Does Monetary Tightening Improve Banking Stability? The Role of Bank Cost Efficiency[*]

Richard Adjei Dwumfour
Curtin University, Australia

Lei Pan[†]
Curtin University & Monash University, Australia

Mark N. Harris
Curtin University, Australia

For Current Version

Working paper — please do not cite

Abstract

Yes, monetary tightening improves banking stability, but stability weakens over the medium run, and bank efficiency shapes this trajectory. The empirical exercise draws on annual data for 3,903 banks across 95 countries over 1996–2024, measuring efficiency via a stochastic metafrontier and identifying policy shocks through Taylor-rule deviations (and, in IV, central bank independence). Fixed-effects estimates and local projections confirm that a one-standard-deviation tightening initially raises Z-scores and reduces non-performing loan growth—yet as borrower distress accumulates, these gains erode at medium horizons. Critically, this deterioration is notably smoother for high-efficiency banks, consistent with stronger risk control. Tightening further compresses net interest margins—most persistently for efficient banks—while credit growth responses remain heterogeneous. An efficiency-augmented New Keynesian DSGE model reproduces both the sign reversal and the cross-bank smoothing pattern. Counterfactual experiments show that the efficiency gap across bank types collapses to zero under homogeneous cost efficiency, and that the impact stabilisation vanishes entirely when the risk-management channel is shut down.

Keywords: Monetary policy; bank stability; risk-taking channel; local projections; stochastic metafrontier; DSGE.

JEL Classification: E52; E58; G21; G28.

---

[*]We would like to thank seminar participants at Bank of Ghana, and Curtin University for helpful comments and suggestions. The usual disclaimer applies.

[†]Corresponding author: `lei.pan@curtin.edu.au`.

1 INTRODUCTION

A central lesson of the past two decades is that monetary policy and financial stability are jointly determined, even when institutional mandates treat them as separate objectives. In the wake of the global financial crisis (GFC) and subsequent episodes of unconventional policy, economists and policymakers have increasingly argued that central banks cannot ignore the financial stability consequences of interest-rate decisions, consistent with calls to prioritise financial stability as a core policy goal [Yellen, 2014] and with the broader notion of a twin objective of monetary and financial stability [Oosterloo and De Haan, 2004]. Yet, achieving this twin objective has largely remained elusive [Borio and Crockett, 2000, Borio, 2005], reflecting both conceptual tensions and empirical disagreements about how monetary policy decisions transmit through bank balance sheets, borrower risk, and the accumulation of vulnerabilities.

The empirical question is therefore direct: does monetary tightening (or easing) improve banking stability, and over what horizons? Indeed, a prominent line of assertion by some economists [Rajan, 2006, Taylor et al., 2010, Altunbas et al., 2014] suggests that accommodative policy in the aftermath of the dot-com bubble bust may have contributed to the build-up of vulnerabilities preceding the GFC, while others [Bernanke, 2010, Svensson, 2010] disagree with this interpretation and instead emphasise regulatory and supervisory failures. Yet, existing empirical evidence does not deliver a consensus. Some studies find that tighter (looser) policy strengthens (weakens) banking stability by restraining (encouraging) risk-taking [Jiménez et al., 2014, Lamers et al., 2019]; others document destabilising effects through higher debt-service burdens, weaker borrower cash flows, and rising credit losses [De Graeve et al., 2008, Altunbas et al., 2014]. Crucially, the sign and timing of these effects appear to vary across institutional settings, financial structures, and horizons.

For example, Lamers et al. [2019] report that expansionary monetary policy is associated with improved bank stability in the euro area, but reduces stability in the U.S., using market-based measures such as long-run marginal expected shortfall (LRMES) in a fixed-effects framework. In contrast, De Graeve et al. [2008] finds that unexpected monetary tightening increases bank distress in Germany, with material heterogeneity in terms of bank characteristics, and estimates a VAR for distress responses. Importantly, even where bank heterogeneity is explicitly studied — typically along dimensions such as capitalisation, liquidity, or funding structure — the role of *cost efficiency* as a fundamental determinant that governs banks' risk-control capacity, loss absorption, and the smoothness of medium-run stability dynamics has received limited direct attention. Consequently, these existing results diverge for three related reasons: (i) context and external validity (with a disproportionate focus on the U.S. and Europe), (ii) the measurement of "stability" (market-based systemic-risk indicators versus accounting- and balance-sheet-based solvency and asset-quality metrics), and (iii) identification and dynamics (static panel designs versus impulse-response approaches that trace horizon-by-horizon transmission).

Against this backdrop, we make two testable predictions primarily. First, we test whether monetary tightening can increase (or decrease) stability and whether this effect is horizon-dependent. Second, we also introduce and test an *efficiency channel* to understand whether the monetary policy-stability relationship is systematically dependent on the level of cost efficiency of banks. To assess these predictions empirically, the paper assembles a global panel of roughly 90 countries and nearly 4,000 banks, which enables a unified assessment of monetary policy transmission to bank stability across diverse banking systems. We also utilise balance sheet-based measures — Z-score and non-performing loans (NPLs) — of banking stability, which have been extensively used in the banking stability literature [Uhde and Heimeshoff, 2009, Houston et al., 2010, Demirgüç-Kunt and Huizinga, 2010, Carretta et al., 2015]. While market-based risk measures (typically constructed from equity returns and market valuations) are informative when available, they are inherently limited to publicly traded (and sufficiently liquid) institutions and therefore offer incomplete coverage in global bank panels that include privately held banks and many emerging-market institutions [Acharya et al., 2012, Brownlees and Engle, 2017]. We also estimate the model using various techniques, including fixed effects regression (FE), local projections (LP), and instrumental variables (IV) approach, to provide robustness to our results.

Overall, we find that monetary tightening improves banking stability. This is consistent across most regions (except for Europe & Central Asia), income groups, especially for high income, low income and lower middle income countries – and for developed and developing countries. Moreover, the results show a systematic horizon dependence: a contractionary monetary policy innovation raises bank stability on impact and at short horizons, but this stabilisation is followed by a deterioration at medium horizons. We provide three key complementary channels to support the interpretation. First, the medium-run decline is systematically smoother for more cost-efficient banks, where cost efficiency is interpreted as a shifter of operating technology and managerial capability linked to monitoring, screening, and loss-mitigation capacity [Berger and Humphrey, 1997, Berger and DeYoung, 1997]. Second, intermediation margins (net interest margins) compress on impact and recover more slowly among high-efficiency banks. Lastly, credit growth declines with meaningful heterogeneity across efficiency groups. Taken together, the evidence suggests an adjustment margin beyond pass-through to intermediation margins: operational and risk-control capacity shapes how monetary tightening propagates into realised asset quality stress.

The analysis builds on, and contributes to, a broad literature showing that policy rates shape bank lending and risk-taking through bank funding conditions, intermediation constraints, and the broader macro-financial environment in which borrower creditworthiness evolves [Bernanke and Gertler, 1995, Kashyap and Stein, 2000, Gertler and Kiyotaki, 2010, Borio and Zhu, 2012, Jiménez et al., 2014]. Rather than imposing a priori that monetary policy is stabilising or destabilising, the empirical results emphasise that both forces can coexist over different horizons; perhaps, this may explain the divergent views in the literature. The key novelty is that we document this horizon dependence within a single empirical design and show that bank cost efficiency is a first-order determinant of the medium-run dynamics even when the response of bank margins moves in the opposite direction.

The paper makes three important contributions. First, it establishes horizon-dependent sign reversal in the response of bank stability to monetary tightening in a large cross-country banking panel, and shows that the subsequent decline is not uniform across banks: cost efficiency robustly predicts the extent of smoothing. Second, it clarifies the relevant channels by jointly tracing stability, net interest margins, and credit growth across efficiency groups, which distinguishes a margin-based mechanism from a loss-mitigation mechanism. Third, it develops an Efficiency-Augmented Banking DSGE (EAB-DSGE) framework that mirrors the empirical findings. The model is a standard New Keynesian core [Woodford and Walsh, 2005, Galí, 2015] which is augmented with a banking block in which: i) a persistent borrower-distress state gradually raises default risk after tightening, and ii) bank risk-management effort which responds on impact, with its effectiveness and cost governed by bank efficiency. The interaction of a slow loss channel and a fast risk-management channel reproduces the sign reversal and efficiency-conditioned smoothing observed in the data, providing a coherent mapping from monetary shocks to stability dynamics.

Finally, the results have direct policy implications, especially for central banks and monetary authorities. If tightening delivers short-run stabilisation but induces delayed fragility through borrower distress, then policy assessments based on near-term indicators may understate risks during tightening cycles. Moreover, the distribution of bank efficiency is not merely a micro-level feature: it conditions aggregate resilience by shaping the time profile of stability and the cross-sectional propagation of monetary policy. This places operational efficiency and risk-management capacity as central state variables for macroprudential monitoring when monetary policy shifts, particularly in global tightening episodes where vulnerabilities can accumulate even as measured stability improves initially.

The remainder of the study is organised as follows. Section 2 provides a brief review of related literature. Section 3 describes the data used in the study, while Section 4 presents the methods used in analysing the data. In Section 5, we present and discuss the empirical results. Section 6 presents the supporting DSGE model, calibrated to the empirical results, which explains the channels of monetary tightening and their impact on banking stability. Section 7 concludes the study and discusses the policy implications of the findings.

2 RELATED LITERATURE

A rich literature studies the determinants of banking stability, spanning both bank-level fundamentals (e.g., capitalisation, funding structure, governance, and operational performance) and macro-financial conditions (notably the monetary policy stance). This paper sits at the intersection of two strands: i) how monetary policy decisions transmit to bank stability over different horizons, and ii) how *cost efficiency* conditions that transmission. We review these strands to motivate the empirical design and to clarify the gaps the study addresses.

A substantial empirical literature investigates whether monetary policy influences bank stability and through which mechanisms. A central framework is the *risk-taking channel*: prolonged accommodation can compress intermediation margins, raise incentives to reach for yield, and soften screening/monitoring—thereby increasing the risk content of bank balance sheets even when contemporaneous measured performance looks benign [Rajan, 2006, Jiménez et al., 2014]. Related arguments emphasise that easy financial conditions can inflate asset prices and collateral values, reducing measured risk and relaxing constraints in a way that may later unwind sharply [Bernanke et al., 1996, Matsuyama, 2007].

Micro-evidence strongly supports the existence of a risk-taking channel, but its implications for stability are not uniform across settings. Using Spanish credit-register data, Jiménez et al. [2014] show that lower policy rates induce banks—especially those with weaker buffers—to expand credit to ex ante riskier borrowers, and that these loans default more subsequently, consistent with a shift in the composition of credit risk rather than only the volume. Complementary evidence from a dollarised financial system is provided by Ioannidou et al. [2015], who exploit the pass-through of U.S. monetary policy to Bolivian banks and document greater risk-taking and softer pricing when U.S. rates fall. At the macro-finance interface, Maddaloni and Peydró [2011] utilise bank lending survey information for the U.S. and the euro area, demonstrating that low short-term rates (monetary policy) are associated with a measurable softening of lending standards—an effect that interacts with securitisation and supervisory environments. Taken together, these results rationalise why monetary accommodation can be associated with increased risk-taking; however, they do not deliver a consensus prediction for stability once one accounts for countervailing forces (e.g., recapitalisation through higher borrower cash flows, valuation effects, and policy backstops), differences in banking structure, and the horizon at which stability is measured.

Importantly, the lack of consensus is not confined to North America and Europe. Beyond these regions, Chen et al. [2017], for instance, analyse over 1000 banks across 29 emerging market economies and reports that an easing of monetary policy increases bank risk (measured by indicators like the loan-loss-provision ratio or Z-score). Moreover, some key bank-specific characteristics have been identified in the previous studies to explain these relationships. For instance, De Graeve et al. [2008] find that these results are dependent on bank size and capitalisation. They find that bank distress responses to monetary policy tightening are largest for small and low-capitalised banks. This is consistent with Kishan and Opiela [2000] who also show the importance of bank size and capitalisation for monetary transmission. Among these characteristics, bank efficiency has not received the necessary attention in this transmission mechanism. This gap is central to our contribution: we study not only whether stability rises or falls after a policy shock, but also whether the *time profile* of stability responses differs systematically by bank efficiency.

Consequently, we review a second strand of literature that examines whether bank efficiency predicts risk and stability. Seminal contributions argue that efficiency contains information about managerial quality and internal controls and therefore should forecast asset quality and risk outcomes. According to the "bad management" hypothesis, cost-inefficient banks are more likely to exhibit weaker monitoring and inferior credit processes, which can lead to problem loans and subsequent instability [Berger and DeYoung, 1997, Kwan and Eisenbeis, 1997]. Subsequent evidence for European banking systems similarly links operational performance to risk outcomes, often finding that more efficient banks display lower risk or superior resilience, though results depend on sample composition, period, and risk proxies [Williams, 2004, Fiordelisi et al., 2011]. At the same time, the evidence is not unidirectional: some studies argue that lower efficiency can coincide with greater stability if it reflects conservative business models or higher capital buffers, implying that the mapping between efficiency and risk may be mediated by balance-sheet structure and strategic choices [Altunbas et al., 2007]. This mixed evidence motivates treating efficiency as an economically meaningful shifter—rather than assuming it is mechanically stabilising in all contexts.

Two limitations of the existing efficiency–stability literature are particularly salient for the present study. First, much of the evidence focuses on *level* relationships (e.g., whether efficiency predicts average risk) rather than on *dynamic* responses to macro-financial shocks. Second, even when monetary policy is part of the conditioning information set, efficiency is rarely modelled as a *channel* that governs banks' adjustment of screening, monitoring, provisioning, and loss absorption—precisely the mechanisms through which policy can generate horizon-dependent effects on stability. Our empirical design targets this channel directly by interacting monetary policy shocks with efficiency groups and by tracing impulse responses across horizons. Moreover, our DSGE framework formalises this economic logic: a fast, efficiency-conditioned risk-management response can stabilise banks on impact, while a slower borrower-distress channel can dominate at medium horizons, generating sign reversal with smoother deterioration among more efficient banks.

In sum, prior research establishes plausible channels through which monetary policy affects bank risk-taking and lending standards, and a separate body of work links efficiency to risk and stability. However, the literature remains divided on whether monetary tightening stabilises or destabilises banks once horizons, institutional settings, and risk measures are taken seriously, and it offers limited direct evidence on cost efficiency as a conditioning channel for the dynamic stability response. This paper addresses these gaps.

3 DATA

We obtain bank-level data from the Osiris database, covering 7,386 listed and unlisted commercial banks across 143 countries over 1996–2024. To ensure reliable frontier estimation in the first-stage efficiency analysis, we exclude countries with fewer than five banks, leaving 7,318 banks in 106 countries for the stochastic metafrontier estimation of cost efficiency. We then proceed to the second stage, where we estimate the effects of monetary policy on bank stability while conditioning on the cost-efficiency measures obtained from the first stage. Owing to data availability for bank-level and macroeconomic controls, the baseline regression sample comprises 3,903 banks across 95 countries.[1] We additionally obtain data on central bank independence from Garriga [2016] and country-level monetary policy stance measures from Müller et al. [2025]. Macroeconomic controls are sourced from the World Bank's World Development Indicators (WDI). All bank-level variables are winsorised at the 1st and 99th percentiles following standard practice [Carretta et al., 2015, Beck et al., 2013, Husted et al., 2020, Adelino et al., 2023].

We discuss the key variables used in our main specification, which examines the impact of monetary policy on banking stability.[2] Following the literature [Roy, 1952, Uhde and Heimeshoff, 2009], we use Z-score as our main measure of banking stability, which measures a bank's distance to insolvency. Our monetary policy variable is constructed using the Taylor-type rule [Taylor, 1993] as used in the literature [Lamers et al., 2019].

We also employ some key bank-specific variables following the literature. First, our key variable of interest is cost efficiency. We employ the stochastic metafrontier technique proposed by Huang et al. [2014], which is discussed in sections 4.1 and 4.2. This is used to estimate the cost efficiency of each bank. Following the extant literature as discussed earlier Berger and Humphrey [1997], Fiordelisi et al. [2011], we expect a positive relationship between cost efficiency and stability. This implies that cost-efficient banks reflect good management, hence exhibit better monitoring and credit processes and will therefore have less risk and be more stable. This underpins our argument that these banks typically enjoy the *efficiency buffer* in the face of policy tightening. We also note that in rare cases, as found by [Altunbas et al., 2007], inefficient banks may be more stable, especially when they are highly capitalised, reflecting a conservative business model with the advantage of higher capital buffers.

Second, we include a measure of bank liquidity as a control variable (Bank liquidity). This is the ratio of liquid assets to total assets. We measure bank size as the log of total assets (Size). The empirical results are ambiguous as to the relationship between bank size and stability. From one perspective, larger banks are perceived as more stable due to their diversification and capital buffers, as well as their ability to enter other markets [Uhde and Heimeshoff, 2009]. On the other hand, larger banks can face the moral hazard problem where they tend to take on more risk. This exposure to higher risk can make these banks unstable. We therefore expect either a positive or a negative relationship between bank size and stability.

We also control for the asset structure of banks, measured as the ratio of fixed assets to total assets (Asset Structure). At the industry/country level, we control for bank concentration, which is the assets of the three largest banks (Bank Concentration). The literature provides evidence of both concentration-stability and concentration-fragility views. In the concentration-stability view, banks in monopolistic banking markets dominated by a few large banks can enjoy higher profits and thus improve their stability. The concentration-fragility view also emphasises the too-big-to-fail argument where larger banks take on more risk, and this risky behaviour is exacerbated by their knowledge that governments normally would bail them out in case of distress [Mishkin, 2016]. We therefore expect either a positive or a negative relationship between bank concentration and stability.

3.1 Summary Statistics

Table 1 summarises the main variables used in the two-stage design and provides a first-pass characterisation of the data's scale and heterogeneity. The stochastic metafrontier variables block indicates a large estimation sample for frontier recovery (49,859 bank-year observations),

[1] The countries and number of banks are listed in Table A.1 in the Online Appendix.

[2] Details of the variables used in the stochastic metafrontier estimations are well discussed under Section 4.2. Other details on the construction of Z-score and our monetary policy variables are also discussed under Section 4.3.

with substantial dispersion in both total costs and output: the mean log total cost is 17.36 (s.d. 3.61) and the mean log gross loans is 15.92 (s.d. 3.66). Input-price ratios also vary meaningfully across banks and countries, consistent with heterogeneous funding and operating environments; in particular, the dispersion in the relative price of borrowed funds (mean 0.20, s.d. 1.45) suggests material cross-sectional and intertemporal shifts in funding conditions that are central to cost-technology estimation. The macro environment is equally heterogeneous: GDP per capita growth averages 1.95 (s.d. 3.49), while inflation is volatile—mean 4.43 with a very large s.d. of 10.81 using CPI (and similarly 4.63 with s.d. 12.43 using the GDP deflator)—which motivates controlling for macro conditions directly in the frontier estimation and in the second-stage stability regressions.

The *efficiency scores* block reveals economically meaningful gaps between country-frontier and metafrontier performance. Average country-frontier cost efficiency is 0.68 (s.d. 0.21), implying that, relative to the best-practice bank within each country technology, the typical bank could reduce costs by roughly one-third for a given output and input prices. By contrast, global metafrontier efficiency averages 0.45 (s.d. 0.18), indicating that once all banks are evaluated against a common best-practice global technology set, the implied cost gap is substantially larger. The intermediate metafrontier benchmarks (regional mean 0.53, development status 0.47, income dynamics 0.49, ECB group 0.46) sit between these two extremes, consistent with persistent cross-country technology differences and the relevance of grouping-based technology sets. This implies that the large dispersion in efficiency may suggest heterogeneous adjustment of banks to monetary policy shocks, with more cost-efficient banks plausibly better positioned to absorb sustained tightening through smoother funding-cost pass-through and more resilient margins.

Competition measures corroborate substantial variation in market power: the Lerner index is 0.52 (s.d. 0.25) at the country benchmark and higher at the global benchmark (0.67, s.d. 0.21), suggesting that market power appears stronger when evaluated against the global metafrontier—consistent with the idea that technological gaps can translate into pricing power differentials. Finally, the key second-stage variables show wide dispersion in stability and meaningful time-series variation in policy. The Z-score averages 26.47 with a large s.d. of 23.65, pointing to pronounced heterogeneity in distance-to-default across banks and over time. Monetary policy stance measures are standardised (means near zero and s.d. around one), which facilitates interpretation of dynamic responses: a one-standard-deviation tightening can be mapped directly into the impulse responses of banking stability. The remaining controls exhibit substantial cross-sectional spread — e.g., liquidity ratios average 35.61 (s.d. 44.40) and concentration (CR3) averages 42.69 (s.d. 19.72) — underscoring that both bank balance-sheet structure and market structure vary sharply across the sample.

Table 1. Summary statistics across blocks

Variable	Symbol	Obs	Mean	Std. Dev.
Panel A: Stochastic metafrontier variables				
Total cost (log)	lcost	49,859	17.36	3.61
Output: Gross Loans (log)	ly1	49,859	15.92	3.66
Price of Capital Input scaled by price of labour (log)	lx2: ln(w2/w1)	49,859	3.9773	1.413
Price of borrowed funds scaled by price of labour (log)	lx3: ln(w2/w1)	49,859	0.20	1.451
GDP per capita growth	GDP p.c. growth	49,859	1.95	3.49
Inflation (CPI) (%)	Inflation	49,859	4.4279	10.8084
Inflation (GDP deflator) (%)	Inflation	49,859	4.633	12.43
Panel B: Efficiency scores				
Cost efficiency (Country frontier)	$\hat{CE}^j$	44,652	0.68	0.21
Cost efficiency (Global metafrontier)	$M\hat{C}E$:(CostEff)-Global	44,652	0.45	0.18
Cost efficiency (Regional metafrontier)	$M\hat{C}E$:(CostEff)-Regional	44,652	0.53	0.22
Cost efficiency (Development-status metafrontier)	$M\hat{C}E$:(CostEff)-Devstat	44,652	0.47	0.19
Cost efficiency (Income-dynamics metafrontier)	$M\hat{C}E$:(CostEff)-IncGrp	44,652	0.49	0.19
Cost efficiency (ECB group metafrontier)	$M\hat{C}E$:(CostEff)-ECBM	44,652	0.46	0.19
Panel C: Competition (Lerner index)				
Lerner index (Country frontier)	$\widehat{L}_{it}^{(g)}$	40,562	0.52	0.25
Lerner index (Global metafrontier)	$\widehat{L}_{it}^{M}$	40,562	0.67	0.21
Panel D: Other variables				
Monetary policy stance (official, standardised)	Policy_j^z	44,269	-0.004	0.96
Monetary policy stance (hybrid, standardised)	Policy_j^z	44,269	-0.003	0.95
Bank stability: (Z-score)	Z-score	44,652	26.47	23.65
Bank stability: NPL growth (log change)	$\Delta \ln(\text{NPL})$	39,227	0.13	1.14
Bank stability: NPL growth minus loan growth	$\Delta \log(\text{NPL}) - \Delta\ell$	39,177	-12.12	3.71
Real GDP (constant 2015, USD billions)	Real GDP	44,652	6866.95	7735.67
Liquidity ratio (%)	Bank Liquidity	44,652	35.61	44.40
Bank size (log of assets)	Size	44,652	16.65	3.56
Asset structure (%)	Asset Structure	44,652	1.72	3.72
Bank concentration (%)	Bank concentration	44,533	42.69	19.72
GDP growth (%)	GDP growth	44,652	2.92	3.20
Inflation (CPI) (%)	Inflation	44,652	3.79	5.11
Institutional Quality	Quality	42,760	0.59	0.85
Net interest margin (%)	NIM	44,649	3.8290	3.6209
Loan growth (log change)	$\Delta\ell$	40,569	12.13	3.59
Loans-to-assets ratio (%)	Loans/TA	44,595	59.52	20.29
Macroprudential: Liquidity	Liquidity	43,816	0.20	0.68
Macroprudential: FX-related limit (LFX)	LFX	43,816	0.02	0.24
Macroprudential: Loan-to-value limit (LTV)	LTV	43,816	0.02	0.40
IV: Central bank independence	CBI	44,188	0.52	0.18

4 EMPIRICAL STRATEGIES

The methods used in estimating the bank efficiency scores are discussed here. We then proceed to specify our main model, which shows the impact of monetary policy on banking stability, also accounting for the role of bank efficiency.

4.1 Estimating bank efficiency – Stochastic Metafrontier (SMF) approach

In this paper, the bank efficiency scores are estimated using the stochastic meta-frontier (SMF) cost function approach of Huang et al. [2014]. While Huang et al. [2014] developed this approach to estimate technical efficiency from a production function, we apply this methodology to a cost function in a similar way as Dwumfour et al. [2022].[3] The advantage of the SMF approach is that it enables the estimation of comparable cost functions for each country. We first estimate the country-specific cost frontier using the Stochastic Frontier (*SF*) and then move on to estimate the metafrontier cost function. In summarising the SMF, suppose that country *j*, its *SF* of the ith decision making unit (DMU) in this case bank, in the tth period is modelled as:

$$C_{jit} = f_t^j\left(\mathbf{X}_{jit}\right) e^{V_{jit}+U_{jit}}, \quad j = 1, 2, \ldots, J: \; i = 1, 2, \ldots, N_j \; ; t = 1, 2, \ldots, T \tag{1}$$

where C_{jit} is the scalar cost and $\mathbf{X}_{jit}$ is the vector of output Y and input prices of the ith bank in country *j*, for the period *t*. The subscript *t* and superscript *j* of the function, $f_t^j(.)$ of the cost frontier indicate that the technologies for the various individual groups may differ at different times. The standard *SF* approach denotes V_{jit} as the statistical noise, while U_{jit} is the term for the cost inefficiency. V_{jit}s are assumed to be $N(0, \sigma_v^{j2})$ and are independent of the $U_{jit}s$ which follow a truncated-normal distribution as $N\left(\mu^j\left(Z_{jit}\right), \; \sigma_u^{j2}\left(Z_{jit}\right)\right)$. Here, the truncation is done at zero and with a mode of $\mu^j\left(Z_{jit}\right)$ where the $Z_{jit}s$ are identified exogenous variables. Here, the cost efficiency (CE) of the bank for the country frontier in this function will be:

$$CE_{it}^j = \frac{f_t^j\left(\mathbf{X}_{jit}\right) e^{V_{jit}}}{C_{jit}} = e^{-U_{jit}} \tag{2}$$

The environmental variables Z_{jit} are exogenous to the banks, even though they are related to the *CE* of the banks in their specific countries.

