# Rate Channeling Through Deposit Menus

*Uniform Anchors, Menu Rotation, and Monetary Transmission*

Prithu Vatsa

Department of Economics and Finance

University of Texas at El Paso

pvatsa@utep.edu

March 2026

**Submission to the NBER Summer Institute 2026**

Prithu Vatsa

Department of Economics and Finance

University of Texas at El Paso

pvatsa@utep.edu

## Disclosure of Conflicts of Interest

The author has no conflicts of interest to disclose.

## Workshop Preference and Justification

This paper is submitted to two workshops:

**1. Monetary Economics** (first preference)

This paper studies how market concentration shapes the transmission of monetary-policy rate changes through the deposit menu. The core question—where in the liability structure policy-rate pass-through appears, and why it varies with local market structure—is a monetary transmission question. The findings on incomplete and heterogeneous pass-through, asymmetric responses over the policy cycle, and the role of the anticipated component of policy speak directly to the Monetary Economics workshop.

**2. Macro, Money and Financial Frictions** (second preference)

The paper also contributes to the literature on financial frictions and banking. The

mechanism operates through heterogeneous depositor frictions (search, switching, attention) interacting with multi-product pricing by banks. The results on deposit recomposition, net-interest-margin effects, and the role of the local competitive frontier connect to macro-finance themes on how bank market power amplifies or attenuates monetary transmission through the funding side of bank balance sheets.

## Abstract

A large literature shows that deposit-rate pass-through is incomplete and weaker in concentrated markets, while a parallel literature shows that large banks often post nearly uniform deposit rates across geographies. This paper argues that both facts survive once deposit pricing is treated as a menu of contracts rather than a single rate. Using proprietary deposit-rate data from a platform banks use for competitive benchmarking, I show that market concentration shapes monetary transmission through rotation of the deposit menu rather than uniform attenuation of a representative rate. Vanilla rates serve as inert anchors, but concentration reshapes the rest of the menu: competitive offers move less, screening premia widen, and relative prices within the same bank diverge across markets. The strongest evidence is that concentration weakens the local competitive frontier rather than simply inducing larger unilateral markdowns by the focal bank. These pricing distortions have real consequences for liability composition and profitability during tightening cycles.



**JEL Classification:** E52, G21, L13, D43

**Keywords:** Deposit pricing, monetary policy transmission, market concentration, deposit menus, pass-through, market power

# 1 Introduction

A large body of literature documents incomplete deposit-rate pass-through and shows that the shortfall is larger in more concentrated markets.[1] This result has become central to modern accounts of monetary transmission through bank funding, the valuation of deposit franchises, and the competitive distortions associated with market power.[2] At the same time, a parallel body of evidence finds that large banks often post nearly uniform deposit rates across their geographic networks.[3] The two facts are routinely read as being in tension. This paper argues that banks do not set a deposit rate; they set a deposit menu—and once that is recognized, uniform anchors and concentration-dependent pass-through are not competing facts but different margins of the same pricing problem.

In equilibrium, banks solve a screening and funding-allocation problem, not a one-price problem. They post rates across contracts that differ in liquidity, maturity, balance thresholds, eligibility, and comparability, and they use that menu to allocate pass-through across depositors with different elasticities. Households, in turn, do not choose "a deposit rate." They allocate balances across the menu subject to attention, search, switching, and liquidity frictions that vary across products, depositor types, and local markets.[4]

Once pricing is multi-product and frictions are heterogeneous, local market structure need not operate through the level of a representative posted rate. It can operate through where pass-through appears in the menu: in liquid rates, in competitive offers, and in the wedges that sort depositors into different instruments. The transmission

[1] See the seminal work of Drechsler et al. (2017) and earlier contributions by Hannan and Berger (1991), Neumark and Sharpe (1992), and Berger and Hannan (1989). More recent structural work includes Wang et al. (2022) and Enkhbold (2023).

[2] Hoffmann et al. (2019), Drechsler et al. (2021), and d'Avernas et al. (2023).

[3] Radecki (1998), Begenau and Stafford (2023), and Granja and Paixão (2026).

[4] Hortaçsu and Syverson (2004) and Egan et al. (2017).

question is therefore not just how much pass-through occurs, but where in the menu it occurs. I call this reallocation of pass-through across contracts *rate channeling*: market structure channels monetary-policy adjustment into specific products, wedges, and depositor segments rather than transmitting it uniformly through a representative posted rate. *Menu rotation* is its within-menu manifestation. Anchor contracts directed at high-friction depositors can remain centrally set and locally inert, consistent with geographic uniformity. For those depositors, the anchor is the economically relevant price because it governs the no-action outcome: for CDs, it is the rate at which balances auto-renew if the depositor takes no action.

Adjustment then appears elsewhere: liquid rates compress when unilateral retention power is high; competitive offers move less when the outside option is weak; and banks steepen the term structure offered to high-balance, rate-elastic depositors while widening complexity premia on odd-maturity and gated contracts.[5] Uniform pricing and concentration-dependent transmission are therefore compatible. Uniformity is a property of the anchor; local market structure shapes the rest of the menu.

The empirical implication is immediate. If the relevant object is a deposit menu rather than a single posted rate, identification requires data that observe the full vector of posted rates at high frequency, within banks across local markets, and over both tightening and easing regimes. I use data from a commercial deposit-benchmarking platform through which U.S. financial institutions benchmark local deposit pricing against competitors and monitor contract menus in real time. That distinction matters because the platform records not a regulatory average, a survey response, or an ex post portfolio yield, but the posted contract prices that institutions actually benchmark, compare, and revise. The empirical advantage is therefore not only granularity, but proximity to the pricing problem itself. The raw panel begins in 2015Q1 and extends

[5]See Gabaix and Laibson (2006) on shrouded attributes and complexity premia.

through 2024Q4. It contains approximately 11 million bank–county–week observations before sample construction and roughly 9 million observations in the main analysis sample, spanning more than 2,100 commercial banks across 2,930 counties and covering approximately 93 percent of U.S. commercial-bank deposits. Within-bank posted-rate dispersion is large: the average within-bank rate range is 143 basis points in the regression sample, direct evidence that a single institution simultaneously prices depositors with different elasticities, balances, liquidity needs, and search frictions. Section 4 describes the data, product taxonomy, and sample restrictions in detail.

Figure 1 illustrates the mechanism before the formal design. Within Bank of America—one of the largest deposit franchises in the United States—the anchor layer of the menu is essentially fixed through the 2022–2023 tightening cycle. The money market account rate remains nearly unchanged across counties despite 525 basis points of Fed tightening, and the vanilla 12-month rack CD behaves similarly. On these contracts, the uniform-pricing view is right: they are centrally posted anchors aimed at high-friction depositors for whom marginal local competition is second-order. But the same institution, in the same quarters, simultaneously activates a different layer of the menu. Its best available 12-month CD offer, dormant at a single value through 2021 and early 2022, fans out across counties after liftoff; by 2023Q2, Bank of America posts 30 distinct county-level values, with a 10th–90th percentile range from 0.05 to 4.51 percent. Promotional specials expand in parallel, and odd-term contracts go further still, posting rates above the bank's own best standard-maturity offer in the same quarter—a complexity premium layered on top of competitive pricing. The point is not that the uniform-pricing view is wrong and the deposits-channel view is right, or vice versa. It is that both are true, but on different margins of the same pricing problem. The anchor stays uniform; the competitive menu rotates.

To organize the mechanism, I use a stylized framework of deposit-menu pricing.

The key object is not a single deposit rate but a menu of contract prices. Some contracts are anchors: default prices that apply under inertia or no action. Others are competitive offers posted to attract or retain more rate-sensitive balances. A useful way to think about the pricing problem is that markdowns vary with the elasticity a bank faces on each contract, and those elasticities need not respond to market structure in the same way across products. I refer to the best rival offer available in a local market as the *local competitive frontier*. This distinction matters because the bank's own best offer is observed in the data, but the economic force of interest is often the frontier it faces.

Two dimensions of market structure matter, and they matter differently. *Own market power*—HHIC, the bank's squared own deposit share—governs retention on liquid, non-maturity products: more captive funding lowers outflow sensitivity and raises the optimal markdown on liquid rates during tightening. *Rival concentration*—$\text{HHI}_{\text{OTH}}$, the residual HHI attributable to competitors—governs the quality of the outside option that disciplines the local competitive frontier. When rivals are concentrated, fewer institutions bid aggressively for elastic depositors, so the best local rival offer moves less with policy. A bank's own best posted offer is then its response to that frontier, not the frontier itself, and can therefore reflect both rival structure and own retention power.

The framework's sharpest empirical implication is therefore comparative rather than scalar. It does not require every product-level effect to sort cleanly into a single channel, especially for contracts that sit between anchor and frontier. It predicts instead that, within the same bank and quarter across geographies, local market structure should reshape relative prices inside the menu. That is the central empirical test of the paper. Vanilla anchors should remain inert. Liquid rates should compress with own concentration. The local competitive frontier and the short end of the

high-balance menu should weaken with rival concentration. And within-menu wedges should move in the directions implied by rate channeling.

The empirical design follows these predictions directly. The vanilla anchor rate provides the falsification test. HHIC should matter most for liquid products and bank-specific markdowns. $\text{HHI}_{\text{OTH}}$ should matter most for the local competitive frontier and the outside option facing rate-elastic depositors. And within-menu wedges should show how compensation is reallocated across liquidity, maturity, and comparability. I estimate product-level pass-through and wedge dynamics in a bank-by-county-by-week panel spanning two tightening cycles, the zero-lower-bound period, and the subsequent easing. The baseline specification absorbs bank-by-quarter fixed effects, market-by-quarter fixed effects, and bank-by-county fixed effects, with standard errors clustered at the bank and county levels. This structure controls for all time-varying bank-wide determinants of deposit pricing and all time-varying market-wide determinants, including local demand shocks and competitor responses. Identification therefore comes from within-bank, within-market, within-quarter variation: within a given bank-quarter, the only remaining source of cross-county variation in pass-through is local market structure.

Four findings establish rate channeling. First, vanilla pass-through is locally inert: both HHIC and $\text{HHI}_{\text{OTH}}$ coefficients are economically small and statistically indistinguishable from zero, sharply rejecting local-demand confounds. Second, MMA pass-through declines with own market power, while rival concentration enters insignificantly. Third, short-maturity high-balance best rates compress with rival concentration, and the high-balance maturity slope steepens. Fourth, a leave-one-out decomposition of the consumer-facing MMA gap shows that the dominant margin is a weaker local frontier rather than a larger unilateral markdown by the focal bank. Because a bank's own best posted offer is an endogenous response to the local frontier, some best-offer level

effects load on both concentration margins. The sharper evidence therefore lies in the frontier and relative-price objects, where the mechanism is observed most directly. Subsequent sections show that the rotation is asymmetric over the policy cycle. It loads most strongly on the anticipated component of policy, while pure-surprise effects are generally weak and information effects are narrower and concentrated in the own-concentration margin. The rotation also generates real deposit recomposition and net-interest-margin effects during the 2022–2023 tightening cycle.

A natural concern is that the deposit menu itself is endogenous, so the within-menu results could reflect selection into offering different products rather than repricing within a stable menu. The design deals with this concern in two ways. First, the main wedge regressions are estimated only when the bank is simultaneously offering both relevant contracts in the same county and quarter, so identification comes from relative pricing conditional on coexistence of the products. Second, the menu-architecture evidence shows that product shares, product counts, and availability move little with the concentration interaction. The main response is therefore not extensive-margin redesign of the menu, but intensive-margin repricing within a menu that is largely already in place. In addition, the bank-by-county and bank-by-quarter fixed effects absorb all permanent menu differences across bank–market pairs and all bank-wide menu changes over time, so identification relies only on within-pair, within-quarter repricing that cannot be driven by systematic menu redesign.

The main point is therefore simple: local concentration does not shift the entire menu uniformly. It changes where pass-through appears inside the menu.

The paper contributes to four literatures. First, it refines the deposits channel by identifying the product-level margins through which market structure operates—within-menu wedges, maturity tilting, and frontier suppression—margins that disappear in one-rate designs. Viewed through this lens, the classic deposit-channel result is the

aggregate implication of a richer menu-pricing problem rather than a competing view. Second, the paper reconciles the deposit-channel and uniform-pricing literatures by showing that uniformity is a property of anchor contracts aimed at high-friction depositors, while local market structure governs the margins relevant for elastic depositors. Third, it provides pricing-side micro-evidence for the depositor heterogeneity emphasized in structural models of deposit demand, connecting demand segmentation to concrete instruments on the posted menu. Fourth, it supplies an antitrust-relevant decomposition showing that the main consumer-facing pricing distortion from concentration operates through market-level suppression of the outside option rather than through unilateral markdowns by a single bank.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 develops the stylized framework and derives the empirical predictions. Section 4 describes the data, product taxonomy, market structure measures, and identification strategy. Section 5 presents the main product-level pass-through and within-menu wedge results. Section 5.3 implements the leave-one-out frontier decomposition. Section 5.5 links menu rotation to deposit flows and bank balance-sheet outcomes. Section 5.6 studies the anticipated-versus-surprise shock decomposition and asymmetry over the policy cycle. Section 5.7 presents robustness tests. Section 6 concludes.

# 2 Related Literature

This paper sits at the intersection of three literatures on deposits, market power, and monetary transmission.

First, a large deposits-channel literature shows that pass-through from policy rates to deposit rates is incomplete and weaker in more concentrated markets. The canonical

statement is Drechsler et al. (2017), who show that when the federal funds rate rises, banks widen deposit spreads, deposits flow out of the banking system, and these effects are stronger in concentrated markets. Earlier empirical work pointed in the same direction. Berger and Hannan (1989) relate higher market concentration to lower deposit rates. Hannan and Berger (1991) document substantial rigidity in deposit pricing in banking, and Neumark and Sharpe (1992) show that banks in concentrated markets are slower to raise deposit rates when market rates rise and faster to cut them when market rates fall. Drechsler et al. (2021) extend the logic to maturity transformation, showing that the deposit franchise hedges interest-rate risk because banks can pay deposit rates that are low and insensitive to market rates. More recent work embeds imperfect competition more explicitly in banks' funding choices. Wang et al. (2022) estimate a dynamic banking model in which monetary policy affects imperfectly competitive banks' funding costs and banks optimize pass-through to borrowers and depositors. Enkhbold (2023) show that the pass-through of policy to mortgage rates depends jointly on deposit-market concentration and wholesale-funding reliance. Eichenbaum et al. (2025) develop a banking model in which household attentiveness to deposit rates is state dependent, generating state-dependent responses of banks' net interest margins and aggregate activity to monetary policy.

Second, a parallel literature asks whether local market structure can operate through posted deposit rates at all. Radecki (1998) provides early evidence that many banks set uniform deposit rates over geographic areas larger than the local county or metropolitan market. Using branch-level data, Begenau and Stafford (2023) argue that U.S. banks, especially the largest, predominantly use uniform deposit-rate setting policies, implying limited within-bank local variation in posted levels. Granja and Paixão (2026) show that after mergers, acquired branches are often brought into line with the acquirer's network-wide pricing schedule, again suggesting limited scope for

local concentration to affect the posted level of a representative deposit rate. Read literally, this evidence sits uneasily with scalar versions of the deposits channel.

Third, a literature on depositor heterogeneity, search frictions, and price discrimination explains why that apparent tension need not be real. Hortaçsu and Syverson (2004) show that search frictions and product differentiation can sustain substantial price dispersion even for financially homogeneous products. Egan et al. (2017) develop a structural model of deposit competition in which insured and uninsured depositors choose among differentiated banks, and the two groups differ sharply in their sensitivity to bank distress. Yankov (2024) shows that costly search can generate deposit-rate dispersion, sticky adjustment, and low and asymmetric pass-through. Egan et al. (2025) document that depositors shop infrequently for deposits, giving rise to dynamic competition for "sleepy" funding. d'Avernas et al. (2023) show that large banks offer better liquidity and financial services but lower deposit rates, and that their customers are less rate elastic than customers of small banks. The framework in this paper is informed by this literature as well as by classic models of second-degree price discrimination and shrouded attributes (Gabaix and Laibson, 2006; Mussa and Rosen, 1978; Tirole, 1988).

The contribution of this paper is to open up the pricing object. Prior work in the deposits channel and the uniform-pricing literature can both be right once the relevant object is not a representative posted rate but the menu of contracts through which banks price heterogeneous depositors. Empirically, the paper separates own concentration from rival concentration within the same bank across geographies and shows which contracts remain inert, which margins reflect unilateral retention power, which reflect the local competitive frontier, and how within-menu wedges reallocate pass-through across products. Those objects are not visible in designs that average deposit pricing to a single scalar rate, and they complement rather than overturn the

earlier literatures.

# 3 A Stylized Framework for Deposit-Menu Pricing

This section introduces a stylized framework designed to organize the empirical objects in the paper. The framework is built around a simple pricing logic: markdowns vary with the elasticity a bank faces on each contract, and those elasticities need not respond to market structure in the same way across products. The goal is to discipline which margins the paper tests and why, not to provide a full structural characterization of deposit demand. None of the paper's identification relies on the framework being literally correct.

The starting point is simple: a bank does not choose a single deposit rate. It chooses a *menu* of contracts indexed by liquidity, maturity, balance threshold, eligibility, and comparability. Let $k \in \mathcal{K}$ index products and let $r_{ikc}$ denote the posted rate that bank $i$ offers on product $k$ in market $c$. The bank chooses the vector $r_i \equiv \{r_{ikc}\}_{k\in\mathcal{K}}$ to fund its balance sheet at minimum cost while retaining relationship balances and competing for more rate-sensitive funds. Write profits as

$$\Pi_i = \sum_{k\in\mathcal{K}} \left[v_k(R) - r_{ikc}\right] D_{ikc}(r_i, r^*_{kc}; \theta_c) - \Phi(r_i), \tag{1}$$

where $R$ is the policy-sensitive outside funding benchmark, $v_k(R)$ is the shadow value of deposits of type $k$, $D_{ikc}$ is product-level deposit demand, $r^*_{kc}$ is the relevant rival offer or market frontier, and $\Phi(r_i)$ captures any costs of operating a richer menu. The central pricing object is therefore the vector $r_i$, not a representative posted rate.

Depositors are heterogeneous in rate sensitivity (cf. Egan et al., 2025; Eichenbaum et al., 2025). One group is relationship-bound, inattentive, or high-friction. These

depositors value convenience, branch proximity, transaction services, or inertia, and they are relatively unlikely to compare offers across institutions. A second group is more rate-sensitive. These depositors search, compare products across banks, and respond to the best available offer in the market. In equilibrium, banks use the menu to allocate different depositor types into different contracts. Standard vanilla products and some liquid accounts serve as anchors for high-friction balances; competitive best offers, high-balance contracts, acquisition products, and less comparable contracts serve as the margins on which banks compete for or screen more elastic funds.

Abstracting from cross-product terms, contract-by-contract pricing can be summarized by a Lerner-style condition:

$$v_k(R) - r_{ikc} = -\frac{D_{ikc}}{\partial D_{ikc}/\partial r_{ikc}}. \tag{2}$$

What matters is that the relevant elasticity differs across products, and that different dimensions of market structure load on different elasticities. I distinguish two such dimensions. The first is *own concentration*, $HHIC_{ic}$, the bank's own contribution to concentration in market $c$. In reduced form, this captures unilateral retention power over relationship-bound balances. The second is *rival concentration*, $HHI_{ic}^{OTH}$, the residual concentration attributable to competing institutions. In reduced form, this captures the quality of the outside option and the competitiveness of the local frontier.

These two objects generate different comparative statics. Higher own concentration raises the mass of captive balances and lowers the retention elasticity of liquid, non-maturity products. The optimal response is weaker rate adjustment on those products when policy tightens. Higher rival concentration, by contrast, weakens the outside option available to rate-sensitive depositors because fewer rivals bid aggressively on the frontier. The optimal response is weaker adjustment in frontier offers and stronger

reliance on within-menu wedges such as steeper maturity schedules, acquisition premia, or less comparable contracts. Anchor products can therefore remain locally inert even when other parts of the menu move sharply with market structure.

A useful way to separate the objects is to write a bank's own best posted offer as

$$r_{ikc}^{best} = r_{kc}^{frontier} - m_{ikc}, \tag{3}$$

where $r_{kc}^{frontier}$ is the local competitive frontier and $m_{ikc}$ is the bank-specific markdown relative to that frontier. Rival concentration primarily shifts the frontier; own concentration primarily shifts the markdown. A bank's own best posted offer can therefore load on both concentration margins even when the outside-option prediction applies most cleanly to the frontier itself. This distinction will matter in the empirical work below.

The key implication is that monetary transmission operates through *menu rotation* rather than through a single posted level. Market structure changes the mapping from the policy rate into the vector of contract rates: some product sensitivities decline, others rise, and within-menu wedges move systematically. The framework yields directional predictions rather than a knife-edge one-product/one-channel mapping. Products that sit between anchor and frontier may load on both own and rival concentration. The cleanest predictions are four directional ones.

1. **Anchor invariance.** Standard vanilla anchor contracts should be locally inert: their responsiveness to policy should load weakly, if at all, on either own concentration or rival concentration.

2. **Liquid-rate compression through own concentration.** Products used to retain liquid balances should respond less to policy when own concentration is

higher, because unilateral retention power lowers the outflow elasticity of those balances.

3. **Outside-option / frontier channel.** The local competitive frontier should respond less to policy when rival concentration is higher, because the outside option available to elastic depositors is weaker. A bank's own best posted offer may also under-respond with own concentration, because a larger captive base reduces the bank's incentive to match that frontier aggressively.

