Dynamic Competition for Sleepy Deposits*

Mark Egan†, Ali Hortaçsu‡, Nathan Kaplan§, Adi Sunderam¶ and Vincent Yao||

November 2025

Abstract

We examine the implications of sleepy deposits and their impact on competition, bank value, and financial stability in the US banking sector. We first document the shopping behavior of depositors using novel data on account openings and closures. Depositors infrequently shop for deposits, with 5–15% of depositors opening a new account each year. Shopping behavior is idiosyncratic: deposit accounts are more likely to be closed due to the depositor either moving or dying than because the depositor switched to a new account offering higher rates or better services. Building on these facts, we develop an empirical model of the supply and demand for "sleepy deposits." In the model, banks face dynamic "invest-versus-harvest" incentives in competing for depositors who shop infrequently. We estimate the model and find that depositor sleepiness accounts for 58% of the average bank's deposit franchise value. Sleepiness softens competition, particularly raising markups and franchise value for banks in low-concentration areas, as well as for banks with either low-quality deposit services or high marginal costs. Sleepiness also creates stability in the banking sector. For two main money center banks in the US, the probability of default after the Federal Reserve's 2022-2023 hiking cycle would have increased to more than 20% in a counterfactual without sleepy depositors.

*We thank Jason Allen, Rob Clark, JF Houde, and seminar and conference participants at the Stanford Institute for Theoretical Economics, the University of Wisconsin IO/Finance Reading Group, Duke University, Federal Reserve Bank of New York, Columbia University, UCLA, Stanford University, and CEBRA for helpful comments. Vincent Yao is an unpaid member of the Fiserv Small Business Index Advisory Council and a paid consultant. Mark Egan is an unpaid member of the Fiserv Small Business Index Advisory Council.

†Harvard Business School & NBER. Email: megan@hbs.edu.

‡University of Chicago & NBER. Email: hortacsu@uchicago.edu.

§Harvard Business School and Department of Economics. Email: nkaplan@hbs.edu.

¶Harvard Business School & NBER. Email: asunderam@hbs.edu.

||J. Mack Robinson College of Business at Georgia State University. Email: wyao2@gsu.edu.

1 Introduction

Bank deposits are often referred to as "sleepy" or "sticky," meaning that they are not frictionlessly re-optimized every period. This sleepiness is thought to be a critical feature of the banking sector, enabling banks to engage in maturity transformation (e.g., Hanson et al., 2015) and sustain their business models. Sleepiness also has important consequences for bank regulation. For example, regulators in the UK have expressed concern that bank deposit rates are insufficiently responsive to short-term interest rates (namely, far less responsive than bank *lending* rates), potentially harming depositors.[1] Moreover, the global financial crisis and the recent run on Silicon Valley Bank (SVB) have emphasized the importance of deposit stickiness for financial stability (Drechsler et al., 2023; Jiang et al., 2024).

In this paper, we study the implications of sleepy deposits for the banking sector using new microdata and within the context of a quantitative dynamic model. We start by documenting new facts about depositor behavior in novel data on account openings and closures. We find that deposit account turnover is low on average and varies in the cross-section of account types. Accounts operated by older depositors have lower turnover, whereas business accounts and larger accounts have higher turnover. Our data also shed light on the reasons for account turnover. Less than 20% of account closures are a result of depositors shopping for better terms; the majority are driven by account inactivity, moving, and death.

We then develop and estimate a dynamic structural model of demand and supply for deposits. In the model, depositors are "sleepy" and only infrequently shop for deposits. This sleepiness gives rise to a dynamic "invest-versus-harvest" tradeoff (Beggs and Klemperer, 1992). Banks face a tension between "investing" in their deposit base by attracting active depositors with higher deposit rates and "harvesting" their sleepy depositors by setting lower deposit rates. It is theoretically ambiguous, and therefore an empirical question, whether this tradeoff raises or lowers equilibrium markups (Dubé et al., 2009; Brown et al., 2024; MacKay and Remer, 2024).

We show that taking depositor sleepiness and dynamic competition seriously has important implications for the banking sector. First, sleepiness raises average markups. Average markups are 53% lower in the absence of depositor sleepiness. Second, under dynamic competition, markups are procyclical, rising more when the short rate rises. Third, accounting for dynamic competition eliminates the increasing relationship between concentration and markups that obtains under static competition. Fourth, sleepiness increases banks' franchise values, accounting for about 58% of the average bank's deposit franchise value. Fifth, sleepiness disproportionately increases bank franchise value for banks with either low product quality or high costs.

Our anonymous bank microdata comes from one of Fiserv's core account processing platforms. Fiserv is a leading provider of payments and financial services technology solutions.

[1] Rafe Uddin, "FCA threatens banks with action as it reviews savings rates," *Financial Times*, September 1, 2023, https://www.ft.com/content/5fac24ff-e9b9-4b03-aede-3dd2507e7f43.

One of its key businesses involves providing financial technology solutions to banks and credit unions, helping them maintain customer deposit and loan accounts. Our dataset includes turnover data (openings and closures) from over 900 banks and credit unions, covering more than 42 million deposit accounts. We observe anonymized turnover data aggregated across different dimensions.

We start by examining the persistence of deposits, focusing on account turnover, measured as new deposit account openings in a year relative to all outstanding accounts. New accounts account for 5-15% of all outstanding accounts, meaning that the average life of an account is 8-9 years. We also document cross-sectional variation in deposit turnover. Turnover is higher among business and trust accounts compared to individual accounts. Larger accounts, measured by account balances, also appear to be more active. In contrast, accounts held by elderly individuals are less active. Somewhat surprisingly, we find that accounts set up for online banking tend to experience less turnover. Part of this appears to be driven by the motivations for switching accounts. We show that moving, which is inherently a less important factor for online depositors, is a key reason why depositors switch accounts.

We interpret this cross-sectional variation in account turnover as being primarily driven by depositor sleepiness. Consistent with this interpretation, we observe that deposit accounts are more often closed due to a depositor either dying or moving than because the depositor moved their funds to a bank with better rates or services.

In the second half of the paper, we develop and estimate a model of demand and supply for deposits that features sleepy depositors. In the model, depositors enter a given period either active ("awake") or inactive ("asleep") following the basic setup of Beggs and Klemperer (1992) and Brown et al. (2024). Active depositors shop for bank accounts and select the bank that maximizes their indirect utility. In contrast, inactive depositors do not shop and simply maintain their bank account from the previous period. We assume that whether a depositor is active or inactive depends on market-specific and aggregate state variables, such as the federal funds rate. In this way, our model can capture the drivers of sleepiness, such as inattention (whether rational or irrational), switching costs, and search costs, in reduced form, although we do not explicitly take a stand on the underlying mechanism.

Banks compete for deposits to maximize their franchise value, defined as the present value of current and future profits. Depositor sleepiness creates invest-versus-harvest incentives on the supply side, which in turn induces dynamic competition on deposit rates. Banks set rates based on their stocks of sleepy deposits, but today's rates affect the stocks of sleepy deposits carried into tomorrow, which in turn shape the rates that banks set tomorrow. Banks collect deposits and earn a spread, competing in each market by setting deposit rates.

The model is quantitative and straightforward to estimate using data. We begin by estimating the share of inactive depositors in each market, parameterizing depositor sleepiness as a function of market-level and aggregate state variables. A key challenge in this estimation is to

determine whether the observed persistence in depositors' choices over banks arises from depositor sleepiness or from persistence in depositor preferences. We address this identification challenge in two ways.

First, we leverage our microdata on account openings and closures. We show that, under the assumptions of our model, new account activity is sufficient to identify the degree of depositor sleepiness. This provides us with a first, direct measure of sleepiness. Second, we examine the annual autocorrelation in bank deposits at the bank-by-county level. A natural concern is that this autocorrelation may be driven by both depositor sleepiness and persistence in branch quality. To control for unobserved quality, we use a control function approach, as outlined in Allen et al. (2024), incorporating cost-shifter instruments used in Whited et al. (2021) and Whited et al. (2022).

We estimate that approximately 94% of depositors are inactive each year, consistent with estimates from practitioners (Bancography, 2014). These estimates are quantitatively similar across our two approaches to measuring sleepiness. Consistent with our reduced-form evidence, we find cross-sectional variation in the degree of sleepiness across depositors. In addition, we find that sleepiness is negatively correlated with the lagged federal funds rate—people wake up when interest rates increase. This evidence is consistent with other work (e.g., Lu et al., 2025) that highlights that depositor alertness tends to increase when the returns to shopping are high.

With our estimates of sleepiness in hand, we next estimate depositor preferences. We use our estimates of sleepiness to calculate each bank's market share among active depositors. The choices of these active depositors allow us to estimate the preferences of depositors. Specifically, we calculate each bank's market share among active depositors and then estimate preferences using market share data following Berry (1994). We find that accounting for depositor sleepiness is critical for estimating depositor's preferences. Depositors appear to be about 33% more elastic after accounting for sleepiness.

Lastly, we estimate the supply-side parameters of our model. The way in which each bank optimally navigates the invest-versus-harvest tradeoff depends both on the demand-side parameters (preferences and sleepiness) and on the way in which rate choices today affect all banks' market shares tomorrow (Beggs and Klemperer, 1992). To accommodate these forces in our estimation, we use the simulation-based approach developed in Bajari et al. (2007). Given our demand-side parameter estimates, a rate-setting policy function mapping states to deposit rates, and a marginal cost function, we forward-simulate deposit rates and market shares over many time periods and alternative paths of play to construct expected franchise values at many points in the state space. We recover marginal costs by finding the parameter values that best rationalize the estimated policy as optimal.

We use our demand- and supply-side model estimates to quantitatively understand the importance of sleepy deposits. We consider counterfactuals in which all depositors are always active, eliminating the invest-versus-harvest tradeoff and rendering competition static. This

comparison not only generates interesting economic insights, it is also policy-relevant. Regulators in several countries, most notably the UK's Financial Conduct Authority (FCA), have sought to increase the extent to which banks pass through interest rates to their depositors.[2] Our static counterfactual is a limiting case of the FCA's proposal in which bank deposit betas are constrained to unity and banks compete with each other on constant spreads as opposed to time-varying ones. By forcing banks to simultaneously make all current and future strategic choices, this policy eliminates the dynamic invest-versus-harvest tradeoff otherwise induced by depositor sleepiness.

