The Shadow Cost of Supervision: Evidence from Prompt Corrective Action Thresholds

Andrew Bird[1] Stephen A. Karolyi[2] Thomas G. Ruchti[3]

March 2026

Abstract

We estimate the shadow cost of bank supervision using bunching at Prompt Corrective Action capital thresholds. Banks distort their pivotal capital ratio by 4.17 percentage points to maintain well-capitalized status, implying an aggregate cost of $79 billion per year ($48–64 billion excluding stress-tested banks). A sufficient-statistics decomposition attributes 27% (1.11 pp) to observable balance sheet channels—dominated by constrained asset growth (0.81 pp)—with the remaining 73% providing an upper bound on soft supervisory burdens—including interaction effects between observable channels—that are invisible in balance sheet data. Neither bunching nor outcome discontinuities appear before PCA's enactment.

JEL Classification: G21, G28, G32

Keywords: Bank supervision, capital requirements, shadow cost, bunching, regression discontinuity, Prompt Corrective Action

[1]Argyros College of Business and Economics, Chapman University, One University Drive, Orange, CA 92866. Email: bird@chapman.edu.

[2]School of Business, George Mason University, 4400 University Drive, MSN 5F5, Fairfax, VA 22030. Email: skarolyi@gmu.edu.

[3]Pamplin College of Business, Virginia Tech, 880 West Campus Drive, Blacksburg, VA 24061. Email: ruchti@vt.edu.

1 Introduction

How much does it cost banks to be supervised? The question has acquired new urgency. All three federal banking agencies are simultaneously rewriting the rules of bank oversight: the OCC is reorganizing into size-based divisions and removing fixed examination requirements, the FDIC is redirecting examiners toward "core, material financial risks," and the Federal Reserve has issued new Supervisory Operating Principles that direct examiners to focus on solvency threats rather than process compliance.[4] These reforms reflect a bipartisan consensus that the cumulative burden of supervision constrains credit supply and reduces bank profitability. Yet no one has measured the magnitude of the burden being reformed. Without an estimate of the private cost of supervision, policymakers cannot evaluate whether proposed reforms are proportionate, and practitioners cannot assess how much profitability they stand to recover.

The setting we exploit is Prompt Corrective Action (PCA)—the backbone of U.S. bank supervision since 1992—which assigns every FDIC-insured bank to one of five capital adequacy categories based on four simultaneous capital ratio thresholds: total risk-based capital ($\geq 10\%$), Tier 1 risk-based capital ($\geq 8\%$), CET1 ($\geq 6.5\%$), and leverage ($\geq 5\%$). A bank must clear *all four* simultaneously to qualify as well capitalized; falling below any single ratio triggers reclassification. The economically relevant margin is well capitalized to adequately capitalized, because consequences escalate sharply at this boundary—brokered deposit prohibition, enhanced examinations, growth restrictions—making the well-capitalized cutoff the critical notch in the supervisory regime. The consequences are severe and escalating. At the well capitalized boundary—where virtually all the bunching occurs—a bank that slips to adequately capitalized loses access to brokered deposits (a high-yield funding source prohibited under 12 CFR 337.6), faces caps on deposit interest rates, and confronts enhanced

[4] OCC News Release 2025-95 (October 6, 2025); Hill, "Rightsizing Regulation to Promote American Opportunity," testimony before the U.S. Senate Committee on Banking, Housing, and Urban Affairs (February 26, 2026); Bowman, "Modernizing Supervision and Regulation: 2025 and the Path Ahead," California Bankers Association (January 7, 2026).

examination frequency and restrictions on branching, acquisitions, and new business lines. At lower categories the consequences grow harsher: undercapitalized banks must submit capital restoration plans and suspend dividends; critically undercapitalized banks must be placed in receivership within 90 days. These are not abstract risks—they are automatic, statutory, and well understood by bank management.

Adjusting capital ratios to stay above these thresholds is costly, and the costs are ordered. The cheapest margin is slowing asset growth—shrinking the denominator of the capital ratio—because it requires only forgoing new lending rather than restructuring existing positions. When growth constraints alone are insufficient, banks turn to costlier margins: retaining more earnings by cutting dividends, shifting loan portfolios toward lower-risk-weight assets, and, as a last resort, issuing new equity—the most expensive margin because it dilutes existing shareholders. Yet the costs of *crossing* the threshold—brokered deposit prohibition, rate caps, enhanced examinations, and the escalating sanctions described above—are costlier still. Under the maintained assumption that banks optimize across these margins (standard in revealed-preference work), they will exhaust the cheapest adjustments before resorting to expensive ones, and they will adjust rather than cross whenever the costs of reclassification exceed the costs of adjustment. The equilibrium we observe—banks stretching their capital ratios well beyond what pure risk-return optimization would dictate—reveals that this condition holds: the crossing costs dominate, and banks respond by bunching above the threshold. The decomposition in Section 6 confirms this cost ordering empirically, showing that banks rely overwhelmingly on growth constraints and only marginally on equity issuance.

What, precisely, is the "shadow cost" we estimate? It is the total private cost that a bank reveals it is willing to bear—through all of the adjustments described above—to avoid reclassification below well-capitalized status. We measure it in capital-ratio units: how many percentage points of capital ratio banks stretch beyond what they would choose absent the threshold. This is not a cost borne only by the handful of banks sitting exactly at the boundary. The bunching estimator captures the behavior of every bank whose unconstrained

optimum would have placed it just below the threshold but who instead pushes itself above—the entire mass of "bunchers" visible as a pile-up in the capital ratio distribution. In dollar terms, if a bank holds Δq extra percentage points of capital ratio, the opportunity cost is the foregone return on the excess equity: $(\Delta q/100) \times \text{RWA} \times r_e$, where RWA is risk-weighted assets and r_e is the cost of equity capital. For a bank with $1 billion in risk-weighted assets and a 10% cost of equity, a 4.17 pp distortion implies an annual shadow cost of $4.2 million. Aggregated across all 16,621 banks in the PCA system, this yields $79 billion per year—a revealed-preference measure backed out from observed behavior, not an assumed or calibrated parameter. Section 6 translates this further into a fraction of bank assets, enabling cross-size comparison.

This paper estimates that shadow cost and decomposes it into its component parts.

Our approach exploits a simple economic insight. If supervision were costless, the distribution of capital ratios near PCA thresholds would be smooth. If instead supervision imposes costs that jump at the threshold, banks will distort their capital ratios to remain on the favorable side, creating a pile-up just above the cutoff and a corresponding hole just below. This "cost technology" works because PCA satisfies three conditions that make it uniquely suited for a bunching design. Two conditions validate the bunching estimate itself: (1) the thresholds are rule-based—statutory and automatic, not discretionary, so the location of the notch is known *ex ante* and identical for every bank; (2) a pre-PCA placebo is available—the 10% threshold carried no regulatory significance before 1992, providing a clean counterfactual era in which to verify that bunching is absent. A third condition enables the decomposition of the total cost into its component parts: (3) rich outcome data—quarterly Call Reports provide approximately 65 outcome variables that allow us to trace the specific channels through which banks respond. Together, these conditions mean that the density distortion at the threshold reads out the magnitude of the cost jump, without requiring assumptions about the shape of the profit function or the competitive structure of banking markets. This revealed-preference logic, first formalized by Kleven and Waseem (2013) and

applied to Dodd-Frank size thresholds by Alvero and Xiao (2023), has not previously been brought to PCA because the setting is more complex: a bank must satisfy four capital ratios simultaneously, so the analyst must determine which threshold is binding. We resolve this by constructing a "pivotal" running variable—the minimum distance to any binding threshold—that reduces the multidimensional problem to a single margin.

Bunching tells us the total cost of supervision, but not how banks bear it. The $79 billion could reflect constrained lending, forced equity retention, lost funding access, or examination burdens that never appear on a balance sheet—and the policy implications differ sharply depending on the answer. To decompose the shadow cost, we pair the bunching analysis with a regression discontinuity design (RDD) at the same threshold. Banks just above the well capitalized cutoff serve as the treatment group; banks just below serve as controls. Because PCA classification is mechanically determined by the capital ratio, and because we verify that the distribution was smooth at this threshold before PCA existed (1990–1992), the design isolates the causal effect of supervision on each bank outcome separately. The RDD estimates reveal that supervision operates as a multi-dimensional constraint: well-capitalized banks grow 5.72 pp slower, hold 11.78 pp fewer loans as a share of assets, earn 0.28 pp higher ROA (by shifting toward safer assets with lower credit losses), hold significantly more brokered deposits, and face different acquisition probabilities—each representing a distinct cost margin through which the shadow cost is borne.

The final step—and the one that distinguishes our approach from prior work—is to convert each channel-specific RDD estimate into a common unit so the pieces can be compared to the whole. The balance sheet accounting identity provides the mapping. Because the capital ratio equals equity divided by risk-weighted assets, any change in the ratio must operate through one of four flows: earnings (ROA), dividend payments, equity issuance, or asset growth. The equity flow identity $\Delta(E/A) \approx \text{ROA} - \text{Div}/A + \text{Issuance}/A - (E/A) \times g_A$ decomposes the change in capitalization exhaustively into these four terms. We estimate the RDD effect on each term, convert each into capital-ratio-equivalent units using scaling pa-

rameters estimated directly from the near-threshold sample (not calibrated), and sum them. This is not a regression R^2 or a structural model—it is an accounting decomposition that asks: of the 4.17 pp total distortion revealed by bunching, how many percentage points can be attributed to each measurable flow? The answer is that these four flows account for 1.11 pp (27% of Δq), with asset growth contributing 0.81 pp (19%), ROA contributing 0.26 pp (6%), and issuance contributing 0.05 pp (1%). The remaining 3.06 pp (73%)—formally rejected as zero with $z = 14.96$—provides an upper bound on the costs that do not map one-to-one through the accounting identity: interaction effects between channels, risk-weight adjustments through portfolio reallocation, and genuinely "soft" supervisory costs—examination burden, reputational stigma, management constraints, and the lost option value of regulatory flexibility—that are invisible in financial statements.

This "bunching + RDD + sufficient statistics" combination has not been implemented before. Existing estimates of supervisory costs are either structural—delivering aggregate costs but requiring functional form assumptions—or reduced-form—identifying specific channels but lacking a common unit for aggregation. Our method bridges this gap by providing both the total and the decomposition with formal inference (Section 2 discusses the relationship to prior work in detail). The shadow cost captures both the regulatory burden and the market-signaling value of well-capitalized status (Berger et al., 2000)—and signaling is not a caveat that weakens the estimate but an integral component of the cost banks bear. The pre-PCA placebo bounds the signaling channel: if the market value of a 10% capital ratio alone drove bunching, we would observe it before 1992, when the threshold carried no regulatory significance. The absence of pre-PCA bunching implies that the regulatory channel dominates, though we cannot fully separate the two.

Four findings anchor the paper. First, the density discontinuity yields $\Delta q = 4.17$ pp (95% CI: $[3.85, 4.53]$), concentrated in total risk-based capital with no significant manipulation of Tier 1 or CET1—banks manage the single most salient ratio. The aggregate shadow cost is \$79 billion per year (\$48–64 billion excluding CCAR-subject banks). The implied

cost function is convex: the empirical magnitudes mirror the economic cost ranking from the introduction—asset growth, the cheapest margin because it requires only forgoing new lending, absorbs the largest share (19.4% of Δq); ROA, reflecting the costlier step of retaining earnings, contributes 6.1%; and equity issuance, the most expensive margin because it dilutes shareholders, contributes just 1.1%. The ordering of empirical magnitudes matches the ordering of economic costs, consistent with increasing marginal costs across successive adjustment channels.

Second, the decomposition attributes 27% of Δq to observable balance sheet flows (95% CI: $[22\%, 31\%]$), dominated by constrained asset growth (19%). The remaining 73%—formally rejected as zero ($z = 14.96$)—bounds the costs that do not map one-to-one through the accounting identity: interaction effects between channels, risk-weight adjustments, and genuinely soft supervisory costs (examination burden, reputational stigma, management constraints, and the lost option value of regulatory flexibility). Reform proposals must address both threshold-based constraints and these soft burdens.

Third, a "smoking gun" test validates the causal interpretation. Banks just above the well capitalized threshold hold significantly more brokered deposits as a share of assets than banks just below, directly reflecting the statutory prohibition on brokered deposits for banks that lose well capitalized status (12 CFR 337.6). The identical test applied to the pre-PCA period (1990–1992) produces a precisely estimated zero. Of 14 level outcomes tested in the pre-PCA period, none is significant at the 5% level, confirming that the effects we document are caused by PCA, not by pre-existing correlations between capital and bank behavior.

Fourth, the shadow cost is concentrated across the size distribution: banks above \$10 billion (fewer than 3% of institutions) bear a cost of 0.32% of assets versus 0.24% for community banks. In percentage-of-assets terms—computed as $(r_e \times \Delta q/100) \times (\text{RWA}/A)$, where RWA/A is the average risk weight—the shadow cost translates from capital-ratio units (pp) through dollar units (\$79B) to a fraction of the balance sheet, bridging the three units used throughout the paper. The size gradient is driven by higher leverage that amplifies the

opportunity cost of excess capital. Threshold-based reforms primarily benefit community banks; the soft supervisory costs that dominate for the largest institutions require different policy instruments.

A growing literature establishes that supervision yields substantial stability benefits—reducing resolution costs (Correia et al., 2025), preventing the persistent real effects of bank failures (Ivanov and Karolyi, 2023), and mitigating crisis-driven credit contractions (Baron et al., 2021)—but has not measured the private cost side. Our $79 billion estimate provides the missing input for a cost-benefit analysis of bank supervision: paired with the stability benefits quantified by Correia et al. (2025) and others, policymakers can evaluate whether the current supervisory regime is proportionate or whether reforms should recalibrate the balance. The 27%/73% decomposition tells policymakers *where* the cost falls—and therefore which reforms will actually reduce it (Section 2 discusses the relationship to this literature in detail).

The paper proceeds as follows. Section 2 reviews the related literature. Section 3 describes the PCA institutional setting, data, and sample construction. Section 4 documents bunching at PCA thresholds and estimates the total shadow cost in capital-ratio units. Section 5 uses the RDD to identify the specific channels through which banks respond to supervision. Section 6 develops the sufficient-statistics decomposition, translates the shadow cost into dollar values, and explores heterogeneity across the bank size distribution. Section 7 discusses welfare implications and concludes.

2 Related Literature

Our paper builds on three strands of research and contributes a framework that addresses limitations in each.

The first strand estimates the costs of bank capital regulation using structural models. Kisin and Manela (2016) use ABCP conduits as a revealed-preference instrument, finding

that a one-percentage-point increase in Tier 1 capital costs participating banks approximately $220 million per year. Begenau (2020) calibrates a DSGE model in which optimal capital requirements trade off reduced risk-taking against reduced liquidity provision. Moreover, our 73% residual implies that any profit-function-based calibration of supervisory costs—including Begenau (2020)'s—can capture at most the observable share: soft supervisory burdens that do not map through the equity flow identity constrain the curvature parameter from above, suggesting that structural estimates may understate the total cost. Conversely, because the 73% residual is an upper bound—absorbing interaction effects alongside genuine soft costs—prior structural estimates that attribute a smaller share to soft costs are consistent with, not in tension with, our decomposition. These papers deliver aggregate cost estimates but require assumptions about the shape of profit functions, competitive structure, and the mapping from capital to lending—and the aggregate depends on functional form choices the data cannot verify. We differ by exploiting bunching at PCA thresholds to estimate the *total* shadow cost without structural assumptions beyond the smoothness of the counterfactual density, and by decomposing it into observable and unobservable channels via RDD. Our $79 billion provides an independent cross-check on Kisin and Manela (2016)'s per-bank estimate without relying on their ABCP conduit instrument, and the dominance of constrained asset growth (19%) over equity issuance (1%) confirms that banks respond to capital requirements primarily by shrinking the denominator of the capital ratio—inconsistent with models that treat the cost of capital regulation as primarily an equity-issuance cost.

The second strand examines the real effects of supervision using reduced-form methods. Correia et al. (2025) shows that supervision reduces FDIC resolution costs by 9% of assets at the expense of credit supply; Granja and Leuz (2024) find that the OTS abolition increased small business lending by 10%; Kashyap et al. (2025) use RDD at India's PCA thresholds to show that intervention reduces borrower defaults by 3.2 percentage points; Correia et al. (2026) establish that bank failures are fundamentals-driven and predictable from public accounting data, validating PCA's capital-based triggers. Each paper identifies

a specific causal effect but cannot aggregate across channels because it lacks a common unit of measurement. Our shadow cost quantifies the total burden, and the 27%/73% decomposition tells us where the balance sheet contraction these papers emphasize sits within the total: it accounts for roughly one-quarter. The remaining three-quarters operates through channels that reduced-form studies, examining one outcome at a time, cannot capture.

