Who Bears the Burden of Bank Capital Regulation?

Buffers, Risk Weights, and the Supply of Credit

Juan Zurita*

This draft: August 19, 2026

Job Market Paper

Click here for the latest version

Abstract

Who bears the burden of bank capital regulation? Using narratively identified capital-requirement shocks and a quarterly panel of U.S. banks (1994–2019), we document two cross-sectional facts: within banks, commercial lending contracts more than real estate lending, reflecting differential risk weights; across banks, institutions with thin voluntary buffers cut lending most. A heterogeneous macro-banking model with precautionary buffers, solved globally and estimated on within-bank impulse responses, rationalises both facts. A one-percentage-point tightening costs 2.4–2.9 percent of credit-market surplus. Risk weights and size tiers redistribute this burden but barely change its aggregate magnitude—a gross transition cost excluding the stability benefits of capital.

JEL classification: E44, E58, G21, G28.

Keywords: Capital Requirements, Capital Buffers, Risk Weights, Credit Supply, Banking Heterogeneity, Macroprudential Policy

*School of Economics, University of Edinburgh. Email: juan.zurita@ed.ac.uk. This paper has greatly benefited from discussions with Corrado Di Guilmi, Mengheng Li, Peter Docherty, Jorge Miranda Pinto, Mario Fiorini, Adeline Delavande, Peter Siminski, David Goldbaum, Timothy Kam, Mei Dong, Shawn Chen, James Graham, and Rustam Jamilov, as well as participants at the Econometric Society Australasian Meeting 2023, the 2022 Financial Markets and Corporate Governance at Monash University and the WEAI Conference Melbourne 2023. I am grateful to Urfat Isayev for excellent research assistance. Any remaining errors are my own responsibility.

1 Introduction

Since the global financial crisis, bank capital requirements have risen substantially, and the new regime is deliberately differentiated: business loans carry higher regulatory risk weights than residential mortgages, and statutory requirements step up at the $10 billion and $50 billion asset thresholds. Both margins presume that regulators know where the burden of capital regulation falls. This paper asks who in fact bears it—a question about the transmission of macroprudential policy whose answer determines both the sectoral composition of any credit contraction and the aggregate cost of the regime. The answer, we find, is organised by two variables that appear on every bank's balance sheet but in no statute: the risk-weight composition of the asset portfolio, which determines which *sectors* absorb the contraction in credit, and the voluntary capital buffer a bank holds above its requirement, which determines which *banks* do. The aggregate cost of a tightening, by contrast, is governed almost entirely by the level of requirements rather than by their design.

We reach this answer in two steps. Empirically, we assemble a quarterly panel of roughly 4,800 U.S. banks from the Call Reports (1994–2019) and trace the narratively identified capital-requirement shocks of Drechsel & Miura (2025) through the cross-section of lending using local projections. Because the identified innovation is an aggregate surprise, our regressions exploit within-bank variation: bank fixed effects and macroeconomic controls absorb the common cycle, lending shows no pre-trends ahead of the identified innovations, and consumer credit—fully risk-weighted but funded against high interest margins—serves as a placebo throughout. Structurally, we develop a heterogeneous macro-banking model in which banks facing funding-market frictions hoard precautionary capital buffers, so the cross-sectional distribution of regulatory slack is an equilibrium object rather than a primitive. The regulatory constraint binds only occasionally and policy functions are kinked at the frontier, so we solve the model globally with deep-learning methods and estimate the requirement process by simulated method of moments on the sixteen within-bank impulse-response moments.

Two facts emerge, one across assets and one across banks. Within banks, commercial and industrial (C&I) lending—which carries the higher regulatory risk weight—contracts about forty percent more than real estate lending after a tightening, with hump-shaped responses that trough three to four years out; consumption and credit-card lending do not respond at any horizon. Across banks, transmission is governed by the *distance to the regulatory frontier*: sorting institutions each quarter into terciles of their lagged regulatory slack, banks in the thinnest-buffer tercile cut C&I lending 1.29 times more than banks

in the thickest. The gradient survives within size classes, and it disappears when tercile membership is frozen at pre-sample values—what matters is the buffer a bank holds when the shock arrives, not a fixed bank type. Scaled through the estimated model, one unit of the identified innovation moves required capital by four basis points, and the innovations are effectively permanent, consistent with post-crisis regulation being a sequence of regime changes rather than transitory shocks.

The economics behind both facts is a single mechanism. A bank that cannot cheaply replace funds at the margin—because its deposit franchise is exhausted and wholesale markets penalise it—self-insures against regulatory breach by holding equity above its requirement, so the cross-sectional distribution of buffers mirrors the distribution of funding frictions. When a tightening arrives, the regulatory constraint binds only for banks whose buffer it exhausts: for them the shadow cost of capital spikes, and the cheapest source of regulatory relief is to shed the assets that carry the highest risk weight—C&I loans—because each dollar cut frees the most required capital. Transmission is therefore state-dependent across banks and asymmetric across sectors for the same underlying reason: the price of regulatory capital differs across balance sheets and across asset classes.

The estimated model reproduces the sign of all sixteen targeted moments, the sectoral ordering at every horizon, and the hump-shaped dynamics. More demandingly, it delivers a low-to-high buffer-tercile response ratio of 1.23 against 1.29 in the data—an *untargeted* moment that validates the buffer mechanism at the model's core. The estimated model then serves as a laboratory for regulatory experiments that the historical record cannot identify. Three quantitative results follow. Converted to a permanent one-percentage-point increase in required capital, lending contracts by 3.5% in C&I, 3.0% in real estate, and 3.1% in total—inside the range of quasi-experimental estimates from UK, euro-area, and Belgian settings. The welfare cost of that tightening is 2.4–2.9% of steady-state credit-market surplus—the sum of borrower and intermediary surplus in the loan market—and essentially all of it flows through the contraction in the scale of intermediation. Counterfactuals that flatten risk weights, or that replace the size-tiered schedule with a uniform requirement raising the same aggregate capital, redistribute the burden but leave its total nearly unchanged. Finally, transmission is non-linear in the state: the buffers that insulated most banks against the moderate innovations of 1994–2019 are a finite resource, and shocks large enough to erode them trigger contractions disproportionate to the shock.

The policy content is immediate. Debates over the architecture of capital requirements—risk weights, size tiers—are, to first order, debates about incidence rather than aggregate cost. Supervisory attention to the distribution of voluntary buffers identifies where a tightening will bind more accurately than

statutory size categories do, and the same logic implies that buffer releases in downturns are most potent for institutions operating near the frontier. Because the framework is partial equilibrium and prices no stability benefits of capital, our welfare figures are gross transition costs: informative about the burden of tightening and its distribution, not about optimal requirement levels.

Our results connect three literatures. First and closest, a quantitative macro-banking literature studies capital regulation in economies with dynamic, heterogeneous banks: Corbae & Erasmo (2021), Begenau (2020), Elenev *et al.* (2021), and Mendicino *et al.* (2020). These papers characterise optimal requirements in general equilibrium; our contribution is to discipline the cross-section directly—the model is estimated on within-bank impulse responses to identified policy innovations and validated on an untargeted cross-sectional moment—which is what allows us to speak to the *incidence* of a given tightening, at the cost of abstracting from default and its stability benefits. Relative to Corbae & Erasmo (2021), who model a single loan sector in a two-bank-type economy with exogenous default, we introduce two sectors carrying distinct risk weights (mortgages at $\omega_M = 0.50$ versus C&I at $\omega_B = 1.00$) so that the model endogenously generates within-bank asset reallocation, estimate the policy process via within-bank local projections rather than calibrating to aggregate moments, and solve the model globally with neural-network function approximation, which allows us to preserve the full cross-sectional distribution of buffers that drives incidence. Second, a quasi-experimental literature measures the credit-supply effects of requirement changes: Aiyar *et al.* (2016) for UK bank-specific requirements, Gropp *et al.* (2020) for the EBA capital exercise, De Jonghe *et al.* (2020a) for Pillar-2 decisions, Jimenez *et al.* (2017) for dynamic provisioning, and Mésonnier & Monk (2015) and Conti *et al.* (2023) at the aggregate level; Acharya *et al.* (2022) provide a particularly granular anatomy of macroprudential transmission using Irish mortgage-level data. Relative to this work, we hold the policy event fixed and trace how the same identified innovation is absorbed across asset classes and across the buffer distribution, and we provide the structural mapping from those responses to aggregate welfare. Third, our funding-capacity evidence connects to the deposits-channel literature of Drechsler *et al.* (2017): banks with strong deposit franchises face little funding risk, and we show that they optimally hold thinner regulatory buffers, so the deposit franchise shapes not only monetary transmission but also the incidence of capital regulation. Methodologically, the global deep-learning solution follows Maliar *et al.* (2021).

The remainder of the paper proceeds as follows. Section 2 presents the data, the identification strategy, and the two cross-sectional facts. Section 3 develops the stylised model and Section 4 the quantitative framework. Section 5 reports the estimation, fit, and counterfactuals, and Section 6 draws the policy

implications. Section 7 concludes.

2 Empirical Evidence

The mechanism this paper studies is balance-sheet-driven. Unlike monetary policy, which influences credit conditions by shifting the base cost of funding, changes in capital requirements alter the quantity of equity required to support a marginal unit of risk-weighted assets, creating a direct link between a bank's capital position and its shadow cost of lending (Van den Heuvel *et al.*, 2002). A tightening forces a bank into a three-margin adjustment: raising new equity, shrinking the asset portfolio, or reallocating it toward lower-risk-weight assets. Because business loans carry higher regulatory weights than mortgages, the last margin makes high-weight sectors the efficient place to seek regulatory relief per dollar of credit curtailed (De Jonghe *et al.*, 2020a; Sivec & Volk, 2021; De Jonghe *et al.*, 2020b)—and because the ability to use the first two margins depends on a bank's funding structure, the incidence of a given tightening depends on the cross-sectional distribution of funding frictions and voluntary buffers. Isolating this balance-sheet transmission from broader monetary cycles is the task of the empirical design below; formalising it is the task of the model in Section 3.

2.1 Data and Stylized Facts

Our primary data source is the Consolidated Reports of Condition and Income (Call Reports) provided by the FDIC. The sample spans from 1994 to 2019 and includes quarterly balance sheet data for approximately 4,800 U.S. commercial banks—effectively the *population* of FDIC-insured institutions, not a sample drawn from it. This granular dataset allows us to decompose total loans and leases into specific categories: real estate, commercial and industrial (C&I), consumption, and credit card loans. We complement these micro-data with macroeconomic indicators from the Federal Reserve Bank of St. Louis (FRED), specifically real GDP growth and the effective federal funds rate.

When two institutions merge, we retain only the surviving entity's post-merger balance sheet; the acquired bank's final pre-merger observation is dropped so that the mechanical jump in the survivor's assets does not contaminate growth rates. All bank-level variables entering regressions are winsorised at the 1st and 99th percentiles of their quarterly cross-sectional distributions, which further attenuates any residual discontinuities from mergers. As a robustness check, we re-estimate the headline local projections after

excluding all bank-quarters within one year of a completed acquisition; the point estimates are virtually unchanged (Online Appendix).

To analyze the transmission of policy across the banking hierarchy, we partition our sample into three cohorts based on total asset size, following the regulatory thresholds established by the Dodd-Frank Act (2010) and the EGRRCPA (2018). We define Small (Community) Banks as those with assets below $10 billion, Medium (Regional) Banks as those between $10 billion and $50 billion, and Large (Systemically Important) Banks as those exceeding $50 billion. These thresholds are not merely statistical conveniences; they reflect the points at which banks become subject to increasingly stringent stress testing, enhanced prudential standards, and closer capital oversight by the Federal Reserve.

For the time-series analysis, we aggregate individual bank-level data to these categories by summing the levels of loans and equity. For regulatory ratios, such as the total risk-based capital ratio, we compute weighted averages where each institution's share of total equity serves as the weight. This ensures that our aggregate measures reflect the systemic capitalization of each sector and that the behavior of systemically important institutions is accurately captured, rather than being obscured by an unweighted average of thousands of smaller entities.

Table I provides summary statistics for our main variables. The high standard deviation in Total Assets ($1.4 billion mean vs. $29.5 billion SD) underscores the extreme heterogeneity in the U.S. banking sector, necessitating our size-based cross-sectional analysis.

Central to the mechanism this paper studies is the role of asset composition and risk-weighting: the "regulatory wedge" created by risk-weighted assets (RWA). Acharya *et al.* (2014) document that regulatory risk weights can diverge substantially from market-implied measures of risk, making the mapping from balance-sheet composition to effective capital charges a first-order determinant of bank behaviour. Figure 1 shows that the RWA-to-total-asset ratio rises systematically with bank equity—larger banks hold portfolios with higher average risk weights, while small banks' regulatory position is driven by portfolio composition (Le Leslé & Avramova, 2012). Following the Global Financial Crisis this wedge has widened, with smaller banks increasingly concentrating in real estate while larger banks maintain diversified exposures. The stylized facts below unpack the components of this wedge.

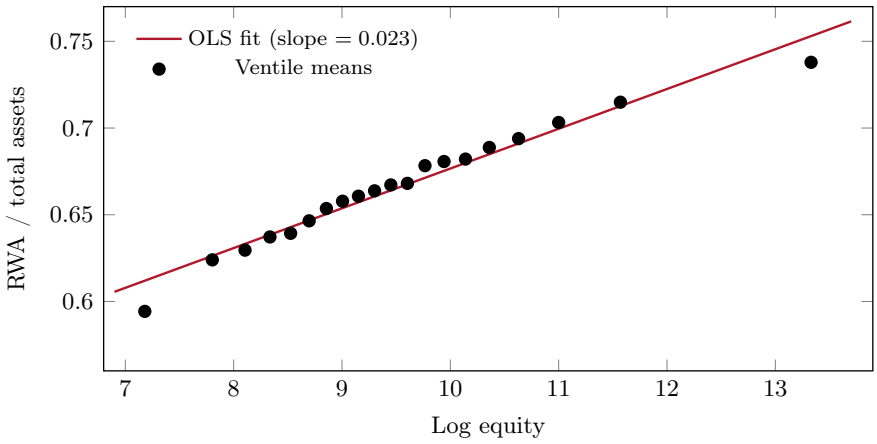


FIGURE 1.—The role of weight effects

Notes: Binned scatter: each point is the mean risk-weighted-assets-to-total-assets ratio within a ventile of log equity (899,168 bank-quarter observations, 1994–2020); the line is the OLS fit on the underlying bank-quarter data (slope 0.023). The positive slope is the "regulatory wedge": larger banks hold portfolios with higher average risk weights, while small banks' regulatory position is driven by portfolio composition. Source: US Call Reports.

Stylized Fact 1: Market Concentration. Figure 2 illustrates a significant structural shift in the U.S. credit market. While the number of institutions has declined, market participation has become increasingly concentrated among large banks. This suggests that the aggregate response to capital requirements is increasingly dominated by the behavior of a few systemically important lenders.

Figure 2 reveals a crucial link between bank size and asset composition. Small banks exhibit high loan-type concentration, with mortgage lending constituting over 90% of their total portfolios. Conversely, large banks maintain significantly more diversified balance sheets, with real estate representing only 40% of their lending. As mortgages and business loans carry different regulatory risk weights (ω), this specialization implies that macroprudential shocks will have inherently asymmetric effects across the banking distribution. Banks with specialized portfolios have fewer margins of adjustment when capital constraints bind, a feature we explore in our empirical identification and quantitative model.

TABLE I
SUMMARY STATISTICS (1994–2019)

Variable	Source	N	Mean	Std. Dev.
Panel A: Balance Sheet (Millions $)				
Total Assets	Call Reports	996 754	1446.19	29 500.00
Total Liabilities	Call Reports	996 754	1320.56	27 000.00
Total Equity	Call Reports	994 569	151.18	2949.80
Panel B: Regulatory Ratios and Asset Shares				
Capital / RWA (%)	Call Reports	948 614	22.61	243.41
Mortgage Loan Share (0-1)	Own Calc.	996 754	0.69	0.23
Business Loans Share (0-1)	Own Calc.	996 754	0.14	0.18

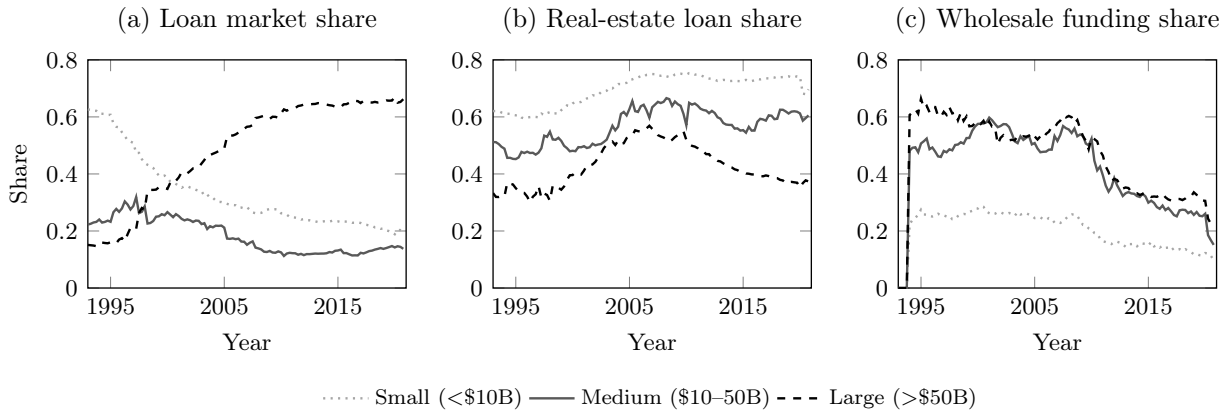


FIGURE 2.— Market concentration, asset specialization and funding structure

Notes: Panel (a) shows each size cohort's share of total bank loans; panel (b) the real-estate share of each cohort's loan portfolio; panel (c) wholesale funding as a share of non-equity funding by cohort. Cohorts split at $10B and $50B of lagged total assets. Quarterly, 1993–2020. Source: US Call Reports.