Our estimates imply that sleepiness contributes significantly to average markups. Markups would be 53% lower on average in the absence of depositor sleepiness, falling from about 68 basis points (bps) to about 32 bps. Moreover, sleepiness changes the cyclicality of markups, making them procyclical. At the zero lower bound, the average markup is about 31 bps. On the other hand, we estimate an average markup of 127 bps in 2023, after short-term interest rates had risen to over 5%. This cyclicality is not simply due to the fact that realized rates systematically undershot expectations throughout most of our sample period. For simple market structures, we can directly compute the equilibrium of our model and demonstrate that markups are cyclical in equilibrium. In other words, invest-versus-harvest incentives vary over the business cycle in a way that generates cyclical markups. When the short rate is high, banks with high market shares harvest their existing deposit bases, raising their markups. This softens competition, allowing banks that are investing in growing their deposit bases to charge higher markups than they otherwise would.

We then show that sleepiness and dynamic competition also change the geographic distribution of markups. In particular, they increase markups in areas that a static view of competition would suggest are competitive, i.e., in low-concentration areas with low Herfindahl-Hirschman indices (HHIs). In these markets, banks endogenously shift towards harvesting their sleepy deposit bases, softening competition for active depositors.

We next turn to examining deposit franchise value. Under our static counterfactual, the average bank's deposit franchise value falls by about 58%. In other words, sleepiness accounts for more than half of the value of the average bank's deposit franchise. We also document a large degree of cross-bank heterogeneity in the dependence of franchise value on sleepiness. In particular, banks with low-quality deposit services and banks with high marginal costs benefit most from sleepiness: moving from the 25th to the 75th percentile of product quality decreases sleepiness dependence by about 10.4 percentage points (62% to 51.6%), whereas moving from the 25th to the 75th percentile of marginal costs increases sleepiness dependence by about 3.4 percentage points (54.9% to 58.3%). Intuitively, sleepiness allows low-quality, high-cost banks to charge relatively high markups without losing too many depositors in the near-term.

[2] Laura Noonan and Lucy Fisher, "UK Regulator Sets Deadline for Banks to Justify Low Interest Rates for Savers," *Financial Times*, July 31, 2023, https://www.ft.com/content/e03abe49-a513-4777-8d90-32e252733690.

Finally, we quantify the consequences of these counterfactual reductions in deposit franchise values for financial stability. For the three largest deposit-taking banks (JP Morgan, Bank of America, and Wells Fargo), we convert franchise value losses into changes in risk-neutral probabilities of default. We estimate substantial increases in default probabilities for all three banks, increasing by about 10 percentage points in normal times and more than 20 percentage points during the 2022-2023 monetary tightening cycle.

We consider two extensions of our model. In the first, we extend our model of banks to allow price discrimination between new and existing depositors by offering short-term bonuses to new customers. This enables banks to partially, though not perfectly, discriminate between active and inactive depositors. The model estimates suggest that, while offering bonuses can be an effective way to attract new customers, banks face large customer acquisition costs. Our estimates suggest that it costs a bank roughly 15 cents to attract an additional dollar of deposits from new depositors. In the second, we allow depositors to make forward-looking choices when they are active. We estimate demand to be more elastic in this specification relative to our baseline results. Although both extensions enrich the model, they do not quantitatively alter our conclusions about how dynamic competition and the corresponding invest-versus-harvest tradeoff affect competition in the banking sector.

The remainder of the paper is as follows. We first conclude the introduction by reviewing the different literatures to which our paper contributes. Section 2 describes our data sources and sample construction, and provides summary statistics of the key variables used in our empirical analysis. Section 3 presents several new stylized facts about depositor behavior that motivate the need for a quantitative model of depositor sleepiness. Section 4 describes our model, which we structurally estimate in Section 5. We use our estimates to analyze policy counterfactuals in Section 6, after which we conclude.

Related literature:

Our paper relates to several strands of literature. First, we contribute to the growing literature at the intersection of banking and industrial organization. Previous studies, including Dick (2008), Egan et al. (2017), Xiao (2020), Whited et al. (2021), Egan et al. (2022), and Wang et al. (2022), have employed standard demand estimation tools from the industrial organization literature to estimate static demand for bank deposits and study the implications for bank stability, monetary policy, and bank value.[3] We extend this line of work by incorporating dy-

[3]These tools have also been applied to the asset side of banks (Buchak et al., 2018; Benetton et al., 2021; Benetton, 2021; Allen et al., 2024). Other related work includes Ackerberg and Gowrisankaran (2006), which analyzes how network externalities shape customer and bank adoption of the ACH payments system; Gowrisankaran and Krainer (2011), which studies the consumer welfare effects of banning ATM surcharges; Aguirregabiria et al. (2024), which examines how bank competition and economies of scope and scale influence the distribution of credit in the United States; and Corbae and D'Erasmo (2020), Corbae and D'Erasmo (2021), Corbae and D'Erasmo (2024), which study banking industry dynamics, including the interactions between bank regulation, competition, and financial stability.

namics into both household demand for deposits and bank competition for those deposits. Our estimation of depositor preferences incorporates information about depositor turnover from new microdata. We show that dynamics substantially alter the competitive equilibrium in the U.S. banking market, raising average markups, protecting low-quality and high-cost banks from competition, and increasing financial stability.

Second, we contribute to the literature in industrial organization on dynamic competition under inertial demand (Farrell and Shapiro, 1988; Beggs and Klemperer, 1992; Padilla, 1995). Recent structural work has highlighted the importance of accounting for dynamics, inattention, and switching costs in both the banking sector (Honka et al., 2017) as well as other retail markets, such as consumer packaged goods (Dubé et al., 2009), digital camcorders (Gowrisankaran and Rysman, 2012), printers (Melnikov, 2013), the gasoline market (MacKay and Remer, 2024), health insurance (Fleitas, 2017; Miller, 2019; Pakes et al., 2021), index funds (Brown et al., 2024), and subscription markets (Einav et al., 2023). Additionally, previous work has emphasized the importance of inattention in shaping consumer mortgage refinancing decisions (Agarwal et al., 2016; Andersen et al., 2020; Zhang, 2022). We show that dynamics significantly alter competitive behavior in the $18 trillion US deposit market. Relative to prior work, our findings offer new insights into the heterogeneous impacts of inertial demand on competition in the cross-section of both firms and markets.

Finally, we relate to the broader literature on bank deposits. One part of this literature studies the sluggish and incomplete response of deposit rates to changes in the federal funds rate (Berger and Hannan, 1991; Neumark and Sharpe, 1992). Drechsler et al. (2017) argue that the sluggish response of rates is driven by bank market power: when the funds rate increases, the opportunity cost of using cash as a substitute is higher, effectively increasing banks' market power and allowing them to raise their deposit spreads. Our model offers a complementary mechanism, generating the sluggish adjustment of deposit rates through changes in banks' invest-versus-harvest incentives as the funds rate changes. Another part of this literature uses microdata to shed light on sources of sleepiness, generally finding that inattention (rational or otherwise) is an important ingredient (Argyle et al., 2025; Lu et al., 2025; Lu and Wu, 2025; Cirelli and Olafsson, 2025). Our model is consistent with several sources of sleepiness, including inattention, and our contribution is to explore the consequences of sleepiness for dynamic competition. A third part of the literature underscores the importance of deposit franchise value for bank runs, particularly during the 2023 Regional Banking Crisis (Drechsler et al., 2023; DeMarzo et al., 2024; Jiang et al., 2024; Blickle et al., 2025). We show that deposit sleepiness substantially raises deposit franchise value, amplifying the importance of the channels these papers study.

2 Data

We combine several different datasets for our analysis. Our primary data sources are: anonymized data on account openings and closures from one of the core bank processing platforms, annual information on deposit volumes at the branch level from the FDIC Summary of Deposits, quarterly information on bank salaries and fixed expenses from the FDIC Call Reports, and weekly information on bank deposit rates at the branch$\times$account-type level from Ratewatch. We also bring in secondary sources of data to supplement portions of the analysis, namely county-level demographics from the Census' American Community Survey.

2.1 Sources

2.1.1 Bank Microdata

Our anonymous microdata comes from several of the core account processing platforms offered by the data provider. The main dataset that we use in our analysis is from one such platform, and includes information on openings, closures, and depositor and account characteristics for 12 million deposit accounts from 89 banks and credit unions. We also use an extended dataset, derived from several of the data provider's platforms, that offers slightly less information about depositor and account characteristics but covers more than 900 banks and credit unions and about 42 million deposit accounts. Using both datasets, we study deposit account openings and closures and how they vary with depositor characteristics.

2.1.2 Summary of Deposits

We obtain Summary of Deposits data from the FDIC website from 1994-2023. In particular, we collect information on branch-level deposit volumes for each bank branch as of June 30th of each year. Because banks have substantial latitude to allocate deposits across branches within county, we aggregate observations up to the bank-county-year level for all of our analyses.

2.1.3 RateWatch

RateWatch is a commercial vendor that collects high-frequency information on banks' deposit rates across both account types and branches. The data that we obtain from RateWatch has coverage from 2001 through the end of our sample period (June of 2023). To map into our Summary of Deposits data, we construct a bank-county-year panel of average deposit rates for several different account types: checking accounts, savings accounts, $10k money market deposit accounts, $25k money market deposit accounts, 3-month Certificates of Deposit, and 12-month Certificates of Deposit. For each rate type, we obtain a single rate for each bank-county-year cell by averaging over all reported rates for the bank's branches within that county from April-June of that year.

Our reported estimates throughout the paper use rates constructed from $10k money market deposit accounts, which we have found to be most reflective of the average rate over banks' deposit mixes (i.e., including their non-interest-bearing deposit volumes).[4]

2.1.4 Call Reports

We collect quarterly data on banks' salaries and fixed expenses from the FDIC Call Reports. These cost data are important for identifying the structural parameters corresponding to our household deposit demand system. We construct two bank-year level cost shifters used for our analysis: (i) cumulative salary expenses from June of the prior year to June of the current year, deflated by number of employees as of June of the prior year; (ii) cumulative non-interest expenses on premises and fixed assets from June of the prior year to June of the current year, deflated by total assets as of June of the prior year.