4. **Within-menu wedge adjustment.** Screening wedges should move systematically with concentration, but not necessarily all on the same margin. Duration- and teaser-related wedges may load most clearly on rival concentration, while complexity premia may load more strongly on own retention power. The common prediction is that banks reallocate compensation away from the most transparent contracts and toward duration, eligibility, or imperfect comparability. The economic reason is simple. Repricing an anchor or liquid default contract changes the remuneration on a large stock of infra-marginal balances. Repricing a targeted contract changes the remuneration only for the balances the bank is trying to attract, retain, or sort. A bank with market power therefore prefers to reallocate compensation across contracts rather than move all rates in parallel.

These predictions motivate the empirical objects directly. Prediction 1 is tested on vanilla 3-month and 12-month rates. Prediction 2 is tested on MMA levels and liquid-product spreads such as MMA–vanilla and MMA–savings. Prediction 3 is tested most cleanly on frontier objects: the leave-one-out best rival offer, the short-horizon high-balance best rate, and bank best offers as endogenous responses to that frontier. Prediction 4 is tested on within-menu wedges: the high-balance maturity slope, the acquisition premium, the odd-term premium, and related relative-price objects inside

the same bank's menu.

Accordingly, the sharpest empirical prediction is comparative rather than scalar. The framework does not require every contract-level effect to sort cleanly into a single concentration margin, especially for products that sit between anchor and frontier. It does, however, predict that, within the same bank and quarter across geographies, local market structure should reshape relative prices inside the menu. That is where the evidence is strongest. Vanilla anchors remain inert. Liquid rates compress with own concentration. The local frontier and the short end of the high-balance menu weaken with rival concentration. And within-menu wedges move in the directions implied by menu rotation.

One further observation makes precise why relative prices are not second-order. The bank's effective deposit funding cost is

$$C_{it} = \sum_{k \in \mathcal{K}} \omega_{ikt}\, r_{ikt}, \tag{4}$$

where $\omega_{ikt}$ is the share of balances allocated to product $k$. Its sensitivity to the policy rate decomposes as

$$\frac{dC_{it}}{dR_t} = \underbrace{\sum_k \omega_{ikt} \frac{dr_{ikt}}{dR_t}}_{\text{pass-through on existing balances}} + \underbrace{\sum_k r_{ikt} \frac{d\omega_{ikt}}{dR_t}}_{\text{composition}}. \tag{5}$$

The first term is pass-through on a fixed portfolio. The second term is composition: once balance shares respond to the menu, changes in relative prices alter funding costs through reallocation. Calling within-menu wedges second-order would be correct only if depositors' product choices were fixed, that is, only under the knife-edge assumption that $d\omega_{ikt}/dR_t = 0$ for all $k$. They are not. The balance-sheet evidence presented below confirms that product shares move with the policy cycle. Menu rotation therefore

enters effective funding costs at first order.

# 4 Data, Variables, and Identification

## 4.1 Data Sources

**Proprietary deposit pricing data.** The primary data source is a commercial deposit-benchmarking platform through which U.S. financial institutions monitor competitor posted rates in local markets and benchmark their own product menus in real time. The raw observation is a posted contract: a specific rate offered by a specific institution on a specific product and balance tier in a given week. That matters for the economics of the data. The platform is not a regulatory filing, a survey response, or an ex post portfolio average. It is a record of the posted contract prices that institutions observe, compare, and revise for competitive benchmarking.

The closest antecedent in the literature is RateWatch (now S&P Global Market Intelligence), used in Drechsler et al. (2017), Wang et al. (2022), Yankov (2024), and much of the prior empirical literature on deposit pricing and monetary transmission. RateWatch collects advertised deposit rates at the branch level on a weekly basis from contacts at nearly 100,000 financial institution locations across the United States. The proprietary platform used in this paper differs in both institutional role and empirical content. It is a deposit-analytics solution whose product suite is built from contributed data covering trillions of dollars in account-level deposit balances and is used by pricing, treasury, and asset-liability teams at participating institutions. The platform offers broad institutional coverage and rich observation of the menu itself—the posted rates that banks and their competitors actually benchmark against one another in real time. Most importantly for identification, within-bank cross-county variation in

posted rates reflects genuine strategic variation rather than survey-measurement noise or administrative lag.

The data extraction used in this paper spans 2015Q1 through 2024Q4—approximately 520 weeks covering two complete tightening cycles (2016–2018 and 2022–2023, the latter the largest in the Federal Reserve's 111-year history), the pandemic-era zero-lower-bound period, and the subsequent easing episode. After sample construction, the panel contains approximately 9 million bank–county–week observations across more than 2,000 commercial banks (approximately 1,200–1,400 in any given year) in 2,930 counties. These institutions hold approximately $17 trillion in domestic deposits, or roughly 93 percent of the $18.8 trillion U.S. commercial banking total (UBPR, 2023Q4). The estimation sample is restricted to banks that operate in at least two counties in a given week: single-county banks contribute no within-bank cross-county variation and therefore provide no identifying information for the main specifications.

**FDIC Summary of Deposits.** Market structure variables are constructed from the FDIC Summary of Deposits, which reports institution-level deposit balances by county as of June 30 each year. Deposit shares are computed at the institution–county–year level and merged into the weekly pricing panel with a one-year lag. This timing makes the concentration measures predetermined with respect to the pricing outcomes they are used to explain and insulates the design from endogenous deposit flows within the policy cycle.

**Federal Reserve policy rate and monetary surprises.** The federal funds target rate is the upper limit of the federal funds target range (FRED series DFEDTARU).[6]

[6]Board of Governors of the Federal Reserve System (US), Federal Funds Target Range—Upper Limit [DFEDTARU], retrieved from FRED, Federal Reserve Bank of St. Louis; `https://fred.stlouisfed.org/series/DFEDTARU`.

Quarterly policy changes, $\Delta FFR_t$, are computed as the change in the end-of-quarter target relative to the previous quarter. Because the upper limit, lower limit, and midpoint of the target range move in parallel, first-differencing makes the series choice immaterial. The shock decomposition uses high-frequency monetary policy surprises from Bauer and Swanson (2023), as updated on the Federal Reserve Bank of San Francisco's monetary policy surprises data page.[7] The raw surprise is constructed from 30-minute changes in money-market futures around FOMC announcements, and the orthogonalized surprise removes the Fed-information component following Bauer and Swanson (2023). Because these surprises are observed at the meeting level, the quarterly shock series used in the paper are formed by summing meeting-level surprises within quarter before decomposing quarterly policy changes into anticipated, information, and pure-surprise components.

**UBPR balance-sheet data.** Bank-quarter balance-sheet outcomes are drawn from the Uniform Bank Performance Report (UBPR), compiled by the Federal Financial Institutions Examination Council (FFIEC) from quarterly Call Report filings. The UBPR provides standardized performance ratios at the bank–quarter level, including annualized deposit yields and funding costs, net interest margin, and deposit-composition breakdowns by category and maturity. These variables are used in the concentration-exposure tests (Section 5.5) to link the local pricing distortions documented in the deposit-menu analysis to realized bank-level funding and profitability outcomes. Three key measures are the cost of total interest-bearing deposits (UBPRE106), the cost of time deposits (UBPRHR55), and net interest income as a percent of average assets (UBPRKX40). Deposit-composition shares and small-denomination time-deposit volumes by remaining maturity are also obtained from the UBPR.

[7] `https://www.frbsf.org/research-and-insights/data-and-indicators/monetary-policy-surprises/`.

## 4.2 Product Taxonomy and Sample Construction

**Retail filter.** The analysis focuses on personal, non-wholesale deposit products. A contract is included in the retail sample if it is in the personal segment, is not a wholesale or interbank instrument, and has a minimum opening balance of no more than $10,000. This restriction isolates the mass-market deposit menu. Products with minimum balances between $10,000 and $100,000 are excluded from both the retail and high-balance samples. The high-balance sample is defined symmetrically as personal, non-wholesale products with minimum balance at or above $100,000, corresponding to the conventional high-balance threshold and the segment most directly exposed to rate shopping and balance reallocation.

**Product categories.** Within the retail sample, each posted rate is assigned to one of six mutually exclusive categories based on contractual features. *Acquisition* products carry new-money restrictions, teaser rates, or promotional terms designed to attract new depositors. *Retention* products are relationship-conditioned offers available only to existing customers who satisfy balance or cross-sell criteria. *Promotional* products are time-limited specials. *Customer optionality* products embed depositor-beneficial options such as no-penalty or bump-rate certificates. *Bank optionality* products embed bank-beneficial options such as callable certificates or rate caps. *Vanilla* products are standard certificates with none of the preceding features. Products with multiple features are assigned according to a strict hierarchy:

$$\text{acquisition} \succ \text{retention} \succ \text{promotional} \succ \text{customer optionality} \succ \text{bank optionality} \succ \text{vanilla}.$$

The vanilla category is therefore the residual anchor contract: the posted rate available without product enhancement, gating, or embedded option value. The underlying

data are far more granular: the raw contract records distinguish senior, junior, student, minor, IRA, and other depositor-segment variants, as well as callable, bump-rate, step-up, no-penalty, and add-on certificate structures—hundreds of distinct product descriptions per institution. The six-category taxonomy aggregates this variation into the economically meaningful distinctions required by the pricing framework.

**Rate construction.** Within each bank–county–week, rates are constructed from the underlying contract-level records using column-maximum aggregation: for a given product category and maturity, the retained rate is the maximum posted rate in the cell. This follows the standard convention that the economically relevant competitive signal is the best available posted offer. The retained rates should therefore be interpreted as the bank's relevant contract prices for a given product tier in a given local market.

The main rate objects are:

- $r_{12}^{\text{van}}$: vanilla 12-month certificate rate;
- $r_{12}^{\text{best}}$: best available 12-month retail rate;
- $r_{3}^{\text{best}}$: best available 3-month retail rate;
- $r^{\text{mma}}$: money market account rate;
- $r^{\text{odd}}$: best available odd-term certificate rate;
- $r_{12}^{\text{jmb,best}}$: best available high-balance 12-month rate;
- $r_{3}^{\text{jmb,best}}$: best available high-balance 3-month rate.

The *high-balance term slope* is defined as

$$slope_{ict}^{\text{jmb}} = r_{12,ict}^{\text{jmb,best}} - r_{3,ict}^{\text{jmb,best}},$$

which measures the term premium embedded in the menu offered to the most rate-

elastic, high-balance depositors.

**Within-menu range.** A useful summary of the menu architecture is the within-bank rate range, defined as the maximum minus the minimum posted rate across all products in a given bank–county–quarter cell. In the estimation sample, this range averages 143 basis points. That is not a descriptive curiosity. It is the observable footprint of a multi-product pricing problem in which a single institution simultaneously prices depositors with different elasticities, balance sizes, liquidity needs, and search frictions.

## 4.3 Market Structure Variables

Let $s_{ict}$ denote institution $i$'s deposit market share in county $c$ in year $t$, computed from the FDIC Summary of Deposits as institution deposits in county $c$ divided by total county deposits. The county-level Herfindahl index is

$$HHI_{ct} = \sum_{j \in \mathcal{B}_{ct}} s_{jct}^2, \tag{6}$$

where $\mathcal{B}_{ct}$ is the set of institutions operating in county $c$ in year $t$.

The paper's central decomposition splits total concentration into two components:

$$HHI_{ct} = \underbrace{s_{ict}^2}_{HHIC_{ict}} + \underbrace{\sum_{j \neq i} s_{jct}^2}_{HHI_{ict}^{OTH}}. \tag{7}$$

The first term, $HHIC_{ict}$, is the focal institution's own contribution to local concentration. Economically, it proxies for unilateral retention power over local balances. The second term, $HHI_{ict}^{OTH}$, is rival concentration. It measures the competitiveness of the outside option faced by rate-sensitive depositors.

Both components are lagged one year before merging into the pricing panel. The

$HHIC$ and $HHI^{OTH}$ values applied to a week in year $t$ are therefore constructed from the June Summary of Deposits of year $t-1$. In the regressions, both variables are standardized to unit variance, so coefficients can be interpreted as basis-point changes in rate adjustment per one-standard-deviation change in the relevant concentration margin for a 100-basis-point policy move.

## 4.4 Empirical Specification

**From weekly posted rates to quarterly outcomes.** The raw pricing panel is at the bank–county–week level. The main regressions operate at the bank–county–quarter level, so the paper maps weekly rates into quarterly outcomes using two transformations matched to the product's pricing structure.

For term products such as certificates of deposit, the paper uses a *quarterly rate accumulator*:

$$\Delta^{\mathrm{DRQ}} r^{p}_{ict} = \sum_{w \in Q_{ict}} \left( r^{p}_{icw} - r^{p}_{ic,w-1} \right), \tag{8}$$

where $Q_{ict}$ contains weeks in quarter $t$ for which the product is observed in consecutive weeks for the same bank–county pair. This is the cumulative within-quarter repricing of a posted contract. If a product is not observed in consecutive weeks, the quarterly change is coded missing rather than zero, so the regression conditions on active product offering rather than treating non-posting as non-response.

For liquid products such as money market accounts and savings, the paper uses the quarter-to-quarter change in the within-quarter mean:

$$\bar{r}^{p}_{ict} = \frac{1}{|W_{ict}|} \sum_{w \in W_{ict}} r^{p}_{icw}, \qquad \Delta^{\mathrm{Dqm}} r^{p}_{ict} = \bar{r}^{p}_{ict} - \bar{r}^{p}_{ic,t-1}. \tag{9}$$

The choice between these two operators is not a measurement convention but an

economic implication of the contractual structure of each product type. A certificate of deposit is a fixed-rate bilateral contract: once originated, its rate is locked for the full term, and the bank cannot reprice it. Monetary transmission for CDs therefore operates through the origination margin—the rate quoted on new contracts today—not through the stock-average rate, which mainly reflects the slow runoff of legacy contracts at historical rates. $\Delta^{\mathrm{DRQ}}$ is designed to isolate this origination-pricing decision. An MMA or savings account, by contrast, is a variable-rate callable instrument: the bank posts a standing rate and can change it at any time, and every depositor in the product earns the new rate as soon as it is updated. For a variable-rate liquid account, there is no separate origination margin: the posted rate *is* the stock rate. The relevant quarterly object is therefore the average remuneration actually applying to the stock of balances, which is what $\Delta^{\mathrm{Dqm}}$ captures.

This distinction also explains why the $\Delta^{\mathrm{DRQ}}$ and $\Delta^{\mathrm{Dqm}}$ estimands can differ sharply for liquid products. For an MMA, $\Delta^{\mathrm{DRQ}}$ measures the net change in the quoted rate over the quarter, whereas $\Delta^{\mathrm{Dqm}}$ measures the average rate actually prevailing over the quarter. Under the saturated fixed-effect design, common bank-level reset decisions on liquid accounts can be largely absorbed by bank×quarter effects, making the endpoint-style $\Delta^{\mathrm{DRQ}}$ operator weak even when differences in average quarterly remuneration remain economically relevant. A weak $\Delta^{\mathrm{DRQ}}$ result for MMA products should therefore not be read as evidence against the liquid-retention mechanism; it indicates that the relevant variation lies in the average quarterly stock rate rather than in the net endpoint change of the weekly posted quote.[8]

[8]Appendix B.1 shows that the liquid HHIC channel is present under $\Delta^{\mathrm{Dqm}}$ but disappears under both end-of-quarter and cumulative-quote operators, consistent with the view that for MMAs the relevant quarterly object is the average stock rate applying to balances rather than the endpoint change in the posted quote.

**Spread and leave-one-out construction.** All within-menu spreads are constructed from raw weekly rates before winsorization and are then winsorized as composite objects at the 1st and 99th percentiles. This avoids mechanically compressing tail spreads by taking differences of separately winsorized components. The distinction matters most for contracts with fat-tailed price distributions, such as high-balance, promotional, and odd-term products.

For the leave-one-out analysis, the market frontier for product $p$ in county $c$ and week $w$ is the highest rival offer:

$$r_{icw}^{p,\mathrm{LOO}} = \max_{j \neq i,\, j \in \mathcal{B}_{cw}} r_{jcw}^{p}, \tag{10}$$

with careful tie handling: a bank tied at the frontier is compared to the same frontier rate, while a sole frontier setter is compared to the second-highest rival. The leave-one-out gap $\ell_{icw}^{p} = r_{icw}^{p} - r_{icw}^{p,\mathrm{LOO}}$ is passed through the same quarterly accumulation pipeline as the level rates.

**Baseline specification.** Let $\Delta r_{ict}^{p}$ denote the quarterly change in product $p$ for bank $i$ in county $c$ and quarter $t$. The baseline specification is

$$\Delta r_{ict}^{p} = \alpha_{it} + \gamma_{mt} + \delta_{ic} + \beta_1^{p}\Big(\Delta FFR_t \times zHHIC_{ict}\Big) + \beta_2^{p}\Big(\Delta FFR_t \times zHHI_{ict}^{OTH}\Big) + \varepsilon_{ict}^{p}, \tag{11}$$

where $\alpha_{it}$ are bank-by-quarter fixed effects, $\gamma_{mt}$ are market-by-quarter fixed effects, and $\delta_{ic}$ are bank-by-county fixed effects. The level effect of $\Delta FFR_t$ does not enter separately because it is a quarter-level scalar absorbed by the bank-by-quarter fixed effects. The interaction terms survive because concentration varies across counties within the same bank-quarter. Identification therefore comes from within-bank, within-

quarter, cross-county differences in how policy adjustment loads onto the menu as local market structure varies.

**What the fixed effects absorb.** Each fixed-effect layer addresses a distinct source of confounding. Bank-by-quarter fixed effects absorb all time-varying bank-wide determinants of pricing, including treasury decisions, wholesale funding conditions, balance-sheet composition, and institution-wide repricing strategy. Market-by-quarter fixed effects absorb all local shocks common to all banks in the same market and quarter, including local demand conditions, macroeconomic news, and competitor responses. Bank-by-county fixed effects absorb time-invariant features of a bank's local presence, including branch footprint, franchise strength, and persistent customer composition.

The identifying variation is therefore narrow and economically meaningful: the same bank, in the same quarter, facing the same aggregate policy shock and the same institution-wide repricing environment, adjusts one local menu differently from another because local competitive structure differs. This is exactly the relevant design if the object of interest is not whether some banks are more aggressive than others, but whether the *same bank across geographies* prices the same menu differently when local market power differs.

**Market definition and fixed-effect assignment.** For money market and savings products, the paper uses a Combined-MSA market fixed effect that assigns urban counties to their MSA and rural counties to a state residual group. This keeps rural observations, which are important for liquid-product results. For certificate, high-balance, and odd-term products, the paper uses MSA-by-quarter fixed effects and restricts to MSA counties, reflecting the geographic scope of competition in those

products and the leave-one-out requirement of multiple active institutions. Standard errors are two-way clustered at the bank and county levels throughout (Cameron et al., 2011; Petersen, 2009).

**Why the joint specification is essential.** The baseline always enters $zHHIC$ and $zHHI^{OTH}$ jointly. This is not a robustness exercise but a theoretical requirement. The two components capture different mechanisms and are mechanically related by construction. Estimating them separately would confound unilateral retention power with the quality of the rival-set outside option. The joint specification therefore matches the distinction developed in Section 3 and keeps the interpretation of the coefficients economically clean.

**Threats to identification.** Three concerns are most natural. The first is that local deposit demand shocks are correlated with concentration. The market-by-quarter fixed effects absorb all shocks shared within a market, and the remaining concern would have to be bank-specific, within-market, and correlated with local share. The paper's internal falsification is the vanilla anchor rate: if omitted local demand were driving the results, it should move anchor contracts as well. Empirically, it does not—while MMA and high-balance best rates load on their respective concentration components, vanilla contracts do not.

The second concern is bank sorting into concentrated markets. Banks with low average responsiveness might simply dominate concentrated counties. The design absorbs bank-by-quarter and bank-by-county fixed effects, so the coefficients are identified from within-bank differences across counties, not from cross-bank comparisons.

The third concern is reverse causality from pricing into concentration. Concentration measures are lagged one year, and the shift-share design below provides an

additional layer of insulation.

**Asymmetry.** To test whether concentration distorts tightening and easing differently, the paper splits policy changes into hike and cut components:

$$\Delta_t^+ = \max(\Delta FFR_t, 0), \qquad \Delta_t^- = \max(-\Delta FFR_t, 0), \tag{12}$$

and estimates directional interactions with both concentration components jointly. Wald tests assess whether the concentration gradient is the same in hiking and easing quarters.

**Shock decomposition.** Following Bauer and Swanson (2023), the quarterly policy change is decomposed into anticipated, information, and pure-surprise components:

$$\Delta FFR_t = \underbrace{\Delta FFR_t - \epsilon_t^{\text{raw}}}_{\text{anticipated}} + \underbrace{\epsilon_t^{\text{raw}} - \epsilon_t^{\text{orth}}}_{\text{information}} + \underbrace{\epsilon_t^{\text{orth}}}_{\text{pure surprise}} . \tag{13}$$

Each component is interacted separately with the two concentration margins. This decomposition distinguishes menu adjustment that loads on expected policy paths from adjustment that responds only to unexpected innovations.