Here, the meta-frontier cost function that is common to all the countries in period *t* is defined as $f_t^M\left(\mathbf{X}_{jit}\right)$. This function is the same for all groups $j = 1, 2, \ldots, J$. The meta-frontier envelopes the individual country-specific frontiers and can be represented as:

$$f_t^j\left(\mathbf{X}_{jit}\right) = f_t^M\left(\mathbf{X}_{jit}\right) e^{U_{jit}^M}, \; \forall j, i, t \tag{3}$$

where $U_{jit}^M \geq 0$. This implies that $f_t^j(.) \geq f_t^M(.)$ and thus the ratio of the meta cost frontier to the *jth* group's cost frontier is defined as the technological gap ratio (TGR), which is represented as:

$$TGR_{It}^j = \frac{f_t^M\left(\mathbf{X}_{jit}\right)}{f_t^j\left(\mathbf{X}_{jit}\right)} = e^{-U_{jit}^M} \leq 1. \tag{4}$$

The fact that each group or country is exposed to certain unique environmental characteristics – both economic and non-economic – accounts for the technological gap. This makes the technological gap component, U_{jit}^M, country-, bank-, and time-specific. Given the observed outputs and inputs, this ratio measures the ratio of the potential minimum cost available at the metafrontier level to the cost function at the country level. The meta-frontier of bank *i* in country *j* at time *t*, $f_t^M\left(\mathbf{X}_{jit}\right)$ can therefore be expressed as:

$$MCE_{jit} \equiv \frac{f_t^M\left(\mathbf{X}_{jit}\right) e^{V_{jit}}}{C_{jit}} = TGR_{it}^j \times CE_{it}^j \tag{5}$$

where MCE_{jit} is therefore the cost efficiency of the bank with respect to the meta cost frontier, $f_t^M(.)$, as opposed to the bank's cost efficiency, CE_{it}^j, with respect to the group-*j* (country) production technology $f_t^j(.)$. The estimated empirical panel framework is hence expressed as follows:

$$M\hat{C}E_{jit} \equiv T\hat{G}R_{it}^j \times C\hat{E}_{it}{}^j \tag{6}$$

[3] See Dwumfour et al. [2022] for a complete derivation of the cost efficiency following the approach of Huang et al. [2014].

where $M\hat{CE}_{jit}$ is the meta cost efficiency of bank i in country j at time t.

4.2 Empirical specification of SMF model

The study employs a two-step approach to estimate the cost efficiency scores based on the *SMF*. First, bank cost efficiency scores are estimated at the country level, thus using the country-specific frontier. Then, the second step is to estimate the final bank cost efficiency scores using the global cost frontier, which is referred to as the metafrontier.[4] The translog cost function is based on the bank intermediation approach, which has been used widely in the literature [Sealey Jr and Lindley, 1977, Hughes and Mester, 1993, Shamshur and Weill, 2019, Dwumfour et al., 2022]. In this approach, banks are modelled to take deposits, convert them to loans using capital and labour. Hence, in the cost function, we use loans as the output. The translog cost function is therefore given by:

$$\begin{aligned}\ln\left(\frac{TC_{it}}{w_{1it}}\right) &= \alpha_0 + \beta_y \ln Y_{it} + \sum_{m=2}^{3} \beta_m \ln\left(\frac{w_{mit}}{w_{1it}}\right) \\ &+ \frac{1}{2}\gamma_{yy}\left(\ln Y_{it}\right)^2 + \sum_{m=2}^{3} \gamma_{ym} \ln Y_{it} \ln\left(\frac{w_{mit}}{w_{1it}}\right) \\ &+ \frac{1}{2}\sum_{m=2}^{3}\sum_{n=2}^{3} \gamma_{mn} \ln\left(\frac{w_{mit}}{w_{1it}}\right) \ln\left(\frac{w_{nit}}{w_{1it}}\right) + v_{it} + u_{it}.\end{aligned} \tag{7}$$

where TC_{it} is total operating cost of bank i at time t; Y_{it} stands for output (total gross loans); $w_{1it}, w_{2it}, w_{3it}$ represent input prices of labour, physical capital, and borrowed funds; v_{it} denotes random error term *i.i.d.* with $\nu_{it} \sim N(0, \sigma_\nu^2)$, independent of regressors; u_{it} is non-negative cost inefficiency term $u_{it} \sim N(\mu(Z_{it}),\ \sigma_u^2(Z_{it}))$ with $u_{it} \perp\!\!\!\perp \nu_{it}$; and $m, n \in \{2, 3\}$ index the normalised input prices.

We define the input prices based on previous studies [Hasan and Marton, 2003, Fries and Taci, 2005, Davies and Tracey, 2014, Shamshur and Weill, 2019, Dwumfour et al., 2022]. The price of labour (w_{1it}) is the ratio of staff expenses to total assets. The price of physical capital (w_{2it}) is the ratio of non-interest expenses to fixed assets. The price of borrowed funds (w_{3it}) is defined as the ratio of interest expense to total assets.

Equation (7) is estimated for each country *j*. Following Berger and Mester [1997], Dietsch and Lozano-Vivas [2000] and Dwumfour et al. [2022], the study includes some bank-, and country-specific environmental variables that may account for technological differences among the banks. Hence, the study follows the approach of Battese and Coelli [1995] which allows for the inclusion of environmental variables such that the truncation of the distribution of the inefficiency term is of the form $N(\mu_{it},\ \sigma^2)$ where $\mu_{it} = z_{it}\delta$, with the z_{it} representing the environmental variables and the vector of unobserved scalar parameters are represented by δ. The bank-specific environmental variables included in the country-level frontier estimations are bank profitability (measured by return on average assets) and the equity ratio. Again, as noted by Fries and Taci [2005], it is essential to incorporate country-level environmental variables to account for heterogeneity in cross-country technology efficiencies and variations in service quality. Failing to account for these would assume that bank efficiency is purely driven by managerial decisions on the composition and scale of inputs. Hence, in the Metafrontier analysis, we include GDP per capita growth and inflation.[5]

[4] We also estimate efficiency scores using different metafrontier based on the regional classification, the income groups based on the World Bank classification, development status (developed *vs* developing) as well as comparing euro area countries under the (ECB *vs* non-ECB members).

[5] The *predict bc* option of the *sfpanel STATA* package is used to generate the efficiency scores (both $C\hat{E}_{it}^{j}$ and

4.3 Model specification: monetary policy on bank stability

We estimate the impact of monetary policy on banking stability following Equation (8) below:[6]

$$ZScore_{i,j,t} = \alpha_i + \mu_j + \delta_t + \alpha_1 Policy^{z}_{j,t-1} + \alpha_2 \,\text{CostEff}_{i,j,t} + \boldsymbol{\gamma}\mathbf{X}^{bank}_{i,j,t} + \boldsymbol{\phi}\mathbf{X}^{country}_{j,t} + \varepsilon_{i,j,t} \quad (8)$$

where i,j,t represents bank i in country j at time t. $Policy^{z}_{j,t-1}$, the monetary policy stance variable for country j at time $t-1$ which we discuss later under this section. The equation models the Z-score for bank i in country j at time t, denoted as $ZScore_{i,j,t}$, which serves as the dependent variable measuring a bank's distance to insolvency or financial stability which is widely used in the banking stability literature [Roy, 1952, Uhde and Heimeshoff, 2009, Houston et al., 2010, Demirgüç-Kunt and Huizinga, 2010, Carretta et al., 2015]. This is calculated as follows:

$$Z_{i,j,t} = \frac{ROAA_{i,j,t} + \frac{E_{i,j,t}}{A_{i,j,t}}}{\sigma(ROAA)_{i,j}} \quad (9)$$

where $ROAA_{i,j,t}$ is the return on average assets for bank i in country j at time t, $E_{i,j,t}$ is equity, $A_{i,j,t}$ is total assets, and $\sigma(ROAA)_{i,j}$ is the standard deviation of $ROAA$ for bank i in country j.[7] A higher Z-score denotes lower default probability.

Again, the right-hand side includes α_i, the bank-specific fixed effect capturing time-invariant heterogeneity across individual banks; μ_j is the country-specific fixed effect that accounts for time-invariant country differences and δ_t is the time-specific fixed effect capturing common shocks across all entities at time t. α_1, the coefficient estimating the impact of the policy change on the Z-score. α_2, the coefficient quantifying the effect of cost efficiency on bank Z-score; $\text{CostEff}_{i,j,t}$ is the cost efficiency measure for bank i in country j at time t (derived from stochastic metafrontier frontier analysis). It quantifies how changes in bank cost efficiency affect their stability. Higher values of CostEff indicate better efficiency in minimising costs for a given level of output or loans. $\boldsymbol{\gamma}$, a vector of coefficients for bank-level controls, $\mathbf{X}^{bank}_{i,j,t}$. These are bank-specific covariates, including bank size, liquidity and asset structure, as used in the literature [Dwumfour et al., 2022]. $\boldsymbol{\phi}$ is a vector of coefficients for country-level controls, $\mathbf{X}^{country}_{j,t}$. Also following the literature [Dwumfour et al., 2022], we include a measure of bank asset concentration ratio based on the top 3 banks. The country controls also include macroeconomic factors, such as GDP growth and inflation, which vary by country and over time. $\varepsilon_{i,j,t}$, the idiosyncratic error term representing unobserved random shocks.

4.3.1 Measurement of monetary policy stance

We construct our monetary policy stance indicator, $Policy_{j,t-1}$, at the country–year level and merge them into the bank-level panel. This is standardised within countries so that regression coefficients reflect the effect of a one-standard-deviation monetary tightening. Following the banking and monetary transmission literature [Altunbas et al., 2014, Jiménez et al., 2014, Altunbas et al., 2018, Lim et al., 2023], we employ the one-year lag of our policy measure to address reverse causality and to account for the delayed transmission of monetary policy to bank risk-taking and financial stability. Higher values of the monetary policy indicator correspond to a tighter policy stance. We describe the construction of our policy measure based on the Taylor rule below:

$M\hat{C}E_{jit}$). This option estimates the cost efficiency scores following Battese and Coelli [1988] via $E\ \{exp\,(\varepsilon)\}$.

[6] We also include one lag of all controls in order to mitigate any possibility of endogeneity.

[7] We also use 3-year rolling standard deviation of ROAA as robustness. The results shown in Appendix E remain robust.

Taylor-Rule deviations. To proxy for discretionary deviations from systematic monetary policy, we calculate deviations of the observed policy rate from a benchmark Taylor-type rule [Taylor, 1993]. Given the annual frequency of the data and the broad cross-country coverage, we adopt coefficients of 0.5 on both the inflation gap and the output gap, which is standard in annual cross-country banking studies [Clarida et al., 2000, Altunbas et al., 2018].

Output gap. The output gap is estimated country-by-country using the Hodrick-Prescott (HP) filter [Hodrick and Prescott, 1997] applied to the logarithm of real GDP with smoothing parameter $\lambda = 6.25$ as suggested by Ravn and Uhlig [2002]:

$$y_{j,t} = 100 \times \Big(\ln Y_{j,t} - \ln Y^*_{j,t}\Big), \tag{10}$$

where $Y^*_{j,t}$ denotes HP-filtered potential output. Positive values indicate that actual output exceeds potential.

Inflation gap. The inflation gap is defined as:

$$\pi^{\text{gap}}_{j,t} = \begin{cases} \pi_{j,t} - \pi^*_j, & \text{if country } j \text{ operates an inflation-targeting regime,} \\ \pi_{j,t} - \widetilde{\pi}_{j,t}, & \text{otherwise,} \end{cases} \tag{11}$$

where $\pi_{j,t}$ is annual CPI inflation, π^*_j is the official midpoint inflation target, and $\widetilde{\pi}_{j,t}$ is the HP-filtered inflation trend ($\lambda = 6.25$) for non-inflation-targeting countries [Altunbas et al., 2018, Lim et al., 2023]. For non-targeting regimes, the inflation gap captures cyclical deviations of inflation from its medium-run trend.

Time-varying neutral real rate. A country-specific time-varying neutral real rate $r^*_{j,t}$ is proxied by applying the HP filter ($\lambda = 6.25$) to the ex-post real policy rate:

$$r^*_{j,t} = \text{HP}\Big(i_{j,t} - \pi_{j,t}\Big), \tag{12}$$

which serves as a reduced-form proxy for medium-run neutral monetary conditions [Lim et al., 2023].

Taylor rule and monetary policy measure. The implied neutral nominal policy rate is given by:

$$i^*_{j,t} = r^*_{j,t} + \pi_{j,t} + 0.5\,\pi^{\text{gap}}_{j,t} + 0.5\,y_{j,t}. \tag{13}$$

Our headline monetary policy stance indicator is the deviation of the observed policy rate from this benchmark:

$$\text{Policy}_{j,t} \equiv i_{j,t} - i^*_{j,t}. \tag{14}$$

This deviation is standardised within countries to obtain $\text{Policy}^{\text{z}}_{j,t}$, and the one-year lag $\text{Policy}^{\text{z}}_{j,t-1}$ is used in our regressions. We also use Taylor-rule deviation, using only official inflation targets (with the inflation gap set to zero for non-targeting countries), as an alternative to provide further robustness.

Figure 1 summarises the cross-country distribution of our hybrid monetary-policy stance — the first quintile (P20), median, and the last quintile (P80) — which highlights both major global episodes and dispersion in policy settings. The series exhibits pronounced accommodation around downturns, with a sharp easing during the early-2000s slowdown and a deep, broad-based policy response during the Global Financial Crisis (GFC), where the lower tail (P20) falls substantially below the median, indicating that a non-trivial subset of countries adopted exceptionally expansionary stances. The subsequent rebound and oscillation through the 2010s are consistent with gradual normalisation and intermittent renewed accommodation, with the upper tail (P80) remaining persistently above the median during several years—suggestive of heterogeneous normalisation speeds across countries. The COVID-19 episode is characterised by an abrupt, outsized surge in accommodation at the top of the distribution (P80 spiking well above the median), followed by a rapid reversal into a markedly restrictive phase in 2021–2022, with the lower tail dropping sharply and the dispersion widening; this pattern is consistent with a synchronized global pivot from pandemic-era support toward disinflationary tightening, while the cross-sectional spread indicates substantial heterogeneity in timing and intensity across countries.

Figure 1. Monetary policy stance, hybrid Taylor rule

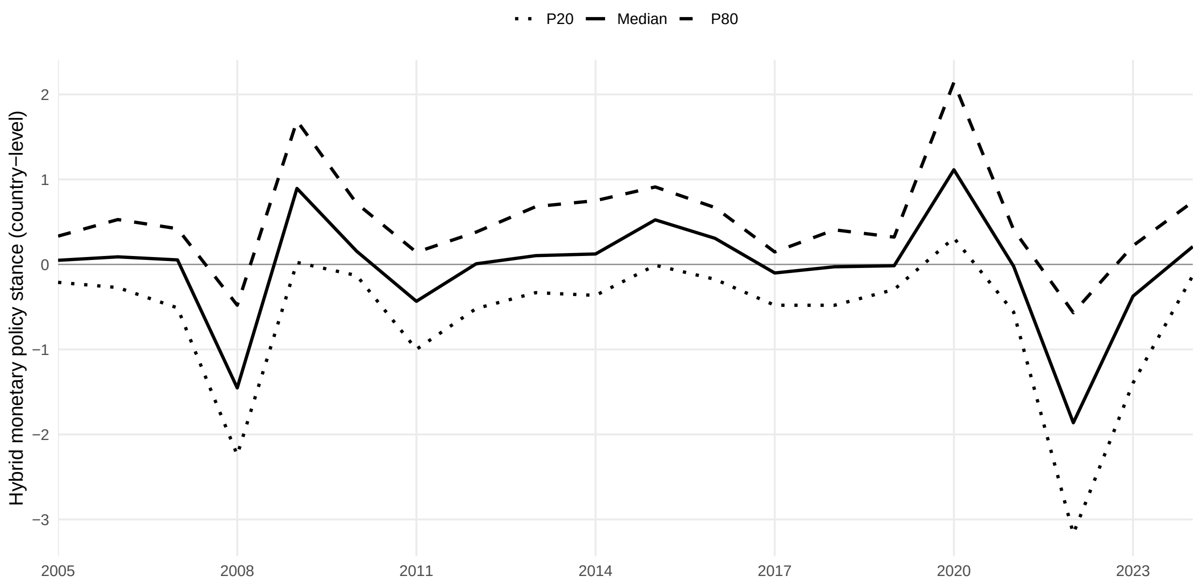

Note: The figure reports the median, the first quintile (P20), and the last quintile (P80) of the monetary policy stance measures estimated following Equation (14). The monetary policy stance measure is given by the difference between the the observed policy rate and the neutral policy rate, calculated using the HP filter with a smoothing parameter $\lambda = 6.25$.

4.4 Monetary policy, cost efficiency, and bank Stability

To examine whether the transmission of monetary policy to bank stability depends on banks' cost efficiency, we embed the interaction between the monetary policy stance and cost efficiency within a local projection (LP) framework. For each forecast horizon $h = 0, 1, \ldots, H$, we estimate:

$$\begin{aligned} ZScore_{i,j,t+h} = \alpha_i + \mu_j + \delta_t + \beta_{1,h}\, \text{Policy}^z_{j,t-1} + \beta_{2,h}\; \text{CostEff}_{i,j,t} \\ + \beta_{3,h}\Big(\text{Policy}^z_{j,t-1} \times \text{CostEff}_{i,j,t}\Big) + \boldsymbol{\gamma}'_h \mathbf{X}^{bank}_{i,j,t} + \boldsymbol{\phi}'_h \mathbf{X}^{country}_{j,t} + \varepsilon_{i,j,t+h}. \end{aligned} \tag{15}$$

Rather than interpreting the horizon-specific coefficients in isolation, inference is based on the impulse response functions implied by Equation (15). Specifically, the dynamic response of bank stability to a one-standard-deviation monetary policy tightening at horizon h, conditional on a given level of cost efficiency, is:

$$\text{IRF}_h(\text{CostEff}) = \beta_{1,h} + \beta_{3,h}\, \text{CostEff}\,. \tag{16}$$

We evaluate these IRFs at selected percentiles of the cost-efficiency distribution (25th, 50th, and 75th percentiles) and trace their evolution over the forecast horizon. This approach allows us to assess how both the magnitude and the smoothness of the response of bank stability to monetary policy shocks vary systematically with efficiency, without imposing dynamic restrictions on the underlying adjustment process.

A more positive IRF at higher cost-efficiency percentiles indicates that efficient banks exhibit greater resilience to monetary tightening, consistent with stronger buffers, superior screening, or enhanced pricing capacity. Conversely, a more negative IRF among high-efficiency banks is consistent with re-optimisation or risk-taking channels, whereby efficient institutions adjust balance sheets more aggressively following policy shocks.

In the baseline specification, we do not condition on lagged bank stability, allowing the LPs to capture the full dynamic response of Z-score to the monetary shock. As a robustness check, we augment the specification with a lagged Z-score.

5 EMPIRICAL RESULTS

Here, we discuss the results of the stochastic metafrontier analysis, based on Equation (7), and the efficiency scores generated using Equation (6). We then proceed to discuss our main results, which examine the impact of monetary policy on banking stability, with a focus on the role of bank efficiency. In this analysis, we also examine potential heterogeneities in our findings based on subsamples of our data. Specifically, we examine differences in regions based on the World Bank classifications. These are East & Asia and Pacific (EAP), Europe & Central Asia (ECA), Latin America & Caribbean (LAC), Middle East & North Africa (MENA), North America (NA), South Asia (SA) and Sub-Saharan Africa (SSA). The second grouping is based on the income groups as classified by the World Bank. These are High-Income Countries (HIC), Low-Income Countries (LIC), Lower-Middle-Income Countries (LMIC), and Upper-Middle-Income Countries (UMIC). We also categorise countries based on their development status, specifically as developed and developing countries. Lastly, given that banks under the European Central Bank (ECB) are the the only monetary Union with common currency in our sample,[8] we group the countries into ECB members and non-ECB members, given the potential dynamics of operating under a monetary union.

5.1 Results of stochastic metafrontier analysis

We first discuss the results from the stochastic metafrontier (SMF) analysis. The results include those related to the global frontier and the income group, as well as the regional, development status and ECB member countries sub-frontier group analysis. The results of the main global and the income group results are shown in Table 2 while the results of the other sub-groups are reported in Tables B.1–B.3 in the Online Appendix.

The global SMF estimates in Table 2 indicate a well-behaved translog cost structure with strong statistical significance. The output elasticity component is dominated by the positive and highly significant coefficient on ln(Loans), suggesting that higher loan output is associated with higher total cost, as expected in a cost frontier setting. The first-order price effects also conform to economic intuition under the normalisation by $w1$: $\ln(w2/w1)$ enters positively while $\ln(w3/w1)$ enters negatively, consistent with relative-price substitutions embedded in the translog specification. The interaction and squared terms further support the notion of curvature

[8]The only other monetary bloc to be considered in our sample is the West African Economic and Monetary Union (WAEMU/UEMOA), a monetary bloc where member nations share the West African CFA Franc (XOF) as a common currency, managed by the regional central bank (BCEAO). However, our sample only has one of the member countries, Côte d'Ivoire; hence, we do not have a separate group for this bloc.

and flexibility, indicating that marginal cost responses vary with output scale and input price configurations, rather than remaining constant.

On the inefficiency side, the μ-equation shows that GDP p.c. growth is associated with lower cost inefficiency (negative and significant), consistent with the view that stronger macroeconomic performance improves banking sector efficiency through demand expansion, balance-sheet strength, and scale economies. Inflation appears to be positively associated with cost inefficiency at the global level, which is plausible given the operational and pricing frictions that inflation can introduce into intermediation and cost management.

Turning to the group-level results, the income splits indicate meaningful heterogeneity in technology and the macro-inefficiency relationship. In particular, the effect of GDP p.c. growth is strongly negative in all four income groups, with particularly large magnitudes for LICs, indicating that macroeconomic improvements may translate into larger efficiency gains where structural constraints are more binding. The Usigma and Vsigma estimates are also uniformly significant in income groups, suggesting that the decomposition of the composite error into inefficiency and noise is empirically relevant across the samples, thereby supporting the appropriateness of the SMF framework over a pooled homogeneous-technology alternative.

The other group-level results presented in Tables B.1–B.3 in the Online Appendix reinforce this pattern. For instance, the regional SMF results presented in Table B.1 reinforce the argument that banks operate under heterogeneous technologies across geographic clusters. The results imply differences in cost structures and the way input-price pressures translate into total cost. For example, the differences in the output coefficients and the interaction terms involving ln(Loans) and input prices signal meaningful cross-regional differences in scale and substitution patterns. The μ-equation also exhibits notable heterogeneity: GDP p.c. growth remains predominantly efficiency-enhancing across regions (negative coefficients in most cases), while the inflation effect is more mixed in sign and magnitude. The ECB split further suggests that the cost structure and macro-inefficiency channel differ materially between ECB and non-ECB systems, which is consistent with differing regulatory, monetary, and competitive environments.

The LR test results in Table 3 with Kodde–Palm critical values provide a formal complement to the coefficient-based narrative. Across the global sample and each classification group (regions, income groups, development status, and ECB membership), the reported *LR_stat* values are substantially larger than the relevant KP critical thresholds at 10%, 5%, and 1%. This implies a strong rejection of the null of no cost inefficiency across all panels, confirming that the one-sided inefficiency component is empirically non-trivial. More importantly for the paper's organising theme, the consistent rejections across panels support the interpretation that technology (and thus the attainable cost frontier) differs meaningfully across banks from different countries, regions, income groups and development status, validating the use of a stochastic metafrontier approach to summarise within-country efficiency and between-group technology gaps in a single coherent framework. Overall, these results highlight the technological and environmental heterogeneities that drive metafrontier gaps.

Table 2. Stochastic metafrontier results – global and income groups

	Global	HIC	LIC	LIMC	UMIC
Frontier					
ln(Loans)	0.794***	0.771***	0.721***	0.862***	0.803***
	(0.003)	(0.004)	(0.011)	(0.009)	(0.015)
ln(w2/w1)	0.062***	0.011	0.508***	0.150***	0.166***
	(0.007)	(0.007)	(0.043)	(0.019)	(0.024)
ln(w3/w1)	-0.025***	0.010	-0.281***	-0.188***	-0.115***
	(0.006)	(0.007)	(0.038)	(0.024)	(0.029)
ln(Loans) $\times$ ln(w2/w1)	0.000	0.005***	-0.026***	-0.007***	-0.007***
	(0.000)	(0.000)	(0.001)	(0.001)	(0.001)
ln(Loans) $\times$ ln(w3/w1)	-0.001***	-0.000	0.005***	0.004***	0.000
	(0.000)	(0.000)	(0.002)	(0.001)	(0.002)
ln(w2/w1) $\times$ ln(w3/w1)	0.007***	-0.002***	0.023***	0.014***	0.036***
	(0.001)	(0.001)	(0.006)	(0.002)	(0.003)
0.5[ln(Loans)]^2	0.012***	0.012***	0.024***	0.008***	0.011***
	(0.000)	(0.000)	(0.001)	(0.000)	(0.001)
0.5[ln(w2/w1)]^2	0.000	-0.009***	-0.009	0.009***	-0.005**
	(0.001)	(0.001)	(0.008)	(0.002)	(0.002)
0.5[ln(w3/w1)]^2	0.002***	0.005***	0.028***	-0.023***	0.037***
	(0.001)	(0.001)	(0.007)	(0.003)	(0.003)
Constant	1.889***	2.207***	1.442***	1.536***	1.949***
	(0.034)	(0.038)	(0.145)	(0.099)	(0.145)
Mu					
GDP p.c. growth	-0.071***	-0.070***	-0.435***	-0.066***	-0.076***
	(0.002)	(0.005)	(0.027)	(0.003)	(0.005)
Inflation	0.002***	-0.593***	0.032***	0.002***	-0.021***
	(0.000)	(0.026)	(0.005)	(0.000)	(0.004)
Usigma					
Constant	-0.408***	-0.138***	-0.480***	-0.405***	-0.112***
	(0.008)	(0.024)	(0.071)	(0.022)	(0.028)
Vsigma					
Constant	-4.321***	-4.750***	-3.786***	-3.258***	-2.890***
	(0.021)	(0.022)	(0.099)	(0.046)	(0.049)
Observations	49,859	29,744	1,779	9,922	8,136
Log Likelihood	-28640.004	-8945.604	-459.696	-6523.118	-6325.623
Wald χ^2	7886434.96	5436008.99	578517.87	627949.26	436108.16

Note: Standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3. Likelihood ratio tests (Kodde–Palm critical values)

Item	N	LL_U	LL_R	LR_stat	KP10	KP5	KP1	Reject–10%	Reject–5%	Reject–1%	p_KP
Global	49,859	-2.86e+04	-4.20e+04	26626.154	1.642	2.706	5.412	1	1	1	0.000
East Asia & Pacific	6,872	-4468.153	-4992.373	1048.440	1.642	2.706	5.412	1	1	1	0.000
Europe & Central Asia	12,392	-7557.901	-9083.080	3050.357	1.642	2.706	5.412	1	1	1	0.000
Latin America & Caribbean	2,826	-1678.136	-2026.449	696.626	1.642	2.706	5.412	1	1	1	0.000
Middle East & North Africa	4,026	-3826.647	-4239.617	825.941	1.642	2.706	5.412	1	1	1	0.000
North America	16,756	33800.153	32592.898	2414.510	1.642	2.706	5.412	1	1	1	0.000
South Asia	5,414	-807.326	-1872.228	2129.804	1.642	2.706	5.412	1	1	1	0.000
Sub-Saharan Africa	1,806	-1348.248	-1560.066	423.636	1.642	2.706	5.412	1	1	1	0.000
High Income	29,744	-8945.604	-1.89e+04	19906.875	1.642	2.706	5.412	1	1	1	0.000
Low Income	1,779	-459.696	-983.758	1048.124	1.642	2.706	5.412	1	1	1	0.000
Lower Middle	9,922	-6523.118	-7963.633	2881.031	1.642	2.706	5.412	1	1	1	0.000
Upper Middle	8,414	-6514.501	-7634.378	2239.754	1.642	2.706	5.412	1	1	1	0.000
Developed	30,243	-9009.575	-1.63e+04	14594.386	1.642	2.706	5.412	1	1	1	0.000
Developing	19,616	-1.57e+04	-2.01e+04	8806.766	1.642	2.706	5.412	1	1	1	0.000
Non-ECB	44,921	-2.40e+04	-3.69e+04	25815.727	1.642	2.706	5.412	1	1	1	0.000
ECB	4,938	-2782.731	-3592.649	1619.836	1.642	2.706	5.412	1	1	1	0.000

Reject: 1 to reject null and 0 fail to reject.