The third strand documents soft supervisory costs. Passalacqua et al. (2021) find that bank inspections cause credit reallocation toward productive firms, implying that some supervisory burden reflects efficient restructuring rather than deadweight loss. Bisetti (2024) provides market-based evidence that regulatory monitoring has positive value for bank shareholders. Our 73% residual—the first formal estimate of the soft cost share, measured with a confidence interval ($[69\%, 78\%]$) rather than assumed in a calibration—provides an upper bound on these costs and disciplines future structural work: any model of supervisory costs must be consistent with the finding that genuinely soft costs dominate the observable balance sheet channels.

Methodologically, we combine the bunching framework of Kleven and Waseem (2013) and Alvero and Xiao (2023) with an RDD in the tradition of Garcia and Steele (2022) and Bonaccorsi di Patti et al. (2023), linked by a sufficient-statistics decomposition (Chetty, 2009) with estimated (not calibrated) conversion parameters and bootstrap inference. Ours is the first paper to decompose a bunching-implied shadow cost into causal channels using a complementary RDD, providing a template for evaluating the cost of regulatory notches in other settings.

3 Institutional Setting and Data

This section describes the PCA framework, data sources, and sample construction.

3.1 The PCA Framework

The Federal Deposit Insurance Corporation Improvement Act of 1991, enacted in response to the savings and loan crisis, established the Prompt Corrective Action framework effective December 1992. The framework was designed to solve a specific problem: during the S&L crisis, regulators had delayed intervention at deteriorating institutions, allowing losses to compound and ultimately costing taxpayers over $130 billion. PCA eliminated this discretion at the margin by mandating specific supervisory responses at predetermined capital thresholds.

PCA classifies banks into five categories based on four capital ratios. A bank must satisfy *all four* thresholds simultaneously to qualify for a given category; falling below any single ratio triggers reclassification regardless of performance on the other three. This "weakest link" rule is central to our empirical design, because it means that for each bank in each quarter, one specific ratio is the binding constraint—the one closest to its threshold.

— Insert Table 1 Here —

The consequences of reclassification are both mandatory and discretionary, and they escalate sharply across categories. At the Well Capitalized boundary—the margin that generates almost all of the bunching we document—mandatory consequences include prohibitions on accepting brokered deposits without an FDIC waiver (12 CFR 337.6), caps on deposit interest rates, and limitations on pass-through deposit insurance. Discretionary consequences include enhanced examination frequency, restrictions on asset growth, branching, and acquisitions, requirements for capital restoration plans, and potential management changes. At lower categories, the consequences grow more severe: Undercapitalized banks face suspension of dividends and mandatory capital restoration plans; Critically Undercapitalized banks must be placed in receivership within 90 days.

A critical feature for our analysis is that the costs of reclassification extend well beyond the explicit statutory restrictions. Supervisors exercise substantial discretion through Matters

Requiring Attention (MRAs) and informal communications that impose compliance costs, constrain strategic decisions, and signal supervisory concern to market participants.[5] These soft costs are invisible in balance sheet data but may constitute a substantial fraction of the total supervisory burden. Our accounting decomposition in Section 6 is designed to separate the observable costs (those that appear as balance sheet changes) from this unobservable residual.

3.2 Data Sources

Our primary data source is the universe of quarterly Call Reports filed by FDIC-insured commercial banks, accessed through the FFIEC Central Data Repository. We obtain balance sheet items, income statement items, and regulatory capital ratios (total risk-based capital, Tier 1, CET1, and leverage, plus the official FDIC PCA classification code). We construct level-based outcome variables from dollar amounts for the RDD and retain growth rates for the decomposition. We supplement these with brokered deposit volumes, branch-level data from the FDIC Summary of Deposits, merger events from the National Information Center, and CRA and SBA lending data for geographic outcomes.

3.3 Sample Construction

The raw data contain 1,665,311 bank-quarter observations across 23,474 unique institutions from 1993Q1 through 2025Q2. We construct the analysis sample in three steps, each motivated by the requirements of our identification strategy.

First, we compute the pivotal running variable. Because PCA classifies banks on the basis of four ratios simultaneously, we need a single measure that captures which ratio is

[5]The October 2025 OCC/FDIC joint proposal to define "unsafe or unsound practice" and restrict MRAs to "material financial risks" signals a policy shift toward constraining this discretion. Joint NPRM, 90 Fed. Reg. 80,456 (October 30, 2025).

binding and by how much. We define the pivotal distance as:

$$X_{it} = \min\left(c_{it}^{\text{Total RBC}} - 10,\ c_{it}^{\text{Tier 1}} - \bar{c}_t^{T1},\ c_{it}^{\text{CET1}} - 6.5\right) \tag{1}$$

where $c^{\text{Total RBC}}$ is the total risk-based capital ratio, $c^{\text{Tier 1}}$ is the Tier 1 risk-based capital ratio, c^{CET1} is the common equity Tier 1 ratio, and $\bar{c}_t^{T1}$ is the time-varying Tier 1 well-capitalized threshold: 6% before 2015 and 8% after Basel III. CET1 enters only from 2015 onward. A bank with $X_{it} > 0$ is well capitalized; $X_{it} < 0$ means at least one threshold is breached. We exclude the leverage ratio because it operates on a different base (total assets rather than risk-weighted assets) and is binding for fewer than 3% of near-threshold bank-quarters. Dropping observations with missing or zero component ratios yields 1,203,240 valid observations.

Second, we construct dollar-level outcome variables directly from balance sheet amounts to ensure the RDD captures stock-based discontinuities; growth-rate outcomes are retained for the decomposition. Third, we merge external data sources using identifier crosswalks, producing a final dataset of 1,665,311 bank-quarter observations with 467 columns. For the RDD, we use the full sample of approximately 1.15 million observations (21,035 unique banks), with standard errors clustered at the bank level.

— Insert Table 2 Here —

Table 2 reports summary statistics. The median bank maintains a total risk-based capital ratio of 15.6%—more than five percentage points above the well-capitalized threshold, consistent with a large shadow cost of approaching it. Approximately 94% of observations are well capitalized. Figure 1 shows that the non–well-capitalized share peaked at 4.7% of banks in 2010, but the asset-weighted share peaked earlier (7.9% in 1993), indicating that reclassification disproportionately affects smaller institutions.

— Insert Figure 1 Here —

3.4 The Pre-PCA Sample

We construct a parallel sample covering the pre-PCA period from 1990 to 1992 (179,755 bank-quarter observations, 15,851 unique institutions), when the 10% threshold carried no formal regulatory significance. We extend this placebo further back to 1984–1989 (406,539 observations, 20,065 institutions) using simulated risk-based capital ratios constructed from Basel I risk weights applied to balance sheet components. Absence of bunching in both pre-PCA windows provides cumulative evidence that the post-PCA effects reflect the regulation rather than pre-existing capital ratio features.

4 Bunching at PCA Thresholds

If PCA supervision imposes costs, those costs should be visible in the distribution of capital ratios: banks with unconstrained optima just below the threshold will distort upward, creating a pile-up above and a hole below. This section documents the bunching, verifies its absence before PCA, and estimates the structural parameter Δq.

4.1 The Distribution of Capital Ratios

Figure 2 plots the distribution of the pivotal running variable for the full post-PCA sample. The pivotal distance equals zero for a bank at exactly the binding well capitalized threshold, positive for banks above, and negative for banks below.

— Insert Figure 2 Here —

Two features are immediately visible: a pronounced excess mass just above zero and a corresponding "missing mass" just below—the hallmark of bunching. The Calonico et al.

(2014) density discontinuity test strongly rejects continuity at the pivotal threshold ($p < 0.001$; Figure 3).[6] Bunching is concentrated in total risk-based capital ($T = 6.00$, $p < 0.001$), while Tier 1 ($T = 1.12$, $p = 0.26$) and CET1 ($T = 0.30$, $p = 0.76$) show no significant discontinuities—banks manage the most salient ratio. Appendix H (Table A.8) reports the full set of tests.

— Insert Figure 3 Here —

4.2 Pre-PCA Density Placebo

The pre-PCA sample tests whether this bunching is inherent to the capital ratio distribution or regulation-driven. Before December 1992, the 10% threshold carried no regulatory significance. The pre-PCA density test fails to reject continuity (1990–1992: $T_{\text{JK}} = -1.14$ for Total RBC, $T_{\text{JK}} = -0.20$ for Tier 1 RBC, both $p > 0.30$). The extended 1984–1989 window likewise shows no evidence of bunching at any simulated threshold (simulated Total RBC: $T = -0.54$, $p = 0.590$; pivotal: $T = -0.22$, $p = 0.826$). The contrast—strong rejection post-PCA, no rejection pre-PCA even during the S&L crisis when capital was under exceptional stress—confirms that bunching is caused by PCA regulation. Appendix Table A.8 reports the complete set of density placebo tests.

4.3 Bunching Estimation

We adopt the local polynomial density discontinuity framework of Cattaneo et al. (2020), which estimates the density separately on each side of the cutoff with data-driven bandwidth selection. Unlike the global polynomial approach of Kleven (2016), local fitting does not require global smoothness and can be applied to the pivotal binding-score, whose min-distance

[6] Standard density tests are applied separately to each ratio; a joint manipulation test for the multivariate running variable (Crippa, 2025) would have greater power but is not yet available.

construction produces a non-smooth density shape. At the well-capitalized threshold, total risk-based capital shows overwhelming bunching ($T = 42.48$, $p < 0.001$), with the density just above exceeding the density just below by $\hat{\rho} = 2.28$. Under locally uniform counterfactual density, the bunching range is $\Delta q = B/\hat{f}^{-}(0)$ pp, measuring the maximum distortion the marginal buncher will incur to avoid reclassification.

— Insert Table 3 Here —

Among individual ratios, total risk-based capital dominates ($T = 42.48$, $p < 0.001$, $\Delta q = 3.32$ pp), consistent with banks managing the most salient and widely reported ratio. Tier 1 capital shows no significant discontinuity ($T = 0.04$, $p = 0.97$) under time-varying thresholds (6% pre-2015, 8% post-Basel III), suggesting that the Tier 1 threshold rarely binds independently, and CET1—available only from 2015 and estimated on roughly one-sixth the sample—likewise shows no discontinuity ($T = -0.40$, $p = 0.69$). The pivotal running variable, which captures the binding constraint across all ratios, yields a larger distortion than total RBC alone ($T = 51.29$, $p < 0.001$, $\Delta q = 4.17$ pp, 95% CI $[3.85, 4.53]$), reflecting the additional distortion required to clear all three ratio thresholds simultaneously. Total risk-based capital remains overwhelmingly the binding constraint.

The implied Δq is the central structural parameter of the paper: it measures the total cost of supervision in capital-ratio units, but tells us nothing about *how* banks bear this cost. Decomposing Δq into specific channels requires estimating the causal effect of PCA classification on individual outcomes.

5 The Channels of Supervision: RDD Estimates

The regression discontinuity design at the well-capitalized threshold identifies the causal effect of PCA classification on each bank outcome separately, providing the channel-specific estimates needed for the decomposition.

5.1 Empirical Specification

Banks just above the threshold are nearly identical to those just below, except the former enjoy unrestricted status while the latter face mandatory restrictions. Any discontinuous change at the threshold is caused by PCA classification. We estimate:

$$Y_{it} = \alpha + \tau \cdot \mathbf{1}[X_{it} \geq 0] + f(X_{it}) + \varepsilon_{it} \tag{2}$$

where Y_{it} is the outcome, X_{it} is the pivotal distance to the threshold, and $f(\cdot)$ is a local polynomial fit separately on each side of the cutoff. The coefficient τ captures the causal effect of well capitalized classification.

We implement the bias-corrected local polynomial estimator of Calonico et al. (2014, 2020) with a triangular kernel, linear local polynomial, and MSE-optimal bandwidth selection.

Because well capitalized status requires simultaneously meeting all four capital ratio thresholds, the PCA design is a multi-score RDD with an AND rule. We follow the binding-score approach of Reardon and Robinson (2012), collapsing the running variable into a single scalar: $X_{it} = \min_r(c_{rit} - \bar{c}_r)$. This is formally equivalent to the distance-to-boundary metric of Cattaneo et al. (2016).[7] Because PCA is a strict AND rule—no supervisory consequence attaches to falling below a single ratio while remaining above the others—the partial-effects concern of Choi and Lee (2018) is mitigated. As a diagnostic, separate-ratio RDDs using only total risk-based capital (Appendix Table A.3) produce qualitatively consistent results.

Three identification assumptions underpin the design: (i) continuity of potential outcomes at the threshold, supported by pre-PCA placebo tests (Section 5.4); (ii) imprecise control over the running variable, supported by the Gerard et al. (2020) bounds analysis (Section 5.5); and (iii) SUTVA, addressed via the Berg et al. (2021) spillover diagnostic (Section 5.5).

[7] Cattaneo et al. (2025) note that distance-based methods can exhibit bias at boundary kinks; in our setting, fewer than 5% of near-threshold observations lie at intersections where two thresholds bind simultaneously.

5.2 Main Results

Tables 4–6 report RDD estimates organized by economic dimension: stock positions (Table 4), flows and strategic decisions (Table 5), and geographic market access (Table 6). Figure 4 presents visual evidence for six key outcomes. The flow effects map directly into the decomposition identity $\Delta c \approx \text{ROA} - \text{Div}/A + \text{Issuance}/A - c \times g_A$.

— Insert Tables 4–5 Here —

— Insert Figure 4 Here —

The most direct test of PCA's mandatory restrictions concerns brokered deposits. Under 12 CFR 337.6, banks that lose well capitalized status are prohibited from accepting brokered deposits without an FDIC waiver. Well-capitalized banks hold 0.37 pp more brokered deposits as a share of assets ($p = 0.024$) and are 20.6 pp more likely to use them at all ($p < 0.001$). The identical test in the pre-PCA period produces no significant discontinuity—the combination of a positive post-PCA effect, a precisely estimated zero pre-PCA, and a clear statutory mechanism provides the strongest single piece of evidence that PCA supervision has real, binding consequences.

The brokered deposit result is embedded in a broader pattern of conservative funding. Well-capitalized banks maintain higher core deposit shares (+7.04 pp), more cash (+4.00 pp), and substantially larger liquid asset buffers (+13.26 pp), all significant at $p < 0.001$. Deposit growth is 6.45 pp lower ($p < 0.001$), and the loan-to-deposit ratio is 5.57 pp lower ($p < 0.001$). This combination of statutory access to brokered deposits with a conservative funding structure suggests that the option value of well-capitalized status extends well beyond any single regulatory provision: banks near the threshold constrain growth and hold liquid buffers as part of a deliberate capital management strategy.

The funding conservatism is accompanied by a substantial portfolio reallocation. Banks above the threshold hold 11.78 pp fewer loans and 8.73 pp more securities as a share of assets ($p < 0.001$ for both), reflecting a shift from higher-yield, higher-risk-weight assets toward safer, more liquid positions that mechanically raise the capital ratio. Within the loan portfolio, well-capitalized banks tilt toward residential real estate (+3.43 pp) and away from C&I (−3.69 pp) and consumer lending (−4.69 pp). NPL ratios are 0.68 pp lower ($p < 0.001$), consistent with de-risking. This portfolio channel operates through the denominator of the capital ratio: replacing loans (typically 100% risk weight) with securities (typically 20%) reduces risk-weighted assets per dollar of total assets, improving capitalization without requiring additional equity.

The portfolio and funding effects are reinforced by a growth constraint that turns out to be the dominant observable channel in the decomposition. Well-capitalized banks have 5.72 pp lower asset growth and 1.39 pp lower equity growth ($p < 0.001$ and $p = 0.006$). Because asset growth falls proportionally more than equity growth, the equity-to-assets ratio rises sharply (+5.13 pp, $p < 0.001$)—banks maintain their capital cushion primarily by constraining the denominator rather than expanding the numerator. Equity issuance is 0.033 pp *lower* for well-capitalized banks ($p = 0.024$), consistent with supervisory recapitalization pressure on banks below the threshold rather than voluntary equity raising above it. This asset growth channel accounts for 19.4% of the total bunching-implied distortion in the decomposition of Section 6.

These patterns are reflected in the profitability evidence. Well-capitalized banks earn 0.28 pp higher ROA ($p < 0.001$) despite a 5.95 pp lower net interest margin ($p < 0.001$). This profitability paradox is explained by the portfolio reallocation documented above: well-capitalized banks accept lower yields in exchange for substantially lower credit losses (charge-off rate −0.34 pp, $p < 0.001$), so the net effect on bottom-line profitability is positive. The higher ROA in turn accelerates equity accumulation through retained earnings, contributing a further 6.1% to the decomposition.