2.2 Sample Construction

We construct a bank×county×year level panel containing: (i) total deposits, from the Summary of Deposits; (ii) average offered deposit rates, from RateWatch; (iii) bank-level cost shifters, from the Call Reports; (iv) numbers of new accounts and total active accounts, from the micro-data; and (v) county-level demographics, from the 2019 American Community Survey.

Because we need to account for the effects of interest accumulation on deposit volumes throughout our empirical analyses, which requires knowledge of deposit rates from the prior year, we restrict our sample to observations for which we have bank-county level deposit rate and deposit volume information for both the current and the prior year.[5] To reduce noise in our forward simulations, we further restrict the sample to bank-county pairs appearing in at least 5 years of data. Finally, we drop any remaining single-bank markets. Our final sample contains 341,395 observations covering 2002-2023.

2.3 Summary Statistics

Table 1 displays summary statistics for the key variables used throughout our empirical analysis. The average deposit rate in our sample is 34 bps, and the average spread between the federal funds rate and deposit rates is 97 bps.[6] The average bank-county observation in our sample has a market share of 16%.

[4] In untabulated results, we verify that both our reduced-form and structural estimates are quantitatively similar under various alternative constructions of deposit rates.

[5] As a result, we have to drop observations from 2001 – the first year for which we have RateWatch data. Our sample will also therefore omit the first year in which a bank operates in a county.

[6] Deposit rates in our sample are lower than what one would infer from deposit interest expenses in the Call Reports – the primary reason for this discrepancy is that we use rates offered on small-balance savings accounts, whereas a large share of call report deposits are certificates of deposit and wholesale deposits, both of which tend to offer significantly higher interest rates.

3 Motivating Evidence

The premise of our model and paper is that depositors are "sleepy"—they are often inactive for long periods and shop infrequently. Here, we use our turnover data to document the persistence in account holdings and understand the potential drivers of this persistence. We examine both account openings and closures.

3.1 Account Openings

We first show that depositors rarely open new accounts, examining account turnover and average account life. We calculate turnover by dividing the total number of new accounts of that type, across all banks in a given year, by the total number of existing accounts of the same type for that year. Average account life is given by the reciprocal of the turnover rate.

Figure 1 displays average account life for different account types. Turnover is similar for checking and savings accounts. New checking and savings accounts make up roughly 8-15% of existing accounts, translating to an average life of about eight years. Turnover for time deposit accounts is higher. New time deposit accounts comprise about 10-20% of all active time deposit accounts, corresponding to an average life of about six years. This higher turnover for time deposits is intuitive, as time deposit accounts have a natural maturity date, which prompts depositors to open and close accounts.

As one would expect, business accounts and accounts with large initial balances exhibit higher turnover, while accounts held by people over the age of 65 feature much lower turnover. Somewhat surprisingly, internet banking accounts exhibit lower turnover. One might expect internet banking customers to be more sophisticated and thus more likely to shop around. However, our subsequent results in Section 3.2 shed light on this: a common reason depositors switch accounts is relocation. Customers without online banking typically need to open a new account when they move to an area in which their current bank does not have any branches, whereas online banking allows them to continue to use their account even in the absence of a nearby branch.

To systematically summarize the cross-sectional drivers of turnover, we classify bank accounts by their key attributes. For our main sample, we cut the data by five characteristics: account type (checking, savings, or time deposits), holder type (person, business, or trust), online banking setup (yes or no), initial balance above $50,000 (roughly the top 10% of accounts), and whether the account-holder is over 65 years old (for personal accounts). For our extended sample, we cut the data by account type (checking, savings, or time deposits), holder type (person or business), and whether the account-holder is over 65 years old (for personal accounts). As a result, we have 72 unique account profiles (3x3x2x2x2) in our main sample and 12 unique account profiles (3x2x2) in our extended sample. Denoting account profiles by i, we calculate account turnover for each profile-by-bank-by-year combination:

$$Turnover_{ijt} = a_t + a_j + X'_{ijt}b + e_{ijt.} \quad (1)$$

We control for the account characteristics used to construct the account profile cells in X_{ijt}. We also include bank fixed effects (a_j), time fixed effects (a_t), and bank-time fixed effects (a_{jt}) in our most stringent specification.

Table 2 displays the regression results corresponding to equation (1). Columns (1)-(4) display results using our main sample, and columns (5)-(8) display results using our extended sample. Consistent with Figure 1, turnover varies by account type. Column (4) indicates that, within a bank-year, savings accounts on average have turnover rates 1.3 percentage points higher than checking accounts, while turnover among time depositors is 5.1 percentage points higher than for individuals with checking accounts. Turnover is 4.8 percentage points lower for account holders aged 65 and older, suggesting that older customers are less likely to shop around for banks. Turnover is 4.1 percentage points higher for business accounts and 7.2 percentage points higher for larger accounts. Our results in the extended sample are qualitatively similar to the results in the main sample, indicating that account turnover activity in our main sample is broadly representative of turnover in the larger population of banks and credit unions.

The low average account turnover we observe could be due to either (i) persistent depositor preferences over banks or (ii) frictions like depositor sleepiness, switching costs, or search costs. The cross-sectional variation in account turnover we document within bank-year in Table 2 seems more consistent with frictional explanations. For instance, consider our result that business and large accounts at the same bank at the same time have higher turnover than other accounts. This is more likely explained by lower search and switching costs than by the idea that these depositors have more volatile preferences. Similarly, older retail depositors likely exhibit lower turnover because they face higher search and switching costs, not because their preferences are more persistent than others.

3.2 Account Closings

We next examine why depositors switch banks, focusing on account closings. Our main sample of microdata provides details on the reasons customers close bank accounts. When an account is closed, the bank typically reports the reason in a standardized format. For 75% of closures, we observe a reason, along with a brief description, which we categorize into five different classifications.[7]

Figure 2 displays a histogram of the reasons why customers close bank accounts. The most common reason is "inactivity," where the bank closes the account on the customer's behalf. Other common reasons include the account no longer being needed by the customer or the

[7]We focus on closures where the depositor provided a reason and exclude those resulting from internal transfers within the bank (i.e., when the depositor closed an account due to opening another within the same bank). In 25% of closures, no reason is provided, and internal transfers account for 18% of all closures.

customer moving or passing away. Collectively, moves and deaths account for over 20% of closures, with half attributed to moving.

Conversely, depositors report switching to a bank offering better rates, services, or fees in only 17% of cases. These findings suggest that, while some consumers actively shop for better banking options, many are inactive and close accounts for more idiosyncratic reasons.

4 Model

We develop and estimate a tractable dynamic model of demand and supply for deposits. A key feature of the model is that we allow the persistence of deposit demand to be driven by both product quality persistence and depositor sleepiness.

In the model, depositors are either awake (i.e., active) or asleep (i.e., inactive) each period, following the basic setup of Beggs and Klemperer (1992) and Brown et al. (2024). Whether a depositor is awake or asleep in a given period is an exogenous process, which can be thought of as a reduced-form approximation to switching costs or inattention. Depositors who are awake re-optimize and select the bank that maximizes their utility, while depositors who are asleep maintain their banking relationship from the previous period.

Banks set deposit rates, dynamically competing for depositors to maximize the franchise value of the bank. Due to depositor inactivity, banks may face "invest-versus-harvest" incentives. On one hand, inactive depositors mechanically make deposit demand less elastic. This reduced elasticity incentivizes banks to offer lower deposit rates and "harvest" their existing depositor base. On the other hand, depositor sleepiness increases the lifetime value of attracting a depositor today. Each depositor the bank attracts is, in expectation, likely to remain a depositor in the future due to sleepiness. This effect of sleepiness on the expected length of the depositor-bank relationship encourages banks to offer higher deposit rates and "invest" in their depositor base.

4.1 Depositors

Each period, depositors are either awake or asleep. Awake depositors select the bank that maximizes their indirect utility, while asleep depositors maintain their banking relationship from the previous period. We model an awake depositor's bank choice as a discrete choice problem following Berry (1994).

4.1.1 Depositor Preferences

Households demand deposit accounts. Active depositors select the bank j from the set $\mathcal{J}_k$ available in market k to deposit their funds. Depositors derive utility from promised deposit rates and other bank services. The indirect utility derived by depositor i from bank j in market

k at time t is given by:

$$u_{ijkt} = \alpha\rho_{jkt} + \delta_{jkt} + \varepsilon_{ijkt}.$$

Without loss of generality and for ease of exposition, we write the utility of a depositor in terms of deposit spreads (ρ_{jkt}) rather than deposit rates. The deposit spread is defined as the difference between the risk-free rate R_t^F and the deposit rate: $\rho_{jkt} := R_t^F - r_{jkt}$. The spread then reflects the "price" that depositors pay for banking services. The parameter $\alpha < 0$ measures the disutility depositors experience from paying higher spreads (i.e., prices).

Depositors also derive utility from banking services, captured by $\delta_{jkt} + \varepsilon_{ijkt}$. The term δ_{jkt} reflects vertical differentiation at the bank-by-market-by-time level: bank quality differences that vary across banks, markets, and over time. For example, Bank of America may offer superior services relative to Citibank in a specific market, for instance due to the location of their branches. The term ε_{ijkt} captures horizontal differentiation in the marketplace. In other words, depositors do not necessarily agree on which bank is the best in a given market.

One could interpret this setup in two ways. Either depositors make myopic choices, with u_{ijkt} capturing the indirect utility received in the current period. Or they make forward-looking choices, with u_{ijkt} capturing the full expected discounted utility of selecting a bank in the current period, acknowledging that they may be inactive and thus unable to reoptimize in subsequent periods. For the purposes of presenting the model, we do not need to take a stand on the interpretation of u_{ijkt}. For estimation and counterfactuals, the interpretation matters. In our baseline specification, we assume that depositors are making a myopic choice, though we explore the case of fully forward-looking demand in Section 5.4.2.