5.1.1 Efficiency scores

We move to discuss the cost efficiency scores generated following Equation (7). We present the map of efficiency scores to illustrate the differences between countries and regions in terms of efficiency, as well as the coverage of our dataset. These are shown in Figures 2 and 3. Figure 2 indicates a clear clustering by level of financial development and institutional capacity. High efficiency tends to be concentrated in advanced and well-supervised banking systems (e.g., North America, Northern/Western Europe, and high-income Asia such as Singapore), consistent with stronger managerial practices, tighter cost discipline, and deeper financial infrastructures that support scale and process standardisation. By contrast, lower efficiency values are more prevalent in parts of Latin America and Sub-Saharan Africa, as well as in several emerging/frontier systems, reflecting greater operational frictions (e.g., higher overheads, weaker intermediation technology, and more volatile macro-financial environments). The dispersion is economically meaningful, as the cross-country range is wide (from roughly 0.17 to 0.60), which motivates our use of a metafrontier approach to ensure that the cross-country efficiency signal captures managerial/operational performance net of technology heterogeneity, rather than simply reflecting differences in country-specific banking technologies.

The trend of regional medians of metafrontier cost efficiency in Figure 3, reveal level-and-trend heterogeneity across the global banking landscape. North America consistently exhibit the highest median efficiency levels (typically around 0.5–0.65), followed by South Asia. East Asia & Pacific and Europe & Central Asia occupy an intermediate range (roughly 0.38–0.48 over most years). By contrast, Latin America & the Caribbean and Sub-Saharan Africa tend to post lower medians (often around 0.30–0.37), indicating systematically larger cost inefficiencies relative to the global best-practice technology. Importantly, aside North America and South Asia where cost efficiency has seen some gradual upward drift consistent with longer-run diffusion of managerial practices and intermediation technologies, while the other regions have seen relatively stable cost efficiency over the period showing persistent cross-region gaps.

Figure 2. Global metafrontier cost efficiency scores

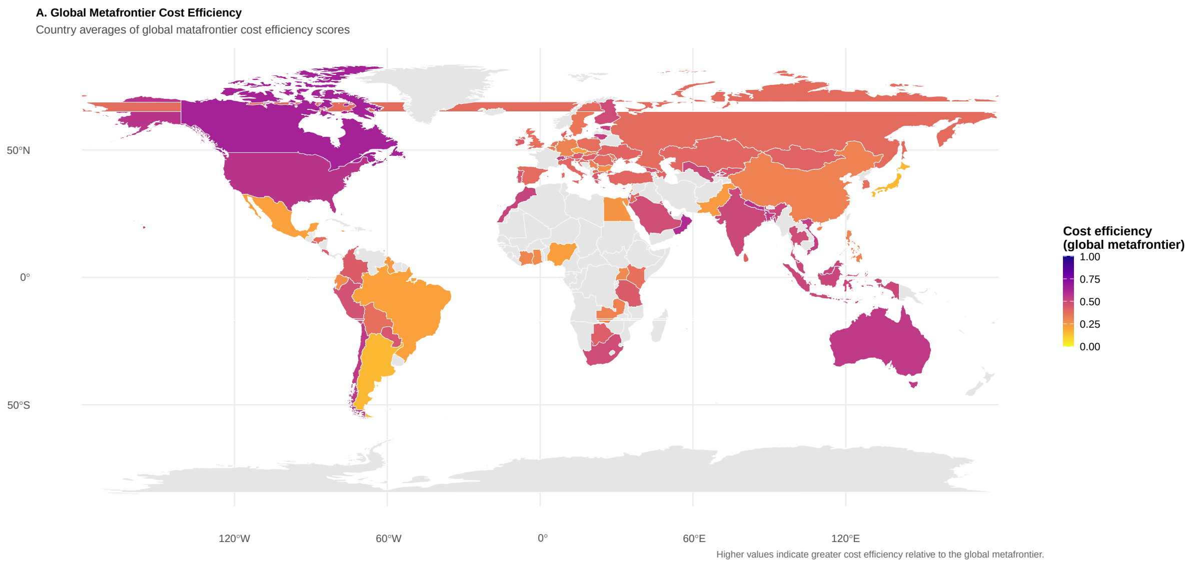

Note: Grey areas represent missing data.

Figure 3. Trend of median metafrontier cost efficiency by region

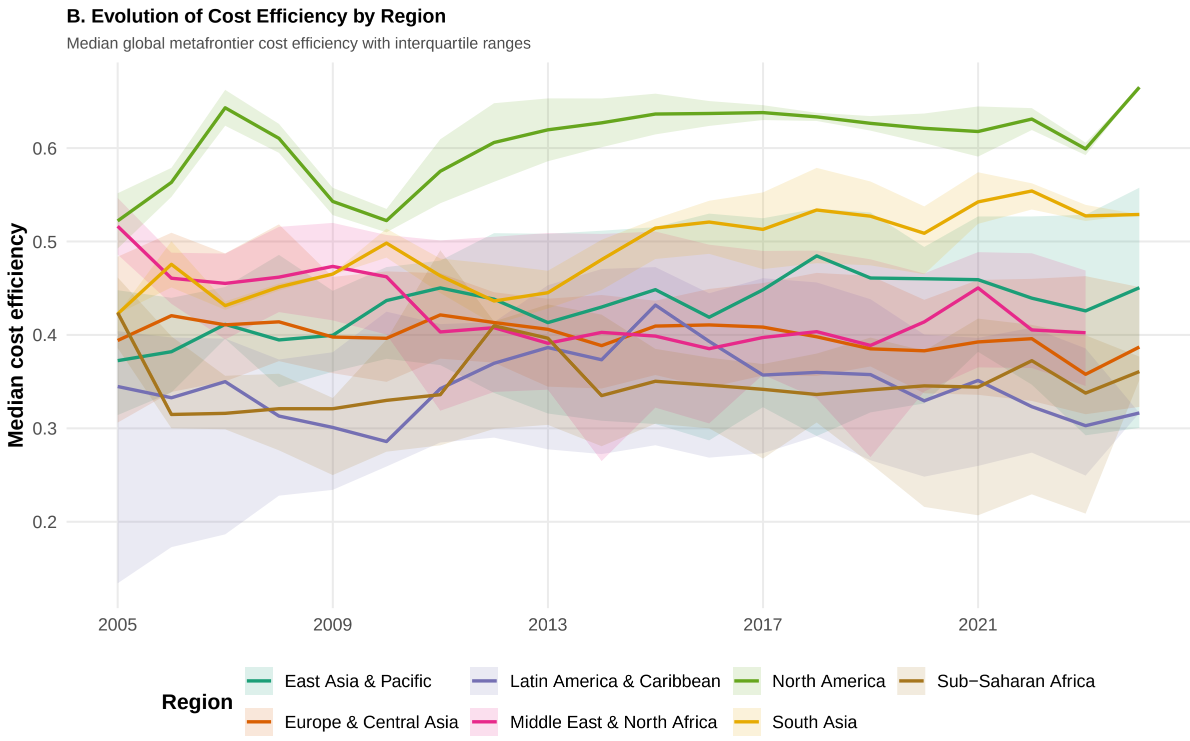

5.2 Monetary policy on banking stability

Table 4 presents the core empirical result of the paper: monetary tightening is associated with higher bank stability, measured by the Z-score, after absorbing a rich set of bank fixed effects, country fixed effects, and common time shocks. Across both measures of policy stance — the hybrid and the official indicator — the estimated coefficient on $\text{Policy}^{z}_{j,t-1}$ is positive and highly statistically significant. This pattern is robust to alternative timing conventions: the contemporaneous specification (Models 1 and 3) and the fully lagged specification (Models 2 and 4) deliver qualitatively identical conclusions, with the lagged models yielding somewhat larger point estimates. Quantitatively, this effect is consequential. For instance, a one standard deviation increase in the policy stance (0.95) leads to a 0.14 (Model 1) increase in Z-score. The baseline evidence, therefore, supports the first pillar of our central hypothesis: at least over the average policy cycle in the sample, tightening shocks tend to strengthen bank stability rather than weaken it.

The role of cost efficiency emerges as the second pillar of the paper's narrative. The coefficient of cost-efficiency is positive and statistically significant in both the contemporaneous and lagged specifications. This suggests that more efficient banks are systematically more stable, consistent with an interpretation of efficiency based on "operational discipline" or "risk management capacity" as discussed earlier [Berger and DeYoung, 1997]. When all predictors are lagged (Models 2 and 4), the cost-efficiency coefficient attenuates but remains statistically significant, which is informative. Cost efficiency is typically persistent, may be partially absorbed by bank fixed effects, and its marginal contribution is more difficult to isolate when measured in the presence of noise and dynamics. This motivates the paper's emphasis on dynamic responses and heterogeneity: efficiency is not only a level shifter of stability, but also a conditioning state that shapes how stability responds over time to a common monetary tightening shock—an implication that the subsequent local-projection IRFs are designed to quantify.

The results of the remaining covariates are consistent with literature. Liquidity and asset structure are positively related to stability, consistent with balance-sheet resilience and portfolio composition buffering shocks. In contrast, larger banks display lower stability in this specification, which is consistent with greater risk-taking capacity (moral hazard problem), complexity, or thinner capital buffers once fixed effects are controlled for [Uhde and Heimeshoff, 2009]. On the macro side, stronger GDP growth and higher institutional quality are robustly stabilising, while inflation is stabilising only in the contemporaneous models and becomes insignificant under lags, consistent with inflation capturing short-run nominal effects that are less predictive once dynamics are accounted for. Taken together, these results motivate the paper's organising proposition: monetary tightening can be stabilising on impact, but the durability and distribution of this effect depend on bank-level capacity—proxied by cost efficiency—which we examine directly in the dynamic, state-contingent local-projection analysis that we discuss subsequently.

Table 4. Impact of monetary policy on banking stability

Policy variable:	Hybrid		Official	
Model:	(1)	(2)	(3)	(4)
$\text{Policy}^{z}_{j,t-1}$	0.1391***	0.1159***	0.1736***	0.1544***
	(0.0389)	(0.0374)	(0.0395)	(0.0381)
Cost efficiency	4.301***	2.621***	4.297***	2.610***
	(0.8219)	(0.7868)	(0.8221)	(0.7872)
Bank liquidity	0.0298***	0.0161***	0.0298***	0.0161***
	(0.0030)	(0.0022)	(0.0030)	(0.0022)
Size	-2.754***	-1.855***	-2.752***	-1.852***
	(0.1980)	(0.1781)	(0.1981)	(0.1781)
Asset structure	0.1087***	0.0576	0.1089***	0.0575
	(0.0405)	(0.0384)	(0.0404)	(0.0383)
Bank Concentration	0.0272**	0.0440***	0.0275**	0.0444***
	(0.0109)	(0.0121)	(0.0109)	(0.0121)
GDP growth	0.0413***	0.0352**	0.0382**	0.0325**
	(0.0153)	(0.0142)	(0.0154)	(0.0142)
Inflation (CPI)	0.0362***	0.0089	0.0357***	0.0086
	(0.0116)	(0.0109)	(0.0116)	(0.0109)
Institutional Quality	2.912***	3.246***	2.849***	3.175***
	(0.7138)	(0.7253)	(0.7125)	(0.7240)
Fixed-effects				
Bank	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes
No. of Banks	3,903	3,773	3,903	3,773
N	42,519	39,170	42,519	39,170
R^2	0.95	0.96	0.95	0.96

Note: Clustered (bank level) standard errors in parentheses.
Lag 1 of all predictors in Models (2) and (4).
****p<0.01, **p<0.05, *p<0.1.*

5.3 Robustness: sub-sample/group analysis

We assess the external validity of our baseline results by allowing the effect of monetary tightening on bank stability to vary systematically across economic and institutional environments. Specifically, we estimate heterogeneous policy effects across the World Bank regions (EAP, ECA, LAC, MENA, NA, SA, and SSA), across income groups (HIC, UMIC, LMIC, LIC), across development status (developed versus developing), and across monetary union membership (ECB versus non-ECB). These splits are informative because the transmission of policy to bank balance sheets depends on the maturity and structure of financial systems, the prevalence of variable-rate credit, the depth of capital markets, the regulatory regime, and the degree to which policy shocks are amplified by macroeconomic conditions. The corresponding estimates are reported in Tables C.1–C.4 in the Online Appendix, and the patterns are robust to also using the group-specific metafrontier cost efficiency measures.[9]

Across most regions, monetary tightening is associated with higher bank stability, with the strongest effects in North America, Sub-Saharan Africa, the Middle East & North Africa, Latin America & the Caribbean and a smaller but still positive effect in East Asia & the Pacific, consistent with tighter conditions improving portfolio risk through stricter underwriting and a more conservative balance-sheet stance. The main exception is Europe & Central Asia (ECA), where the estimated effect is negative. The ECA result is plausible if the dominant transmission margin is asset-quality and valuation losses rather than improved risk selection: sharp rate increases can raise debt-service burdens for leveraged borrowers, elevate defaults among marginal firms and households, and generate mark-to-market losses on securities holdings, while competitive banking structures and exposure to slow-repricing (often fixed-rate) assets can compress profitability and slow internal capital generation in the short run. Moreover, institutional features—such as tight capital and liquidity requirements and the joint operation of macroprudential and monetary tightening—can weaken the mapping from policy rates to measured stability, and region-specific macro shocks (including energy and geopolitical disturbances) may coincide with tightening and amplify borrower distress. Taken together, the ECA evidence indicates that the net stability effect of tighter policy is regime-dependent: it is positive when tighter conditions primarily reduce risk-taking, but can turn negative when valuation and credit-loss channels dominate.

The income-group and development-status splits reinforce the baseline conclusion that monetary tightening is, on average, stabilising, while clarifying where the stabilisation is most pronounced. The positive effect is present across all income groups and appears strongest in low-income countries, with a similarly larger effect in developing economies relative to developed ones. A plausible interpretation is that in lower-income and developing settings, periods of accommodative policy may be more closely associated with rapid credit expansion, weaker screening, and higher marginal borrower risk; tightening, therefore, delivers larger improvements in portfolio quality and risk discipline. At the same time, the ECB split highlights that policy effects are not uniform across monetary regimes: we find significant stabilising effects in non-ECB countries but no statistically meaningful effect for ECB members as a group. This pattern is consistent with the notion that in monetary unions, country-level bank outcomes depend not only on the common policy rate but also on cross-country differences in financial structure, sovereign-bank linkages, and the interaction of common monetary policy with heterogeneous macro conditions; aggregation across member states can therefore attenuate average effects. Across all subsamples, cost efficiency remains positively associated with stability, underscoring that the efficiency-stability nexus is not a sample-specific artefact but rather a pervasive margin shaping banks' resilience. Overall, the subgroup evidence supports the central theme of the

[9]Tables D.1–D.8 in the Online Appendix.

paper — tightening is typically stabilising, and efficiency strengthens resilience — while also emphasising that the sign and magnitude of policy effects are state– and structure–dependent, with ECA representing an empirically and economically coherent regime in which valuation and borrower-solvency channels can dominate.

5.4 Monetary policy on banking stability (interaction with cost efficiency)

We next evaluate whether the stabilising effect of monetary tightening is systematically heterogeneous across banks with different levels of cost efficiency. The central empirical object is the interaction between the policy stance and bank-level cost efficiency. In line with our conceptual framework, this interaction is intended to capture an *efficiency-buffer* mechanism: conditional on the same monetary policy shock, more cost-efficient banks should be able to preserve stability more effectively because they manage risk and operating costs better, adjust balance-sheet policies with less disruption, and face a lower marginal cost of screening, monitoring, and internal control.

An essential point for interpretation is that the interaction term captures how this marginal policy effect changes with efficiency. A positive interaction is consistent with efficiency, amplifying the stabilising component of tightening and/or attenuating its destabilising component when tightening becomes prolonged. Importantly, in our empirical strategy, we capture non-linearity by emphasising a dynamic response: the impulse response of stability can rise initially and subsequently decline over the horizon, even when the contemporaneous regression is locally linear. In this sense, the interaction is interpreted through the entire IRF path rather than through a static curvature restriction in the policy variable.

Our local-projection IRFs shown in Figure 4 directly operationalise this logic. The left panel (i) follows Equation (15) while the right panel (ii) follows the same equation but includes one lag of the dependent variable (Z-score). From Panel (i), conditioning on cost efficiency (at the 25th, 50th, and 75th percentiles), we find that a tightening shock increases stability on impact and in the short run, but the response begins to weaken and eventually turns negative as the horizon extends. Crucially, the decline in stability is markedly smoother for high-efficiency banks (cost efficiency at 75th percentile): the post-peak deterioration is less steep, and the adjustment path is less volatile relative to low-efficiency banks (cost efficiency at 75th percentile). These results are similar in Panel (ii), where it takes after 4 years for stability to reach the negative territory for high cost-efficient banks, compared to 3 years for low cost-efficient banks. This is precisely the pattern predicted by the efficiency-buffer hypothesis: more efficient institutions do not merely start from higher stability; they also exhibit greater resilience as the tightening episode persists. In economic terms, efficiency appears to shift the effective "turning point" of the stability response outward and dampen the sensitivity of stability to prolonged policy restraint, consistent with superior risk governance and operational flexibility rather than simple margin expansion.

Figure 4. Local projections responses of banking stability (Z-score), conditional on bank cost efficiency

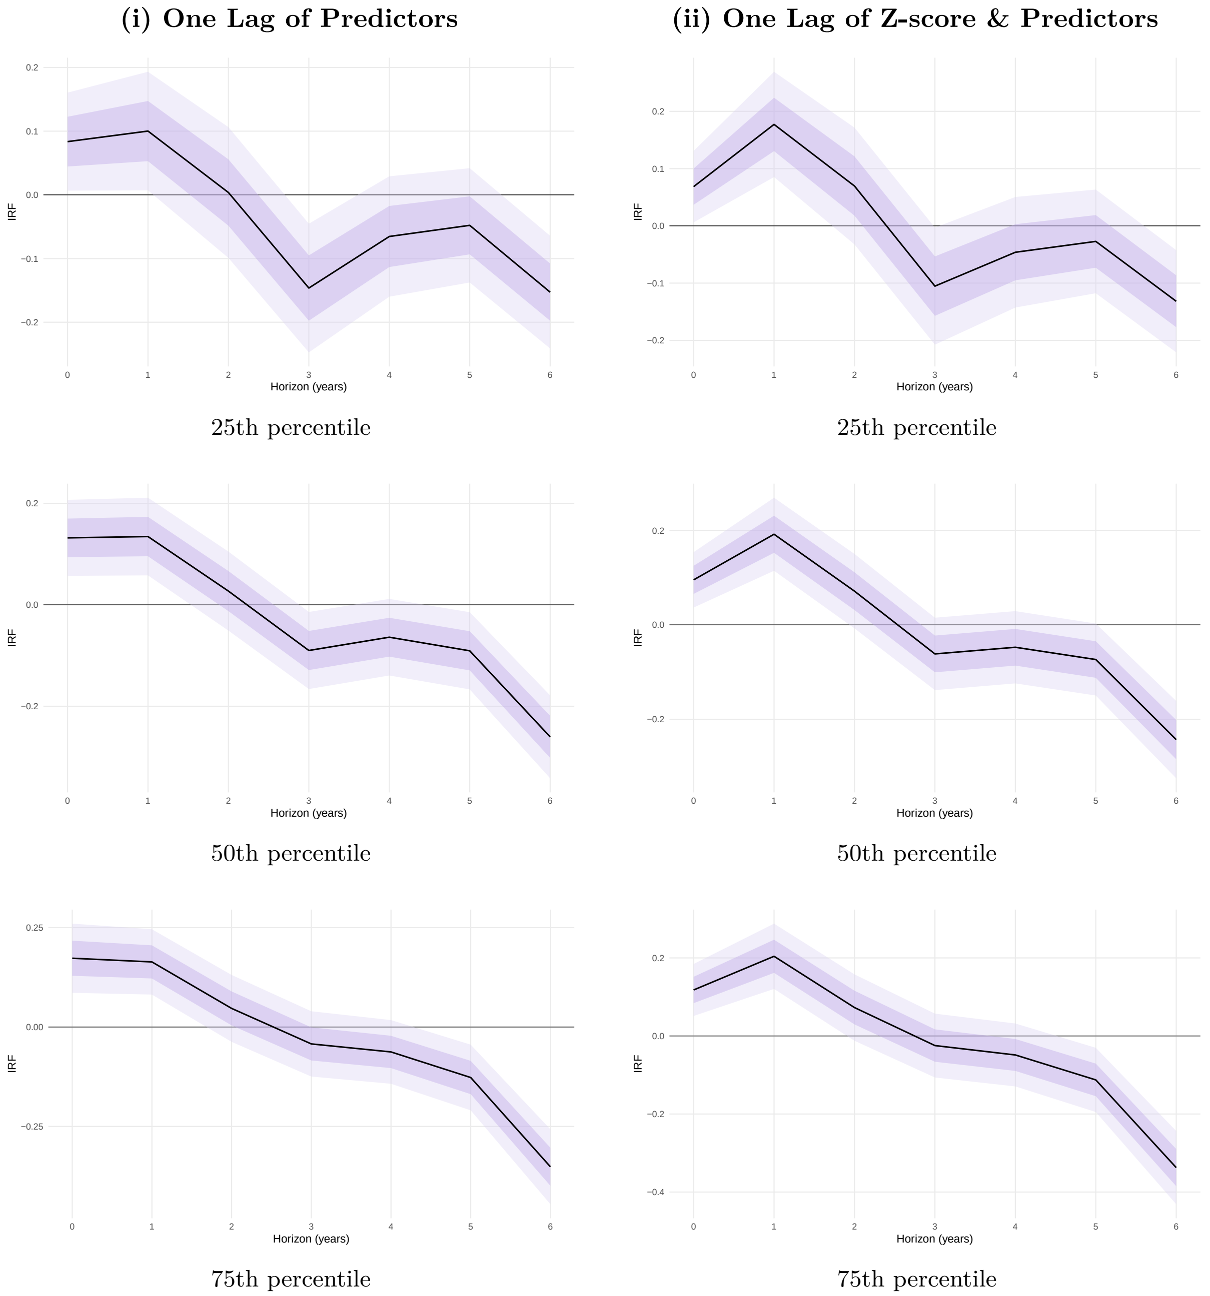

Note: The figure plots local projections responses of banking stability (Z-score) to a one-standard-deviation monetary policy shock, conditional on bank cost efficiency (evaluated at the 25th, 50th, and 75th percentiles). The lighter and darker bands represent 68% and 95% error bands, respectively. Column (i) includes 1 lag of all predictors; column (ii) includes one lag of Z-score and predictors.

5.5 Robustness: controlling for macroprudential policies

To strengthen the identification of monetary policy effects on banking stability and to further mitigate potential omitted variable bias, we control for time-varying macroprudential policies. Macroprudential tools are frequently deployed in conjunction with monetary policy to address financial stability risks; hence, it is important for us to control for these tools.

We draw on the updated iMaPP database of Alam et al. [2025], which provides detailed monthly policy action indicators for 17 macroprudential instruments across a broad set of advanced and emerging market economies. Each tightening action is coded as +1, each loosening as −1, and no change or neutral actions as 0, while the sum of these actions gives the cumulative policy decision. Following Alam et al. [2025], we aggregate these monthly indicators to the annual frequency by summing actions within each year, yielding net cumulative annual tightening (positive values) or loosening (negative values).

Our preferred macroprudential controls are: i) *Liquidity requirements (Liquidity)*: This index captures measures aimed at mitigating systemic liquidity and funding risks, including liquidity coverage ratios, net stable funding ratios, liquid asset ratios, core funding ratios, and non-currency-specific external debt restrictions. ii) *Limits on foreign exchange positions (LFX)*: This index includes limits on net or gross open foreign exchange positions, FX exposures, FX funding restrictions, and currency mismatch regulations.

These two instruments are selected because Alam et al. [2025] and Cerutti et al. [2017] identify liquidity requirements as the most frequently used macroprudential tool in advanced economies and limits on foreign exchange positions as the predominant instrument in emerging market and developing economies (EMDEs). Given the global scope of our sample, which includes both advanced and emerging markets, these controls are particularly relevant for capturing the dominant macroprudential responses in each country group.

As a further robustness check, we alternatively replace LFX with limits on loan-to-value ratios (LTV), which cover caps on residential and commercial mortgages as well as other secured loans (including "speed limits" on high-LTV lending). LTV restrictions are among the most targeted borrower-based instruments, providing an additional dimension of macroprudential tightening that operates through credit demand rather than bank liquidity or currency risk.

By including these macroprudential indices (entered contemporaneously and lagged where appropriate), we isolate the effects of monetary policy from concurrent regulatory actions aimed at financial stability. The results, reported in Tables 5 and 6, confirm that our baseline findings on the impact of monetary policy tightening on bank stability remain robust and, if anything, are strengthened (coefficients of the monetary policy variables increased slightly) after accounting for these key macroprudential policy dimensions.

The results from our LP regressions following the interactions between monetary policy and cost efficiency are also presented in Figure 5. The results are consistent with our earlier findings on the *efficiency buffer*, where high-cost-efficient banks exhibit a smoother stability response to monetary policy shocks over the horizon.