Stylized Fact 2: Asset Specialization. The importance of this sectoral divide is further illuminated by the structural divergence in bank portfolios over the last three decades. Figure 3 illustrates the "Asset Gap"—defined as the difference between real estate and business credit as a proportion of total lending—across the banking hierarchy.

Since the mid-1990s, a trend that accelerated sharply following the Global Financial Crisis, small and medium-sized banks have experienced a persistent widening of this gap. This indicates an increasing concentration in real estate at the expense of commercial and industrial (C&I) lending. In contrast, large banks have maintained a significantly narrower gap, reflecting more diversified balance sheets and a broader range of lending activities.

From a regulatory standpoint, this specialization is a critical determinant of policy transmission. Because real estate and business loans carry distinct risk weights (ω_{RE} vs. ω_{BIZ}), a bank's capacity to buffer a capital shock depends heavily on its portfolio flexibility and funding capacity. While a diversified large institution can pivot between asset classes to optimize its risk-weighted assets (RWA), specialized small banks face narrower margins of adjustment when capital constraints bind. This empirical specialization provides the foundational rationale for the heterogeneous "sectoral lending channel" that we explore in both our empirical identification and quantitative model.

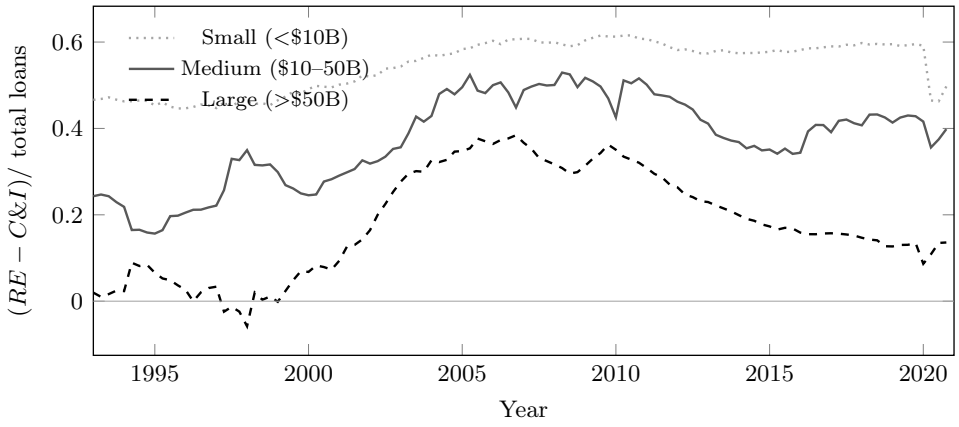


FIGURE 3.— Difference between real estate and business credit

Notes: This figure shows the difference between real estate and business credit as a proportion of total lending for small, medium and large banks, 1993–2020. Cohorts split at $10B/$50B of lagged total assets. Source: US Call Reports.

Stylized Fact 3: The Regulatory Wedge and Capital Volatility. The distinction between a bank's accounting leverage and its regulatory capital position is central to understanding the impact of macroprudential shocks. Figure 4 compares the equity-to-asset ratio (the leverage ratio, panel a) with the equity-to-risk-weighted asset ratio (the regulatory capital ratio, panel b) for the three bank cohorts.

Two key observations emerge from this comparison. First, while small banks often appear more highly capitalized than large banks on an unweighted basis (panel a), this advantage is compressed when accounting for risk-weighting (panel b). This confirms that the regulatory burden is not uniformly distributed: the transition from total assets to risk-weighted assets (RWA) creates a "regulatory wedge" that is most binding for institutions with specialized, high-risk-weight portfolios.

Second, the equity-to-RWA ratio exhibits significantly higher volatility than the simple leverage ratio. This discrepancy underscores the significance of the risk-weighting mechanism (ω) as a driver of capital-requirement shocks. Because the denominator of the regulatory ratio is sensitive to shifts in both

asset composition and regulatory risk-classification, banks must manage their portfolios dynamically to maintain compliance. The higher sensitivity of the weighted ratio suggests that banks with smaller equity or less diversified portfolios have fewer degrees of freedom to absorb unanticipated changes in capital requirements, a mechanism we formalize in our theoretical framework.

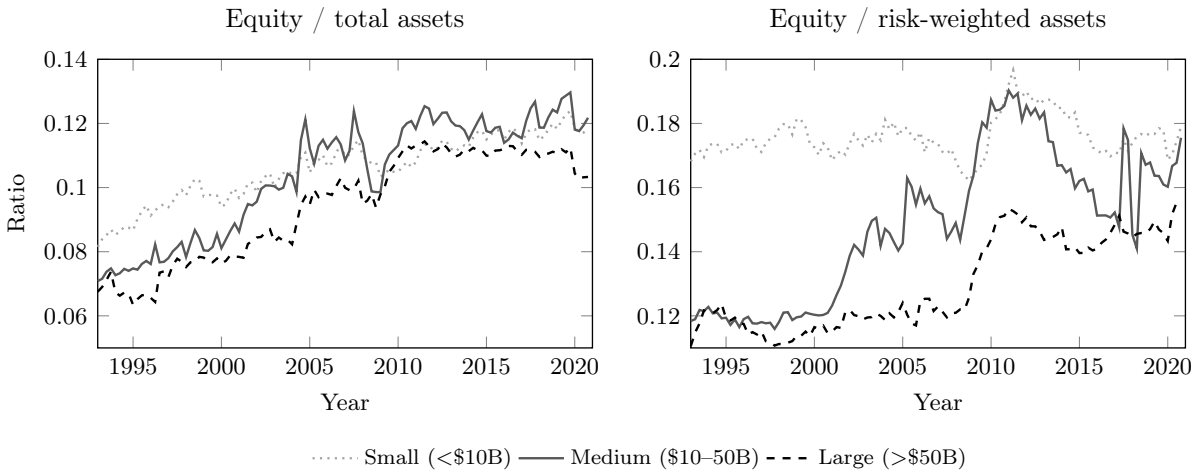


FIGURE 4.—Leverage and regulatory capital ratios by bank size

Notes: Panel (a) plots the equity-to-total-asset ratio (leverage ratio) and panel (b) the equity-to-risk-weighted-asset ratio (equity-weighted within cohort) for small, medium and large banks, 1993–2020. Source: US Call Reports.

2.2 The Evolution of Capital Regulation in the United States

Capital regulation in the United States has undergone a structural transformation over the past four decades, shifting from a microprudential focus on individual bank solvency to a macroprudential framework centered on systemic stability. This evolution provides a rich source of variation for identifying the impact of regulatory shocks on the lending channel.

The modern era of U.S. capital regulation began with the 1992 implementation of the Basel I Accord, which established uniform minimum capital ratios based on credit-risk-weighted assets. However, Basel I's coarse risk buckets eventually proved insufficient for capturing the complexities of off-balance-sheet exposures and market risk. While the Basel II framework attempted to introduce more granular, internal ratings-based (IRB) sensitivities in the mid-2000s, its implementation was interrupted by the 2008 financial crisis. The crisis revealed fundamental fragilities in bank capital structures-namely, an over-reliance on short-term wholesale funding and insufficient high-quality loss-absorbing capital.

In the wake of the crisis, the regulatory landscape shifted toward the *quantity* and *quality* of capital.

The Dodd-Frank Act (2010) and the subsequent U.S. implementation of Basel III (2013) introduced the Common Equity Tier 1 (CET1) requirements, capital conservation buffers, and a supplementary leverage ratio. Most importantly for our study, these reforms introduced "tailored" regulation: systemically important financial institutions (SIFIs) became subject to annual stress tests and higher capital surcharges, while the Economic Growth, Regulatory Relief, and Consumer Protection Act (EGRRCPA, 2018) later eased the compliance burden for community banks.

Table II summarizes the major regulatory events used in our empirical analysis. Following the identification strategy of Drechsel & Miura (2025), we focus on the "first mention" of these legislative acts in official speeches rather than their final enactment dates. This approach captures the pre-emptive behavior of banks as they adjust their balance sheets in anticipation of new requirements. As shown in the final column of Table II, banks exhibit significant equity capital accumulation between the initial announcement and the formal passing of the acts. For instance, equity capital grew by an average of 39% following the first mentions of the Riegle-Neal Act and by 27% surrounding the Dodd-Frank Act.

This pre-emptive accumulation of equity capital is central to our mechanism. It suggests that regulatory announcements trigger a significant balance sheet reorganization long before the legal enforcement date. In the following sections, we utilize these announcement-day surprises to identify the causal impact of capital requirement shifts on the cross-section of bank lending.

TABLE II
FINANCIAL REGULATION AND EQUITY CAPITAL ACCUMULATION

Legislative Act	**Passed**	**1st Mention**	**Mean Δ Equity**
Riegle-Neal Interstate Banking Act	09/13/1994	09/15/1992	+39%
Gramm-Leach-Bliley Act	11/12/1999	04/19/1998	+26%
Dodd-Frank Wall Street Reform Act	07/21/2010	07/24/2008	+27%
EGRRCPA (Regulatory Relief Act)	05/24/2018	09/28/2016	+7%

Notes: This table reports key U.S. legislative acts and the dates of the speeches in which they were first publicly introduced, following Drechsel & Miura (2025). Mean Δ Equity represents the average growth in bank equity between the first mention and enactment.

2.3 The Anatomy of the Mechanism: Funding Capacity and Regulatory Slack

While the macroeconomic history in the previous section establishes that regulatory announcements induce equity accumulation, the cross-sectional capacity to respond to these shifts depends fundamentally on a bank's liability structure. To open the black box of this mechanism before estimating our causal macroeconomic models, we utilize our quarterly bank-level data to examine the joint distribution of

regulatory cushions, funding alternatives, and dynamic adjustment choices.

Stylized Fact 4: The Funding-Buffer Trade-Off. Our theoretical framework posits that banks hold voluntary capital buffers as a precautionary cushion against regulatory breaches. Institutions with deep, flexible funding sources can afford to hold minimal excess capital because they can seamlessly raise external funding or issue equity under stress. Conversely, banks facing significant funding market frictions must hoard larger precautionary capital buffers to self-insure.

To document this in the data, we define a bank's continuous *Funding Capacity* as the share of total liabilities funded by core retail deposits and non-interest-bearing (NIB) deposits; its complement is the share of wholesale funding and other managed liabilities.[1] We define *Regulatory Slack* (D_{it}) as the distance between a bank's total risk-based capital ratio and the prevailing regulatory minimum requirement ($r_{it} - \kappa_t$), measured in percentage points. We formalize this relationship using the following descriptive panel specification:

$$D_{it} = \alpha_i + \gamma_t + \beta_1 \; Funding \; Capacity_{it} + \beta_2 \; \log(\text{Assets}_{it}) + \varepsilon_{it} \tag{1}$$

where γ_t are quarter fixed effects and α_i are bank fixed effects, included in the most demanding variant. Because a bank's funding model is a slow-moving structural characteristic, the two sets of fixed effects answer different questions: without α_i, β_1 compares banks with different liability structures; with α_i, it is identified only from within-bank changes in the funding mix.

Table III reports both, together with a pure between-banks specification estimated on bank-level time averages. The funding-buffer trade-off is a sharp *cross-sectional* regularity: in the pooled and between specifications, β_1 is negative and highly significant—a one-standard-deviation increase in the core-deposit share of liabilities (≈ 0.10) is associated with 2.0–2.6 percentage points less regulatory slack, roughly 15% of the median buffer of 15.6 percentage points. Banks whose liabilities are dominated by sticky retail and non-interest-bearing deposits—a stable funding franchise that provides insurance in itself—operate with structurally thinner voluntary buffers, while banks exposed to wholesale funding and its rollover risk self-

[1]The economic rationale behind this metric follows the deposit-franchise view of bank funding: core retail and non-interest-bearing deposits are sticky and withdrawal-insensitive, so a high funding-capacity ratio indicates a liability structure that is largely insulated from funding-market stress. When subjected to a capital or macroprudential shock, such institutions face little rollover risk and can absorb the adjustment without triggering escalating marginal funding penalties. Conversely, a low ratio indicates reliance on wholesale and other managed liabilities, which are rate-sensitive and must be rolled over in markets whose access can deteriorate precisely when capital is scarce. Facing this funding risk—and a deeply costly recourse to friction-laden equity markets—such institutions optimally self-insure by maintaining larger precautionary capital buffers ex-ante. When those buffers erode, the liability-side rigidity forces the adjustment onto the asset side, leaving their aggregate lending supply highly sensitive to unexpected contractions.

insure by carrying larger cushions. Within banks, the gradient is far smaller ($\beta_1 = -2.3$), precisely because the funding mix barely moves at the quarterly frequency: the buffer is an equilibrium object shaped by the bank's structural funding model, not a margin adjusted quarter to quarter—the interpretation our model formalises. The size gradient, by contrast, is robustly negative in every specification: larger institutions, with diversified funding and equity-market access, allow their regulatory slack to compress toward the frontier.

TABLE III
DISTANCE TO CONSTRAINT: PANEL ESTIMATES

Dependent Variable: *Regulatory Slack (pp)*	**(1)** Pooled	**(2)** Between banks	**(3)** Within
Funding Capacity	-19.44***	-25.97***	-2.27**
	(1.73)	(1.95)	(1.05)
Log Assets	-2.80***	-2.70***	-6.27***
	(0.09)	(0.10)	(0.16)
Quarter fixed effects	Yes	—	Yes
Bank fixed effects	No	—	Yes
Observations	935,234	16,117	935,056
R^2	0.104	0.124	0.693

Notes: Estimates of equation (1). The dependent variable is regulatory slack in percentage points, winsorised at the 1st/99th percentiles. Funding capacity is the core-plus-NIB deposit share of total liabilities. Column (2) is estimated on bank-level time averages (one observation per bank) with robust standard errors; columns (1) and (3) cluster standard errors by bank. Sample: quarterly bank panel, 1993Q2–2019Q4. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

This systematic variation in buffer hoarding constitutes the critical pre-condition driving the asymmetric aggregate transmission of macroprudential policy. Consequently, the aggregate credit elasticity of a policy innovation is determined not merely by its nominal magnitude, but by whether the innovation is sufficiently large to bind the endogenous constraint of the marginal institution. Highly capitalized institutions operating deep within their respective regulatory frontiers absorb the policy shock via liability-side reconfigurations, leaving credit supply unaltered. Conversely, institutions operating in close proximity to the regulatory frontier-predominantly small and mid-tier lenders bound by localized retail deposit constraints-experience binding asset-side adjustments. Ultimately, it is this cross-sectional heterogeneity in the distance-to-constraint distribution, rather than a uniform balance-sheet exposure, that generates the size-dependent and sector-specific credit supply contractions documented in Section 2.4. In the quantitative framework of Section 4, funding capacity maps inversely into the funding-market friction that generates buffers endogenously: banks facing steeper marginal costs of replacing funds hold more precautionary equity, so the empirical funding-capacity gradient is the cross-sectional fingerprint of

the same mechanism the model formalises.

Stylized Fact 5: Buffer Heterogeneity Cuts Across Size Tiers. Figure 5 displays the distribution of regulatory slack within each of the three statutory size cohorts. Small banks—the most numerous group—span the full range of buffer positions, from near-zero slack to buffers exceeding 20 percentage points; medium and large institutions cluster closer to the regulatory frontier but still exhibit meaningful dispersion. The key implication is that statutory size is a poor proxy for regulatory tightness: many small banks operate with ample buffers and are effectively insulated from moderate capital shocks, while some medium-tier institutions sit close to the constraint. This within-tier heterogeneity motivates our empirical strategy of conditioning on the buffer distribution directly, rather than on asset-size thresholds alone.

