of interest (e.g., Chodorow-Reich 2014; Jiménez et al. 2020; Khwaja and Mian 2008). In this section, we demonstrate the effects of the previously documented contamination on such regressions.

Like these studies, we proceed in two steps. The primary step assesses the impact of an exogenous shock on credit growth. Our quasi-natural experiment is inspired by Core et al. (2025). We use the large and sudden monetary policy tightening that began in July 2022, with the ECB raising rates by 450 basis points over the subsequent fourteen months. This rise was largely unanticipated in its timing and magnitude. Consequently, firms’ pre-determined exposure to interest rate fluctuations—through their choice of floating versus fixed rate contracts and debt maturity structure—can be treated as exogenous with respect to the subsequent tightening episode. In this context, our aim is thus to identify the differential change in credit growth between firms that borrow at a floating rate versus those that borrow at a fixed rate, while controlling for credit demand. For each quarter, h , during 2022Q3-2023Q4, we estimate the following cross-sectional regression:

$$\ln(Q_{fb,2022Q2+h}) - \ln(Q_{fb,2022Q2}) = \iota_b^h + z_{c,s}^h + \pi^h \mathbf{1}_{f,pre}[exposed] + \sum_{j=1}^h \vartheta_j^h \cdot demand_{fb,2022Q2+j} + \epsilon_{fb}^h \quad (45)$$

When we control for demand effects, $demand_{fb,t}$, we expect the coefficient π^h to measure the supply effect on loan quantity. ι_b^h captures the bank fixed effect, absorbing heterogeneous exposure to monetary tightening across banks due to their deposit franchise (e.g., Drechsler et al. 2017), interest rate risk exposure (e.g., Gomez et al. 2021), or security losses (e.g., Greenwald et al. 2024) that do not vary at the relationship level. $z_{c,s}^h$ are country-sector fixed effects, accounting for conditions in a firm’s sector unrelated to the nature of its credit arrangements. This regression considers firms with differential “exposure” to contractionary monetary policy (based on whether the firm’s borrowing is at a variable rate or needs to be rolled over), much like exposure-based designs such as Khwaja and Mian (2008) or Jiménez et al. (2014). These regressions are cross-sectional, but the expanding horizon spanning 6 quarters from 2022Q3 to 2023Q4 uncovers the cumulative supply impact during a monetary

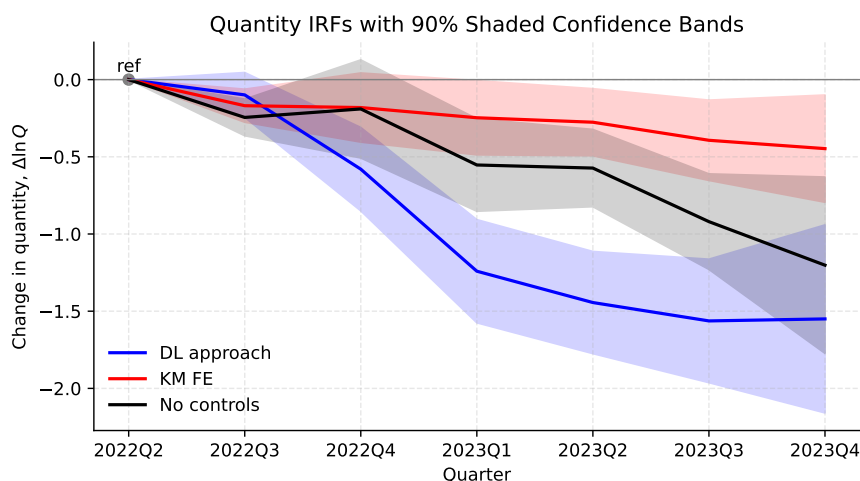


Figure 3: Dynamic Effect of Tightening on Quantities

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.

tightening episode.

Figure 3 plots the results. First, the black line (with shaded 90% interval) shows that, without controlling for demand effects (so $demand_{fb,t}$ is omitted), firms exposed to rising interest rates exhibit significantly lower credit growth from 2022Q2 onwards.¹⁶ But to what extent is this driven by the supply channel versus the demand channel? The blue line plots the effect controlling for our relationship-level demand innovations (u_{fb}^d). The response is significantly more negative, more than twice as large in magnitude from 2022Q4 to 2023Q2.¹⁷ This suggests that an *offsetting* positive demand channel was at play, with firms borrowing at a floating rate demanding more credit. These findings are consistent with banks viewing firms with lower exposure to risks from rising interest rates as more attractive borrowers, while firms with greater exposure seek to secure credit before rates rise higher. The red line, estimated using firm-time effects to control for demand, as in Khwaja and Mian (2008) (“KM”) (so $demand_{fb,t}$ is $d_{f,t}$), offers a stark contrast. Instead of a more pronounced effect, it shows considerable *attenuation* relative to the black line (total effect), with a barely

¹⁶Standard errors are double clustered at the firm and bank levels.

¹⁷Our demand innovations are relationship-specific. Using the firm-level average demand innovation yields nearly identical results.

significant effect at most horizons. In line with the results in Section 4.3, this indicates that firm-time effects are contaminated by supply variation, meaning that controlling for them instead of our identified demand innovations washes out the effect of the supply channel we attempt to capture. Empirically, these results demonstrate that the contraction of credit supply to interest rate risk-exposed firms led to a significant reduction in their borrowing from 2022 onward.

Following the model of papers like Khwaja and Mian (2008), Chodorow-Reich (2014), Cingano et al. (2016), or Jiménez et al. (2020), we now document how this reduction in borrowing transmits to real firm outcomes. We consider regressions of the form

$$Y_{f,2021Q4+h} - Y_{f,2021Q4} = \tilde{i}_{b^*(f)}^h + \tilde{z}_{c,s}^h + \tilde{\pi}^h \mathbf{1}_{pre}[exposed] + \sum_{j=1}^h \tilde{\vartheta}_j^h \cdot demand_{f,2021Q4+j} + \tilde{\epsilon}_f^h, \quad h \in \{4, 8\} \quad (46)$$

where $Y_{f,t}$ is a firm outcome. Firm outcomes are observed annually, so we consider the one-year and two-year impact (i.e. $h \in \{4, 8\}$) relative to the latest pre-tightening recorded value (2021Q4). $b^*(f)$ is firm f 's main lender (thus, $\tilde{i}_{b^*(f)}^h$ is the main lending bank's fixed effect), and $\tilde{z}_{c,s}^h$ is a country-sector fixed effect. Credit demand is a set of 4 or 8 quarterly demand controls, depending on whether we assess the one-year or two-year impact. $\tilde{\pi}^h$ is the coefficient of interest. It captures how real economic outcomes during a monetary tightening differ for ex-ante more exposed firms, all else equal.