4.1.2 Depositor Sleepiness

Depositors enter each period as either active or inactive. We let the indicator variable D_{it} denote whether a depositor i is active at time t. Whether a depositor is active depends on the latent variable D_{it}^*, where:

$$D_{it}^* = \mathbf{S}_{k(i)t}'\Gamma + X_{it}'\Theta + \eta_{it}, \qquad D_{it} = \mathbb{1}(D_{it}^* > 0). \tag{2}$$

The term $\mathbf{S}_{k(i)t}$ is a vector of market-level and aggregate state variables (e.g., the federal funds rate) that influence whether a depositor in market k is active at time t. The parameter Γ measures how these state variables impact whether a depositor is awake or asleep. For example, depositors may be more likely to become active when aggregate interest rates are high. The term X_{it} is a vector of depositor characteristics that impact whether a depositor is awake or asleep, such as the depositor's age. Lastly, the term η_{it} is a depositor-specific idiosyncratic shock, generating heterogeneity in which depositors are active versus inactive.

The share of depositors that are inactive in a given market k at time t is given by

$$\phi_{kt} = 1 - \mathbb{E}[D_{it}|\mathbf{S}_{k(i)t}],$$

where we take the expectation over the idiosyncratic shocks η_{it} and depositor demographics X_{it}. When we estimate the model, we estimate the parameter ϕ_{kt} as a function of market-level and aggregate state variables $\mathbf{S}_{kt}$.

Our model of depositor activity through the parameter ϕ_{kt} can be thought of as a reduced-form approach that captures several underlying frictions, such as inattention or switching costs. For example, a depositor may be rationally inattentive until the federal funds rate is high enough to justify acquiring more information and searching for a new bank. Our model captures these effects through the term $\mathbf{S}'_{kt}\Gamma$. Similarly, the parameter ϕ_{kt} may also capture switching costs, as depositors may not switch banks until the return from doing so is sufficiently large to offset these costs.[8]

4.2 Banks

Banks engage in dynamic oligopolistic competition to maximize the present discounted value of all current and future profits. Each period t, bank j collects deposits in market k at interest rate r_{jkt} and invests them in a constant-returns-to-scale technology $R_{jt} \sim \mathcal{N}(\mu_{jt}, \sigma_{jt})$.

Banks face two types of marginal costs. First, they must pay the promised deposit rate r_{jkt} to each depositor. Second, they incur a constant marginal cost c_{jkt} of collecting deposits. This cost consists of both explicit expenses for servicing depositors (e.g., variable wages for bank tellers) and any implicit benefits or costs of acquiring an additional deposit.

The bank's flow profits, denoted π_{jt}, at time t are given by:

$$\pi_{jt} = \sum_{k \in \mathcal{K}} Dep_{jkt}((R_{jt} - R_t^F) + \rho_{jkt} - c_{jkt}).$$

The bank collects deposits Dep_{jkt} in each market k, and $\mathcal{K}$ denotes the set of markets. The marginal profit associated with collecting an additional dollar of deposits for bank j in market k at time t is given by $R_{jt} - R_t^F + \rho_{jkt} - c_{jkt}$, where $R_{jt} - R_t^F$ captures excess lending returns over the short rate, $\rho_{jkt} = R_t^F - r_{jkt}$ captures the benefit of funding at the deposit rate rather than the short rate (i.e., the deposit spread), and c_{jkt} captures the non-deposit-related marginal cost of an additional dollar of deposits.

The above assumes that banks must offer the same deposit rate to all depositors in a market. In reality, banks may price-discriminate between their existing ("incumbent") depositors, many

[8] Of course, the exact moment of the rate distribution that drives depositor behavior depends on the exact micro-foundation. In a switching cost model, the maximum rate matters, while in a search model, conditional expected rates matter.

of whom are inactive in a given period, and depositors they are trying to win over from other banks, all of whom must by construction be active in that period. In Appendix E, we develop and estimate an extension in which banks can offer separate rates in each period to their incumbent and new depositors, potentially incurring different marginal costs when servicing new versus incumbent depositors. This extension allows banks to partially, though not perfectly, price discriminate across active and inactive depositors. The model estimates suggest that, while offering bonuses can be an effective way to attract new customers, banks face large customer acquisition costs.

Bank j's franchise value at time t, denoted V_{jt}, is given by the present discounted value of all current and future profits:

$$V_{jt} = \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^s \pi_{js}\Big],$$

where the bank discounts future profits by the discount factor $\beta < 1$.

4.3 Equilibrium

Following Bajari et al. (2007), we study stationary Markov Perfect Equilibria in pure strategies. Under this solution concept, each bank's equilibrium choices in each market are functions only of: (i) the current values of the game's payoff-relevant states, which are commonly observed by all market participants; and (ii) an i.i.d. private cost shock drawn at the start of each period. These assumptions allow us to avoid explicitly estimating banks' first order conditions. Instead, we rely simply on the assumption that their actions yield higher franchise value than other possible actions.

4.3.1 Depositors

In equilibrium, all active depositors select the bank that maximizes their indirect utility, while inactive depositors maintain their choice from the previous period. Following the demand estimation literature, we assume that the depositor-bank-market-time specific utility shocks (ε_{ijkt}) are distributed as Type-1 Extreme Value (Berry, 1994; Berry et al., 1995). This depositor optimization implies that the market share of bank j in market k at time t among active depositors (i.e., the probability an active depositor selects bank j in market k at time t) is given by:

$$s_{jkt}^{Active} = \frac{\exp(\alpha\rho_{jkt} + \delta_{jkt})}{\sum_{\ell \in \mathcal{J}_k} \exp(\alpha\rho_{ikt} + \delta_{\ell kt})}, \tag{3}$$

where s_{jkt}^{Active} denotes the bank's market share among active depositors.

Inactive depositors select the same bank as in the previous period. Thus, total deposits

collected by bank j in market k at time t are given by:

$$Dep_{jkt} = \underbrace{(1-\phi_{kt})M_{kt}s_{jkt}^{Active}}_{\text{Active Depositor Demand}} + \underbrace{\phi_{kt}(1+r_{jkt-1})Dep_{jkt-1}}_{\text{Inactive Depositor Demand}}, \tag{4}$$

The first term on the right hand side represents active depositor demand, which is equal to the share of depositors who are awake, $1-\phi_{kt}$ multiplied by the size of the deposit market, M_{kt}, and the bank's active market share, s_{jkt}^{Active}. The second term represents inactive depositor demand, which is equal to the deposits collected by the bank in the previous period, $Deposits_{jkt-1}$, grossed up by the interest they earned $(1+r_{jkt-1})$, multiplied by the fraction of depositors who are inactive, ϕ_{kt}.

4.3.2 Banks

Banks compete for deposits by playing a dynamic, Markovian rate-setting game: they formulate strategies as functions of current, payoff-relevant states. In each market k, the state space consists of:

1. *exogenous public states* $\mathbf{\Omega}_{kt}$: objects that are invariant to banks' actions (e.g., the federal funds rate, bank-specific quality δ_{jkt}).
2. *endogenous public states* $\mathbf{W}_{kt}$: objects that are functions of banks' past actions (e.g., market shares in k)
3. *exogenous private states* $\boldsymbol{\chi}_{kt}$: i.i.d. cost shocks drawn by each bank in the market at the start of each period. We assume that the distribution of these shocks is common over banks, markets, and time, and that it is known by all participants.

Let $\mathbf{S}_{kt} \coloneqq (\mathbf{\Omega}_{kt}, \mathbf{W}_{kt})$ denote the *ex-ante* state of the game for a given market and period (i.e., before cost shocks are realized). Bank j's franchise value from choosing a spread-setting strategy $\sigma_j : \mathbf{S}_{kt} \times \chi_{jkt} \to \rho_{jkt}$, given its conjecture of other banks' rate-setting strategies $\boldsymbol{\sigma}_{-j}$, follows the recursion:

$$\begin{aligned} V(\mathbf{S}_{kt}, \boldsymbol{\sigma}) = {} & \mathbb{E}_{\boldsymbol{\chi}_{jt}, R_{jt}}[Dep(\mathbf{S}_{kt}, \boldsymbol{\sigma})(R_{jt} - R_t^F + \sigma_j(\mathbf{S}_{kt}, \chi_{jt}) - c(\mathbf{S}_{kt}, \chi_{jt})) \,|\, \mathbf{S}_{kt}] \\ & + \beta \mathbb{E}_{\mathbf{S}_{kt+1}, \boldsymbol{\chi}_{t+1}}[V(\mathbf{S}_{kt+1}, \boldsymbol{\sigma}) \,|\, \boldsymbol{\sigma}, \mathbf{S}_{kt}] \end{aligned} \tag{5}$$

where $Dep(\cdot)$ re-casts equation (4) as a function of the state.[9] Under a well-defined state transition process: $T : \mathbf{S}_{kt} \times \boldsymbol{\rho}_{kt} \to \mathbf{S}_{kt+1}$, the recursion in equation (5) is well-defined.

[9] That is,

$$Dep(\mathbf{S}_{kt}, \boldsymbol{\sigma}) \coloneqq \phi(\mathbf{S}_{kt}) M(\mathbf{S}_{kt}) s^{Active}(\mathbf{S}_{kt}, \boldsymbol{\sigma}) + (1-\phi(\mathbf{S}_{kt}))(1+\sigma(S_{kt-1}, \chi_{t-1}))Dep(\mathbf{S}_{kt-1}, \boldsymbol{\sigma})$$

Importantly, our specification of the private shocks $\boldsymbol{\chi}_{kt}$ simplifies the problem considerably: shocks can affect banks' current actions and current costs, but by the Markov property these are the only channels through which they can affect future outcomes of the game. Moreover, because the shocks are i.i.d., we need not condition on past realizations in the continuation values in equation (5), which would be difficult given that the shocks are unobservable.[10]

In equilibrium, all banks must prefer their current strategy to any other Markovian strategy, holding fixed conjectures about opponents' play:

$$V(\mathbf{S}, \boldsymbol{\sigma}) \geq V(\mathbf{S}, \sigma_j', \boldsymbol{\sigma}_{-j}) \quad \forall\, \mathbf{S} \in \mathcal{S},\ \forall\, \sigma_j',\ \forall\, j \in \mathcal{J},$$

In our empirical analysis, we restrict our attention to symmetric equilibria.[11]

4.4 Discussion: Dynamic Versus Static Competition

Equation (5) can be used to illustrate intuitive differences between static and dynamic competition. Consider a single market of constant size $M = 1$ that is populated by J banks. Each bank j enters the current period with lagged market share $s_{j,-}$ and has constant marginal cost c_j and constant product quality δ_j. Suppose that there is no aggregate risk (i.e., the short rate R_t^F is constant).