**Concentration-exposure design.** To connect menu rotation to bank-level deposit flows and balance-sheet outcomes, the paper constructs a Bartik-style concentration-exposure measure using predetermined geographic footprint shares (Goldsmith-Pinkham et al., 2020). Let $\omega_{ic,0}$ denote bank $i$'s share of pre-sample deposits in county $c$, constructed from the FDIC Summary of Deposits averaged over 2010–2014. The weight is built in two steps: branch-level deposits are first summed to the bank–county–year level, then averaged across years, so that a multi-branch presence in a dense urban

county is not over- or under-represented relative to a single-branch rural presence. The bank-level exposure is

$$\widetilde{HHI}_{it} = \sum_{c \in \mathcal{C}_i} \omega_{ic,0} \cdot HHI_{ct}, \tag{14}$$

where $\mathcal{C}_i$ is the set of counties in which bank $i$ operates. The identifying assumption is that pre-sample footprint shares are predetermined with respect to subsequent bank-quarter outcomes, conditional on fixed effects. The bank-quarter test regression takes the form

$$Y_{it} = \alpha_i + \delta_t + \sum_{s=0}^{2} \beta_s \Big( \Delta \text{FFR}_{t-s} \times z(\widetilde{HHI}_i) \Big) + \varepsilon_{it}, \tag{15}$$

where $Y_{it}$ is a bank-quarter outcome (deposit growth, composition share, yield, or net interest margin from the UBPR), $z(\widetilde{HHI}_i)$ is the standardized Bartik concentration exposure from (14), $\alpha_i$ and $\delta_t$ are bank and quarter fixed effects, and two lags of the interaction are included to capture delayed adjustment. The cumulative effect $\sum_{s=0}^{2} \hat{\beta}_s$ is tested by linear combination. Standard errors are clustered by bank.

The empirical design is built around a narrow claim. The paper does not argue that local market structure moves all deposit rates equally. It argues that local market structure changes *which* contracts absorb policy adjustment and *how* that adjustment is distributed across the menu. The methodology is designed exactly for that claim: high-frequency posted prices, within-bank cross-county comparisons, decomposition of concentration into own and rival components, and a primary focus on within-menu wedges and frontier objects rather than on a single representative rate.

## 4.5 Summary Statistics

The analysis panel spans roughly 1,600 unique banks observed across approximately 2,775 counties over 40 quarters. Table 1 reports descriptive statistics organized by observational level. Panel A covers bank–quarter variables: product-level rates, menu-composition shares, and menu-dispersion measures. Panel B covers county–quarter market structure: HHI, the own-concentration and rival-concentration components, leave-one-out frontier gaps, and county characteristics. Panel C covers bank–county–quarter variables at the regression level: within-menu spreads, slopes, and dispersion. Panel D covers the quarterly federal funds rate, its first difference, and the shock decomposition. The anticipated component accounts for most quarterly policy variation, but the pure-surprise and information components remain economically meaningful, with quarterly standard deviations on the order of 5–10 basis points. Panel E covers bank–quarter balance-sheet variables from the UBPR used in the concentration-exposure tests: deposit yields and costs, net interest margin, and deposit-composition shares.

Several features of the data are worth noting before turning to the regressions. First, the rate hierarchy is preserved on average and consistent with the paper's mechanism: the mean best available 12-month rate exceeds the vanilla 12-month rack rate, confirming that the local frontier lies above the anchor. Second, the high-balance 12-month best rate sits below the retail best rate, but the high-balance maturity slope is substantially positive, indicating that banks embed a term premium in the high-balance menu. Third, the odd-term premium above the best standard-maturity rate is economically large, consistent with a complexity premium directed at depositors who compare products across institutions using standard-maturity benchmarks. Fourth, the within-bank rate range averages 143 basis points in the regression sample, underscoring that scalar pass-through measures miss the central object of interest.

The median county has 8 competing banks. County-level HHI shows substantial dispersion, providing meaningful cross-sectional variation in market structure. The leave-one-out gap—the difference between a bank's own rate and the best rival rate in the same county—is negative on average for money market accounts and standard CDs, indicating that the typical bank sits below the local frontier on its most transparent products. That is precisely the environment in which changes in the quality of the outside option can matter most.

# 5 Results

## 5.1 Product-Level Pass-Through: Anchors, Liquid Products, and the Competitive Frontier

The central prediction of the paper concerns *where* policy adjustment appears within the menu, not whether every individual contract-level coefficient is estimated with equal precision in levels. Tables 2 and 3 should be read together. Table 2 shows which product levels move. Table 3 shows how the same bank's relative prices change across geographies. The latter is the sharper test of the mechanism.

Anchor contracts are locally inert. The stylized framework predicts that vanilla CDs—the standard rack rate with no product enhancement, gating, or embedded option—should be invariant to local market structure because these products serve high-friction depositors whose retention does not depend on the competitive environment (Prediction 1). The data confirm this prediction cleanly. For vanilla certificates, the joint coefficients on own concentration and rival concentration are small and statistically indistinguishable from zero. The 3-month vanilla CD loads at $\hat{\beta}^{\text{HHIC}} = 0.003$ and $\hat{\beta}^{\text{OTH}} = -0.033$; the 12-month vanilla is similarly negligible ($\hat{\beta}^{\text{HHIC}} = -0.147$;

$\hat{\beta}^{\text{OTH}} = -0.015$). None of these coefficients approaches conventional significance levels. This null result is not merely a non-finding—it is economically important. It serves as a built-in falsification test: if the concentration interaction terms were proxying for omitted local demand conditions, those conditions would shift the entire posted menu, including vanilla anchors. The fact that vanilla contracts remain inert while other products on the same menu move sharply rules out such confounds.

The cleanest non-null level margin is liquid pricing. The framework predicts that products used to retain liquid balances should respond less to policy when own concentration is higher, because unilateral retention power lowers the outflow elasticity of captive deposits (Prediction 2). Money market account adjustment declines with own concentration, not rival concentration, exactly as predicted. In the joint specification (Table 2, Panel C, column 3), a one-standard-deviation increase in HHIC lowers MMA pass-through by $0.156^{***}$ basis points per 100-basis-point policy move, while the coefficient on rival concentration is economically small and insignificant. This is the unilateral retention-power margin: banks with a larger captive local base undertransmit on liquid balances because the depositors holding MMA balances in those markets are disproportionately relationship-bound and unlikely to leave even when rates lag.

A second clear level margin is the short end of the high-balance frontier. The framework predicts that competitive best offers aimed at rate-sensitive depositors should respond less when rival concentration is higher, because the outside option itself is weaker (Prediction 3). The 3-month high-balance best rate loads strongly on rival concentration (Table 2, Panel C, column 5). The joint coefficient on $HHI^{\text{OTH}} \times \Delta\text{FFR}$ is $-0.778^{***}$ basis points per 100-basis-point policy move, with no comparable loading on HHIC. When the rival set is itself concentrated—that is, when depositors have few outside options—banks reduce short-maturity high-balance pass-through sharply. This

pattern is consistent with a weaker local frontier: when the rival set is concentrated, banks face less pressure to match aggressive outside offers at the short end of the high-balance menu. The leave-one-out decomposition below sharpens this interpretation by showing, for selected products, that the outside-option channel dominates unilateral markdown.

The best available 12-month CD—the most aggressively priced standard-maturity contract in each local menu—also compresses with concentration, though less cleanly on one margin alone. In the joint specification, the own-share component is $-0.663^{**}$ and the rival component is $-0.425$ (insignificant). This is exactly the kind of contract on which a mixed result is economically plausible: a best 12-month CD sits between anchor and frontier, reflecting both unilateral markdown incentives and the quality of the outside option. The paper does not require such contracts to sort perfectly into only one channel.

This mixed loading is exactly what the framework allows. The bank's own best 12-month offer is not a pure frontier object; it is the bank's endogenous response to the local frontier. In the notation of Section 3, $r_{ikc}^{best} = r_{kc}^{frontier} - m_{ikc}$, where rival concentration shifts the frontier and own concentration shifts the markdown relative to that frontier. Because both objects move, the bank's own best 12-month rate need not load as cleanly on rival concentration as the frontier itself or the short end of the high-balance menu.

Two additional level products move in the direction suggested by selective compensation on screening contracts. The 12-month high-balance best rate loads positively on rival concentration, and odd-term CDs also rise with rival concentration. These level effects should not be read as the cleanest test of the screening margin, because screening is observed most directly in relative prices inside the menu. Their role is instead to show that contracts farther from the transparent anchor do not uniformly

compress with concentration. Some products absorb compensation selectively as banks redirect pass-through across the menu rather than transmitting it uniformly through a single posted rate.

Taken together, the level results establish the paper's first claim. Uniform anchors and concentration-dependent transmission are not contradictory facts. They are distinct margins of the same pricing problem: the anchor remains flat, liquid products load on own retention power, the short competitive frontier loads most clearly on rival structure, and intermediate contracts reflect mixtures of both. The broad pattern matches the framework's directional predictions, but not in a knife-edge one-product/one-channel way. That is precisely why the within-menu wedge evidence below is the sharper test of the mechanism.

## 5.2 Within-Menu Wedges: Compression, Screening, and Lock-In

The framework's sharpest prediction is relative, not scalar: within the same bank and quarter across geographies, local market structure should change the prices that matter for the allocation of elastic balances. The within-menu spreads are therefore not auxiliary evidence. They are the most direct test of menu rotation, because they ask whether the internal geometry of the menu moves with local concentration—inside the same institution, at the same time, under the same aggregate policy stance.

The MMA–vanilla 3-month spread compresses sharply with own concentration (Table 3, Panel C, column 1): the joint-specification coefficient on HHIC is $-0.230^{***}$ basis points per 100-basis-point policy move, while rival concentration is negligible. What does this mean concretely? When the Fed raises rates by 100 bps, a bank in a county at the 90th percentile of own concentration compresses the MMA–vanilla gap by

roughly 0.21 bps more than the same bank in a county at the 10th percentile. Over the cumulative 525 bps tightening cycle, that wedge compounds to approximately 1.1 bps of differential liquid-rate compression inside the same institution. Concentration therefore suppresses the return paid on liquid balances relative to the anchor inside the same bank. The MMA–savings spread behaves similarly ($\hat{\beta}^{\text{HHIC}} = -0.139^{**}$), indicating that the compression extends across the liquid tier. The distortion operates through the relative price of liquidity inside the bank's own contract menu: depositors holding liquid balances in concentrated markets receive systematically less compensation for their deposits, relative to the vanilla anchor, than depositors in the same bank's less concentrated markets.

The high-balance 3-month premium over the vanilla 3-month CD narrows with both own and rival concentration ($\hat{\beta}^{\text{OTH}} = -0.451^{**}$, $\hat{\beta}^{\text{HHIC}} = -0.821^{**}$). This is not a contradiction. The short high-balance tier sits exactly where the framework predicts overlapping forces: banks with greater local retention power have less reason to pay up for liquid high-balance balances, and banks facing a weaker rival set have less need to defend the frontier at the short end. The 3-month high-balance best rate is the product most vulnerable to immediate competitive switching by large-balance, rate-elastic depositors. When *either* dimension of local market power rises the bank can afford to pay less on this contract. This mixed loading is economically natural. It is exactly what one would expect for a short-horizon contract that lies between a pure liquid-retention margin and a pure frontier object.

The high-balance term slope is the cleanest rival-structure wedge (Table 3, Panel C, column 3). The 12-month minus 3-month high-balance spread widens by $2.708^{***}$ basis points per 100-basis-point policy move as rival concentration rises, with essentially no loading on own concentration. Concentration steepens the high-balance term structure by suppressing the short end—the contract most exposed to competitive

switching—more strongly than the long end, thereby raising the relative compensation for duration. This is lock-in through rival structure, not simply a common pass-through shortfall. A rate-sensitive depositor in a concentrated market faces a choice: accept a compressed short-term rate or commit to a 12-month contract to earn a rate closer to the policy rate. The steeper slope serves as a screening device that sorts depositors by their willingness to trade liquidity for yield.

The odd-term premium behaves differently. The odd-term minus best 12-month spread ($\hat{\beta}^{\text{HHIC}} = 1.595^{**}$) rises primarily with own concentration. This is consistent with the use of product complexity as a screening device: banks with greater local retention power widen the premium on non-standard maturities (7-month, 11-month, 13-month contracts) and sort attentive depositors into less comparable instruments. Because odd-term CDs cannot be easily compared across banks at standard benchmark maturities, the depositor who chooses an odd-term product has, in effect, opted out of the standard rate-comparison frontier. The bank can therefore set rates on these products with less direct competitive pressure. The odd-term level and the odd-term premium need not load on the same concentration margin, because the premium is defined relative to the bank's own best standard-maturity benchmark, which itself moves with local concentration. The acquisition premium over the vanilla 12-month rate ($\hat{\beta}^{\text{OTH}} = 0.740^{***}$), by contrast, loads more clearly on rival concentration, which is sensible for a contract designed to attract new money from the margin most directly exposed to outside offers.

These relative-price results are the core evidence in the paper. If concentration merely attenuated a common pass-through factor, the within-menu spreads would be approximately invariant to concentration. They are not. Liquid margins compress. High-balance slopes steepen. Acquisition and odd-term premia widen. The relative prices inside the same bank's menu change systematically across geographies. That

is menu rotation. The pattern is systematic rather than knife-edge: liquid spreads load most clearly on own concentration, the high-balance term slope and acquisition premium load most clearly on rival concentration, odd-term premia load more strongly on own concentration, and some short high-balance spreads load on both. That mix is exactly what one would expect if products differ in the extent to which they serve as anchors, frontier offers, or screening devices.

## 5.3 The Mechanism: Market-Level Frontier Suppression Dominates Own Markdowns

The preceding results show where concentration reshapes the menu, but they do not yet isolate the channel through which the depositor's outside option deteriorates. This subsection asks a narrower question. When concentration changes the rate a depositor faces, is that mainly because the focal bank marks down more aggressively relative to its rivals, or because the best rival offer in the market is itself lower?

To answer that question, Table 4 compares the bank's own posted rate with the best rival posted rate in the same local market. This leave-one-out comparison is a targeted diagnostic of the outside-option channel. It does not explain every coefficient in Tables 2 and 3. Rather, it asks, for selected consumer-facing products, whether concentration worsens the terms faced by depositors primarily through a lower local frontier or through larger unilateral markdowns by the focal bank.

Before turning to the decomposition, it helps to be explicit about the object being measured. For each product, I compare the bank's own posted rate with the best rival posted rate in the same local market. The resulting leave-one-out gap is the bank-versus-frontier gap, defined as own rate minus best rival rate. In levels this object is typically negative, because most banks sit below the best local rival offer. In

changes, a positive concentration coefficient does not mean that the focal bank falls farther behind the frontier. It means the gap becomes less negative because the rival frontier falls more than the bank's own rate. The consumer-facing harm is therefore a weaker outside option—a lower frontier—rather than necessarily a larger unilateral markdown by the focal bank.

For money market accounts (Table 4, Panel A, column 2), the decomposition is stark. The own-rate response to the concentration interaction is economically modest ($\hat{\gamma} = -0.100^{**}$): the focal bank does mark down slightly relative to where it would price in a less concentrated market. But the market frontier itself collapses ($\hat{\theta} = -3.350^{***}$), generating a leave-one-out gap of $\hat{\beta} = 3.250^{***}$. The frontier-suppression share is 97 percent. The dominant mechanism is not that a given bank marks down more aggressively than its rivals; it is that the frontier facing depositors—the best rate available from any rival in the same county—is lower to begin with.

The number-of-rivals control absorbs part of the gap. Including $N_{\text{rivals}} \times \Delta\text{FFR}$ reduces the gap coefficient, and the control itself is strongly negative. Fewer rivals mechanically suppress the frontier. But the residual concentration effect remains large and statistically overwhelming, indicating that the result is not merely a count-of-rivals effect.

The same logic appears across the products reported in Table 4, though the decomposition is most striking for the MMA and vanilla 12-month contracts. For the vanilla 12-month CD, the own-rate response is negligible ($\hat{\gamma} = -0.055$, insignificant), but the frontier drops sharply ($\hat{\theta} = -2.110^{***}$), generating a gap of $2.055^{***}$—again driven almost entirely by frontier suppression. For the 3-month high-balance best rate, the own-rate and frontier both move, but the frontier response is again larger. In every case, the outside option deteriorates more than the focal bank's own rate changes. The worsening terms depositors face reflect a lower local frontier rather than

larger unilateral markdowns by the focal bank.

Panel B enters rival and own concentration simultaneously in the gap equation; coefficients measure how each concentration margin changes the bank-versus-frontier gap. Unlike Panel A, Panel B does not separately decompose that gap into frontier suppression versus unilateral markdown.

Figure 3 summarizes the decomposition visually. For the products in Table 4, the frontier component is generally larger than the own-rate component, especially for the MMA and vanilla 12-month contract. The LOO exercise therefore sharpens the outside-option interpretation of the paper's consumer-facing gap results: when concentration changes the rates depositors face, it does so primarily because the local frontier is lower, not merely because the focal bank marks down more aggressively. This does not replace the own-concentration channel identified in Tables 2 and 3. Rather, it complements it. Own concentration still matters for the bank's own pricing of liquid balances and some within-menu wedges, while the LOO decomposition shows that the market-wide outside option is a quantitatively dominant margin in the rates depositors actually face.

## 5.4 Economic Magnitudes over the Hiking Cycle

The per-unit coefficients in Tables 2 and 3 capture the effect of a one-standard-deviation concentration difference on pass-through for a single 100-basis-point policy move. But the economically relevant quantity is the cumulative wedge over a full tightening cycle. Panel D of Table 3 performs this translation by multiplying the dominant joint-specification coefficient for each spread by the P10-to-P90 concentration gap and the cumulative 525-basis-point rate hike over 2022Q1–2024Q2.

The magnitudes are economically meaningful. The high-balance term slope widens

by an estimated 29.3 basis points between the 10th and 90th percentiles of rival concentration—nearly a third of a percentage point over the cycle. This is the accumulated price of lock-in: a rate-sensitive depositor in a concentrated market must commit to a 12-month high-balance CD to earn compensation that would be available at the 3-month maturity in a less concentrated market. The acquisition premium widens by 8.0 bps and the odd-term premium by 7.7 bps. On the liquid side, the MMA–vanilla spread compresses by 1.1 bps and the short high-balance premium by 4.0 bps as own concentration rises. Concentration therefore rotates the menu in an economically relevant way: it suppresses liquidity and short-horizon competition while steepening the reward to lock balances into duration.

Figure 2 traces the cumulative within-menu wedge quarter by quarter as the tightening cycle progresses. The wedge opens monotonically through the hiking phase, flattens once the federal funds rate reaches its plateau, and remains elevated through the holding period. The pattern confirms that the rotation accumulates gradually with policy tightening rather than reflecting a single quarter or an outlier.

## 5.5 Real Consequences for Funding Composition

If banks in concentrated markets compress liquid rates and widen screening margins, rate-elastic depositors should respond by reallocating balances. That reallocation need not happen instantaneously: depositors face switching costs, search frictions, and CD lock-in, and bank treasury functions adjust wholesale funding and balance-sheet composition with a lag. The concentration-exposure design therefore includes two lags of the $\Delta\text{FFR} \times z(\widetilde{HHI})$ interaction, and the cumulative effect $\sum_{s=0}^{2} \hat{\beta}_s$ is the primary object of interest.

Table 5 links the pricing distortions to bank-quarter outcomes using the Bartik

concentration-exposure design (Section 4.4), asking whether differential pre-sample exposure to concentrated markets predicts differential balance-sheet outcomes.

**Deposit flows and funding costs (Panel A).** Total deposit growth is essentially unresponsive to concentration exposure ($-0.024$, insignificant). This is consistent with the mechanism: menu rotation reallocates *which types* of deposits a bank retains even if total deposit growth moves little. But the composition shifts materially. High-balance deposit growth—the most rate-elastic component—falls cumulatively by $0.668^{**}$ percentage points per standard deviation of exposure. Core-deposit growth, by contrast, rises by $0.104^{**}$ percentage points. Concentrated banks retain price-insensitive, relationship-bound balances while losing large, rate-sensitive balances—exactly the balance-sheet footprint predicted by the pricing mechanism: compressed liquid rates and widened screening margins push elastic depositors out and retain inattentive core balances.

The funding-cost results are equally consistent. Deposit yields decline cumulatively by $1.068^{***}$ basis points, the cost of time CDs falls by $1.863^{***}$ basis points, and net interest margins widen by $1.508^{***}$ basis points. Concentrated banks pay less for their deposit base because the menu rotation compresses rates on the products that account for the bulk of retail balances. The menu-level pricing distortion translates directly into measurable profitability gains at the bank level.

**Deposit composition and time-deposit maturity (Panel B).** Deposit-composition shares tell the same story from a different angle. The core deposit share rises by $0.764^{**}$ percentage points cumulatively, the high-balance share falls by $0.333^{**}$ percentage points, and brokered deposits—a wholesale funding source that substitutes for lost rate-sensitive retail deposits—decline on impact, though the cumulative brokered-share

effect is imprecisely estimated. Within small-denomination time deposits—defined as time certificates with face value at or below $250,000, the FDIC insurance cap—short-maturity CD balances drop sharply: the contemporaneous effect on sub-one-year small time CDs is $-2{,}694^{***}$ thousand dollars per standard deviation of concentration exposure, with a cumulative decline of $-482^{***}$ thousand dollars. Longer-maturity small CDs also fall contemporaneously but recover cumulatively, consistent with banks steering remaining time-deposit balances into longer durations.

These are the realized balance-sheet consequences of menu rotation. Concentration changes who stays, who leaves, and which contract types absorb the adjustment. By compressing liquid and short-term rates while allowing screening margins to expand, concentrated banks push rate-sensitive, high-balance funding out and retain inattentive core balances. The profitability gain—a wider net interest margin—is the bank-level counterpart of the lower compensation received by depositors in concentrated markets.

## 5.6 Expected Policy versus Surprise; Tightening versus Easing

The total quarterly $\Delta \text{FFR}_t$ bundles three economically distinct components. Following Bauer and Swanson (2023), we decompose it into (i) an *anticipated* component already priced into fed-funds futures; (ii) a *pure monetary-policy surprise* orthogonal to macroeconomic news; and (iii) an *information* component capturing revised beliefs about fundamentals. Table 6 interacts each with both concentration margins.