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 first three columns report results for 2022, and the second three for 2023, thus measuring effects at different horizons. Columns (1) and (4) report results without controlling for demand, (2) and (5) control for the firm-time fixed effect, as in Khwaja and Mian (2008), and (3) and (6) control for our demand innovation, aggregated at the firm level. In each regression, the coefficient of interest is that on an indicator for whether the firm borrows at a floating interest rate or via short-term loans (and is thus exposed to rising rates). For total assets, the effects in 2022 are negative and highly significant across all columns, suggesting that the effect was largely supply-driven. In 2023, however, the effect controlling for our demand measure is nearly 1.5 times as large as that controlling for the AKM measure. For both periods, the results controlling for the AKM measure show an attenuated effect relative to not controlling for any demand measure at all, again indicating that the AKM demand measure is contaminated by supply effects. The results for turnover are similar, with our demand measure delivering slightly stronger effects than the specification with no demand

control, and the AKM-based estimate appearing attenuated. For the credit to total assets ratio, the 2022 effects are smaller in magnitude. They remain significant (at the 10% level) without a demand control and using our measure, but lose significance when the KM demand measure is included. The former coefficients are more than twice as large in magnitude as the latter. For 2023, all effects are highly significant, but the coefficient using our measure is again more than 1.5 times that using the AKM measure.

In sum, we find strong evidence of heterogeneity in the credit supply effect of the 2022 monetary tightening on firms. To do so, we controlled for firms' credit demand, which Table 1 shows can be far from homogeneous across relationships. Starting from 2022Q2, we document a significant contraction in credit supply to firms that borrow at a floating rate and are thus exposed to interest rate risk, relative to those firms that borrow at a fixed rate. The size of the gap increases through the second half of 2023, when it flattens out, along with interest rates. This contraction in credit supply transmits starkly to firm outcomes: firms borrowing at floating rates or relying on short-term credit experience significantly lower total assets, turnover, and credit to total assets ratios than those borrowing at fixed rates or long-term.

Before moving to a final application, we take stock of the findings. The preceding results have highlighted the risks of adopting AKM-type methodologies in data where the identifying assumptions are unlikely to hold. We documented formal statistical rejections of these assumptions for 25 of 27 country-periods and reported evidence of considerable heterogeneity in recovered demand and supply innovations, which is ignored by such methodologies. We demonstrated the consequences of these assumptions failing in simple regressions of interest rates on AKM-style demand (firm-time effects) and supply (bank-time effects) measures. Both demand and supply measures were associated with lower interest rates, indicating a mix of the true underlying demand and supply shocks. However, the ultimate object of interest for researchers in this empirical literature tends to be the impact of credit conditions on firm outcomes. Therefore, we have importantly shown how relying on firm-time effects to control for demand – as in Khwaja and Mian (2008) – leads to a significant underestimate of the differential credit supply effects of borrowing floating versus fixed rate following the start of monetary contraction in 2022. These contaminated demand controls similarly lead to attenuated estimates of the effect of the associated contraction in credit supply on firms' total assets, turnover, and credit to assets ratio during this episode, leading to quantitatively (and at times qualitatively) different conclusions about the real effects.

Table 3: Real Effects

	(1)	(2)	(3)	(4)	(5)	(6)
	ln(Total Assets)					
Floating or short-term borrower	-0.695*** (0.182)	-0.611*** (0.171)	-0.726*** (0.168)	-1.065*** (0.245)	-0.795*** (0.225)	-1.163*** (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*** (0.226)	-0.664*** (0.221)	-0.758*** (0.212)	-1.272*** (0.306)	-1.087*** (0.303)	-1.359*** (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* (0.150)	-0.129 (0.146)	-0.322* (0.167)	-1.258*** (0.349)	-0.723*** (0.269)	-1.280*** (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 (46), 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.

4.5 Heterogeneous Transmission of Monetary Policy

Finally, we exploit the *heterogeneity* in our demand and supply innovations to present novel results on the transmission of monetary policy. While these innovations, as relationship-level shock measures, should not be associated with the level effect of aggregate monetary policy shocks, (and indeed, this is guaranteed by demeaning price and quantity data by quarter), they may still be associated with the heterogeneous effects of monetary policy across relationships. That is, our demand and supply innovations may reflect the heterogeneous transmission of monetary policy. To study this transmission, we estimate regressions of the form

$$u_{fb,t}^i = \varphi^i \cdot share_{fb,t-1} + \pi_{mp}^i \cdot mp_t \times share_{fb,t-1} + \pi_{cbi}^i \cdot cbi_t \times share_{fb,t-1} + \text{fixed effects} \dots \quad (47)$$

The key coefficients in these regressions are the interaction coefficients, π_{mp}^i and π_{cbi}^i (where $i \in \{demand, supply\}$), which measure how demand and supply innovations covary with relationship-level observables following monetary policy shocks. mp_t and cbi_t refer to monetary policy shocks and central bank information shocks from Jarociński and Karadi (2020). We include and interact observables ($share_{fb,t-1}$) from the prior quarter for each relationship-level covariate to use pre-determined values. In the first three columns of Table 4, the dependent variable is the demand innovation, and in the second three, it is the supply innovation. We consider, in turn, setting $share_{fb,t}$ to the share of fixed rate loans between f and b , or the share of collateralized loans between the pair, and finally a specification including both. The former variable relates to the literature on the floating rate channel of monetary policy (Gürkaynak et al. 2022; Ippolito et al. 2018), whereas the latter links to the role of loan types in the bank lending channel (see, e.g., Ivashina et al. 2022). We find that, for a contractionary monetary policy shock, a higher share of fixed rate loans is associated with a significantly lower demand for credit. This is consistent with firms that have already secured credit at a fixed, lower rate not wishing to initiate new borrowing at the new, higher interest rate. For a central bank information shock signaling good news about the economy, firms with a higher share of fixed rate loans demand more credit in anticipation of higher interest rates in the future. Conversely, whether a firm's loans are collateralized does not explain variation in demand following either type of shock. Turning to supply, we find that, for a contractionary monetary policy shock, a larger share of fixed rate loans is associated with a significantly higher supply of credit. Such firms are less exposed to interest rate risk in the face of rising rates, and are thus perceived as less risky borrowers and more attractive to lenders. For a central bank information shock signaling good news about the economy, firms with a higher share of fixed rate loans receive lower credit supply, as banks are more eager

to lend at floating rates in anticipation of interest rates rising in the near future. Having a higher share of collateralized loans is associated with a higher supply of credit following a contractionary monetary policy shock. This is consistent with banks prioritizing asset-based rather than cash-flow-based borrowers when credit conditions tighten.

These estimates provide novel insights into the transmission of monetary policy. Both contractionary policy and central bank information are associated with starkly different demand from firms that are exposed (by borrowing at a floating rate) compared to those who have secured credit at the prevailing lower rate. Banks' supply decisions differ markedly along the same margins, and they may also seek out borrowers who are able to provide collateral. More precisely, these findings connect to and extend the following strands of the monetary transmission literature. First, our demand-side results complement the floating rate channel of Ippolito et al. (2018), who document that firms with more unhedged bank debt display stronger sensitivity to monetary policy. We extend their analysis by showing that this channel operates at the relationship level and manifests asymmetrically across demand and supply: while firms with floating-rate exposure reduce demand following tightening, consistent with anticipatory behavior, banks simultaneously increase relative supply to fixed-rate borrowers, viewing them as less risky. Second, our supply-side results on collateral are consistent with models emphasizing the collateral channel (e.g., Chaney et al. 2012; Kiyotaki and Moore 1997), showing that banks prioritize asset-based lending when conditions tighten. Third, the heterogeneous responses we document have implications for the literature on firm heterogeneity in monetary transmission (e.g., Gertler and Gilchrist 1994; Ottonello and Winberry 2020). Unlike these studies, which focus on firm-level characteristics, we show that heterogeneity exists within firms across their credit relationships – a dimension that existing models do not capture.