Differentiating equation (5), bank j's equilibrium markup is given by :

$$\underbrace{\rho_j - c}_{\text{markup}} = \underbrace{\frac{-1}{\alpha(1 - s_j^{Active})}}_{\text{static pricing rule}} + \underbrace{\frac{\phi}{1-\phi}\frac{s_{j,-}}{s_j^{Active}}\frac{-1}{\alpha(1 - s_j^{Active})}}_{\text{harvesting incentive}} + \underbrace{\beta \nabla V_j(\mathbf{s} \mid \boldsymbol{\delta}, \mathbf{c}) \begin{pmatrix} \frac{s_1^{Active}}{1 - s_j^{Active}} \\ \vdots \\ -1 \\ \vdots \\ \frac{s_K^{Active}}{1 - s_j^{Active}} \end{pmatrix}}_{\text{investing incentive}} \tag{6}$$

The first term on the right hand side of equation (6) reflects the standard, statically optimal "inverse-elasticity" pricing rule. The second term reflects the bank's incentive to harvest its installed base of sleepy deposits. The relative strength of the harvesting incentive is increasing in sleepiness ϕ. Moreover, the harvesting incentive is always non-negative: at the static equilibrium outcome, if $\phi > 0$, the harvesting incentive exerts upward pressure on markups.

The third term captures the bank's incentive to invest in growing its deposit base. The

[10] Berry and Compiani (2023) develop a way to accommodate serially correlated unobservables in settings such as ours, but their approach comes at the cost of no longer point-identifying the parameters of the policy function.

[11] It is, in principle, feasible to relax this assumption by estimating separate rate-setting policy functions across banks or markets. However, the second stage of the Bajari et al. (2007) procedure has been shown to be highly sensitive to policy function approximation error. We therefore prefer to pool observations when estimating the policy function to reduce noise.

relative strength of bank j's investment incentive is a function of three objects: (i) the discount rate; (ii) the sensitivity of bank j's value function to the vector of lagged market shares; and (iii) the sensitivity of market shares to bank j's spread.

The importance of the discount rate is intuitive, as the benefits from investing are realized in the future. The sensitivity of the value function to lagged market shares arises because a fraction of today's depositors will be asleep tomorrow. If all depositors were awake, our model would become static and this sensitivity would be zero. The final object on which the investment incentive depends, the sensitivity of market shares to bank j's deposit spread, is important because it scales the afore-mentioned change in the value function. Given the logit structure of our demand system, this object is simply the vector of diversion ratios from bank j to all other banks in the market.

Equation (6) shows that it would be difficult to obtain closed-form results relating equilibrium markups to model primitives. Strategic interactions across time periods, encoded by the term $\nabla V(\mathbf{s} \mid \boldsymbol{\delta}, \mathbf{c})$, are shaped by all banks' product qualities and marginal costs. The richness of these strategic interactions can overturn intuitions obtained from static models, highlighting the importance of accurately modeling the dynamic nature of deposit market competition.

5 Estimation

A key strength of our model is that it is dynamic, quantitative, and relatively straightforward to estimate. The parameters to estimate are the sleepiness function $\phi(\mathbf{S}_{kt})$, depositor preferences $(\alpha, \boldsymbol{\delta}_{jkt})$, and bank marginal costs $(\mathbf{c}_{jkt})$. We estimate the model in three steps.

First, we estimate depositor sleepiness. Second, given our estimates of depositor sleepiness, we estimate the preferences of active depositors following Berry (1994). Third, to estimate marginal costs, we employ the two-step approach developed by Bajari et al. (2007). We flexibly estimate the mapping from banks' observed states to the deposit spreads they set. Marginal cost parameters are then estimated by exploiting the equilibrium condition that the observed spread-setting policy function maximizes all banks' expected franchise values.

5.1 Depositor Sleepiness

In our estimation, we allow sleepiness to vary depending on depositor type and market characteristics as motivated by our reduced-form analysis in Section 3. Depositor types, denoted ℓ, are defined by a vector of demographics $\mathbf{X}_{\ell kt}$. For depositor type ℓ, bank j, county k, and year t, total deposits are given by the following identity:

$$Dep_{\ell jkt} = \underbrace{(1 - \phi(\mathbf{S}_{kt}, \mathbf{X}_{\ell kt}))M_{\ell kt}s_{jkt}^{Active}}_{\text{Active Depositor Demand}} + \underbrace{\phi(\mathbf{S}_{kt}, \mathbf{X}_{\ell kt})(1 + R_{t-1}^{F} - \rho_{jkt-1})Dep_{\ell jkt-1}}_{\text{Inactive Depositor Demand}} \tag{7}$$

where $M_{\ell kt}$ refers to the total market size within depositor type l and s_{jkt}^{active} refers to bank j's market share among active depositors.[12] We estimate depositor sleepiness jointly using two sets of moments, one from our microdata on account turnover and the second the autocorrelation of branch-level deposits.

5.1.1 Account-Level Turnover

The challenge in estimating equation (7) from our microdata is that we do not directly observe active demand. However, our data on new account openings allow us to compute active demand under our model assumptions. New account openings at bank j at time t equal the number of depositors who were active at t and selected bank j, but did not previously bank with bank j, such that:

$$New\,Dep_{\ell jkt} = (1 - \phi(\mathbf{S}_{kt}, \mathbf{X}_{\ell kt}))M_{\ell kt}s_{jkt}^{Active}(1 - s_{jkt-1}),$$

which implies

$$\text{Active Depositor Demand} = New\,Dep_{\ell jkt}\frac{1}{1 - s_{jkt-1}}.$$

We are relying on the assumption of our model that, conditional on average bank utility, depositors' idiosyncratic preferences ε_{ijkt} are independent over time. Otherwise, depositors previously at bank j who are awake in period t would have a different propensity to select bank j than other depositors awake in the same period.[13]

With this assumption, all objects in equation (7) are observed. To proceed with estimation, we divide by deposits, $Dep_{\ell jkt}$ and rearrange to obtain:

$$1 - \frac{New\,Dep_{\ell jkt}}{Dep_{\ell jkt}}\frac{1}{1 - s_{jkt-1}} = \phi(\mathbf{S}_{kt}, \mathbf{X}_{\ell kt})\frac{Dep_{\ell jkt-1}}{Dep_{\ell jkt}} + e_{\ell jkt} \tag{8}$$

where $e_{\ell jkt}$ represents measurement error that is "outside" the model. Equation (8) shows that depositor sleepiness closely maps to the reduced-form evidence presented in Section 3. The dependent variable is a scaled version of our measure of deposit turnover, $\frac{New\,Dep_{\ell jkt}}{Dep_{\ell jkt}}$, where the structural model informs the appropriate scaling parameter.

When estimating these moment conditions using our micro-data, we measure total and new deposits in terms of deposit accounts, defined at the depositor-type level.[14] We define depositor types as demographic cells based on whether the individual is over age 65, opened the account with an initial balance less than $50,000, and signed up for internet banking. For example,

[12] Market shares are constant across depositor types, given our assumptions on active depositor preferences.

[13] As we discuss below, Figure 3 shows that the model closely matches an empirical moment that distinguishes sleepiness from persistent preferences: the pace at which banks entering a new market accumulate market share.

[14] This explains why equation (8) omits the term $(1 + R_{t-1}^{F} - \rho_{jkt-1})$ that appears in (7). (7) is in dollar units, while (8) measures numbers of accounts.

one cell would represent accounts held by individuals over 65, with an initial balance below $50,000, who enrolled in internet banking.

To reduce noise and simplify the estimation, we set the right-hand side regressor of equation (8), $\frac{Dep_{\ell jkt-1}}{Dep_{\ell jkt}}$, to .95. That is, we assume a growth rate of about 5%. We also aggregate equation (8) to the bank-cell-year level (i.e., aggregate across markets), replacing s_{jkt-1} with s_{jt-1}, the bank's market share across all of its markets.

Finally, we parameterize $\phi(\cdot)$ to be linear in states and further assume that that $\phi(\cdot)$ only depends on the lagged federal funds rate and depositor characteristics: $\phi(\mathbf{S}_{kt}, \mathbf{X}_{\ell kt}) = \boldsymbol{\Upsilon}_1'\mathbf{S}_t + \boldsymbol{\Upsilon}_2'\mathbf{X}_{\ell t}$. Thus, our estimated model more closely captures depositor inattention, rational or otherwise, rather than search or switching costs — in search and switching cost models, $\phi(\cdot)$ depends on endogenous variables, not just exogenous state variables. Our assumption fits our motivating evidence in Section 3.2, which showed that account closures were mostly driven by idiosyncratic reasons rather than search. This distinction has a limited impact on the estimation step, but it influences the real-world policies that correspond to the counterfactuals we consider.[15]

5.1.2 Autocorrelation in Branch Deposits

To further discipline our estimation, we use bank-county level panel data on deposit volumes from the Summary of Deposits. Equation (7) suggests that in this data we could estimate $\phi(\mathbf{S}_{kt}, \mathbf{X}_{kt}) = \boldsymbol{\Upsilon}_1'\mathbf{S}_{kt} + \boldsymbol{\Upsilon}_2'\mathbf{X}_{kt}$ by running the following regression:

$$Dep_{jkt} = (\boldsymbol{\Upsilon}_1'\mathbf{S}_{kt} + \boldsymbol{\Upsilon}_2'\mathbf{X}_{kt})(1 + R^F_{t-1} - \rho_{jkt-1})Dep_{jkt-1} + \nu_{jkt}.$$

However, this regression is only unbiased if lagged deposits are uncorrelated with active demand (i.e., the first term in equation (7)). This assumption does not hold because of vertical differentiation across banks. If some banks provide superior average deposit services, these banks will have higher deposits and higher spreads both yesterday and today, regardless of depositor sleepiness.