This decomposition is a mechanism test, not merely a robustness check. If the concentration gradient reflects deliberate menu management, the anticipated component should load most cleanly: menu rotation is a planned response to the forward funding cost path, not a mechanical reaction to a single announcement. The pure surprise may

also matter but with less precision, because many menu adjustments require lead time. The information component is a useful benchmark—banks have less reason to reprice deposit menus in response to updated macro beliefs holding the policy rate fixed.

The anticipated component is indeed the strongest margin. For the MMA, the anticipated interaction with HHIC is $-0.067^{***}$, and the best 12-month CD shows the same pattern (anticipated HHIC = $-0.374^{**}$). The short-maturity high-balance best rate loads on anticipated OTH at $-0.367^{***}$. Pure-surprise effects are generally weak: none of the seven products show a statistically significant HHIC interaction, and only the odd-term and high-balance 12-month products load on OTH, consistent with menu rotation being planned rather than reflexive. The information component occupies a middle ground—it is narrower than the anticipated channel but not uniformly null. The MMA interaction is $-0.049^{***}$, suggesting that concentrated banks also condition liquid-product pricing on revised macro beliefs; the remaining products are insignificant, indicating that the information channel is concentrated in the own-concentration margin for liquid deposits rather than operating broadly across the menu. Overall, the decomposition supports forward-looking menu management: the anticipated component carries the main concentration gradient, pure surprises contribute little, and information effects are real but narrow.

**Tightening versus easing.** If concentration operated only through mechanical stickiness, one would expect symmetric sluggishness in both directions. If banks actively exploit captive depositor bases, the concentration effect should differ across hiking and easing phases: slower to raise during tightening, slower to lower during easing.

Table 8 confirms the latter. For the MMA, the HHIC interaction is $-0.116^{**}$ during tightening: concentrated banks raise MMA rates more slowly. During easing,

the coefficient reverses sign ($0.219^{**}$): concentrated banks also lower MMA rates more slowly. The Wald statistic rejects equality ($\chi^2 = 11.9$, $p < 0.001$). A purely mechanical sticky-rate interpretation would not imply that the concentration gradient changes so sharply with the direction of policy. The best 12-month CD and short-maturity high-balance best rate show parallel asymmetries, while vanilla CDs remain inert in both directions, confirming anchor invariance separately in hiking and easing phases.

## 5.7 Robustness and Identification

Table 7 addresses three threats to the baseline identification strategy. The first is that the horse-race interaction could be driven by arbitrary within-bank geographic sorting rather than local competitive conditions. The second is that concentration could predict rate changes even in the absence of a policy impulse, in which case the baseline interaction would be picking up a secular gradient rather than differential pass-through. The third is that concentration could proxy for county characteristics—especially population and income—that are themselves associated with depositor elasticity. The three diagnostics below show that the main results do not appear to be driven by any of these alternatives.

**Permutation test (Panel A).** The horse-race specification is identified from within-bank, within-quarter, cross-county variation in concentration. If the result were driven by arbitrary county assignment within the same bank-quarter rather than by local market structure, it should survive a placebo in which concentration is randomly reassigned across the bank's counties. To test this, I randomly reassign the concentration measures across a bank's counties within each bank-quarter, preserving the within-cell distribution of concentration while destroying its mapping to the actual county. The baseline specification is then re-estimated on the permuted interactions.

Permuted coefficients are small and insignificant across all products. The result therefore requires the correct county-level assignment and is not an artifact of within-bank geographic sorting.

**Zero-quarter test (Panel B).** The baseline is a differential pass-through design: it asks whether the response to policy differs with local concentration. A natural placebo is therefore to ask whether concentration predicts rate changes in quarters with essentially no policy movement. I restrict the sample to quarters with $|\Delta \text{FFR}_t| < 5$ basis points and enter the concentration measures directly, asking whether concentration has any systematic relationship with rate changes when there is no meaningful policy impulse to transmit. It does not. Coefficients are small and insignificant across products. The concentration gradient is therefore activated by policy movement rather than reflecting a general secular relationship between concentration and deposit-rate dynamics.

**Orthogonalization (Panel C).** The most substantive remaining concern is that county-level concentration may proxy for observable county characteristics that are themselves correlated with depositor elasticity. Large, high-income counties tend to have lower HHI and may also have more active rate shopping and stronger outside options. To address this concern, I orthogonalize both $\text{HHI}^{\text{OTH}}$ and HHIC with respect to county population and median household income, then re-standardize the residuals and re-estimate the baseline specification. The key coefficients survive. The HHIC coefficient on the MMA remains strongly negative, the rival-concentration coefficient on the short-maturity high-balance best rate remains strongly negative, and vanilla anchor rates remain insignificant. The main channels therefore do not appear to be driven by the most obvious county-level demographic gradient correlated with

concentration.

Taken together, the three diagnostics support the baseline interpretation. The result requires the correct county-level assignment, activates when there is a policy impulse to transmit, and survives purging concentration of the population-and-income gradient. Appendix B reports additional robustness checks, including state-by-quarter fixed effects and alternative dependent-variable definitions; the baseline results are stable.

## 5.8 Menu Architecture

Concentration reshapes menu *pricing* much more than menu *composition* (Table 9). Product shares—vanilla, odd-term, promotional, and high-balance—are largely invariant to the concentration interaction. The number of products on the menu does not change in an economically meaningful way, and the probability of offering a promotional, odd-term, or high-balance product is similarly unresponsive. The one margin that approaches significance is the within-menu rate range, consistent with the rotation mechanism narrowing overall rate dispersion as liquid products compress. The same menu is being priced differently—a repricing of the same menu rather than a change in the product set.

## 5.9 Interpretation

Across the results, four conclusions emerge. First, anchor products are locally inert, ruling out the view that concentration simply shifts all posted rates and providing an internal falsification against omitted local demand. Second, liquid products are governed primarily by own local retention power: the bank's captive deposit base, not the structure of the rival set, determines how much policy is transmitted through the

liquid margin. Third, the competitive frontier and the short end of the high-balance menu are governed primarily by rival structure and outside-option quality. Fourth, these pricing distortions have equilibrium consequences for the funding mix and bank profitability, shifting composition away from rate-elastic balances and toward captive core deposits.

The paper's claim is not that every product sorts cleanly into only one theoretical bucket. The claim is that within the same bank across geographies, local market structure changes the relative prices inside the deposit menu. Concentration compresses liquid compensation where retention is easiest, weakens the frontier where rivals compete less aggressively, and steepens the maturity structure offered to the most rate-elastic depositors. The competitive cost is predominantly a market-equilibrium effect—suppression of the outside option—rather than unilateral conduct by any single institution.

# 6 Conclusion

This paper studies how local deposit-market concentration shapes the transmission of monetary policy once the relevant pricing object is not a representative deposit rate, but the bank's posted contract menu. A large literature has shown that deposit-rate adjustment is incomplete and weaker in concentrated markets. A parallel literature has shown that large banks often post uniform rates across their geographic networks. Read through a scalar lens, those findings appear to point in opposite directions. The main argument of this paper is that the contradiction is only apparent. It arises from treating deposit pricing as if the relevant sufficient statistic were a single posted rate. Once banks are understood as pricing a menu, and households as choosing across that menu subject to heterogeneous frictions, both facts can hold at once.

The central empirical finding is that concentration does not uniformly reduce the response of all contracts. It rotates adjustment across the menu. Anchor contracts remain locally inert. Liquid products load on the bank's own local retention power. Competitive frontier offers load on the structure of rival competition. Screening instruments—odd-term premia, acquisition premia, and the term structure offered to high-balance, rate-elastic depositors—adjust in ways that reallocate compensation across liquidity, maturity, and comparability rather than simply lowering a single price. The relevant transmission margin is therefore not a scalar shortfall, but the structure of the menu and the wedges inside it.

This perspective also changes the competitive interpretation of concentration. The leave-one-out decomposition shows that the dominant channel is not an idiosyncratic markdown by the focal bank relative to an otherwise competitive market. It is suppression of the frontier itself. In concentrated markets, all banks move less on the contracts that matter most for rate-sensitive depositors. The outside option deteriorates as a market-equilibrium outcome. That distinction is first-order for merger review and antitrust policy in banking. If the competitive harm from concentration operated primarily through unilateral conduct, remedies aimed at the pricing behavior of a dominant bank might be sufficient. The evidence here instead suggests that the main distortion runs through the competitiveness of the rival set and the quality of the frontier available to depositors.

The results also imply that monetary policy is transmitted through deposits in a more segmented way than the conventional scalar view admits. The deposit menu determines which households bear the adjustment to policy and through which contracts. Captive depositors receive little adjustment on anchor and liquid products. More rate-elastic depositors are steered toward contracts that impose duration, complexity, or eligibility conditions rather than being paid uniformly at the margin. Deposit-side

transmission is therefore also a problem of screening, selection, and contract design. The shock decomposition suggests that menu adjustment is primarily forward-looking: anticipated policy changes matter most, pure-surprise effects are generally weak, and information effects are narrower and concentrated in the own-concentration margin.

The framework used in the paper is intentionally stylized. It is not a formal structural model of the full deposit market, and the paper does not claim to estimate all of the underlying primitives of deposit demand. Its contribution is different. It establishes a robust reduced-form fact with a clear economic structure: local market concentration reshapes monetary transmission by altering the competitive and screening margins of the deposit menu. That fact is visible only in data that observe the menu at the frequency and granularity at which banks actually price deposits.

The broader lesson is simple. Deposit pricing is not a one-price problem. Banks price menus, households choose across them under frictions, and market structure governs which contracts absorb policy adjustment. Once that is recognized, incomplete transmission and uniform posted anchors are not rival descriptions of the world. They are different margins of the same equilibrium pricing problem. The scalar view treats relative pricing inside the menu as secondary. That view is incorrect. Once depositors reallocate across contracts, frontier prices and within-menu premia become first-order determinants of who bears policy adjustment and what the bank pays for funding.

# References


Bauer, M.D. and E.T. Swanson (2023). A reassessment of monetary policy surprises and high-frequency identification. *NBER Macroeconomics Annual* 37(1), 87–155.

Begenau, J. and E. Stafford (2023). Uniform rate setting, branch expansion, and the deposit channel. *Harvard Business School Working Paper*, December.

Berger, A.N. and T.H. Hannan (1989). The price-concentration relationship in banking. *Review of Economics and Statistics* 71(2), 291–299.

Cameron, A.C., J.B. Gelbach, and D.L. Miller (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics* 29(2), 238–249.

d'Avernas, A., A.L. Eisfeldt, C. Huang, R. Stanton, and N. Wallace (2023). The deposit business at large vs. small banks. *NBER Working Paper* No. 31865.

Drechsler, I., A. Savov, and P. Schnabl (2017). The deposits channel of monetary policy. *Quarterly Journal of Economics* 132(4), 1819–1876.

Drechsler, I., A. Savov, and P. Schnabl (2021). Banking on deposits: Maturity transformation without interest rate risk. *Journal of Finance* 76(3), 1091–1143.

Egan, M., A. Hortaçsu, and G. Matvos (2017). Deposit competition and financial fragility: Evidence from the U.S. banking sector. *American Economic Review* 107(1), 169–216.

Egan, M., A. Hortaçsu, N.A. Kaplan, A. Sunderam, and V. Yao (2025). Dynamic competition for sleepy deposits. *NBER Working Paper* No. 34267.

Eichenbaum, M.S., F. Puglisi, S. Rebelo, and M. Trabandt (2025). Banks and the state-dependent effects of monetary policy. *NBER Working Paper* No. 33523.

Enkhbold, A. (2023). Monetary policy transmission, bank market power, and wholesale funding reliance. *Bank of Canada Staff Working Paper* 2023–35.

Gabaix, X. and D. Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121(2), 505–540.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review* 110(8), 2586–2624.

Granja, J. and N. Paixão (2026). Bank consolidation and uniform pricing. *Journal of Financial Economics* 176, 104204.

Hannan, T.H. and A.N. Berger (1991). The rigidity of prices: Evidence from the banking industry. *American Economic Review* 81(4), 938–945.

Hoffmann, P., S. Langfield, F. Pierobon, and G. Vuillemey (2019). Who bears interest rate risk? *Review of Financial Studies* 32(8), 2921–2954.

Hortaçsu, A. and C. Syverson (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. *Quarterly Journal of Economics* 119(2), 403–456.

Mussa, M. and S. Rosen (1978). Monopoly and product quality. *Journal of Economic Theory* 18(2), 301–317.

Neumark, D. and S.A. Sharpe (1992). Market structure and the nature of price rigidity: Evidence from the market for consumer deposits. *Quarterly Journal of Economics* 107(2), 657–680.

Petersen, M.A. (2009). Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies* 22(1), 435–480.

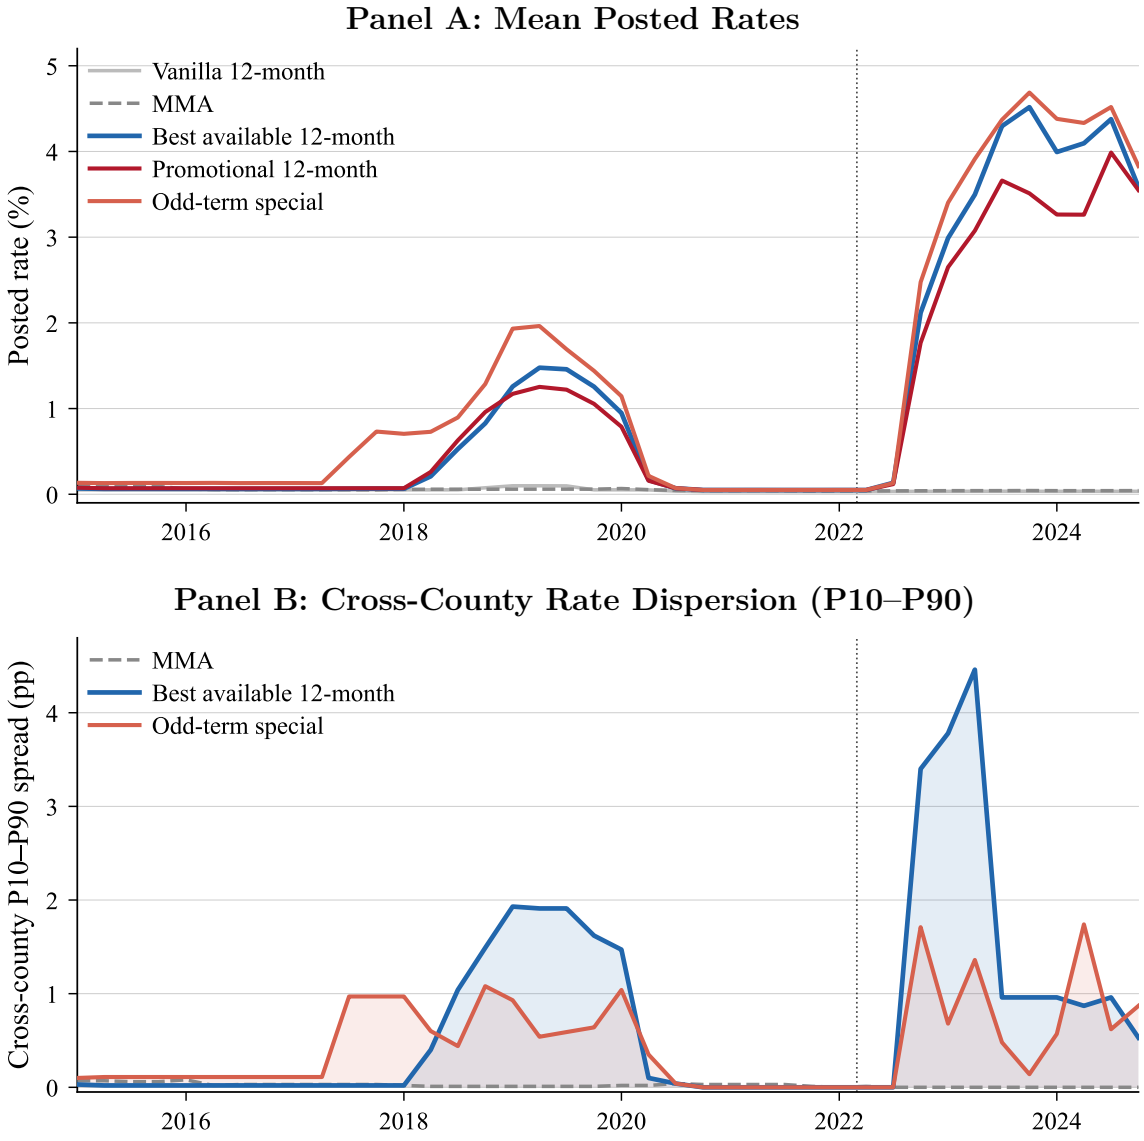
Radecki, L.J. (1998). The expanding geographic reach of retail banking markets. *Federal Reserve Bank of New York Economic Policy Review* 4(2), 15–34.

Tirole, J. (1988). *The Theory of Industrial Organization.* MIT Press, Cambridge, MA.

Wang, Y., T.M. Whited, Y. Wu, and K. Xiao (2022). Bank market power and monetary policy transmission: Evidence from a structural estimation. *Journal of Finance* 77(4), 2093–2141.

Yankov, V. (2024). In search of a risk-free asset: Search costs and sticky deposit rates. *Journal of Money, Credit and Banking* 56(5), 1053–1098.

Figure 1: Deposit Menu Pricing at Bank of America during the 2022–2023 Tightening Cycle




*Notes:* The figure provides an illustrative case study of posted deposit rates at Bank of America, a large multi-market institution. Panel A plots mean posted rates across all counties served by the bank for each deposit product. Anchor rates (the money market account and the vanilla 12-month rack CD) remain nearly flat over time despite 525 basis points of Federal Reserve tightening. Competitive-frontier products (the best available 12-month CD, promotional specials, and odd-maturity contracts) rise sharply. Panel B plots the 10th-to-90th percentile cross-county spread for each product, measuring geographic price dispersion within the same bank. Anchor products show virtually no cross-county variation, while competitive instruments fan out sharply after liftoff. The contrast illustrates the menu-rotation mechanism documented in the paper: anchor contracts remain uniformly posted across markets while screening and acquisition instruments absorb local competitive pressure. The figure is descriptive rather than part of the identification strategy.

Table 1: Summary Statistics

| | Mean | SD | P10 | Median | P90 | $N$ |
|---|---|---|---|---|---|---|
| **Panel A: Bank–Quarter** | | | | | | |
| Vanilla 12m CD rate (pp) | 0.792 | 1.007 | 0.107 | 0.399 | 2.129 | 38,372 |
| Best available 12m CD rate (pp) | 1.128 | 1.280 | 0.150 | 0.500 | 3.456 | 39,966 |
| Money market account rate (pp) | 0.574 | 0.799 | 0.073 | 0.250 | 1.530 | 38,550 |
| High-balance 3m best rate (pp) | 0.478 | 0.860 | 0.050 | 0.176 | 1.239 | 9,960 |
| High-balance 12m best rate (pp) | 0.940 | 1.152 | 0.144 | 0.449 | 2.783 | 11,583 |
| Odd-term best rate (pp) | 1.381 | 1.371 | 0.150 | 0.850 | 3.922 | 33,313 |
| Vanilla share of CD menu | 0.769 | 0.229 | 0.435 | 0.833 | 1.000 | 42,255 |
| Odd-term share of CD menu | 0.204 | 0.174 | 0.000 | 0.174 | 0.450 | 42,255 |
| High-balance CD share of CD menu | 0.039 | 0.115 | 0.000 | 0.000 | 0.180 | 42,255 |
| Menu dispersion (pp) | 0.539 | 0.428 | 0.130 | 0.405 | 1.226 | 42,333 |
| **Panel B: County–Quarter** | | | | | | |
| HHI (×100) | 26.20 | 17.10 | 10.75 | 21.67 | 47.33 | 105,402 |
| Own concentration, HHIC (×100) | 6.87 | 15.32 | 0.43 | 2.15 | 14.50 | 105,159 |
| Rival concentration, $\text{HHI}^{\text{OTH}}$ (×100) | 19.23 | 11.62 | 8.81 | 16.92 | 32.68 | 104,843 |
| Number of banks in county | 12.28 | 14.26 | 3 | 8 | 25 | 105,402 |
| LOO gap: MMA (pp) | −0.573 | 0.726 | −1.754 | −0.250 | 0.000 | 90,193 |
| LOO gap: Best 12m CD (pp) | −0.534 | 0.586 | −1.403 | −0.317 | 0.000 | 90,209 |
| LOO gap: Vanilla 12m CD (pp) | −0.439 | 0.712 | −1.386 | −0.123 | 0.000 | 81,289 |
| County deposits ($B) | 5.784 | 39.19 | 0.167 | 0.718 | 6.769 | 105,402 |
| Population | 113,026 | 346,409 | 6,808 | 29,438 | 234,671 | 114,689 |
| Median household income ($) | 57,564 | 16,438 | 39,990 | 54,805 | 78,192 | 114,689 |
| **Panel C: Bank–County–Quarter (Regression Level)** | | | | | | |
| MMA − vanilla 3m spread (pp) | 0.255 | 1.107 | −0.449 | 0.010 | 1.400 | 354,720 |
| High-balance term slope, 12m−3m (pp) | 0.501 | 1.208 | 0.000 | 0.100 | 1.950 | 193,797 |
| Odd-term − best 12m spread (pp) | 0.360 | 0.828 | −0.050 | 0.050 | 1.250 | 463,218 |
| Within-bank rate range (pp) | 1.429 | 1.410 | 0.050 | 0.990 | 4.003 | 671,531 |
| Menu dispersion (pp) | 0.506 | 0.465 | 0.060 | 0.354 | 1.269 | 675,891 |
| **Panel D: Quarterly Aggregates** | | | | | | |
| Federal funds rate (bps) | 198.5 | 190.4 | 25.0 | 137.5 | 537.5 | 40 |
| Δ FFR (bps) | 10.9 | 48.7 | −50.0 | 0.0 | 50.0 | 39 |
| Anticipated ΔFFR (bps) | 10.3 | 43.8 | −32.7 | 6.2 | 35.9 | 39 |
| Pure surprise (bps) | 1.5 | 5.3 | −4.9 | 0.0 | 8.2 | 40 |
| Information shock (bps) | −1.1 | 7.2 | −3.6 | 0.0 | 2.7 | 40 |
| **Panel E: Bank–Quarter UBPR Outcomes** | | | | | | |
| Deposit yield (bps) | 80.3 | 82.8 | 11.0 | 53.0 | 198.0 | 42,088 |
| Cost of time CDs (bps) | 142.7 | 115.6 | 25.0 | 111.0 | 338.0 | 42,082 |
| Net interest margin (bps) | 299.4 | 115.0 | 64.0 | 316.0 | 399.0 | 42,091 |
| Core deposit share (%) | 91.16 | 8.41 | 81.76 | 93.49 | 98.39 | 42,091 |
| High-balance deposit share (%) | 4.06 | 4.01 | 0.72 | 2.83 | 8.80 | 42,091 |
| Brokered deposit share (%) | 3.53 | 6.66 | 0.00 | 0.49 | 10.36 | 42,091 |
| Small time CD ≤1yr ($M) | 495.2 | 3,067.0 | 21.9 | 97.0 | 703.5 | 42,091 |
| Small time CD >1yr ($M) | 180.3 | 959.3 | 7.0 | 43.2 | 260.9 | 42,091 |

*Notes:* Panel A reports bank–quarter–level deposit rates and menu-composition variables. Rates are in percentage points; shares are fractions of the bank's CD product count in a given county–quarter. Panel B reports county–quarter market structure variables and leave-one-out (LOO) gap measures. HHI, HHIC, and $\text{HHI}^{\text{OTH}}$ are scaled so that 100 corresponds to a monopoly. The LOO gap equals the bank's own rate minus the best rival rate in the county. Panel C reports within-menu spread, slope, and rate-range measures at the regression level. Panel D reports quarterly monetary policy variables in basis points: the federal funds rate, its first difference, and the Bauer–Swanson shock decomposition into anticipated, pure-surprise, and information components. Meeting-level surprises are summed within quarter before decomposing each quarterly ΔFFR into anticipated, information, and pure-surprise components. Panel E reports bank–quarter UBPR balance-sheet variables used in the concentration-exposure tests (Table 5): annualized yields and costs in basis points, deposit-composition shares as percentages of total deposits, and small-denomination time-deposit balances in millions of dollars (≤$250,000 face value; UBPRM023 = remaining maturity under one year, UBPRM024 = over one year). The sample spans 2015Q1–2024Q4 (40 quarters). Raw data coverage includes 2,930 counties; the regression sample covers roughly 2,775 counties after imposing the product-specific estimation restrictions (approximately 1,600 unique banks).