From a policy perspective, these results highlight the distributional consequences of monetary policy. The incidence of tightening depends on the pre-existing structure of credit contracts: firms with floating-rate or uncollateralized borrowing face not only higher debt servicing costs but also relatively lower credit supply, as banks endogenously reallocate toward safer exposures. 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. Moreover, the distribution of rate exposure is itself shaped by macroprudential policy: collateral requirements, interest rate risk regulation, and borrower-based measures all influence the prevalence of fixed versus floating-rate contracts and secured versus unsecured lending. Our findings thus point to an underappreciated interaction between macroprudential and monetary policy—choices made during benign conditions about credit contract structure determine the distributional incidence of

monetary tightening when it arrives.

Insights of the type reported in the first six columns of Table 4 could not be recovered without our methodology. That is because we assess the determinants of demand and supply at the relationship level, and previous measures of demand and supply exist only at the firm or bank level. To illustrate the ways in which our methodology extends the existing literature, the final two columns of Table 4 report two additional regressions. Column (7) considers credit growth as the dependent variable, with the same right hand side as columns (3) and (6). This parallels the frontier for relationship-level regressions in the existing literature. Regressing quantity alone confounds both demand and supply effects, and as a result, there is only one significant interaction, monetary policy with the share of collateralized loans. However, because credit volume is simultaneously determined by demand and supply, there is no way to know whether this represents a supply or demand channel. Column (8) reflects the first contribution of this paper relative to the credit supply literature: the idea that prices ought be studied alongside quantities in order to disentangle demand and supply effects. In a regression of the change in interest rates on the same right hand side, all coefficients are highly significant. Combined with the estimates in column (7), this provides the econometrician with important information. For instance, focusing on the interaction of monetary policy with the share of collateralized loans, the positive coefficient for quantity and the negative coefficient for price indicates that there must be an effect operating through the supply channel. However, it says nothing of the magnitude of the effect, since this pattern could be consistent with an even larger supply effect that is offset by a smaller demand effect. Similar interpretations apply to the other pairs of coefficients. In contrast, columns (3) and (6), using our relationship-level innovations as dependent variables, highlight the value of our novel identification strategy. Only in these regressions, with identified demand and supply measures, are we able to actually quantify both the relationship-level demand and supply effects underlying the previous quantity and price coefficients. Focusing, for instance, on the interaction of monetary policy with the share of fixed rate loans, we now learn that a null effect on quantity and a significant negative effect on price mask both a large negative demand effect and a large positive supply effect.

5 Conclusion

This paper proposes a new methodology to study relationship-level effects or shocks, addressing multiple limitations of the seminal AKM framework. The responses of observables to those effects or shocks are identified without distributional assumptions, exploiting the structure of the covariances of observables across an agent's relationships. The associated

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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Demand innovation			Supply innovation			Credit growth	Change in Interest Rate
Share of fixed rate loans (0.014 (0.012)		0.014 (0.012)	-0.018*** (0.007)		-0.018*** (0.007)	-0.004 (0.003)	0.042*** (0.012)
MP × Share of fixed rate loans	-0.537*** (0.152)		-0.538*** (0.153)	0.674*** (0.155)		0.692*** (0.153)	0.086 (0.057)	-1.091*** (0.143)
CBI × Share of fixed rate loans	0.961*** (0.176)		0.960*** (0.175)	-0.385*** (0.118)		-0.392*** (0.118)	-0.084 (0.053)	1.082*** (0.159)
Share of collateralized loans		0.015*** (0.006)	0.015*** (0.005)		-0.009 (0.008)	-0.009 (0.009)	0.006 (0.010)	0.025*** (0.005)
MP × Share of collateralized loans		-0.026 (0.071)	-0.064 (0.075)		0.395*** (0.067)	0.431*** (0.068)	0.131** (0.066)	-0.324*** (0.078)
CBI × Share of collateralized loans		0.068 (0.092)	0.087 (0.083)		-0.012 (0.073)	-0.028 (0.070)	0.048 (0.063)	0.271*** (0.088)
Adjusted R^2	0.08	0.08	0.08	0.04	0.04	0.04	0.00	0.13

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.

closed-form estimator is simple, as is inference.

Empirically, we present a wealth of novel results for European credit markets. The slopes of credit demand and supply curves have been relatively stable within countries over recent years, while varying dramatically across countries. Importantly, we reject key AKM identifying assumptions for almost all country-periods. There is considerable heterogeneity in both demand and supply shocks that previous methodologies cannot recover. We document substantial bias in AKM-style demand and supply estimates, most clearly manifest in a positive demand effect being associated with a reduction in interest rates. Using AKM-style estimates to control for demand accordingly attenuates the effect of the post-2022 monetary policy contraction on the credit supply to exposed firms essentially to zero, while our demand measure leads to a significant and large effect over the following six quarters. This differential contraction in credit manifests in significantly worse real outcomes for exposed firms, which would again be underestimated by relying on AKM-style estimates. Finally, there is considerable heterogeneity in the distributions of demand and supply shocks following monetary policy announcements, which is associated with key firm characteristics. These findings can be helpful to discipline theoretical models, suggest identifying assumptions in less data-rich settings, and inform policy debates.

On the theoretical side, future work could consider the assumptions required to extend the methodology to study agents with a single relationship (i.e. single-bank borrowers) or the extensive margin. It may also be possible to add further observables, for instance loan maturity, to generate over-identifying restrictions (and thus make model assumptions testable) or to impose structure on time variation in the coefficient matrix when multiple periods of data are available. Extending the methodology to capture aggregate (time series) variation (as in Amiti and Weinstein 2018) is appealing, but appears challenging since any interesting aggregate impulses would likely appear as proximal causes for both supply and demand.

While we apply our approach to identify demand and supply shocks for corporate credit, myriad other settings abound. Application to a wealth of other financial settings is straightforward, as is that to various trade and production network settings. The classic AKM setting of worker- and employer-effects, now possible at the relationship level, appears fruitful, as do myriad applications in industrial organization using rich consumer data.