Finally, PCA classification shapes corporate strategy through direct statutory channels. Well-capitalized banks are 0.3 pp more likely to make acquisitions ($p = 0.017$) and 7.3 pp more likely to pay dividends ($p < 0.001$). Under FDICIA, banks below well-capitalized status require prior regulatory approval for acquisitions, branch openings, and new business activities, while dividend payments face escalating restrictions. The strategic option value embedded in well-capitalized status—the right to grow through acquisition without pre-clearance, to pay dividends without restriction, and to pursue new lines of business—is a component of the shadow cost that does not appear in balance sheet flow identities but imposes real costs on near-threshold banks who must forgo these opportunities.

5.3 Market Access and Service Provision

Beyond the balance sheet, PCA supervision reshapes the geographic scope and community reach of bank lending. Table 6 documents causal effects on branch networks, CRA geographic coverage, and SBA small business lending—outcomes that do not map through the equity flow identity but reflect the broader franchise costs of PCA status.

Well-capitalized banks maintain larger branch networks (+2.50 branches, $p = 0.004$) yet exhibit *lower* net branch change (-0.031, $p < 0.001$). This "network freeze" pattern is consistent with the regulatory mechanism: banks below the well-capitalized threshold face heightened scrutiny for *any* branch changes—both openings and closings—under FDICIA's prior-approval requirements. The restriction effectively freezes the geographic footprint, while well-capitalized banks freely optimize by closing underperforming locations. The net effect is that well-capitalized banks maintain larger networks through the option to manage them actively.

CRA geographic coverage shows even larger discontinuities: well-capitalized banks serve substantially more CRA assessment areas (+56 counties, $p < 0.001$), reflecting the broader geographic reach enabled by unrestricted branching and acquisition authority. In the SBA lending channel, well-capitalized banks are 7.3 pp more likely to participate in SBA lending

programs ($p < 0.001$), and conditional on participation, they originate loans with 0.30 pp lower interest rates and 9.51 pp lower default rates. The lower default rates are consistent with well-capitalized banks' ability to be more selective, while the interest rate effect may reflect both better credit quality and the competitive advantage of maintaining SBA-preferred-lender status—a designation that requires regulatory good standing.

These market access effects are economically distinct from the balance sheet channels. They capture the provision of banking services—small business credit, branch access, geographic coverage—that affect communities rather than the bank's own capital ratio. Their presence underscores that the shadow cost of PCA supervision extends beyond private bank costs to the real economy.

— Insert Table 6 Here —

5.4 Pre-PCA Placebo Tests

We estimate Equation (2) on the 1990–1992 sample for all level outcomes available in pre-PCA Call Report data.[8]

— Insert Table 7 Here —

Table 7 shows that none of the outcomes tested is significant at the 5% level. The pre-PCA coefficient on equity-to-assets is -0.007 ($p = 0.80$) versus $+5.13$ ($p < 0.001$) post-PCA; loans-to-assets: $+0.77$ ($p = 0.17$) versus -11.78 ($p < 0.001$). The extended 1984–1989 placebo is consistent: of 13 outcomes tested, one is significant at 5% (dividends/assets, $p = 0.046$)—within the expected false-positive rate. The absence of systematic pre-period discontinuities across both placebo windows supports a causal interpretation.

[8]Growth-rate outcomes require more than 12 quarters. Market access outcomes cannot be tested: FDIC Summary of Deposits begins in 1994, CRA disclosures in 1996, and SBA records are incomplete before 1993.

5.5 Robustness

We assess sensitivity across several dimensions, focusing on the most economically important threats to identification.

All core estimates are stable across bandwidth choices from 50% to 200% of the MSE-optimal bandwidth, robust to asymmetric bandwidth selection and local quadratic polynomials (Appendix Table A.5), and consistent when using total risk-based capital alone as the running variable (Appendix Table A.3).

The density discontinuity confirms that some banks manipulate their capital ratios, but this is precisely the bunching behavior we measure, not a threat to identification. Under the Gerard et al. (2020) partial identification approach with the economically plausible monotone assumption, key outcomes retain bounds that exclude zero: equity/assets ($[+4.11, +6.15]$), loans/assets ($[-18.9, -4.6]$), and NPL ratio ($[-1.33, -0.04]$); Appendix Table A.2 reports the full set. Competitive spillovers are addressed using the Berg et al. (2021) diagnostic (Appendix Table A.6): core treatment effects are unchanged, and any bias runs toward zero. Sub-period analysis confirms stability across pre-GFC, GFC, and post-Dodd-Frank eras, with the GFC equity/assets estimate an order of magnitude larger (+5.29 pp versus +0.34 pp pre-GFC), consistent with the threshold becoming far more consequential under capital pressure (Appendix Table A.4).

6 Decomposition and Dollar Valuation of the Shadow Cost

The bunching analysis delivers the *total* shadow cost in capital-ratio units ($\Delta q = 4.17$ pp); the RDD delivers *channel-specific* causal effects in each outcome's native units. The balance sheet identity $c = (E/A)/\omega$ provides the mapping: we convert each RDD estimate into capital-ratio-equivalent units, compare the sum to Δq, and attribute the remainder to an unobserved residual. This is an accounting decomposition following the sufficient-

statistics approach of Chetty (2009), not a structural estimation. The key innovation is that we *estimate* conversion parameters directly from the near-threshold sample and propagate uncertainty via parametric bootstrap.

The term "shadow cost" is used in three nested senses: (i) S in capital-ratio units (this section); (ii) the *dollar shadow cost* S_i, the bank-specific opportunity cost of holding excess equity (Section 6.5); and (iii) the *aggregate shadow cost*, the system-wide sum (Section 6.6).

6.1 The Shadow Cost of a Regulatory Notch

The bunching framework of Kleven and Waseem (2013) provides the theoretical link between observed behavior and the shadow cost. Consider a bank choosing its capital ratio c to maximize profits:

$$\max_c \ \pi(c) - S \cdot \mathbf{1}[c < \bar{c}] \tag{3}$$

where $\pi(c)$ is the profit function (decreasing in c because equity is costly relative to debt), $\bar{c}$ is the PCA threshold, and S is the shadow cost of reclassification. The regulatory notch induces banks with unconstrained optima in $[\bar{c} - \Delta q, \bar{c}]$ to bunch at $\bar{c}$. The indifference condition for the marginal buncher—located at $c^* = \bar{c} - \Delta q$—defines the shadow cost as the foregone profit from distorting the capital ratio:

$$S = \int_{\bar{c}-\Delta q}^{\bar{c}} \pi'(c)\, dc \tag{4}$$

Under a quadratic profit function, $S = \frac{\theta}{2}\Delta q^2$, where θ is the curvature.[9] Under the exponential specification of Alvero and Xiao (2023), $S = \theta(e^{\Delta q/100} - 1 - \Delta q/100)e^{\bar{c}/100}$. Both imply convexity in Δq. Crucially, the *bunching range* Δq is estimated from the CJM density discontinuity (Section 4), but the *profit curvature* θ is not—it cancels out of the decomposition shares, which depend only on the ratio of each channel's contribution to the total, and enters only the dollar valuation, where it is absorbed by the observed ROE (Section 6.5).

[9] $\pi'(c) = -\theta(c - c^*)$. Integrating from $\bar{c} - \Delta q$ to $\bar{c}$: $S = \frac{\theta}{2}\Delta q^2$.

Two caveats apply: S captures the *private* cost, not the social cost, and the interpretation requires that the profit function is well approximated near the threshold.

6.2 The Accounting Decomposition

The total distortion—from the density discontinuity estimate—is Δq percentage points of capital ratio. The question is: how much operates through each observable channel? The capital ratio identity $c = E/\text{RWA}$ implies:

$$\frac{\Delta c}{c} = \frac{\Delta E}{E} - \frac{\Delta \ell}{\ell} - \frac{\Delta \omega}{\omega} \tag{5}$$

where E is equity, ℓ is total assets, and ω is the average risk weight. Any change in the capital ratio must operate through one of these three margins: raising equity, shrinking assets, or reducing risk weights.

Each RDD estimate $\hat{\tau}_j$ captures the causal effect of PCA classification on a specific channel j, measured in the channel's native units (e.g., percentage points of asset growth). We convert each into capital-ratio-equivalent units Δc_j by evaluating the identity at near-threshold bank means:

$$\Delta c_j^{\text{asset}} = \bar{c} \times \frac{|\hat{\tau}_j|}{100} \times \frac{\bar{A}}{\overline{\text{RWA}}}, \qquad \Delta c_j^{\text{NPL}} = \bar{c} \times \frac{|\hat{\tau}_j|}{\bar{\omega}} \times \lambda_{RW}, \qquad \Delta c_j^{\text{equity}} = \bar{c} \times \frac{|\hat{\tau}_j|}{100} \tag{6}$$

We estimate the scaling parameters directly from the near-threshold subsample (± 2 pp, $n = 209{,}815$ bank-quarters) rather than calibrating them (Panel B of Table 8). The mean risk weight $\bar{\omega} = 0.728$ slightly exceeds the commonly assumed 0.70, and the deposit-to-RWA ratio $\bar{D}/\overline{\text{RWA}} = 1.18$ is 24% higher than the standard 0.95, reflecting community banks' deposit reliance.

The decomposition restricts to equity flow identity channels—asset growth, ROA, dividends, and equity issuance—which map through the numerator of $c = (E/A)/\omega$. Risk-weight changes from portfolio reallocation operate through the denominator and are conservatively

assigned to the residual.

The total bunching range is then decomposed:

$$\Delta q = \underbrace{\sum_{j \in \mathcal{S}} \Delta c_j}_{\text{Observable}} + \underbrace{\Delta q - \sum_{j \in \mathcal{S}} \Delta c_j}_{\text{Residual}} \tag{7}$$

where $\mathcal{S}$ is the set of channels significant at $p < 0.05$. The decomposition is additive by construction; the residual absorbs genuinely unobservable costs and any interaction effects between channels.

6.3 Decomposition Results

— Insert Table 8 Here —

Table 8 reports the results. The decomposition maps each RDD estimate *strictly one-to-one* against the four terms in the equity flow identity $\Delta(E/A) \approx \text{ROA} - \text{Div}/A + \text{Issuance}/A - (E/A) \times g_A$, converting each to capital-ratio-equivalent units using the estimated conversion parameters. To avoid double-counting, the table includes exactly one RDD estimate per identity term: asset growth (g_A) for the denominator channel, ROA (NI/A) for profitability, equity issuance as a fraction of assets for the issuance channel, and dividends as a fraction of assets for the payout channel.[10]

The dominant channel operates through constrained growth. Asset growth is 5.72 pp lower for well-capitalized banks ($p < 0.001$), contributing 0.81 pp to Δq (19.4%): slower asset growth reduces risk-weighted assets relative to equity, mechanically raising the capital ratio. Deposit growth (-6.45 pp, $p < 0.001$) confirms this channel from the liability side and

[10] Earlier iterations of this table summed redundant proxies for the same economic flows—e.g., both asset growth and deposit growth (which captures the liability side of the same balance sheet contraction), multiple ROA specifications, and both extensive- and intensive-margin measures of equity issuance. Because these alternative measures reflect the same underlying adjustment, summing them violates the accounting identity and overstates the observable share. The corrected decomposition restricts to one measure per identity term and reports the alternative specifications as corroborating evidence.

produces a similar capital-ratio contribution (+0.76 pp), but is not added to the identity sum because deposit contraction and asset contraction reflect the same underlying balance sheet shrinkage.

Profitability contributes a further 0.26 pp (6.1%): ROA is 0.18 pp higher for well-capitalized banks ($p < 0.001$), accelerating equity accumulation through retained earnings. An alternative ROA specification using lagged denominators yields a similar magnitude (+0.27 pp), confirming the robustness of this channel. Equity issuance as a fraction of assets contributes 0.05 pp (1.1%). The dividend channel is economically negligible (0.00 pp), though the probability of paying a dividend is significantly higher for well-capitalized banks (+0.073, $p < 0.001$), consistent with the statutory benefit of well-capitalized status. Bootstrap standard errors and confidence intervals are computed from 2,000 parametric draws that jointly resample both the RDD estimates and the conversion parameters, propagating all sources of uncertainty.[11]

Beyond the identity channels, Table 8 reports corroborating evidence. The brokered deposit prohibition contributes +0.052 pp ($\approx 1.2\%$ of Δq) but is not added to the identity sum to avoid overlap with asset growth. Market access channels (SBA, CRA) are statistically significant but contribute negligible capital-ratio-equivalent adjustments because these lending volumes are small relative to total assets; their importance is distributional, documenting that supervision costs extend to the communities banks serve. Strategic action channels (M&A) capture the option-value dimension of PCA status.

The identity-consistent observable channels account for 26.6% of Δq (SE = 2.3%, 95% CI: [22.4%, 31.4%]). A formal test overwhelmingly rejects the hypothesis that observable channels fully explain the distortion ($z = 14.96$, $p < 0.001$). The residual of 73.4% provides an upper bound on costs that do not map through the equity flow identity—examination

[11] The parametric bootstrap draws the RDD estimates ($\hat{\tau}$) and the cross-sectional conversion parameters (e.g., risk weights, asset-to-RWA ratios) from independent normal distributions. This independence assumption is reasonable because the conversion parameters are estimated cross-sectionally on the ±2 pp near-threshold sample, whereas the RDD parameters are local threshold effects estimated from a different source of identifying variation.

burden, reputational stigma, management constraints, and the lost option value of regulatory flexibility. As a specification check, performing the decomposition with stock-position rather than flow RDD estimates yields a lower observable share, as expected: flows map directly into the $\Delta(E/A)$ identity that governs capital ratio dynamics.

6.4 The Unobserved Residual

The residual—73.4% (95% CI: [68.6%, 77.6%])—is economically large. Because the decomposition is strictly additive, the residual absorbs genuinely unobservable "soft" costs, non-linear interaction effects between observable channels, and channels that affect the capital ratio through multiple simultaneous mechanisms. It therefore provides an *upper bound* on soft supervisory costs—examination burden, reputational stigma, the option value of regulatory flexibility, and management constraints imposed through discretionary actions.

This has a direct policy implication. Observable channels—dominated by the growth constraint—account for roughly one-quarter of the shadow cost and are directly targeted by threshold adjustments (e.g., reducing the CBLR from 9% to 8%). The remaining three-quarters operates through channels that threshold adjustments cannot address; proposals that constrain supervisory discretion—such as requiring MRAs be tied to material financial risks—target this residual directly.

6.5 Dollar Valuation

To translate the bunching-implied distortion into dollar values, we note that a bank at the margin holds excess equity of $\Delta E_i = (\Delta q/100) \times \text{RWA}_i$. The opportunity cost is:

$$S_i = r_e \times \frac{\Delta q}{100} \times \text{RWA}_{i,t-1} \tag{8}$$

where r_e is the cost of equity capital and RWA is lagged one quarter. Our baseline uses a flat 10% market cost of equity, consistent with the long-run U.S. banking average. RWA

is derived from $E/(c/100)$, where E is total equity and c is the total risk-based capital ratio. The only estimated parameter is Δq; everything else is observable. As a robustness check, using each bank's own lagged ROE (winsorized at $\pm 50\%$) yields similar aggregates (Appendix Table A.1).

We compute Equation (8) for 1,012,170 bank-quarter observations (16,621 unique institutions, 1993–2025).

6.6 Dollar Results: Size, Category, and Aggregation

— Insert Table 9 Here —

Table 9 reports the results by PCA category and bank size. The shadow cost is economically significant at all sizes: the smallest banks (<\$100M) bear a median annual cost of \$110,000, while banks above \$10 billion exceed \$77 million per year. The cost as a percentage of assets is *increasing* in bank size (0.24% for banks under \$100M vs. 0.32% for banks over \$10B), because larger banks tend to have lower capital ratios and hence larger RWA per dollar of equity. Undercapitalized banks face higher costs as a share of assets than well-capitalized banks, for the same reason.

Our primary aggregate estimate applies the bunching-implied $\Delta q = 4.17$ pp to near-threshold banks using a 10% market cost of equity, yielding \$79 billion per year.[12] This figure represents approximately 0.3% of the banking system's \$23 trillion in total assets. Excluding CCAR-subject banks using a time-varying size threshold (Section 6.8) yields a lower bound of \$48 billion attributable specifically to PCA supervision. For comparison, Kisin and Manela (2016) estimate the shadow cost of a one-percentage-point Tier 1 increase at \$220 million per year for participating banks; our estimate captures the full PCA regime and the entire banking system.

[12] Near-threshold is defined as ± 3 pp of the 10% well-capitalized threshold. With a uniform cost of equity, the distinction between the near-threshold and full-sample estimates is negligible (\$80 billion); the near-threshold figure is our preferred specification because it restricts attention to banks for whom the bunching interpretation is strongest.