Table 2: Product-Level Pass-Through of Monetary Policy by Market Concentration

| | Van 3m (1) | Van 12m (2) | MMA (3) | Best 12m (4) | HB 3m (5) | Odd-Term (6) | HB 12m (7) |
|---|---|---|---|---|---|---|---|
| **Panel A: Total HHI** | | | | | | | |
| $z_{\text{HHI}} \times \Delta\text{FFR}$ | −0.019 | −0.080 | −0.139*** | −0.685** | −0.910*** | 0.365 | 1.172* |
| | (0.051) | (0.124) | (0.042) | (0.305) | (0.301) | (0.299) | (0.619) |
| **Panel B: Own Concentration (HHIC)** | | | | | | | |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | 0.004 | −0.134 | −0.161*** | −0.593** | −0.098 | −0.113 | −0.149 |
| | (0.045) | (0.132) | (0.045) | (0.287) | (0.355) | (0.326) | (0.920) |
| **Panel C: Horse Race (Rival and Own Concentration Jointly)** | | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | −0.033 | −0.015 | −0.041 | −0.425 | −0.778*** | 0.539* | 1.182* |
| | (0.054) | (0.112) | (0.034) | (0.318) | (0.222) | (0.275) | (0.640) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | 0.003 | −0.147 | −0.156*** | −0.663** | −0.345 | −0.002 | 0.124 |
| | (0.046) | (0.140) | (0.044) | (0.292) | (0.373) | (0.338) | (0.961) |
| Bank × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Bank × County | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| MSA × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Within $R^2$ | 0.3040 | 0.1080 | 0.6570 | 0.1070 | 0.1410 | 0.1290 | 0.0650 |
| Observations | 443,145 | 487,762 | 439,461 | 505,806 | 167,002 | 460,357 | 173,987 |
| Bank clusters | 1,406 | 1,589 | 1,509 | 1,599 | 393 | 1,427 | 456 |
| County clusters | 2,740 | 2,773 | 2,719 | 2,774 | 1,075 | 2,730 | 1,084 |

*Notes:* Each coefficient reports the interaction of standardized local market concentration with the quarterly change in the federal funds rate (ΔFFR). The dependent variable is the quarterly change in the posted rate for the deposit product indicated in the column header, measured in basis points using the product-specific quarterly transformation described in Section 4.4. Panel A uses total HHI. Panel B uses own concentration (HHIC). Panel C includes both rival concentration ($z_{\text{HHI}^{\text{OTH}}}$) and own concentration ($z_{\text{HHIC}}$) simultaneously. The coefficient therefore measures how a common policy-rate change loads differentially on local concentration within the same bank and quarter. The level effect of $\Delta\text{FFR}_t$ is absorbed by bank × quarter fixed effects. MSA × Quarter fixed effects use Combined-MSA (CMSA) definitions for columns (1)–(4) and (6), and standard MSA definitions for high-balance products in columns (5) and (7). CMSA markets assign rural counties to state-level groups. Standard errors in parentheses, clustered two-way by bank and county. $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

Table 3: Within-Menu Spread Responses to Monetary Policy by Market Concentration

| | MMA − Van 3m (1) | HB 3m − Van 3m (2) | HB Slope (12m−3m) (3) | Odd − Best 12m (4) | Acq − Van 12m (5) | MMA − Savings (6) |
|---|---|---|---|---|---|---|
| **Panel A: Total HHI** | | | | | | |
| $z_{\text{HHI}} \times \Delta\text{FFR}$ | $-0.188^{***}$ | $-0.594^{**}$ | $2.798^{***}$ | $1.177^{**}$ | $0.680^{***}$ | $-0.145^{***}$ |
| | (0.052) | (0.262) | (0.606) | (0.496) | (0.220) | (0.043) |
| **Panel B: Own Concentration (HHIC)** | | | | | | |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | $-0.234^{***}$ | $-0.664^{*}$ | $-0.017$ | $1.466^{**}$ | $0.087$ | $-0.134^{**}$ |
| | (0.062) | (0.372) | (1.076) | (0.619) | (0.254) | (0.057) |
| **Panel C: Horse Race** | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | $-0.049$ | $-0.451^{**}$ | $2.708^{***}$ | $0.674$ | $0.740^{***}$ | $-0.058^{*}$ |
| | (0.039) | (0.230) | (0.575) | (0.437) | (0.200) | (0.033) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | $-0.230^{***}$ | $-0.821^{**}$ | $0.595$ | $1.595^{**}$ | $0.276$ | $-0.139^{**}$ |
| | (0.064) | (0.393) | (1.181) | (0.625) | (0.276) | (0.057) |
| Bank × Quarter | Yes | Yes | Yes | Yes | Yes | Yes |
| Bank × County | Yes | Yes | Yes | Yes | Yes | Yes |
| MSA × Quarter | Yes | Yes | Yes | Yes | Yes | Yes |
| Within $R^2$ | 0.6760 | 0.0360 | 0.0320 | 0.0500 | $0.2110^{\dagger}$ | 0.6510 |
| Observations | 402,209 | 141,646 | 148,099 | 299,664 | 169,181 | 434,299 |
| Bank clusters | 1,320 | 329 | 350 | 1,142 | 795 | 1,484 |
| County clusters | 2,672 | 1,065 | 1,066 | 1,184 | 2,586 | 2,714 |

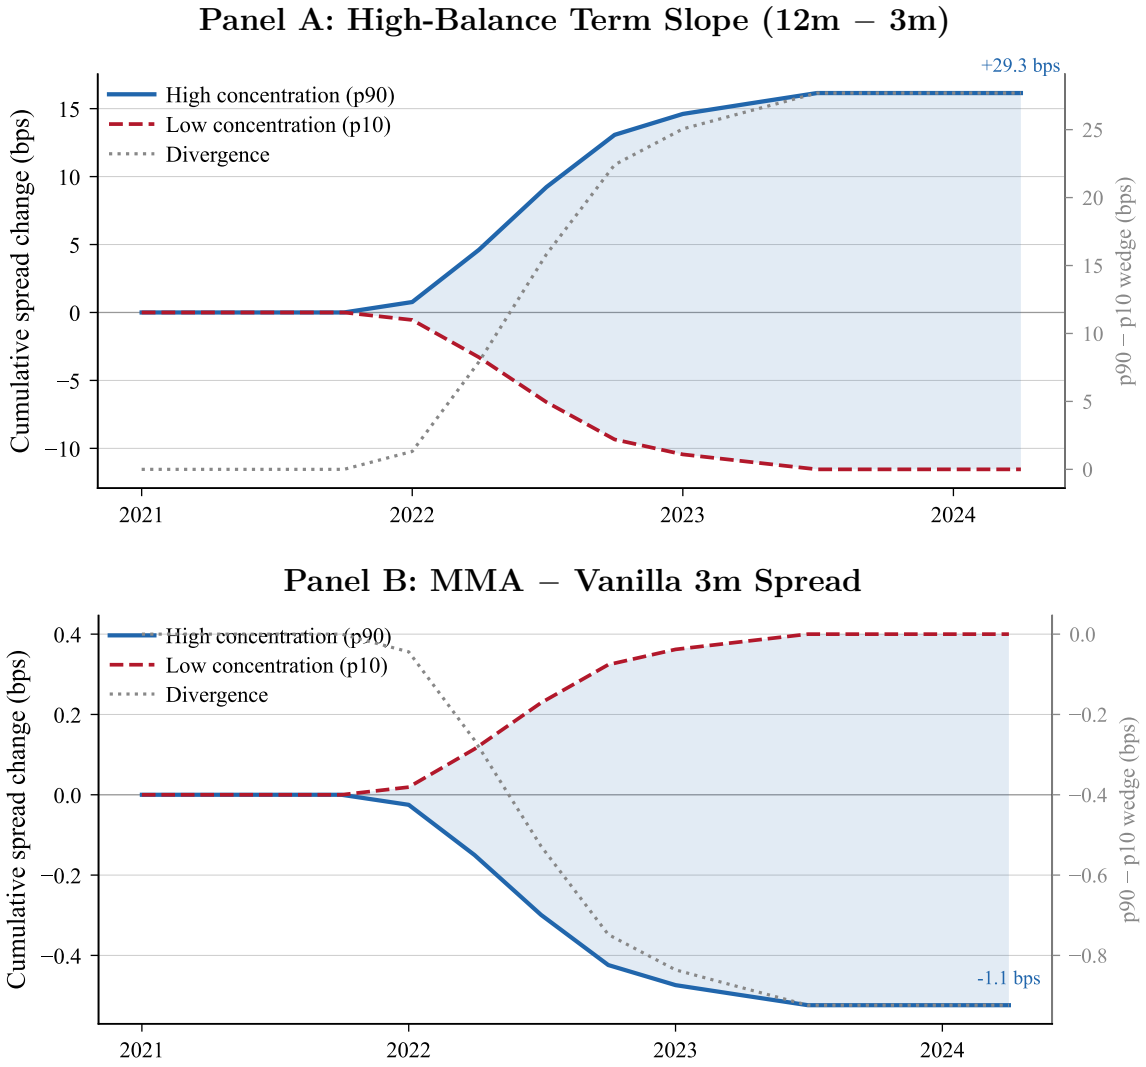
**Panel D: Implied Cumulative Spread Wedge, P10 to P90 of County-Level Concentration**

| Spread | Channel | $\hat{\beta}$ (bps/pp) | SE | $t$ | $z^{p10}$ | $z^{p90}$ | Cum. wedge (bps) |
|---|---|---|---|---|---|---|---|
| MMA − Van 3m | HHIC | $-0.230^{***}$ | 0.064 | −3.62 | −0.42 | 0.50 | −1.1 |
| HB 3m − Van 3m | HHIC | $-0.821^{**}$ | 0.393 | −2.09 | −0.42 | 0.50 | −4.0 |
| HB slope (12m−3m) | $\text{HHI}^{\text{OTH}}$ | $2.708^{***}$ | 0.575 | 4.71 | −0.90 | 1.16 | +29.3 |
| Odd − Best 12m | HHIC | $1.595^{**}$ | 0.625 | 2.55 | −0.42 | 0.50 | +7.7 |
| Acq − Van 12m | $\text{HHI}^{\text{OTH}}$ | $0.740^{***}$ | 0.200 | 3.70 | −0.90 | 1.16 | +8.0 |
| MMA − Savings | HHIC | $-0.139^{**}$ | 0.057 | −2.47 | −0.42 | 0.50 | −0.7 |

Cumulative wedge $= \hat{\beta} \times (z^{p90} - z^{p10}) \times \sum_t \Delta\text{FFR}_t$, where $\sum_t \Delta\text{FFR}_t = 5.25$ pp over 2022Q1–2024Q2. Coefficients are from Panel C (horse race). Each row uses the concentration margin on which the corresponding wedge loads most clearly in the joint specification. The objective is descriptive magnitude, not a structural welfare calculation.

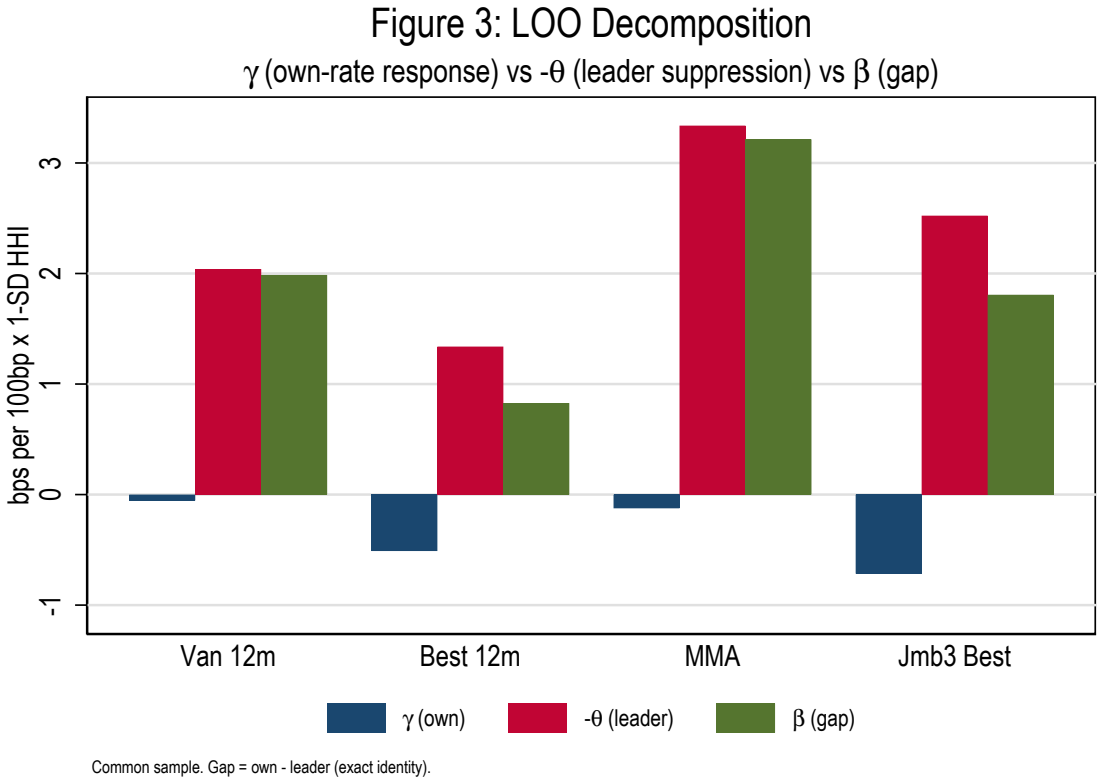
*Notes:* Panels A–C report the interaction of standardized concentration with ΔFFR. The dependent variable is the quarterly change in the within-menu spread indicated in the column header, measured in basis points. "HB Slope" is the 12-month minus 3-month high-balance best-rate spread. "Odd − Best 12m" is the odd-term premium over the best standard-maturity 12-month CD. "Acq − Van 12m" is the teaser premium on acquisition CDs over the vanilla 12-month rate. As in Table 2, the coefficient measures how a common policy-rate change loads differentially on local concentration within the same bank and quarter. MSA × Quarter uses CMSA definitions for columns (1), (4)–(6) and standard MSA for columns (2)–(3). Panel D translates the horse-race coefficients from Panel C into implied cumulative spread wedges over the 2022Q1–2024Q2 hiking cycle. Standard errors in parentheses; clustered two-way by bank and county. $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$. $^{\dagger}$ Simple $R^2$ reported in place of within $R^2$, which was $\leq 0$ in this specification.

Figure 2: Cumulative Within-Menu Wedges by Market Concentration, 2021Q1–2024Q2




*Notes:* Each panel plots the predicted cumulative within-menu spread at the 90th percentile (solid line) and 10th percentile (dashed line) of the relevant county-level concentration distribution. The shaded region between the two lines visualizes the concentration-driven wedge: it measures how much the cumulative spread diverges between a high-concentration and a low-concentration county as policy tightening accumulates. The wedge itself is also plotted on the right axis (short-dashed line). Formally, the wedge equals $\hat{\beta} \times (z^{p90} - z^{p10}) \times \sum_{s \leq t} \Delta\text{FFR}_s$, where $\sum_t \Delta\text{FFR}_t = 5.25$ pp over the 2022Q1–2024Q2 hiking cycle. Coefficients are from the horse-race specification (Table 3, Panel C). Panel A uses the rival-concentration coefficient on the high-balance 12m−3m slope; Panel B uses the own-concentration coefficient on the MMA−vanilla 3m spread. In concentrated markets, the high-balance term slope steepens (the shaded region fans out upward) while the liquid spread compresses (the shaded region fans out downward). Implied magnitudes correspond to Table 3, Panel D.

Figure 3: Leave-One-Out Decomposition: Own Rate versus Frontier Components




*Notes:* For each product, the figure displays the three components of the leave-one-out decomposition: the own-rate response ($\hat{\gamma}$), the negative of the frontier response ($-\hat{\theta}$), and the resulting gap ($\hat{\beta} = \hat{\gamma} - \hat{\theta}$). In every case, the frontier response dominates the own-rate response, indicating that the competitive distortion from concentration operates primarily through suppression of the outside option rather than unilateral markdowns by the focal bank. Estimates correspond to Table 4, Panel A.