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>
- 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), 1–60. <https://www.jstor.org/stable/26372543>
- 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>
- 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
- 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>
- Gottfries, A., & Jarosch, G. (2023). *Dynamic monopsony with large firms and noncompetes* (Working Paper No. 31965). National Bureau of Economic Research. <https://doi.org/10.3386/w31965>
- 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>
- Jiménez, G., Ongena, S., Peydró, J.-L., & Saurina, J. (2014). Hazardous times for monetary policy: What do twenty-three million bank loans say about the effects of monetary policy on credit risk-taking? *Econometrica*, 82(2), 463–505. <https://doi.org/https://doi.org/10.3982/ECTA10104>
- 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>
- Kaas, L., & Kircher, P. (2015). Efficient firm dynamics in a frictional labor market. *American Economic Review*, 105(10), 3030–60. <https://doi.org/10.1257/aer.20131702>
- 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>
- 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](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*

$$u_{fb}^d = e_{fb}^d + v_{fb}^d \quad (48)$$

$$u_{fb}^s = e_{fb}^s + v_{fb}^s. \quad (49)$$

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

$$\lim_{F, B \rightarrow \infty} \frac{B}{N_{FF}^2} \sum_{b=1}^B \text{var} \left(\sum_{f' \neq f} \text{vech}(v_{fb} v_{f'b}') \right) \text{ and } \lim_{F, B \rightarrow \infty} \frac{F}{N_{BB}^2} \sum_{f=1}^F \text{var} \left(\sum_{b' \neq b} \text{vech}(e_{fb} e_{fb}') \right) \quad (50)$$

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_{fb}^d and e_{fb}^s are firm-specific demand and supply components (correlated over banks, and v_{fb}^d and v_{fb}^s 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.

$$\lim_{F,B \rightarrow \infty} \frac{N}{FB} = \kappa \in (0, 1], \quad N \equiv \sum_{b=1}^B F_b = \sum_{f=1}^F B_f; \quad (51)$$

2.

$$\frac{B}{F^2} \rightarrow 0 \text{ as } F, B \rightarrow \infty; \quad (52)$$

3.

$$\frac{F}{B^2} \rightarrow 0 \text{ as } F, B \rightarrow \infty. \quad (53)$$

The first point of Assumption 4 is relatively mild; it guarantees that there is a non-vanishing share of firms (banks) whose number of connections increases proportionally with B (F). 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 $F, B \rightarrow \infty$.

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

Proof.

$$\text{var}(\text{vech}(S_{FF})) = \text{var} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \text{vech}(\eta_{fb} \eta'_{f'b}) \right) \quad (54)$$

$$= N_{FF}^{-2} \text{var} \left(\text{vech} \left(\sum_{b=1}^B \sum_{f' \neq f} \eta_{fb} \eta'_{f'b} \right) \right) \quad (55)$$

$$= N_{FF}^{-2} \sum_{b=1}^B \sum_{b'=1}^B \text{cov} \left(\text{vech} \left(\sum_{f' \neq f} \eta_{fb} \eta'_{f'b} \right), \text{vech} \left(\sum_{\tilde{f} \neq \tilde{f}'} \eta_{\tilde{f}b'} \eta'_{\tilde{f}'b} \right)' \right) \quad (56)$$

$$= N_{FF}^{-2} \sum_{b=1}^B \sum_{f' \neq f} \sum_{\tilde{f} \neq \tilde{f}'} \text{cov} \left(\text{vech}(\eta_{fb} \eta'_{f'b}), \text{vech}(\eta_{\tilde{f}b'} \eta'_{\tilde{f}'b})' \right) \quad (57)$$

$$+ 2N_{FF}^{-2} \sum_{b' \neq b} \sum_{f' \neq f} \text{cov} \left(\text{vech}(\eta_{fb} \eta'_{f'b}), \text{vech}(\eta_{f'b'} \eta'_{f'b'})' \right) \quad (58)$$

$$\equiv \omega_{FF}, \quad (59)$$

where in (58) terms are dropped for which $(\tilde{f}, f) \neq (f, f')$ since these are zero by the structure and independence conditions in Assumption 3. We now derive the asymptotic behaviour of this variance. Note that the first summation contains

$$\frac{1}{4} \sum_{b=1}^B (F_b(F_b - 1))^2 \quad (60)$$

terms, while the second summation contains

$$\sum_{b' \neq b} \sum_{f' \neq f} \mathbf{1}_{fb} \mathbf{1}_{f'b'} \mathbf{1}_{f'b} \mathbf{1}_{f'b'} \quad (61)$$

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

$$\lim_{F, B \rightarrow \infty} \frac{1}{F^2 B} \sum_{b=1}^B \frac{F_b(F_b - 1)}{2} = \tilde{\delta} \in (0, 1/2] \quad (62)$$

and thus

$$\lim_{F, B \rightarrow \infty} \left(\frac{1}{F^2 B} \sum_{b=1}^B \frac{F_b(F_b - 1)}{2} \right)^2 = \tilde{\delta}^2 \in (0, 1/4], \quad (63)$$

and additionally

$$\lim_{F, B \rightarrow \infty} \frac{1}{F^4 B} \sum_{b=1}^B \frac{F_b^2 (F_b - 1)^2}{4} = \tilde{\gamma} \in (0, 1/4]. \quad (64)$$

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

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

$$q_{FF}(\boldsymbol{\eta}, \theta) = \text{vech}(S_{FF} - A\Lambda_{FF}A') \quad (65)$$

to $E[q_{FF}(\boldsymbol{\eta}, \theta)]$ 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

$$\|q_{FF}(\boldsymbol{\eta}, \theta) - E[q_{FF}(\boldsymbol{\eta}, \theta)]\| = \|S_{FF} - \Sigma_{FF}\|, \quad (66)$$

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

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:

$$\omega_{FB} = \text{cov} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \text{vech}(\eta_{fb} \eta'_{f'b}), \frac{1}{N_{BB}} \sum_{\tilde{f}=1}^F \sum_{\tilde{b} \neq b'} \text{vech}(\eta_{\tilde{f}b'} \eta'_{\tilde{f}\tilde{b}})' \right) \quad (67)$$

$$= \frac{4}{N_{FF} N_{BB}} \sum_{f' \neq f} \sum_{b' \neq b} \text{cov} \left(\text{vech}(\eta_{fb} \eta'_{f'b}), \text{vech}(\eta_{f'b} \eta'_{f'b})' \right). \quad (68)$$

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

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

$$\eta_{fb,i} \eta'_{f'b,j} \equiv \alpha_{b,ff',ij} + \beta_{b,ff',ij} + \zeta_{b,ff',ij} + \xi_{b,ff',ij}, \quad (69)$$

where

$$\alpha_{b,ff',ij} = (A_{i1} e_{fb}^d + A_{i2} e_{fb}^s) (A_{j1} e_{f'b}^d + A_{j2} e_{f'b}^s) \quad (70)$$

$$\beta_{b,ff',ij} = (A_{i1} v_{fb}^d + A_{i2} v_{fb}^s) (A_{j1} v_{f'b}^d + A_{j2} v_{f'b}^s) \quad (71)$$

$$\zeta_{b,ff',ij} = (A_{i1} e_{fb}^d + A_{i2} e_{fb}^s) (A_{j1} v_{f'b}^d + A_{j2} v_{f'b}^s) \quad (72)$$

$$\xi_{b,ff',ij} = (A_{i1} v_{fb}^d + A_{i2} v_{fb}^s) (A_{j1} e_{f'b}^d + A_{j2} e_{f'b}^s). \quad (73)$$

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

$$\frac{\sqrt{B}}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \alpha_{b,ff',ij} = \frac{\sqrt{B}F^2}{N_{FF}} \sum_{b=1}^B \frac{1}{F^2} \sum_{f' \neq f} \alpha_{b,ff',ij} = \frac{\sqrt{B}F^2}{N_{FF}} \sum_{b=1}^B O_p(F^{-1}), \quad (74)$$

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

$$\frac{\sqrt{B}F^2}{N_{FF}} \sum_{b=1}^B O_p(F^{-1}) = O\left(\frac{BF^2}{N_{FF}}\right) O_p\left(B^{\frac{1}{2}}F^{-1}\right) \rightarrow 0 \quad (75)$$

jointly in F, B , where the limit follows since $O\left(\frac{BF^2}{N_{FF}}\right) \rightarrow O(1)$ by (62) and $B^{\frac{1}{2}}/F \rightarrow 0$ by Assumption 4.