Figure 6 plots local-projection impulse responses of bank stability (Z-score) to a one-standard-deviation monetary tightening shock, conditional on bank cost efficiency (25th, 50th, and 75th percentiles), while controlling for macroprudential policies that operate through liquidity requirements and loan-to-value (LTV) limits. Three findings stand out. First, the qualitative pattern of short-run stabilisation followed by medium-run deterioration remains intact after conditioning on LTV policy. Across all efficiency percentiles and in both columns, a tightening shock raises Z-score on impact and over short horizons, but this initial improvement steadily unwinds as the horizon extends, with the response eventually turning negative at medium horizons. This persistence of sign reversal under LTV controls strengthens the interpretation that the medium-run decline is not an artefact of omitted borrower-based macroprudential tightening, but rather reflects the delayed propagation of tighter monetary conditions into realised asset-quality stress.

Second, the *efficiency-buffer mechanism* remains clearly visible. The post-peak decline is systematically smoother for banks at the 75th percentile of cost efficiency than for those at the 25th percentile: high-efficiency banks exhibit a less steep deterioration and a more gradual transition toward negative territory. Put differently, conditioning on LTV restrictions does not eliminate cross-bank heterogeneity in the stability path; instead, it reinforces the idea that operational and risk-control capacity governs how quickly (and how sharply) the delayed fragility component dominates the initial stabilisation effect.

Third, the comparison across columns (i) and (ii) indicates that the main message is robust to dynamics in the dependent variable. Adding lagged Z-score (column (ii)) changes the shape of the IRFs mechanically—by absorbing persistence in stability—but the core ordering across efficiency groups and the horizon-dependent weakening of the tightening effect remain. Overall, Figure 6 therefore corroborates the central narrative: monetary tightening tends to strengthen measured bank stability in the near term, yet this benefit is not durable, and cost efficiency materially dampens the medium-run deterioration even after accounting for borrower-side macroprudential policy via LTV limits.

Table 5. Impact of monetary policy on banking stability controlling for macroprudential policies (Liquidity and LFX) – FE results

Policy variable:	Hybrid		Official	
Model:	(1)	(2)	(3)	(4)
$\text{Policy}^{z}_{j,t-1}$	0.1367***	0.1057***	0.1724***	0.1435***
	(0.0400)	(0.0382)	(0.0406)	(0.0391)
Cost efficiency	4.239***	2.549***	4.236***	2.540***
	(0.8304)	(0.7940)	(0.8306)	(0.7944)
Bank liquidity	0.0300***	0.0164***	0.0299***	0.0164***
	(0.0030)	(0.0023)	(0.0030)	(0.0023)
Size	-2.768***	-1.863***	-2.766***	-1.861***
	(0.2016)	(0.1806)	(0.2016)	(0.1807)
Asset structure	0.1120***	0.0591	0.1121***	0.0591
	(0.0417)	(0.0393)	(0.0417)	(0.0393)
Bank Concentration	0.0251**	0.0418***	0.0254**	0.0421***
	(0.0110)	(0.0122)	(0.0110)	(0.0122)
GDP growth	0.0414**	0.0445***	0.0380**	0.0416***
	(0.0171)	(0.0154)	(0.0172)	(0.0155)
Inflation (CPI)	0.0311**	0.0043	0.0307**	0.0041
	(0.0126)	(0.0119)	(0.0126)	(0.0119)
Institutional Quality	2.656***	3.118***	2.588***	3.048***
	(0.7361)	(0.7488)	(0.7348)	(0.7475)
Macroprudential: Liquidity	-0.1781***	-0.1978***	-0.1755***	-0.1948***
	(0.0457)	(0.0449)	(0.0457)	(0.0449)
Macroprudential: LFX	0.2998***	0.3066***	0.2954***	0.3031***
	(0.1083)	(0.1081)	(0.1083)	(0.1080)
Fixed-effects				
Bank	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes
No. of Banks	3,816	3,687	3,816	3,687
N	41,690	38,426	41,690	38,426
R^2	0.95	0.96	0.95	0.96

Note: Lag 1 of all predictors in Models (2) and (4). LFX: Limits on FX positions. Clustered (bank level) standard errors in parentheses.
****p<0.01, **p<0.05, *p<0.1.*

Table 6. Impact of monetary policy on banking stability controlling for macroprudential policies (Liquidity and LTV) – FE results

Policy variable:	Hybrid		Official	
Model:	(1)	(2)	(3)	(4)
$\text{Policy}^{z}_{j,t-1}$	0.1243***	0.0942**	0.1612***	0.1332***
	(0.0399)	(0.0379)	(0.0405)	(0.0388)
Cost efficiency	4.302***	2.601***	4.298***	2.590***
	(0.8307)	(0.7938)	(0.8309)	(0.7943)
Bank liquidity	0.0300***	0.0163***	0.0300***	0.0163***
	(0.0030)	(0.0023)	(0.0030)	(0.0023)
Size	-2.773***	-1.870***	-2.772***	-1.868***
	(0.2014)	(0.1805)	(0.2014)	(0.1806)
Asset structure	0.1121***	0.0584	0.1123***	0.0584
	(0.0416)	(0.0391)	(0.0416)	(0.0390)
Bank Concentration	0.0246**	0.0413***	0.0250**	0.0416***
	(0.0110)	(0.0122)	(0.0110)	(0.0122)
GDP growth	0.0530***	0.0553***	0.0493***	0.0521***
	(0.0167)	(0.0151)	(0.0168)	(0.0152)
Inflation (CPI)	0.0367***	0.0100	0.0362***	0.0096
	(0.0125)	(0.0118)	(0.0125)	(0.0118)
Institutional Quality	2.634***	3.099***	2.567***	3.029***
	(0.7337)	(0.7471)	(0.7325)	(0.7460)
Macroprudential: Liquidity	-0.1926***	-0.2117***	-0.1897***	-0.2083***
	(0.0451)	(0.0447)	(0.0452)	(0.0447)
Macroprudential: LTV	-0.2578***	-0.2532***	-0.2508***	-0.2461***
	(0.0922)	(0.0852)	(0.0922)	(0.0852)
Fixed-effects				
Bank	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes
No. of Banks	3,816	3,687	3,816	3,687
N	41,690	38,426	41,690	38,426
R^2	0.95	0.96	0.95	0.96

Note: Lag 1 of all predictors in Models (2) and (4). LTV: Limits on Loan-to-Value Ratio
Clustered (bank level) standard errors in parentheses.
****p<0.01, **p<0.05, *p<0.1.*

Figure 5. Local Projections responses of banking stability (Z-score), controlling for macroprudential policies (Liquidity and LFX)

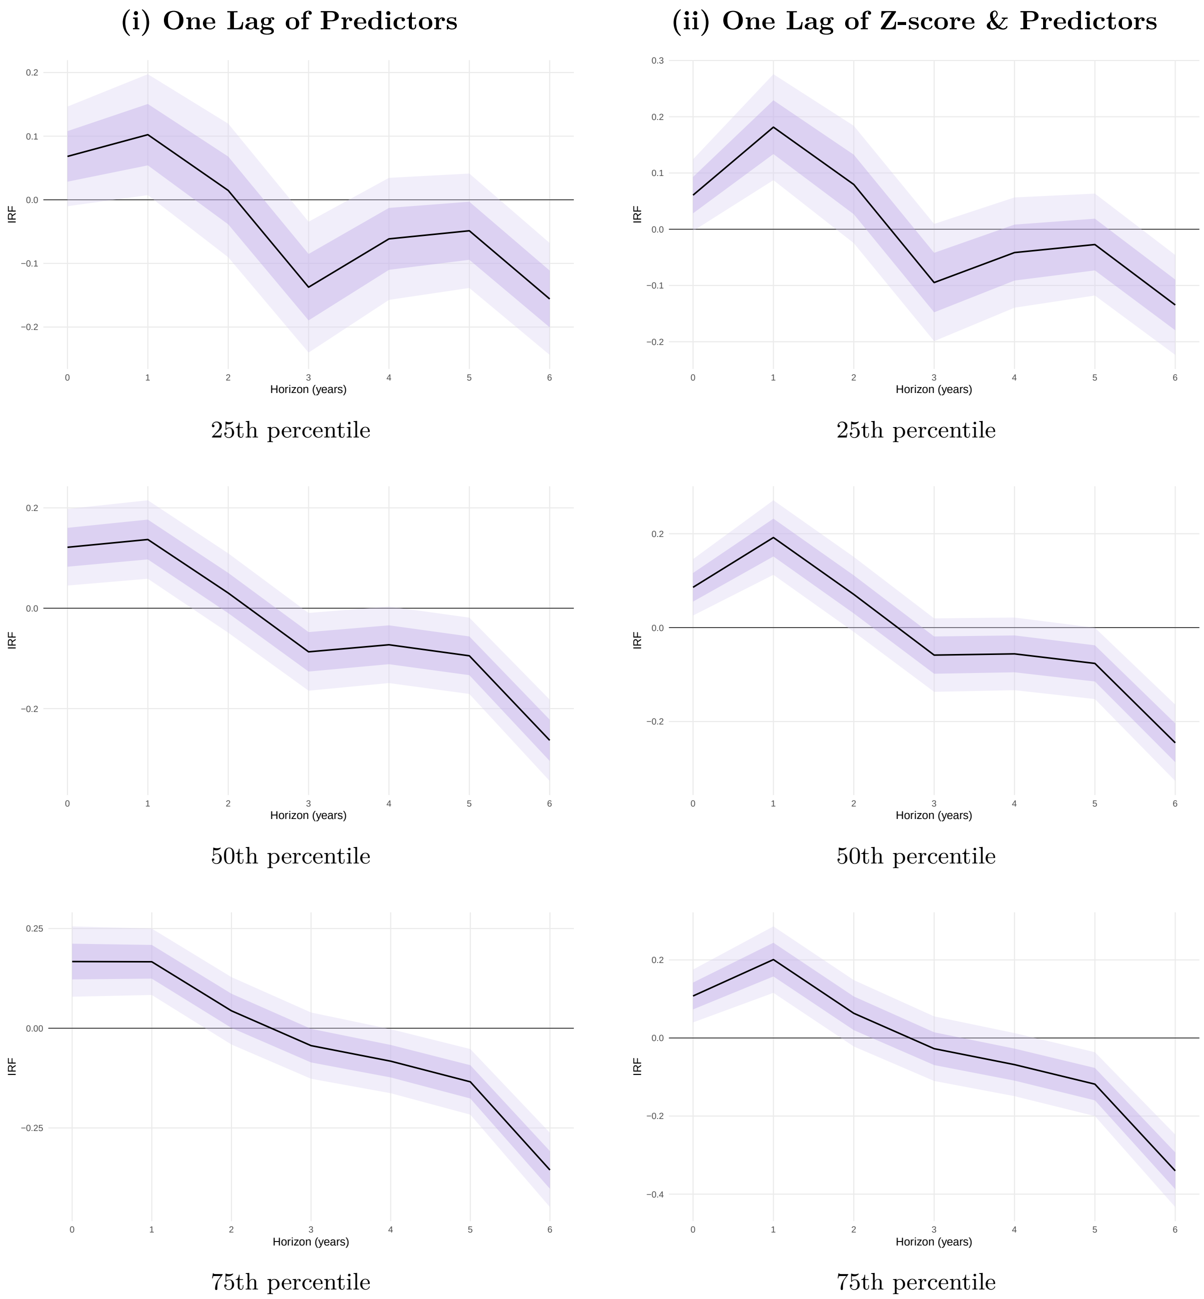

Note: The figure plots local projections responses of banking stability (Z-score) to a one-standard-deviation monetary policy shock, conditional on bank cost efficiency (evaluated at the 25th, 50th, and 75th percentiles), controlling for macroprudential policies (Liquidity and LFX). The lighter and darker bands represent 68% and 95% error bands, respectively. Column (i) includes 1 lag of all predictors; column (ii) includes one lag of Z-score and predictors.

Figure 6. Local projections responses of banking stability (Z-score), controlling for macroprudential policies (Liquidity and LTV)

(i) One Lag of Predictors | **(ii) One Lag of Z-score & Predictors**

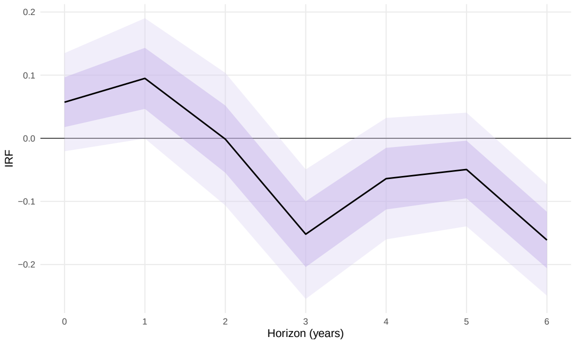

25th percentile

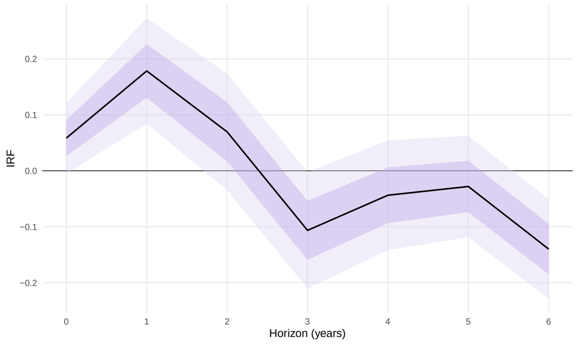

25th percentile

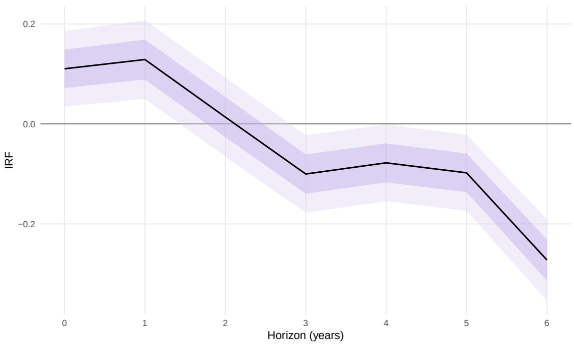

50th percentile

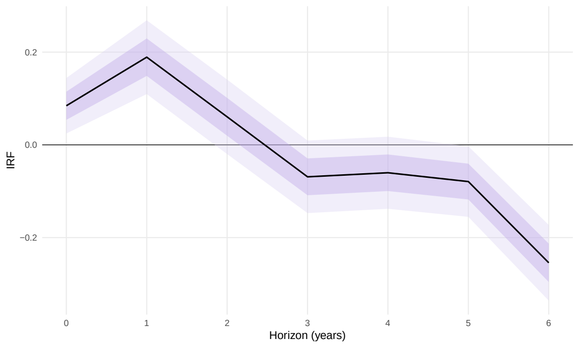

50th percentile

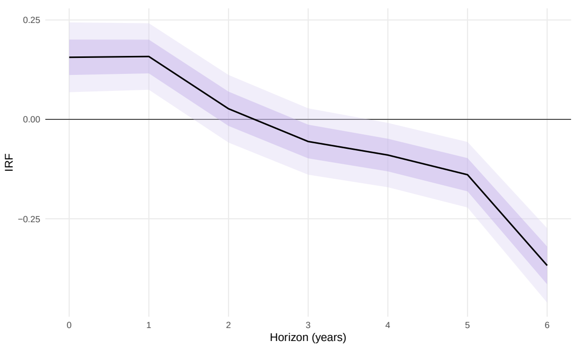

75th percentile

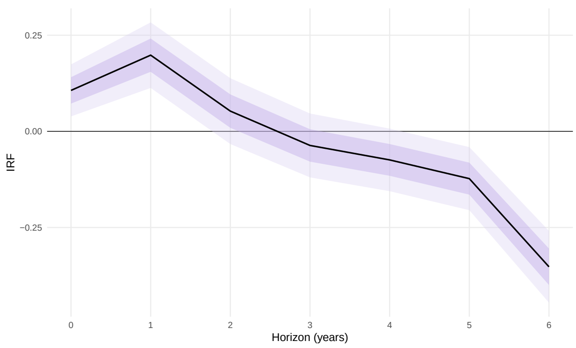

75th percentile

Note: The figure plots local projections responses of banking stability (Z-score) to a one-standard-deviation monetary policy shock, conditional on bank cost efficiency (evaluated at the 25th, 50th, and 75th percentiles), controlling for macroprudential policies (Liquidity and LTV). The lighter and darker bands represent 68% and 95% error bands, respectively. Column (i) includes 1 lag of all predictors; column (ii) includes one lag of Z-score and predictors.

5.6 Channel analysis: banking intermediation channel (bank spread)

This section examines the mechanisms by which monetary policy shocks affect bank stability and how this transmission varies with cost efficiency. We focus on the banking spread — measured by net interest margin (NIM) — as the primary channel, because it directly captures the repricing wedge between interest income on assets and interest expense on liabilities following tightening episodes. Complementary channels are analysed in sections 5.7 and H in the Online Appendix: we document: i) the response of credit growth and ii) portfolio-adjustment behaviour (loan share reallocation), which provide additional evidence on banks' balance-sheet management but are not treated as the main organising mechanism in the baseline discussion.

We begin by estimating reduced-form regressions of NIM on the monetary policy stance measure and the bank-level and macro controls, as defined in Equation (8). We then implement local projections following Equation (15) to trace the dynamic response of NIM to a policy tightening, conditional on cost-efficiency percentiles. The regression evidence in Tables 7–9 indicates that tightening compresses bank spreads on average, with the magnitude generally becoming more negative once macroprudential controls are included.

The LP evidence in Figures 7–9 shows that at short horizons ($h = 0$–3) the NIM response is negative across the 25th, 50th, and 75th percentiles of cost efficiency, consistent with asymmetric repricing: funding costs (particularly wholesale and rate-sensitive deposits) adjust rapidly, while loan yields and securities returns reprice more sluggishly due to fixed-rate exposures, contractual rigidities, interest-rate caps, and competitive constraints.

These results suggest that conditioning the stability response to policy shocks on cost efficiency reveals a systematic pattern of heterogeneity. Although differences are not always statistically significant, lower-efficiency banks exhibit a faster and smoother NIM reversal relative to higher-efficiency banks, while high-efficiency banks experience a more persistent NIM compression (lasting up to roughly three years) before recovering. Our interpretation is that, conditional on a common country-level policy shock, less efficient banks restore margins more quickly through stronger pass-through and strategic repricing (e.g., faster loan-rate adjustment, wider spreads, reallocation toward higher-yield assets, or a funding mix that reprices differently). By contrast, high-efficiency banks display a more muted NIM recovery, consistent with greater competitive discipline, a liability structure that transmits tightening more strongly into interest expense, and/or an active choice to prioritise asset quality over yield chasing [Rajan, 2006]. Importantly, this pattern aligns with our core stability results: efficient banks remain more stable despite weaker NIM recovery. Consequently, the heterogeneous stability responses to tightening are unlikely to be driven by differential margin buffers. This evidence supports the risk-management and operational discipline (screening and monitoring quality, provisioning practices, and cost control) interpretation of cost efficiency: efficient banks attenuate the translation of tightening into credit losses (and thus into Z-score deterioration), consistent with their lower marginal cost of monitoring and superior underwriting/portfolio discipline.

Table 7. Channel analysis: impact of monetary policy on banking spread (NIM) – FE results

Policy variable:	Hybrid		Official	
Model:	(1)	(2)	(3)	(4)
$\text{Policy}^{z}_{j,t-1}$	-0.0368***	-0.0344***	-0.0369***	-0.0364***
	(0.0135)	(0.0133)	(0.0131)	(0.0130)
Cost efficiency	1.516***	0.8358***	1.518***	0.8383***
	(0.3309)	(0.2908)	(0.3309)	(0.2908)
Bank liquidity	-0.0012	-0.0016	-0.0012	-0.0016
	(0.0010)	(0.0011)	(0.0010)	(0.0011)
Size	-0.1155*	-0.3321***	-0.1158*	-0.3325***
	(0.0701)	(0.0643)	(0.0701)	(0.0643)
Asset structure	-0.0551	-0.0665	-0.0551	-0.0665
	(0.0516)	(0.0450)	(0.0516)	(0.0450)
Bank Concentration	-0.0021	-0.0014	-0.0021	-0.0014
	(0.0032)	(0.0033)	(0.0032)	(0.0033)
GDP growth	0.0110**	0.0088*	0.0112**	0.0090*
	(0.0054)	(0.0049)	(0.0055)	(0.0049)
Inflation (CPI)	0.0564***	0.0525***	0.0565***	0.0526***
	(0.0097)	(0.0089)	(0.0097)	(0.0089)
Institutional Quality	-0.3016	-0.1815	-0.2962	-0.1733
	(0.2068)	(0.2049)	(0.2068)	(0.2048)
Fixed-effects				
Bank	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes
No. of Banks	3,903	3,773	3,903	3,773
N	42,516	39,168	42,516	39,168
R^2	0.84	0.85	0.84	0.85

Note: Lag 1 of all predictors in Models (2) and (4).
Clustered (bank level) standard errors in parentheses.
****p<0.01, **p<0.05, *p<0.1.*

Table 8. Channel analysis: impact of monetary policy on banking spread (NIM), controlling for macroprudential policies (Liquidity and LFX) – FE results

Policy variable:	Hybrid		Official	
Model:	(1)	(2)	(3)	(4)
$\text{Policy}^{z}_{j,t-1}$	-0.0442***	-0.0460***	-0.0443***	-0.0473***
	(0.0138)	(0.0136)	(0.0134)	(0.0133)
Cost efficiency	1.519***	0.8384***	1.521***	0.8415***
	(0.3351)	(0.2947)	(0.3350)	(0.2947)
Bank liquidity	-0.0013	-0.0016	-0.0013	-0.0016
	(0.0010)	(0.0011)	(0.0010)	(0.0011)
Size	-0.1306*	-0.3400***	-0.1308*	-0.3405***
	(0.0709)	(0.0651)	(0.0708)	(0.0650)
Asset structure	-0.0576	-0.0702	-0.0576	-0.0702
	(0.0521)	(0.0458)	(0.0521)	(0.0458)
Bank Concentration	-0.0030	-0.0024	-0.0031	-0.0025
	(0.0032)	(0.0033)	(0.0032)	(0.0033)
GDP growth	0.0138**	0.0135**	0.0140**	0.0138***
	(0.0061)	(0.0053)	(0.0061)	(0.0053)
Inflation (CPI)	0.0556***	0.0517***	0.0556***	0.0518***
	(0.0107)	(0.0097)	(0.0107)	(0.0097)
Institutional Quality	-0.3830*	-0.2655	-0.3761*	-0.2560
	(0.2141)	(0.2121)	(0.2140)	(0.2120)
Macroprudential: Liquidity	-0.0517***	-0.0687***	-0.0520***	-0.0690***
	(0.0178)	(0.0183)	(0.0178)	(0.0183)
Macroprudential: LFX	-0.0246	-0.0463	-0.0232	-0.0448
	(0.0398)	(0.0388)	(0.0398)	(0.0388)
Fixed-effects				
Bank	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes
No. of Banks	3,816	3,687	3,816	3,687
N	41,687	38,424	41,687	38,424
R^2	0.84	0.85	0.84	0.85

Note: Lag 1 of all predictors in Models (2) and (4). LFX: Limits on FX
Clustered (bank level) standard errors in parentheses.
****p<0.01, **p<0.05, *p<0.1.*

Table 9. Channel analysis: impact of monetary policy on banking spread (NIM) , controlling for macroprudential policies (Liquidity and LTV) – FE results

Policy variable:	Hybrid		Official	
Model:	(1)	(2)	(3)	(4)
$\text{Policy}^{z}_{j,t-1}$	-0.0445***	-0.0453***	-0.0447***	-0.0468***
	(0.0137)	(0.0135)	(0.0133)	(0.0132)
Cost efficiency	1.519***	0.8332***	1.521***	0.8364***
	(0.3353)	(0.2946)	(0.3353)	(0.2946)
Bank liquidity	-0.0013	-0.0016	-0.0013	-0.0016
	(0.0010)	(0.0011)	(0.0010)	(0.0011)
Size	-0.1307*	-0.3395***	-0.1309*	-0.3400***
	(0.0708)	(0.0650)	(0.0708)	(0.0650)
Asset structure	-0.0576	-0.0702	-0.0576	-0.0701
	(0.0521)	(0.0458)	(0.0521)	(0.0458)
Bank Concentration	-0.0030	-0.0024	-0.0031	-0.0024
	(0.0032)	(0.0033)	(0.0032)	(0.0033)
GDP growth	0.0139**	0.0128**	0.0141**	0.0130**
	(0.0062)	(0.0054)	(0.0063)	(0.0055)
Inflation (CPI)	0.0559***	0.0515***	0.0559***	0.0516***
	(0.0108)	(0.0098)	(0.0108)	(0.0098)
Institutional Quality	-0.3811*	-0.2614	-0.3741*	-0.2519
	(0.2143)	(0.2123)	(0.2142)	(0.2122)
Macroprudential: Liquidity	-0.0516***	-0.0675***	-0.0519***	-0.0679***
	(0.0180)	(0.0185)	(0.0180)	(0.0185)
Macroprudential: LTV	-0.0086	0.0122	-0.0092	0.0113
	(0.0258)	(0.0256)	(0.0258)	(0.0256)
Fixed-effects				
Bank	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes
No. of Banks	3,816	3,687	3,816	3,687
N	41,687	38,424	41,687	38,424
R^2	0.84	0.85	0.84	0.85

Note: Lag 1 of all predictors in Models (2) and (4). LTV: Limits on Loan-to-Value Ratio
Clustered (bank level) standard errors in parentheses.
****p<0.01, **p<0.05, *p<0.1.*

Figure 7. Local projections responses of net interest margin (NIM) to monetary policy shock

(i) One Lag of Predictors | **(ii) One Lag of NIM & Predictors**

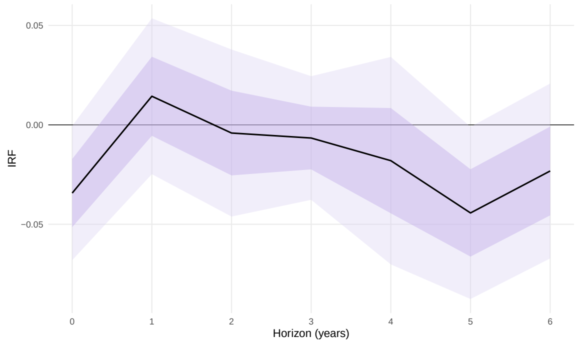

25th percentile

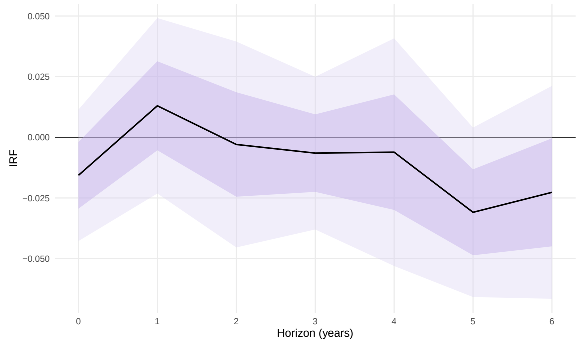

25th percentile

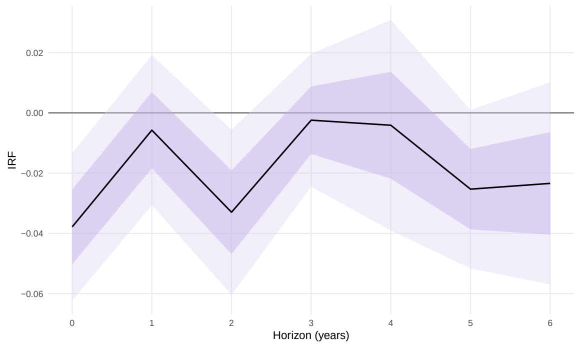

50th percentile

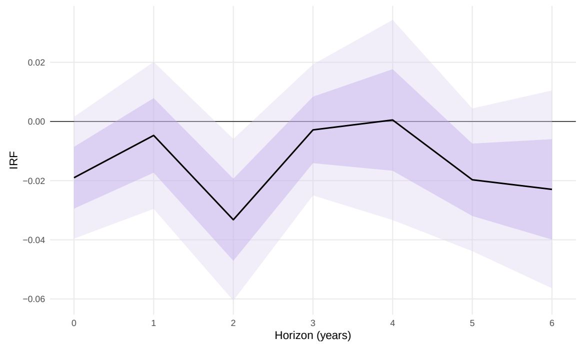

50th percentile

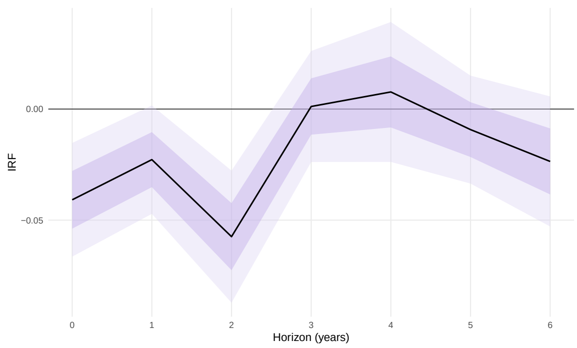

75th percentile

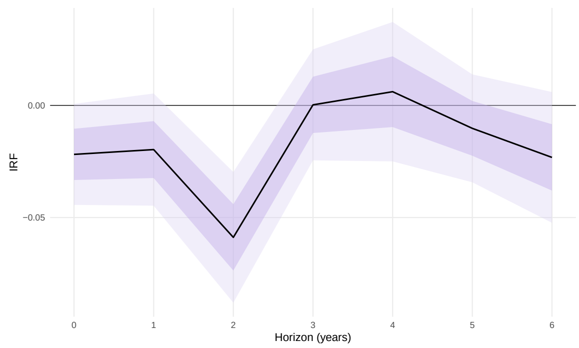

75th percentile

Note: The figure plots local projections responses of net interest margin (NIM) to a one-standard-deviation monetary policy shock, conditional on bank cost efficiency (evaluated at the 25th, 50th, and 75th percentiles). The lighter and darker bands represent 68% and 95% error bands, respectively. Column (i) includes 1 lag of all predictors; column (ii) includes one lag of NIM and predictors.