Applying the identity-consistent decomposition shares to the $79 billion aggregate, observable channels that map one-to-one through the equity flow identity account for approximately $21 billion (26.6%) and the residual for $58 billion (73.4%).[13] The residual remains economically large and statistically significant ($z = 14.96$), providing an upper bound on the costs of supervision that do not map strictly through the accounting identity.

Figure 5 plots the annual time series. The aggregate cost grew from $51 billion in 1993 to $231 billion in 2025, driven by consolidation that concentrates risk-weighted assets on fewer, larger banks (Panel B: 13,593 banks to 2,861, with mean per-bank costs rising from $3,750K to $81,000K). The signed-ROE specification dips during the financial crisis (2008–2010) while the market-rate specification does not, revealing a natural countercyclicality: during crises, the opportunity cost of excess capital falls precisely when the probability of needing it rises.

6.7 Size Heterogeneity and Policy Counterfactuals

The concentration of shadow costs has immediate implications for reform design. Community banks (<$1 billion) account for 89% of institutions but a modest fraction of aggregate shadow costs; the proposed CBLR reduction from 9% to 8% would lower their shadow cost by roughly $1/4.17 \approx 24\%$, though the binding constraint for most community banks (median capital ratio 15.6%) is examination intensity, not the threshold itself. Regional banks ($1–$10 billion) face mean costs of $8 million per year and may benefit most from tightening MRA materiality standards. The largest banks (>$10 billion) bear a disproportionate share despite comprising fewer than 3% of institutions; threshold-based reforms have little effect on these banks, supervised under the full Basel framework, and the soft costs captured in the residual are likely largest here (see Correia et al., 2025, who find that stricter supervision reduces FDIC resolution costs by 9% of assets).

[13] Decomposition dollar amounts are computed by applying the estimated channel shares (26.6% observable, 73.4% residual) to the near-threshold aggregate shadow cost.

6.8 CCAR-Subject Banks and the Binding Constraint

The largest banks—subject to CCAR and DFAST—face a binding constraint that is typically the Stressed Capital Buffer (SCB), not the static PCA 10% threshold, so applying Δq uniformly may overstate the PCA-specific cost for these institutions. We address this with a time-varying size exclusion: no exclusion before 2013 (when CCAR did not exist), exclusion above \$50 billion from 2013–2018 (Dodd-Frank Section 165), and above \$100 billion from 2019 onward (EGRRCPA).[14] The CCAR-adjusted aggregate is \$48 billion per year, compared to the unadjusted \$79 billion. A static \$100 billion cutoff yields \$40 billion but understates PCA costs in the pre-CCAR era when PCA was the binding constraint for all banks. The CCAR-adjusted figure represents a lower bound on the aggregate shadow cost attributable specifically to PCA supervision.

6.9 Sensitivity Analysis

Figure 6 reports heatmaps of the bunching estimates across 20 specifications (5 polynomial orders $\times$ 4 fitting ranges). All 20 sharp specifications produce negative excess mass (-0.23 to -2.57), confirming that bunching is invariant to specification choice. Fuzzy Δq estimates range from 1.6 to 6.3 pp but are sensitive to polynomial order at orders ≥ 8; our preferred CJM local polynomial estimate of $\Delta q = 4.17$ pp (95% CI $[3.85, 4.53]$) avoids global smoothness assumptions. Substituting the confidence interval bounds yields a dollar cost range of \$73–\$86 billion.[15] The decomposition *shares* are invariant to the profit-function specification and to θ, because they depend only on each channel's ratio to the total bunching range; only the dollar *level* is affected.

[14] The binding stress-test capital constraint—the Stressed Capital Buffer—was specific to CCAR-subject institutions above \$50 billion. Using \$10 billion as the exclusion threshold would be too conservative.

[15] Scaling linearly: $79 \times 3.85/4.17 \approx 73$ and $79 \times 4.53/4.17 \approx 86$.

7 Conclusion

This paper asks how much it costs banks to be supervised, and answers using the behavior of banks themselves. By combining bunching estimation with a regression discontinuity design at Prompt Corrective Action capital thresholds, we estimate that banks distort their pivotal capital ratio by 4.17 percentage points to maintain well capitalized status, at a conservative aggregate cost of $79 billion per year ($48–64 billion excluding CCAR-subject banks). A sufficient-statistics decomposition with estimated conversion parameters and bootstrap inference attributes 27% of this cost to observable balance sheet channels that map one-to-one through the equity flow identity—dominated by constrained asset growth and equity retention—and shows that the remaining 73% provides an upper bound on the soft supervisory costs that current reforms target directly.

The findings have three implications for the ongoing restructuring of bank supervision. First, the brokered deposit result—a large increase in usage at the threshold, with a precisely estimated zero before PCA existed—establishes that PCA's mandatory restrictions are sharply binding. Banks are not simply holding capital buffers out of prudence; they are actively managing their balance sheets in response to specific regulatory consequences. This validates both the causal interpretation of our estimates and the policy premise that supervisory burden is real and measurable.

Second, the decomposition into observable and unobservable costs provides a roadmap for reform. Proposals to lower capital thresholds (the CBLR reduction) and modify risk weights (Basel III endgame) address the hard costs—the quarter we can measure. Proposals to constrain supervisory discretion (materiality standards for MRAs, the elimination of reputation risk, Supervisory Operating Principles tied to solvency threats) address the soft costs—the three-quarters we cannot. The sufficient-statistics decomposition makes clear which reforms target which margin: threshold changes and growth constraints address the observable $21 billion (27%), while supervisory discretion reforms target the unobservable $58 billion residual (73%).

Third, the concentration of the burden across the size distribution implies that the current reform agenda, which emphasizes community bank relief, addresses only a fraction of the total cost. The largest banks bear the highest shadow cost both in absolute terms and as a share of assets, yet threshold-based reforms have little effect on institutions supervised under the full Basel framework. For these banks, the binding constraint is examination intensity and supervisory discretion, not the capital ratio itself.

An important distinction separates private cost from social cost. Our shadow cost measures what banks are willing to pay to avoid reclassification—not the welfare loss from PCA supervision. To the extent that supervision corrects agency problems or reduces negative externalities from bank failure (Correia et al., 2025; Passalacqua et al., 2021; Bisetti, 2024), the social cost may be substantially lower, and some of the shadow cost may reflect the elimination of socially inefficient risk-taking rather than deadweight loss. The private cost represents one side of a cost-benefit calculation; the stability benefits—including reduced failure probability (Ivanov and Karolyi, 2023) and mitigated credit contractions (Baron et al., 2021)—remain to be quantified in a unified welfare framework.

One limitation deserves emphasis. The dollar formula treats Δq as excess equity held against RWA, but to the extent that banks achieve part of the capital ratio improvement through denominator management (slower growth reducing RWA) rather than numerator accumulation, the effective equity cost is lower while the lending distortion is larger. The decomposition quantifies this tradeoff: asset growth accounts for the largest single observable channel contribution.

The "bunching + RDD" methodology provides a template for evaluating regulatory notches in other settings. The approach requires a threshold that generates observable bunching, a rich set of outcome variables for the RDD, and an accounting identity that maps channel-specific effects into a common unit. These conditions are satisfied by the SIFI designation threshold ($250 billion), the stress testing threshold ($100 billion), and international analogs such as India's PCA framework (Kashyap et al., 2025). As the de-

bate over bank supervision continues, estimates grounded in revealed-preference behavior and disciplined by pre-period placebo tests offer a foundation for evidence-based regulatory design.

References

Alvero, A. and K. Xiao (2020): "Fuzzy bunching," *Working Paper, Columbia University.*

——— (2023): "The real effects of the Dodd-Frank SIFI threshold," *Working Paper.*

Baron, M., E. Verner, and W. Xiong (2021): "Banking Crises Without Panics," *Quarterly Journal of Economics*, 136, 51–113.

Begenau, J. (2020): "Capital Requirements, Risk Choice, and Liquidity Provision in a Business-Cycle Model," *Journal of Financial Economics*, 136, 355–378.

Berg, T., M. Reisinger, and D. Streitz (2021): "Spillover Effects in Empirical Corporate Finance," *Journal of Financial Economics*, 142, 1109–1127.

Berger, A. N., S. M. Davies, and M. J. Flannery (2000): "Comparing Market and Supervisory Assessments of Bank Performance: Who Knows What When?" *Journal of Money, Credit and Banking*, 32, 641–667, see also Berger et al. (2008) in *Oxford Handbook of Banking.*

Bisetti, E. (2024): "The Value of Regulators as Monitors: Evidence from Banking," *Management Science*, 70, 8464–8483.

Bonaccorsi di Patti, E., M. Moscatelli, and S. Pietrosanti (2023): "The Impact of Bank Regulation on the Cost of Credit: Evidence from a Discontinuity in Capital Requirements," *Journal of Financial Intermediation*, 55.

Calonico, S., M. D. Cattaneo, and M. H. Farrell (2020): "Optimal bandwidth choice for robust bias-corrected inference in regression discontinuity designs," *Econometrics Journal*, 23, 192–210.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014): "Robust nonparametric confidence intervals for regression-discontinuity designs," *Econometrica*, 82, 2295–2326.

Cattaneo, M. D., M. Jansson, and X. Ma (2020): "Simple local polynomial density estimators," *Journal of the American Statistical Association*, 115, 1449–1455.

Cattaneo, M. D., L. Keele, R. Titiunik, and G. Vazquez-Bare (2016): "Interpreting Regression Discontinuity Designs with Multiple Cutoffs," *Journal of Politics*, 78, 1229–1248.

Cattaneo, M. D., R. Titiunik, and R. Yu (2025): "Estimation and Inference in Boundary Discontinuity Designs," *Working Paper*.

Chetty, R. (2009): "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1, 451–488.

Choi, J. Y. and M.-J. Lee (2018): "Regression Discontinuity with Multiple Running Variables Allowing Partial Effects," *Political Analysis*, 26, 258–274.

Correia, S., S. Luck, and E. Verner (2025): "Supervising Failing Banks," Tech. rep., NBER Working Paper No. 34343.

——— (2026): "Failing Banks," *Quarterly Journal of Economics*, 141, 147–204.

Crippa, L. (2025): "Joint Manipulation Testing in Multidimensional Regression Discontinuity Designs," *Working Paper*.

Garcia, R. E. and S. Steele (2022): "Stress Testing and Bank Business Patterns: A Regression Discontinuity Study," *Journal of Banking & Finance*, 135.

Gerard, F., M. Rokkanen, and C. Rothe (2020): "Bounds on Treatment Effects in Regression Discontinuity Designs with a Manipulated Running Variable," *Quantitative Economics*, 11, 839–870.

Granja, J. and C. Leuz (2024): "The Death of a Regulator: Strict Supervision, Bank Lending, and Business Activity," *Journal of Financial Economics*, 158.

Ivanov, I. T. and S. A. Karolyi (2023): "Fighting Failure: The Persistent Real Effects of Resolving Distressed Banks," *Working Paper*.

Kashyap, N., S. Mahapatro, and P. L. Tantri (2025): "Preventing Defaults in Response to Deteriorating Bank Health: The Prompt Corrective Action Approach," *Contemporary Accounting Research*.

Kisin, R. and A. Manela (2016): "The Shadow Cost of Bank Capital Requirements," *Review of Financial Studies*, 29, 1780–1820.

Kleven, H. J. (2016): "Bunching," *Annual Review of Economics*, 8, 435–464.

Kleven, H. J. and M. Waseem (2013): "Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan," *Quarterly Journal of Economics*, 128, 669–723.

Passalacqua, A., P. Angelini, F. Lotti, and G. Soggia (2021): "The Real Effects of Bank Supervision: Evidence from On-Site Bank Inspections," Tech. rep., Bank of Italy Working Paper No. 1349.

Reardon, S. F. and J. P. Robinson (2012): "Regression Discontinuity Designs with Multiple Rating-Score Variables," *Journal of Research on Educational Effectiveness*, 5, 83–104.

Table 1: PCA Capital Thresholds and Supervisory Consequences

Panel A: FDICIA 1991 Thresholds (1993–2014)

	Total RBC	Tier 1 RBC	CET1	Leverage
Well Capitalized	$\geq 10\%$	$\geq 6\%$	—	$\geq 5\%$
Adequately Capitalized	$\geq 8\%$	$\geq 4\%$	—	$\geq 4\%$
Undercapitalized	$< 8\%$	$< 4\%$	—	$< 4\%$
Sig. Undercapitalized	$< 6\%$	$< 3\%$	—	$< 3\%$

Panel B: Basel III Thresholds (2015–present)

	Total RBC	Tier 1 RBC	CET1	Leverage
Well Capitalized	$\geq 10\%$	$\geq 8\%$	$\geq 6.5\%$	$\geq 5\%$
Adequately Capitalized	$\geq 8\%$	$\geq 6\%$	$\geq 4.5\%$	$\geq 4\%$
Undercapitalized	$< 8\%$	$< 6\%$	$< 4.5\%$	$< 4\%$
Sig. Undercapitalized	$< 6\%$	$< 4\%$	$< 3\%$	$< 3\%$

Notes: A bank must meet all thresholds simultaneously to qualify for a given category; falling below any single ratio triggers reclassification. Panel A shows FDICIA (1991) thresholds using three capital ratios. Panel B shows revised thresholds effective January 2015 under Basel III, which raised the Tier 1 RBC well-capitalized threshold from 6% to 8% and introduced Common Equity Tier 1 (CET1) as a fourth binding constraint. Total RBC and Leverage thresholds are unchanged. Mandatory consequences of losing Well Capitalized status include restrictions on brokered deposits (12 CFR 337.6), interest rate caps on deposits, and limitations on pass-through insurance. Discretionary consequences include enhanced examination frequency, restrictions on asset growth, branching, and acquisitions, mandatory capital restoration plans, and potential management changes.

Table 2: Summary Statistics

Variable	N	Mean	SD	P10	P50	P90
Panel A: Capital Ratios						
Total Risk-Based Capital (%)	1,203,239	21.65	36.18	10.97	15.56	29.70
Tier 1 Risk-Based Capital (%)	1,203,240	20.55	36.28	9.81	14.44	28.64
Common Equity Tier 1 (%)	210,257	19.43	40.23	—	14.16	26.68
Panel B: Bank Characteristics						
Asset Growth (%)	1,625,982	8.31	16.19	−4.14	5.22	21.71
Equity Growth (%)	1,620,619	7.54	18.57	−5.72	6.36	20.06
Deposit Growth (%)	1,614,022	8.42	17.78	−4.68	4.93	22.20
Dividend Payment (0/1)	855,813	0.46	0.50	0.0	0.0	1.0
NPL Ratio (%)	1,583,529	1.42	2.46	0.02	0.69	3.37

Notes: Bank-quarter level observations, 1993Q1–2025Q2. Growth rates are year-over-year, winsorized at the 1st and 99th percentiles. CET1 available from 2014 onward. NPL ratio is nonperforming loans as a percentage of total loans. PCA classification: Well Capitalized 1,131,595 (94.1%), Adequately Capitalized 27,772 (2.3%), Undercapitalized 7,497 (0.6%), Significantly Undercapitalized 36,376 (3.0%); see Figure 1 for the time series.