We adopt a control function approach to address this issue, regressing spreads on bank-level cost shifters to recover latent demand:

$$\text{Control Function:} \quad \rho_{jkt} = \lambda \mathbf{Z}_{jt} + \underbrace{v_{jkt}}_{\text{Latent Demand}} \tag{9}$$

$$\text{Second Stage:} \quad Dep_{jkt} = (\boldsymbol{\Upsilon}_1'\mathbf{S}_{kt} + \boldsymbol{\Upsilon}_2'\mathbf{X}_{kt})(1 + R^F_{t-1} - \rho_{jkt-1})Dep_{jkt-1} + H(v_{jkt}) + \iota_{jkt} \tag{10}$$

[15] For instance, consider a counterfactual setting $\phi = 0$. If ϕ reflects inattention, the most relevant policy might be an awareness campaign about rates. If ϕ reflects switching costs, the most relevant policies would facilitate opening and closing accounts.

We use two cost shifters, $\mathbf{Z}_{jt}$: (1) salaries, calculated as wage expenses over the past four quarters divided by the lagged number of employees, and (2) fixed expenses, summed over the past four quarters and divided by lagged assets.

This approach requires two key identifying assumptions.[16] First, the control function must be correctly specified: $\mathbb{E}[v_{jkt} \mid \mathbf{Z}_{jt}] = 0$. Second, latent demand must be correctly specified:

$$H(v_{jkt}) + \iota_{jkt} = \nu_{jkt}$$
$$\mathbb{E}[(1 + R^F_{t-1} - \rho_{jkt-1}) Dep_{jkt-1} \cdot \iota_{ijkt}] = 0$$

Under these assumptions, we will consistently estimate the parameters $\boldsymbol{\Upsilon}$.

Finally, to facilitate comparability across markets with different amounts of deposits, we estimate equation (10) in terms of market shares rather than the level of deposits.

5.1.3 Results

We report our estimates of depositor sleepiness in Table 3. Columns (1)-(2) rely solely on depositor cell-level microdata (our main sample), columns (3)-(4) present estimates using only the bank-county level data, and columns (5)-(6) combine both sets of moments.

Column (1) shows that average inertia implied by our microdata is 86.8%. In other words, we estimate that just under 87% of depositors are inactive at a given point in time.[17] Column (2) allows inertia to vary as a function of the federal funds rate and account-level demographics. Mirroring our reduced-form findings, we find that older account holders and those with internet banking are substantially more inertial, whereas account holders with larger balances are substantially less inertial. Inertia also falls as the short rate rises: a one percentage point increase in the short rate reduces the fraction of depositors who are inactive by 0.4 percentage points.

Columns (3) and (4) present control function estimates using only the Summary of Deposits data. In column (3), we find that an additional dollar of lagged deposits increases current deposits by 97.7¢. Thus, we find broadly similar inertia estimates leveraging different sources of variation across two different data sets. Column (4) adds an interaction between lagged deposits and the lagged federal funds rate. As in column (2), we find that inertia declines with the short rate.

Finally, columns (5) and (6) report GMM estimates using both sets of moments and find broadly similar results.

[16] We can relax these assumptions in our estimation by flexibly controlling for increasingly higher-order terms of the cost shifters in the control function and the residuals in the second stage.

[17] By comparison, Brown et al. (2024) find that about 87% of retail index fund investors are inactive.

5.1.4 Sleepiness versus Persistent Unobserved Preference Heterogeneity: Evidence from Market Entry

We now provide a brief validation of our assumption that inertia is driven by sleepiness rather than by persistent unobserved preference heterogeneity. In our model, we allow observed and unobserved bank quality to be correlated over time; however, we assume that conditional on being awake, unobserved preference heterogeneity is not persistent (i.e., the depositor-specific utility shocks, ε_{ijkt}, are i.i.d. across periods). In general, it can be difficult to determine whether the persistence in deposits arises from inattention or from persistent unobserved preference heterogeneity. A key distinction between these two frictions, which we exploit, lies in the speed of adjustment to new entrants. Persistent preference heterogeneity implies that an entrant should immediately attain its steady-state market share. In contrast, depositor sleepiness implies that the entrant can access only $1-\phi_t \approx 6.5\%$ of the market in each period. As a result, a bank is limited in how quickly it can raise deposits, even if it offers high-quality services and competitive deposit rates.

In Figure 3, we compare how fast entrant banks achieve their market shares at year 10 in the data to their speed in a simulation of our model with our estimated preferences and degree of sleepiness. One can think of this as an untargeted moment from the model. The two series are very similar, indicating that depositors are sleepy. Moreover, the similarity in the curvature of the paths in the data and our simulated model corroborates our estimated degree of sleepiness.

5.2 Depositor Preferences

5.2.1 Estimation

We estimate the preferences of depositors following Berry (1994). Equation (3) implies we can estimate depositor preferences using the following regression:

$$\log(s_{jkt}^{Active}) = \alpha\rho_{jkt} + \delta_j + \mu_{kt} + e_{jkt}. \tag{11}$$

Persistent unobserved heterogeneity in product quality that is not absorbed by bank fixed effects will again bias OLS estimates of equation (11). In addition, firms may take into account idiosyncratic demand shocks that we cannot observe when setting prices.

We use the same control function approach as above to address these endogeneity concerns. The identifying assumptions are that: (i) fluctuations in salaries and fixed expenses affect banks' spread-setting decisions; and (ii) spreads are the only channel through which these cost shifters affect market shares. In other words, it cannot be the case that the cost shifters are systematically related to unobserved demand shocks.

To estimate equation (11) we need each bank's market share among active depositors, which may differ from observed overall market shares. Following Brown et al. (2024), we convert

overall market shares to active market shares using our sleepiness estimates from Section 5.1:

$$Dep_{jkt}^{Active} := \max\left\{0, Dep_{jkt} - \boldsymbol{\Upsilon}' \begin{pmatrix} \mathbf{S}_{kt} \\ \tilde{\mathbf{X}}_{kt} \end{pmatrix} (1+R_{t-1}^{F}-\rho_{jkt-1}) Dep_{j,k,t-1}\right\}, \qquad s_{jkt}^{Active} := \frac{Dep_{jkt}^{Active}}{\sum_{\ell \in \mathcal{J}_k} Dep_{\ell kt}^{Active}}$$

where $\tilde{\mathbf{X}}_{kt}$ converts account-level demographics in the data into implied county-level demographics in the American Community Survey.[18][19]

5.2.2 Results

We report our demand estimates in Table 4. To demonstrate the importance of accounting for depositor sleepiness, columns (1) and (2) report estimates assuming that all depositors are active (i.e., using overall market shares),[20] while columns (3) and (4) report estimates accounting for sleepiness (i.e., using active market shares).

Comparing the estimates in columns (2) and (4), we see that accounting for sleepiness results in a 34% higher estimated elasticity of demand.[21] Intuitively, given a change in rates and change in total deposit quantities, the implied elasticity is higher when only a fraction of deposits active and therefore able to react.

In Appendix 5.4.2, we re-estimate demand assuming that depositors are forward-looking. We show that this amounts to adding an additional product characteristic to the demand system: the discounted expected utility of remaining asleep and continuing to bank with today's active choice for many periods. We find qualitatively similar estimates under forward-looking and myopic demand, though the former imply more elastic demand than the latter. More importantly, both estimates indicate that, once we take sleepiness into account, depositors are significantly more elastic than suggested by estimates generated under the assumption that all depositors are always active.

5.3 Bank Costs

We collect each bank's asset returns and marginal cost of producing deposits, $R_{jt} - c(\mathbf{S}_{kt}, \chi_{jt})$, into a single variable, which we term "net marginal costs."[22] We estimate these costs using the

[18] In particular, to extrapolate the account-level demographics used to estimate sleepiness in column (7) of table 3 to the full Summary of Deposits sample, we map each demographic variable to the nearest corresponding demographic in the American Community Survey.

[19] Our model implies that active demand should always be non-negative. However, because we estimate a high degree of inertia and deposit growth is volatile at the bank-county level, active demand is negative in about 14% observations. We drop these observations from the sample when estimating the active demand system.

[20] The estimates in columns (1) and (2) are similar to estimates in prior work.

[21] Elasticities are calculated assuming a market share of 10% and a deposit rate equal to the sample median, which is approximately 15 bps.

[22] If external finance is costly, then our estimated net costs will reflect the marginal risk-adjusted return to lending–part of the value of deposits arises from the fact that they are the marginal source of funding for loans. Conversely, if Modigliani and Miller (1958) holds and external finance is frictionless,then our estimated net cost will not capture

two-step approach developed in Bajari et al. (2007) (henceforth "BBL"). The logic underlying BBL is to bypass solving for the equilibrium of the model by instead assuming that observed bank actions reflect equilibrium play. Section 5.3.1 explains how we parameterize marginal costs and characterizes the state space of the game. Section 5.3.2 specifies and estimates banks' spread-setting policy function. Section 5.3.3 specifies transition functions for the state variables. Section 5.3.4 forward-simulates paths of play to calculate expected franchise values under both the estimated policy function as well as many alternative policy functions, and uses these estimates to recover marginal costs.

5.3.1 Parameterization and State Space

We parameterize net marginal costs to be a linear function of a common intercept, the current short rate, the cost shifters, and the private cost shocks:

$$c_{jkt} \coloneqq \omega + \zeta R_t^F + \gamma' \mathbf{Z}_{jt} + \chi_{jkt}$$

We allow net costs to vary with the short rate for two reasons. First, lending returns may vary with the short rate. Second, it allows us to isolates variation in markups driven by variation in the strength of invest-versus-harvest incentives over time.[23]

Given this cost specification and the Markov concept, the state space in market k at time t consists of: (i) the current short rate R_t^F; (ii) lagged short rate R_{t-1}^F; (iii) lagged deposit shares;[24] (iv) cost shifters (salaries and fixed expenses) $\mathbf{Z}_{kt}$; (v) product qualities $\boldsymbol{\delta}_{kt}$; (vi) county demographics; and (vii) private cost shocks $\boldsymbol{\chi}_{kt}$.

5.3.2 Policy Function Estimation

The first step of the BBL estimation procedure requires us to recover the policy function, characterizing the reduced-form mapping from states to actions. Since we have continuous state and action spaces, we adopt a flexible parametric specification as in Ryan (2012) and Matvos and Seru (2014).[25] Our baseline policy function specification is a first-order polynomial of all

returns to lending because banks essentially operate the deposit and lending businesses separately. This distinction will become important in our counterfactual analyses on financial stability. In particular, if Modigliani and Miller (1958) held, we would need to bring in additional information on the value of a bank's lending franchise in order to properly characterize its counterfactual probability of default.