Figure 8. Local projections responses of net interest margin (NIM) to monetary policy shock, controlling for macroprudential policies (Liquidity and LFX)

(i) One Lag of Predictors

(ii) One Lag of NIM & Predictors

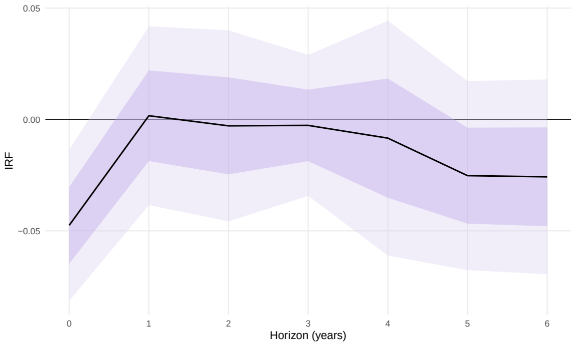

25th percentile

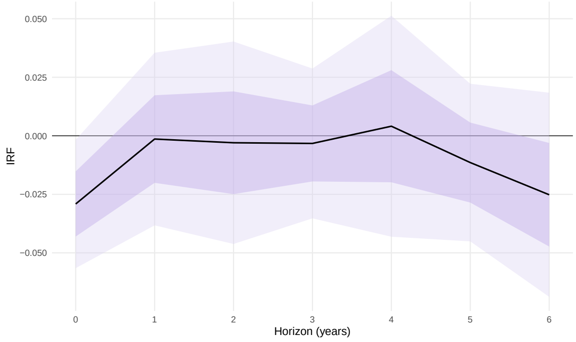

25th percentile

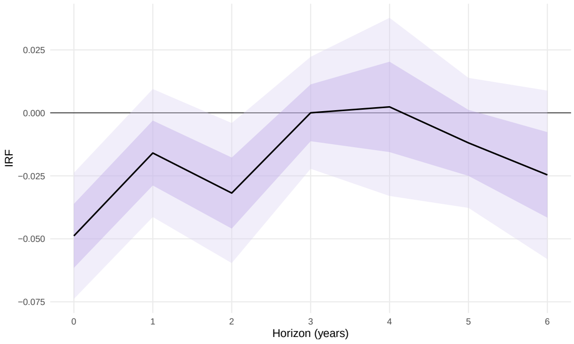

50th percentile

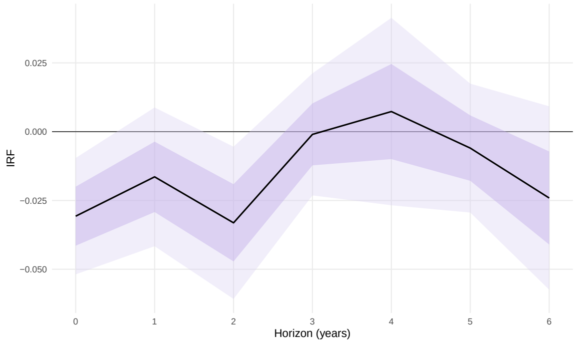

50th percentile

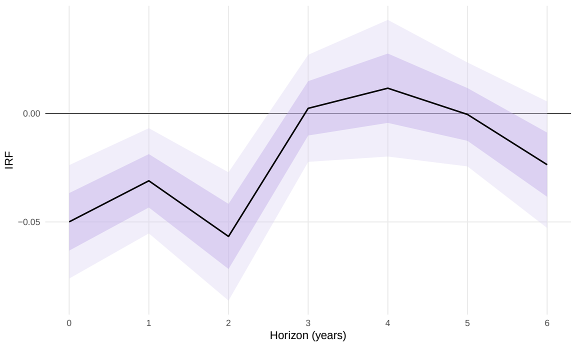

75th percentile

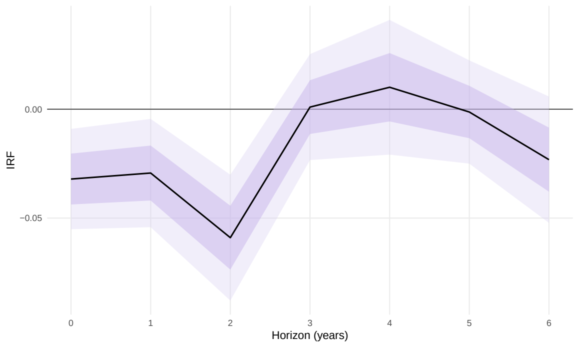

75th percentile

Note: The figure plots local projections responses of net interest margin (NIM) to a one-standard-deviation monetary policy shock, conditional on bank cost efficiency (evaluated at the 25th, 50th, and 75th percentiles), controlling for macroprudential policies (Liquidity and LFX). The lighter and darker bands represent 68% and 95% error bands, respectively. Column (i) includes 1 lag of all predictors; column (ii) includes one lag of NIM and predictors.

Figure 9. Local projections responses of net interest margin (NIM) to monetary policy shock, controlling for macroprudential policies (Liquidity and LTV)

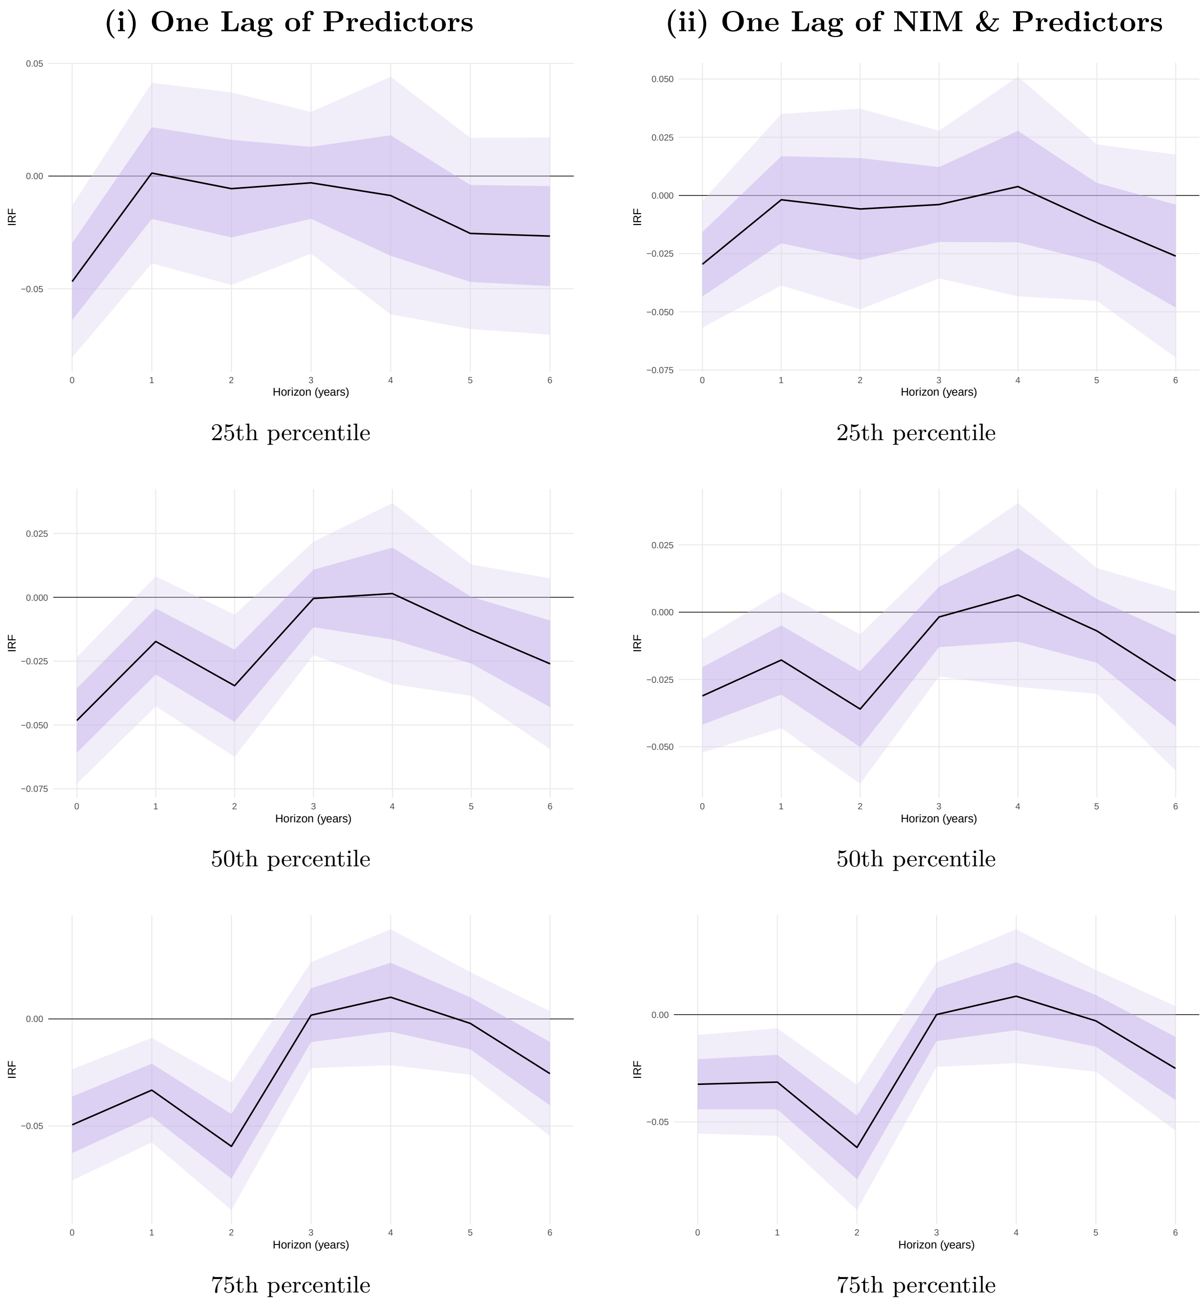

Note: The figure plots local projections responses of net interest margin (NIM) to a one-standard-deviation monetary policy shock, conditional on bank cost efficiency (evaluated at the 25th, 50th, and 75th percentiles), controlling for macroprudential policies (Liquidity and LTV). The lighter and darker bands represent 68% and 95% error bands, respectively. Column (i) includes 1 lag of all predictors; column (ii) includes one lag of NIM and predictors.

5.7 Complementary Channel: Credit Growth

As complementary evidence on balance-sheet adjustment, we examine the credit-growth response to monetary tightening, conditioning on bank cost efficiency. Credit growth captures the lending margin through which policy shocks propagate to bank balance sheets—via shifts in loan supply, funding conditions, and internal constraints. We estimate local projections in which the dependent variable is bank credit growth (log changes, $\Delta\ell$), and the key regressor is the monetary policy stance shock interacted with cost efficiency, controlling for the same bank-level and macro covariates as in Equation (15). Figures G.1–G.3 (under section G in Online Appendix) report impulse responses evaluated at the 25th, 50th, and 75th percentiles of the efficiency distribution.

The impact responses are heterogeneous across cost-efficiency percentiles. In baseline specifications without macroprudential controls, credit growth increases on impact for low-efficiency banks ($P25$) but contracts for high-efficiency banks ($p75$). Quantitatively, the impact response is positive for $P25$ (e.g., $I\hat{R}F^{P25}_{\Delta\ell}(0) \approx 0.007$), while it is negative for $P75$ (e.g., $I\hat{R}F^{P75}_{\Delta\ell}(0) \approx -0.015$), with the median group also contracting (e.g., $I\hat{R}F^{P50}_{\Delta\ell}(0) \approx -0.005$). This divergence is consistent with different adjustment speeds: less efficient banks may initially sustain lending to preserve income when funding conditions tighten (e.g., delayed repricing, relationship-lending rigidities, or slower internal constraint adjustment), whereas more efficient banks appear to react more promptly by tightening underwriting and contracting credit growth, consistent with faster balance-sheet discipline.

Beyond impact, the dynamics are not monotone, which helps anticipate the asset-quality evidence discussed next. At short horizons, the contraction among high-efficiency banks is not persistent—for example, $P75$ credit growth turns slightly positive at $h = 1$ (≈ 0.002) and is near zero around $h = 4$–5, whereas low-efficiency banks display a clearer medium-run contraction (e.g., $P25$ falls to ≈ -0.017 at $h = 2$ and to ≈ -0.030 by $h = 6$). The key takeaway is that cost efficiency maps into the timing and smoothness of lending adjustment: high-efficiency banks exhibit a more front-loaded response (initial contraction and earlier stabilisation), while low-efficiency banks display a more inertial profile with slower correction and a more persistent medium-run decline.

These credit-growth patterns complement the main stability findings. They indicate that the "efficiency buffer" operates partly through differential adjustment frictions in loan supply — with more efficient banks moving earlier to contain balance-sheet risk — even though our core interpretation emphasises risk-management discipline as the primary stabilising force on impact. This sequencing is important for interpreting the subsequent NPL dynamics: early credit restraint and tighter standards can coincide with near-term improvements in measured stability, while delayed borrower stress and loan-loss realisations can still emerge at medium horizons even when lending growth has already adjusted.

5.8 Robustness: Bank competition heterogeneity

We complement the cost-efficiency analysis with an examination of bank competition as a conditioning variable. Using bank-level marginal costs and Lerner indices derived from a two-step stochastic metafrontier framework—country-specific cost frontiers plus a global metafrontier estimated on the fitted systematic component—we evaluate the local-projection IRFs at three Lerner-index percentiles. The on-impact stabilising effect of monetary tightening is positive across all competition regimes and is somewhat larger and more persistent for banks with higher market power, consistent with a charter-value/competition-fragility view of bank resilience. The full derivation, Lerner-index world maps, regional trends, and Lerner-conditioned IRFs are reported under section F in the Online Appendix.

5.9 Robustness: using NPL as a measure of stability

While the primary stability outcome in the paper is Z-score, it is useful to complement it with a direct measure of asset quality based on nonperforming loans (NPLs). NPLs are more tightly linked to realised borrower distress and loan performance. Examining NPL dynamics alongside the Z-score helps distinguish a "buffer/margins" channel (which can mechanically support composite stability measures in the short run) from a "credit-loss" channel (which may emerge with a delay as repayment stress materialises). This distinction is particularly relevant in tightening cycles, where the earlier results indicate horizon-dependent sign reversals in stability and heterogeneity in terms of cost efficiency.

Let $\Delta \ln \mathrm{NPL}_k(h)$ denote the local-projection impulse response of log NPL growth at horizon h for efficiency group $k \in \{25, 50, 75\}$, and let $\Delta \ln \mathrm{L}_k(h)$ (also denoted as $\Delta\ell_k$) denote the corresponding impulse response of log loan growth. We follow Equation (15) and (16) but use these measures as the outcome variables. Because the NPL ratio is NPL/Loans, the log-change identity implies:

$$\Delta \ln \left(\frac{\mathrm{NPL}}{\mathrm{L}}\right)_k (h) \; = \; \Delta \ln \mathrm{NPL}_k(h) \; - \; \Delta \ln \mathrm{L}_k(h). \tag{17}$$

Accordingly, we report: i) NPL growth, $\Delta \ln \mathrm{NPL}_k(h)$, and ii) the difference object, $\Delta \ln \mathrm{NPL}_k(h) - \Delta \ln \mathrm{L}_k(h)$, which is a good measure of NPL-ratio growth that nets out denominator movements and avoids the mechanical ambiguity of ratio levels when both numerator and denominator move contemporaneously. The results are shown in Figures I.1 and I.2 under section I in the Online Appendix. The local-projection evidence reveals a clear two-phase pattern. First, at short horizons ($h = 0, 1$), a contractionary monetary policy innovation reduces NPL growth across the distribution, with the effect particularly pronounced for the median and high-efficiency groups (e.g., $\Delta \ln \mathrm{NPL}_{P50}(0) \approx -0.027$ and $\Delta \ln \mathrm{NPL}_{P75}(0) \approx -0.039$, Panel (i) Figure I.1). Over the same horizons, loan growth is already weak or negative for $P50$ and $P75$ (and close to zero for $P25$) as we discussed earlier. Consistent with Equation (17), the implied NPL-ratio growth is therefore negative at impact and at $h = 1$ for all groups, indicating an initial improvement in asset-quality dynamics in ratio terms.

Second, at medium horizons ($h = 2, 3$), NPL growth turns sharply positive and is precisely estimated across all percentiles (e.g., $\Delta \ln \mathrm{NPL}_{P50}(2) \approx 0.042$ and $\Delta \ln \mathrm{NPL}_{P75}(2) \approx 0.053$), while loan growth remains subdued and typically negative around the same period as noted earlier. The difference measure (ratio-growth measure) consequently rises strongly and significantly at $h = 2$–3 for each group, implying that NPLs begin to grow faster than loans. This medium-run deterioration is economically important: because the ratio-growth measure controls for the denominator, it indicates that the subsequent worsening is not merely a mechanical consequence of slower credit expansion, but reflects a genuine increase in the intensity of problem-loan accumulation relative to the loan book. At longer horizons ($h \geq 4$), the responses display partial reversals and oscillations, which are consistent with a combination of loan-loss recognition and resolution (charge-offs and write-downs), balance-sheet repair, and intertemporal re-optimisation of lending following the initial tightening episode.

Taken together, the Z-score and NPL results align with a fast-versus-slow propagation mechanism. Tightening can improve measured stability (Z-score) and reduce NPL growth initially — consistent with an immediate risk-management/underwriting response and reduced risk-taking — but borrower repayment stress and credit-loss realisations materialise with a delay, raising NPL accumulation (and NPL-ratio growth) at medium horizons. This two-speed dynamics provides an empirical bridge to the DSGE mechanism developed later: a fast risk-management channel that improves near-term measured stability, and a slow borrower-distress channel that generates delayed deterioration in asset quality and, ultimately, in stability.

5.10 Robustness: instrumental variables–two-stage least squares (IV-2SLS)

The study employs the two-stage least squares (2SLS) instrumental variables (IV) approach to further strengthen the identification of monetary policy stance. Given the second stage model as in Equation (8) using the predicted monetary policy stance, $\hat{Policy}^{z}_{j,t-1}$, following from the first stage regression as specified below:

$$Policy^{z}_{j,t-1} = \alpha_i + \mu_j + \delta_t + \beta_2\,\mathrm{CBI}_{j,t} + \gamma'\mathrm{Controls}_{ijt} + \xi_{j,t} \tag{18}$$

where all variables are as defined earlier. Here, we instrument monetary policy variable using a measure of central bank independence,$\mathrm{CBI}_{j,t}$ sourced from Garriga [2016]. It is an index that aggregates 16 legal indicators into four main categories: tenure of the bank's governor, an indicator related to policy formulation, an indicator related to the central bank's objectives, and an indicator related to the limitation on lending to the government. The index ranges from 0 to 1, with higher values suggesting more independence. We use the aggregate measure, which takes the average of the four indicators. The advantage of the CBI from Garriga [2016] is that the dataset is comprehensive, covering approximately 182 countries, unlike previous studies that focused on developed countries and provided limited samples from developing countries. It is by far the most extensive dataset that computes the Cukierman, Webb and Neyapty (CWN) index, which is updated yearly from 1970 to 2024.

Identification strategy We exploit cross-country and over-time variation in central bank independence (CBI) as an instrument for monetary policy. Conceptually, CBI captures the central bank's capacity to control monetary instruments [Bernhard, 2002] or, equivalently, the set of legal and institutional constraints on the government's ability to influence the conduct of monetary policy [Garriga, 2016]. We use CBI as an instrument because it is theoretically well grounded, highly correlated with the monetary policy stance, and offers a credible exclusion restriction. The exclusion restriction is also plausible. Conditional on bank, country, and year fixed effects, the degree of central bank independence is unlikely to have a direct effect on banking-sector stability, except through its influence on monetary policy choices. CBI has no direct effect on banks' risk-taking or balance-sheet decisions once monetary policy is controlled for. Formally, we assume that CBI is predetermined with respect to bank-level shocks and affects banking stability only through monetary policy, implying that CBI is orthogonal to the structural error term, $\varepsilon_{i,j,t}$. The results are presented in Tables J.1 to J.3 under section J in the Online Appendix.

Table J.1 reports the baseline two-stage least squares (2SLS) estimates of the impact of monetary policy on banking stability, where the monetary policy stance is instrumented using central bank independence (CBI). Across all specifications, the coefficient on the lagged monetary policy stance is positive and statistically significant, regardless of whether the Hybrid or Official policy measure is used. This suggests that exogenous monetary tightening, as indicated by higher CBI, is associated with increased bank stability. The magnitude of the effect is economically meaningful and larger in specifications that include lagged controls (Models (2) and (4)), suggesting that accounting for persistence in bank characteristics strengthens the stabilising role of monetary tightening. The first-stage results confirm the instrument's relevance: CBI is strongly and negatively associated with the monetary policy stance, consistent with the notion that more independent central banks implement tighter policy frameworks. The first-stage F-statistics comfortably exceed conventional thresholds (10), alleviating concerns about weak instruments and supporting a causal interpretation of the second-stage estimates.

Tables J.2 and J.3 extend the baseline analysis by explicitly controlling for macroprudential policy instruments, focusing on liquidity-based measures and limits on foreign exchange (LFX) and loan-to-value (LTV) ratios. Importantly, the positive and significant effect of monetary tightening on bank stability remains robust to the inclusion of these policy tools, both in magnitude and significance. This suggests that the stabilising effect of monetary policy is not merely capturing the operation of macroprudential regulation but reflects an independent transmission channel. Among the macroprudential controls, LFX measures are consistently positive and significant, indicating that restrictions on foreign currency exposures complement monetary tightening in enhancing bank resilience. By contrast, liquidity- and LTV-based macroprudential tools display weaker and less precisely estimated effects, particularly in lagged specifications. Overall, the 2SLS evidence implies that monetary tightening, when plausibly exogenous and insulated from reverse causality via CBI, contributes to higher banking stability even in regulatory environments where macroprudential policies are actively deployed.

5.11 Testing non-linearity in policy stance

To further assess the economically plausible state dependence of stability response to monetary policy stance, we complement our earlier estimations with a parsimonious quadratic term in the policy stance measure. This exercise is intended as a diagnostic of functional form—testing for diminishing returns or threshold effects consistent with nonlinear balance-sheet channels—rather than as a replacement for the main empirical design. Our baseline specifications remain the workhorse because they deliver a transparent average marginal effect that is directly comparable across local projections and the linearised DSGE mapping. The results are presented in Tables K.1–K.3 and Figures K.1–K.3 under section K in the Online Appendix.

Across the tables, the policy stance enters with a positive linear term and a negative quadratic term, implying a concave (inverted-U) relationship between tightening and banking stability: the marginal effect of a higher policy stance is positive at relatively accommodative-to-moderate settings but declines monotonically as policy becomes more restrictive, turning negative beyond the estimated inflection point (see Figures K.1–K.3). Economically, this pattern is consistent with a "discipline" phase in which moderate tightening improves stability (e.g., via stronger screening/risk management and reduced risk-taking) followed by a "strain" phase in which further tightening compresses margins and worsens borrower stress, thereby eroding stability. The implied turning points are stable across specifications and controls: for the hybrid measure they lie between 0.356 and 0.666 (Models (2)–(1) across the baseline and macroprudential controls), and for the official measure between 0.507 and 0.785 (Models (4)–(3)), indicating that the peak stabilising effect occurs at moderately restrictive stances and that, beyond this region, additional tightening reduces stability; the downward-sloping marginal-effect profiles reinforce that interpretation.