Table 3: Density Discontinuity at the Well Capitalized Threshold

Running Variable	T_{JK}	p-value	$\hat{f}^-(0)$	$\hat{f}^+(0)$	$\hat{\rho}$	Δq (pp)	N
Total RBC (10%)	42.48***	< 0.001	0.0212	0.0483	2.275	3.32 (0.12)	1,203,239
Tier 1 RBC (6%/8%)	0.04	0.966	0.0033	0.0033	1.003	—	1,203,240
CET1 (6.5%)	−0.40	0.687	0.0008	0.0006	0.847	—	210,257
Pivotal (binding score)	51.29***	< 0.001	0.0211	0.0486	2.306	4.17 (0.19)	1,203,251

Notes: Post-PCA period (1993–2025). Density discontinuity estimates follow Cattaneo et al. (2020) using local polynomial density estimation with jackknife standard errors via `rddensity`. T_{JK} is the jackknife test statistic for $H_0 : f(0^-) = f(0^+)$. $\hat{f}^-(0)$ and $\hat{f}^+(0)$ are the estimated densities just below and above the threshold; $\hat{\rho} = \hat{f}^+(0)/\hat{f}^-(0)$ is the density ratio. The bunching range $\Delta q = (\hat{f}^+(0) - \hat{f}^-(0)) \times h_R/\hat{f}^-(0)$, where h_R is the CJM-optimal bandwidth on the right side, measures the capital ratio distortion in percentage points under a locally uniform counterfactual density. All running variables use time-varying thresholds: the Tier 1 RBC well-capitalized cutoff is 6% pre-2015 and 8% post-Basel III; CET1 (6.5%) applies only post-2015. The pivotal running variable is the minimum distance to well-capitalized status across all applicable ratios. Pre-PCA placebo density tests for all running variables are reported in Appendix Table A.8. Bootstrap standard errors for Δq from 200 replications. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 4: RDD Estimates: Balance Sheet Positions at the Well Capitalized Threshold

Outcome	$\hat{\tau}_{BC}$	Robust SE	p-val.	BW
Panel A: Capital				
Equity / Assets (%)	+5.129***	(0.119)	< 0.001	0.18
Panel B: Asset Composition				
Loans / Assets (%)	−11.783***	(0.503)	< 0.001	0.23
Securities / Assets (%)	+8.731***	(0.414)	< 0.001	0.30
Cash / Assets (%)	+4.001***	(0.206)	< 0.001	0.26
Liquid Assets / Assets (%)	+13.263***	(0.490)	< 0.001	0.21
Panel C: Loan Portfolio				
CRE / Loans (%)	+0.045*	(0.023)	0.054	0.96
Residential RE / Loans (%)	+3.425***	(0.697)	< 0.001	0.65
C&I / Loans (%)	−3.693***	(0.438)	< 0.001	0.61
Consumer / Loans (%)	−4.692***	(0.424)	< 0.001	0.69
Panel D: Funding				
Deposits / Assets (%)	−3.091***	(0.273)	< 0.001	1.11
Brokered Deposits / Assets (%)	+0.365**	(0.162)	0.024	0.80
Core Deposits / Total Deposits (%)	+7.044***	(0.427)	< 0.001	0.28
Panel E: Asset Quality				
NPL Ratio (%)	−0.685***	(0.071)	< 0.001	1.05
Loan Loss Reserves / Loans (%)	−0.499***	(0.056)	< 0.001	1.16

Notes: Bias-corrected RDD estimates with CERT-clustered robust standard errors from `rdrobust` (Calonico et al., 2014, 2020). Running variable is the pivotal distance to the Well Capitalized threshold (minimum across Total RBC, Tier 1 RBC, and CET1). Full sample (∼1.15 million bank-quarters, 21,035 banks); no subsampling. MSE-optimal bandwidths. Asset composition and funding use total assets as denominator; loan portfolio shares use total loans; core deposit share uses total deposits. Brokered Deposits/Assets is denominated in assets because the statutory restriction (12 CFR 337.6) operates on absolute volumes; Core Deposits/Total Deposits captures within-deposit composition. Average risk weight $\omega = \text{RWA/Assets}$ is computed as (Equity/Assets)/(Total RBC ratio/100). Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 5: RDD Estimates: Flows and Strategic Decisions at the Well Capitalized Threshold

Outcome	$\hat{\tau}_{BC}$	Robust SE	p-val.	BW
Panel A: Equity Flows (intensive, per dollar of assets)				
ROA (%)	+0.279***	(0.027)	< 0.001	0.86
Dividends / Assets (%)	+0.040***	(0.012)	< 0.001	1.10
Equity Issuance / Assets (%)	−0.033**	(0.014)	0.024	0.92
Panel B: Equity Flows (extensive, 0/1 indicators)				
Dividend Payment (0/1)	+0.073***	(0.009)	< 0.001	1.10
Equity Issuance (0/1)	−0.021**	(0.010)	0.039	0.78
Panel C: Strategic Actions				
Acquisition Made (0/1)	+0.003**	(0.001)	0.017	0.96
Any M&A Activity (0/1)	+0.003**	(0.001)	0.014	0.93
Panel D: Balance Sheet Growth				
Asset Growth (%)	−5.721***	(0.534)	< 0.001	0.37
Deposit Growth (%)	−6.447***	(0.566)	< 0.001	0.38
Panel E: Operations				
Net Interest Margin (%)	−5.954***	(0.628)	< 0.001	0.53
Efficiency Ratio (%)	−1.560***	(0.531)	0.003	0.75

Notes: Bias-corrected RDD estimates with CERT-clustered robust standard errors from `rdrobust`. Running variable is the pivotal distance to the Well Capitalized threshold. Full sample (∼1.15 million bank-quarters, 21,035 banks); no subsampling. Panel A reports equity flow measures (per dollar of assets) that enter the $\Delta(E/A)$ decomposition (Equation 5). Panel B reports the corresponding extensive margins (0/1 indicators). The sign pattern—positive dividend *payment* but negative dividend/assets, and negative issuance *propensity* but positive issuance/assets—reflects composition and supervisory pressure: well-capitalized banks pay dividends more often but in smaller amounts per dollar of assets, while banks below the threshold face supervisory recapitalization pressure that raises total equity issuance volume per dollar of assets. Panel C reports binary indicators for corporate decisions constrained by PCA's strategic expansion restrictions. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6: RDD Estimates: Market Access and Service Provision at the Well Capitalized Threshold

Outcome	$\hat{\tau}_{BC}$	Robust SE	p-val.	BW
Panel A: Branch Network (FDIC Summary of Deposits, annual)				
Branch Count	+2.503***	(0.867)	0.004	0.86
Counties Served	+0.751***	(0.187)	< 0.001	0.91
Net Branch Change (qtrly)	−0.031***	(0.008)	< 0.001	1.51
Panel B: CRA Geographic Coverage and Lending (annual, large institutions)				
Counties with CRA Lending	+56.01***	(10.46)	< 0.001	0.65
Net CRA Counties Entered	+9.12***	(2.33)	< 0.001	0.97
CRA Small Business Lending ($M)	+269.7***	(77.8)	< 0.001	1.01
Rural Lending Share (%)	+0.873***	(0.187)	< 0.001	0.62
Outside-Assessment-Area Share (%)	+1.301***	(0.164)	< 0.001	0.70
Panel C: SBA Small Business Lending				
Any SBA Lending (0/1)	+0.073***	(0.006)	< 0.001	1.02
SBA Loan Count (per quarter)	+0.835***	(0.213)	< 0.001	0.89
SBA Average Interest Rate (%)	−0.299**	(0.118)	0.011	0.76
SBA Default Rate (%)	−9.511***	(2.096)	< 0.001	1.10
SBA Guarantee Share (%)	−7.985***	(0.891)	< 0.001	0.78

Notes: Bias-corrected RDD estimates with CERT-clustered robust standard errors from `rdrobust`. Running variable is the pivotal distance to the Well Capitalized threshold. Panel A uses annual FDIC Summary of Deposits data; Branch Count and Counties Served reflect year-end geographic footprint, while Net Branch Change is measured quarterly. Panel B uses annual CRA small business lending disclosures (required for institutions with assets $\geq$ \$376M at the relevant year-end threshold); CRA Small Business Lending is total originations of small business loans (commercial loans $\leq$\$1M), in millions of dollars. "Outside-Assessment-Area Share" captures the fraction of CRA lending conducted outside a bank's formally designated geographic assessment area, reflecting geographic reach beyond the core market. Panel C uses SBA 7(a) loan data collapsed to bank-quarter totals; "SBA Guarantee Share" is the government-guaranteed percentage of outstanding SBA loans, which is inversely related to borrower credit quality—lower guarantee shares indicate lenders perceive lower credit risk. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7: Pre-PCA Placebo: RDD Estimates at the Well Capitalized Threshold (1990–1992)

Outcome	$\hat{\tau}_{BC}$	Robust SE	p-val.	BW
Panel A: Capital				
Equity / Assets (%)	−0.007	(0.029)	0.799	1.08
Panel B: Asset Composition				
Loans / Assets (%)	+0.774	(0.564)	0.170	0.94
Cash / Assets (%)	−0.275	(0.159)	0.083	1.21
Liquid Assets / Assets (%)	−0.328	(0.507)	0.517	1.03
Securities / Assets (%)	−0.063	(0.500)	0.900	1.09
Panel C: Loan Portfolio				
CRE / Loans (%)	−0.028	(0.058)	0.629	1.10
Residential RE / Loans (%)	+0.465	(0.936)	0.619	0.94
C&I / Loans (%)	−0.417	(0.517)	0.420	0.85
Consumer / Loans (%)	+0.016	(0.516)	0.976	0.90
Panel D: Funding				
Deposits / Assets (%)	−0.116	(0.266)	0.663	0.75
Brokered Deposits / Assets (%)	+0.086	(0.047)	0.069	0.77
Core Deposits / Total Deposits (%)	+0.204	(0.272)	0.454	1.14
Panel E: Asset Quality				
NPL / Loans (%)	−0.001	(0.019)	0.954	1.24
Loan Loss Reserves / Loans (%)	+0.016	(0.023)	0.492	1.02
Net Chargeoffs / Loans (%)	−0.001	(0.019)	0.954	1.24
Panel F: Risk Weights				
Avg Risk Weight (ω)	+0.005	(0.004)	0.255	1.29
Panel G: Equity Flows & Operations				
ROA (%)	−0.007	(0.020)	0.718	0.72
Dividends / Assets (%)	+0.002	(0.010)	0.824	0.96
Dividend Payment (0/1)	−0.004	(0.017)	0.799	0.85
Equity Issuance (0/1)	+0.021	(0.015)	0.166	0.61
Net Interest Margin (%)	+0.004	(0.042)	0.921	0.99

Notes: Pre-PCA placebo RDD estimates (1990–1992, $N \approx 176{,}000$ bank-quarters). Bias-corrected estimates with CERT-clustered robust standard errors from `rdrobust`, matching the specification in Tables 4–5. Panels A–G report all level outcomes available in the 1990–1992 Call Report data; growth-rate outcomes are excluded because 12 quarters provides insufficient variation. Market access outcomes (Table 6) cannot be tested pre-PCA: FDIC Summary of Deposits data begins in 1994, CRA disclosures begin in 1996, and SBA loan records have incomplete pre-1993 coverage. None of the outcomes in Panels A–G is significant at 5%; only brokered deposits/assets is marginally significant ($p = 0.069$), attributable to pre-existing FIRREA restrictions from 1989. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 8: Shadow Cost Decomposition: $\Delta c \approx \text{ROA} - \text{Div}/A + \text{Issuance}/A - c \times g_A$

Panel A: Identity-Consistent Decomposition

Identity Term	$\hat{\tau}$ (RDD)	Δc (pp)	Boot. SE	95% CI	Share of Δq
$c \times g_A$: Asset growth	-5.721^{***}	$+0.807$	(0.076)	[0.654, 0.951]	19.4%
ROA (NI/A)	$+0.181^{***}$	$+0.255$	(0.032)	[0.192, 0.319]	6.1%
Issuance / Assets	-0.033^{**}	$+0.046$	(0.020)	[0.006, 0.085]	1.1%
Dividends / Assets	$+0.040^{***}$	$+0.000$	(0.000)	[0.000, 0.000]	0.0%
Sum observable		**+1.109**	**(0.083)**	**[0.946, 1.272]**	**26.6%**
Unobserved residual		+3.059	(0.205)	[2.666, 3.481]	73.4%
Total Δq		**4.168**	**(0.185)**		**100.0%**
H_0: residual = 0		$z = 14.96$, $p < 0.001$			

Panel B: Estimated Conversion Parameters

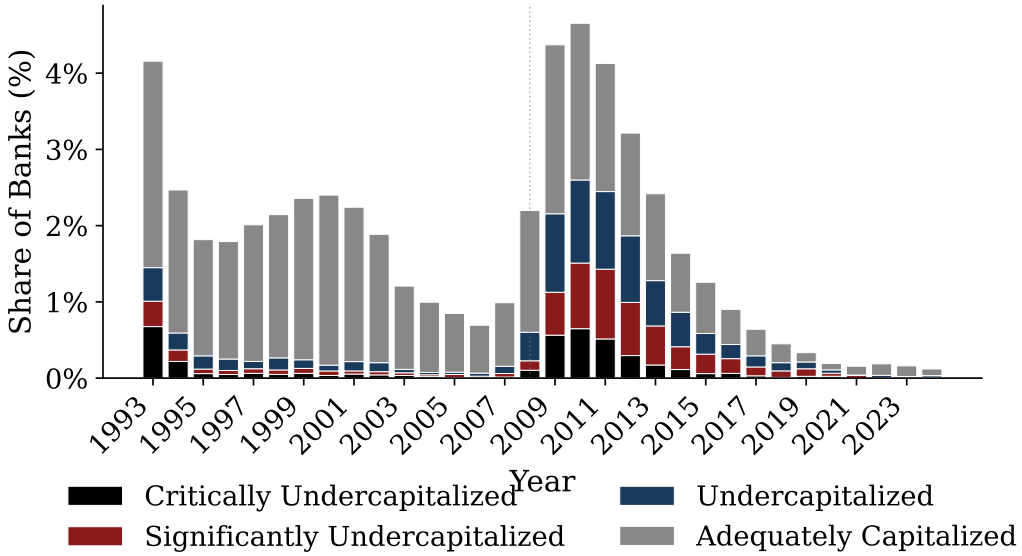
Parameter	Estimate	(SE)	Source
$\bar{c}$ (threshold ratio)	0.1000	–	Regulatory
$\bar{\omega}$ (mean risk weight)	0.7282	(0.0003)	Sample mean
$\bar{A}/\overline{\text{RWA}}$ (assets to risk-weighted)	1.4113	(0.0006)	$1/\bar{\omega}$
$\bar{D}/\overline{\text{RWA}}$ (deposits to risk-weighted)	1.1836	(0.0006)	Sample mean
Near-threshold sample ($\pm$2pp)	209,815 bank-quarters		

Notes: Panel A decomposes the bunching-implied distortion $\Delta q = 4.168$ pp into capital-ratio-equivalent contributions via the equity flow identity $\Delta(E/A) \approx \text{ROA} - \text{Div}/A + \text{Issuance}/A - (E/A) \times g_A$, using exactly one RDD estimate per identity term (no double-counting). Asset growth operates through the denominator (slower growth shrinks RWA, raising the ratio); ROA, issuance, and dividends operate through the numerator (equity accumulation). RDD estimates from bias-corrected local linear estimation (pivotal running variable, full sample, CERT-clustered SEs). The 26.6% observable share reflects identity-consistent channels only; the 73.4% residual captures supervisory costs that do not flow through the equity accumulation identity—including examination burdens, the option value of unrestricted strategic expansion, and any misalignment between accounting equity and market-value capital. Panel B reports conversion parameters from the near-threshold sample ($\pm$2 pp). Bootstrap SEs and CIs from 2,000 parametric draws jointly resampling RDD estimates and conversion parameters. Stars: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

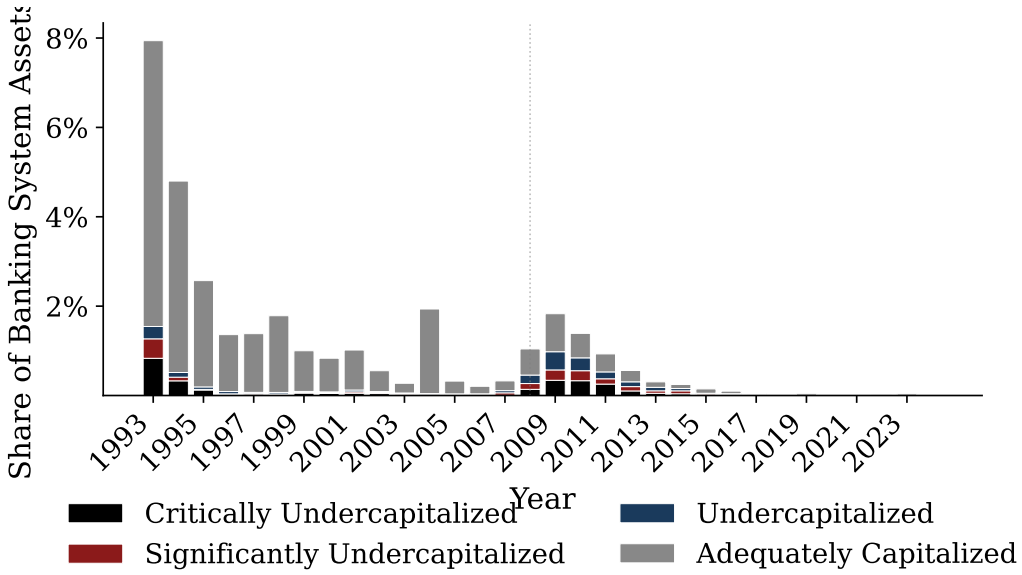
Table 9: Dollar Shadow Costs by PCA Category and Bank Size

	Banks	Mean ROE (%)	Shadow Cost ($/yr) Mean	Median	% of Assets
Panel A: By PCA Category					
Well Capitalized	16,191	9.6	$5,009K	$286K	0.26
Adequately Capitalized	3,266	0.4	$1,384K	$344K	0.29
Undercapitalized	955	−23.6	$997K	$410K	0.27
Sig. Undercapitalized	3,306	−10.5	$1,105K	$420K	0.26
Panel B: By Bank Size					
< $100M	11,987	8.5	$122K	$110K	0.24
$100M–$500M	10,084	9.7	$590K	$496K	0.26
$500M–$1B	3,553	9.9	$1,933K	$1,840K	0.28
$1B–$10B	2,381	10.5	$8,306K	$5,517K	0.29
> $10B	400	11.3	$273,376K	$76,692K	0.32
Near-threshold (primary)	**10,596**				**$79B/yr**
All banks	16,275				$80B/yr
CCAR-adjusted (time-varying)					$48B/yr
CCAR-adjusted (static $100B)					$40B/yr

Notes: Shadow costs computed from Equation (8) using $\Delta q = 4.17$ pp from the CJM density discontinuity estimator and a 10% market cost of equity. RWA derived from equity / (capital ratio / 100). All dollar amounts in thousands except aggregates. "Banks" is the number of unique institutions with valid observations. Aggregate figures are time-series averages of annual cross-sectional totals (1993–2025). "Near-threshold" restricts to banks within ±3 pp of the 10% well-capitalized threshold. "CCAR-adjusted (time-varying)" excludes banks above the applicable stress-testing threshold: no exclusion pre-2013, >$50B for 2013–2018 (Dodd-Frank CCAR), and >$100B for 2019+ (post-EGRRCPA). "CCAR-adjusted (static)" applies the $100B cutoff throughout for comparability.