Table 4: Leave-One-Out Frontier Decomposition

| | Van 12m (1) | MMA (2) | Best 12m (3) | HB 3m (4) |
|---|---|---|---|---|
| **Panel A: Total HHI Decomposition** | | | | |
| Own rate ($\hat{\gamma}$) | $-0.055$ | $-0.100^{**}$ | $-0.512$ | $-0.657^{***}$ |
| | (0.140) | (0.041) | (0.347) | (0.253) |
| Frontier rate ($\hat{\theta}$) | $-2.110^{***}$ | $-3.350^{***}$ | $-1.311^{***}$ | $-2.531^{***}$ |
| | (0.430) | (0.292) | (0.380) | (0.485) |
| LOO gap ($\hat{\beta} = \hat{\gamma} - \hat{\theta}$) | $2.055^{***}$ | $3.250^{***}$ | $0.799^{*}$ | $1.875^{***}$ |
| | (0.416) | (0.283) | (0.464) | (0.525) |
| **Panel B: Horse-Race Gap Decomposition** | | | | |
| Gap $\times$ $z_{\text{HHI}^{\text{OTH}}}$ | $1.520^{***}$ | $2.505^{***}$ | $0.754^{*}$ | $1.352^{***}$ |
| | (0.401) | (0.263) | (0.431) | (0.521) |
| Gap $\times$ $z_{\text{HHIC}}$ | $1.925^{***}$ | $2.751^{***}$ | 0.384 | $2.070^{***}$ |
| | (0.362) | (0.194) | (0.546) | (0.632) |
| **Panel C: Rival-Count Control** | | | | |
| LOO gap (baseline, from Panel A) | $2.055^{***}$ | $3.250^{***}$ | $0.799^{*}$ | $1.875^{***}$ |
| | (0.416) | (0.283) | (0.464) | (0.525) |
| LOO gap (controlling $N_{\text{rivals}} \times \Delta\text{FFR}$) | $1.326^{***}$ | $2.734^{***}$ | 0.560 | $1.395^{**}$ |
| | (0.405) | (0.264) | (0.473) | (0.544) |
| $N_{\text{rivals}} \times \Delta\text{FFR}$ | $-1.127^{***}$ | $-0.580^{***}$ | $-0.278^{**}$ | $-1.953^{***}$ |
| | (0.175) | (0.060) | (0.139) | (0.332) |
| Bank $\times$ Quarter | Yes | Yes | Yes | Yes |
| Bank $\times$ County | Yes | Yes | Yes | Yes |
| MSA $\times$ Quarter | Yes | Yes | Yes | Yes |
| Within $R^2$ (gap eq.) | 0.0360 | 0.5410 | 0.0630 | 0.0610 |
| Observations | 471,842 | 428,754 | 494,533 | 217,305 |
| Bank clusters | 1,568 | 1,500 | 1,585 | 430 |
| County clusters | 2,546 | 2,476 | 2,553 | 1,899 |

*Notes:* Panel A reports the own-rate response ($\hat{\gamma}$), the leave-one-out frontier response ($\hat{\theta}$), and the implied gap response ($\hat{\beta} = \hat{\gamma} - \hat{\theta}$). Because the gap is defined as own rate minus best rival rate, a positive gap coefficient means the gap becomes less negative because the frontier falls more than the bank's own rate. Panel B enters rival and own concentration simultaneously in the gap equation; coefficients measure how each concentration margin changes the bank-versus-frontier gap. Unlike Panel A, Panel B does not separately decompose that gap into frontier suppression versus unilateral markdown. Panel C adds the number of rivals interacted with $\Delta$FFR to absorb the mechanical effect of having fewer competitors. MSA $\times$ Quarter uses CMSA definitions for all columns. Standard errors in parentheses; clustered two-way by bank and county. $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

Table 5: Real Balance-Sheet Consequences: Bartik Concentration-Exposure Design

**Panel A: Bank-Quarter Deposit Flows and Yield/Cost Outcomes**

| | Total Dep. (×100) (1) | High-Bal Dep. (×100) (2) | Core Dep. (×100) (3) | Brokered Dep. (×100) (4) | Dep. Yield (bps) (5) | CD Cost (bps) (6) | NIM (bps) (7) |
|---|---|---|---|---|---|---|---|
| Contemporaneous ($\hat{\beta}_0$) | 0.031 | −0.004 | 0.043 | 2.063 | −0.439 | −0.742 | 2.069* |
| | (0.043) | (0.256) | (0.045) | (2.049) | (0.335) | (0.665) | (1.155) |
| Cumulative ($\sum_{s=0}^{2} \hat{\beta}_s$) | −0.024 | −0.668** | 0.104** | 0.756 | −1.068*** | −1.863*** | 1.508*** |
| | (0.044) | (0.266) | (0.048) | (1.545) | (0.315) | (0.376) | (0.535) |
| Lags 1–2 included | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Bank FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Quarter FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Within $R^2$ | 0.1810 | 0.0540 | 0.2190 | 0.0340† | 0.1250 | 0.0960 | 0.0090† |
| Observations | 37,117 | 37,117 | 37,117 | 37,117 | 37,113 | 37,111 | 37,117 |
| Bank clusters | 1,475 | 1,475 | 1,475 | 1,475 | 1,475 | 1,474 | 1,475 |

**Panel B: Deposit Composition Shares and Small-Denomination Time-Deposit Volumes**

| | Core | High-Bal | Brokered | Raw Δ ($000s) | | Growth Rate (%) | |
|---|---|---|---|---|---|---|---|
| | Share (pp) (1) | Share (pp) (2) | Share (pp) (3) | Sm. CD <1yr (4) | Sm. CD >1yr (5) | Sm. CD <1yr (6) | Sm. CD >1yr (7) |
| Contemporaneous | 5.634*** | −1.190** | −2.994*** | −2,694*** | −972*** | −0.505*** | −0.613*** |
| | (1.324) | (0.585) | (1.047) | (627) | (196) | (0.131) | (0.162) |
| Cumulative | 0.764** | −0.333** | −0.350 | −482*** | 3.6 | −0.030 | 0.016 |
| | (0.309) | (0.136) | (0.261) | (146) | (47.0) | (0.028) | (0.043) |
| Bank FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Quarter FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Within $R^2$ | 0.0770 | 0.0870 | 0.0150 | 0.0600 | 0.0480 | 0.1110 | 0.0110 |
| Observations | 39,931 | 39,931 | 39,931 | 39,931 | 39,931 | 39,931 | 39,918 |
| Bank clusters | 1,621 | 1,621 | 1,621 | 1,621 | 1,621 | 1,621 | 1,621 |

*Notes:* Both panels use the Bartik-style concentration-exposure design described in Section 4.4, which interacts predetermined geographic footprint shares from the 2010–2014 Summary of Deposits with local concentration and the national policy cycle. Panel A reports contemporaneous and cumulative effects on bank-quarter deposit volumes and yield/cost measures from the UBPR; deposit-volume coefficients (columns 1–4) are $\Delta \log(1 + \text{volume})$ multiplied by 100, so they approximate percentage-point growth rates. Columns 5–7 are first-differences of UBPR annualized rates in basis points: deposit yield (UBPRE106), cost of time CDs (UBPRHR55), and net interest margin (UBPRKX40). Lags 1 and 2 are included but not reported. Panel B shows contemporaneous and cumulative effects on UBPR deposit-composition shares (columns 1–3: constructed as category/total deposits, first-differenced, in percentage points) and small-denomination time-deposit volumes (columns 4–7). "Small" time deposits are those with face value ≤$250,000 (the FDIC insurance cap): UBPRM023 is time deposits ≤$250K with remaining maturity under one year; UBPRM024 is time deposits ≤$250K with remaining maturity over one year. Columns 4–5 report the raw quarterly first-difference in thousands of dollars; columns 6–7 report $\Delta \log(1 + \text{volume}) \times 100$, so those coefficients read directly as percentage-point quarterly growth rates (e.g., a coefficient of −0.505 means that a one-standard-deviation increase in Bartik-weighted concentration reduces quarterly volume growth by 0.505 percentage points). Standard errors in parentheses; clustered by bank. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$. † Simple $R^2$ reported in place of within $R^2$, which was ≤ 0 in this specification.

Table 6: Shock Decomposition: Anticipated, Pure-Surprise, and Information Components

| | Van 3m (1) | Van 12m (2) | MMA (3) | Best 12m (4) | HB 3m (5) | Odd-Term (6) | HB 12m (7) |
|---|---|---|---|---|---|---|---|
| *Anticipated Component* | | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times$ Antic | −0.026 | −0.032 | −0.025 | −0.259 | −0.367*** | 0.275** | 0.624** |
| | (0.027) | (0.053) | (0.018) | (0.160) | (0.098) | (0.134) | (0.305) |
| $z_{\text{HHIC}} \times$ Antic | −0.001 | −0.077 | −0.067*** | −0.374** | −0.179 | 0.037 | 0.035 |
| | (0.022) | (0.064) | (0.022) | (0.146) | (0.153) | (0.167) | (0.162) |
| *Pure Surprise Component* | | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times$ Surp | 0.032 | −0.017 | 0.026 | 0.221 | −0.047 | 0.048 | −0.079 |
| | (0.043) | (0.057) | (0.029) | (0.206) | (0.199) | (0.206) | (0.285) |
| $z_{\text{HHIC}} \times$ Surp | 0.053 | 0.039 | −0.027 | 0.415* | −0.178 | 0.010 | 0.328 |
| | (0.033) | (0.064) | (0.027) | (0.215) | (0.291) | (0.186) | (0.349) |
| *Information Component* | | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times$ Info | 0.016 | 0.056 | −0.020 | −0.010 | −0.120 | −0.063 | −0.064 |
| | (0.020) | (0.042) | (0.023) | (0.100) | (0.074) | (0.088) | (0.261) |
| $z_{\text{HHIC}} \times$ Info | 0.008 | −0.082 | −0.049*** | −0.199* | −0.153 | −0.184* | −0.291* |
| | (0.017) | (0.084) | (0.019) | (0.120) | (0.108) | (0.096) | (0.176) |
| Bank × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Bank × County | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| MSA × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Within $R^2$ | 0.3040 | 0.1080 | 0.6570 | 0.1070 | 0.1410 | 0.1290 | 0.0650 |
| Observations | 443,145 | 487,762 | 439,461 | 505,806 | 167,002 | 460,357 | 173,987 |
| Bank clusters | 1,406 | 1,589 | 1,509 | 1,599 | 393 | 1,427 | 456 |
| County clusters | 2,740 | 2,773 | 2,719 | 2,774 | 1,075 | 2,730 | 1,084 |

*Notes:* The quarterly change in the federal funds rate is decomposed into anticipated, pure-surprise, and information components following Bauer and Swanson (2023). Meeting-level surprises are summed within quarter; the anticipated component is then $\Delta\text{FFR}_t$ minus the quarterly raw-surprise sum, the pure surprise is the quarterly orthogonalized-surprise sum, and the information shock is the difference between the raw and orthogonalized sums. Each coefficient reports the interaction of the indicated shock component with the relevant standardized concentration measure. Within each shock component, both $z_{\text{HHI}^{\text{OTH}}}$ and $z_{\text{HHIC}}$ are entered simultaneously (horse-race specification); the three shock-component regressions are estimated separately. Fixed effects and controls are otherwise identical to Table 2, Panel C. MSA × Quarter uses CMSA definitions for columns (1)–(4) and (6), and standard MSA definitions for high-balance products in columns (5) and (7). Within $R^2$ and $N$ are from the anticipated-component regression; the surprise and information regressions have slightly larger samples because those shock series cover a broader set of quarters than the anticipated component, which requires both $\Delta\text{FFR}_t$ and the raw-surprise sum. Standard errors in parentheses; clustered two-way by bank and county. $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

Table 7: Robustness and Identification Diagnostics

| | Van 3m (1) | Van 12m (2) | MMA (3) | Best 12m (4) | HB 3m (5) | Odd-Term (6) | HB 12m (7) |
|---|---|---|---|---|---|---|---|
| **Panel A: Permutation Test (Randomized Concentration)** | | | | | | | |
| $z^{\text{perm}}_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.022 | −0.102 | −0.074 | 0.358 | 0.201 | 0.122 | 0.095 |
| | (0.044) | (0.108) | (0.061) | (0.352) | (0.378) | (0.338) | (0.351) |
| Within $R^2$ | 0.2980 | 0.1110 | 0.6670 | 0.1070 | 0.2150 | 0.1270 | 0.0560 |
| Observations | 381,995 | 422,501 | 376,865 | 439,135 | 131,863 | 397,227 | 137,247 |
| **Panel B: Zero-Quarter Test ($\lvert\Delta\text{FFR}\rvert < 5$ bps)** | | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}}$ | 0.111 | 0.783 | 0.156 | 0.511 | −0.384 | 0.347 | 1.219 |
| | (0.135) | (0.525) | (0.158) | (0.768) | (0.532) | (0.733) | (1.827) |
| $z_{\text{HHIC}}$ | −0.275 | −0.066 | −0.127 | −1.840 | 2.590 | −1.053 | −1.068 |
| | (0.240) | (0.607) | (0.170) | (1.331) | (3.264) | (1.043) | (2.621) |
| Observations | 203,796 | 223,578 | 201,284 | 232,031 | 76,605 | 209,586 | 79,532 |
| **Panel C: Concentration Orthogonalized to Demographics** | | | | | | | |
| $z^{\perp}_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | −0.016 | 0.039 | 0.003 | −0.186 | −0.723*** | 0.486* | 1.123 |
| | (0.056) | (0.108) | (0.034) | (0.315) | (0.209) | (0.255) | (0.715) |
| $z^{\perp}_{\text{HHIC}} \times \Delta\text{FFR}$ | 0.018 | −0.122 | −0.159*** | −0.554* | −0.124 | −0.145 | −0.200 |
| | (0.044) | (0.130) | (0.046) | (0.285) | (0.341) | (0.321) | (0.890) |
| Bank × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Bank × County | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| MSA × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Clustering | Bank × County (two-way) | | | | | | |

*Notes:* Panel A randomly permutes local concentration across a bank's counties within each bank–quarter, breaking the link between a branch's actual county-level competitive environment and the concentration measure while preserving the within-bank-quarter marginal distribution. The table reports the permuted rival-concentration coefficient from the horse-race specification for compactness. Panel B restricts the sample to quarters with no change in the federal funds rate ($\lvert\Delta\text{FFR}\rvert < 5$ bps) and tests whether the *level* of concentration predicts rates. Panel C orthogonalizes both $z_{\text{HHI}^{\text{OTH}}}$ and $z_{\text{HHIC}}$ against county population and median household income, removing any confound from county demographics. MSA × Quarter uses CMSA definitions for columns (1)–(4) and (6), and standard MSA definitions for high-balance products in columns (5) and (7). Standard errors in parentheses; clustered two-way by bank and county. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 8: Asymmetric Pass-Through: Tightening versus Easing Quarters

| | Van 3m (1) | Van 12m (2) | MMA (3) | Best 12m (4) | HB 3m (5) | Odd-Term (6) | HB 12m (7) | Acq 12m (8) |
|---|---|---|---|---|---|---|---|---|
| *Rival Concentration ($z_{\text{HHI}^{\text{OTH}}}$)* | | | | | | | | |
| Hike × $z_{\text{HHI}^{\text{OTH}}}$ | −0.004 | 0.071 | −0.051 | −0.221 | −0.678** | 0.901*** | 1.451* | −0.522 |
| | (0.055) | (0.136) | (0.053) | (0.346) | (0.326) | (0.340) | (0.743) | (0.528) |
| Cut × $z_{\text{HHI}^{\text{OTH}}}$ | 0.093 | 0.187 | 0.018 | 0.820 | 0.967 | 0.186 | −0.671 | −1.213*** |
| | (0.130) | (0.243) | (0.096) | (0.699) | (0.849) | (0.439) | (0.902) | (0.371) |
| Wald $F$ (hike = cut) | 0.50 | 0.20 | 0.68 | 1.86 | 5.62** | 1.58 | 2.69 | 3.27* |
| | [0.48] | [0.66] | [0.41] | [0.17] | [0.02] | [0.21] | [0.10] | [0.07] |
| *Own Concentration ($z_{\text{HHIC}}$)* | | | | | | | | |
| Hike × $z_{\text{HHIC}}$ | −0.001 | −0.079 | −0.116** | −0.244 | 0.105 | −0.062 | 1.276 | 0.419 |
| | (0.047) | (0.131) | (0.049) | (0.457) | (0.402) | (0.467) | (1.184) | (0.512) |
| Cut × $z_{\text{HHIC}}$ | −0.007 | 0.264 | 0.219** | 1.361** | 1.067* | −0.079 | 1.705 | 0.014 |
| | (0.093) | (0.336) | (0.086) | (0.636) | (0.626) | (0.469) | (1.510) | (0.435) |
| Wald $F$ (hike = cut) | 0.00 | 1.02 | 11.90*** | 6.64*** | 1.42 | 0.00 | 0.05 | 0.48 |
| | [0.96] | [0.31] | [0.00] | [0.01] | [0.24] | [0.98] | [0.83] | [0.49] |
| Bank × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Bank × County | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| MSA × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Within $R^2$ | 0.3040 | 0.1080 | 0.6570 | 0.1070 | 0.1410 | 0.1290 | 0.0650 | 0.0300 |
| Observations | 443,145 | 487,762 | 439,461 | 505,806 | 167,002 | 460,357 | 173,987 | 207,708 |
| Bank clusters | 1,406 | 1,589 | 1,509 | 1,599 | 393 | 1,427 | 456 | 905 |
| County clusters | 2,740 | 2,773 | 2,719 | 2,774 | 1,075 | 2,730 | 1,084 | 2,613 |

*Notes:* The table reports separate interactions for tightening quarters (hike: $\Delta\text{FFR} > 0$) and easing quarters (cut: $\Delta\text{FFR} < 0$), entered simultaneously within each regression. The Wald $F$-statistic tests the null that the hike and cut coefficients are equal; $p$-values in brackets. Stars on the Wald statistic reflect significance of the $p$-value. MSA × Quarter uses CMSA definitions for all columns except (5) and (7), which use standard MSA. Standard errors in parentheses; clustered two-way by bank and county. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 9: Menu Architecture: Product Shares, Complexity, and Availability

| | Panel A: Product Shares | | | | Panel B: Architecture | | Panel C: Availability | | |
|---|---|---|---|---|---|---|---|---|---|
| | Vanilla Share (1) | High-Bal CD (2) | Odd-Term (3) | Promo Share (4) | $N$ Prod. (5) | Rate Range (6) | Has Promo (7) | Has Odd-Term (8) | Has HB (9) |
| *Horse Race (Rival and Own Concentration Jointly)* | | | | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.011 | 0.001 | −0.010 | −0.008* | −0.745 | −0.085 | −0.042* | −0.041 | −0.000 |
| | (0.012) | (0.002) | (0.012) | (0.004) | (0.463) | (0.172) | (0.022) | (0.027) | (0.029) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.003 | −0.000 | −0.020* | −0.006 | 0.087 | −0.388* | 0.008 | −0.046* | 0.003 |
| | (0.012) | (0.002) | (0.011) | (0.006) | (0.466) | (0.207) | (0.030) | (0.028) | (0.028) |
| Bank × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Bank × County | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| MSA × Quarter | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Within $R^2$ | 0.4780 | 0.8360 | 0.2500 | 0.4840 | 0.0880 | 0.2000 | 0.3750 | 0.2340 | 0.1190 |
| Observations | 512,262 | 512,262 | 512,262 | 512,262 | 518,142 | 517,393 | 512,999 | 512,999 | 512,999 |
| Bank clusters | 1,588 | 1,588 | 1,588 | 1,588 | 1,619 | 1,614 | 1,592 | 1,592 | 1,592 |
| County clusters | 2,774 | 2,774 | 2,774 | 2,774 | 2,778 | 2,776 | 2,776 | 2,776 | 2,776 |

*Notes:* Each observation is a bank–county–quarter. The dependent variables are menu-composition and availability measures. Panel A: product-count shares (fractions of the bank's CD menu in a given county–quarter). Panel B: number of distinct products offered and the within-bank rate range. Panel C: indicator variables for the presence of specific product types. All specifications use the horse-race decomposition with both $z_{\text{HHI}^{\text{OTH}}}$ and $z_{\text{HHIC}}$ entered simultaneously. MSA × Quarter uses CMSA definitions. Standard errors in parentheses; clustered two-way by bank and county. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table 10: Variable Definitions

| Variable | Definition |
|---|---|
| *Deposit Products* | |
| Vanilla CD | Standard certificate of deposit at a benchmark maturity (3m or 12m) with no special features, relationship requirements, or minimum-balance thresholds beyond the bank's standard. Serves as the pricing anchor. |
| Best available CD | Highest posted rate at a given maturity across all CD products offered by the bank in a county, including promotional, acquisition, and odd-term CDs. |
| High-balance (HB) CD | CD requiring a minimum deposit above the bank's standard threshold (typically \$100,000+). Proxies for large-denomination ("jumbo") deposits. Available primarily at larger, urban-market banks. |
| Odd-term CD | CD with a non-standard maturity (e.g., 7m, 11m, 13m, 17m) that does not match benchmark maturities. Used as a screening or retention device. |
| Acquisition CD | Promotional CD offered to new customers or new money only, typically at above-market rates for a limited time. |
| MMA | Money market account — a liquid savings product with check-writing privileges. Rates are administratively reset. |
| Savings | Standard passbook savings account. Serves as the lowest-yield liquid product in the deposit menu. |
| *Within-Menu Spreads* | |
| MMA − Van 3m | Money market rate minus vanilla 3-month CD rate. Captures substitutability between liquid and short-term products. |
| HB slope (12m−3m) | High-balance 12-month best rate minus high-balance 3-month best rate. Measures the maturity incentive within the high-balance tier. |
| Odd − Best 12m | Odd-term best rate minus best available 12-month CD rate. Captures the screening premium from non-standard maturities. |
| Acq − Van 12m | Acquisition 12-month rate minus vanilla 12-month rate. Measures the teaser premium for new-money CDs. |
| *Market Structure* | |
| HHI | Herfindahl–Hirschman Index of deposit concentration (branch-level SOD), scaled so 100 = monopoly. |
| HHIC | Own-bank contribution to county HHI ($= s_i^2 \times 10{,}000$). Measures the bank's own market power. |
| $\text{HHI}^{\text{OTH}}$ | Rival concentration: HHI − HHIC. Captures the competitive environment facing the bank. |
| LOO gap | Bank's own rate minus the best rival rate in the county (leave-one-out frontier). |
| $\widetilde{HHI}$ (Bartik) | Bank-level concentration exposure: weighted average of county-level HHI using pre-sample (2010–2014) branch-deposit footprint shares $\omega_{ic,0}$ from the FDIC Summary of Deposits. See Equation (14). |

Table 10: Variable Definitions (continued)

| Variable | Definition |
|---|---|
| *Monetary Policy* | |
| $\Delta$FFR | Quarterly change in the effective federal funds rate (basis points). |
| Anticipated $\Delta$FFR | $\Delta\text{FFR}_t$ minus the within-quarter sum of raw meeting-level surprises from Bauer and Swanson (2023). Captures the expected component of the quarterly policy change. |
| Pure surprise | Within-quarter sum of orthogonalized meeting-level surprises. Measures the unexpected policy innovation purged of the Fed-information channel. |
| Information shock | Difference between the within-quarter sums of raw and orthogonalized meeting-level surprises. Captures the central-bank information channel: the component of the rate surprise correlated with revised macro expectations. |
| *Bank-Quarter Deposit Flows (Table 5, Panel A)* | |
| Total deposits (UBPR2200) | $\Delta \log(1 + \text{total deposits}) \times 100$. Quarterly growth rate of aggregate deposits, from UBPR. |
| High-balance deposits (UBPRK222) | $\Delta \log(1+\text{high-balance deposits}) \times 100$. Quarterly growth rate of time deposits exceeding \$250,000 with remaining maturity of one year or less. |
| Core deposits (UBPRK434) | $\Delta \log(1 + \text{core deposits}) \times 100$. Quarterly growth rate of core deposits: all transaction accounts + MMDA + other savings + time deposits $\leq$\$250,000 $-$ fully insured brokered deposits. |
| Brokered deposits (UBPR2365) | $\Delta \log(1 + \text{brokered deposits}) \times 100$. Quarterly growth rate of deposits placed through brokers or deposit-listing services. |
| *Bank-Quarter Yield, Cost, and Composition Outcomes (UBPR)* | |
| Deposit yield (UBPRE106) | Interest on all interest-bearing deposits divided by average interest-bearing deposits, annualized. FFIEC UBPR. Reported in basis points. |
| Cost of time CDs (UBPRHR55) | Interest on all time deposits divided by average time deposits, annualized. FFIEC UBPR. Reported in basis points. |
| Net interest margin (UBPRKX40) | Net interest income as a percent of average assets, annualized. FFIEC UBPR. Reported in basis points. |
| Core deposit share (UBPRK434/UBPR2200) | Core deposits as a percentage of total deposits. First-differenced. |
| High-balance deposit share (UBPRK222/UBPR2200) | Time deposits exceeding \$250,000 with remaining maturity of one year or less, as a percentage of total deposits. First-differenced. |
| Brokered deposit share (UBPR2365/UBPR2200) | Brokered deposits as a percentage of total deposits. First-differenced. |
| Small time CD $\leq$1yr (UBPRM023) | Time deposits with face value $\leq$\$250,000 and remaining maturity under one year, from Call Report Schedule RC-E. In thousands of dollars. |
| Small time CD >1yr (UBPRM024) | Time deposits with face value $\leq$\$250,000 and remaining maturity over one year, from Call Report Schedule RC-E. In thousands of dollars. |

# A Appendix Framework: Deposit-Menu Pricing and Market Structure

This appendix expands the stylized framework in Section 3 of the main text. The framework is built around a simple pricing logic: markdowns vary with the elasticity a bank faces on each contract, and those elasticities need not respond to market structure in the same way across products. Its role is interpretive rather than identificational. The empirical design does not depend on the framework being literally correct. The purpose of the appendix is to make explicit why the paper studies anchors, liquid products, frontier objects, and within-menu wedges separately, and why those objects need not all load on market structure in the same way.