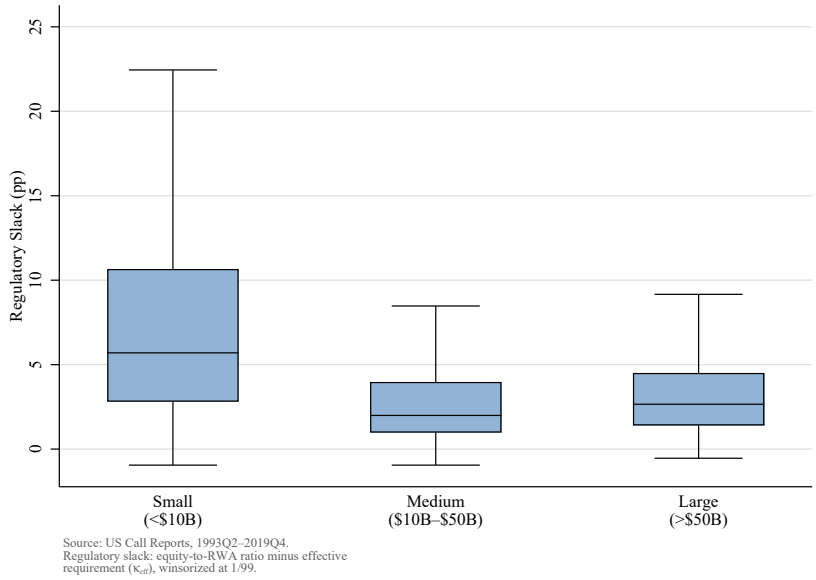


FIGURE 5.—Regulatory slack spans 20+ pp among small banks but compresses near the frontier for larger institutions

Notes: Each box plot displays the interquartile range, median, and 10th/90th percentile whiskers of regulatory slack ($D_{it} = r_{it} - \kappa_t$, in percentage points) within the small (<$10B; ≈880,000 bank-quarters), medium ($10B–$50B; ≈3,200 bank-quarters), and large (>$50B; ≈1,500 bank-quarters) cohorts. The wide dispersion among small banks reflects both economic heterogeneity in funding models and the depth of the community-bank panel; the narrow boxes for medium and large institutions confirm that these cohorts cluster near the regulatory frontier. Sample: quarterly bank panel, 1993Q2–2019Q4. Source: US Call Reports.

2.4 Empirical Strategy

Guided by the cross-sectional regularity documented in Section 2.3-namely, that precautionary capital buffers are systematically larger for institutions facing funding market frictions, and that the distance to the regulatory frontier, rather than bank size per se, dictates the intensity of credit adjustments-we now implement a causal macroeconomic framework. To capture how lending responds to an unexpected tightening of capital requirements across heterogeneous intermediaries, we identify exogenous shifts in the regulatory capital target and trace their dynamic effects on the cross-section of credit supply. U.S. banking law, including the Dodd-Frank Act and subsequent tailoring legislation, applies asset thresholds of \$10B and \$50B to distinguish community, regional, and systemically important institutions; the cross-sectional dispersion in precautionary buffers documented in Table III, however, cuts across these statutory tiers, which motivates conditioning the analysis on the buffer distribution itself. Our subsequent empirical analysis decomposes these credit supply paths across residential real estate, commercial business, consumption, and credit card lending segments.

We employ the Local Projection (LP) method (Jordà, 2005). LPs are particularly well-suited for our analysis as they accommodate exogenously identified shocks and provide a flexible framework for examining heterogeneous responses across different credit types and bank sizes without imposing the rigid dynamic structure of a standard VAR (Plagborg-Møller & Wolf, 2021).

Our benchmark system is defined by the vector $x_t = [l_t^{RE}, l_t^{CI}, e_t, g_t, i_t]'$, where l_t^j represents log lending in sector j (both sectors enter every equation, so each response is conditioned on the other sector's history), e_t is log equity, g_t is real GDP growth, and i_t is the federal funds rate. We estimate the impulse responses at horizon h using the following specification:

$$x_{t+h} - x_{t-1} = \alpha_h + \beta_h \text{shock}_t^{\text{Cap}} + \sum_{p=1}^{6} \Psi_{h,p} x_{t-p} + \varepsilon_{t+h} \tag{2}$$

where $\text{shock}_t^{\text{Cap}}$ is the unanticipated capital requirement surprise identified via the narrative and sign-restriction approach of Drechsel & Miura (2025). The shock series is available from 2002Q1, so the effective estimation sample is 2002Q1–2019Q4; lending and control histories extend back to 1994. Following the AIC criterion, we include six lags of the control variables to account for macroeconomic momentum; the results are robust to four-lag and eight-lag alternatives (Online Appendix). Aggregate specifications are estimated by OLS with Newey–West standard errors (lag $h + 6$); bank-level specifications include

bank fixed effects, winsorise log lending at the 1st/99th percentiles, and cluster standard errors by bank. Because the identified shock is an aggregate series, the bank-level regressions exploit cross-sectional heterogeneity in exposure, with the macroeconomic controls absorbing the common cycle.

Three checks support the identification. First, a pre-trend placebo re-estimating equation (2) at horizons $h \in \{-4, \ldots, -1\}$ yields coefficients indistinguishable from zero ($|t| < 1.5$ throughout): lending does not anticipate the identified innovations. Second, the shock series itself is serially uncorrelated at the quarterly frequency (first-order autocorrelation -0.47, mean ≈ 0), consistent with its construction as a sequence of surprises; the near-unit-root persistence estimated in Section 5 is a property of the requirement *level* that the surprises cumulate into, not of the surprises themselves. Third, consumption and credit-card lending—asset classes carrying full risk weights but funded against high interest margins—show no systematic response at any horizon, indicating that the shock transmits through the regulatory valuation of the balance sheet rather than through aggregate credit demand.

2.5 Results: Credit Type and Bank Size Heterogeneity

Figure 6 presents the local projection impulse response functions (IRFs) of four distinct credit categories following a positive, one-standard-deviation shock to macroprudential capital requirements. A clear risk-weighting hierarchy of structural responses emerges across the various credit instruments. Business lending (C&I) exhibits the most pronounced and persistent contraction, confirming that regulatory tightening induces banks to actively de-risk their balance sheets by compressing exposures to assets with higher regulatory risk weights.

Conversely, real estate lending displays notable structural resilience, remaining relatively stable over the forecast horizon. This trajectory suggests a classic portfolio "flight to quality," wherein institutions preferentially reallocate or maintain capital allocations in asset classes characterized by historically lower regulatory capital charges and stable collateral foundations.

Meanwhile, consumption and credit card lending channels prove highly inelastic, showing statistically negligible volume responses. This finding is particularly informative because consumer loans carry the same 100% risk weight as C&I credit under Basel III, so the regulatory *gross* cost per dollar of each asset class is identical. The fact that banks nonetheless contract C&I but not consumer credit implies that the *net* regulatory cost—gross capital cost minus the interest-margin buffer—is the economically relevant object: the elevated spreads on consumer credit lines provide sufficient cash-flow buffers for banks to

internalise the higher capital charge without volume rationing, whereas the thinner margins on business lending leave the balance-sheet channel operative.

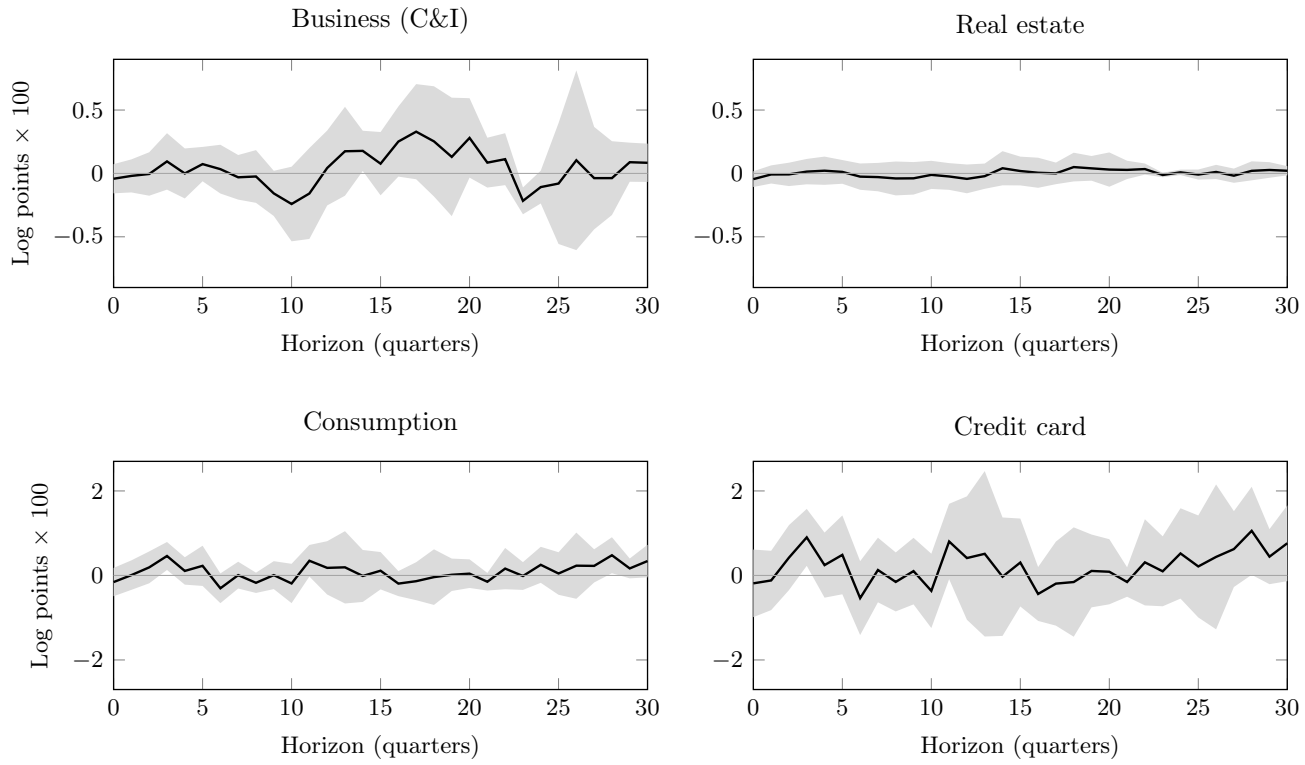


FIGURE 6.—Local projection IRFs of credit supply to a capital requirement shock

Notes: Aggregate local-projection impulse responses of the four credit aggregates to the identified capital-requirement innovation (equation 2); shaded areas are 95% Newey–West confidence bands. Panels within each row share a common vertical scale; the bottom-row (consumer credit) responses are statistically indistinguishable from zero at all horizons. Sample 2002Q1–2019Q4; lag order $p = 6$. Aggregates are unweighted sums over the bank panel. Source: US Call Reports.

Furthermore, our cross-sectional analysis reveals stark heterogeneity in the transmission of capital shocks across the banking sector, organised by the *distance to the regulatory frontier* rather than by asset size per se. Sorting banks each quarter into within-quarter terciles of their lagged regulatory slack $D_{i,t-1}$ and re-estimating the within-bank local projections tercile by tercile (Figure 7), institutions operating with thin voluntary buffers contract lending significantly more following a tightening: the average C&I response over the first twenty quarters is -0.17% per unit innovation in the bottom tercile against -0.13% in the top tercile—a low-to-high ratio of 1.29—with troughs of -0.53% versus -0.41% at $h \approx 15$, and C&I contracting more than real estate *within every tercile*—the sectoral asymmetry that the stylised model of Section 3 attributes to risk-weight-driven asset substitution. This buffer gradient is precisely what the structural framework predicts, and quantitatively so: the estimated model's impulse responses

by initial buffer tercile imply a low-to-high C&I response ratio of 1.23—closely matching the empirical 1.29—and the model reproduces the sign and sectoral ordering of the gradient (Section 3). The gradient is robust to alternative definitions of the regulatory buffer: measuring slack relative to the leverage ratio instead of the risk-based ratio, using fixed capital-ratio cutoffs rather than within-quarter terciles, and varying the bandwidth of the "near-frontier" zone all yield low-to-high response ratios in the range 1.2–1.4 (Online Appendix). Banks near the frontier adjust on the asset side, while banks with ample voluntary slack absorb more of the shock through internal capital reconfiguration—a state-dependence we formalise below as the "Insulated Zone."

Two robustness exercises sharpen the interpretation. First, the gradient reflects the *current* regulatory state rather than a permanent bank type: when tercile membership is frozen at each bank's pre-sample (1994–2001) average slack, the gradient disappears—as expected if what matters is the buffer a bank holds when the shock arrives, the state variable of the model, rather than a fixed characteristic. Second, a continuous interaction of the shock with lagged slack, conditioning on size-decile-specific shock exposure, is not robust; the slack distribution has an extreme right tail of very small institutions with near-unit capital ratios, and the tercile design is insulated from this tail while a linear interaction is not. We accordingly read the cross-section through the tercile split. Because the identified innovation is an aggregate surprise with no pre-trends (Section 2.4), sorting on the lagged buffer one quarter ahead of an unanticipated aggregate shock leaves little scope for anticipatory selection. The bank fixed effects in the within-bank specification absorb any permanent component of risk appetite; the remaining concern is that *time-varying* risk appetite could correlate with buffer position. The frozen-slack test addresses this directly: because the gradient vanishes when tercile membership is fixed at pre-sample averages, the cross-sectional signal cannot be attributed to a slow-moving bank type—including persistent differences in portfolio riskiness—but must instead reflect the time-varying buffer state, as the model's mechanism requires. Finally, the headline responses are not an artifact of ending the estimation sample in 2019: re-estimating the small-cohort local projections on the full extent of the identified shock series (2002Q1–2023Q2), either controlling for a 2020–21 pandemic dummy interacted with the shock or excluding pandemic-dated shocks outright, leaves the target moments essentially unchanged—mean C&I responses of -0.142% per unit shock over $h = 1$–20 under both treatments, against -0.145% in the baseline window, and -0.106% versus -0.105% for real estate.

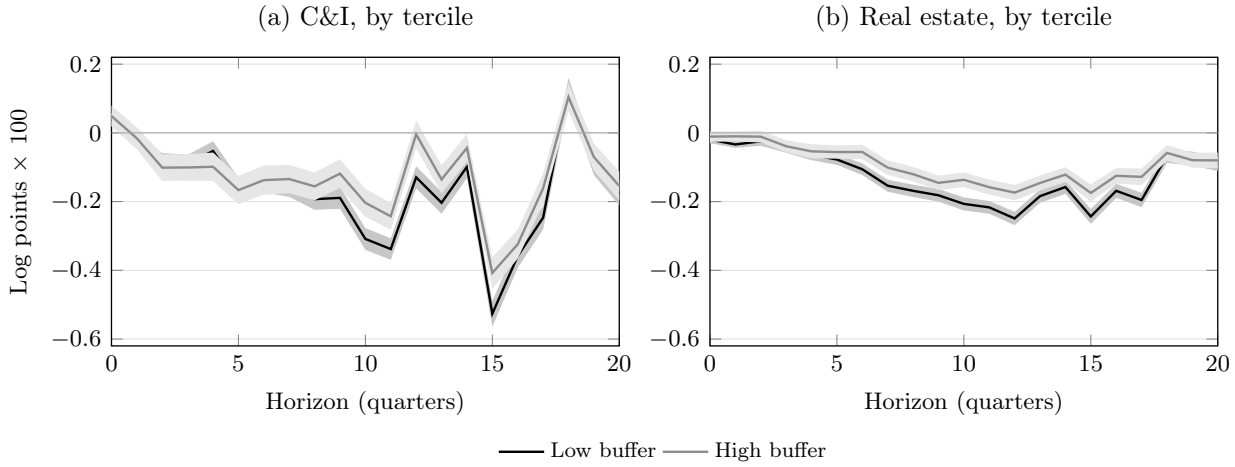


FIGURE 7.— Credit responses by capital-buffer tercile: data vs. model

Notes: Within-bank local projections (bank fixed effects, standard errors clustered by bank; 90% bands shaded) estimated separately for banks in the bottom (darker) and top (lighter) within-quarter terciles of lagged regulatory slack $D_{i,t-1}$; sample 2002Q1–2019Q4, lag $p = 6$. Low-buffer banks contract more in both sectors, with the gap larger for C&I (mean low-to-high response ratio 1.29×) than for RE (1.37×). The estimated model reproduces the C&I gradient (1.23×, untargeted). Responses in log points (×100) per unit of the identified capital-requirement innovation.

We emphasise that statistical power is concentrated among the several thousand community and regional banks that constitute the vast majority of the panel; the small club of institutions above the $50B threshold—roughly forty in the final sample year—is too small to sign its cohort-level responses with confidence (Figure 8). This is a structural feature of U.S. banking data, not a specification choice: the cross-section of SIFIs is an order of magnitude smaller than the community and regional tiers, and no feasible identification strategy based on aggregate capital-requirement shocks can overcome this limitation within the current regulatory regime. The empirical identification therefore leverages the rich cross-sectional variation in capital buffers across the mid-to-large bank distribution where statistical power is sufficient. Crucially, the buffer gradient—the paper's main cross-sectional result (1.29× low-to-high response ratio)—is estimated across the *full* panel, not within any single size cohort, and the tercile sort pools banks of all sizes by their regulatory slack. The gradient's strength comes precisely from the thousands of community and regional institutions that span the entire range of buffer positions (Figure 5), while SIFI estimates, though directionally consistent, contribute little to the precision of the aggregate result. These findings imply that the aggregate transmission of capital regulation is conditional on the cross-sectional distribution of voluntary capital buffers within the banking sector. Notably, this state-dependence does not extend to consumer loan portfolios; consumption and credit card allocations remain uniformly unresponsive to regulatory innovations throughout the cross-section.