Panel A: Share of Banks Below Well Capitalized



Panel B: Share of Banking System Assets Below Well Capitalized

Figure 1: PCA Category Distribution Over Time (Excluding Well Capitalized)

Notes: **Panel A**: Fraction of banks classified below Well Capitalized in each year. **Panel B**: Fraction of banking system assets in each below-well-capitalized category. Stacked bars from bottom: Critically Undercapitalized (darkest), Significantly Undercapitalized, Undercapitalized, and Adequately Capitalized (lightest). Well Capitalized banks (≈94% of observations) are omitted to reveal variation in the lower categories. The vertical dashed line marks 2008. Data: FDIC Call Reports, 1993–2024.

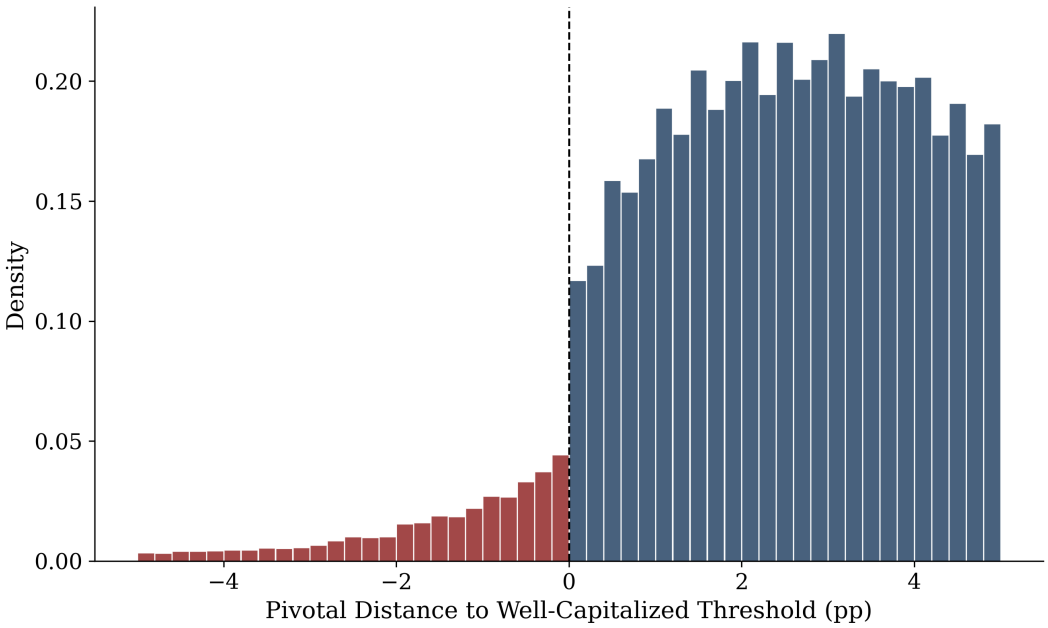


Figure 2: Distribution of Pivotal Distance to Well Capitalized Threshold

Notes: Distribution of the minimum distance across all four capital ratios (Total RBC, Tier 1, CET1, Leverage) to their respective Well Capitalized thresholds. The pivotal distance captures the binding constraint: the ratio closest to its threshold. Bin width = 0.2 percentage points. The dashed line marks the Well Capitalized threshold (distance = 0). The visible excess mass just above zero indicates strategic bunching.

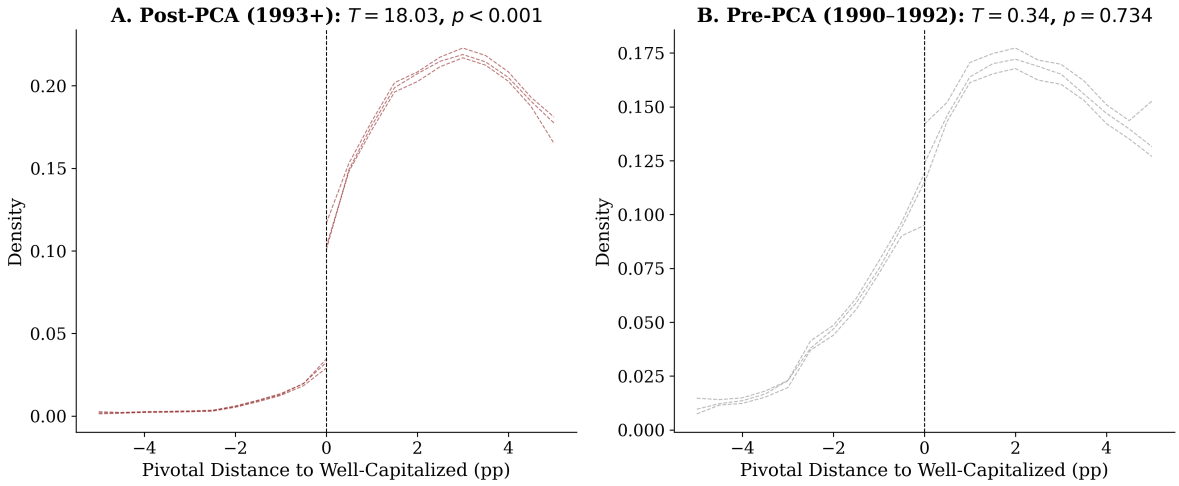


Figure 3: Density Discontinuity Test at the Well Capitalized Threshold

Notes: Cattaneo, Jansson, and Ma (2020) density discontinuity test at the Well Capitalized threshold. Panel A shows the post-PCA period (1993–2025): the test strongly rejects continuity ($p < 0.001$), confirming significant bunching above the threshold. Panel B shows the pre-PCA placebo period (1990–1992): the test fails to reject continuity, indicating no bunching before PCA was enacted. Histograms show binned frequencies; solid lines are local polynomial density estimates from `rdplotdensity`, fit separately on each side of the threshold using data-driven bandwidths. Dashed lines show 95% confidence bands. T-statistics and p-values are shown in each panel title.

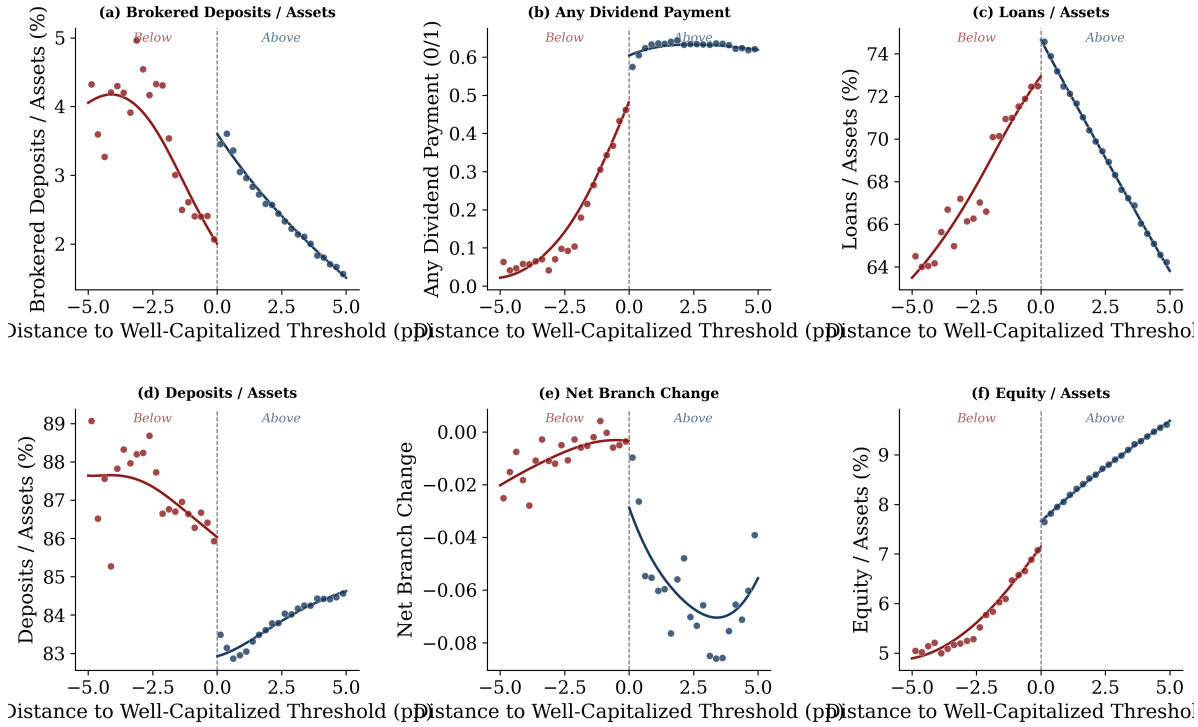
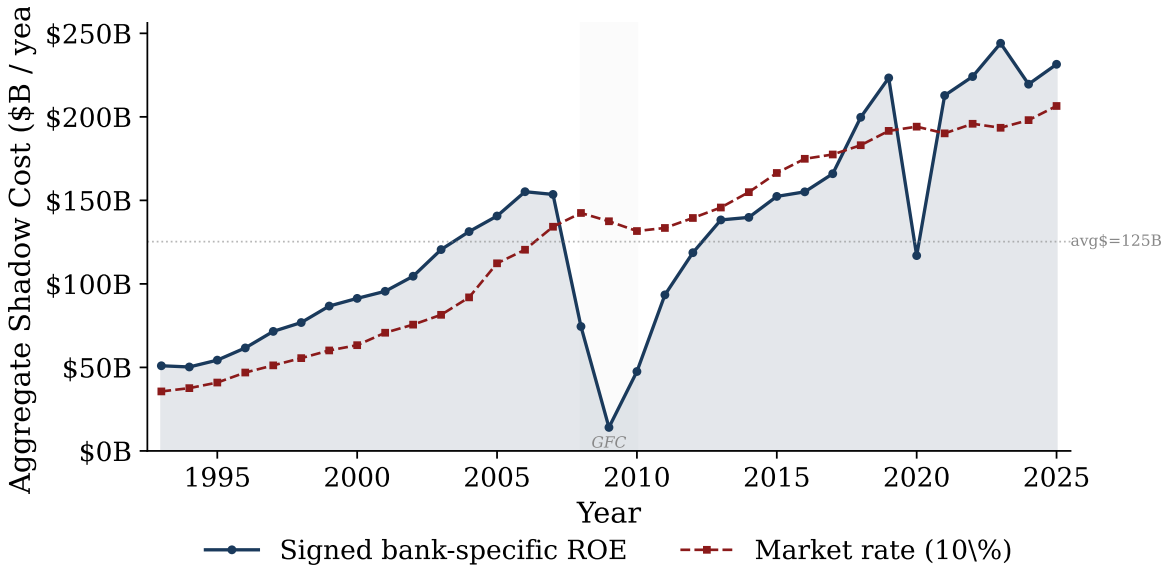
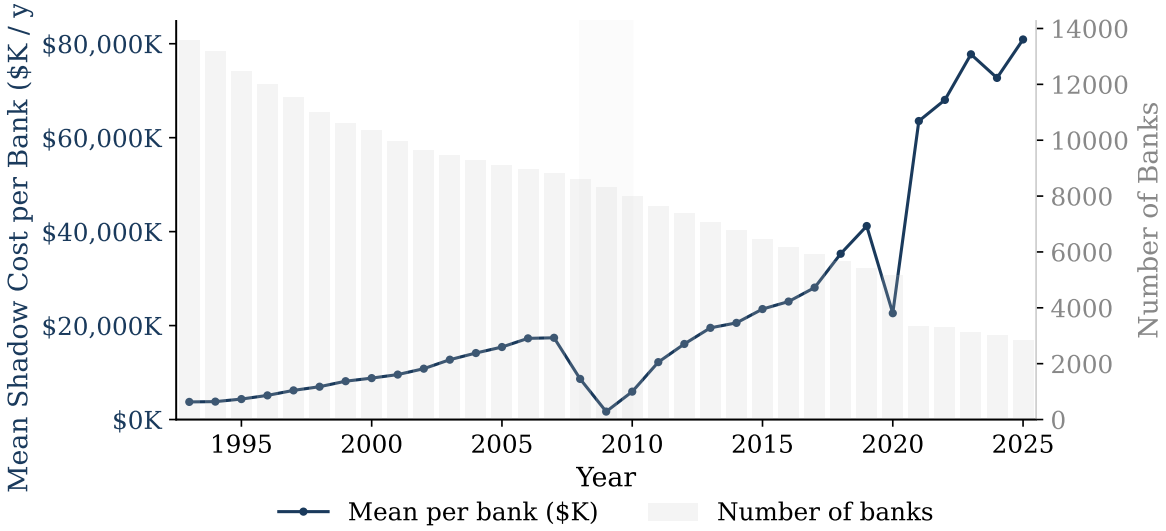


Figure 4: RDD Discontinuity Plots: Key Bank Outcomes at Well Capitalized Threshold

Notes: Each panel plots binned means of the outcome variable against the pivotal distance to the Well Capitalized threshold. The dashed vertical line marks the threshold (distance = 0). Local polynomial fits are estimated separately on each side of the cutoff. Outcomes represent key supervisory channels: (a) brokered deposits (mandatory restriction under 12 CFR 337.6), (b) dividend payments, (c) loan share of assets, (d) deposit funding share, (e) branch network changes, and (f) equity capitalization.



Panel A: Aggregate Annual Shadow Cost



Panel B: Per-Bank Shadow Cost and Industry Consolidation

Figure 5: Annual Shadow Cost of PCA Supervision, 1993–2025

Notes: **Panel A**: Aggregate annual shadow cost of PCA supervision under two specifications: signed bank-specific ROE (dark blue solid) and a flat 10% market cost of equity (dark red dashed), both using $\Delta q = 4.17$ pp applied to near-threshold banks. The dotted horizontal line marks the time-series average of the market-rate specification. The sharp dip in 2008–2010 reflects negative ROE during the financial crisis; the market-rate specification is insensitive to profitability cycles. **Panel B**: Mean per-bank shadow cost (left axis, $K/year) and number of active banks (right axis, gray bars). Industry consolidation reduces the number of banks while concentrating the aggregate shadow cost on fewer, larger institutions. The paper's primary estimate of $79 billion/year (Table 9) uses the 10% market rate, near-threshold banks.