## A.1 Pricing Logic

A bank $i$ in market $c$ posts a menu of deposit contracts $k \in \mathcal{K}$, where products differ by liquidity, maturity, balance threshold, eligibility, and comparability. Let $r_{ikc}$ denote the posted rate on product $k$. The bank chooses the vector $r_i = \{r_{ikc}\}_{k \in \mathcal{K}}$ to fund its balance sheet at minimum cost while retaining relationship balances and competing for more rate-sensitive funds. A convenient reduced-form representation of profits is

$$\Pi_i = \sum_{k \in \mathcal{K}} \left[ v_k(R) - r_{ikc} \right] D_{ikc}(r_i, r_{kc}^{frontier}; \theta_c) - \Phi(r_i), \tag{A.1}$$

where $R$ is the policy-sensitive outside funding benchmark, $v_k(R)$ is the shadow value of deposits of type $k$, $r_{kc}^{frontier}$ is the relevant local rival offer or competitive frontier, $\theta_c$ indexes local depositor frictions, and $\Phi(r_i)$ captures any costs of operating a richer menu. The central pricing object is therefore the vector $r_i$, not a representative posted rate. Depositors are heterogeneous in rate sensitivity (cf. Egan et al., 2025;

Eichenbaum et al., 2025). One group is relationship-bound, inattentive, or high-friction. These depositors value convenience, branch proximity, transaction services, or inertia, and they are relatively unlikely to compare offers across institutions. A second group is more rate-sensitive. These depositors search, compare products across banks, and respond to the best available offer in the market. In equilibrium, banks use the menu to allocate different depositor types into different contracts. Standard vanilla products and some liquid accounts serve as anchors for high-friction balances; competitive best offers, high-balance contracts, acquisition products, and less comparable contracts serve as the margins on which banks compete for or screen more elastic funds. The framework does not solve the bank's full pricing problem. It uses only a simple monotone pricing logic: the wedge on a contract is larger when the effective elasticity the bank faces on that contract is lower. Let

$$\delta_{ikc} \equiv v_k(R) - r_{ikc}, \tag{A.2}$$

denote the wedge between the shadow value of deposits of type $k$ and the posted rate on that contract, and let $\varepsilon_{ikc}^{eff}$ denote the effective elasticity relevant for product $k$. The only comparative static the framework needs is

$$\frac{\partial \delta_{ikc}}{\partial \varepsilon_{ikc}^{eff}} < 0. \tag{A.3}$$

Contracts facing less elastic demand sustain larger wedges. The empirical question is therefore which dimensions of market structure shift the relevant elasticity for which products.

## A.2 Depositor Margins and the Two Concentration Objects

The paper's concentration decomposition corresponds to two distinct pricing margins.

**Own concentration.** Let $HHIC_{ic}$ denote the bank's own contribution to local concentration. In the framework, this is a reduced-form proxy for unilateral retention power over relationship balances. For liquid, non-maturity products, higher own concentration lowers the effective retention elasticity:

$$\frac{\partial \varepsilon_{ikc}^{ret}}{\partial HHIC_{ic}} < 0, \qquad \text{for liquid products.} \tag{A.4}$$

As the bank's captive local base becomes larger, it faces less outflow pressure on balances that can be retained through convenience, inertia, or relationship services.

**Rival concentration.** Let $HHI_{ic}^{OTH}$ denote residual concentration attributable to competing institutions. In the framework, this is a reduced-form proxy for the quality of the outside option available to rate-sensitive depositors. For competitive products, rival concentration weakens the local frontier:

$$\frac{\partial^2 r_{kc}^{frontier}}{\partial R \, \partial HHI_{ic}^{OTH}} < 0. \tag{A.5}$$

When rivals are concentrated, fewer institutions bid aggressively for elastic depositors, so the best available local offer moves less with policy. These two objects need not have symmetric effects. Own concentration primarily changes the bank's unilateral retention margin. Rival concentration primarily changes the quality of the outside option that disciplines competitive pricing.

## A.3 From the Local Frontier to the Bank's Own Best Offer

A useful distinction is between the *local competitive frontier* and the bank's *own best posted offer*. Write

$$r_{ikc}^{best} = r_{kc}^{frontier} - m_{ikc}, \tag{A.6}$$

where $m_{ikc}$ is the bank-specific markdown relative to the local frontier. This markdown captures the extent to which the bank chooses not to match the best outside offer. Own concentration is assumed to affect this margin through unilateral retention power:

$$\frac{\partial^2 m_{ikc}}{\partial R \, \partial HHIC_{ic}} > 0. \tag{A.7}$$

A larger captive base reduces the bank's incentive to match the frontier aggressively when policy tightens. Equations (A.5) and (A.7) imply that the bank's own best posted offer can load on *both* concentration margins:

$$\frac{\partial^2 r_{ikc}^{best}}{\partial R \, \partial HHI_{ic}^{OTH}} < 0, \qquad \frac{\partial^2 r_{ikc}^{best}}{\partial R \, \partial HHIC_{ic}} < 0. \tag{A.8}$$

This is why the paper does not require every best-offer product to sort cleanly into a single channel. The outside-option prediction applies most cleanly to the frontier itself. A bank's own best posted rate is its endogenous response to that frontier and can therefore reflect both rival structure and unilateral retention power.

## A.4 Mapping Contracts to Empirical Objects

The framework groups observed contracts into four broad categories.

**Anchor products.** These are transparent, standard contracts directed at high-friction depositors. Their effective elasticity is low for reasons largely orthogonal to

local market structure at the relevant margin. Hence

$$\frac{\partial^2 r_{ikc}}{\partial R \, \partial HHIC_{ic}} \approx 0, \qquad \frac{\partial^2 r_{ikc}}{\partial R \, \partial HHI_{ic}^{OTH}} \approx 0. \tag{A.9}$$

The empirical counterparts are vanilla 3-month and 12-month certificates. For CDs, these are the contracts that govern the default auto-renewal outcome if the depositor takes no action.

**Liquid products.** These products retain balances that can leave quickly. Their relevant margin is unilateral retention. As own concentration rises, the bank faces less outflow pressure from these balances, so the contract responds less to policy:

$$\frac{\partial^2 r_{ikc}}{\partial R \, \partial HHIC_{ic}} < 0, \qquad \text{for liquid products.} \tag{A.10}$$

The empirical counterparts are MMA levels and liquid-product spreads such as MMA–vanilla and MMA–savings.

**Frontier objects.** These are the prices on which banks compete most directly for rate-sensitive depositors. Their relevant discipline is the outside option. The clean prediction is therefore about the local frontier and contracts that sit very close to it:

$$\frac{\partial^2 r_{kc}^{frontier}}{\partial R \, \partial HHI_{ic}^{OTH}} < 0. \tag{A.11}$$

Empirically, this maps most naturally into the leave-one-out best rival offer, the short-horizon high-balance best rate, and bank best offers understood as endogenous responses to the frontier rather than pure frontier objects themselves.

**Within-menu wedges.** These are relative-price objects inside the same bank's menu. Their purpose is to absorb policy changes through maturity, eligibility, or reduced comparability rather than through the most transparent contract. The common prediction is not that every wedge loads on the same concentration margin, but that wedges move systematically with market structure as banks reallocate compensation across products. A generic representation is

$$\frac{\partial^2 W_{ikc}}{\partial R \, \partial Z_{ic}} \neq 0, \tag{A.12}$$

where $W_{ikc}$ is a within-menu spread and $Z_{ic}$ is the concentration margin most relevant for that wedge. One parsimonious mapping consistent with the evidence is given by three leading examples:

$$\frac{\partial^2 \left( r_{ikc}^{HB,12m} - r_{ikc}^{HB,3m} \right)}{\partial R \, \partial HHI_{ic}^{OTH}} > 0, \tag{A.13}$$

$$\frac{\partial^2 \left( r_{ikc}^{Acq,12m} - r_{ikc}^{Van,12m} \right)}{\partial R \, \partial HHI_{ic}^{OTH}} > 0, \tag{A.14}$$

$$\frac{\partial^2 \left( r_{ikc}^{Odd} - r_{ikc}^{Best,12m} \right)}{\partial R \, \partial HHIC_{ic}} > 0. \tag{A.15}$$

These sign restrictions are not meant as a full partition of the entire menu, nor as a claim that every wedge must load on only one concentration margin. They are organizing reduced-form statements that guide the empirical analysis and summarize the cleanest empirical patterns.

## A.5 Why Relative Prices Are First-Order

Within-menu wedges matter for effective funding costs because depositors can reallocate across products. Once balance shares respond to relative prices inside the menu, changes in those wedges affect the bank's funding cost through both repricing and composition. Let the bank's effective deposit funding cost be

$$C_{it} = \sum_{k \in \mathcal{K}} \omega_{ikt} r_{ikt}, \tag{A.16}$$

where $\omega_{ikt}$ is the share of balances allocated to product $k$. Its sensitivity to the policy rate is

$$\frac{dC_{it}}{dR_t} = \underbrace{\sum_k \omega_{ikt} \frac{dr_{ikt}}{dR_t}}_{\text{pass-through on existing balances}} + \underbrace{\sum_k r_{ikt} \frac{d\omega_{ikt}}{dR_t}}_{\text{composition}}. \tag{A.17}$$

The first term is pass-through holding the product mix fixed. The second term is composition: once balance shares respond to the menu, changes in relative prices affect funding costs through reallocation. Within-menu wedges are therefore not second-order unless product allocation is fixed, that is, unless $d\omega_{ikt}/dR_t = 0$ for all $k$. A compact way to summarize the mechanism is to define the vector of product-specific pass-through responses,

$$\beta_{ic} = \left( \frac{\partial r_{ikc}}{\partial R} \right)_{k \in \mathcal{K}}. \tag{A.18}$$

A scalar approach asks how concentration changes one representative element of this vector. The framework instead predicts that concentration changes its *shape*. Menu rotation is this change in shape: own concentration compresses liquid-product adjustment, rival concentration weakens the frontier, and within-menu wedges adjust so that policy is redistributed across contracts rather than transmitted uniformly through a single rate.

## A.6 Predictions and Empirical Counterparts

The empirical analysis in the main text is organized around four directional predictions.

1. **Anchor invariance.** Standard vanilla anchor contracts should be locally inert. This is tested on vanilla 3-month and 12-month rates.

2. **Liquid-retention channel.** Liquid products should under-respond to policy when own concentration is higher. This is tested on MMA levels and liquid-product spreads.

3. **Outside-option / frontier channel.** The local competitive frontier should under-respond to policy when rival concentration is higher. A bank's own best posted offer may also under-respond with own concentration because it is an endogenous response to that frontier. This is tested most cleanly on the leave-one-out frontier, the short-horizon high-balance best rate, and bank best offers interpreted relative to the frontier.

4. **Within-menu wedge adjustment.** Relative prices inside the menu should move systematically with concentration as banks reallocate compensation away from the most transparent contracts and toward duration, eligibility, or imperfect comparability. This is tested on the high-balance maturity slope, acquisition premia, odd-term premia, and related spreads.

The framework therefore does not predict that every individual contract-level coefficient must sort cleanly into a single channel. It predicts something sharper and more relevant for the data: within the same bank and quarter across geographies, local market structure should change where adjustment appears inside the menu. That is the sense in which the empirical results are interpreted as evidence of menu rotation.

# B Fixed-Effect Saturation and Identification

The baseline specification in the main text absorbs bank × quarter, market × quarter, and bank × county fixed effects. This appendix studies how the main horse-race coefficients behave as those controls are introduced progressively. The purpose of the exercise is not to show invariance to underspecified regressions. Sparse specifications can be materially confounded because they leave common bank-wide repricing shocks and market-wide demand or competitor shocks in the error term. The relevant question is narrower: once the design absorbs those bank-time and market-time shocks, do the coefficient patterns for the key objects emphasized in the main text emerge and remain broadly stable under further saturation? They do.

For liquid products such as MMAs, the baseline outcome uses the quarter-to-quarter change in the within-quarter mean ($\Delta^{\text{Dqm}}$), following the contractual logic in Section 4.4 of the main text.

Table A1 first replaces the baseline CMSA × quarter fixed effect with state × quarter, a coarser geographic control. The main coefficient patterns remain similar, indicating that the results are not driven by the specific choice of CMSA market definition.

Tables A2 and A3 report the fixed-effect saturation exercise on common samples defined by the fully saturated specification for each outcome. This isolates the effect of progressively absorbing bank-time and market-time shocks from changes in sample composition. Starting from sparse specifications, each successive block adds controls that absorb a distinct source of nuisance variation until reaching the fully saturated baseline. The important comparison is not between the fully saturated design and the sparsest possible specification per se. Rather, it is whether the horse-race coefficients become sharper and then remain stable once the design absorbs the variation that is

orthogonal to the paper's mechanism.

This is the broad pattern for the key objects emphasized in the main text: precision often improves as the fixed-effect structure is saturated. That improvement has a straightforward economic interpretation. The identifying variation in the paper is local and relative: within the same bank-quarter, across counties with different concentration. Sparse specifications leave large bank-wide repricing shocks and market-wide demand or competitor shocks in the error term. Adding bank $\times$ quarter and market $\times$ quarter fixed effects removes exactly that nuisance variation. Precision therefore improves when the fall in residual outcome variance is larger than the shrinkage in the residualized concentration interaction. When coefficients move materially across sparse specifications, the natural interpretation is not that the result is fragile, but that omitted bank-time or market-time shocks were confounding the less saturated design.

The ladder is especially informative for three objects. First, vanilla anchor rates remain weakly related to concentration throughout, and become cleaner nulls under the fully saturated design. Second, the MMA own-concentration coefficient becomes sharper once bank-wide repricing shocks are absorbed, consistent with the idea that the liquid-retention channel is identified from local deviations around a common bank-level pricing policy. Third, the rival-concentration coefficients on frontier and screening objects become more informative once market-wide shocks are absorbed, consistent with the outside-option channel operating through local competitive conditions rather than through broad regional movements in deposit demand.

Table A4 then reports the residual standard deviation of the outcome and the two concentration interactions after partialling out each FE set, clarifying why precision can improve under the more saturated designs. As bank $\times$ quarter and market $\times$ quarter fixed effects absorb bank-wide repricing and market-wide demand shocks, the residual

standard deviation of the outcome falls, while enough local identifying variation remains in the concentration interactions. Table A5 reports the same decomposition for the hiking subsample.

Figures A1 and A2 visualize the common-sample exercise. The top panels plot the horse-race coefficients with 95% confidence intervals across the fixed-effect ladder. The lower panels report the residual dispersion of the dependent variable (normalized to FE1 = 1.0) after partialling out the corresponding fixed effects, showing how rapidly outcome noise declines as the design absorbs bank-time and market-time shocks. Figure A2 repeats the exercise on the hiking subsample ($\Delta\text{FFR} > 0$), where the same broad pattern holds.

Taken together, the saturation exercise supports the baseline identification strategy. The key coefficients do not survive merely because the specifications are flexible. They survive because the flexibility is economically well targeted: it absorbs the bank-wide and market-wide shocks that should not identify a local concentration effect, while leaving intact the within-bank, across-market variation through which the paper's mechanism operates.

Table A1: Robustness: State × Quarter Fixed Effects

| | Van 12m (1) | MMA (2) | Best 12m (3) |
|---|---|---|---|
| *Panel A: Total HHI* | | | |
| $z_{\text{HHI}} \times \Delta\text{FFR}$ | $-0.022$ | $-0.122^{**}$ | $-0.562^{*}$ |
| | (0.100) | (0.050) | (0.298) |
| *Panel B: Own Concentration (HHIC)* | | | |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | $-0.103$ | $-0.149^{***}$ | $-0.613^{**}$ |
| | (0.130) | (0.042) | (0.277) |
| *Panel C: Horse Race (Rival and Own Concentration Jointly)* | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.014 | $-0.054$ | $-0.344$ |
| | (0.098) | (0.048) | (0.320) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | $-0.102$ | $-0.145^{***}$ | $-0.638^{**}$ |
| | (0.132) | (0.042) | (0.281) |
| Bank × Quarter | Yes | Yes | Yes |
| Bank × County | Yes | Yes | Yes |
| State × Quarter | Yes | Yes | Yes |
| Within $R^2$ | 0.111 | 0.658 | 0.109 |
| Observations | 487,820 | 439,568 | 505,859 |
| Bank clusters | 1,589 | 1,509 | 1,599 |
| County clusters | 2,773 | 2,719 | 2,774 |

*Notes:* The table replicates the baseline specification from Table 2 of the main text using state × quarter fixed effects instead of CMSA × quarter. The purpose is to test whether the main horse-race patterns depend on the specific market definition used in the baseline. The coefficients remain quantitatively similar, especially for the MMA own-concentration channel and the best-12-month own-concentration channel, indicating that the results are not driven by the CMSA partition. Standard errors in parentheses; clustered two-way by bank and county. $^{*}p < 0.10$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table A2: Common-Sample FE Ladder: Horse-Race Coefficients

| | Van 12m | Best 12m | Odd − Best 12m Spread | HB Slope (12m−3m) | MMA | MMA − Van 12m Spread |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| *FE1: Bank + County* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.672** | 2.784*** | −0.383 | 1.509 | 0.812** | −0.024 |
| | (0.262) | (1.060) | (0.569) | (1.148) | (0.338) | (0.308) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.079 | −1.022 | 0.654* | −1.502 | −0.225 | −0.626** |
| | (0.249) | (0.832) | (0.360) | (1.048) | (0.343) | (0.288) |
| $N$ | 490,460 | 508,066 | 443,013 | 209,956 | 441,125 | 435,415 |
| *FE2: Bank × County* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.662** | 2.765*** | −0.380 | 1.438 | 0.781** | −0.049 |
| | (0.265) | (1.073) | (0.576) | (1.161) | (0.339) | (0.308) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.094 | −1.070 | 0.653* | −1.503 | −0.207 | −0.612** |
| | (0.253) | (0.844) | (0.364) | (1.058) | (0.344) | (0.288) |
| $N$ | 490,257 | 507,895 | 442,723 | 209,612 | 440,892 | 435,152 |
| *FE3: Bank + CMSA × Qtr* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.013 | −0.514 | 0.126 | 2.099*** | −0.078 | 0.018 |
| | (0.143) | (0.340) | (0.292) | (0.529) | (0.163) | (0.150) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.151 | −1.403*** | 0.645* | −0.448 | −0.567*** | −0.790*** |
| | (0.192) | (0.358) | (0.359) | (0.653) | (0.199) | (0.289) |
| $N$ | 490,408 | 508,020 | 442,945 | 209,170 | 441,018 | 435,305 |
| *FE4: Bank × County + Qtr* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | −0.159 | −1.069** | −0.131 | 0.895 | −0.150 | 0.262 |
| | (0.202) | (0.497) | (0.299) | (0.734) | (0.250) | (0.290) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.216 | −1.745*** | 0.674* | −0.818 | −0.562*** | −0.765*** |
| | (0.221) | (0.429) | (0.358) | (0.551) | (0.218) | (0.291) |
| $N$ | 490,257 | 507,895 | 442,723 | 209,612 | 440,892 | 435,152 |
| *Full: Bank × Qtr + CMSA × Qtr + Bank × County* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | −0.013 | −0.420 | 0.712** | 1.879*** | −0.041 | −0.107** |
| | (0.112) | (0.318) | (0.291) | (0.582) | (0.034) | (0.052) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.146 | −0.660** | 0.570* | −0.004 | −0.156*** | −0.198*** |
| | (0.140) | (0.292) | (0.315) | (0.551) | (0.044) | (0.053) |
| $N$ | 487,535 | 505,682 | 439,684 | 207,129 | 439,461 | 433,635 |
| Clustering | Bank × County (two-way) | | | | | |