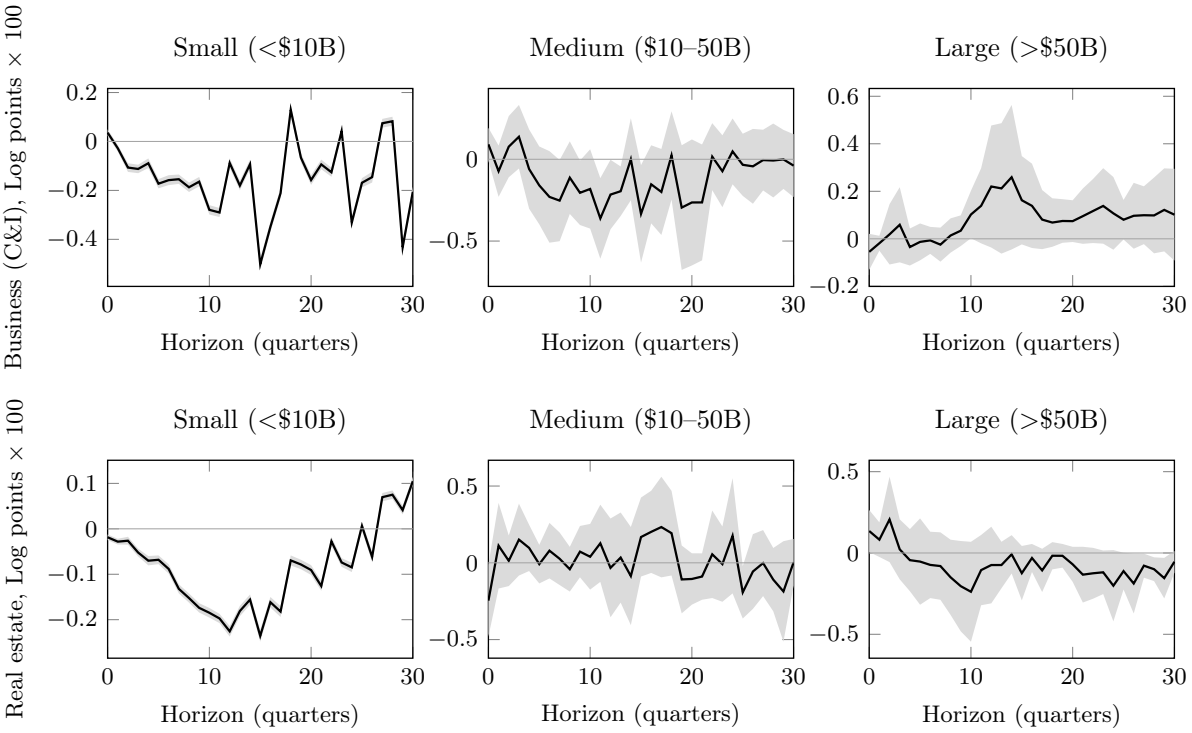


FIGURE 8.— Within-bank LP impulse responses by bank size cohort

Notes: Within-bank local projections (bank fixed effects, standard errors clustered by bank) of C&I and real estate lending by size cohort (split at \$10B/\$50B of lagged assets). Shaded areas are 90% confidence bands. Lag $p = 6$; sample 2002Q1–2019Q4. Note the scale: the small-bank panels are precisely estimated; the large-bank cohort (≈40 institutions) is not. Source: US Call Reports.

To reconcile why well-buffered institutions exhibit attenuated credit supply paths while thinly buffered institutions contract significantly, we introduce a conceptual framework governed by what we term the "Insulated Zone." Historical macroprudential innovations identified in our sample period have been small-to-moderate in magnitude. Banks that maintain large, voluntary ex-ante capital buffers above statutory minimums—predominantly deposit-funded institutions that self-insure against funding-market frictions (Table III)—found these realized historical policy shocks insufficient to deplete their surplus equity. Consequently, such lenders remained within this Insulated Zone, absorbing regulatory changes via internal capital configurations without being forced to adjust their asset mix. As shown in the structural exercise below, the severe non-linear credit contraction predicted for currently well-buffered institutions is a latent vulnerability—one that remains dormant under historical policy variations but rapidly manifests under large, unprecedented counterfactual shocks that completely erode their voluntary buffers.

3 A Stylized Model of Capital Requirements and Credit Allocation

The empirical analysis establishes three robust stylized facts: (i) an aggregate contraction in bank credit following an exogenous capital-requirement shock, (ii) a disproportionate, immediate contraction in business (C&I) lending alongside a resilient mortgage supply, and (iii) a pronounced cross-sectional non-linearity organised by the ex-ante distribution of voluntary capital buffers, for which bank size and funding structure are the observable correlates. This section presents a stylized partial equilibrium model to rationalize these joint dynamics. The framework formally isolates the structural interactions between liability-side funding liquidity constraints, endogenous precautionary capital buffers, and asset-specific regulatory risk weights. Throughout, the primitive source of heterogeneity is a bank's access to marginal funding—not its balance-sheet scale per se—consistent with the panel evidence of Section 2.3.

3.1 Environment and Liability Structure

We consider a static model of a banking sector populated by a continuum of risk-neutral banks indexed by $i \in [0, 1]$, following a standard modelling choice in the macro-banking literature (e.g., Corbae & Erasmo, 2021; Begenau, 2020). Banks differ in their initial equity endowments e_i, structural liability capacity, and baseline asset composition. Lenders intermediate funds by extending credit across two distinct, risk-rated asset classes: residential mortgages (A_i^M) and commercial business loans (A_i^B). Total balance sheet assets are defined by $A_i = A_i^M + A_i^B$.

The asset portfolio is financed via a combination of bank equity and liabilities. To map empirical funding heterogeneities, we partition bank liabilities into two distinct categories:

1. **Core Deposits ($d_{i,C}$):** Stable, low-cost retail funding. While community institutions rely extensively on core deposits as their primary liability franchise, their aggregate supply is locally inelastic and capped by an absolute capacity constraint determined by the bank's geographic footprint:

$$d_{i,C} \leq \bar{d}_{i,C} \tag{3}$$

2. **Wholesale Funding ($d_{i,W}$):** Marginal, market-priced wholesale liabilities utilized to close the balance sheet identity. For any given asset allocation, wholesale funding is determined residually

via the accounting constraint:

$$d_{i,W} = A_i^M + A_i^B - e_i - d_{i,C} \tag{4}$$

Accessing marginal wholesale markets incurs a quadratic friction cost given by $\frac{\gamma_i}{2} d_{i,W}^2$, where the parameter γ_i dictates bank-specific structural access to capital markets. We assume an asymmetric friction distribution such that $\gamma_{\text{small}} \gg \gamma_{\text{large}} \to 0$. This parameterization captures the empirical reality that large, diversified banking organizations possess highly elastic access to wholesale funding lines, whereas smaller community banks face high entry barriers, informational asymmetries, and steep placement penalties in non-deposit funding markets.

3.2 The Regulatory Capital Frontier and Endogenous Buffers

Banks are subject to a binding risk-based capital regulation. The regulatory authority mandates that total equity must represent at least a fraction κ of total Risk-Weighted Assets (RWA_i). Risk-weighted exposures are defined as $RWA_i = \omega_M A_i^M + \omega_B A_i^B$, where commercial business loans carry a structurally higher regulatory risk weight than residential mortgages ($\omega_B > \omega_M$).

We define a bank's absolute distance to its regulatory capital frontier as the capital buffer D_i:

$$D_i = e_i - \kappa \left(\omega_M A_i^M + \omega_B A_i^B \right) \tag{5}$$

Banks hold $D_i \geq 0$ as an endogenous insurance cushion against unexpected structural shocks. If a macroprudential policy innovation tightens the capital regime ($\Delta\kappa > 0$) such that the constraint binds ($D_i < 0$), the institution incurs a severe regulatory non-compliance penalty.

Because low-friction institutions ($\gamma_i \to 0$, empirically the large, diversified organizations) possess near-perfect marginal funding liquidity, they can adjust equity or liabilities ex-post to manage capital shortfalls, and consequently optimize thin ex-ante cushions ($D_i \approx 0$). In contrast, high-friction banks—those whose retail deposit ceiling $\bar{d}_{i,C}$ binds and who face steep wholesale costs γ_i—cannot rapidly raise liability capital, and endogenously hoard large precautionary buffers ($D_i > 0$) ex-ante. Note the precise role of the deposit franchise here, which reconciles this mechanism with the panel evidence of Section 2.3: what disciplines the buffer is the *marginal* funding source. A bank whose stable deposit franchise comfortably covers its funding needs faces little rollover risk and holds a thin buffer; a bank that funds—or, were its buffer breached, would have to fund—at the wholesale margin self-insures. Size matters only through its

correlation with γ_i and with the tightness of the deposit ceiling.

This non-linear threshold dynamic bridges our structural theory with the empirical patterns documented in Section 2.5. Figure 9 illustrates this state dependency via the mapping of the *Insulated Zone*, defined on the joint distribution of funding frictions and buffers:

$$\text{Insulated Zone} \equiv \left\{ (\gamma_i, D_{i,t}) \;\middle|\; D_{i,t} > \bar{D}_{\text{critical}}(\gamma_i) \right\}, \tag{6}$$

where the critical threshold $\bar{D}_{\text{critical}}(\gamma_i)$ is increasing in the funding friction: the harder it is to replace funds at the margin, the larger the buffer required for a given policy innovation to leave lending unaffected.

When policy shocks (ϵ_t^κ) are of mild-to-moderate magnitude—matching the historical variations isolated in our Local Projections—high-friction banks (empirically, predominantly small, deposit-dependent institutions) operate deep within this zone precisely because they have chosen high baseline capital distances ($D_{i,t}$). Because these regulatory innovations fail to push $D_{i,t}$ below the critical threshold $\bar{D}_{\text{critical}}(\gamma_i)$, the shadow cost of the regulatory capital constraint remains entirely non-binding ($\lambda_i = 0$), leaving their credit supply smoothly insulated. However, the theory reveals that this empirical resilience is conditional rather than absolute. Under larger, unprecedented counterfactual policy shocks that breach the frontier of the Insulated Zone, the funding rigidities that motivated the buffer in the first place ($\gamma_i \gg 0$) bind acutely. Once the precautionary buffer is fully exhausted, the asset side becomes hypersensitive, triggering the sharp, non-linear drop in credit supply predicted by the model. The attenuated empirical response of well-buffered banks is therefore not evidence of structural immunity, but rather proof that historical innovations have simply never been large enough to push them out of their Insulated Zone.

3.3 Bank Optimization and General First-Order Conditions

The representative bank selects its optimal multi-asset credit allocation to maximize total economic profits Π_i, net of linear asset management costs ψ_j and quadratic wholesale frictions:

$$\max_{A_i^M, A_i^B, d_{i,C}} \quad \sum_{j \in \{M,B\}} r_j A_i^j - \sum_{j \in \{M,B\}} \frac{\psi_j}{2} (A_i^j)^2 - r_D d_{i,C} - r_W d_{i,W} - \frac{\gamma_i}{2} d_{i,W}^2 \tag{7}$$

subject to the accounting identity, the retail deposit capacity constraint, and the regulatory capital constraint:

$$e_i - \kappa \left(\omega_M A_i^M + \omega_B A_i^B \right) \geq 0 \tag{8}$$

Let λ_i denote the Lagrange multiplier associated with the regulatory capital requirement, representing the shadow cost of regulatory capital. Substituting the wholesale identity directly into the objective function, the interior first-order conditions with respect to asset class $j \in \{M, B\}$ yield:

$$r_j - r_W - \psi_j A_i^j - \gamma_i \left(A_i^M + A_i^B - e_i - d_{i,C} \right) = \lambda_i \kappa \omega_j \tag{9}$$

Equation (9) demonstrates the structural channels driving our empirical cross-sectional results. For low-friction banks, an increase in regulatory stringency ($\Delta\kappa > 0$) yields a minor impact on credit because near-frictionless capital market access ($\gamma_i \to 0$) keeps the shadow cost of capital λ_i low. For high-friction banks, however, the macroprudential innovation directly confronts severe funding rigidities ($\gamma_i \gg 0$). Once a shock is sufficiently large to breach the ex-ante precautionary cushion D_i, the inability to tap external financing causes an acute spike in λ_i, forcing a sharp, non-linear contraction in loan volumes.

3.4 The sectoral lending channel and Risk-Weight Multipliers

When a macroprudential tightening shock binds ($\lambda_i > 0$), the cross-sectional portfolio reallocation path is governed by the structural risk-weight wedge ($\omega_B > \omega_M$), operating through two distinct margins that depend crucially on the institution's baseline diversification and asset size:

1. **The Asset Substitution Margin:** This margin is driven primarily by large and medium banking institutions, which possess the requisite balance sheet diversification (holding a substantial 60% non-mortgage asset baseline) to actively reallocate credit. For such a diversified banking institution, a marginal increase in the regulatory shadow tax λ_i acts asymmetrically across asset tiers. Differentiating the optimal asset-mix ratio with respect to the capital constraint demonstrates that:

$$\left| \frac{\partial A_i^B}{\partial \kappa} \right| > \left| \frac{\partial A_i^M}{\partial \kappa} \right| > 0 \tag{10}$$

   Because commercial business lending carries an elevated risk weight ($\omega_B > \omega_M$), it imposes a steeper marginal capital cost per dollar of credit extended. To satisfy the regulatory frontier while minimizing volume liquidations, diversified large and medium institutions preferentially shed high-weight C&I exposures while preserving real estate credit paths.

   Crucially, because small commercial banks are highly specialized from the outset-with real estate portfolios routinely exceeding 90% of total lending-they lack the structural capacity to utilize this

Asset Substitution Margin. Consequently, small banks cannot easily substitute across sectors; their response to macroprudential innovations is instead governed by whether their larger ex-ante capital buffers allow them to remain entirely within their voluntary *Insulated Zone*.

2. **The Scale Deleveraging Multiplier:** If an institution is forced to pass beyond its buffer and must compress its balance sheet to defend its capital ratio, the regulatory relief generated per dollar of loan liquidation scales directly with the asset's risk weight:

$$\frac{\partial RWA_i}{\partial A_i^B} = \omega_B > \omega_M = \frac{\partial RWA_i}{\partial A_i^M} \tag{11}$$

While low-risk-weight real estate assets require massive nominal balance sheet liquidations to yield marginal regulatory relief, high-risk-weight business portfolios (ω_B) represent the immediate, highly efficient margin of balance sheet optimization. This mechanism rationalizes our empirical Local Projection results, where commercial business portfolios at diversified larger cohorts bear the sharpest and most rapid supply adjustments, while specialized small institutions remain insulated from such sharp non-linear scale contractions under moderate historical shocks.

3.5 Cross-Sectional Heterogeneity and Non-linear Aggregation

The aggregate response of the macroeconomy is ultimately determined by the cross-sectional density of banks for whom the capital frontier actively binds. Figure 9 provides the core graphical intuition of the framework.

Large institutions run with thin buffers but leverage highly elastic funding channels to reallocate asset compositions fluidly. Conversely, small community banks remain completely insulated from minor policy changes due to their substantial ex-ante capital buffers (the "Insulated Zone" in Panel B). However, once a macroprudential innovation is large enough to exhaust these cushions, their localized liability rigidities force an acute, non-linear drop in credit supply concentrated heavily in high-risk-weight business segments.

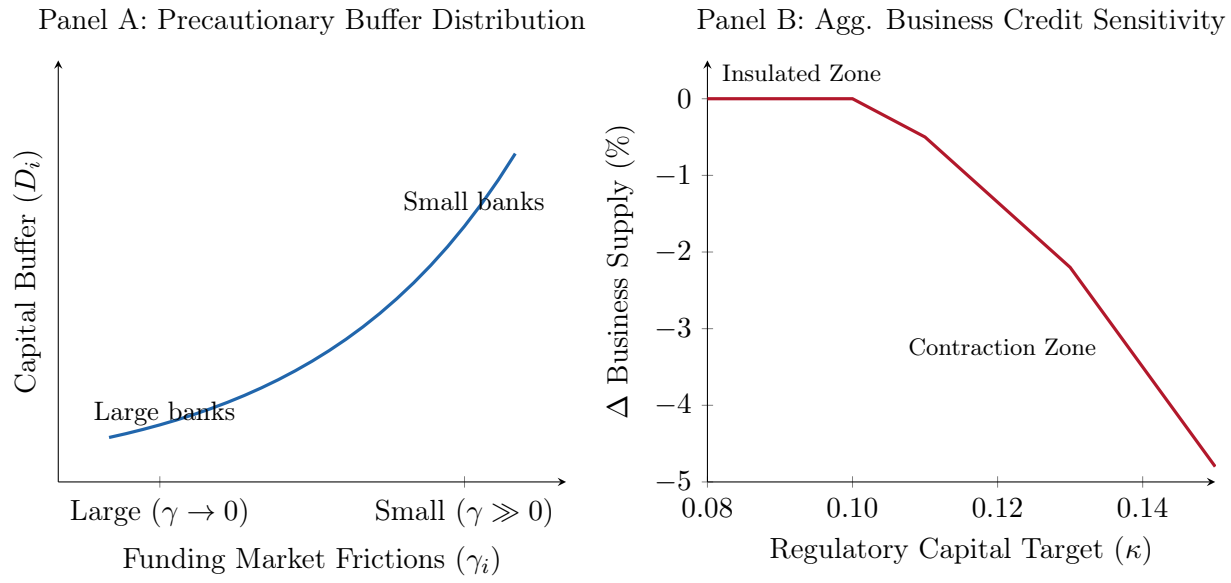


FIGURE 9.—Funding capacity, buffer selection, and aggregate credit sensitivity

Notes: Panel A illustrates the endogenous choice of ex-ante capital buffers as a function of capital market funding frictions. Panel B highlights the non-linear macroeconomic response as capital tightening breaches cross-sectional buffers.