Turning to $\beta_{b,ff',ij}$, $E[\beta_{b,ff',ij}] = E[\eta_{i,fb,i} \eta'_{f'b,j}] = \Sigma_{FF,ij}$, since v_{fb}^d and $v_{f'b}^s$ are the only

source of correlation across firms. Then

$$\sqrt{B} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \right) \quad (76)$$

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,

$$\frac{\sqrt{B}}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \beta_{b,ff',ij} - \Sigma_{FF,ij} \xrightarrow{d} \psi_{b,ij} \sim \mathcal{N}(0, V_{ij}), \quad (77)$$

where

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

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

$$\sqrt{B} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \right) \xrightarrow{d} \Psi_b \sim \mathcal{N}(0, V). \quad (79)$$

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

$$\sqrt{B} \left(\frac{1}{N_{FF}} \sum_{b=1}^B \sum_{f' \neq f} \zeta_{b,ff',ij} - E[\zeta_{b,ff',ij}] \right) = \sqrt{B} O_p \left(N_{FF}^{-\frac{1}{2}} \right), \quad (80)$$

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

$$\sqrt{B} \text{vech} \left(S_{FF,ij} - (A \Lambda_{FF} A')_{ij} \right) \xrightarrow{d} \mathcal{N}(0, V_{ij}) \quad (81)$$

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

$$\begin{pmatrix} \sqrt{B}q_{FF}(\boldsymbol{\eta}, \theta) \\ \sqrt{F}q_{BB}(\boldsymbol{\eta}, \theta) \end{pmatrix} \xrightarrow{d} \begin{pmatrix} \Psi_b \\ \Psi_f \end{pmatrix} \sim \mathcal{N}(0, \mathbf{W}), \quad (82)$$

where

$$\mathbf{W} = \begin{bmatrix} \mathbf{W}_{FF} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_{BB} \end{bmatrix}. \quad (83)$$

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),

$$\tilde{\Phi}(\hat{\theta} - \theta) \xrightarrow{d} \mathcal{N}(0, \mathbf{W}), \quad (84)$$

where

$$\tilde{\Phi} = \left(\begin{bmatrix} \sqrt{B} & 0 \\ 0 & \sqrt{F} \end{bmatrix} \otimes I_3 \right) \hat{\Phi}, \quad \hat{\Phi} = \left. \frac{\partial q(\boldsymbol{\eta}, \theta)}{\partial \theta'} \right|_{\hat{\theta}}. \quad (85)$$

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

$$\mathbf{W}_{FF} = \lim_{F, B \rightarrow \infty} \frac{B}{N_{FF}^2} \sum_{b=1}^B \text{var} \left(\sum_{f' \neq f} \beta_{b, ff'} - \text{vech}(\Sigma_{FF}) \right) \quad (86)$$

$$= \lim_{F, B \rightarrow \infty} \sum_{b=1}^B \frac{B(F_b(F_b - 1)/2)^2}{N_{FF}^2} \text{var} \left(\frac{1}{F_b(F_b - 1)/2} \sum_{f' \neq f} \beta_{b, ff'} - \text{vech}(\Sigma_{FF}) \right) \quad (87)$$

$$= \frac{\tilde{\gamma}}{\tilde{\delta}^2} \tilde{\mu} \lim \text{var} \left(\frac{1}{F_b(F_b - 1)/2} \sum_{f' \neq f} \beta_{b, ff'} - \text{vech}(\Sigma_{FF}) \right) \quad (88)$$

$$= \frac{\tilde{\gamma}\tilde{\mu}}{\tilde{\delta}^2} \Omega_{FF}, \quad (89)$$

where $\Omega_{FF} \equiv \lim_{F \rightarrow \infty} \text{var} \left(\frac{1}{F_b(F_b - 1)/2} \sum_{f' \neq f} \beta_{b, ff'} - \text{vech}(\Sigma_{FF}) \right)$ 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 $\frac{B(F_b(F_b - 1)/2)^2}{N_{FF}^2} \rightarrow 0$, so their contribution to the variance in (88) is asymptotically negligible. For $\tilde{\mu}B$, the scaled summations are asymptotically independently and identically distributed.

Define

$$\hat{\mathbf{W}}_{FF} = \frac{B^2}{N_{FF}^2} \frac{1}{B} \sum_{b=1}^B \left(\sum_{f' \neq f} \text{vech}(\eta_{fb} \eta'_{f'b}) - \text{vech}(S_{FF}) \right) \left(\sum_{f' \neq f} \text{vech}(\eta_{fb} \eta'_{f'b}) - \text{vech}(S_{FF}) \right)' \quad (90)$$

Adopting the decomposition in (70)-(73), $\hat{\mathbf{W}}_{FF}$ 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. $\frac{1}{F^2} \sum_{f' \neq f} \alpha_{b,ff'} = O_p(F^{-1}) \rightarrow 0$, $\frac{1}{F^2} \sum_{f' \neq f} \zeta_{b,ff'} = O_p(F^{-1}) \rightarrow 0$, $\frac{1}{F^2} \sum_{f' \neq f} \xi_{b,ff'} = O_p(F^{-1}) \rightarrow 0$. Therefore, second powers involving those terms are asymptotically negligible. The only non-vanishing terms are those depending exclusively on $\beta_{b,ff'}$.

First, denote

$$\frac{1}{F(F-1)/2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(\Sigma_{FF}) \xrightarrow{d} \nu_b, \quad (91)$$

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

$$\nu_b \sim \nu, \quad E[\nu] = 0, \quad \text{var}(\nu) = \Omega_{FF}, \quad (92)$$

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

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

$$\frac{B}{N_{FF}^2} \sum_{b=1}^B \left(\sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(S_{FF}) \right) \left(\sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(S_{FF}) \right)' \quad (93)$$

$$= \frac{1}{B} \sum_{b=1}^B \frac{B^2 (F_b (F_b - 1) / 2)^2}{N_{FF}^2} \left(\frac{1}{F_b (F_b - 1) / 2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(S_{FF}) \right) \quad (94)$$

$$\times \left(\frac{1}{F_b (F_b - 1) / 2} \sum_{f' \neq f} \beta_{b,ff'} - \text{vech}(S_{FF}) \right)' \quad (95)$$

$$\xrightarrow{p} \lim_{F, B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B \frac{B^2 (F_b (F_b - 1) / 2)^2}{N_{FF}^2} \nu_b \nu_b' = \tilde{\mu} \frac{\tilde{\gamma}}{\tilde{\delta}^2} \Omega_{FF} + (1 - \tilde{\mu}) \times 0 = \frac{\tilde{\gamma} \tilde{\mu}}{\tilde{\delta}^2} \Omega_{FF}, \quad (96)$$