*Notes:* The table reports the horse-race coefficients on rival concentration ($z_{\text{HHI}^{\text{OTH}}}$) and own concentration ($z_{\text{HHIC}}$) as the fixed-effect structure is progressively saturated. For each outcome, the estimation sample is restricted to the common set of observations with non-missing values required by the fully saturated specification. Remaining differences in $N$ across rows reflect fixed-effect-specific singleton dropping rather than sample redefinition. "FE1" includes bank and county fixed effects separately. "FE2" replaces these with bank × county. "FE3" combines bank fixed effects with CMSA × quarter. "FE4" combines bank × county with quarter fixed effects. "Full" is the baseline specification with bank × quarter, CMSA × quarter, and bank × county. Standard errors in parentheses; clustered two-way by bank and county. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table A3: Common-Sample FE Ladder: Horse-Race Coefficients, Hiking Subsample

| | Van 12m | Best 12m | Odd − Best 12m Spread | HB Slope (12m−3m) | MMA | MMA − Van 12m Spread |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| *FE1: Bank + County* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.817** | 3.898*** | 0.168 | 3.300** | −0.515 | −1.036* |
| | (0.357) | (1.233) | (1.290) | (1.560) | (0.509) | (0.553) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | 0.022 | 0.192 | 0.901 | 0.281 | 0.488 | 0.087 |
| | (0.268) | (0.653) | (0.633) | (1.221) | (0.307) | (0.364) |
| $N$ | 199,097 | 206,201 | 179,384 | 83,933 | 178,974 | 176,504 |
| *FE2: Bank × County* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.802** | 3.976*** | 0.283 | 3.084* | −0.762 | −1.206** |
| | (0.389) | (1.314) | (1.431) | (1.579) | (0.561) | (0.609) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.046 | −0.681 | 0.746 | 0.387 | 0.789** | 0.260 |
| | (0.328) | (0.768) | (0.762) | (1.308) | (0.374) | (0.447) |
| $N$ | 198,041 | 205,184 | 178,234 | 83,419 | 177,762 | 175,282 |
| *FE3: Bank + CMSA × Qtr* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.084 | −0.190 | 0.586 | 2.836*** | −0.150 | −0.108 |
| | (0.140) | (0.352) | (0.361) | (0.755) | (0.093) | (0.108) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | 0.038 | −0.459 | 0.399 | 0.540 | −0.003 | −0.246* |
| | (0.152) | (0.404) | (0.389) | (1.012) | (0.117) | (0.146) |
| $N$ | 199,088 | 206,195 | 179,386 | 83,615 | 178,944 | 176,469 |
| *FE4: Bank × County + Qtr* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | −0.372 | −1.144 | 0.745 | 1.386 | −0.440** | 0.390 |
| | (0.266) | (0.722) | (0.627) | (1.135) | (0.197) | (0.327) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.362 | −2.222*** | 0.720 | 0.529 | −0.185 | −0.197 |
| | (0.316) | (0.621) | (0.733) | (0.914) | (0.315) | (0.436) |
| $N$ | 198,041 | 205,184 | 178,234 | 83,419 | 177,762 | 175,282 |
| *Full: Bank × Qtr + CMSA × Qtr + Bank × County* | | | | | | |
| $z_{\text{HHI}^{\text{OTH}}} \times \Delta\text{FFR}$ | 0.068 | −0.264 | 1.290*** | 3.119*** | −0.099 | −0.101 |
| | (0.131) | (0.438) | (0.486) | (0.885) | (0.092) | (0.093) |
| $z_{\text{HHIC}} \times \Delta\text{FFR}$ | −0.062 | −0.821 | 0.288 | 2.115*** | −0.119 | −0.141 |
| | (0.191) | (0.516) | (0.688) | (0.805) | (0.080) | (0.107) |
| $N$ | 196,875 | 204,198 | 176,892 | 82,295 | 177,148 | 174,628 |
| Clustering | Bank × County (two-way) | | | | | |

*Notes:* The table replicates the common-sample fixed-effect ladder (Table A2) on the hiking subsample ($\Delta\text{FFR} > 0$). The common sample is the set of observations with non-missing values required by the fully saturated hiking specification; remaining differences in $N$ across rows reflect singleton dropping. Standard errors in parentheses; clustered two-way by bank and county. * $p < 0.10$; ** $p < 0.05$; *** $p < 0.01$.

Table A4: Residual Dispersion of Outcome and Regressors After FE Partialling

| | Van 12m (1) | Best 12m (2) | Odd − Best 12m Spread (3) | HB Slope (12m−3m) (4) | MMA (5) | MMA − Van 12m Spread (6) |
|---|---|---|---|---|---|---|
| *FE1: Bank + County* | | | | | | |
| SD($\tilde{Y}$) | 49.29 | 124.79 | 111.21 | 122.96 | 33.74 | 42.01 |
| SD($\tilde{X}^{OTH}$) | 51.830 | 52.154 | 52.117 | 50.305 | 52.219 | 52.131 |
| SD($\tilde{X}^{HHIC}$) | 47.433 | 48.115 | 47.396 | 44.180 | 47.500 | 47.362 |
| *FE2: Bank × County* | | | | | | |
| SD($\tilde{Y}$) | 48.80 | 123.72 | 109.91 | 121.99 | 33.64 | 41.89 |
| SD($\tilde{X}^{OTH}$) | 51.364 | 51.703 | 51.636 | 49.920 | 51.803 | 51.700 |
| SD($\tilde{X}^{HHIC}$) | 47.010 | 47.685 | 46.984 | 43.964 | 47.109 | 46.967 |
| *FE3: Bank + CMSA × Qtr* | | | | | | |
| SD($\tilde{Y}$) | 48.13 | 119.88 | 108.95 | 117.71 | 29.28 | 39.76 |
| SD($\tilde{X}^{OTH}$) | 44.044 | 44.380 | 44.159 | 42.256 | 44.348 | 44.243 |
| SD($\tilde{X}^{HHIC}$) | 45.154 | 45.856 | 45.094 | 40.454 | 45.368 | 45.286 |
| *FE4: Bank × County + Qtr* | | | | | | |
| SD($\tilde{Y}$) | 48.48 | 120.78 | 109.56 | 121.23 | 30.32 | 41.33 |
| SD($\tilde{X}^{OTH}$) | 50.027 | 50.344 | 50.351 | 48.716 | 50.402 | 50.293 |
| SD($\tilde{X}^{HHIC}$) | 46.988 | 47.667 | 46.960 | 43.880 | 47.087 | 46.944 |
| *Full: Bank × Qtr + CMSA × Qtr + Bank × County* | | | | | | |
| SD($\tilde{Y}$) | 43.15 | 109.39 | 100.72 | 112.47 | 18.28 | 23.40 |
| SD($\tilde{X}^{OTH}$) | 41.036 | 41.405 | 41.332 | 39.999 | 41.247 | 41.139 |
| SD($\tilde{X}^{HHIC}$) | 41.163 | 41.872 | 41.226 | 37.754 | 41.647 | 41.556 |

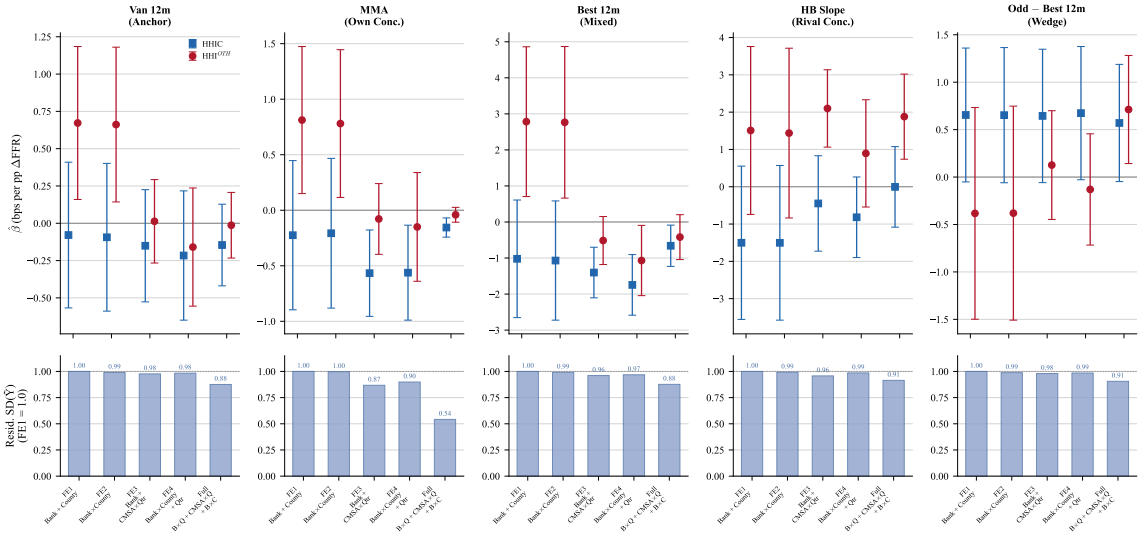
*Notes:* The table reports the standard deviation of the dependent variable and the two concentration interactions after partialling out the corresponding fixed effects. Each column uses the common sample defined by the fully saturated specification for that outcome (same samples as Table A2). SD($\tilde{Y}$) is the residual standard deviation of the outcome in basis points. SD($\tilde{X}^{OTH}$) and SD($\tilde{X}^{HHIC}$) are the residual standard deviations of the rival- and own-concentration interactions, respectively. Fixed effects are partialled out using the same high-dimensional fixed-effect absorption procedure used in the main regressions. The table makes explicit why precision can improve under more saturated designs: bank-wide and market-wide FEs absorb substantial outcome noise while retaining enough local identifying variation in the concentration interactions.

Table A5: Residual Dispersion of Outcome and Regressors After FE Partialling, Hiking Subsample

| | Van 12m (1) | Best 12m (2) | Odd − Best 12m Spread (3) | HB Slope (12m−3m) (4) | MMA (5) | MMA − Van 12m Spread (6) |
|---|---|---|---|---|---|---|
| *FE1: Bank + County* | | | | | | |
| SD($\tilde{Y}$) | 48.17 | 130.94 | 113.91 | 126.41 | 35.56 | 43.01 |
| SD($\tilde{X}^{OTH}$) | 56.525 | 56.842 | 56.192 | 55.103 | 56.735 | 56.727 |
| SD($\tilde{X}^{HHIC}$) | 47.451 | 47.991 | 46.814 | 41.959 | 46.999 | 46.887 |
| *FE2: Bank × County* | | | | | | |
| SD($\tilde{Y}$) | 47.15 | 128.76 | 111.46 | 124.16 | 35.28 | 42.74 |
| SD($\tilde{X}^{OTH}$) | 52.713 | 53.034 | 52.528 | 52.670 | 52.959 | 52.942 |
| SD($\tilde{X}^{HHIC}$) | 41.658 | 42.157 | 41.094 | 38.096 | 41.317 | 41.173 |
| *FE3: Bank + CMSA × Qtr* | | | | | | |
| SD($\tilde{Y}$) | 47.05 | 125.99 | 112.04 | 122.08 | 30.10 | 40.38 |
| SD($\tilde{X}^{OTH}$) | 57.424 | 57.861 | 57.482 | 56.340 | 57.942 | 57.884 |
| SD($\tilde{X}^{HHIC}$) | 56.568 | 57.358 | 56.366 | 50.287 | 56.761 | 56.680 |
| *FE4: Bank × County + Qtr* | | | | | | |
| SD($\tilde{Y}$) | 46.77 | 125.68 | 111.22 | 123.80 | 31.15 | 41.88 |
| SD($\tilde{X}^{OTH}$) | 49.771 | 50.055 | 49.649 | 50.043 | 49.935 | 49.904 |
| SD($\tilde{X}^{HHIC}$) | 41.615 | 42.117 | 41.047 | 37.955 | 41.271 | 41.126 |
| *Full: Bank × Qtr + CMSA × Qtr + Bank × County* | | | | | | |
| SD($\tilde{Y}$) | 40.86 | 113.69 | 101.88 | 114.53 | 16.85 | 21.91 |
| SD($\tilde{X}^{OTH}$) | 40.886 | 41.235 | 40.950 | 41.211 | 40.775 | 40.732 |
| SD($\tilde{X}^{HHIC}$) | 36.536 | 37.071 | 36.200 | 32.728 | 36.443 | 36.348 |

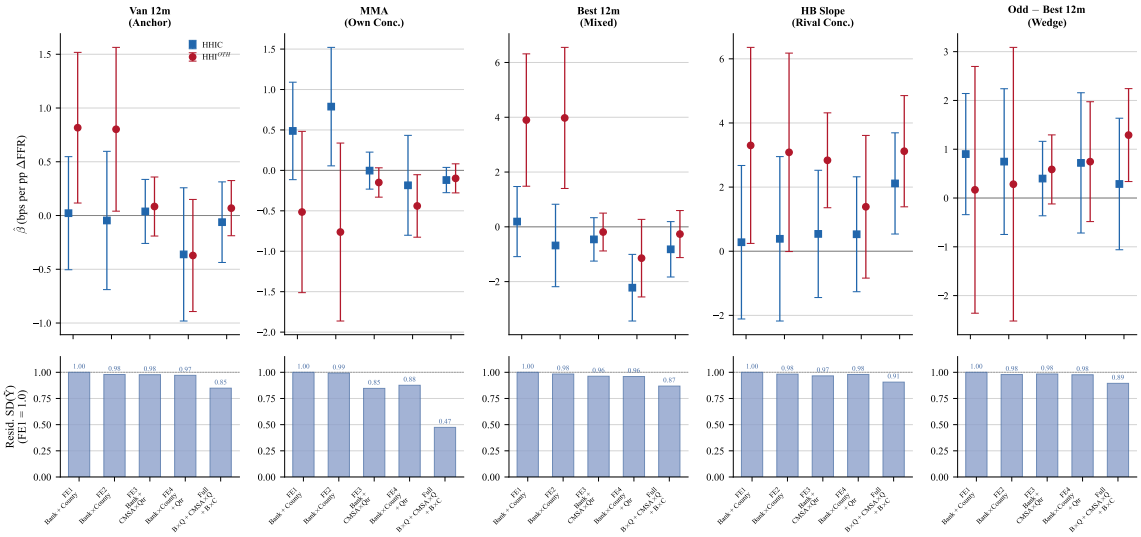
*Notes:* The table replicates Table A4 on the hiking subsample ($\Delta$FFR $> 0$). Each column uses the common sample defined by the fully saturated hiking specification for that outcome (same samples as Table A3). SD($\tilde{Y}$) is the residual standard deviation of the outcome in basis points. SD($\tilde{X}^{OTH}$) and SD($\tilde{X}^{HHIC}$) are the residual standard deviations of the rival- and own-concentration interactions, respectively. Fixed effects are partialled out using the same high-dimensional fixed-effect absorption procedure used in the main regressions.

Figure A1: Fixed-Effect Saturation: Common-Sample Coefficient Paths and Residual Variation




*Notes:* Each panel plots the horse-race coefficients on rival concentration ($z_{\text{HHI}^{\text{OTH}}}$, red circles) and own concentration ($z_{\text{HHIC}}$, blue squares) for the indicated deposit object under progressively saturated fixed-effect structures, estimated on the common sample defined by the fully saturated specification. Error bars represent 95% confidence intervals. The lower bars report the residual standard deviation of the dependent variable (normalized to FE1 = 1.0) after partialling out the corresponding fixed effects. The figure makes visible why precision can improve under more saturated designs: outcome noise falls substantially while enough local identifying variation remains in the concentration interactions. Source data: Tables A2 and A4.

Figure A2: Fixed-Effect Saturation: Common-Sample Coefficient Paths and Residual Variation, Hiking Subsample




*Notes:* The figure replicates Figure A1 on the hiking subsample ($\Delta\text{FFR} > 0$), estimated on the common sample defined by the fully saturated hiking specification. As in the full sample, the more saturated designs absorb bank-wide and market-wide shocks, reducing residual noise in the outcome while local identifying variation in the concentration interactions remains. Source data: Tables A3 and A5.

## B.1 Validating the Liquid-Rate Operator: DQM versus DRQ

Section 4.4 of the main text motivates the use of $\Delta^{\text{Dqm}}$ for liquid products on contractual grounds. This appendix validates that choice directly in the data. The key distinction is that $\Delta^{\text{Dqm}}$ and $\Delta^{\text{DRQ}}$ are not alternative measurements of the same quarterly object. For a variable-rate account such as an MMA or savings account, $\Delta^{\text{Dqm}}$ measures the change in the average posted rate actually applying to the stock of balances over the quarter, whereas $\Delta^{\text{DRQ}}$ measures the net cumulative change in the weekly posted quote. If the economically relevant liquid-rate margin is the stock rate rather than the endpoint change in the quote, the two operators need not produce the same coefficient even under the same identification design.

To assess this difference, I estimate the horse-race specification from the main text on a common sample for three quarterly operators: (i) the change in the within-quarter mean, $\Delta^{\text{Dqm}}$; (ii) the end-of-quarter level change, $\Delta^{\text{EOQ}}$; and (iii) the cumulative within-quarter quote change, $\Delta^{\text{DRQ}}$. The comparison is carried out for the MMA level and for the liquid spread MMA − vanilla. Table A6 reports the resulting coefficients.

Table A7 complements the comparison with path diagnostics. If $\Delta^{\text{Dqm}}$ is the relevant liquid operator, the data should show that within-quarter path variation matters: quarters can exhibit little endpoint change while still displaying meaningful within-quarter movement in the posted stock rate. The diagnostics therefore report the share of bank–county–quarters in which the end-of-quarter change is near zero but the within-quarter range is nontrivial, as well as the share of near-flat policy quarters in which the within-quarter path still moves.

The validation exercise yields a sharp result. For both the MMA level and the MMA–vanilla spread, the own-concentration channel is strong under $\Delta^{\text{Dqm}}$ and essentially absent under both $\Delta^{\text{EOQ}}$ and $\Delta^{\text{DRQ}}$. The estimated HHIC coefficient remains large

Table A6: Liquid-Operator Validation: Common-Sample Horse-Race Coefficients

| | MMA Level | | | MMA − Vanilla Spread | | |
|---|---|---|---|---|---|---|
| | $\Delta^{\mathrm{Dqm}}$ | $\Delta^{\mathrm{EOQ}}$ | $\Delta^{\mathrm{DRQ}}$ | $\Delta^{\mathrm{Dqm}}$ | $\Delta^{\mathrm{EOQ}}$ | $\Delta^{\mathrm{DRQ}}$ |
| $z_{\mathrm{HHI^{OTH}}} \times \Delta\mathrm{FFR}$ | −0.045 | 0.074 | 0.220 | −0.035 | −0.063 | 0.265 |
| | (0.035) | (0.127) | (0.191) | (0.048) | (0.125) | (0.329) |
| $z_{\mathrm{HHIC}} \times \Delta\mathrm{FFR}$ | −0.156*** | −0.028 | −0.030 | −0.236*** | −0.083 | −0.005 |
| | (0.044) | (0.112) | (0.231) | (0.075) | (0.126) | (0.290) |
| Observations | 437,527 | 437,527 | 437,527 | 354,518 | 354,518 | 354,518 |

*Notes:* The table compares three quarterly operators for liquid products on a common sample under the same saturated fixed-effect design. $\Delta^{\mathrm{Dqm}}$ is the quarter-to-quarter change in the within-quarter mean posted rate, $\Delta^{\mathrm{EOQ}}$ is the change in the end-of-quarter posted rate, and $\Delta^{\mathrm{DRQ}}$ is the cumulative within-quarter change in the weekly posted quote. The concentration-sensitive HHIC channel is present under the quarter-mean operator $\Delta^{\mathrm{Dqm}}$ but not under the end-of-quarter or cumulative-quote operators. This indicates that, for liquid accounts, the relevant quarterly object is the average stock rate applying over the quarter rather than the net endpoint change in the weekly posted quote. All regressions absorb bank × quarter, market × quarter, and bank × county fixed effects and cluster two-way by bank and county. Coefficients are reported in basis points per 100-basis-point change in the federal funds rate per one-standard-deviation increase in the relevant concentration measure. $^{*}\,p < 0.10$; $^{**}\,p < 0.05$; $^{***}\,p < 0.01$.

and negative for the quarter-mean operator, but is close to zero and statistically insignificant for the endpoint and cumulative-quote operators. This pattern implies that the concentration-sensitive liquid margin in the data is not the net change in the weekly posted quote, but the average quarterly stock rate actually applying to balances. The path diagnostics support this interpretation directly: a substantial fraction of quarters exhibit little end-of-quarter change despite meaningful within-quarter movement in the MMA path, and the correlation between $\Delta^{\mathrm{Dqm}}$ and $\Delta^{\mathrm{DRQ}}$ is low. The three operators are therefore not interchangeable quarterly summaries of the same object. For liquid accounts, $\Delta^{\mathrm{Dqm}}$ captures the relevant stock-rate margin.

Table A7: Path Diagnostics for Liquid Products

| | MMA Level | MMA − Vanilla Spread |
|---|---|---|
| Share: EOQ change $\approx 0$ but nontrivial within-quarter path | 0.288 | 0.230 |
| Share: near-flat FFR quarter with nontrivial within-quarter path | 0.385 | 0.350 |
| Correlation: $\Delta^{\mathrm{Dqm}}$ vs $\Delta^{\mathrm{EOQ}}$ | 0.447 | 0.558 |
| Correlation: $\Delta^{\mathrm{Dqm}}$ vs $\Delta^{\mathrm{DRQ}}$ | 0.207 | 0.134 |
| Median within-quarter range (bps) | 17.0 | 12.0 |
| Observations | 437,527 | 354,518 |

*Notes:* The table reports descriptive diagnostics for the same common samples used in Table A6. "EOQ change $\approx 0$" means that the absolute end-of-quarter change is at most 1 basis point; "nontrivial within-quarter path" means that the within-quarter range of the relevant posted rate exceeds 1 basis point. "Near-flat FFR quarter" means $|\Delta \mathrm{FFR}_q| < 5$ basis points. The purpose is to show that within-quarter path variation in liquid products can remain economically meaningful even when the endpoint change is limited, so that a quarter-mean operator can capture liquid-rate exposure that an endpoint-style operator misses.