4 Quantitative Model: A Heterogeneous Macro-Banking Framework

While the stylized framework in Section 3 provides the core qualitative mechanics of the sectoral lending channel, this section develops a fully specified quantitative partial-equilibrium framework to assess the macroeconomic magnitude of these structural transformations. We transition from a static cross-sectional environment to an infinitely lived, calibrated heterogeneous-agent model where the joint distribution of balance sheet scale and regulatory slack is disciplined by high-frequency U.S. banking data.

Scope and interpretation of the partial-equilibrium framework. The model abstracts from three channels that a general-equilibrium extension would incorporate: (i) borrower substitution toward non-bank intermediaries (shadow banks, bond markets), which would attenuate the real-side impact of bank credit contraction; (ii) equilibrium loan-rate adjustment, whose direction is ambiguous because both demand and supply shift; and (iii) aggregate-demand feedback through investment, which could amplify the contraction. By shutting down these channels, the framework isolates the *bank supply mechanism*—the joint determination of sectoral credit allocation and the cross-sectional buffer gradient—that is the paper's central object of study. The estimated surplus losses of approximately 2.4–2.9% per percentage-point tightening should accordingly be read as a focused aggregate benchmark for bank-channel

credit contraction, not as a full-economy welfare statement. In particular, the extent to which non-bank intermediaries absorb displaced lending remains an open empirical wedge that could substantially dampen or, if accompanied by regulatory arbitrage and fragility in the shadow sector, redirect the costs estimated here; we leave this extension to future research. Importantly, the magnitudes the model produces (-3.5% C&I, -3.0% RE, -3.1% total per $+1$pp) sit squarely within the range reported by quasi-experimental studies that implicitly capture some of these equilibrium responses (Section 5.4), suggesting that the partial-equilibrium approximation does not produce implausible estimates of the bank-channel effect.

4.1 Mapping Theory to Data: Measuring the Regulatory Buffer

To calibrate the initial state space of the banking sector in our quantitative exercise, we map the theoretical *Distance to Constraint* ($D_{i,t}$) to bank-level regulatory filings using quarterly Consolidated Reports of Condition and Income (Call Reports). For each individual bank i at quarter t, the empirical distance to the regulatory capital frontier is calculated as:

$$D_{i,t} = \left(\frac{\text{Total Capital}_{i,t}}{\text{RWA}_{i,t}} \right) - \kappa_{i,t} \tag{12}$$

where *Total Capital* (`RCFD7205`) and Risk-Weighted Assets (*RWA*, `RCFD7206`) are the exact empirical counterparts to the bank equity state $e_{i,t}$ and the risk-weighted denominator within our structural framework. To map our macro-finance model directly to our empirical cohort designs, the regulatory capital minimum $\kappa_{i,t}$ is formalized as a size-dependent step function dictated by the bank's lagged total asset scale ($A_{i,t-1}$). This setup directly captures the discrete regulatory thresholds established by the Dodd-Frank Act and modern Systemically Important Financial Institution (SIFI) capital surcharges:

$$\kappa_{i,t} = \begin{cases} \kappa_t & \text{if } A_{i,t-1} < \$10\text{B} \\ \kappa_t + \bar{\kappa}_{\text{DFA}} & \text{if } \$10\text{B} \leq A_{i,t-1} \leq \$50\text{B} \\ \kappa_t + \bar{\kappa}_{\text{SIFI}} & \text{if } A_{i,t-1} > \$50\text{B} \end{cases} \tag{13}$$

where κ_t represents the baseline macroprudential requirement, set to a benchmark of 10% in accordance with Prompt Corrective Action (PCA) guidelines for well-capitalized institutions.

Crucially, to fully pin down the regulatory landscape and ensure structural identification, the institutional increments are calibrated using empirical statutory requirements from the post-crisis supervisory

frameworks. We set the Dodd-Frank mid-tier compliance increment to $\bar{\kappa}_{\text{DFA}} = 0.015$ (1.5%) to reflect the enhanced prudential standards and stress-testing capital conservation buffers imposed on regional banking organizations. For larger, systemically important institutions, the surcharge increment is pinned down at $\bar{\kappa}_{\text{SIFI}} = 0.025$ (2.5%), which structurally embeds the average baseline G-SIB and domestic SIFI capital surcharges documented in historical regulatory schedules.

This empirical step function serves two purposes. First, it defines the measurement of the regulatory buffer $D_{i,t}$ used throughout the empirical analysis, so that the distance to the frontier is computed against each bank's *own* statutory requirement. Second, it pins down the tiered schedule evaluated in the policy counterfactuals of Section 5, where the size-differentiated regime is compared with a uniform requirement raising the same aggregate capital. The dynamic model itself is solved under the baseline requirement κ_t; solving the full step function jointly with a realistic size distribution is computationally demanding and left to future work (Section 7).

4.2 Relationship to the Literature

Our structural formulation of the regulatory buffer $D_{i,t}$ aligns with a well-established macro-banking literature that positions capital slack as a primary determinant of aggregate credit supply. Specifically, Kisin & Manela (2016) and Kovner & Van Tassel (2022) identify the shadow cost of capital regulations by observing behavioral distortions in bank asset structures near regulatory minimums. More recently, Couaillier *et al.* (2025) document a significant "Do Not Cross" effect, demonstrating that banks maintaining low distances to regulatory buffers aggressively restrict credit to avoid the market stigma, capital conservation restrictions, and supervisory penalties associated with a technical breach. Our quantitative approach advances this literature by demonstrating that for specialized or liquidity-constrained lenders, this deleveraging behavior is structurally amplified by the risk-weighting wedge ($\omega_B > \omega_M$) and concentrates heavily within high-risk-weight asset classes like commercial and industrial (C&I) business credit.

To evaluate the aggregate transmission dynamics, systemic welfare profiles, and transitional paths of the sectoral lending channel, we embed our stylized mechanics into a dynamic heterogeneous-agent macro-banking framework populated by a continuum of infinitely lived, heterogeneous banking firms. The economy is subject to aggregate macroprudential policy innovations—the only aggregate disturbance in the model—alongside idiosyncratic asset-valuation and credit-depreciation shocks. Restricting the aggregate state to a single policy shock is a deliberate parsimony choice: our identification strategy

isolates the capital-requirement channel, and introducing additional aggregate disturbances (e.g., TFP or funding-cost shocks) would multiply the aggregate state vector without improving identification of the mechanism we estimate. The idiosyncratic shocks already generate the requisite cross-sectional dispersion in buffers; what the single aggregate shock buys is a clean mapping from the estimated policy process $(\rho_\kappa, \sigma_\kappa)$ to the within-bank impulse responses that discipline the model. Market clearing in each credit sector is imposed via the linear inverse demand schedules introduced in Section 4.6.

The primary computational challenge stems from the fact that the aggregate state vector of the economy includes the continuous, time-varying joint distribution of bank equity, total liabilities, and sector-specific asset concentrations, denoted by $\mu_t(e_i, A_i^M, A_i^B)$. Because the transmission of macroprudential shocks is highly non-linear—depending explicitly on the precise mass of institutions operating near discrete regulatory boundaries—standard aggregate-moment approximations following the classical Krusell & Smith (1998) paradigm are structurally insufficient. To overcome the curse of dimensionality while preserving cross-sectional granularity, we deploy a global non-linear solution method utilizing deep neural networks to parameterize both individual policy functions and the endogenous market-clearing asset pricing operators, building upon recent advances in deep macro-finance Maliar *et al.* (2021).

4.3 The Dynamic Bank Optimization Problem

Banks are indexed by $i \in [0, 1]$. At the initialization of quarter t, bank i enters the period with an endogenous equity state $e_{i,t}$, a portfolio of outstanding residential mortgages $A_{i,t-1}^M$, and commercial business lines $A_{i,t-1}^B$. Lenders face asset-specific idiosyncratic valuation shocks $\varepsilon_{i,t}^M$ and $\varepsilon_{i,t}^B$ that capture localized defaults, prepayment variations, and capital depreciation margins.

The specialized banking firm maximizes the expected discounted flow of net dividend distributions to its risk-neutral equity holders:

$$V(e_{i,t}, A_{i,t-1}^M, A_{i,t-1}^B; \Omega_t) = \max_{\{A_{i,t}^M, A_{i,t}^B, d_{i,t}^W, c_{i,t}\}} \mathbb{E}_t \sum_{\tau=0}^{\infty} \beta^\tau \Lambda_{t,\tau} c_{i,t+\tau} \tag{14}$$

where $\beta \in (0, 1)$ is the subjective discount factor and $c_{i,t} \geq 0$ denotes current dividend payouts. Because the equity holders who own the financial intermediaries are risk neutral, the stochastic discount factor satisfies $\Lambda_{t,\tau} = 1$; we carry the notation for generality. The aggregate macroeconomic state vector is defined as $\Omega_t \equiv (\mu_t, \kappa_t)$, where κ_t represents the baseline macroprudential regulatory capital target.

The optimization space is structurally bounded by the bank's dynamic balance sheet identity and capital accumulation law. We explicitly distinguish between stable, retail core deposits ($d_{i,t}^C$) capped at a localized geographic franchise limit ($\bar{d}_{i,C}$), and marginal wholesale funding ($d_{i,t}^W$) which incurs a size-dependent quadratic friction (γ_i):

$$A_{i,t}^M + A_{i,t}^B = e_{i,t} + d_{i,t}^C + d_{i,t}^W \tag{15}$$

Assuming banks exhaust cheaper core deposit capacity ($d_{i,t}^C = \bar{d}_{i,C}$), the dynamic equity law of motion evaluates:

$$\begin{aligned} e_{i,t+1} = (1-\delta_e)e_{i,t} + \sum_{j\in\{M,B\}} \left[(r_{j,t} - \varepsilon_{i,t+1}^j) A_{i,t}^j - \frac{\psi_j}{2}(A_{i,t}^j)^2 \right] \\ - r_{d,t}\bar{d}_{i,C} - r_{W,t} d_{i,t}^W - \frac{\gamma_i}{2}(d_{i,t}^W)^2 - c_{i,t} - \tau_t e_{i,t} \end{aligned} \tag{16}$$

where ψ_j is the quadratic portfolio adjustment cost penalising rapid reshuffling of sector-j assets, δ_e is the quarterly equity depreciation rate capturing operational costs and franchise decay, $r_{d,t}$ is the retail deposit rate, $r_{W,t}$ is the wholesale funding rate, and $\tau_t e_{i,t}$ represents a time-varying systemic insurance levy scaled to the bank's equity footprint to fund entrant capital injections. The optimization space is strictly bounded by the size-dependent macroprudential regulatory capital boundary:

$$\omega_M A_{i,t}^M + \omega_B A_{i,t}^B \leq \frac{e_{i,t}}{\kappa_{i,t}(\text{Scale}_{i,t})} \tag{17}$$

where the regulatory threshold environment $\kappa_{i,t}(\cdot)$ is modeled as a step function dictated by total assets at the start of the period ($\text{Scale}_{i,t} \equiv A_{i,t-1}^M + A_{i,t-1}^B$), ensuring regulatory cohort assignments are fixed during current portfolio rebalancing steps and avoiding sub-derivative discontinuities.

The first-order conditions yield forward-looking stochastic Euler equations that govern the optimal allocation of credit across sectors. Future streams are discounted using $\Lambda_{t,t+1}$ (equal to one under the maintained risk neutrality of bank owners). Let $\mu_{i,t} \geq 0$ represent the non-negative Kuhn-Tucker multiplier on the dividend non-negativity constraint ($c_{i,t} \geq 0$), and let $\lambda_{i,t} \geq 0$ represent the multiplier associated with the regulatory capital frontier. The optimal credit allocation rules explicitly incorporate the risk-weight wedge ($\omega_B > \omega_M$) alongside the marginal cost of external wholesale funding:

$$\beta \mathbb{E}_t \left[\Lambda_{t,t+1}(1+\mu_{i,t+1}) \left(r_{M,t} - \varepsilon_{i,t+1}^M - \psi_M A_{i,t}^M - r_{W,t} - \gamma_i d_{i,t}^W \right) \right] = \lambda_{i,t}\omega_M \tag{18}$$

$$\beta \mathbb{E}_t \left[\Lambda_{t,t+1}(1+\mu_{i,t+1}) \left(r_{B,t} - \varepsilon^B_{i,t+1} - \psi_B A^B_{i,t} - r_{W,t} - \gamma_i d^W_{i,t} \right) \right] = \lambda_{i,t} \omega_B \tag{19}$$

Looking intertemporally, the shadow value of internal equity resolves into an intuitive asset pricing equation:

$$1 + \mu_{i,t} = \frac{\lambda_{i,t}}{\kappa_{i,t}} - \tau_t + \left(1 - \delta_e + r_{W,t} + \gamma_i d^W_{i,t}\right) \cdot \beta \mathbb{E}_t \left[\Lambda_{t,t+1}(1+\mu_{i,t+1}) \right] \tag{20}$$

where both multipliers satisfy standard complementarity slackness conditions:

$$\lambda_{i,t} \geq 0, \quad \frac{e_{i,t}}{\kappa_{i,t}} - (\omega_M A^M_{i,t} + \omega_B A^B_{i,t}) \geq 0, \quad \lambda_{i,t} \left[\frac{e_{i,t}}{\kappa_{i,t}} - (\omega_M A^M_{i,t} + \omega_B A^B_{i,t}) \right] = 0 \tag{21}$$

$$\mu_{i,t} \geq 0, \quad c_{i,t} \geq 0, \quad \mu_{i,t} c_{i,t} = 0 \tag{22}$$

Finally, to maintain a stationary cross-sectional measure μ_t and account for severe idiosyncratic realizations, banks whose equity falls below zero ($e_{i,t} < 0$) are resolved. They exit the market and are immediately replaced by a new entrant endowed with a small positive equity stock. To ensure aggregate accounting integrity within the financial sector, the total equity injected into these new entrants is funded via the insurance levy τ_t deducted pro-rata from the aggregate dividend distributions of surviving solvent financial intermediaries, capturing a systemic internal restructuring pool.

4.4 Deep Learning Solution Algorithm

To solve the dynamic system defined by equations (18)–(22) across an infinite-dimensional state space, we parameterize the individual choice rules and the structural shadow price functions using deep neural networks. Conventional value-function iteration (VFI) is infeasible here because the state space is high-dimensional—individual states $(e_{i,t}, A^M_{i,t-1}, A^B_{i,t-1})$ plus the compressed cross-sectional distribution Ω_t—and the occasionally binding capital and dividend constraints introduce kinks that break the smoothness assumptions of projection methods. The neural-network approach can be viewed as a global weighted-residual method in which the network parameterises the residual function and stochastic gradient descent replaces the fixed collocation grid, inheriting the global accuracy guarantees of weighted-residual methods while scaling to higher-dimensional problems (Maliar *et al.*, 2021). Online Appendix Table OA.II reports robustness to network architecture (depth, width, activation function); training time for the baseline 4×256 architecture is approximately 6 hours on an NVIDIA A100 GPU. Let $S_{i,t} \equiv (e_{i,t}, A^M_{i,t-1}, A^B_{i,t-1}, \Omega_t)$ assemble the comprehensive vector of individual and aggregate states.

We parameterize a multi-layer perceptron (MLP) mapping function $\mathcal{A}(S_{i,t};\Theta)$ governed by optimization weights and biases Θ. This neural network simultaneously approximates the optimal multi-asset policy rules for the current period and the non-linear regulatory shadow cost schedules:

$$\begin{bmatrix} \hat{A}^{M}_{i,t}(S_{i,t};\Theta) \\ \hat{A}^{B}_{i,t}(S_{i,t};\Theta) \\ \hat{c}_{i,t}(S_{i,t};\Theta) \\ \hat{\lambda}_{i,t}(S_{i,t};\Theta) \\ \hat{\mu}_{i,t}(S_{i,t};\Theta) \end{bmatrix} = \mathcal{A}(S_{i,t};\Theta) \tag{23}$$

To maintain computational tractability without loss of critical threshold information, the continuous distribution μ_t is compressed into a multi-dimensional state vector containing its first four statistical moments, augmented explicitly by a localized kernel-density mass parameter evaluating the exact density of institutions clustered within a 1% bandwidth of the \$10B and \$50B asset boundaries. Intertemporal transitions of this compressed vector are mapped via an embedded neural network operator parameterized as a perceived law of motion (PLM), $\Omega_{t+1} = \Gamma_{\Theta}(\Omega_t, Z_{t+1}, \kappa_{t+1})$, tracking law-of-motion updates inside conditional expectations.