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. □

ONLINE APPENDIX

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

$$\hat{A} = \hat{P}(\tilde{A}\tilde{Y})\tilde{A}\tilde{Y}, \quad (97)$$

where $P(\cdot)$ is a signed permutation matrix and $\tilde{Y}(\cdot)$ is a diagonal matrix imposing a scale normalization. In practice, we set $\tilde{Y} = |\tilde{\Lambda}_{FF}|^{\frac{1}{2}}$ (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.¹⁸ A signed permutation matrix is one in which a subset of columns may be negative. We choose $\hat{P}(\tilde{A}\tilde{Y})$ 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 $P\tilde{A}\tilde{Y} - \begin{pmatrix} \text{std}(\Delta r_{fb}) & -\text{std}(\Delta r_{fb}) \\ \text{std}(\Delta l_{fb}) & \text{std}(\Delta l_{fb}) \end{pmatrix}$. 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,

$$\hat{P}(\tilde{A}\tilde{Y}) = \operatorname{argmin}_{P \in \mathcal{P}} \sqrt{\sum_{i=1}^2 \sum_{j=1}^2 \left(P\tilde{A}\tilde{Y} - \begin{pmatrix} \text{std}(\Delta r_{fb}) & -\text{std}(\Delta r_{fb}) \\ \text{std}(\Delta l_{fb}) & \text{std}(\Delta l_{fb}) \end{pmatrix} \right)_{ij}^2}, \quad (98)$$

¹⁸Using 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 Bias of AKM in simple example

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

$$\tilde{d}_1 = d_1 + e_1 + (s_2 - a) \quad (99)$$

$$\tilde{d}_2 = d_2 + e_2 + (s_2 - a) \quad (100)$$

$$\tilde{s}_1 = s_1 - s_2 + (a - (-a))s_1 - s_2 + 2a. \quad (101)$$

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, $s_2 = 0, u_{f2}^d = d_f$, 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

$$\Delta r_{fb} = u_{fb}^d - u_{fb}^s, \quad (102)$$

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,

$$\tilde{d}_1 = (d_1 - d_2) - (e_1 - e_2) \quad (103)$$

$$\tilde{s}_1 = -s_1 + a + (d_2 - e_2) \quad (104)$$

$$\tilde{s}_2 = -s_2 - a + (d_2 - e_2). \quad (105)$$

The commentary is qualitatively unchanged from that in the text.

C.2 Matching functions compatible with Assumption 2

In this appendix, we illustrate the conditions needed for Assumption 1 to be compatible with assortative matching. We focus on the labour market example, where observations are sequential, but similar arguments could be applied to financial settings where match formation and destruction are considered (in contrast to our leading application, where we study the intensive margin). First, we assume that u_{ij}^w and $u_{ij'}^e$ are independent, for all $j \neq j'$, and likewise u_{ij}^w and $u_{i'j}^e$ are independent, for all $i \neq i'$. This is a primitive, defining (and identifying) assumption of the match-specific effects. We focus on Assumption 1.3 for concreteness, but a similar argument applies to 1.4. Let \bar{D} in a conditional expectation denote the remaining network structure besides that otherwise specified. For instance, in $E[\cdot | D_{ij}, \bar{D}]$, $\bar{D} = D_{-ij}$. The additional condition required is either

$$E[u_{ij'}^e | D_{ij'} = 1, D_{ij} = 1, \bar{D}] = E[u_{ij'}^e | D_{ij'} = 1] \quad (106)$$

or

$$E[u_{ij}^w | D_{ij} = 1, D_{ij'} = 1, \bar{D}] = E[u_{ij}^w | D_{ij} = 1]. \quad (107)$$

If the former holds, then

$$E[u_{ij}^w u_{ij'}^e | D_{ij} = 1, D_{ij'} = 1, \bar{D}] = E[E[u_{ij'}^e | u_{ij}^w, D_{ij} = 1, D_{ij'} = 1, \bar{D}] u_{ij}^w | D_{ij} = 1, D_{ij'} = 1, \bar{D}] \quad (108)$$

$$= E[E[u_{ij'}^e | D_{ij} = 1, D_{ij'} = 1, \bar{D}] u_{ij}^w | D_{ij} = 1, D_{ij'} = 1, \bar{D}] \quad (109)$$

$$= E[E[u_{ij'}^e | D_{ij'} = 1] u_{ij}^w | D_{ij} = 1, D_{ij'} = 1, \bar{D}] \quad (110)$$

$$= 0, \quad (111)$$

where the first equality uses the law of iterated expectations, the second uses the definitional independence of u_{ij}^w and $u_{ij'}^e$, the third uses (106), and the last uses Assumption 1.2. If

instead the condition in (107) holds, a similar argument can be made by conditioning on u_{ij}^e in the first step. The condition in (106) can be interpreted as saying that the relative number of matches in which employer j' or worker i is involved contains no predictive information for the employer match-specific effect $u_{ij'}^e$ above and beyond the fact that they have chosen to match in the first place. This seems plausible, since the number of matches an employer experiences over the sample can be determined far more directly by myriad factors like firm size, industry, and aggregate conditions than through the level of employer match-specific effects (which would require a very strong correlation structure to contain a signal strong enough to retain predictive power through the indicators in D). Indeed, a model in which employers make sequential offers until the required number of vacancies is filled, see for instance Kaas and Kircher (2015) or Gottfries and Jarosch (2023), is compatible with this assumption. The condition in (107) requires the same with respect to the relative matches of i and j and u_{ij}^w .

A related argument can also be used to justify the conditional expectations in Assumption 1.3-4 in applications like ours to credit markets, even though we focus on the intensive margin. In this setting, unconditional orthogonality or independence of supply and demand effects across relationships is definitional, or can be seen as a natural consequence of the very concept of structural shocks. However, even after focusing on the intensive margin, “matching” may play a role, since changes in a continuing relationship are only observed *conditional on that relationship continuing*. Given particularly adverse shocks on either side, the relationship could disappear from the sample. The fact that this does not happen could imply something jointly about the values for u_{fb}^d and $u_{f'b}^s$, or u_{fb}^d and $u_{f'b}^s$. However, by the same argument as above, provided either

$$E [u_{f'b}^s | D_{f'b} = 1, D_{fb} = 1, \bar{D}] = E [u_{f'b}^s | D_{f'b} = 1] \quad (112)$$

or

$$E [u_{fb}^d | D_{fb} = 1, D_{f'b} = 1, \bar{D}] = E [u_{fb}^d | D_{fb} = 1]. \quad (113)$$

holds, Assumption 1.3 is satisfied, and analogously for Assumption 1.4. The former says that the number of relationships bank b' and firm f have has no predictive content for the level of bank b' 's supply shock to f (beyond the fact that the lending relationship continues). This is consistent with some banks simply being larger providers of credit than others for institutional or historical reasons, rather than because they experience more positive supply shocks in a given period. Similarly, the second condition says that the number of relationships bank b and firm f have has no predictive content for the level of firm f 's demand shock to b (beyond the fact that the lending relationship continues), consistent with a firm's total

borrowing level being determined, for instance, by its practical need for credit, rather than the precise coincidence of its demand shocks with banks' supply shocks in a given period.

C.3 Partialing out covariates

In this section, we show that covariates, X_{fb} , can be partialled out of the regression model in (42), under some additional assumptions.