Because $Policy^{z}_{j,t-1}$ is standardised, the inflection points are naturally interpreted in standard-deviation units of the stance distribution. Using the sample standard deviations for the standardised series (hybrid ≈ 0.95, official ≈ 0.96), the turning points correspond to roughly 0.38–0.70 s.d. for the hybrid measure (0.356/0.95 to 0.666/0.95) and 0.53–0.82 s.d. for the official measure (0.507/0.96 to 0.785/0.96). To translate these thresholds into basis points, note that a one-standard-deviation move in the raw (non-standardised) stance series equals $\sigma^{\text{raw}}_{\text{Policy}}$ in percentage points, i.e. $100\,\sigma^{\text{raw}}_{\text{Policy}}$ basis points. With $\sigma^{\text{raw}}_{\text{Policy}} = 0.7872$ percentage points (i.e. 78.72 bp), the turning point in basis points is:

$$\text{Turning point (bp)} \;=\; (100\,\sigma^{\text{raw}}_{\text{Policy}}) \times \frac{x^*}{\text{sd}(Policy^z)} \;=\; 78.72 \times \frac{x^*}{\text{sd}(Policy^z)}.$$

Hence the implied peaks occur at approximately 29.5–55.2 bp for the hybrid measure (from $78.72 \times 0.356/0.95$ to $78.72 \times 0.666/0.95$) and 41.6–64.4 bp for the official measure (from $78.72 \times 0.507/0.96$ to $78.72 \times 0.785/0.96$). Interpreted as policy-rate-equivalent magnitudes, these values indicate that the stabilising region is reached at moderate tightening (tens of basis points relative to the sample's typical policy-rate variation), after which marginal tightening becomes progressively less stabilising and eventually destabilising; consistent with that, Models

(2) and (4) deliver lower x^*, implying an earlier onset of the adverse marginal effects when all covariates are lagged.

6 A DSGE MODEL OF MONETARY POLICY, BANK COST EFFICIENCY, AND FINANCIAL STABILITY

6.1 Model Environment

Time is discrete and indexed $t = 0, 1, 2, \ldots$. The economy contains: a representative household; a continuum of monopolistically competitive intermediate-goods firms with Calvo pricing; three bank *types* indexed by cost-efficiency percentile $k \in \{25, 50, 75\}$; a representative borrower (or, equivalently, three borrower segments, one per bank type, that we aggregate); and a monetary authority that follows a Taylor rule.

Households. The representative household maximises

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{C_t^{1-\sigma} - 1}{1-\sigma} - \frac{N_t^{1+\varphi}}{1+\varphi} \right] \tag{19}$$

over consumption C_t and labour supply N_t, where $\beta \in (0,1)$ is the discount factor, $\sigma > 0$ is the inverse of the intertemporal elasticity of substitution, and $\varphi > 0$ is the inverse Frisch elasticity of labour supply. The household is subject to the nominal budget constraint

$$P_t C_t + \sum_k D_{k,t} = W_t N_t + \sum_k R^D_{k,t-1} D_{k,t-1} + \Pi_t + T_t, \tag{20}$$

where P_t is the price level, W_t is the nominal wage, $D_{k,t}$ are deposits placed at bank type k, $R^D_{k,t}$ is the gross deposit rate, Π_t denotes profits rebated from firms and banks, and T_t is a lump-sum transfer. Households draw utility from a Drechsler-style CES aggregate of the three deposit varieties:

$$D_t \equiv \left(\sum_k \omega_k^{1/\eta_D} D_{k,t}^{(\eta_D - 1)/\eta_D} \right)^{\eta_D/(\eta_D - 1)}, \qquad \eta_D > 1, \tag{21}$$

where D_t is the composite deposit index, $\omega_k > 0$ are share parameters (normalised to sum to one), and η_D is the baseline elasticity of substitution across deposit types. Bank-specific deposit demand is therefore

$$D_{k,t} = \omega_k \, (R^D_{k,t}/R^D_t)^{\eta_{D,k}} \, D_t, \tag{22}$$

where $\eta_{D,k}$ is bank-specific because cost efficiency relaxes the elasticity of household substitution towards type k (better technology means smoother access to depositors).

The household's first-order conditions are derived in Appendix L.1. The consumption Euler equation, after log-linearising around the zero-inflation steady state, yields the IS curve:

$$\hat{x}_t = \mathbb{E}_t[\hat{x}_{t+1}] - \frac{1}{\sigma}\Big(\hat{R}_t - \mathbb{E}_t[\hat{\pi}_{t+1}]\Big), \tag{23}$$

where $\hat{x}_t$ is the output gap. The intratemporal labour supply condition pins down the real wage as $W_t/P_t = C_t^{\sigma} N_t^{\varphi}$. The deposit-demand equation (22) follows from the household's cost-minimisation over the CES deposit aggregate (21) (see Appendix L.1).

Firms and prices. The intermediate goods sector is standard New Keynesian with Calvo pricing, delivering the linearised NKPC

$$\pi_t = \beta\, \mathbb{E}_t[\pi_{t+1}] + \kappa\, x_t, \tag{24}$$

where x_t is the output gap and κ is the slope. See Appendix L.2 for the derivation.

Monetary policy. The central bank follows a Taylor rule with interest-rate smoothing,

$$R_t = \rho_R R_{t-1} + (1-\rho_R)(\phi_\pi \pi_t + \phi_x x_t) + \varepsilon_t^R, \tag{25}$$

where R_t is the nominal policy rate, $\rho_R \in (0,1)$ is the smoothing coefficient, $\phi_\pi > 1$ and $\phi_x \geq 0$ are the inflation and output-gap response coefficients, and ε_t^R is a monetary policy innovation following an AR(1) process,

$$\varepsilon_t^R = \rho_\varepsilon \varepsilon_{t-1}^R + \sigma_R e_t^R, \tag{26}$$

with persistence $\rho_\varepsilon \in [0,1)$, shock standard deviation $\sigma_R > 0$, and $e_t^R \sim$ i.i.d.$(0,1)$. The persistence $\rho_\varepsilon > 0$ is essential for the slow distress channel to accumulate.

6.2 Borrower Problem and Default

A representative borrower-segment per bank type holds a stock of debt $B_{k,t}$ (measured in deviation from steady state) and undertakes a project with idiosyncratic productivity $\omega_i \sim F$, lognormal with mean one and variance σ_ω^2. The aggregate gross return on capital is R_t^k (in equilibrium proportional to the output gap), Q_{t-1} is the price of capital, and K_{t-1} is the physical capital stock. Default occurs whenever realised cash flow falls below the contractual debt-service obligation:

$$\omega_i \cdot R_t^k \cdot Q_{t-1} K_{t-1} < R_{k,t-1}^L\, B_{k,t-1}, \tag{27}$$

which defines a threshold $\bar{\omega}_{k,t}$ such that the default probability is

$$p_{k,t} = F(\bar{\omega}_{k,t}). \tag{28}$$

Bank monitoring intensity $m_{k,t}$ shifts the volatility of ω_i:

$$\sigma_{\omega,k}(m) = \sigma_\omega \exp(-\zeta_k m_{k,t}), \tag{29}$$

so that higher monitoring shrinks the tail and reduces $p_{k,t}$. Linearising (28) around the steady state and the threshold, and using the smoothed debt-service ratio d_t (defined below) as a proxy for the threshold deviation,

$$\hat{p}_{k,t} = \pi_d\, \hat{d}_t \; - \; \chi_k\, \hat{m}_{k,t}, \tag{30}$$

where the composite parameters are

$$\pi_d \equiv \frac{f(\bar{\omega})\, \bar{\omega}}{F(\bar{\omega})} \quad \text{(CSV elasticity at the threshold)}, \tag{31}$$

$$\chi_k \equiv \pi_d \cdot \zeta_k \cdot \bar{\omega} \cdot \sigma_\omega \cdot \text{(scale factor)}. \tag{32}$$

The mapping shows that χ_k rises in the technology parameter ζ_k, which we identify with cost efficiency: high-k banks have larger ζ_k and so enjoy a larger marginal default-reduction return to monitoring.

Borrower debt-rollover problem. Debt amortises at rate, δ_b, per period, and new borrowing is determined by a debt-Euler equation that we approximate (linearly) by

$$\hat{b}_t = (1-\delta_b)\,\hat{b}_{t-1} + \delta_b\Big(\eta_y \hat{y}_t - \eta_R \hat{R}^L_t\Big), \tag{33}$$

where $\eta_y, \eta_R > 0$ are the income and rate semi-elasticities and $\hat{R}^L_t$ is the bank-share-weighted average loan rate. Distress—the smoothed debt-service ratio—follows

$$\hat{d}_t = \rho_d\,\hat{d}_{t-1} + (1-\rho_d)\,(\hat{R}^L_t + \hat{b}_t - \hat{y}_t) \;-\; \phi_{\Delta\ell}\,\bar{\ell}_{t-1}, \tag{34}$$

where $\bar{\ell}_{t-1} \equiv \frac{1}{3}\sum_k \hat{\ell}_{k,t-1}$ is the lagged aggregate lending deviation. The $\phi_{\Delta\ell}$ term is a disciplined reduced-form proxy for the credit-crunch channel: when banks contract lending in period $t-1$, borrowers in period t face tighter credit conditions and cannot refinance maturing debt as easily, which worsens the debt-service burden and raises distress. While this channel is motivated by the financial accelerator literature [Bernanke et al., 1996], we note that a fully microfounded version would require an explicit borrower refinancing decision that is beyond the scope of the current linearised framework. Using the lagged (not contemporaneous) lending level avoids simultaneity while capturing the one-period propagation delay. Distress is therefore an endogenous state driven by loan rates, borrower debt, aggregate income, and the availability of bank credit.

6.3 The Bank's Problem

Bank type k chooses lending volume $L_{k,t}$, the loan interest rate $R^L_{k,t}$, the deposit interest rate $R^D_{k,t}$, monitoring intensity $m_{k,t}$, and dividend pay-out $\mathrm{Div}_{k,t}$ to maximise

$$V_{k,t} = \max\; \mathbb{E}_t \sum_{j=0}^{\infty} \Lambda_{t,t+j}\,\mathrm{Div}_{k,t+j}, \tag{35}$$

where $\Lambda_{t,t+j} = \beta^j (C_t/C_{t+j})^{\sigma}$ is the household stochastic discount factor, subject to:

(i) Balance-sheet identity $L_{k,t} = D_{k,t} + E_{k,t}$.

(ii) Capital adequacy $E_{k,t} \geq \bar{\kappa}\,L_{k,t}$, with multiplier $\xi_{k,t} \geq 0$.

(iii) Loan demand $L_{k,t} = (R^L_{k,t}/R^L_t)^{-\epsilon_L}\,L_t$ from the borrower aggregator, where R^L_t is the aggregate loan rate, L_t is aggregate loan demand, and $\epsilon_L > 1$ is the CES elasticity of substitution across bank types.

(iv) Deposit supply equation (22), with bank-specific elasticity $\eta_{D,k}$.

(v) Equity accumulation $E_{k,t} = (1-\delta_e)E_{k,t-1} + (1-\delta_{div})\Pi^B_{k,t}L_{k,t-1} - \mathrm{Div}_{k,t}$, where $\delta_e \in (0,1)$ is the equity depreciation rate, $\delta_{div} \in (0,1)$ is the dividend pay-out ratio, and the per-loan profit is

$$\Pi^B_{k,t} = R^L_{k,t-1}\Big(1 - \lambda_{LGD}\,p_{k,t}\Big) - (1-\bar{\kappa})R^D_{k,t-1} - c^m_k(m_{k,t}). \tag{36}$$

(vi) Risk-management technology operating and adjustment costs

$$c^m_k(m_{k,t}) = \frac{\kappa_m}{2\theta_k} m^2_{k,t} + \frac{\varphi_m}{2\theta_k}(m_{k,t} - m_{k,t-1})^2. \tag{37}$$

(vii) Rotemberg adjustment costs on $R^L_{k,t}$, $R^D_{k,t}$ and $L_{k,t}$:

$$\frac{\phi_{L,k}}{2}(R^L_{k,t}-R^L_{k,t-1})^2 \;+\; \frac{\phi_{D,k}}{2}(R^D_{k,t}-R^D_{k,t-1})^2 \;+\; \frac{\gamma_L}{2}\Big(L_{k,t}-L_{k,t-1}\Big)^2. \tag{38}$$

First-order conditions. The Lagrangian is

$$\mathcal{L}_{k,t} = \mathbb{E}_t \sum_j \Lambda_{t,t+j}\Big[\mathrm{Div}_{k,t+j} - \xi_{k,t+j}(\bar{\kappa}L_{k,t+j} - E_{k,t+j}) - \text{adj costs}\Big]. \tag{39}$$

The FOC for $\mathrm{Div}_{k,t}$ pins down the multiplier on equity at the household SDF. The remaining FOCs are stated in the working-rate units used in the linearised system.

(LP-FOC) Loan-pricing FOC. Differentiating $\mathcal{L}$ with respect to $R^L_{k,t}$ and using $\partial L_{k,t}/\partial R^L_{k,t} = -\epsilon_L L_{k,t}/R^L_{k,t}$ (CES loan demand), the frictionless desired loan rate satisfies

$$R^{L*}_{k,t}\Big(1-\lambda_{LGD}\bar{p}_k\Big) \;=\; \frac{\epsilon_L}{\epsilon_L-1}\Big[(1-\bar{\kappa})\bar{R}^D_k + \bar{\kappa}\bar{R}^E + c^m_k(\bar{m}_k) + \bar{\xi}_k\bar{\kappa}\Big]. \tag{40}$$

Linearising around steady state and absorbing the constants gives

$$\hat{R}^{L*}_{k,t} \;=\; \hat{R}_t \;+\; \omega_p\,\hat{p}_{k,t} \;+\; \omega_m\,\hat{m}_{k,t} \;-\; \omega_e\,(\widehat{car}_{k,t} - \widehat{car}_{\mathrm{avg},t}), \tag{41}$$

where $\widehat{car}_{\mathrm{avg},t} \equiv \frac{1}{3}\sum_j \widehat{car}_{j,t}$, $\omega_p \equiv \frac{\epsilon_L}{\epsilon_L-1}\lambda_{LGD}$ (markup times LGD), ω_m is a monitoring-cost pass-through coefficient, and $\omega_e \equiv \bar{\xi}\bar{\kappa}/\overline{MR}$ is the capital-constraint feedback. The ω_e term loads on the *deviation* of bank k's capital ratio from the cross-sectional average. When a bank is relatively undercapitalised ($\widehat{car}_{k,t} < \widehat{car}_{\mathrm{avg},t}$), the shadow cost of its capital constraint is higher than average and it charges a higher loan rate. This creates a financial accelerator loop in the spirit of Gertler and Kiyotaki [2010]: equity losses tighten the capital constraint, raising loan rates, which deepens borrower distress and default, which further erodes equity. Crucially, under homogeneous efficiency all types are identical, so $\widehat{car}_{k,t} = \widehat{car}_{\mathrm{avg},t}$ and the ω_e term vanishes—ensuring that the zero-gap result (Proposition 3) holds without additional parameter restrictions. With Rotemberg adjustment cost (38), the actual loan rate follows a static-Rotemberg partial-adjustment law (the static approximation is valid at annual frequency):

$$\hat{R}^L_{k,t} \;=\; \rho_{L,k}\,\hat{R}^L_{k,t-1} \;+\; (1-\rho_{L,k})\,\hat{R}^{L*}_{k,t}, \qquad \rho_{L,k} \equiv \frac{\phi_{L,k}}{1+\phi_{L,k}}. \tag{42}$$

High-cost-efficiency banks have larger $\phi_{L,k}$ (relationship lending makes loan repricing costlier), hence larger $\rho_{L,k}$ and a smaller *impact* loan beta $1-\rho_{L,k}$.

(DP-FOC) Deposit-pricing FOC (Drechsler). Differentiating with respect to $R^D_{k,t}$, taking deposit supply (22) as given,

$$R^{D*}_{k,t}\left(1+\frac{1}{\eta_{D,k}}\right) \;=\; R_t \;-\; \xi_{k,t}\bar{\kappa}\,\frac{\partial D_{k,t}}{\partial R^D_{k,t}}\Big/\Big(D_{k,t}\eta_{D,k}\Big) \;+\; \text{adj cost terms}. \tag{43}$$

Ignoring second-order capital-relief terms, the linearised desired deposit rate is the Drechsler markdown:

$$\hat{R}^{D*}_{k,t} \;=\; \beta^*_{D,k}\,\hat{R}_t, \qquad \beta^*_{D,k} \equiv \frac{\eta_{D,k}}{\eta_{D,k}+1}. \tag{44}$$

Banks with more deposit market power (smaller $\eta_{D,k}$) charge a larger markdown (smaller $\beta^*_{D,k}$). We identify high cost efficiency with high $\eta_{D,k}$, so $\beta^*_{D,75} > \beta^*_{D,50} > \beta^*_{D,25}$: efficient banks pay closer to the policy rate and earn smaller deposit franchise rents. With Rotemberg adjustment cost,

$$\hat{R}^D_{k,t} \;=\; \rho_{D,k}\,\hat{R}^D_{k,t-1} \;+\; (1-\rho_{D,k})\,\beta^*_{D,k}\,\hat{R}_t, \qquad \rho_{D,k} \equiv \frac{\phi_{D,k}}{1+\phi_{D,k}}. \tag{45}$$

At annual frequency we set $\rho_{D,k} \approx 0$ (deposits are repriced within the year). The endogenous deposit beta varies with bank type via $\eta_{D,k}$ and is the source of an explicit monetary-policy-conditioned NIM compression. This captures the economically important insight that deposit market power is a crucial determinant of bank profitability, risk-taking, and responsiveness to monetary policy [Drechsler et al., 2017].

(L-FOC) Lending FOC. Differentiating with respect to $L_{k,t}$ and substituting the expected per-loan profit (36),

$$\mathbb{E}_t[\Lambda_{t,t+1}\Pi^B_{k,t+1}] - \xi_{k,t}\bar{\kappa} - \gamma_L(L_{k,t} - L_{k,t-1}) \;=\; 0. \tag{46}$$

Linearising the expected-profit term and substituting the multiplier $\xi_{k,t}$ from the slackness condition $E_{k,t} - \bar{\kappa}L_{k,t} = 0$ when the constraint binds,

$$\hat{l}_{k,t} \;=\; \rho_l\,\hat{l}_{k,t-1} + (1-\rho_l)\Big[a_e\,\hat{e}_{k,t} - a_{r,k}\,\hat{R}^L_{k,t} + a_y\,\hat{y}_t - a_{m,k}\,\hat{m}_{k,t}\Big], \tag{47}$$

where $\rho_l \equiv \gamma_L/(1+\gamma_L)$ is the portfolio-adjustment ratio, $a_e \approx 1/\bar{\kappa}$ when the capital constraint binds, $a_{r,k} = \epsilon_L \cdot$ (steady-state share) comes from CES loan demand, a_y is the output elasticity of loan demand, and $a_{m,k}$ captures the monitoring-induced tightening of credit standards. The financial accelerator operates through the loan-pricing equation (41): when $\omega_e > 0$, equity erosion raises the loan rate, which contracts lending through the CES demand elasticity $a_{r,k}$, which reduces future profit, which further erodes equity. The impact lending response of type k is *entirely* endogenous to the structural channels, with no ad-hoc shock-loaded terms.

(M-FOC) Risk-management FOC. Differentiating with respect to $m_{k,t}$ and using the augmented marginal benefit from the risk-taking channel—banks increase the expected return to monitoring when monetary conditions tighten ($\lambda_{LGD}\chi_k \to \lambda_{LGD}\chi_k(1+\eta\varepsilon^R_t)$, where $\eta > 0$ is the semi-elasticity of the monitoring benefit to the policy innovation),

$$\lambda_{LGD}\chi_k(1+\eta\varepsilon^R_t) - \frac{\kappa_m}{\theta_k}m_{k,t} - \frac{\varphi_m}{\theta_k}(m_{k,t} - m_{k,t-1}) \;=\; 0, \tag{48}$$

where $\kappa_m > 0$ is the operating cost parameter, $\varphi_m > 0$ is the adjustment cost parameter, and $\theta_k > 0$ is the cost-efficiency index (higher θ_k reduces both cost components proportionally). Linearising gives the partial-adjustment monitoring rule:

$$\hat{m}_{k,t} \;=\; \rho_m\hat{m}_{k,t-1} + \nu_k\,\hat{\varepsilon}^R_t, \qquad \rho_m \equiv \frac{\varphi_m}{\kappa_m+\varphi_m}, \quad \nu_k \equiv \frac{\lambda_{LGD}\eta\theta_k\chi_k}{\kappa_m+\varphi_m}. \tag{49}$$

The composite parameter ρ_m is the adjustment-cost ratio (higher values mean slower monitoring adjustment), and ν_k is the impact monitoring loading, which rises in cost efficiency θ_k, delivering the heterogeneous monitoring response. The linearised monitoring cost slope is $c_{m,k} \equiv \kappa_m/\theta_k$, which appears in the ROA identity below.

6.4 Equity, Profitability, and the Z-Score

Definitional identities. Before deriving the dynamic equations, we define the spread, net interest margin, capital ratio, and credit growth identities that appear in the model:

$$\hat{s}_{k,t} = \hat{R}^L_{k,t} - \hat{R}_t, \qquad \text{(spread over policy rate)} \tag{50}$$

$$\hat{\mu}_{k,t} = \hat{R}^L_{k,t} - \hat{R}^D_{k,t}, \qquad \text{(net interest margin, NIM)} \tag{51}$$

$$\widehat{car}_{k,t} = \hat{e}_{k,t} - \hat{l}_{k,t}, \qquad \text{(capital-adequacy ratio)} \tag{52}$$

$$\Delta\hat{l}_{k,t} = \hat{l}_{k,t} - \hat{l}_{k,t-1}. \qquad \text{(credit growth)} \tag{53}$$

The NIM captures the margin between what banks earn on loans and pay on deposits. The capital ratio $\widehat{car}_{k,t}$ measures the deviation of $E_{k,t}/L_{k,t}$ from steady state; the cross-sectional average $\widehat{car}_{\text{avg},t} \equiv \frac{1}{3}\sum_k \widehat{car}_{k,t}$ enters the loan-pricing FOC (41).

Equity dynamics. From the equity accumulation constraint and the dividend FOC,

$$\hat{e}_{k,t} = \rho_e\, \hat{e}_{k,t-1} + \phi_e\, \widehat{\text{roa}}_{k,t}, \qquad \rho_e = 1 - \delta_e, \quad \phi_e = (1 - \delta_{div})\bar{\Pi}^B/\bar{\kappa} \text{ (rescaled).} \tag{54}$$

The per-unit ROA is the structural identity

$$\widehat{\text{roa}}_{k,t} = \hat{R}^L_{k,t} - \lambda_{LGD}\hat{p}_{k,t} - (1-\bar{\kappa})\hat{R}^D_{k,t} - c_{m,k}\hat{m}_{k,t}. \tag{55}$$

Structural Z-score decomposition. Following the empirical definition of Z-scores as a composite of profitability, capitalisation, and earnings volatility, we decompose the model Z-score impulse into five separately weighted channels, each with a structural interpretation:

$$\boxed{\hat{z}_{k,t} \;=\; \underbrace{w_{\text{NIM}}\,\hat{\mu}_{k,t}}_{\text{(i) flow margin}} \;-\; \underbrace{w_{\text{loss}}\,\lambda_{LGD}\,\hat{p}_{k,t}}_{\text{(ii) credit-loss tail}} \;-\; \underbrace{w_{\text{cost}}\,c_{m,k}\,\hat{m}_{k,t}}_{\text{(iii) monitoring cost}} \;+\; \underbrace{w_{\text{CAR}}\,(\hat{e}_{k,t} - \hat{l}_{k,t})}_{\text{(iv) capital ratio}} \;+\; \underbrace{w_{\text{vol},k}\,\hat{m}_{k,t}}_{\text{(v) volatility reduction}}} \tag{56}$$

The weights have structural interpretations tied to the steady-state ROA volatility $\bar{\sigma} \equiv \text{sd(ROA)}$ and the steady-state Z-score level $\bar{Z} \equiv (\overline{\text{ROA}} + \overline{\text{CAR}})/\bar{\sigma}$:

$$\begin{aligned}
w_{\text{NIM}} &= \tfrac{1}{\bar{\sigma}} \cdot \text{(NIM share of ROA)}, \\
w_{\text{loss}} &= \tfrac{1}{\bar{\sigma}} \cdot \text{(credit-loss share of ROA)}, \\
w_{\text{cost}} &= \tfrac{1}{\bar{\sigma}} \cdot \text{(monitoring-cost share of ROA)}, \\
w_{\text{CAR}} &= \tfrac{1}{\bar{\sigma}} \quad \text{(direct effect of the capital ratio on Z)}, \\
w_{\text{vol},k} &= \bar{Z} \cdot \zeta_{\sigma,k}, \quad \text{where } \zeta_{\sigma,k} \equiv -\partial \ln \sigma_\omega / \partial m_k \text{ is the monitoring-to-volatility elasticity.}
\end{aligned}$$

The decomposition makes channels (i)–(v) inspectable for any horizon and any bank type. The fast monitoring/volatility channel dominates impact stabilisation. The slow loss-tail channel dominates the medium-horizon decline. The sign reversal in $z_{k,t}$ arises from the competition between these two structural mechanisms; the cross-sectional ordering arises because ν_k, χ_k, $\zeta_{\sigma,k}$ all rise in θ_k.

6.5 Non-Performing Loans

NPLs follow a Markov loan-quality transition with hazards Performing→NPL at rate $\pi_{PN}\exp(\delta_{PN} d_t)$ and NPL→Resolved at rate $\pi_{NP,k}\exp(\delta_{NP} m_{k,t})$. Linearising the steady-state stock around

$\overline{NPL}_k = \pi_{PN} L/(\pi_{PN} + \pi_{NP,k})$,

$$\hat{npl}_{k,t} = \rho_{npl,k}\hat{npl}_{k,t-1} + (1-\rho_{npl,k})\Big(\iota_d \hat{d}_t - \iota_{m,k}\hat{m}_{k,t}\Big). \tag{57}$$

The composite parameters satisfy $\rho_{npl,k} = 1 - \pi_{NP,k} - \pi_{PN}(L/\overline{NPL}_k - 1)$, $\iota_d = \delta_{PN}\,\pi_{PN}/(\pi_{PN} + \pi_{NP,k})$, and $\iota_{m,k} = \delta_{NP}\,\pi_{NP,k}/(\pi_{PN} + \pi_{NP,k})$. The NPL response is fully endogenous, driven by the interaction of borrower distress and bank monitoring behaviour.

6.6 Equilibrium

A competitive equilibrium is a set of allocations and prices $\{C_t, N_t, D_t, Y_t, R_t, \Pi_t, W_t\}$, $\{B_{k,t}, d_t, m_{k,t}, p_{k,t}, R^L_{k,t}, R^D_{k,t}, \mu_{k,t}, \mathrm{roa}_{k,t}, e_{k,t}, l_{k,t}, z_{k,t}, npl_{k,t}\}_{k\in\{25,50,75\}}$, and shadow prices $\xi_{k,t}$ such that: households satisfy (19)–(22); firms satisfy Calvo pricing implying (24); the central bank follows (25)–(26); for each k banks satisfy (42), (45), (47), (49), (54), (56), (57); borrowers satisfy (33)–(34), (30); and the goods market clears.