The optimal parameters Θ^* are found by minimizing a structural loss that sums the squared Euler-equation residuals for both asset classes and equity, plus a soft-penalty term enforcing the Kuhn-Tucker complementarity conditions (21)–(22). Market clearing in each credit sector is imposed analytically via the closed-form equilibrium rate in equation (27), replacing the numerical inner loop of Maliar *et al.* (2021) with an exact solution at each forward pass. The network is optimized globally via AdamW with cosine annealing. The full loss function, the complementary-slackness penalty derivation, and a detailed implementation summary are reported in the Online Appendix.

4.5 Data and Estimation Strategy

We parameterize and estimate the quantitative model using a unified hybrid two-step strategy to anchor the structural equations within empirical macro-finance matrices.

4.5.1 Calibration of Structural Parameters

Standard macroeconomic constants (e.g., household risk aversion, core asset depreciation rates, and baseline Basel III regulatory weights $\omega_M = 0.50$ and $\omega_B = 1.00$) are calibrated at a quarterly frequency to standard benchmark values from the banking literature. The quadratic adjustment cost scale parameters (ψ_M, ψ_B) are explicitly calibrated to target the long-run historical balance sheet portfolio concentrations documented within our Call Report sample.

4.5.2 Deep Simulated Method of Moments (SMM)

The remaining structural parameters governing the stochastic policy environment—specifically the persistence and volatility parameters of the macroprudential capital policy process $(\rho_\kappa, \sigma_\kappa)$—are structurally estimated via the *Simulated Method of Moments* (SMM). We model κ_t as an AR(1) process, the simplest specification that accommodates the near-unit-root persistence visible in the cumulated shock series while remaining parsimoniously identified from the sixteen LP coefficients. A richer specification—such as a Markov regime-switching process distinguishing tightening episodes from quiescent periods—would be economically appealing, but the two regimes would introduce at least four additional parameters (persistence, volatility, and transition probabilities for each regime). With an SMM objective built on sixteen local-projection coefficients that are themselves noisy, these parameters are not separately identified at conventional significance levels; we therefore maintain the parsimonious AR(1) and note that the near-unit-root estimate ($\hat{\rho}_\kappa \approx 0.99$) already captures the high persistence that a regime-switching specification would attribute to the tightening state.

Following the standard approach in the banking literature, the standard deviation of idiosyncratic credit depreciation shocks, σ_ε, is calibrated externally to isolate the microeconomic baseline of asset quality, thereby avoiding identification multi-collinearity during global optimization. Specifically, σ_ε is fixed to the time-series average of the cross-sectional standard deviation of quarterly net charge-off rates by asset class, computed from the full bank panel and demeaned by the quarterly aggregate to cleanly isolate the idiosyncratic component. This calibration represents the natural empirical counterpart to our structural model's idiosyncratic shock setup following Mendicino *et al.* (2020).

Let $\theta \equiv (\rho_\kappa, \sigma_\kappa)$ collect the two structural parameters slated for estimation. The SMM estimator searches

for the parameter vector that minimizes the distance between the empirical and simulated targets:

$$\hat{\theta} = \arg\min_{\theta} \left[\hat{m}_{\text{data}} - m_{\text{model}}(\theta)\right]' \mathbf{W} \left[\hat{m}_{\text{data}} - m_{\text{model}}(\theta)\right] \tag{24}$$

where $\hat{m}_{\text{data}}$ is a sixteen-dimensional empirical target vector comprising the within-bank local-projection coefficients from Section 2 for both C&I and real estate lending at horizons $\mathcal{H} = \{1, 2, 4, 6, 8, 12, 16, 20\}$ quarters (8 horizons $\times$ 2 sectors = 16 moments). The target therefore spans the full dynamic profile of the sectoral lending channel—both the higher-risk-weight C&I segment and the lower-risk-weight real estate segment—at short, medium, and long horizons.

Heuristically, the two parameters within θ map into complementary features of this sixteen-dimensional target. The policy persistence ρ_κ is primarily identified by the decay rate of the credit contraction across medium-to-long horizons in both sectors; a more persistent shock sustains the regulatory shadow cost $\lambda_{i,t}$ across multiple periods, generating slower impulse response function (IRF) reversion. Conversely, the policy volatility scale σ_κ is pinned down by the short-horizon amplitude of the IRFs, which scales proportionally with the impact-period size of the regulatory innovation. Using both C&I and real estate moments simultaneously disciplines the model's cross-sectoral allocation: the risk-weight wedge $\omega_B > \omega_M$ must generate a differential contraction consistent with the data at every horizon.

The structural weighting matrix $\mathbf{W}$ is the optimal variance-covariance matrix derived via bootstrapping our empirical sample, block-diagonal across the C&I and real estate blocks. For each candidate evaluation vector θ, the deep learning policy network is re-trained to compute the sixteen model-implied conditional moments $m_{\text{model}}(\theta)$, maintaining rigorous consistency between structural parameterization and the global solution paths.

4.6 Loan Demand Schedules and Partial Equilibrium Rates

To close the partial equilibrium model and compute the market-clearing asset returns $(r_{M,t}, r_{B,t})$ that appear in the bank optimisation problem (18)–(22), we introduce reduced-form inverse loan demand schedules for each credit market $j \in \{M, B\}$. The focus of this paper is on the supply-side transmission of capital requirement shocks; the demand block is therefore kept parsimonious, consistent with the partial equilibrium scope of the model.

4.6.1 Inverse Demand Specification

We posit linear inverse demand functions of the form:

$$r_{j,t} \;=\; \theta_j \;-\; \eta_j \, L_t^j, \qquad j \in \{M, B\}, \tag{25}$$

where $L_t^j \equiv \frac{1}{N}\sum_{i=1}^{N} \hat{A}_{i,t}^j$ denotes aggregate lending in sector j per unit of banking mass, $\theta_j > 0$ is the demand intercept (the choke rate when aggregate lending is zero), and $\eta_j > 0$ is the inverse demand elasticity (slope of the inverse demand curve). This specification admits a standard linear demand interpretation: borrowers' willingness to pay for credit falls with aggregate loan supply, so that a contraction in bank lending following a capital shock generates a partial offsetting rise in loan rates.[2]

Market clearing requires that aggregate loan supply equals aggregate loan demand in each sector at every date t:

$$\frac{1}{N}\sum_{i=1}^{N} \hat{A}_{i,t}^j(\mathbf{S}_{i,t};\Theta) \;=\; L_t^j(r_{j,t}), \qquad j \in \{M, B\}, \tag{26}$$

where $\hat{A}_{i,t}^j(\mathbf{S}_{i,t};\Theta)$ are the neural-network policy functions defined in Section 4.4. Substituting (25) into (26) and solving for the equilibrium rate yields:

$$r_{j,t}^{*} \;=\; \theta_j \;-\; \eta_j \cdot \frac{1}{N}\sum_{i=1}^{N} \hat{A}_{i,t}^j(\mathbf{S}_{i,t};\Theta). \tag{27}$$

Equation (27) is evaluated inside the deep learning forward pass at each training iteration: given the current network weights Θ, aggregate supply is computed, the equilibrium rate is updated, and the updated rate feeds back into the individual Euler conditions (18)–(19) in the subsequent iteration.

4.6.2 Calibration of Demand Parameters

The four demand parameters $(\theta_M, \eta_M, \theta_B, \eta_B)$ are calibrated externally rather than estimated, consistent with the partial equilibrium scope of the model. We proceed in two steps.

[2]The linearity assumption is a deliberate simplification. Because identification rests on the supply-side capital-requirement shock ε_t^{κ} rather than demand innovations, the functional form of $L_t^j(\cdot)$ does not affect the causal estimates reported in Section 2. Its role is limited to (i) closing the market-clearing block of the quantitative model and (ii) enabling the welfare calculations in Section 4.7.

Step 1: Steady-state intercepts θ_j. Let $\bar{L}^j$ and $\bar{r}_j$ denote the long-run sample averages of aggregate lending and the implied lending rate in sector j, computed from the Call Reports panel over 1994–2019. The steady-state condition requires $\bar{r}_j = \theta_j - \eta_j \bar{L}^j$, so that once η_j is fixed the intercept is pinned down by:

$$\theta_j \;=\; \bar{r}_j \;+\; \eta_j \, \bar{L}^j, \qquad j \in \{M, B\}. \tag{28}$$

Step 2: Demand elasticities η_j. The elasticities are fixed at values drawn from the external literature and cross-validated against the Federal Reserve Senior Loan Officer Opinion Survey (SLOOS). For the mortgage sector, we set η_M to match the pass-through elasticity reported by Adelino, Schoar and Severino (2012) for residential mortgage demand. For the commercial sector, we set η_B to match the C&I loan demand elasticity estimated in Jiménez, Ongena, Peydró and Saurina (2014). Formally:

$$\eta_j \;=\; -\,\frac{\partial \log L^j}{\partial \log r_j} \cdot \frac{\bar{r}_j}{\bar{L}^j}, \tag{29}$$

where the semi-elasticities on the right-hand side are taken from the respective reduced-form estimates. We verify robustness by re-solving the model under $\eta_j \times \{0.5,\, 1.0,\, 2.0\}$; the model-implied IRF profiles and welfare rankings are insensitive to this variation, confirming that the supply-side results are not driven by the demand parameterisation.

4.6.3 Aggregate Rate Pass-Through

A direct implication of equations (25)–(27) is that a capital-requirement tightening generates a partial rate pass-through effect: the reduction in aggregate loan supply pushes equilibrium rates upward even as individual banks reduce volumes. The aggregate rate response at horizon h satisfies:

$$\frac{\partial r^*_{j,t+h}}{\partial \varepsilon^\kappa_t} \;=\; -\,\eta_j \cdot \frac{\partial L^j_{t+h}}{\partial \varepsilon^\kappa_t}, \tag{30}$$

so that the magnitude of rate pass-through is proportional to the quantity contraction documented in Section 2.5 and scales with the inverse demand slope η_j. Because $|\partial L^B / \partial \varepsilon^\kappa| > |\partial L^M / \partial \varepsilon^\kappa|$ by the risk-weight mechanism, C&I lending rates rise more sharply on impact than mortgage rates, generating an asymmetric cost-of-credit effect across sectors that amplifies the real consequences of macroprudential

Table IV
Calibration and Estimation Summary

Parameter	Description	Value	Source / Target
Panel A: Household and Macro			
β	Quarterly discount factor	0.99	Standard; annualised $\beta^4 \approx 0.96$
σ	Household CRRA coefficient (SDF)	1.00	Log utility; standard in banking literature
Panel B: Bank Technology			
δ_e	Quarterly equity depreciation rate	0.0010	Dynamic consistency: equity-Euler stationarity at $\beta = 0.99$
ψ_M	Mortgage portfolio adjustment cost	0.0073	Calibrated: targets RE loan share ≈ 0.69 (Table I)
ψ_B	C&I portfolio adjustment cost (baseline)	0.05	Preset; admissible range (0.014, 0.142), sensitivity at {0.02, 0.10}
Panel C: Regulatory Capital			
ω_M	Residential mortgage risk weight	0.50	Basel standardised approach; Section 4
ω_B	Commercial & industrial loan risk weight	1.00	Basel standardised approach; Section 4
κ_{base}	Baseline PCA capital minimum	0.10	Prompt Corrective Action well-capitalised threshold
$\bar{\kappa}_{\text{DFA}}$	DFA add-on ($10B–$50B banks)	0.015	Fed enhanced prudential standards; post-2010Q3
$\bar{\kappa}_{\text{SIFI}}$	SIFI/G-SIB add-on (>$50B banks)	0.025	12 CFR 217.403; phased in 2016Q1–2019Q4
Panel D: Funding and Liability			
γ_{large}	Wholesale funding friction, large banks	≈ 0	Paper assumption; near-frictionless capital market access
γ_{small}	Wholesale funding friction, small banks	0.0769	Calibrated: empirical funding structure (Table III)
$\bar{d}_C$	Core-deposit capacity (USD thousands)	447 ,499	Balance-sheet identity at the calibrated steady state
τ_t	Systemic insurance levy rate	0.0002	FDIC DIF assessment (quarterly decimal)
Panel E: Loan Demand (Option B, Section 4.6)			
η_M	Inverse demand slope, mortgages	6.497 e-8	Eq. (29); Adelino *et al.* (2025) semi-elasticity $= -2$
η_B	Inverse demand slope, C&I loans	3.773 e-7	Eq. (29); Jiménez *et al.* (2014) semi-elasticity $= -4$
θ_M	Mortgage demand intercept	0.0462	Eq. (28): $\bar{r}_M + \eta_M \bar{L}_M$; $\bar{r}_M = 0.0154$, $\bar{L}_M = 473{,}578$
θ_B	C&I loan demand intercept	0.0847	Eq. (28): $\bar{r}_B + \eta_B \bar{L}_B$; $\bar{r}_B = 0.0169$, $\bar{L}_B = 179{,}618$
Panel F: Idiosyncratic Credit Shocks (externally calibrated)			
$\bar{\sigma}_{\varepsilon,M}$	SD of idiosyncratic RE charge-off rate	0.0010	Online Appendix; cross-sectional SD of demeaned $\varepsilon^M_{i,t}$, 1994–2019
$\bar{\sigma}_{\varepsilon,B}$	SD of idiosyncratic C&I charge-off rate	0.0041	Online Appendix; cross-sectional SD of demeaned $\varepsilon^B_{i,t}$, 1994–2019
Panel G: Macroprudential Policy Process (SMM-estimated)			
ρ_κ	Persistence of capital requirement innovation	0.9987 (0.0038)	SMM: medium-horizon persistence of the within-bank IRFs
σ_κ	Volatility of capital requirement innovation	0.00040 (0.00002)	SMM: amplitude of the within-bank IRFs
Panel H: Neural Network Architecture (Section 4.4)			
L	Number of hidden layers	4	MLP; Section 4.4
d	Nodes per hidden layer	256	MLP; Section 4.4
activation	Activation function	Swish	Gradient stability near regulatory thresholds
optimizer	Training algorithm	AdamW	Cosine annealing learning rate schedule

Notes: All dollar-denominated variables are in thousands of USD (as reported in the Call Reports) at quarterly frequency. Interest rates are quarterly decimals (annual percentages divided by 400). Risk weights (ω_M, ω_B) and capital requirements (κ) follow the Basel III standardised approach. The regulatory capital minimum $\kappa_{i,t}$ is a size-dependent step function of lagged total assets (Section 4): κ_{base} for assets below $10B, $\kappa_{\text{base}} + \bar{\kappa}_{\text{DFA}}$ for $10B–$50B, and $\kappa_{\text{base}} + \bar{\kappa}_{\text{SIFI}}$ above $50B. Loan demand intercepts (θ_M, θ_B) are the choke rates at which aggregate demand equals zero, derived analytically from steady-state moments and calibrated elasticities via Eq. (28). Parameters marked "—" are calibrated numerically within the model to match the empirical targets stated in the Source column (Section 4.5.1). $\bar{\sigma}_{\varepsilon,j}$ is calibrated externally following Mendicino *et al.* (2020) as the time-series mean of the within-period cross-sectional standard deviation of quarterly net charge-off rates, demeaned to isolate idiosyncratic variation from the common aggregate component. The SMM weighting matrix is $\mathbf{W} = \hat{\Sigma}^{-1}$, block-diagonal across the C&I and RE blocks, each block estimated by a cluster (bank-level) bootstrap of the within-bank LP coefficients (100 replications), for the impulse responses at horizons $\mathcal{H} = \{1, 2, 4, 6, 8, 12, 16, 20\}$; see the Online Appendix for the formal identification argument. Standard errors for the SMM estimates (in parentheses) are asymptotic GMM; ρ_κ is estimated in the interior of the parameter space but within one standard error of a unit root, with 95% profile confidence set $\rho_\kappa \geq 0.994$. Sources: Adelino *et al.* (2025) for the residential mortgage demand semi-elasticity; Jiménez *et al.* (2014) for the C&I loan demand semi-elasticity; Mendicino *et al.* (2020) for the idiosyncratic shock calibration strategy; Drechsel & Miura 2025 for the capital-requirement shock series.

tightening.

4.7 Steady State, Welfare, and Policy Counterfactuals

4.7.1 Stationary Equilibrium

A stationary recursive equilibrium of the model consists of: (i) a time-invariant policy function $\mathbf{A}(\mathbf{S}_i; \Theta^*)$ mapping individual states to credit allocations, dividends, and shadow costs; (ii) a stationary distribution μ^* over individual bank states (e_i, A_i^M, A_i^B); and (iii) constant equilibrium rates (r_M^*, r_B^*) such that the market-clearing conditions (26) hold at μ^*. In the stationary equilibrium the regulatory capital requirement is fixed at its long-run baseline $\kappa^* = \kappa_{\text{base}}$, and the insurance levy τ^* is set to balance the entrant recapitalisation fund. The stationary distribution μ^* is obtained as the ergodic fixed point of the law of motion $\Omega_{t+1} = \Gamma_\Theta(\Omega_t, Z_{t+1}, \kappa_{t+1})$ evaluated at $\kappa_{t+1} = \kappa^*$ and $Z_{t+1} = 0$.