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

$$E[u_{fb}|X_{fb}, D] = E[u_{fb}|D]. \quad (114)$$

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 $\text{var}(Au_{fb}X'_{fb}|D) < \infty$.

Proof. By the usual algebra, making the dependence on network structure D explicit, the OLS estimator of Γ is

$$\hat{\Gamma} = \Gamma + \frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B Au_{fb}X'_{fb}D_{fb} \left(\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B X_{fb}X'_{fb}D_{fb} \right)^{-1}. \quad (115)$$

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

$$E \left[\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B Au_{fb}X'_{fb}D_{fb} \right] = \frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B AE [E[u_{fb}|X, D] X'_{fb}|D] \quad (116)$$

$$= \frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B AE [E[u_{fb}|D] X'_{fb}D_{fb}|D] \quad (117)$$

$$= 0, \quad (118)$$

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

is

$$\text{var} \left(\frac{1}{N} \sum_{f=1}^F \sum_{b=1}^B Au_{fb}X'_{fb}|D \right) \quad (119)$$

$$= \frac{1}{N^2} \sum_{f=1}^F \sum_{b=1}^B \sum_{f'=1}^F \sum_{b'=1}^B \text{cov} (Au_{fb}X'_{fb}, X_{f'b'}Au'_{f'b'}|D) \quad (120)$$

$$\leq \frac{F^2B^2}{N^2} \frac{1}{F^2B^2} \sum_{f=1}^F \sum_{b=1}^B (F+B-1) \text{var} (Au_{fb}X'_{fb}|D) \quad (121)$$

$$= \frac{F^2B^2}{N^2} \frac{(F+B-1)}{FB} \text{var} (Au_{fb}X'_{fb}|D) \rightarrow 0, \quad (122)$$

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 $F, B \rightarrow \infty$, so the denominator converges in probability to $E [X_{fb}X'_{fb}D_{fb}]$. Finally, applying Slutsky's Theorem yields $\hat{\Gamma} \xrightarrow{p} \Gamma$. \square

C.4 Comovement of supply and demand

In this Appendix, we show that orthogonal structural shocks to supply and demand can still accommodate dependence in the quantity supplied and the quantity demanded, potentially reconciling the macroeconomic notion of orthogonal structural shocks with perhaps more familiar application-specific concepts.

Consider the model for quantities supplied and demanded,

$$q_{fb}^s = \beta_s p_{fb} + \gamma_s \lambda_{fb} + u_{fb}^s \quad (123)$$

$$q_{fb}^d = \beta_d p_{fb} + \gamma_d \lambda_{fb} + u_{fb}^d, \quad (124)$$

where $\beta_s > 0$, $\beta_d < 0$, and λ_{fb} is some factor that potentially varies at the relationship level (for instance, the length of the relationship), or at the firm or bank level (for instance, a new investment opportunity becoming available to the firm). The shocks, u_{fb}^s, u_{fb}^d , however, remain mutually orthogonal. In this model, both quantity supplied and demanded can respond directly to the common factor λ_{fb} , inducing correlation in partial equilibrium. Demand can respond directly to u_{fb}^d , while supply responds indirectly through price adjustments, inducing further correlation between supply and demand; analogously, u_{fb}^s directly impacts supply and indirectly impacts demand. While demand and supply shocks remain orthogonal, quantities supplied and demanded are correlated and must be identical in equilibrium.

Solving for the equilibrium price also allows us to map between the setting in (123) and the coefficients on the structural shocks in A . Equating supply and demand to clear the market and solving for price yields

$$p_{fb}^* = \frac{(\gamma_s - \gamma_d)\lambda_{fb} + u_{fb}^s - u_{fb}^d}{\beta_d - \beta_s}. \quad (125)$$

Substituting in q_{fb}^s and q_{fb}^d yields

$$q_{fb}^* = \frac{-\beta_s}{\beta_d - \beta_s}u_{fb}^d + \frac{\beta_d}{\beta_d - \beta_s}u_{fb}^s + \frac{(\beta_d\gamma_s - \beta_s\gamma_d)}{\beta_d - \beta_s}\lambda_{fb}. \quad (126)$$

Doing so makes clear that any factors common to supply and demand, λ_{fb} , ought to be partialled out of quantity and price before embarking on our identification strategy. After doing so – such that the term in λ_{fb} can be ignored – we obtain

$$p_{fb}^* = \frac{-1}{\beta_d - \beta_s}u_{fb}^d + \frac{1}{\beta_d - \beta_s}u_{fb}^s \quad (127)$$

$$q_{fb}^* = \frac{-\beta_s}{\beta_d - \beta_s}u_{fb}^d + \frac{\beta_d}{\beta_d - \beta_s}u_{fb}^s, \quad (128)$$

or $\eta_{fb} = Au_{fb}$, with

$$A = \begin{bmatrix} \frac{1}{\beta_s - \beta_d} & \frac{-1}{\beta_s - \beta_d} \\ \frac{\beta_s}{\beta_s - \beta_d} & \frac{-\beta_d}{\beta_s - \beta_d} \end{bmatrix}. \quad (129)$$

These expressions make clear how our framework can be consistent with models of supply and demand that include simultaneous endogenous responses.

D Monte Carlo Study

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

$$A = \begin{bmatrix} 7.5101 & -1.5346 \\ 0.1670 & 2.2644 \end{bmatrix}. \quad (130)$$

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:

$$u_{fb}^i = z_f^i + z_b^i + z_{fb}^i, \quad i \in \{d, s\}, \quad f = 1, \dots, F, \quad b = 1, \dots, B, \quad (131)$$

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 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.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

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.

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 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.

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)¹⁹.

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

¹⁹In 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.

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 are biased, 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 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 quantify the relative importance of the fixed effects bias 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 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

Country	Pandemic			Inflation			Tightening		
	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).