[23] A large body of recent work has argued that deposit markups are strongly procyclical, e.g., Drechsler et al. (2017, 2021, 2023); Xiao (2020); Whited et al. (2021); Wang et al. (2022). In these papers, the mechanism driving procyclical markups is that the banking sector has more market power when rates are higher because holding cash, the liquid outside option, becomes more expensive. This force is left out of our model, and our parameterization will allow its effects to flow through to costs instead of markups.

[24] We use market shares instead of deposits to ensure that spread-setting behavior does not mechanically change as the market grows in size in the forward simulation.

[25] In order to identify the policy function, we require a scalar unobservable assumption: the vector of unobservable shocks in the policy function is one-dimensional. To satisfy this, we assume that unobservable shocks reflect shocks to the marginal cost of servicing deposits. Given that we recover demand shocks from the estimated demand system

states, which we report in Table 5. Following Ryan (2012), we also include sums of competitors' characteristics (costs, lagged market shares) as regressors.

5.3.3 State Transitions

We close the model by specifying laws of motions over the state space. Endogenous states, by construction, follow deterministic transitions: lagged market shares are updated in accordance with realized market outcomes in the prior period. We assume that banks believe their product qualities δ_{jkt} and costs $\mathbf{Z}_{jt}$ will remain constant at current levels in perpetuity. We make this assumption both for computational ease and because our primary mechanism of interest is unrelated to expected or realized variation in these objects.

We construct a stochastic process for the short rate using instantaneous forward rate estimates published on the Federal Reserve's website.[26] At time t and for time horizons $s - t \in \{1, \cdots, 30\}$, we simulate rates using:

$$R_s^F \sim \mathcal{N}(\mathcal{F}_{t,s}, \hat{\text{Var}}(R^F)),$$

where $\mathcal{F}_{t,s}$ is the estimated instantaneous forward rate at time t, $s - t$ years hence, and where $\hat{\text{Var}}(R^F)$ is the sample variance of the federal funds rate over our sample period.[27] For all time horizons $s - t \in \{31, T\}$, we use the estimated 30-year forward rate.

5.3.4 Forward Simulation

Given the estimated policy function and the state transition processes, we can forward-simulate play to compute the expected franchise value at any point in the state space. Furthermore, we know that the equilibrium policy function maximizes all banks' franchise values. The BBL insight is to invert this condition: we assume that banks are acting optimally, which implies that alternative spread-setting policies should generate (weakly) lower expected franchise values than the true policy function.

To operationalize this idea, we construct 60 such alternative policies. 20 perturb the responsiveness of deposit spreads to the short rate, 20 perturb their responsiveness to lagged market share, and 20 are multiplicative deviations (i.e., multiplying the spread implied by the

and control for them in the policy function, this is a reasonable assumption. Coupled with our assumption that private cost shocks χ are mean-zero and i.i.d. over banks, markets, and time, this obviates the need to estimate the cost shock distribution. We can instead proceed by simply treating the empirical distribution of policy function residuals as an estimate of the equilibrium distribution of policy shocks (i.e., changes to banks' chosen spreads driven by cost shocks) and sample from this distribution in our forward simulation.

[26] Series SVENF published at https://www.federalreserve.gov/data/yield-curve-tables/feds200628_1.html. Accessed 11/19/2024.

[27] This procedure admits a type of time-inconsistency: paths of simulated rates are invariant to idiosyncratic shocks drawn in a given period. We have experimented with alternative processes for interest rates, such as the Cox et al. (1985) process, and obtained very similar results.

estimated policy function by a constant). All perturbations are normal draws centered around the estimated policy function. For each alternative, we perturb the policy function of a single randomly-selected bank in each county-year, holding the policy function of its competitors fixed at the original estimated policy. Each policy alternative thus generates one inequality per county-year, namely that the selected bank is better off at the estimated policy function than at the perturbed one. Following BBL, we estimate marginal costs by minimizing the expected sum of squared violations over all inequalities. We compute expectations over 1000 paths of: (i) interest rates drawn from instantaneous forward rates and (ii) private shocks drawn from the empirical distribution of policy function residuals. We set the discount rate $\beta = 0.9$ and forward-simulate out to 50 periods in the future. Further details are given in Appendix Section A.

To see the intuition for the procedure, consider an alternative policy function that marginally increases spreads. This policy function and our estimated demand system tell us how deposit spreads and quantities jointly evolve in the forward simulation, allowing us to compute the marginal revenue associated with the alternative. Discounting and averaging over simulations thus gives us expected discounted marginal revenue. If this present value of marginal revenue is positive, it must be that expected discounted marginal costs are high. Otherwise, the observed policy function would not be optimal.

Such deviations differentially impact each of the four cost terms at different points in the state space, allowing us to separately identify each parameter. For instance, if an alternative policy raises marginal revenue particularly in high-rate periods within our simulations, it must be that costs load heavily on rates (i.e., ζ is high). Otherwise, the observed policy function would not be optimal. Similarly, if an alternative policy raises expected discounted marginal revenue particularly for banks with high values of the cost shifters $\mathbf{Z}_{jt}$, then $\boldsymbol{\gamma}$ must be high.

5.3.5 Results

Our marginal cost estimates are presented in Table 6, and the time series of implied one-period markups, $\rho_{jt} - c_{jt}$, is presented in Panel (a) of Figure 4. The average markup in our sample is about 68 bps. Moreover, markups exhibit significant cyclicality: average markups exceeded 125 bps in 2005-2007 and 2022-2023, while they hovered between 20 and 45 bps from 2008 to 2015.

These extended periods of low markups largely reflect the fact that interest rates were surprisingly low for much of the sample period, systematically undershooting forward curve-implied expectations. In other words, banks did not expect low markups to persist as long as they did. To illustrate this point, we next study a measure of "expected markups," defined as the constant markup that a bank would have to charge, holding fixed its expectations about short rates and deposit volumes, to obtain the franchise value that it does with the time-varying markups it actually charges. The red line in Panel (b) of Figure 4 plots the time series of

expected markups. The average expected markup in our sample is also about 68 bps, but it varies substantially less throughout the sample. In other words, at times like 2009 when current markups were low, banks still expected to earn significantly higher markups going forward.

5.4 Extensions

We explore two extensions of the model. First, we allow banks to price-discriminate between new and existing depositors by offering short-term bonuses to new customers. Offering bonuses to new depositors enables banks to partially, though not perfectly, price-discriminate between active and inactive depositors. Second, we allow active depositors to make forward-looking, rather than myopic, choices about where to deposit their funds. These extensions enhance the model but do not quantitatively alter our conclusions about how dynamic competition and the corresponding invest-versus-harvest tradeoff affect competition in the banking sector.

5.4.1 Promotional Rates for New Depositors

In our baseline model, banks pay the same deposit rates to incumbent depositors (i.e., those who were with the bank in the prior period) and to new depositors. In Appendix E, we develop and estimate an extension in which banks can offer different rates to new versus existing depositors (i.e., promotional rates). We also allow the marginal costs of serving incumbents and new customers to differ to reflect customer acquisition costs.

Absent frictions, banks would offer aggressive rates to active depositors and never offer positive rates to inactive ones, thereby investing in active customers while maximally harvesting inactive ones. Such price discrimination is infeasible because banks cannot observe which of their incumbent depositors are active. In practice, and in our model extension, banks instead offer separate rates to new and existing depositors, which allows them to imperfectly price-discriminate across active and inactive depositors. New depositors, by definition, are all active and thus have more elastic demand than a bank's existing depositor base, which is instead made up of a mix of active and inactive customers. Consistent with this, we observe in the data that banks offer aggressive promotional rates to new customers.

It is straightforward to extend and estimate our model to incorporate promotional rates for new depositors. First, we estimate separate policy functions for incumbent and new spreads. For simplicity and to maximize comparability with our baseline results, we assume that our original policy function estimates capture the choice of incumbent spreads. We proxy for spreads offered to new depositors as the baseline spread plus the gap between the "promotional" rates offered on certificates of deposit (CD) from the Ratewatch data and non-promotional CD rates of similar maturity.[28] We report these estimates in Appendix Table A1.

[28]The majority of our sample does not have information on promotional rates, so we estimate the "new spread" policy function in the subset of observations that do report promotional rates.

To estimate marginal costs, we then run an augmented forward simulation in which we now simulate forward both policy functions, as well as policy shocks with respect to both spreads. We construct perturbed polices as in the baseline case, though for each type of deviation, we perturb "incumbent" spreads in half of the runs and "new" spreads in the other half.

We report our estimates in Table A2, and we plot the time-varying distribution of estimated markups charged to incumbent and new depositors in Figure A4. Our estimates indicate that the marginal cost of serving new depositors is high relative to that of incumbents, suggesting that banks face high customer acquisition costs. In particular, we find that banks spend about \$0.15 upfront to raise an additional dollar of new deposits (an implied markup of -1500 bps), while they earn sizable markups on their existing depositor base (134 bps, on average). Although this extension provides new insights into bank pricing behavior, acquisition costs, and the role of inertia, it does not significantly change our conclusions about the importance of accounting for depositor sleepiness.

5.4.2 Forward-Looking Active Depositors

We also consider an extension that allows active depositors to make forward-looking choices. In our baseline model, depositors are myopic: they do not take into account that they are likely to be asleep, and therefore unable to re-optimize, for several periods following an active decision. In this extension, we re-estimate our deposit demand system under the assumption that depositors are forward-looking. Forward-looking behavior complicates estimation in the sense that the depositor's problem becomes dynamic:

$$V_{it}(D_{it}=1) = \max_{j \in \mathcal{K}} \Big\{ \alpha \rho_{jkt} + \delta_{jkt} + \varepsilon_{ijkt} + \beta \mathbb{E}[V_{it+1}(D_{it+1}) \,|\, j] \Big\}$$

Decomposing the continuation value into the discounted present value of utility conditional on remaining asleep and the discounted present value conditional on making a new active choice, we obtain:

$$V_{it}(D_{it}=1) = \max_{j \in \mathcal{K}} \Big\{ \alpha \rho_{jkt} + \delta_{jkt} + \varepsilon_{ijkt} + \sum_{s=t+1}^{\infty} (\beta\phi)^{s-t} \mathbb{E}[\alpha \rho_{jks} + \delta_{jks}] + \sum_{s=t+1}^{\infty} (\beta\phi)^{s-t} \frac{1-\phi}{\phi} \mathbb{E}[V_{is}(1,j) \,|\, j] \Big\}$$

The key insight is that the second component of the continuation value *does not depend on the depositor's choice today* – conditional on waking up in any period, all depositors have the same expected utility over that period's choice set, regardless of the bank they used in the previous period. We can therefore difference this term out across all banks in the market without altering choice probabilities. As a result of this differencing, there are no longer any recursive terms in this expression. The only remaining difference between the forward-looking and myopic cases is that for the former, the discounted present value of future utility from staying asleep and remaining at the bank chosen today essentially behaves like an additional characteristic of that

bank's current deposit services.