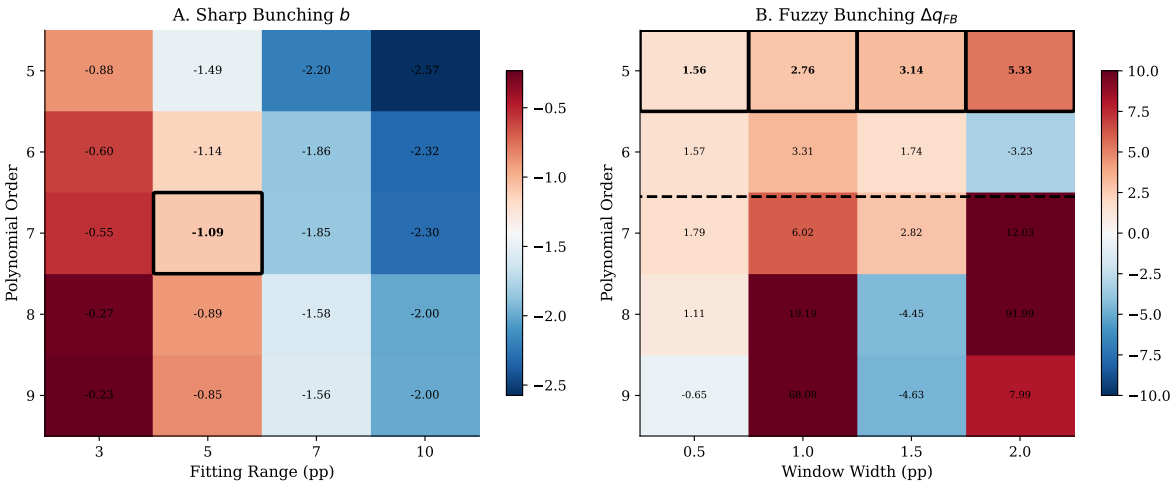


Figure 6: Sensitivity of Bunching Estimates to Specification Choice

Notes: Each cell reports the bunching estimate for a given combination of polynomial order (rows) and fitting range or window width (columns), using total risk-based capital at the 10% well capitalized threshold. Panel A: Sharp normalised excess mass b (negative = missing mass below threshold). All 20 specifications produce negative values, confirming strategic bunching regardless of specification choice. The bold-bordered cell marks the baseline (poly = 7, fit range = 5 pp). Panel B: Fuzzy bunching range Δq in percentage points. The dashed box marks the stable specification region (poly = 5–6); the bold-bordered cell marks the baseline (poly = 5). Color scale clipped to $[-10, 10]$ to preserve contrast; values outside this range reflect overfitting of the counterfactual density at high polynomial orders.

Appendix A Robustness: Market Cost of Equity

Our baseline dollar shadow cost uses a flat 10% market cost of equity as the opportunity cost of holding excess capital. This choice avoids the interpretive difficulties that arise with bank-specific ROE—particularly for negative-ROE banks where the formula is mechanically strained—while remaining close to the cross-sectional mean ROE of 9.6%. We assess robustness by comparing the baseline against absolute-value ROE (treating magnitude as a proxy for financial distress intensity). If the aggregate shadow cost is similar across specifications, the key findings are not driven by the cost-of-equity assumption.

The results confirm robustness. The baseline 10% market rate yields $80 billion per year (all banks) and the absolute-value ROE specification yields $89 billion—a range of only 11%. These comparisons support the paper's primary estimate of $79 billion (near-threshold banks, 10% market cost of equity). The time-varying CCAR-adjusted estimate of $48 billion (Section 6.8) provides a lower bound attributable specifically to PCA supervision, while the static $100 billion cutoff yields $40 billion.

Table A.1: Shadow Cost Robustness: Bank-Specific ROE vs. Fixed Market Cost of Equity

	Market Rate (10%)	ROE-Based (|ROE|)
Panel A: All Banks		
N banks	16,275	16,275
Total ($B/yr)	80	89
Mean per bank ($K/yr)	4,930	5,481
Panel B: Near-Threshold (±3 pp)		
N banks	10,596	10,596
Total ($B/yr)	**79**	96
Mean per bank ($K/yr)	7,450	9,065
Panel C: CCAR-Adjusted (excl. >$100B banks)		
N banks (near-thresh.)	10,539	10,539
Total ($B/yr)	**40**	—

Notes: Both columns use the formula $S_i = \text{COE} \times (\Delta q/100) \times \text{RWA}_{i,t-1}$ with $\Delta q = 4.17$ pp from the CJM density discontinuity estimator. "Market Rate (10%)" uses a flat 10% cost of equity for all banks (baseline); "ROE-Based (|ROE|)" uses the absolute value of each bank's lagged annualized ROE (winsorized at ±50%). Panel B restricts to the near-threshold sample (±3 pp of the 10% well-capitalized threshold). Panel C further excludes CCAR-subject banks (>$100B in assets) for whom the Stressed Capital Buffer, not the PCA threshold, is binding. The Panel B Market Rate figure ($79B) is the paper's primary aggregate estimate. Aggregates are time-series averages of annual cross-sectional totals (1993–2025).

Appendix B Robustness: Manipulation Bounds

The density discontinuity at the well capitalized threshold confirms that some banks manipulate their capital ratios—but this manipulation is precisely the bunching behavior we measure, not a threat to identification per se. The concern is whether banks that sort above the threshold (manipulators) have systematically different treatment effects from those that do not. Gerard et al. (2020) provide a framework for bounding the RDD treatment effect under manipulation. We apply their approach using the friction parameter $\alpha = 0.47$ from the fuzzy bunching model as the manipulation fraction: approximately 47% of near-threshold banks actively manage their capital ratios.

We compute two sets of bounds for the 14 core outcomes that span Tables 4–5 and include all decomposition channels. Worst-case bounds allow manipulators' treatment effects to take any value in the outcome's support (1st–99th percentile range); these are intentionally conservative and serve as an outer envelope. Monotone bounds impose the economically plausible restriction that manipulators' effects are bounded by the interquartile range of the non-manipulators' outcome distribution. The worst-case bounds are wide for all outcomes, as expected given the large manipulation fraction, but the monotone bounds are substantially tighter. Under the monotone assumption, several key outcomes have monotone bounds that exclude zero: equity/assets ($[+4.11, +6.15]$), loans/assets ($[-18.92, -4.64]$), cash/assets ($[+2.11, +5.89]$), liquid assets ($[+6.24, +20.28]$), and NPL ratio ($[-1.33, -0.04]$). The pattern is consistent: the core supervisory channels operate in the same direction regardless of the bounding assumption. These results support the causal interpretation of the main RDD estimates.

Table A.2: RDD Estimates with Manipulation Bounds (Gerard, Rokkanen, and Rothe 2020)

Outcome	Point Est.	Worst LB	Worst UB	Mono. LB	Mono. UB
Capital & Asset Composition					
Equity / Assets (%)	+5.129	−0.296	+10.55	+4.111	+6.147
Loans / Assets (%)	−11.78	−38.58	+15.02	−18.92	−4.644
Cash / Assets (%)	+4.001	−7.713	+15.72	+2.113	+5.889
Liquid Assets / Assets (%)	+13.26	−12.84	+39.36	+6.243	+20.28
CRE / Loans (%)	+0.045	−2.680	+2.769	+0.045	+0.045
Funding					
Deposits / Assets (%)	−3.091	−21.88	+15.69	−7.317	+1.134
Brokered Deposits / Assets (%)	+0.365	−12.50	+13.23	−0.630	+1.361
Core Deposits / Deposits (%)	+7.044	−17.60	+31.69	+1.345	+12.74
Quality & Risk					
NPL Ratio (%)	−0.685	−6.024	+4.654	−1.329	−0.041
Flows & Operations					
ROA (%)	+0.181	−1.781	+2.143	−0.110	+0.472
Dividend Payment (0/1)	+0.073	−0.397	+0.543	−0.397	+0.543
Equity Issuance (0/1)	+0.078	−0.392	+0.548	−0.392	+0.548
Asset Growth (%)	−5.721	−53.20	+41.75	−11.52	+0.079
NIM (%)	−5.954	−71.38	+59.47	−13.06	+1.152

Notes: Point estimates from the full-sample RDD specification (Tables 4–5). Manipulation fraction $\alpha = 0.47$ from the fuzzy bunching estimator. Worst-case bounds allow manipulators' treatment effects to take any value in the outcome's support (1st–99th percentile range); monotone bounds assume manipulators' effects are bounded by the interquartile range of the outcome distribution. The 15 core outcomes span every panel of Tables 4–5 and include all decomposition channels.

Appendix C Robustness: Total Risk-Based Capital as Running Variable

Our baseline specification uses the pivotal running variable—the minimum distance across all four capital ratios to their respective thresholds—to capture the "weakest link" feature of PCA classification. An alternative is to use only the total risk-based capital ratio, which is the binding constraint for the majority of near-threshold bank-quarters and the ratio that shows the strongest density discontinuity. Comparing the two specifications tests whether the pivotal construction, which mixes different margins of capital management, produces results that differ from the single-ratio design.

The comparison reveals broad consistency on the key outcomes. Equity/assets, deposits, dividends, and ROA are significant under both specifications with the same sign. The pivotal design produces substantially larger magnitudes for equity/assets (+5.13 pp vs. +0.42 pp) and loans/assets (−11.78 pp vs. +2.34 pp), reflecting the narrow bandwidth near the sharp bunching point in the pivotal distribution. These differences are economically interpretable: the pivotal design concentrates on banks at the exact binding threshold, where behavioral responses are most intense, while the total RBC design uses a wider bandwidth that averages across a broader range of banks. The qualitative conclusions on dividends and ROA are unchanged across specifications.

Table A.3: RDD Estimates: Pivotal vs. Total Risk-Based Capital Running Variable

Outcome	Pivotal (Baseline) $\hat{\tau}$	p-value	Total RBC Only $\hat{\tau}$	p-value
Capital & Asset Composition				
Equity / Assets (%)	+5.129	<0.001	+0.420	<0.001
Loans / Assets (%)	−11.78	<0.001	+2.336	<0.001
Cash / Assets (%)	+4.001	<0.001	+0.741	<0.001
Liquid Assets / Assets (%)	+13.26	<0.001	+2.981	<0.001
CRE / Loans (%)	+0.045	0.054	+0.087	0.002
Funding				
Deposits / Assets (%)	−3.091	<0.001	−2.135	<0.001
Brokered Dep. / Assets (%)	+0.365	0.024	+0.151	0.041
Core Dep. / Deposits (%)	+7.044	<0.001	+1.823	<0.001
Quality & Risk				
NPL Ratio (%)	−0.685	<0.001	−0.312	<0.001
Avg Risk Weight	−0.069	<0.001	−0.028	<0.001
Flows & Operations				
ROA (%)	+0.181	<0.001	+0.083	0.001
Dividend Payment (0/1)	+0.073	<0.001	+0.071	<0.001
Asset Growth (%)	−5.721	<0.001	−2.318	<0.001
NIM (%)	−5.954	<0.001	−1.447	<0.001

Notes: Both specifications use `rdrobust` with MSE-optimal bandwidth, triangular kernel, and CERT-clustered SEs on the full sample (∼1.15M bank-quarters). "Pivotal" uses the minimum distance across Total RBC, Tier 1, and CET1 with time-varying thresholds; "Total RBC Only" uses the Total Risk-Based Capital ratio distance to 10%. The 15 outcomes match those in Tables A.5–A.4. The pivotal specification produces larger magnitudes because the sharp bunching at the pivotal threshold yields a narrow optimal bandwidth (∼0.18 pp), concentrating on banks very close to the binding constraint. The total RBC specification uses a wider bandwidth and averages over a broader neighborhood.

Appendix D Robustness: Sub-Period Analysis

The PCA framework has operated through three distinct regulatory regimes: the pre-crisis period (1993–2007), when supervision was comparatively relaxed; the Global Financial Crisis and its aftermath (2008–2014), when supervisory intensity peaked; and the post-Dodd-Frank period (2015–2025), when stress testing and CCAR supplanted many PCA-specific constraints for the largest banks. If the shadow cost reflects PCA supervision specifically rather than aggregate macroeconomic conditions, we should see effects in all three periods, potentially with variation in magnitude reflecting the intensity of the supervisory environment.

Table A.4 reports sub-period estimates for all 15 core outcomes across three regulatory regimes. The results reveal large and economically interpretable variation across regimes. The equity/assets discontinuity is present in all three periods—+0.34 pp pre-GFC, +5.29 pp during the GFC, and +0.87 pp post-Dodd-Frank—but the GFC estimate is an order of magnitude larger, consistent with the extreme capital pressures of the crisis making the well-capitalized threshold far more consequential.

The most striking crisis-period patterns appear in portfolio composition. Loans/assets reverses sign: +2.14 pp pre-GFC but −11.23 pp during the GFC ($p < 0.001$), indicating that well-capitalized banks dramatically reduced loan exposure during the crisis while constrained banks could not restructure. Liquid assets show a corresponding +13.89 pp GFC effect ($p < 0.001$, versus −1.50 pp pre-GFC), and cash holdings shift from −0.39 pp pre-GFC to +1.82 pp during the GFC ($p < 0.001$): well-capitalized banks built liquidity buffers while constrained banks could not.

The funding channels also reverse during the GFC. Brokered deposits shift from +1.20 pp pre-GFC to −1.77 pp during the crisis ($p < 0.001$), while core deposits shift from −2.19 pp to +6.30 pp ($p < 0.001$). This is consistent with a flight-to-quality mechanism: well-capitalized banks attracted core depositors during the crisis and reduced reliance on volatile brokered funding, while banks near the PCA threshold experienced depositor withdrawals and turned

to brokered deposits. Post-Dodd-Frank, all funding channels lose significance ($p > 0.10$), consistent with enhanced liquidity regulation (LCR, NSFR) reducing the marginal effect of PCA-specific funding restrictions.

The dividend discontinuity is largest during the GFC ($+0.332$, $p < 0.001$), when PCA's dividend restrictions were most binding. The pre-GFC estimate is also significant ($+0.054$, $p = 0.004$), while the post-Dodd-Frank estimate is negligible (-0.015, $p = 0.874$), consistent with CCAR/DFAST stress testing supplanting PCA's dividend constraint. Asset growth effects are significant across all three periods, consistent with growth constraints being a persistent channel through which banks maintain capital ratios.

Table A.4: RDD Estimates by Regulatory Regime

Outcome	Pre-GFC (1993–2007) $\hat{\tau}$	p	GFC (2008–2014) $\hat{\tau}$	p	Post-DF (2015–2025) $\hat{\tau}$	p
Capital & Asset Composition						
Equity / Assets (%)	+0.343***	0.000	+5.291***	0.000	+0.874*	0.097
Loans / Assets (%)	+2.139***	0.000	−11.23***	0.000	+5.570**	0.012
Cash / Assets (%)	−0.387***	0.008	+1.821***	0.001	−1.083	0.392
Liquid Assets / Assets (%)	−1.497***	0.000	+13.89***	0.000	−3.614	0.127
CRE / Loans (%)	+0.080**	0.011	−0.019	0.739	−0.084	0.484
Funding						
Deposits / Assets (%)	−2.050***	0.000	−3.292***	0.000	−1.197	0.233
Brokered Dep. / Assets (%)	+1.201***	0.000	−1.767***	0.000	+1.191	0.135
Core Dep. / Deposits (%)	−2.188***	0.000	+6.303***	0.000	−1.800	0.171
Quality & Risk						
NPL Ratio (%)	−0.304***	0.000	−4.261***	0.000	−0.949	0.243
Avg Risk Weight	+0.037***	0.000	−0.082***	0.000	+0.050	0.148
Flows & Operations						
ROA (%)	+0.069***	0.003	+0.935***	0.000	+0.286**	0.014
Dividend (0/1)	+0.054***	0.004	+0.332***	0.000	−0.015	0.874
Equity Issuance (0/1)	+0.057***	0.000	−0.150***	0.000	+0.057	0.424
Asset Growth (%)	+1.544**	0.042	+2.968***	0.007	+4.438*	0.052
NIM (%)	+0.862	0.396	+4.322***	0.000	+6.949**	0.036
N	284,864		109,969		73,654	
Banks	11,668		8,300		3,842	

Notes: Full-sample `rdrobust` RDD estimates at the Well Capitalized threshold with MSE-optimal bandwidth, triangular kernel, and CERT-clustered SEs, split by regulatory regime. The 15 core outcomes span every panel of Tables 4–5 and include all decomposition channels. N is the total bank-quarter observations available in each period; effective sample sizes vary by outcome and bandwidth. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix E Robustness: Bandwidth Selection and Polynomial Order

Our baseline RDD specification uses a symmetric MSE-optimal bandwidth with a local linear polynomial ($p = 1$). Because the data density is highly asymmetric—many more banks above the well capitalized threshold than below—the symmetric bandwidth may not be optimal. We assess robustness by comparing four specifications: (1) symmetric bandwidth with linear polynomial (baseline), (2) asymmetric bandwidths optimized separately for each side with linear polynomial, (3) symmetric bandwidth with quadratic polynomial ($p = 2$), and (4) asymmetric bandwidths with quadratic polynomial.

Table A.5 reports the results for the 15 core outcomes. Signs are uniformly stable across all four specifications for every outcome. Significance is likewise robust: outcomes that are significant in the baseline remain significant under alternative bandwidths and polynomial orders. Point estimates vary modestly across specifications, as expected, but the economic magnitudes are consistent. The asymmetric bandwidth specification tends to produce slightly larger effective samples on the above-threshold side, reflecting the data density asymmetry, but this does not materially alter the estimates. The quadratic polynomial produces marginally wider confidence intervals, consistent with the bias-variance tradeoff, but no qualitative conclusions change. These results confirm that our findings are not sensitive to the bandwidth selection algorithm or polynomial order.