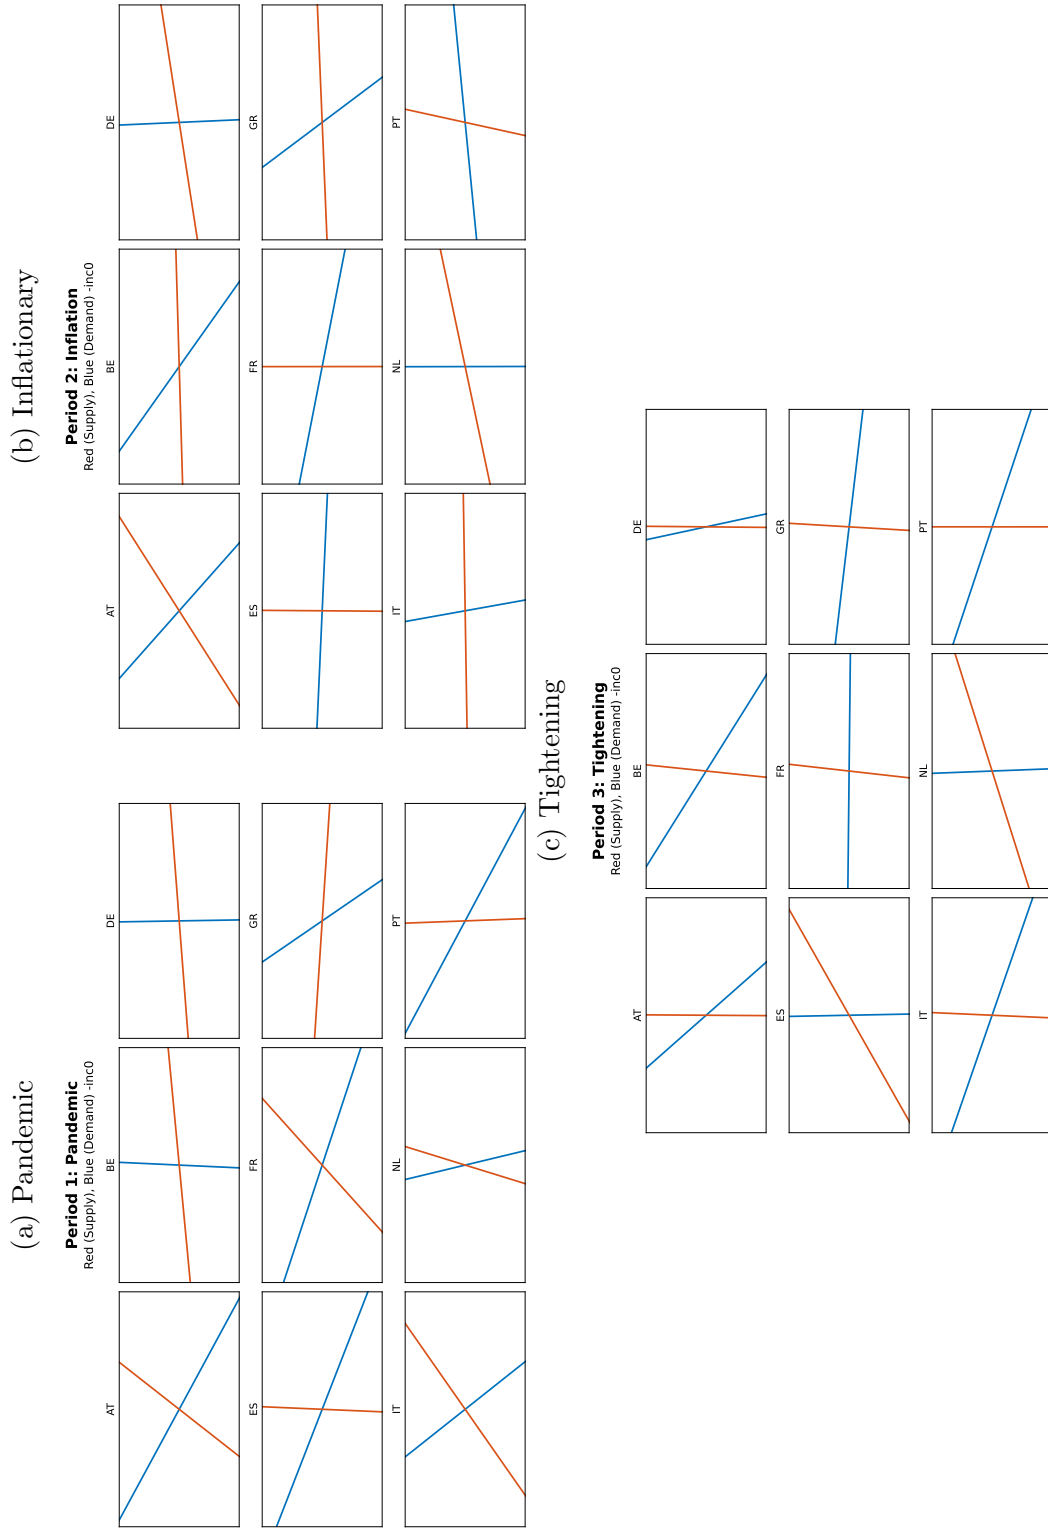
E.2 Supply and demand curves over time

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 Evidence of bias in AKM innovations

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. An alternative approach to assess and quantify the scope for bias is regressing the firm-time (bank-time) fixed effects on our demand and supply innovations. This analysis is reported in Table [E.6](#).

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



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

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*** (0.000)			0.234*** (0.000)
Average Supply innovation	-0.189*** (0.000)		-0.263*** (0.000)	
AKM Firm-Time FE		-0.483*** (0.002)	1.179*** (0.003)	
Firm-time average of AKM Bank-Time FE		-0.741*** (0.007)		-1.102*** (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*** (0.017)			0.499*** (0.017)
Average Supply innovation (b,t)	-0.362*** (0.008)		-0.406*** (0.009)	
Bank-Time average of AKM Firm-Time FE		-0.401*** (0.129)	0.223* (0.117)	
AKM Bank-Time FE		-0.439*** (0.077)		-2.256*** (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 \times 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.

Table E.6: Fixed effects: contribution of demand and supply

	(1)	(2)	(3)
	Firm-Time FE in AKM		
Average Demand Innovation	3.726*** (0.006)	3.576*** (0.007)	
Average Supply Innovation	6.063*** (0.006)		5.973*** (0.007)
Observations	11539414	11539414	11539414
Adjusted R-squared	0.52	0.13	0.38
SE-cluster	Firm	Firm	Firm
	Bank-Time FE in AKM		
Average Demand Innovation	1.834*** (0.123)	1.602*** (0.128)	
Average Supply Innovation	1.731*** (0.118)		1.552*** (0.122)
Observations	31148	31148	31148
Adjusted R-squared	0.28	0.20	0.22
SE-cluster	Bank	Bank	Bank

Notes: The table reports coefficients from regressions projecting estimated two-way fixed effects on our identified demand and supply innovations. In the top panel, the dependent variable is the firm-time fixed effect from an AKM-style decomposition at the bank-firm-quarter level, regressed on firm-time averages of the relationship-level demand and supply innovations, aggregated across all relationships of firm f in quarter t . In the bottom panel, the dependent variable is the bank-time fixed effect from the same decomposition, regressed on corresponding bank-time averages of the demand and supply innovations, aggregated across all borrowing relationships of bank b in quarter t . Column (1) includes both innovations jointly; column (2) includes only the demand innovation; column (3) includes only the supply innovation. All specifications include country \times time fixed effects. Standard errors are clustered at the firm level in the top panel and at the bank level in the bottom panel. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Acknowledgements

We are grateful to 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, and Ugo Albertazzi (discussant), Diana Bonfim, Max Bruche (discussant), Gabe Chodorow-Reich, Ferre De Graeve, Hans Degryse, Aureo de Paula, Martina Jasova (discussant), Peter Jiang, Edoardo Rainone (discussant), Enrico Sette, and Andrei Zeleneev for helpful discussions. 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.

Olivier De Jonghe

European Central Bank, Frankfurt am Main, Germany; National Bank of Belgium, Brussels, Belgium; Tilburg University, Tilburg, The Netherlands; email: olivier.dejonghe@nbb.be

Daniel Lewis (corresponding author)

University College London, London, United Kingdom; Centre for Economic Policy Research, London, United Kingdom; email: daniel.j.lewis@ucl.ac.uk

© European Central Bank, 2026

Postal address 60640 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

PDF

ISBN 978-92-899-7901-6

ISSN 1725-2806

doi:10.2866/2974560

QB-01-26-139-EN-N