The linearised system contains 54 endogenous variables (5 macro, 3 borrower-side, 15×3 per-type bank variables, and 1 cross-sectional average $\widehat{car}_{\mathrm{avg},t}$) and 54 equations.

6.7 Analytical Properties of the Model

Before turning to the quantitative analysis, we establish several analytical results about the model's equilibrium properties. These results hold for any parameter configuration satisfying the Blanchard–Kahn conditions and the sign restrictions stated in the assumptions.

Assumption 1 (Efficiency ordering)**.** *The cost-efficiency parameter θ_k satisfies $\theta_{25} < \theta_{50} < \theta_{75}$, which implies $\chi_{25} \le \chi_{50} \le \chi_{75}$ (default-reduction elasticities) and that the composite impact-stabilisation product $w_{\mathrm{vol},k}\,\nu_k$ is increasing in k: $w_{\mathrm{vol},25}\nu_{25} < w_{\mathrm{vol},50}\nu_{50} < w_{\mathrm{vol},75}\nu_{75}$. Note that the individual factors ν_k and $w_{\mathrm{vol},k}$ need not each be monotone in k; the ordering of the Z-score impact response depends on their product.*

Assumption 2 (Positive channel weights)**.** *All Z-score channel weights satisfy $w_{\mathrm{NIM}} \ge 0$, $w_{\mathrm{loss}} > 0$, $w_{\mathrm{CAR}} \ge 0$, $w_{\mathrm{vol},k} > 0$, and $\lambda_{LGD} > 0$.*

Lemma 1 (Impact monitoring response)**.** *Following a unit monetary-policy tightening ($e^R_t = 1$), the impact monitoring response is $\hat{m}_{k,0} = \nu_k \sigma_R > 0$ for all k, with $\nu_k > 0$.*

Proposition 1 (Impact stabilisation under sufficient conditions)**.** *Define the net impact stabilisation channel for type k as*

$$\mathcal{S}_k \;\equiv\; w_{\mathrm{vol},k}\,\nu_k\sigma_R \;-\; w_{\mathrm{loss}}\,\lambda_{LGD}\Big(\pi_d\hat{d}_0 - \chi_k\nu_k\sigma_R\Big) \;+\; w_{\mathrm{NIM}}\,\hat{\mu}_{k,0},$$

where $\hat{d}_0 = (1-\rho_d)(\hat{R}^L_0 + \hat{b}_0 - \hat{y}_0) - \phi_{\Delta\ell}\bar{\ell}_{-1}$ is the impact distress (with $\bar{\ell}_{-1} = 0$). If $\mathcal{S}_k > 0$, then $\hat{z}_{k,0} > 0$ (impact stabilisation). A sufficient condition is

$$\Big(w_{\mathrm{vol},k} + w_{\mathrm{loss}}\lambda_{LGD}\chi_k\Big)\,\nu_k\sigma_R \;>\; w_{\mathrm{loss}}\lambda_{LGD}\pi_d\hat{d}_0 \;-\; w_{\mathrm{NIM}}\hat{\mu}_{k,0}. \tag{58}$$

Under Assumption 1, the left-hand side is strictly increasing in θ_k (through ν_k, χ_k, and $w_{\mathrm{vol},k}$). On the right-hand side, $\hat{\mu}_{k,0}$ is more negative for higher-efficiency types (because they have larger deposit markdowns $\beta^*_{D,k}$, hence stronger NIM compression), so $-w_{\mathrm{NIM}}\hat{\mu}_{k,0}$ is less negative—the destabilising NIM term is smaller *for more efficient banks. Thus if (58) holds for $k = 25$ (weakest monitoring and largest NIM drag), it holds a fortiori for $k \in \{50, 75\}$.*

Proposition 2 (Sign reversal at medium horizons)**.** *If $\rho_d > \rho_m$ (distress is more persistent than monitoring) and $w_{\text{loss}}\lambda_{LGD}\pi_d > 0$, then for each k there exists a finite horizon $H_k^* > 0$ such that $\hat{z}_{k,h} < 0$ for all $h \geq H_k^*$. Moreover, $H_{25}^* \leq H_{50}^* \leq H_{75}^*$: efficient banks reverse later.*

Proposition 3 (Efficiency gap and its source)**.** *Let $\Theta = (\nu_k, \chi_k, \beta_{D,k}, \rho_{L,k}, a_{m,k}, a_{r,k}, w_{\text{vol},k}, \rho_{npl,k}, \iota_{m,k})$ collect all type-specific parameters. If $\Theta_{25} = \Theta_{50} = \Theta_{75}$ (homogeneous efficiency), then $\hat{z}_{25,h} = \hat{z}_{50,h} = \hat{z}_{75,h}$ for all $h \geq 0$: the efficiency gap is identically zero.*

Proposition 4 (Necessity of the risk-management channel)**.** *If $\nu_k = 0$ for all k (banks cannot adjust monitoring in response to the shock), then $\hat{z}_{k,0} < 0$ for all k whenever $w_{\text{loss}}\lambda_{LGD}\pi_d\hat{d}_0 > w_{\text{NIM}}\hat{\mu}_{k,0}$.*

Lemma 2 (Financial accelerator amplification)**.** *For $\omega_e > 0$, the medium-horizon Z-score decline is strictly deeper than with $\omega_e = 0$, holding all other parameters fixed, provided that the model satisfies BK conditions under both parameter configurations.*

All proofs are collected in Appendix L.5.

6.8 Calibration

Table 10 summarises the calibration. Each composite parameter is presented next to the deeper primitive(s) it represents. The minimum-distance estimator (Section 6.10) re-tunes a subset. Where possible, we discipline calibrated values using the bank panel from the empirical analysis. We construct implied bank-level lending rates (interest on loans / net loans) and deposit rates (interest on customer deposits / customer deposits) from raw Osiris balance-sheet items and verify key transmission parameters by cost-efficiency tercile. The deposit markdown ordering ($\beta_{75}^D > \beta_{50}^D > \beta_{25}^D$) is confirmed by the panel, and NIM persistence estimates (0.58–0.63) bracket the loan-rate stickiness values. The Z-score weights are cross-checked against a variance decomposition of the observed Z-score into its ROA and capital-ratio components, and ϕ_e is consistent with the algebraic identity $(1-\delta_{div})\bar{\Pi}^B/\bar{\kappa}$ evaluated at sample means. A sensitivity analysis confirms that IRF patterns are robust to $\pm 20\%$ perturbations of ρ_m and χ_k.

Table 10. Structural calibration and minimum-distance estimates. For MD-estimated parameters, the "Value" column shows calibration → estimate.

Parameter	Description	Value	Source
A. New Keynesian core			
β, σ, κ	discount, IES^{-1}, NKPC slope	0.99, 1.5, 0.10	Standard
ρ_R, ϕ_π, ϕ_x	Taylor rule	0.80, 1.50, 0.10	Standard
ρ_ε	shock persistence	0.50	Calibrated
σ_R	shock size	0.15 → 0.05	MD est.
B. Borrower technology			
δ_b	amortisation rate	0.15	Calibrated
η_R, η_y	debt-Euler semi-elast.	1.00, 1.00	Normalised
ρ_d	distress smoothing	0.85 → 0.96	MD est.
$\phi_{\Delta\ell}$	credit-crunch amplifier	0.10 → 0.44	MD est.
C. CSV default			
π_d	CSV elasticity	0.50 → 0.99	MD est.
λ_{LGD}	loss given default	0.45	BGG
χ_k	monitoring-vol elast.	0.38, 0.40, 0.55	Calibrated; robust ±20%
D. Risk-management technology			
ρ_m	adj-cost ratio	0.40	Calibrated; robust ±20%
ν_k	monitoring loading	0.35, 0.55, 0.80 → 1.32, 3.83, 3.32	MD est.
$c_{m,k}$	cost slope	0.05	Calibrated
E. Loan-pricing FOC			
ω_p	markup×LGD	0.50	Implied
ω_m	monitoring pass-through	0.05	Calibrated
ω_e	capital-constraint feedback	0.25 → 0.21	MD est.
$\rho_{L,k}$	loan-rate stickiness	0.30, 0.35, 0.40	Panel NIM persistence
F. Deposit pricing (Drechsler et al. 2017)			
$\beta^*_{D,k}$	deposit markdown	0.92, 0.94, 0.95	DSS; panel confirms ordering
G. Equity accumulation			
ρ_e	equity persistence	0.90	Standard
ϕ_e	flow loading	0.15	$(1-\delta_{\text{div}})\bar{\pi}^B/\bar{\kappa}$
$\bar{\kappa}$	SS capital ratio	0.10	Basel III
H. Lending FOC			
ρ_l	portfolio stickiness	0.60	Calibrated
a_e, a_y, a_R	equity, output, policy elast.	0.50, 0.05, 0.00	Calibrated
$a_{r,k}$	own-rate semi-elast.	0.15, 0.25, 0.35	Calibrated
$a_{m,k}$	monitoring→tightening	0.05, 0.10, 0.15 → 0.00, 0.00, 0.10	MD est.
I. Z-score decomposition			
$w_{\text{NIM}}, w_{\text{loss}}, w_{\text{cost}}, w_{\text{CAR}}$	ROA weights	0.15, 1.80, 0.00, 0.10	ΔZ decomposition
$w_{\text{vol},k}$	vol channel	0.50, 0.60, 0.70 → 0.10, 0.19, 0.42	MD est.
J. NPL hazards			
$\rho_{npl,k}$	NPL persistence	0.80, 0.70, 0.65	Panel AR(1) range
$\iota_d, \iota_{m,k}$	inflow loadings	0.10, (0.22, 0.15, 0.12)	Panel; calibrated

6.9 Impulse Response Analysis

We subject the calibrated economy to a one-standard-deviation monetary policy tightening ($e^R_t = 1$) and trace the responses of macro aggregates and the heterogeneous banking block over a ten-period horizon.

Macro responses. Figure 10 displays the impulse responses of output, inflation, the policy rate, and the borrower distress index. The transmission follows the standard New Keynesian mechanism: the policy rate rises on impact and decays at rate ρ_R; the IS curve transmits the tightening into a persistent output contraction; inflation falls via the Phillips curve; and borrower distress accumulates gradually as the debt-service burden rises, peaking several periods after the initial shock. The slow build-up of distress is the key input into the banking block, as it drives up default probabilities and erodes loan performance with a lag.

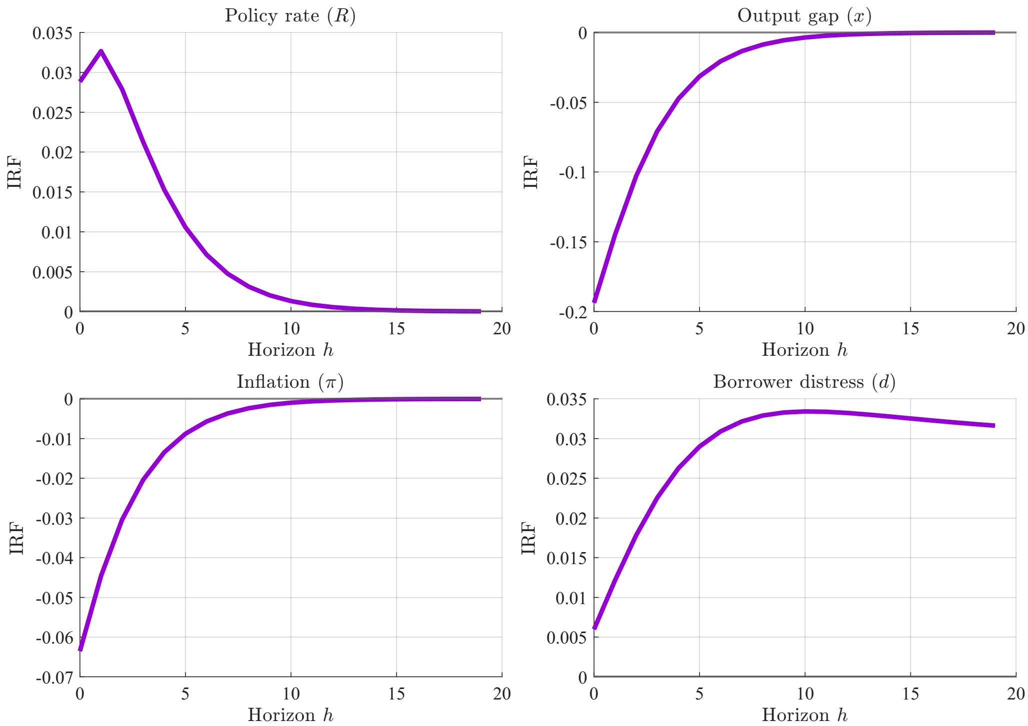

Figure 10. Macro impulse responses to a one-standard-deviation monetary policy tightening.

Bank stability. Figure 11 plots the Z-score response for each cost-efficiency group. On impact, all three bank types experience a stability improvement driven by the fast risk-management channel: the monetary tightening raises monitoring intensity, which compresses the volatility of loan returns and boosts the Z-score. The improvement is strongest for the P75 group, whose superior cost efficiency generates the largest monitoring response ($\nu_{75} > \nu_{50} > \nu_{25}$). Over the medium horizon, the slow build-up of borrower distress dominates and the Z-score declines, eventually turning negative. The sign reversal occurs earliest for the least efficient banks. This pattern—efficiency-conditioned impact smoothing followed by heterogeneous reversal—is a robust prediction of the structural model and mirrors the empirical local-projection evidence.

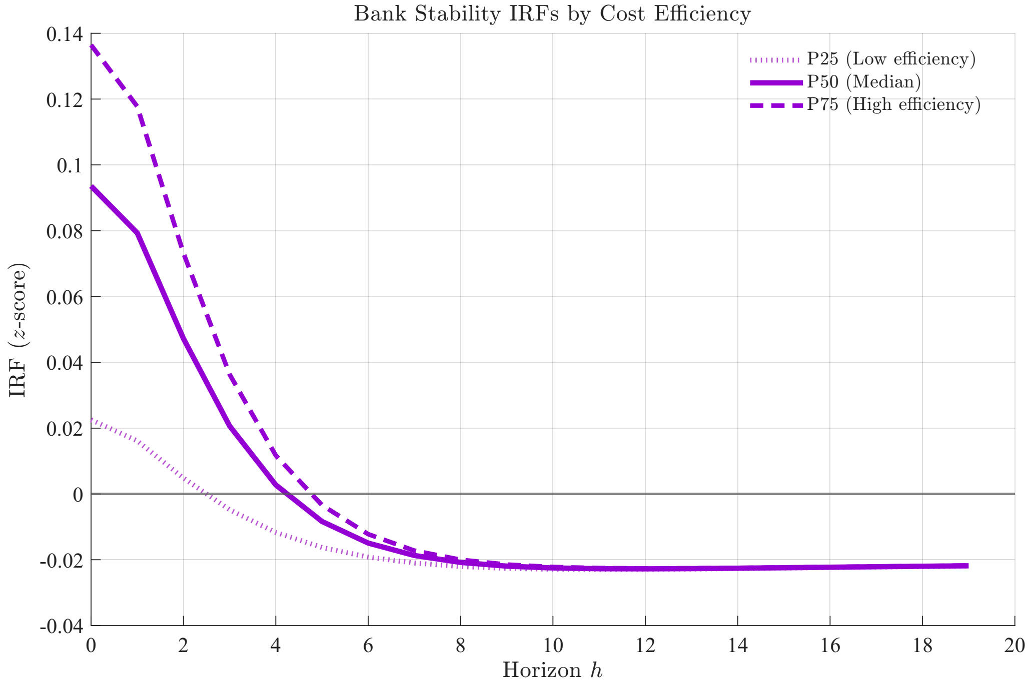

Figure 11. Bank stability index (Z-score) by cost-efficiency group following a monetary tightening.

Net interest margin. Figure 12 shows the net interest margin response. All bank types experience NIM compression on impact because loan rates are sticky (Rotemberg friction) while the policy rate rises immediately, increasing funding costs. The compression is most pronounced for the P75 group, whose higher deposit beta $\beta^*_{D,75}$ means that deposit rates pass through more of the policy-rate increase. As loan rates gradually adjust upward, NIMs partially recover, but the recovery is incomplete because the concurrent rise in default probabilities erodes expected loan income. The heterogeneous NIM dynamics are a key channel through which cost efficiency shapes the profitability component of the Z-score.

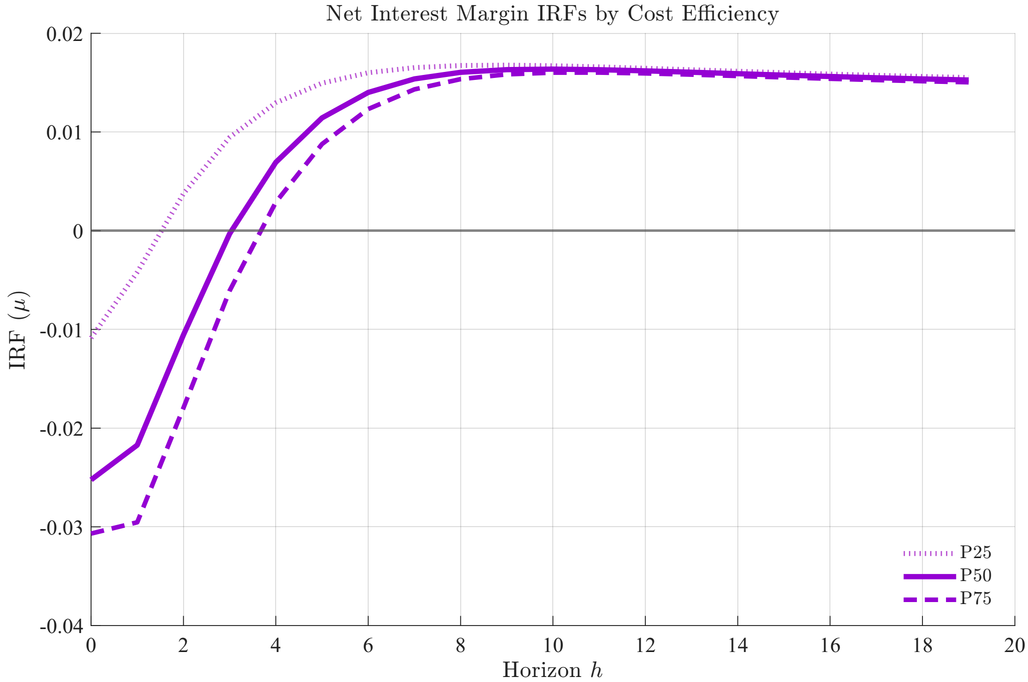

Figure 12. Net interest margin response by cost-efficiency group.

Credit growth. Figure 13 presents the credit growth response. All bank types contract lending following the tightening, consistent with the standard bank lending channel. The contraction is heterogeneous: cost-efficient banks (P75) cut lending more aggressively on impact because their stronger monitoring response tightens credit standards ($a_{m,75} > a_{m,50} > a_{m,25}$). However, the P75 contraction is also shorter-lived, as their superior equity accumulation provides a faster capital buffer to support renewed lending. The cross-sectional pattern in credit growth is consistent with the empirical finding that efficient banks exhibit more active balance-sheet management in response to policy shocks.

Efficiency gap. Figure 14 plots the efficiency gap, defined as $z_{75,t} - z_{25,t}$, over the impulse-response horizon. The gap is front-loaded: it peaks on impact, reflecting the large differential in the monitoring/volatility channel across bank types. Over subsequent periods the gap narrows as the common distress channel—which is symmetric across types—becomes dominant. The gap is mean-reverting, converging to zero as all bank types return to steady state. The front-loaded and mean-reverting profile of the gap underscores that the cross-sectional dispersion in stability is a short-run phenomenon driven by differential risk-management responses, not by permanent structural advantages.

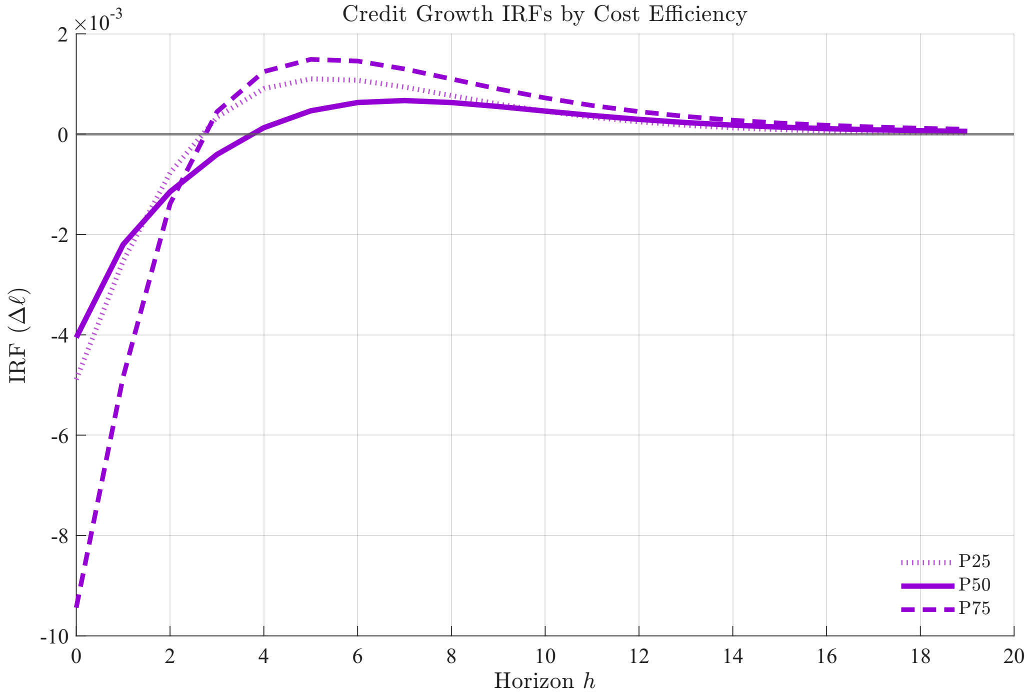

Figure 13. Credit growth by cost-efficiency group following a monetary tightening.

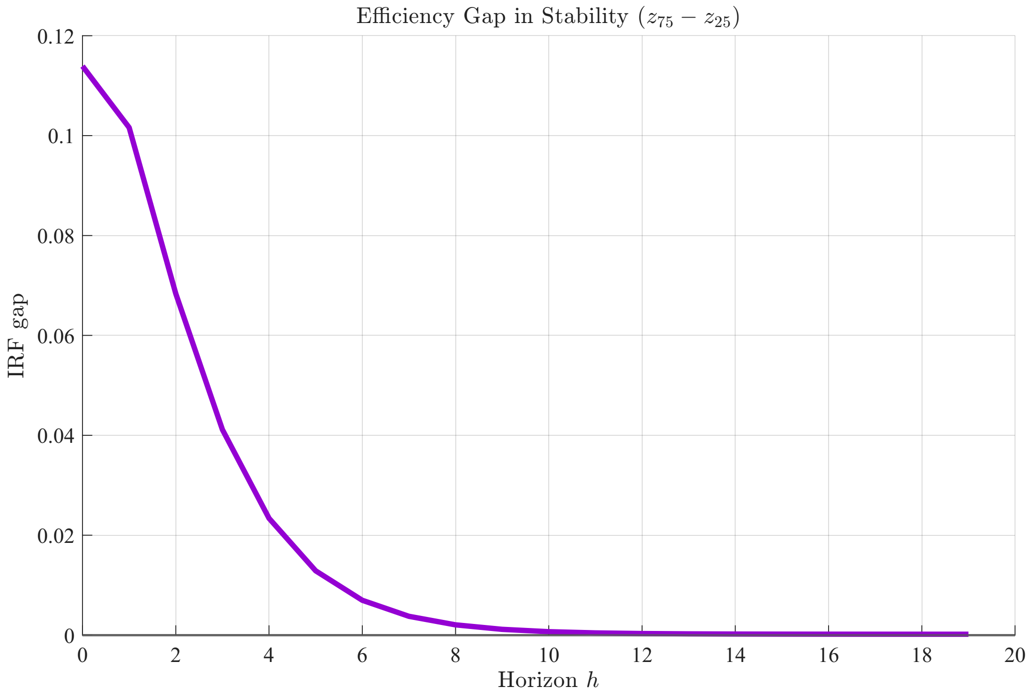

Figure 14. Efficiency gap ($z_{75} - z_{25}$) following a monetary tightening.

Z-score decomposition. Figure 15 decomposes the impact Z-score response into its five structural channels for each bank type. Two channels dominate the positive impact response. The credit-loss channel (ii) is the largest contributor for cost-efficient banks (P50 and P75): the strong monitoring response pushes the default probability below its steady state immediately, so that the term $-w_{\text{loss}}\lambda_{LGD}\hat{p}_k$ enters positively. The volatility-reduction channel (v) reinforces this effect, contributing positively across all three types. For the least efficient group (P25), the two channels are roughly comparable in magnitude. The flow-margin channel (i) contributes negatively due to NIM compression, but is small relative to the loss-reduction and volatility channels. The capital-ratio channel (iv) is near zero on impact because equity is a slow-moving state variable. At medium horizons (not shown), the loss channel reverses sign and becomes the dominant negative force as default probabilities rise with the persistent distress state, while the monitoring-related channels decay as m_k adjusts back toward steady state. This structural decomposition clarifies why efficient banks—with stronger monitoring responses ν_k that suppress both default risk and earnings volatility on impact—enjoy stronger initial stabilisation.

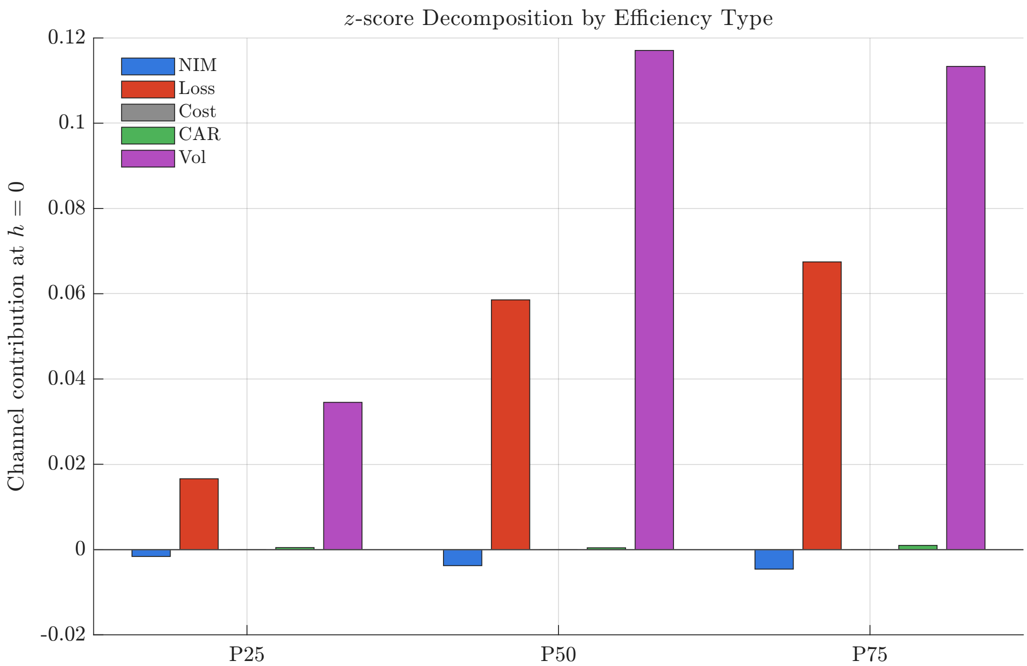

Figure 15. Structural decomposition of the impact Z-score response into five channels by bank type.