4.7.2 Welfare Measure

We measure the aggregate welfare cost of a capital-requirement shock using a consumption-equivalent variation (CEV). Let V^* denote the aggregate bank franchise value in the stationary equilibrium and V^κ the corresponding value following a permanent one-percentage-point tightening of the capital requirement from κ^* to $\kappa^* + 0.01$. The aggregate CEV Λ satisfies:

$$V^\kappa(\mu^*;\, \kappa^* + 0.01) \;=\; V^*((1 - \Lambda)\, \mu^*;\, \kappa^*)\,, \tag{31}$$

so that $\Lambda > 0$ represents the fraction of steady-state bank dividends that households would be willing to forgo in order to avoid the regulatory tightening. Because dividend flows map directly to household income in this partial equilibrium framework, Λ provides a lower-bound estimate of the aggregate welfare cost of the policy innovation.

We decompose Λ into a *sectoral allocation cost* Λ^{alloc}, arising from the distortion in the C&I-to-mortgage lending ratio induced by the risk-weight wedge ($\omega_B > \omega_M$), and a *scale cost* Λ^{scale}, arising from the contraction in total credit supply:

$$\Lambda \;=\; \underbrace{\Lambda^{\text{alloc}}}_{\text{sectoral distortion}} \;+\; \underbrace{\Lambda^{\text{scale}}}_{\text{aggregate credit contraction}} \;+\; \mathcal{O}\big(\Delta\kappa^2\big)\,. \tag{32}$$

This decomposition maps directly onto the two structural channels identified in Section 3: the Asset Substitution Margin generates Λ^{alloc}, while the Scale Deleveraging Multiplier generates Λ^{scale}.

We implement the CEV under two complementary value concepts. Our baseline measures the credit-market surplus generated by intermediation—the area under each sector's inverse loan demand schedule above the marginal funding cost—which falls unambiguously when credit contracts. As a robustness check we also compute the franchise value directly, as the expected discounted stream of bank dividends along a stabilised forward simulation of the estimated policy, with dividends given by the stationary residual of the equity law of motion. The two metrics agree in sign and magnitude (Section 5), bracketing the welfare object.

Three caveats delimit the exercise. First, the framework is partial equilibrium: dividend and surplus flows map to household income, but there is no feedback to factor prices or aggregate demand. Second, the model assigns no benefit to capital—default and crisis-prevention channels are absent—so Λ is a *gross* cost of tightening, not a net policy evaluation; in particular, the dividend externalities emphasised by Acharya *et al.* (2017)—whereby individually optimal capital distributions erode systemic resilience—are precisely the kind of benefit-side channel our framework is silent on. Third, the CEV is evaluated at the estimated, effectively permanent innovation process; transitory experiments would yield smaller costs.

Several general-equilibrium channels, absent from the framework, could attenuate or amplify the partial-equilibrium supply contractions measured here. On the demand side, borrowers denied bank credit may substitute toward non-bank financial intermediaries—shadow banks, private credit funds, or bond markets—partially offsetting the real-side impact of the bank-lending contraction (see Buchak *et al.*, 2018, for evidence that fintech and shadow-bank lenders absorbed a large share of the post-crisis decline in bank mortgage origination, though substitution is likely smaller for relationship-intensive C&I credit). On the supply side, equilibrium loan rates would adjust: if aggregate loan supply falls, the clearing rate rises, crowding out marginal borrowers but also raising bank spreads, which partly restores lending incentives. Finally, the credit contraction could feed back into aggregate demand through reduced investment and employment, further depressing loan demand in a contractionary spiral. The net direction is ambiguous and depends on the elasticity of borrower substitution and the strength of the aggregate-demand multiplier. By holding these channels fixed, our partial-equilibrium estimates isolate the pure supply mechanism and provide a transparent upper bound on the banking-sector adjustment cost.

4.7.3 Policy Counterfactuals

We evaluate two counterfactual policy experiments that exploit the heterogeneous transmission documented in Sections 3 and 4.

Counterfactual 1: Uniform vs. size-differentiated capital requirements. The baseline step function $\kappa_{i,t}(\text{Scale}_{i,t})$ in equation (17) assigns higher capital requirements to larger institutions. We compare this to a counterfactual uniform requirement $\tilde{\kappa}$ that raises the same aggregate regulatory capital in the stationary equilibrium:

$$\tilde{\kappa} \;\text{ s.t. }\; \int \tilde{\kappa}\, e_i \, \mathrm{d}\mu^* \;=\; \int \kappa_i(\text{Scale}_i)\, e_i \, \mathrm{d}\mu^*. \tag{33}$$

A comparison of the welfare cost Λ under the baseline size-differentiated rule and the uniform rule $\tilde{\kappa}$ identifies the efficiency gain (or loss) from regulatory tailoring. Given the Insulated Zone mechanism documented in Section 3.5, we expect the uniform rule to generate a larger aggregate C&I credit contraction relative to the size-differentiated baseline, because it removes the precautionary buffer advantage of small institutions. Because the dynamic model is solved under the baseline requirement (Section 4), this counterfactual compares stationary equilibria under the two schedules rather than re-solving the full step function along the transition; the approximation affects the cross-sectional *incidence* of the comparison, not the aggregate scale result, which is pinned down by the level of aggregate capital that equation (33) holds fixed by construction.

Counterfactual 2: Risk-weight flattening. We evaluate the effect of reducing the risk-weight wedge by setting $\tilde{\omega}_B = \tilde{\omega}_M = 0.75$ (the arithmetic average of the baseline weights), holding κ fixed at κ^*. This experiment isolates the contribution of the sectoral distortion channel Λ^{alloc} to total welfare costs by shutting down the Asset Substitution Margin while preserving the scale effect. The counterfactual aggregate credit contraction under flattened weights satisfies:

$$\frac{\partial \tilde{L}^B_{t+h}}{\partial \varepsilon^\kappa_t} \;\approx\; \frac{\partial \tilde{L}^M_{t+h}}{\partial \varepsilon^\kappa_t} \quad \forall\, h, \tag{34}$$

so that sectoral IRF asymmetry disappears. The difference between the baseline and flattened-weight welfare costs provides a structural estimate of the dollar value of the sectoral distortion created by

differential risk-weighting under Basel III.

5 Quantitative Results

5.1 Estimation: matching the within-bank impulse responses

We estimate the policy process $\theta = (\rho_\kappa, \sigma_\kappa)$ by simulated method of moments on the within-bank local-projection coefficients of Section 2 at horizons $\mathcal{H} = \{1, 2, 4, 6, 8, 12, 16, 20\}$ for C&I and real estate lending (16 moments). The model-implied moments are computed by simulating the estimated policy along a panel impulse path: the capital requirement jumps by σ_κ and decays at rate ρ_κ, and the response propagates through the portfolio adjustment lags and the equity law of motion, so the model impulse responses are hump-shaped like their empirical counterparts. Because the only aggregate disturbance in the model is the requirement innovation, the local-projection coefficient on the simulated panel coincides with the impulse response.

Figure 10 and Table V report the fit. The estimated innovation is small and effectively permanent ($\hat{\sigma}_\kappa = 0.00040$, s.e. 0.00002; $\hat{\rho}_\kappa = 0.9987$, s.e. 0.0038—an interior estimate within one standard error of a unit root, with a 95% profile confidence set of $\rho_\kappa \geq 0.994$), consistent with the narrative content of the identified series: post-crisis capital regulation changes were regime changes, not transitory shocks. The model reproduces the sign of all sixteen moments, the sectoral ordering—C&I contracts more than real estate at every horizon—and the hump-shaped transmission with troughs at three to four years. Formally, the Hansen J-statistic rejects the two-parameter process ($J = 1822$ under the bootstrap weighting matrix, 14 degrees of freedom). This is a mechanical consequence of the precision with which the LP coefficients are estimated on a panel of roughly 4,800 banks: the optimal weighting matrix assigns very large weight to even economically small deviations, so the J-test has high power against any parsimonious structural model.[3] The economically relevant misfit is concentrated in the depth of the mid-horizon troughs: the model generates a C&I-to-RE response ratio of roughly 1.2, against approximately 1.4 in the data, so the pure risk-weight wedge accounts for the ordering and the bulk—though not all—of the observed sectoral asymmetry. The residual asymmetry is not an artifact of the preset C&I adjustment cost: re-deriving

[3]The same tension between statistical rejection and economic fit arises whenever SMM or GMM targets are estimated on large micro panels. Bloom (2009) matches firm-level investment moments with a small set of uncertainty parameters and reports formal rejection alongside close economic fit; David *et al.* (2016) encounter the identical issue fitting bilateral trade and asset-pricing moments from cross-country panels. In each case the J-test rejects because the moments are estimated too precisely for the model to match exactly, not because the model misses the economically relevant patterns.

the calibration under alternative ψ_B values shows, first, that the steady-state targets jointly discipline (ψ_M, ψ_B, γ)—only $\psi_B \in (0.014, 0.142)$ admits positive frictions—and second, that the constrained-regime sectoral asymmetry $s_B - s_M$ remains large and negative across the admissible range (from -3.5 at $\psi_B = 0.02$ to -0.8 at $\psi_B = 0.10$; -2.2 at the baseline), while flattening the risk weights collapses it to zero at any ψ_B. The asymmetry is a risk-weight phenomenon, with ψ_B governing only its amplitude.

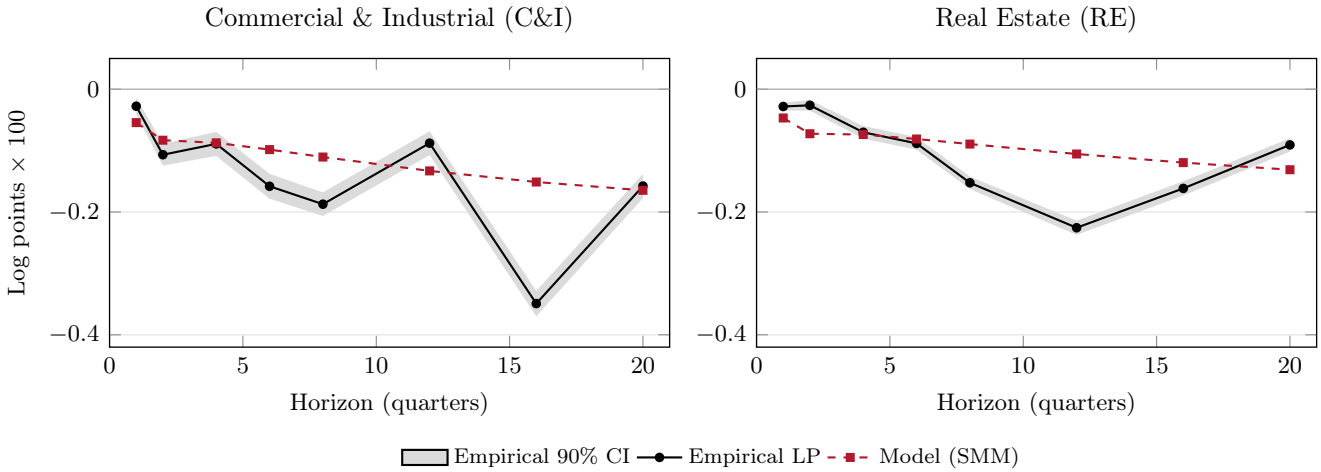


FIGURE 10.—Sectoral credit responses: model vs. within-bank local projections

Notes: Empirical within-bank local projections on the bank panel (bank fixed effects, standard errors clustered by bank; 90% confidence bands shaded) versus the SMM-estimated model's simulated impulse path. Horizons in quarters; responses in log points (×100) per unit of the identified capital-requirement innovation.

TABLE V
SMM MOMENT FIT

Sector	Horizon h	$\hat{b}_h$ (data)	s.e.	b_h (model)	t
C&I	1	-0.028	0.008	-0.054	+3.20
	2	-0.107	0.011	-0.083	-2.20
	4	-0.089	0.012	-0.087	-0.14
	6	-0.158	0.012	-0.098	-4.85
	8	-0.187	0.012	-0.110	-6.57
	12	-0.088	0.012	-0.133	+3.84
	16	-0.349	0.013	-0.151	-15.73
	20	-0.158	0.012	-0.165	+0.58
RE	1	-0.028	0.004	-0.047	+4.87
	2	-0.026	0.005	-0.073	+8.80
	4	-0.070	0.006	-0.074	+0.59
	6	-0.088	0.006	-0.081	-1.15
	8	-0.152	0.006	-0.089	-10.29
	12	-0.226	0.007	-0.106	-17.02
	16	-0.162	0.007	-0.120	-5.97
	20	-0.091	0.007	-0.131	+6.16
Hansen $J = 1822.15$, dof $= 14$, $\chi^2_{0.05} = 23.68$, $p = 0.000$					

Notes: Empirical within-bank local-projection coefficients $\hat{b}_h$ (log points $\times 100$) versus model-implied responses at the SMM estimate; t is the studentised residual. With panel-precise moments the J-test has high power; see the text for the economic reading of the rejection.

5.2 Solution accuracy

The stationary and κ-augmented policies are approximated with the deep-learning method of Section 4.4. Following best practice for global solution methods, Figure 11 maps the Euler-equation residuals across the state space rather than reporting only an average. Accuracy is highest in the interior of the buffer distribution and deteriorates toward the regulatory frontier, which is where the constraint-driven nonlinearity operates; we therefore read the model's cross-sectional predictions near the frontier as conservative. Online Appendix Table OA.II demonstrates that these accuracy patterns are robust to the neural-network architecture: varying the number of hidden layers (2–6), width (128–512 neurons per layer), and activation function (ReLU, GELU, Tanh) leaves the median Euler-equation RMSE within ± 0.02 log units of the baseline and does not alter the qualitative conclusions.

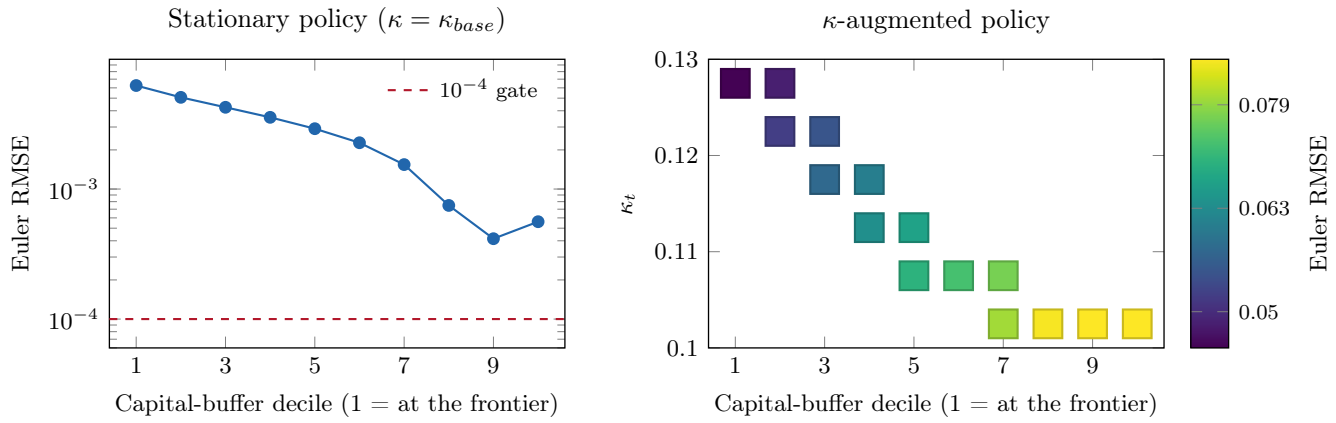


FIGURE 11.—Solution accuracy across the state space

Notes: Left: Euler-equation RMSE of the stationary policy by capital-buffer decile (decile 1 = banks at the regulatory frontier), evaluated on the policy's ergodic distribution; the dashed line marks the 10^{-4} accuracy gate. The spike near the frontier reflects the kink in optimal policy functions at the occasionally binding regulatory constraint, which is the hardest region for the neural-network approximation to resolve; all deciles remain below 10^{-2}. Right: residual map of the κ-augmented policy over the (buffer, κ_t) plane.

5.3 The gross cost of tightening and policy counterfactuals

Table VI reports the consumption-equivalent cost of a permanent one-percentage-point tightening and the policy counterfactuals of Section 4.7.2. The headline cost is $\Lambda = 2.9\%$ of steady-state credit-market surplus under the baseline metric and $\Lambda_V = 2.4\%$ under the franchise-value metric; both decompose additively into scale and allocation channels.