Table A.5: RDD Robustness: Bandwidth Selection and Polynomial Order

	Sym. BW, $p=1$ (Baseline)	Asym. BW, $p=1$	Sym. BW, $p=2$	Asym. BW, $p=2$
Capital & Asset Composition				
Equity / Assets (%)	$+5.129^{***}$	$+4.977^{***}$	$+5.065^{***}$	$+4.938^{***}$
Loans / Assets (%)	-11.783^{***}	-12.411^{***}	-11.802^{***}	-12.088^{***}
Cash / Assets (%)	$+4.001^{***}$	$+4.008^{***}$	$+4.049^{***}$	$+4.050^{***}$
Liquid Assets / Assets (%)	$+13.263^{***}$	$+13.946^{***}$	$+13.194^{***}$	$+13.679^{***}$
CRE / Loans (%)	$+0.045^{*}$	$+0.053^{**}$	$+0.038$	$+0.038$
Funding				
Deposits / Assets (%)	-3.091^{***}	-3.069^{***}	-3.132^{***}	-3.138^{***}
Brokered Dep. / Assets (%)	$+0.365^{**}$	$+0.238$	$+0.334^{*}$	$+0.281$
Core Dep. / Deposits (%)	$+7.044^{***}$	$+7.347^{***}$	$+7.077^{***}$	$+7.157^{***}$
Quality & Risk				
NPL Ratio (%)	-0.685^{***}	-0.681^{***}	-0.691^{***}	-0.685^{***}
Avg Risk Weight	-0.069^{***}	-0.082^{***}	-0.067^{***}	-0.080^{***}
Flows & Operations				
ROA (%)	$+0.181^{***}$	$+0.188^{***}$	$+0.192^{***}$	$+0.191^{***}$
Dividend Payment (0/1)	$+0.073^{***}$	$+0.074^{***}$	$+0.072^{***}$	$+0.072^{***}$
Equity Issuance (0/1)	$+0.078^{***}$	$+0.078^{***}$	$+0.081^{***}$	$+0.081^{***}$
Asset Growth (%)	-5.721^{***}	-6.777^{***}	-5.656^{***}	-6.616^{***}
NIM (%)	-5.954^{***}	-6.959^{***}	-6.172^{***}	-6.902^{***}

Notes: Bias-corrected RDD estimates ($\hat{\tau}_{BC}$) with CERT-clustered robust standard errors from `rdrobust`. "Sym. BW" uses `bwselect='mserd'` (single MSE-optimal bandwidth for both sides); "Asym. BW" uses `bwselect='msetwo'` (separate bandwidths optimized for each side of the threshold). $p = 1$ is local linear; $p = 2$ is local quadratic. Full sample (∼468K bank-quarters for most outcomes). The 15 core outcomes span every panel of Tables 4–5 and include all decomposition channels. Baseline column reproduces estimates from those tables. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix F Competitive Spillover Diagnostic

The stable unit treatment value assumption (SUTVA) requires that one bank's PCA classification does not affect other banks' outcomes. In concentrated local banking markets, however, a bank that loses well-capitalized status and contracts lending could create opportunities for rival banks, biasing our estimates. We address this concern using three complementary specifications that augment the RDD with measures of treatment intensity, following Berg et al. (2021).

The first specification (Panel B) adds the system-level treatment fraction Frac_t—the national proportion of banks below well-capitalized in quarter t—and its interaction with the bank-level treatment indicator. This captures time-varying exposure to aggregate supervisory pressure. The second specification (Panel C) uses bank-specific branch-weighted geographic exposure $\text{Frac}_{it}^{\text{branch}}$: for each county-year, we compute the fraction of banks operating below the well-capitalized threshold; bank i's exposure is the weighted average of these county fractions, with weights equal to the number of branches bank i operates in each county. The third specification (Panel D) uses deposit-weighted geographic exposure $\text{Frac}_{it}^{\text{dep}}$, replacing branch counts with branch-level deposits from the FDIC Summary of Deposits as weights. The deposit-weighted measure captures the economic magnitude of competitive overlap more precisely than simple branch counts.

All three specifications yield consistent findings: the core treatment effects ($\hat{\tau}$) remain highly significant with identical signs across all four outcomes and all panels. The deposit-weighted specification (Panel D) produces negligible magnitude changes: +0.3% for equity/assets, +1.4% for loans/assets, and +4.6% for ROA. The dividend estimate shows a somewhat larger shift (+14.9%), but remains highly significant ($p < 0.001$). These results confirm that the findings are robust to the most economically meaningful measure of local competitive exposure. The interaction terms ($\hat{\gamma}$) are often statistically significant but economically interpretable as correlated local conditions rather than direct spillovers: counties with more distressed banks tend to be experiencing local economic weakness that affects all

banks, not just those near the threshold.

The theoretical direction of spillover bias is outcome-specific. For credit supply outcomes (loans/assets, asset growth, branch changes), competitive reallocation could *inflate* the estimated gap: if a below-threshold bank contracts lending, a well-capitalized rival in the same market gains loan demand and expands, pushing Y_C up and making $\hat{\tau} = Y_T - Y_C$ larger in absolute value than the partial-equilibrium effect. For capital ratios (equity/assets) and profitability (ROA), the bias direction is ambiguous—competitive reallocation does not mechanically raise control banks' capital ratios. For regulatory channels (brokered deposits), the effect is statutory rather than competitive, so SUTVA concerns are minimal. The empirical evidence in Panels B–D is the decisive test: the core $\hat{\tau}$ estimates are stable across all three measures of local treatment intensity, with the most economically meaningful specification (deposit-weighted, Panel D) producing negligible magnitude changes. This stability implies that, regardless of the theoretical bias direction, spillovers do not materially affect our estimates in practice.

Table A.6: Spillover Diagnostic: Three Specifications

	Equity / Assets (%)	Loans / Assets (%)	Dividend (0/1)	ROA (%)
Panel A: Standard RDD				
$\hat{\tau}$	+2.080***	−3.170***	+0.136***	+0.220***
	(0.049)	(0.239)	(0.011)	(0.017)
Panel B: + System-Level Fraction				
$\hat{\tau}$	+2.034***	−4.059***	+0.103***	+0.124***
	(0.053)	(0.286)	(0.013)	(0.027)
$\hat{\beta}$ (System frac.)	−2.416***	−25.728***	−1.039***	−3.678***
	(0.703)	(5.773)	(0.208)	(0.827)
$\hat{\gamma}$ ($D_i\times$ Sys. frac.)	+1.457**	+33.349***	+1.210***	+3.419***
	(0.700)	(5.705)	(0.209)	(0.827)
$\Delta\hat{\tau}$ (%)	−2.2	−28.0	−24.1	−43.8
Panel C: + Branch-Weighted Exposure				
$\hat{\tau}$	+2.043***	−3.177***	+0.147***	+0.222***
	(0.052)	(0.285)	(0.012)	(0.020)
$\hat{\beta}$ (Branch exp.)	−1.904***	−0.826	−0.298***	−0.619***
	(0.243)	(2.656)	(0.088)	(0.176)
$\hat{\gamma}$ ($D_i\times$ Branch exp.)	−3.126***	−1.765	−1.308***	−1.654***
	(0.330)	(2.714)	(0.094)	(0.176)
$\Delta\hat{\tau}$ (%)	−1.8	−0.2	+8.4	+0.7
Panel D: + Deposit-Weighted Exposure				
$\hat{\tau}$	+2.085***	−3.127***	+0.156***	+0.230***
	(0.051)	(0.275)	(0.011)	(0.020)
$\hat{\beta}$ (Deposit exp.)	−1.312***	+0.022	−0.150*	−0.539***
	(0.228)	(2.635)	(0.088)	(0.179)
$\hat{\gamma}$ ($D_i\times$ Deposit exp.)	−3.440***	−2.087	−1.369***	−1.809***
	(0.349)	(2.706)	(0.098)	(0.182)
$\Delta\hat{\tau}$ (%)	+0.3	+1.4	+14.9	+4.6
N	468,487	468,487	468,487	468,487
Clusters	14,091	14,091	14,091	14,091

Notes: Panel A reports standard local linear RDD estimates within ±5 pp of the Well Capitalized threshold with triangular kernel and CERT-clustered standard errors. Panel B augments with Frac_t, the national fraction of banks below well-capitalized in quarter t. Panel C uses bank-specific branch-weighted geographic exposure: $\text{Frac}_{it}^{\text{branch}}$ is the weighted average of county-level below-threshold fractions, with weights equal to the number of branches bank i operates in each county (from FDIC Summary of Deposits). Panel D uses deposit-weighted exposure: $\text{Frac}_{it}^{\text{dep}}$ uses branch-level deposits as weights instead of branch counts. $\Delta\hat{\tau}$ is the percentage change in the treatment effect relative to Panel A. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix G Density Discontinuity Test: Adequately Capitalized Threshold

The well capitalized threshold is the primary focus because it generates the strongest density discontinuity and carries the most operationally significant supervisory consequences. For completeness, Table A.7 applies the same CJM density discontinuity test (Table 3) to the adequately capitalized threshold (Total RBC $\geq$ 8%, Tier 1 RBC $\geq$ 6%, CET1 $\geq$ 4.5%, Leverage $\geq$ 4%). The adequately capitalized threshold is the lower boundary of regulatory compliance: banks that fall below it are classified as undercapitalized and face the most severe PCA restrictions.

Table A.7: Density Discontinuity Test at the Adequately Capitalized Threshold

Running Variable	T_{JK}	p-value	$\hat{f}^{-}(0)$	$\hat{f}^{+}(0)$	$\hat{\rho}$	N
Total RBC ($\geq 8\%$)	-0.47	0.637	0.0083	0.0079	0.951	1,203,239
Tier 1 RBC ($\geq 6\%$)	-0.88	0.378	0.0073	0.0066	0.905	1,203,240
CET1 ($\geq 4.5\%$)	0.27	0.784	0.0009	0.0010	1.092	210,257
Pivotal (min-distance)	-0.61	0.541	0.0078	0.0072	0.924	1,203,251

Notes: Cattaneo, Jansson, and Ma (2020) jackknife density discontinuity test statistic T_{JK} and associated p-values at the adequately capitalized threshold. Estimated densities immediately below and above the 8%/6%/4.5% thresholds; $\hat{\rho} = \hat{f}^{+}(0)/\hat{f}^{-}(0)$ is the density ratio. None of the four running variables shows a significant density discontinuity at this threshold ($p > 0.37$ in all cases), in sharp contrast to the well-capitalized threshold (Table 3). This is consistent with banks spending little time at the adequately capitalized margin: institutions that lose well-capitalized status typically either recapitalize quickly or continue deteriorating toward undercapitalization, generating a sparse and smooth density near the 8% boundary. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix H Pre-PCA Density Discontinuity Tests: All Running Variables

Table A.8 reports CJM density discontinuity test statistics for the pivotal running variable and each of its three components across two pre-PCA placebo windows. The 1990–1992 window covers the period just before PCA enactment when all four capital ratios were

measurable. The 1984–1992 window extends further back and uses simulated risk-based capital ratios constructed from Call Report balance sheet components using Basel I risk weights (see note *a*). For comparison, the post-PCA statistics are reproduced from Table 3. The absence of significant density discontinuities in either pre-PCA window, for any running variable, confirms that strategic bunching at the well-capitalized threshold is a post-PCA phenomenon driven by the supervisory consequences introduced by FDICIA.

Table A.8: Pre-PCA Density Discontinuity Placebo Tests: All Running Variables

	Post-PCA (1993–2025)		Pre-PCA Placebo (1990–1992)		Extended Pre-PCA Placebo (1984–1992)	
Running Variable	T_{jk}	p-value	T_{jk}	p-value	T_{jk}	p-value
Pivotal (min-distance)	51.29***	< 0.001	-0.87	0.385	-0.22[a]	0.826
Total RBC ($\geq 10\%$)	42.48***	< 0.001	-1.14	0.255	-0.54[a]	0.590
Tier 1 RBC ($\geq 6\%$)	0.04	0.966	1.66	0.097	1.57[a]	0.116
CET1 ($\geq 6.5\%$)[b]	0.30	0.765	—		—	

Notes: Cattaneo, Jansson, and Ma (2020) jackknife density discontinuity test statistic T_{jk} and associated p-values. Running variable is each capital ratio's distance to its Well Capitalized threshold (Pivotal = minimum across all applicable ratios). Post-PCA statistics reproduced from Table 3; sample sizes 1,203,239–1,203,251. Pre-PCA (1990–1992): 176,143–176,266 observations. Extended Pre-PCA (1984–1992): 459,752–506,853 observations depending on running variable. Stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

[a] Risk-based capital ratios for 1984–1992 include both the final FDICIA year (1992) and the prior Basel I simulation period. For 1984–1989, ratios are constructed using Basel I risk weights applied to Call Report balance sheet components ($W_{\text{loans}} = 1.0$, $W_{\text{securities}} = 0.20$, $W_{\text{cash}} = 0.0$, $W_{\text{other}} = 0.50$); these represent simulated thresholds that were not legally binding until December 1992. For 1990–1992, risk-based capital ratios are directly reported. The pivotal distance for 1984–1989 uses the same simulation.

[b] CET1 was introduced under Basel III effective January 2015 (6.5% well-capitalized threshold). CET1 is not available for any pre-PCA period and cannot be tested in either placebo window.

Appendix I Fuzzy Bunching Estimates: Well Capitalized Threshold

Table A.9 reports fuzzy bunching estimates following Alvero and Xiao (2020) for all four running variables at the well-capitalized threshold during the post-PCA period. The fuzzy estimator relaxes the sharp estimator's assumption that all banks with unconstrained optima in $[\bar{c}-\Delta q, \bar{c}]$ relocate exactly to the threshold; instead it estimates the non-responding fraction

α jointly with Δq. We use the fuzzy estimator rather than the sharp estimator because sharp estimates for CET1 and Tier 1 RBC exhibit estimation difficulties (implausible values, near-singular systems) that are absorbed by the fuzzy friction parameter. Note that the paper's primary Δq estimate uses the CJM density discontinuity framework (Table 3); the fuzzy estimates here provide a structural decomposition of active versus frictions-driven bunching that complements the reduced-form density test.

Table A.9: Fuzzy Bunching Structural Decomposition: Well Capitalized Threshold (1993–2025)

Running Variable	Δq (pp)	SE	α	$1-\alpha$	Poly	Bandwidth	N
Post-PCA Period (1993–2025)							
Pivotal (min-distance)	6.30	(0.50)	0.533	0.467	5	5.0	785,025
Total RBC ($\geq 10\%$)	3.71	(0.13)	0.411	0.589	5	5.0	773,598
Tier 1 RBC ($\geq 6\%$)	4.67	(0.74)	1.000	0.000	5	5.0	559,809
CET1 ($\geq 6.5\%$)[a]	—	—	—	—	5	5.0	84,289

Notes: Fuzzy bunching estimates follow Alvero and Xiao (2020) with polynomial order 5 and a 5 pp estimation window on each side of the cutoff. Δq is the bunching range in percentage points; α is the friction parameter (fraction of non-responding banks); $1-\alpha$ is the fraction of active bunchers. SE from 50 bootstrap replications. The primary estimate for the paper ($\Delta q = 4.17$ pp) is from the CJM density discontinuity estimator (Table 3), which has better small-sample properties for the pivotal running variable; fuzzy bunching estimates provide corroborating evidence on the structural interpretation.

For Tier 1 RBC, $\alpha = 1.0$ indicates that all observed excess mass above the threshold is attributed to frictions (e.g., round-number heaping, reporting conventions), with zero structural active bunchers—consistent with the CJM density test result ($T_{jk} = 0.04$, $p = 0.97$), which fails to reject continuity. This reflects that the Tier 1 RBC threshold rarely binds independently given the Total RBC threshold dominates.

Pre-PCA placebo (1990–1992) CJM density tests are reported in Table A.8; fuzzy bunching estimates for the pre-PCA window are not presented because the CJM tests confirm no significant excess mass to estimate. The adequately capitalized threshold bunching estimates are reported in Table A.7.

[a] CET1 ($\geq 6.5\%$) bunching estimates are unreliable due to the limited post-2014 sample ($N = 84{,}289$, roughly one-fifteenth of the full sample). The sharp estimator yields $\Delta q = -9.3$ pp (dispersal) and the fuzzy estimator yields $\Delta q = 69.3$ pp with SE $> 10^6$, reflecting near-singular density estimation at this sample size. The CJM density test ($T_{jk} = 0.30$, $p = 0.77$) confirms no significant bunching at the CET1 threshold; this is consistent with Total RBC remaining the binding constraint and CET1 carrying independent regulatory significance primarily for large, complex institutions outside the near-threshold community bank sample.