6.10 Minimum-Distance Estimation

We estimate a subset of the structural parameters by minimum distance, matching the model-implied impulse responses to the empirical local-projection (LP) IRF targets. The moment vector $\hat{g}$ stacks the LP IRFs of $z_k, \mu_k, \Delta\ell_k$ for $k \in \{25, 50, 75\}$ at horizons $h = 0, \ldots, 6$, giving 63 moments. The objective is

$$\hat{\theta} = \arg\min_{\theta} \left(\hat{g} - g(\theta)\right)' W \left(\hat{g} - g(\theta)\right), \tag{59}$$

with $W = \text{diag}(1/\hat{se}_h^2)$ using the LP horizon-by-horizon standard errors. We estimate the 14-parameter subset

$$\theta = (\sigma_R,\ \pi_d,\ \rho_d,\ \nu_{25}, \nu_{50}, \nu_{75},\ w_{\text{vol},25}, w_{\text{vol},50}, w_{\text{vol},75},\ a_{m,25}, a_{m,50}, a_{m,75},\ \omega_e,\ \phi_{\Delta\ell})'.$$

The financial accelerator parameters ω_e and $\phi_{\Delta\ell}$ are estimated because they govern medium-horizon amplification dynamics that are essential for matching the deep troughs in the LP targets at $h = 5$–6. The common Z-score channel weights (w_{NIM}, w_{loss}) are fixed at their structural calibration values to preserve the balance between the flow-margin and credit-loss channels, while the type-specific volatility-channel weights $w_{\text{vol},k}$ are estimated because they control the cross-sectional heterogeneity in impact stabilisation. The remaining parameters are fixed at their structural calibration values (Table 10).

Results. The minimised weighted sum of squared errors is 369.4, representing a 45.8% reduction from the initial calibration value of 682.1. The improvement is concentrated in the impact and reversal horizons where the re-estimated monitoring loadings ν_k, volatility-channel weights $w_{\mathrm{vol},k}$, and distress parameters (π_d, ρ_d) bring the model closer to the LP point estimates. The model captures the qualitative sign reversal in the Z-score at medium horizons, the NIM compression on impact with the correct cross-sectional ordering ($|z_{75}| > |z_{50}| > |z_{25}|$), and the heterogeneous credit contraction. Sign reversals occur at $h = 3$ (P25), $h = 5$ (P50), and $h = 5$ (P75), broadly matching the LP evidence. Of the 63 horizon-specific moments, 38 (60%) are matched in sign.

Monitoring and credit tightening. A notable feature of the estimates is that the monitoring-to-credit-tightening loadings $a_{m,25}$ and $a_{m,50}$ are driven to their lower bound of zero, while only $a_{m,75}$ remains positive. This reflects the LP evidence on credit growth: P25 banks actually *expand* lending on impact ($\Delta\ell_{25}(0) = +0.007$), while P50 banks contract only mildly (-0.005) and P75 banks contract most sharply (-0.015). A positive $a_{m,k}$ would push lending down whenever monitoring rises, working against the P25 pattern. The estimator drives $a_{m,25}$ and $a_{m,50}$ to their lower bound, revealing that less efficient banks increase monitoring after tightening (the risk-taking channel) but do not translate this into tighter credit standards—other channels (loan-rate repricing via $a_{r,k}$ and equity erosion via a_e) are sufficient to generate the observed contraction. Only the most efficient banks ($k = 75$) have the organisational capacity to link monitoring intelligence directly to lending decisions, consistent with the broader theme that cost efficiency governs not just the level but also the operational integration of risk-management responses.

Identification and fit. We view the minimum-distance exercise as providing discipline over the baseline calibration—not as delivering tightly identified point estimates of every structural parameter. The most reliable inference concerns *qualitative* comparisons: the ordering of the composite impact-stabilisation product $w_{\mathrm{vol},k}\nu_k$ across types, the sign and timing of the Z-score reversal, and the direction of the 46% fit improvement from baseline.

Qualitatively, the model reproduces the key LP patterns: impact stabilisation, the cross-sectional ordering $z_{75} > z_{50} > z_{25}$, NIM compression, and the medium-horizon sign reversal. Quantitatively, discrepancies remain at longer horizons ($h = 5$–6), where the LP troughs are deeper than the model can generate within a first-order linearised, single-shock framework.

6.11 Counterfactual Experiments

To assess the independent contribution of each structural mechanism, we conduct five counterfactual experiments at the baseline calibration vector (Table 10). In each experiment, a single mechanism is modified while the rest of the calibration is held fixed.

Table 11. Counterfactual impacts (period $h = 0$) and the efficiency gap

Experiment	$z_{25}(0)$	$z_{50}(0)$	$z_{75}(0)$	$z_{75} - z_{25}$
(1) Baseline	+0.0106	+0.0432	+0.1026	+0.0920
(2) No deposit market power $\beta_{D,k} = 1$	+0.0095	+0.0424	+0.1019	+0.0924
(3) Homogeneous efficiency (avg parameters)	+0.0483	+0.0483	+0.0483	**0.0000**
(4) No risk-management channel $\nu_k = 0$	−0.0317	−0.0326	−0.0334	−0.0017
(5) Higher capital requirement $\bar{\kappa} = 0.15$	+0.0106	+0.0432	+0.1027	+0.0920

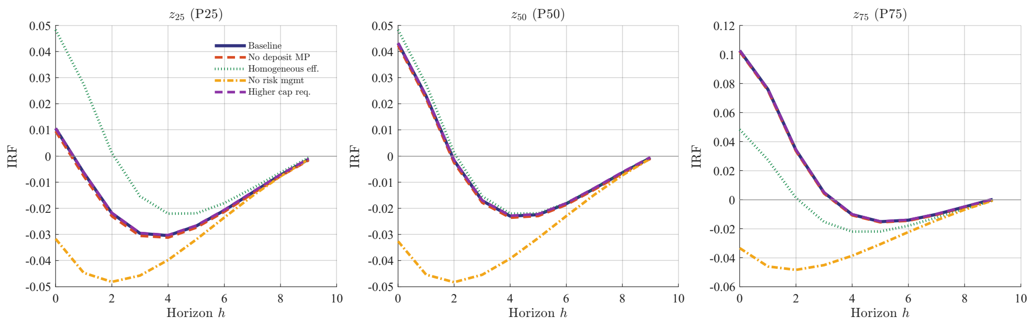

Figure 16. Bank stability (Z-score) under structural counterfactuals.

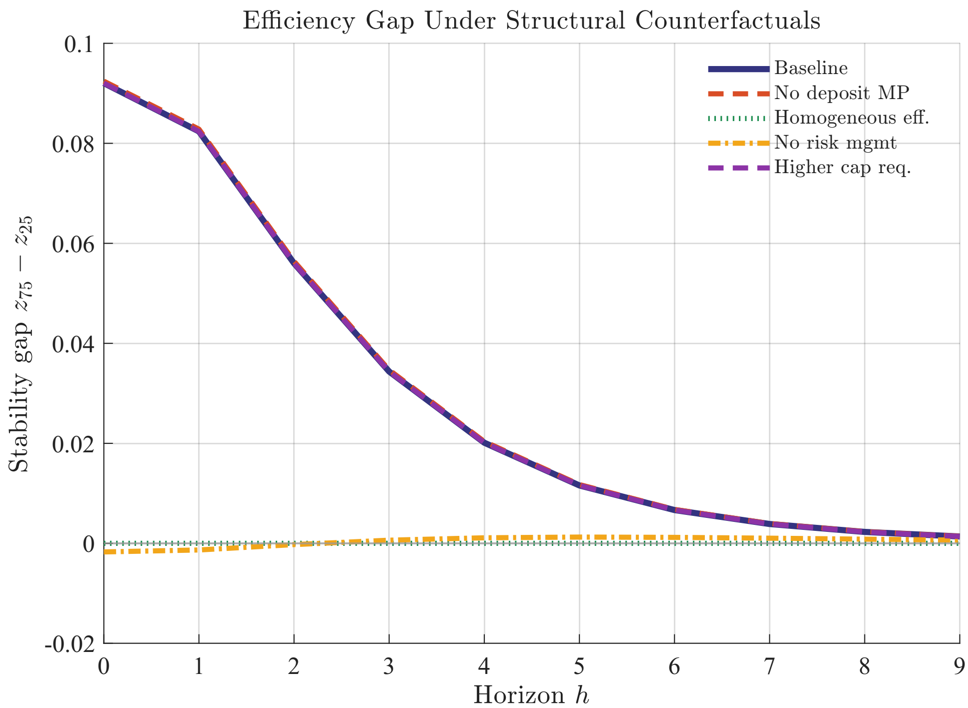

Figure 17. Efficiency gap ($z_{75} - z_{25}$) under structural counterfactuals.

Table 11 reports the impact Z-score responses and the efficiency gap for each experiment, Figure 16 plots the full IRF paths, and Figure 17 traces the efficiency gap over the horizon. Four main findings emerge.

- **Homogeneous efficiency (Experiment 3) collapses the efficiency gap to exactly zero.** The cross-sectional ordering of stability that motivates the entire paper is generated, in the model, by the heterogeneity in $(\nu_k, \chi_k, a_{m,k}, \rho_{L,k}, \beta_{D,k}, w_{\text{vol},k})$. When all of these are averaged, banks become indistinguishable, and the gap vanishes at every horizon.

- **Switching off the risk-management channel (Experiment 4) flips the impact response negative for all three types.** The empirical *impact stabilisation*—monetary tightening initially raises the Z-score—is therefore a structural consequence of the risk-taking channel of monetary policy operating through bank monitoring. Without it, banks see their stability deteriorate immediately as borrower distress builds.

- **Removing deposit market power (Experiment 2) deepens NIM compression**: $\mu_k(0)$ moves from $(-0.009, -0.019, -0.034)$ to $(-0.016, -0.024, -0.038)$. This is the Drechsler insight in reverse: the deposit franchise (markdown $\beta^*_{D,k} < 1$) acts as a partial cushion against funding-cost increases. The stability gap widens slightly, indicating that deposit market power has a modest stabilising role.

- **Higher capital requirements (Experiment 5) have small near-term effects** because the constraint is not strictly binding in steady state and the linearised lending rule loads on $\hat{e}_{k,t}$ rather than the multiplier directly. This is a known feature of the linearised CSV/capital-constraint approach and would be mitigated by a non-linear solution.

Mapping to the empirical LP evidence. Taken together, the structural impulse responses and counterfactual experiments replicate the two central empirical facts documented in the local-projection analysis: monetary tightening compresses net interest margins—most persistently for high-efficiency banks—yet the stability response is more favourable and smoother for these banks. In the model, this wedge arises because cost efficiency primarily determines the strength and effectiveness of the risk-management response, thereby reducing default risk and mitigating expected losses. The counterfactuals confirm that cross-sectional resilience is driven primarily by a loss-mitigation (operational-buffer) channel rather than by the faster recovery of post-tightening margins. The homogeneous-efficiency experiment demonstrates that the efficiency gap is attributable to heterogeneous bank primitives, and the no-risk-management experiment shows that the impact stabilisation documented in the LP evidence is a structural consequence of the risk-taking channel operating through monitoring. Thus the DSGE framework provides a disciplined structural rationale for the horizon-dependent sign reversals and the efficiency-conditioned smoothing observed in the reduced-form estimates.

7 CONCLUSION AND POLICY DISCUSSIONS

This paper re-examines the relationship between monetary policy and banking stability through a lens that has received limited direct attention in the empirical transmission literature: bank cost efficiency as a fundamental governing risk-control capacity, loss absorption, and the smoothness of medium-run stability dynamics. The motivating gap is that existing work on the risk-taking channel and bank heterogeneity has largely focused on balance-sheet buffers (capitalisation, liquidity, or funding structure), while treating operational efficiency either as secondary or as an outcome rather than a state variable that conditions the stability response to monetary policy shocks. We address this gap by integrating a cross-country efficiency framework into the identification and interpretation of monetary policy effects on bank stability, and by explicitly linking reduced-form evidence to a structural mechanism.

Empirically, we implement a two-stage design. In the first stage, we recover bank-level cost efficiency using a stochastic translog cost frontier, and then construct metafrontier efficiency measures that map country-specific technologies to broader technology sets (global and group-based benchmarks). This provides a comparable cross-country measure of the managerial/operational components of bank resilience, which is applicable across heterogeneous banking systems. In the second stage, we estimate the effect of monetary policy on banking stability, controlling for bank, time, and country fixed effects. We find that, overall, monetary tightening increases banking stability, with the results robust across different sub-samples, estimation techniques, and controls for macroprudential policies. Moreover, we also estimate the dynamic effects of monetary policy on bank stability and its underlying channels using local projections (LP), allowing for flexible impulse responses and horizon-by-horizon inference. These results demonstrate that monetary policy shocks have economically significant effects on stability and key intermediating margins, and that these effects are heterogeneous, consistent with cost efficiency shaping banks' capacity to manage risk and smooth stability dynamics in the medium term.

Finally, we develop a DSGE model with bank heterogeneity to provide a structural interpretation of the reduced-form facts. The model combines a standard New Keynesian core with a slow-moving borrower distress state and a fast risk-management channel, whose responsiveness and effectiveness increase with cost efficiency. In the model, a monetary tightening compresses net interest margins on impact through repricing frictions, while simultaneously inducing risk-management adjustments that reduce default probabilities and improve the credit composition. The interaction of these forces generates stability dynamics that match the qualitative pattern observed in the data and clarifies the conditions under which stability improvements are front-loaded versus delayed. The structural exercise, therefore, serves two roles: it rationalises why sign reversals can occur at particular horizons in the reduced-form estimates, and it formalises the interpretation of cost efficiency as a foundational determinant governing the smoothness of the medium-run stability response.

The policy implications are immediate. First, monetary policy and financial stability cannot be evaluated through a single monotone "tightening is stabilising" or "tightening is destabilising" narrative; the net effect depends on the regime and the composition of bank balance sheets, as well as on banks' operational capacity to manage risk. Second, supervisory assessments should treat cost efficiency as more than an operational performance metric: it contains information about risk-control capability and shock absorption that is relevant for macro-financial resilience. This implies that macroprudential policy and supervision can improve the stability consequences of monetary tightening by strengthening the operational foundations of risk management—through governance, monitoring technologies, and the organisational capacity to reprice, rebalance, and provision promptly. Third, because efficiency and market power are distinct and benchmark-dependent concepts in a cross-country setting, policy frameworks should avoid conflating competitive conditions with operational resilience. The evidence supports the design of macroprudential buffers and supervisory intensity to account for efficiency heterogeneity, even among banks with similar capital and liquidity positions.

Overall, the paper presents a unified empirical-structural account of how monetary policy is transmitted to banking stability when banks differ in cost efficiency. By combining a cross-country stochastic metafrontier measurement of efficiency with dynamic LP inference, IV identification, and a mechanism-consistent DSGE interpretation, we provide evidence that efficiency conditions the stability response to policy in a meaningful way over the medium term. The broader implication is that improving the operational and managerial foundations of intermediation can strengthen the stabilising component of the risk-taking channel and reduce the likelihood that tightening episodes propagate into fragility through valuation losses and credit deterioration.

REFERENCES

Viral Acharya, Robert Engle, and Matthew Richardson. Capital shortfall: A new approach to ranking and regulating systemic risks. *American Economic Review*, 102(3):59–64, 2012.

Manuel Adelino, Miguel A Ferreira, Mariassunta Giannetti, and Pedro Pires. Trade credit and the transmission of unconventional monetary policy. *The Review of Financial Studies*, 36(2): 775–813, 2023.

Zohair Alam, Adrian Alter, Jesse Eiseman, Gaston Gelos, Heedon Kang, Machiko Narita, Erlend Nier, and Naixi Wang. Digging deeper—evidence on the effects of macroprudential policies from a new database. *Journal of Money, Credit and Banking*, 57(5):1135–1166, 2025.

Yener Altunbas, Santiago Carbo, Edward PM Gardener, and Philip Molyneux. Examining the relationships between capital, risk and efficiency in european banking. *European Financial Management*, 13(1):49–70, 2007.

Yener Altunbas, Leonardo Gambacorta, and David Marques-Ibanez. Does monetary policy affect bank risk-taking? *International Journal of Central Banking*, pages 95–135, 2014.

Yener Altunbas, Mahir Binici, and Leonardo Gambacorta. Macroprudential policy and bank risk. *Journal of International Money and Finance*, 81:203–220, 2018.

George E Battese and Tim J Coelli. Prediction of firm-level technical efficiencies with a generalized frontier production function and panel data. *Journal of econometrics*, 38(3): 387–399, 1988.

George Edward Battese and Tim J Coelli. A model for technical inefficiency effects in a stochastic frontier production function for panel data. *Empirical Economics*, 20(2):325–332, 1995.

Thorsten Beck, Olivier De Jonghe, and Glenn Schepens. Bank competition and stability: Cross-country heterogeneity. *Journal of Financial Intermediation*, 22(2):218–244, 2013.

Allen N Berger and Robert DeYoung. Problem loans and cost efficiency in commercial banks. *Journal of Banking & Finance*, 21(6):849–870, 1997.

Allen N Berger and David B Humphrey. Efficiency of financial institutions: International survey and directions for future research. *European Journal of Operational Research*, 98(2):175–212, 1997.

Allen N Berger and Loretta J Mester. Inside the black box: What explains differences in the efficiencies of financial institutions? *Journal of Banking & Finance*, 21(7):895–947, 1997.

Ben Bernanke, Mark Gertler, and Simon Gilchrist. The financial acceletator and the flight to quality. *The Review of Economics and Statistics*, 78(1):1–15, 1996.

Ben S. Bernanke. Monetary policy and the housing bubble. Speech at the Annual Meeting of the American Economic Association, Atlanta, GA, January 2010. URL https://www.federalreserve.gov/newsevents/speech/bernanke20100103a.htm. Accessed: 2025-03-08.

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

William Bernhard. *Banking on Reform: Political Parties and Central Bank Independence in the Industrial Democracies*. University of Michigan Press, 2002.

Claudio Borio. The search for the elusive twin goals of monetary and financial stability. *National Institute Economic Review*, 192(1):84–101, 2005. Keynote lecture at the IGIDR Sixth Annual Conference on Money and Finance in the Indian Economy, 25–27 March 2004.

Claudio Borio and Andrew Crockett. In search of anchors for financial and monetary stability. *Greek Economic Review*, 20(2):1–14, 2000.

Claudio Borio and Haibin Zhu. Capital regulation, risk-taking and monetary policy: a missing link in the transmission mechanism? *Journal of Financial Stability*, 8(4):236–251, 2012.

Christian Brownlees and Robert F Engle. Srisk: A conditional capital shortfall measure of systemic risk. *The Review of Financial Studies*, 30(1):48–79, 2017.

Alessandro Carretta, Vincenzo Farina, Franco Fiordelisi, Paola Schwizer, and Francesco Saverio Stentella Lopes. Don't stand so close to me: The role of supervisory style in banking stability. *Journal of Banking & Finance*, 52:180–188, 2015.

Eugenio Cerutti, Stijn Claessens, and Luc Laeven. The use and effectiveness of macroprudential policies: New evidence. *Journal of Financial Stability*, 28:203–224, 2017.

Minghua Chen, Ji Wu, Bang Nam Jeon, and Rui Wang. Monetary policy and bank risk-taking: Evidence from emerging economies. *Emerging Markets Review*, 31:116–140, 2017.

Richard Clarida, Jordi Gali, and Mark Gertler. Monetary policy rules and macroeconomic stability: evidence and some theory. *The Quarterly Journal of Economics*, 115(1):147–180, 2000.

Richard Davies and Belinda Tracey. Too big to be efficient? the impact of implicit subsidies on estimates of scale economies for banks. *Journal of Money, Credit and Banking*, 46(s1): 219–253, 2014.

Ferre De Graeve, Thomas Kick, and Michael Koetter. Monetary policy and financial (in) stability: An integrated micro–macro approach. *Journal of Financial Stability*, 4(3):205–231, 2008.

Asli Demirgüç-Kunt and Harry Huizinga. Bank activity and funding strategies: The impact on risk and returns. *Journal of Financial Economics*, 98(3):626–650, 2010.

Michel Dietsch and Ana Lozano-Vivas. How the environment determines banking efficiency: A comparison between french and spanish industries. *Journal of Banking & Finance*, 24(6): 985–1004, 2000.

Itamar Drechsler, Alexi Savov, and Philipp Schnabl. The deposits channel of monetary policy. *The Quarterly Journal of Economics*, 132(4):1819–1876, 2017.

Richard Adjei Dwumfour, Eric Fosu Oteng-Abayie, and Emmanuel Kwasi Mensah. Bank efficiency and the bank lending channel: new evidence. *Empirical Economics*, 63(3):1489–1542, 2022.

Franco Fiordelisi, David Marques-Ibanez, and Phil Molyneux. Efficiency and risk in european banking. *Journal of Banking & Finance*, 35(5):1315–1326, 2011.

Steven Fries and Anita Taci. Cost efficiency of banks in transition: Evidence from 289 banks in 15 post-communist countries. *Journal of Banking & Finance*, 29(1):55–81, 2005.

Jordi Galí. *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications.* Princeton University Press, 2015.

Ana Carolina Garriga. Central bank independence in the world: A new data set. *International Interactions*, 42(5):849–868, 2016.

Mark Gertler and Nobuhiro Kiyotaki. Financial intermediation and credit policy in business cycle analysis. In *Handbook of Monetary Economics*, volume 3, pages 547–599. Elsevier, 2010.

Iftekhar Hasan and Katherin Marton. Development and efficiency of the banking sector in a transitional economy: Hungarian experience. *Journal of Banking & Finance*, 27(12):2249–2271, 2003.

Robert J Hodrick and Edward C Prescott. Postwar us business cycles: an empirical investigation. *Journal of Money, Credit, and Banking*, pages 1–16, 1997.

Joel F Houston, Chen Lin, Ping Lin, and Yue Ma. Creditor rights, information sharing, and bank risk taking. *Journal of Financial Economics*, 96(3):485–512, 2010.

Cliff J Huang, Tai-Hsin Huang, and Nan-Hung Liu. A new approach to estimating the metafrontier production function based on a stochastic frontier framework. *Journal of Productivity Analysis*, 42:241–254, 2014.

Joseph P Hughes and Loretta J Mester. A quality and risk-adjusted cost function for banks: Evidence on the "too-big-to-fail" doctrine. *Journal of Productivity Analysis*, 4(3):293–315, 1993.

Lucas Husted, John Rogers, and Bo Sun. Monetary policy uncertainty. *Journal of Monetary Economics*, 115:20–36, 2020.

Vasso Ioannidou, Steven Ongena, and José-Luis Peydró. Monetary policy, risk-taking, and pricing: Evidence from a quasi-natural experiment. *Review of Finance*, 19(1):95–144, 2015.

Gabriel Jiménez, Steven Ongena, José-Luis Peydró, and Jesús Saurina. 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, 2014.

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

Ruby P Kishan and Timothy P Opiela. Bank size, bank capital, and the bank lending channel. *Journal of Money, Credit and Banking*, pages 121–141, 2000.

Simon Kwan and Robert A Eisenbeis. Bank risk, capitalization, and operating efficiency. *Journal of Financial Services Research*, 12(2):117–131, 1997.

Martien Lamers, Frederik Mergaerts, Elien Meuleman, and Rudi Vander Vennet. The trade-off between monetary policy and bank stability. *International Journal of Central Banking*, 15(2): 1–42, 2019.

Ivan Lim, Jens Hagendorff, and Seth Armitage. The effects of regulatory office closures on bank behavior. *Journal of Money, Credit and Banking*, 2023.

Angela Maddaloni and José-Luis Peydró. Bank risk-taking, securitization, supervision, and low interest rates: Evidence from the euro-area and the us lending standards. *the review of financial studies*, 24(6):2121–2165, 2011.

Kiminori Matsuyama. Credit traps and credit cycles. *American Economic Review*, 97(1):503–516, 2007.

Frederic S Mishkin. *The economics of money, banking, and financial markets*. Pearson education, 2016.

Karsten Müller, Chenzi Xu, Mohamed Lehbib, and Ziliang Chen. The global macro database: A new international macroeconomic dataset. Working Paper 33714, National Bureau of Economic Research, April 2025. URL http://www.nber.org/papers/w33714.

Sander Oosterloo and Jakob De Haan. Central banks and financial stability: a survey. *Journal of Financial Stability*, 1(2):257–273, 2004.

Raghuram G Rajan. Has finance made the world riskier? *European Financial Management*, 12 (4):499–533, 2006.

Morten O Ravn and Harald Uhlig. On adjusting the hodrick-prescott filter for the frequency of observations. *Review of Economics and Statistics*, 84(2):371–376, 2002.

Andrew Donald Roy. Safety first and the holding of assets. *Econometrica*, pages 431–449, 1952.

Calvin W Sealey Jr and James T Lindley. Inputs, outputs, and a theory of production and cost at depository financial institutions. *The Journal of Finance*, 32(4):1251–1266, 1977.

Anastasiya Shamshur and Laurent Weill. Does bank efficiency influence the cost of credit? *Journal of Banking & Finance*, 105:62–73, 2019.

Lars EO Svensson. Inflation targeting. In *Handbook of Monetary Economics*, volume 3, pages 1237–1302. Elsevier, 2010.

John B Taylor. Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, volume 39, pages 195–214. Elsevier, 1993.

John B Taylor et al. Getting back on track: macroeconomic policy lessons from the financial crisis. *Federal Reserve Bank of St. Louis Review*, 92(3):165–176, 2010.

André Uhde and Ulrich Heimeshoff. Consolidation in banking and financial stability in europe: Empirical evidence. *Journal of Banking & Finance*, 33(7):1299–1311, 2009.

Jonathan Williams. Determining management behaviour in european banking. *Journal of Banking & Finance*, 28(10):2427–2460, 2004.

Michael Woodford and Carl E Walsh. Interest and prices: Foundations of a theory of monetary policy. *Macroeconomic Dynamics*, 9(3):462–468, 2005.

J Yellen. Monetary policy and financial stability. the 2014 michel camdessus central banking lecture. *International Monetary Fund, Washington, DC*, 2014.