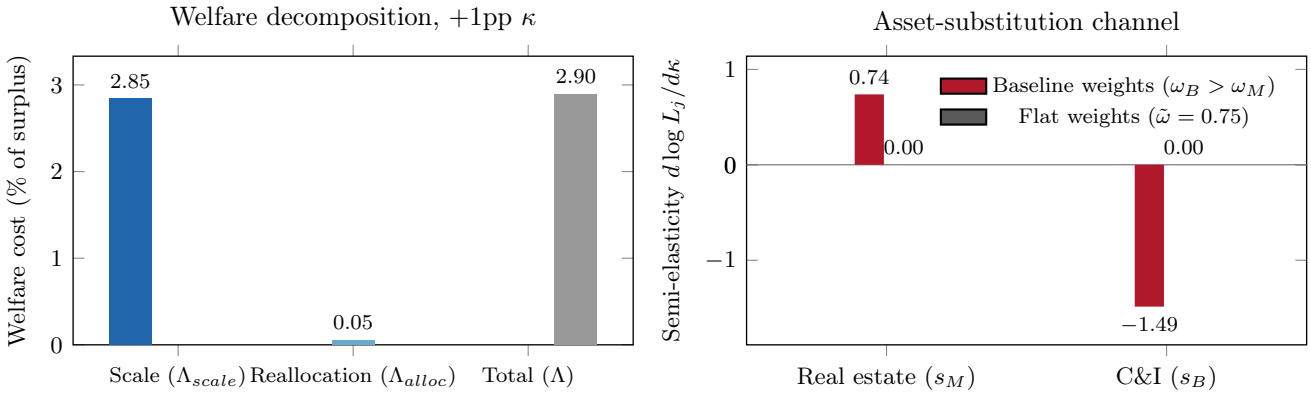


FIGURE 12.—Welfare decomposition and the asset-substitution channel

Notes: Left: decomposition of the consumption-equivalent (CEV) welfare cost of a +1pp tightening into scale and reallocation channels (eq. 32), in percent of steady-state credit-market surplus. Right: constrained-regime semi-elasticities of sectoral lending with respect to the capital requirement (Counterfactual 2). The positive real-estate response is the substitution effect: when the requirement tightens, constrained banks reallocate toward the low-risk-weight sector while cutting high-risk-weight C&I lending. Under flattened weights ($\tilde{\omega} = 0.75$) both semi-elasticities collapse to zero (the gray bars appear at zero by design), confirming that the Asset Substitution Margin vanishes entirely when the risk-weight wedge is removed.

TABLE VI
GROSS COST OF A +1PP TIGHTENING AND COUNTERFACTUALS

Experiment	Quantity	Value
CEV of +1pp κ tightening (eq. 37–38)		
Total welfare cost	Λ	2.898%
scale channel	Λ_{scale}	2.846%
allocation channel	Λ_{alloc}	0.054%
additivity check	$\Lambda_{scale} + \Lambda_{alloc}$	2.900%
Credit contraction (total / RE / C&I)	ΔL	-3.09% / -3.02% / -3.50%
CEV from the franchise value V (discounted dividends)		
Total welfare cost	Λ_V	2.404%
scale channel	$\Lambda_{V,scale}$	2.392%
allocation channel	$\Lambda_{V,alloc}$	0.026%
additivity check	sum	2.419%
CF1: uniform vs. size-differentiated κ (eq. 39)		
Differentiated rule	Λ	2.682%
Uniform $\tilde{\kappa} = 0.1246$	Λ	2.680%
CF2: flat risk weights $\tilde{\omega} = 0.75$ (eq. 40)		
Baseline sectoral asymmetry	$s_B - s_M$	-2.224
Flat-weight sectoral asymmetry	$s_B - s_M$	-0.000

Two lessons follow. First, the aggregate welfare cost is a scale phenomenon: Λ^{scale} accounts for essentially the entire cost under both metrics, while the allocation term is an order of magnitude smaller. The *level*

of capital requirements is what matters for the aggregate cost of regulation. Second, risk-weight design governs incidence, not the aggregate: flattening risk weights (CF2) eliminates the sectoral asymmetry in lending responses—the semi-elasticity gap $s_B - s_M$ collapses from -2.2 to zero—leaving the aggregate cost nearly unchanged, and the uniform-requirement counterfactual (CF1) shows that size-tailoring the requirement redistributes, but barely reduces, the aggregate burden. Differentiated risk weights and tiered requirements are therefore best understood as distributional instruments: they choose which sectors and which banks absorb a given tightening.

5.4 Magnitudes in the context of the literature

Table VII places the estimated magnitudes alongside the quasi-experimental evidence cited above, converting each study's headline result, where possible, to the effect of a one-percentage-point increase in required capital. The model's stationary comparison—a 3.5% contraction in C&I lending, 3.0% in real estate, and 3.1% in total credit per +1pp—sits squarely inside the range spanned by the reduced-form literature: above the intensive-margin flow effects estimated from the EBA capital exercise (Mésonnier & Monk, 2015) and Belgian Pillar-2 decisions (De Jonghe *et al.*, 2020a), close to the linearised medium-run response in the Italian narrative evidence (Conti *et al.*, 2023), and below the UK bank-level estimates of Aiyar *et al.* (2016), which capture a regime with frequent, bank-specific requirement changes. The trough responses implied by our local projections are larger, as expected of peak effects: scaling the within-bank IRFs through the SMM mapping (one unit of the identified innovation moves the requirement by $\hat{\sigma}_\kappa = 4.0$ basis points) implies peak contractions of roughly 9% for C&I and 6% for real estate per percentage point, an out-of-sample linear extrapolation that should be read as an upper bound. We note that the identified shock of Drechsel & Miura (2025) is scaled in high-frequency bank equity price surprises rather than requirement points, so all such conversions necessarily run through the model's estimated mapping.

On the welfare side, our headline cost—$\Lambda \approx 2.4$–2.9% of steady-state credit-market surplus for a permanent +1pp tightening, essentially all of it through the scale channel—is consistent in spirit with both ends of the cited literature. Kisin & Manela (2016) estimate a modest *private* shadow cost of tighter requirements for large U.S. banks ($220 million per year in aggregate, about 0.4% of profits per bank), reflecting banks' ability to arbitrage the constraint at the margin; our partial-equilibrium surplus measure is larger precisely because it prices the induced credit contraction rather than the compliance cost. Conversely, structural analyses such as Mendicino *et al.* (2020) emphasise that transition costs of

the kind we measure erode about a quarter of the long-run welfare *gains* from safer banks—a reminder that our Λ is a gross cost that abstracts from the stability benefits of capital (Section 4.7.2). In the same spirit, Begenau (2020) shows that equilibrium welfare can *rise* with higher requirements once reduced bank-default probabilities are priced in; our framework is deliberately silent on this benefit side, so the cost figures we report should be understood as inputs to, not conclusions about, the optimal-requirement question.

TABLE VII
CREDIT-SUPPLY EFFECTS OF CAPITAL REQUIREMENTS: THIS PAPER VS. THE CITED LITERATURE

Study	**Design / setting**	**Effect of +1pp required capital**
Aiyar *et al.* (2016)	UK bank-level requirement changes, 1998–2007	Lending growth 6.5–7.2pp lower
Mésonnier & Monk (2015)	EBA 2011–12 capital exercise, euro area	Annualised loan growth 1.2–1.6pp lower
De Jonghe *et al.* (2020a)	Pillar-2 decisions, Belgian credit register, 2013–15	Corporate credit growth $\approx$ 0.2–0.4pp per quarter lower
Conti *et al.* (2023)	Narrative BVAR, Italy (EBA/CA episodes)	Loan supply -1.5% after 2 years per +40bp Tier 1 ($\approx -3.8\%$ per pp, linearised)
This paper (model)	SMM-estimated policy, permanent +1pp κ	C&I -3.5%, RE -3.0%, total credit -3.1% (stationary comparison)
This paper (LPs)	Within-bank LPs scaled through the SMM mapping (one shock unit $\mapsto \hat{\sigma}_\kappa = 4.0$bp of κ)	Trough responses $\approx -8.8\%$ (C&I, h=16) and $\approx -5.7\%$ (RE, h=12); out-of-sample linear extrapolation

Notes: Literature estimates as reported by the original studies, converted to a +1pp change in required capital where the original units differ; conversions are linear. Flow effects (growth-rate reductions over months–quarters) and level effects (log-level responses at a horizon) are not directly comparable; the table reports each in its native form. This paper's model numbers compare stationary equilibria under κ_{base} and $\kappa_{base} + 0.01$.

6 Discussion

Three policy lessons follow from the analysis. First, the aggregate cost of capital regulation is governed by the *level* of requirements, not by their design. A permanent one-percentage-point tightening costs 2.4–2.9% of steady-state credit-market surplus, and virtually all of that cost flows through the contraction in the scale of intermediation. Flattening risk weights or replacing the size-tiered schedule with a uniform requirement of equal aggregate capital leaves the total burden nearly unchanged. Debates about the architecture of requirements are therefore, to a first order, debates about incidence rather than aggregate cost.

Second, incidence operates along two dimensions that regulators can read directly off bank balance sheets.

Across asset classes, the risk-weight wedge concentrates the contraction in high-risk-weight lending: C&I credit bears roughly forty percent more of the adjustment than real estate, in the data and in the model, so tightening episodes tax business lending–intensive intermediation most heavily. Across banks, the distance to the regulatory frontier determines who adjusts: thinly buffered institutions cut credit sharply while well-buffered lenders absorb historical innovations almost entirely within their Insulated Zone. Notably, this state-dependence is not a size story—voluntary buffers vary widely within every size tier—which suggests supervisory attention should track the *buffer distribution*, consistent with the "do not cross" behaviour documented for banks approaching their buffers (Couaillier *et al.*, 2025), rather than statutory size categories alone.

Third, the muted historical response of well-buffered banks is not evidence of immunity. In the model, the same precautionary buffers that insulated lenders against the moderate, if persistent, requirement changes of 1994–2019 are a finite resource: counterfactual shocks large enough to erode them trigger contractions that are sharply non-linear in the shock size. The apparent resilience of the small-bank sector is thus a latent vulnerability, and extrapolating historical semi-elasticities to unprecedented tightenings will understate their effects. The same logic, run in reverse, speaks to the current policy debate on buffer usability in downturns—the countercyclical capital buffer (CCyB) debate—where releases are most potent precisely for the institutions operating near the frontier.

These conclusions come with the caveats of Section 4.7.2: the framework is partial equilibrium and prices no stability benefits of capital, so the welfare figures are gross costs of tightening, informative about transition burdens and their distribution rather than about optimal requirement levels.

7 Conclusion

This paper studies how capital requirements transmit to the supply of credit when banks are heterogeneous. Empirically, narratively identified capital-requirement shocks traced through a quarterly panel of U.S. banks reveal a sectoral lending channel—within banks, high-risk-weight C&I lending contracts significantly more than real estate lending, while consumer credit does not respond—and a cross-sectional gradient organised by the distance to the regulatory frontier rather than bank size per se. Structurally, a heterogeneous macro-banking model in which funding frictions generate endogenous precautionary buffers, solved globally with deep-learning methods and estimated by simulated method of moments on the within-bank impulse responses, reproduces the ordering, the hump-shaped dynamics, and the bulk of the

magnitude of these responses with a two-parameter, effectively permanent requirement process.

The quantitative message is that the welfare cost of tightening is a scale phenomenon: differentiated risk weights and size-tiered requirements determine which sectors and which banks bear the contraction but contribute little to its aggregate size, while the buffer distribution determines the economy's exposure to future, larger tightenings. We emphasise once more that these are gross transition costs measured in a partial-equilibrium framework that prices no stability benefits of capital: nothing in our results speaks to whether requirements are too high or too low, only to who bears the burden of changing them. Extending the framework to general equilibrium with a priced default channel—so that the costs measured here can be weighed against the stability benefits of capital—and solving the fully size-dependent regulatory step function with a realistic size distribution remain natural next steps.

References

Acharya, Viral V., Engle, Robert, & Pierret, Diane. 2014. Testing Macroprudential Stress Tests: The Risk of Regulatory Risk Weights. *Journal of Monetary Economics*, **65**, 36–53.

Acharya, Viral V., Le, Hanh T., & Shin, Hyun Song. 2017. Bank Capital and Dividend Externalities. *The Review of Financial Studies*, **30**(3), 988–1018.

Acharya, Viral V., Bergant, Katharina, Crosignani, Matteo, Eisert, Tim, & McCann, Fergal. 2022. The Anatomy of the Transmission of Macroprudential Policies. *The Journal of Finance*, **77**(5), 2533–2575.

Adelino, Manuel, Schoar, Antoinette, & Severino, Felipe. 2025. Credit supply and house prices: Evidence from mortgage market segmentation. *Journal of Financial Economics*, **163**, 103958.

Aiyar, Shekhar, Calomiris, Charles W, & Wieladek, Tomasz. 2016. How does credit supply respond to monetary policy and bank minimum capital requirements? *European Economic Review*, **82**, 142–165.

Begenau, Juliane. 2020. Capital requirements, risk choice, and liquidity provision in a business-cycle model. *Journal of Financial Economics*, **136**(2), 355–378.

Bloom, Nick. 2009. The Impact of Uncertainty Shocks. *Econometrica*, **77**(3), 623–685.

Buchak, Greg, Matvos, Gregor, Piskorski, Tomasz, & Seru, Amit. 2018. Fintech, Regulatory Arbitrage, and the Rise of Shadow Banks. *Journal of Financial Economics*, **130**(3), 453–483.

Conti, Antonio M, Nobili, Andrea, & Signoretti, Federico M. 2023. Bank capital requirement shocks: A narrative perspective. *European Economic Review*, **151**, 104254.

Corbae, Dean, & Erasmo, Pablo D. 2021. Capital Buffers in a Quantitative Model of Banking Industry Dynamics. *Econometrica*, **89**(6), 2975–3023.

Couaillier, Cyril, Reghezza, Alessio, dâAcri, Costanza Rodriguez, & Scopelliti, Alessandro. 2025. How to release capital requirements in an economic downturn? Evidence from euro area credit register. *Journal of Financial Intermediation*, **63**, 101148.

David, Joel M., Henriksen, Espen, & Simonovska, Ina. 2016. The Risky Capital of Emerging Markets. *Journal of Finance*, **71**(5), 2107–2144.

De Jonghe, Olivier, Dewachter, Hans, & Ongena, Steven. 2020a. Bank capital (requirements) and credit supply: Evidence from pillar 2 decisions. *Journal of Corporate Finance*, **60**, 101518.

De Jonghe, Olivier, Dewachter, Hans, Mulier, Klaas, Ongena, Steven, & Schepens, Glenn. 2020b. Some borrowers are more equal than others: Bank funding shocks and credit reallocation. *Review of Finance*, **24**(1), 1–43.

Drechsel, Thomas, & Miura, Ko. 2025. The macroeconomic effects of bank regulation: New evidence from a high-frequency approach.

Drechsler, Itamar, Savov, Alexi, & Schnabl, Philipp. 2017. The Deposits Channel of Monetary Policy. *Quarterly Journal of Economics*, **132**(4), 1819–1876.

Elenev, Vadim, Landvoigt, Tim, & Van Nieuwerburgh, Stijn. 2021. A macroeconomic model with financially constrained producers and intermediaries. *Econometrica*, **89**(3), 1361–1418.

Gropp, Reint, Guettler, Andre, & Saadi, Vahid. 2020. Public bank guarantees and allocative efficiency. *Journal of Monetary Economics*, **116**, 53–69.

Jiménez, Gabriel, Ongena, Steven, Peydró, José-Luis, & Saurina, Jesús. 2014. Hazardous times for monetary policy: What do twenty-three million bank loans say about the effects of monetary policy on credit risk-taking? *Econometrica*, **82**(2), 463–505.

Jimenez, Gabriel, Ongena, Steven, Peydró, José-Luis, & Saurina, Jesús. 2017. Macroprudential Policy, Countercyclical Bank Capital Buffers, and Credit Supply: Evidence from the Spanish Dynamic Provisioning Experiments. *Journal of Political Economy*, **125**(6), 2126–2177.

Jordà, Òscar. 2005. Estimation and inference of impulse responses by local projections. *American economic review*, **95**(1), 161–182.

Kisin, Roni, & Manela, Asaf. 2016. The Shadow Cost of Bank Capital Requirements. *The Review of Financial Studies*, **29**(7), 1780–1820.

Kovner, Anna, & Van Tassel, Peter. 2022. Evaluating regulatory reform: Banks' cost of capital and lending. *Journal of Money, Credit and Banking*, **54**(5), 1313–1367.

Krusell, Per, & Smith, Jr, Anthony A. 1998. Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy*, **106**(5), 867–896.

Le Leslé, Vanessa, & Avramova, Ms Sofiya. 2012. *Revisiting risk-weighted assets*. International Monetary Fund.

Maliar, Lilia, Maliar, Serguei, & Winant, Pablo. 2021. Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, **122**, 76–101.

Mendicino, Caterina, Nikolov, Kalin, Suarez, Javier, & Supera, Dominik. 2020. Bank capital in the short and in the long run. *Journal of Monetary Economics*, **115**, 64–79.

Mésonnier, J-S, & Monk, A. 2015. Heightened bank capital requirements and bank credit in a crisis: the case of the 2011 EBA Capital Exercise in the euro area. *Rue de la Banque*.

Plagborg-Møller, Mikkel, & Wolf, Christian K. 2021. Local projections and VARs estimate the same impulse responses. *Econometrica*, **89**(2), 955–980.

Sivec, Vasja, & Volk, Matjaž. 2021. Bank response to policy-related changes in capital requirements. *The Quarterly Review of Economics and Finance*, **80**, 868–877.

Van den Heuvel, Skander J, *et al.* 2002. The bank capital channel of monetary policy. *The Wharton School, University of Pennsylvania, mimeo*, 2013–14.