To proceed with estimation, we construct depositor expectations of future utility using the forward-simulated spreads from our baseline estimation.[29] We present our estimates using this approach in Appendix Table A3. Our results are qualitatively similar to the baseline demand estimates, although we find that demand is more elastic when we assume active depositors are forward-looking rather than myopic.

6 Implications for Markups and Franchise Values

In this section, we use counterfactuals to illustrate the insights that emerge from modeling dynamic competition for deposits. We compare the estimates from our dynamic model to counterfactuals in which all depositors are always active. In these counterfactuals, banks no longer face invest-versus-harvest incentives and competition is static à la Bertrand-Nash.

We show that sleepiness and dynamic competition generate five main effects. First, they significantly raise average markups. Second, markups become procyclical, rising when the short rate rises. Third, sleepiness and dynamic competition particularly increase markups in areas that a static view of competition would suggest are competitive, i.e., in low-concentration areas with low Herfindahl-Hirschman indices (HHIs). Fourth, sleepiness increases banks' franchise values and lowers their default probabilities. Finally, the franchise values of banks with either low product-quality or high costs increase most.

Our comparison between dynamic and static competition not only generates interesting economic insights, it is also policy-relevant. In recent years, regulators in several countries have proposed and implemented regulations that seek to constrain the ways in which banks set deposit rates. Perhaps most prominently, the UK's Financial Conduct Authority (FCA) has engaged in a policy initiative to increase the pass-through of short-term interest rates into deposit rates. In Appendix C, we show that the limiting case of this initiative—a policy that requires banks to *fully* pass through changes in short rates—is isomorphic to eliminating depositor sleepiness.[30] The intuition is that, by forcing full pass-through and constraining banks to choose a single spread for both the current and all future periods, the regulator shuts down the invest-versus-harvest tradeoff otherwise induced by depositor sleepiness: a bank's choice of spread affects both its active market share today *and its active market share tomorrow*. In other words, the

[29] This comes with the caveat that these simulated spreads are based on the assumption that banks price off of depositors' myopic preferences. Strictly speaking, banks set rates as a function of depositor preferences and lagged market shares, such that future spreads are a complicated, nonlinear function of α, δ_j, and e_{jkt}. To construct depositor expectations consistent with banks pricing off of their forward-looking preferences, we would thus need to perform a global search over the space of demand-side parameters, where for each conjecture of parameter values we would have to forward-simulate over aggregate rates and policy shocks, construct the expectation over future spreads, and then run the above regression. This is computationally infeasible.

[30] This equivalence exactly holds in the absence of a zero lower bound (ZLB) on deposit rates, but we show in the appendix that the results are very similar quantitatively whether or not we impose a ZLB.

bank cannot invest today and harvest tomorrow; it either invests in all periods or harvests in all periods.

6.1 Markups in Dynamic Equilibrium

We begin by comparing the behavior of markups in our dynamic model versus the static counterfactual.

6.1.1 Level

Panel (b) of Figure 4 studies the average level of markups. It compares expected markups in our dynamic model to those in the static counterfactual. While average expected markups in the dynamic model are about 68 bps, in the counterfactual they are about 32 bps. Thus, eliminating banks' invest-versus-harvest incentives reduces equilibrium expected markups by an average of 36 bps, about 53%. Intuitively, in the dynamic model, banks face a less elastic set of depositors than they do in the static counterfactual.

Panel (b) of Figure 4 also illustrates the distribution of markups across banks in the dynamic model and static counterfactual. The dark grey band displays the 25th-75th percentiles, and the light grey band displays the 10th-90th sample percentiles.[31] For most years, the 10th percentile markup in the dynamic model is above the 90th percentile markup in the static model.

6.1.2 Cyclicality

We next turn to the cyclicality of markups. The green dots in Panel (a) of Figure 5 display binned sample averages of estimated markups as a function of the federal funds rate. As discussed in the previous section, markups display substantial cyclicality, averaging about 30 bps when the short rate is 0% and 125 bps when the short rate is 5%.

At first, this cyclicality may seem puzzling: we do not include an outside option in our demand system, and we model depositor elasticities as being constant over time. Moreover, to the extent that depositor sleepiness decreases when the short rate increases, as we found in Table 3, one would expect markups to be *counter*-cyclical. As equation (6) shows, harvest incentives and thus markups are lower when sleepiness is lower.

To see the intuition, we first consider what a static model delivers, depicted in the blue dots. In our demand system, a static model can only deliver cyclical markups as a mechanical byproduct of the zero lower bound (ZLB) on deposit rates. That is, when deposit rates are stuck at zero, increases in the short rate raise deposit spreads one-for-one (i.e., deposit rates remain at zero). If the ZLB constraint does not bind, then banks charge the optimal static markup, which is constant. The blue dots in Panel (a) display this property. When the short rate is at zero,

[31] Percentiles are calculated across county-bank observations each year.

banks are bound by the ZLB and therefore charge markups below their optimal unconstrained levels; at higher levels of the short rate, banks simply charge this constant markup.

Under dynamic competition, this logic breaks down because of the invest-versus-harvest tradeoff. When the short rate is high, the harvest incentive leads banks that have accumulated high market shares to keep their deposit rates at or near zero. This softens the competition facing banks with little accumulated market share that want to invest. These banks can undercut the harvesting banks while still charging relatively high markups.

One may worry that the procyclical markups we estimate in panel (a) are not a property of the dynamic model; they may simply reflect the fact that our estimated policy function is strongly procyclical. To address this concern, we directly solve for the Markov Perfect equilibrium of our model at our estimated parameter values in a hypothetical market with three symmetric banks. The red dots plot the equilibrium policy function across different short rates.[32] As is evident in the figure, simulated and estimated markups display similar degrees of cyclicality, suggesting that dynamic competition does indeed drive our results.

6.1.3 Market Concentration

Panel (b) of Figure 5 illustrates a second key property of our dynamic model: markups are essentially constant as a function of market concentration, as measured by the Herfindahl-Hirschman index (HHI). As before, this result obtains both in our estimated model and in our simulated dynamic equilibrium. In contrast, the static benchmark delivers the expected result that markups are substantially higher in more concentrated markets. In the data, deposit spreads are nearly invariant to market concentration (see Appendix Figure A5), corroborating our findings.[33]

How does dynamic competition alter the relationship between concentration and competition? Consider a market in which one bank holds the majority of market share, making concentration high in this market. This bank will choose to harvest and offer uncompetitive rates, but the other banks in the market will optimally respond by competing aggressively to win over the large bank's incumbent depositors.

6.2 Effects on Franchise Values

We next analyze the effect of the static counterfactual on the value of the deposit franchise. In Panel (a) of Figure 6, we plot the time series of aggregate deposit franchise value under our baseline estimates (orange line) and under the counterfactual policy (teal line). The static

[32] The dots correspond to the markups that a bank with 10% lagged market share sets, when one competitor has 10% lagged share and the other 80%.

[33] It could still be the case that more concentrated markets feature banks with lower marginal costs, in which case markups would be increasing in market concentration even when deposit spreads do not. We note that the flat relationship between deposit spreads and market concentration obtains even after controlling for observable measures of bank operating costs, cutting against this possibility.

counterfactual reduces the aggregate value of the deposit franchise by about 54%. In dollar terms, about $1 trillion of franchise value would have been lost if the FCA policy had been implemented in 2023, compared to an aggregate market value of equity in the banking sector of about $2.3 trillion.

Perhaps even more striking than the high average dependence of the deposit franchise on sleepy deposits is the degree of heterogeneity. As Panel (b) demonstrates, about 15% of banks derive more than 80% of franchise value from sleepiness whereas about 3% of banks derive less than 20% of franchise value from sleepy demand. The next section analyzes the drivers of this heterogeneity in greater detail.

6.2.1 Bank Heterogeneity

In panel (a) of Figure 7, we show bin-scatter plots of each bank's share of franchise value lost under the counterfactual policy against deposit product quality δ_j. We find that banks with the lowest product quality are the most negatively affected by the policy. The intuition underlying this result is that sleepiness dampens the returns to having higher product quality. This is because the vast majority of depositors in a given period are asleep and thus not choosing their banks as a function of product quality. In equilibrium, this allows smaller banks to charge relatively higher markups without losing too much market share.

In panel (b) of Figure 7, we show bin-scatter plots of each bank's share of franchise value lost under the counterfactual policy against its marginal cost, netting out the common cost component from the short rate, ζR_t. We find that the least cost-efficient banks are much more negatively affected in the counterfactual. Depositor sleepiness enables firms with high operating costs to pass their expenses through to depositors.

6.2.2 Effects on Financial Stability

A different way to understand the effects of the static counterfactual on franchise value is to examine risk-neutral probabilities of default. Consider a bank's decision to default or continue operating in the future, in the style of Leland (1994). The bank will default if and only if the loss incurred today is larger than the present value of future operating profits. In the language of our model, recalling that lending returns follow $R_{jt} \sim \mathcal{N}(\mu_{jt}, \sigma_{jt}^2)$:

$$\text{Default if: } -(R_{jt} - \mu_{jt} - r_{jt} - c_{jt})Dep_{jt} > \beta V_{jt+1}$$

Integrating out R_{jt} and rearranging, we obtain:

$$\text{Probability of Default} = \Phi\left(\frac{\frac{-\beta V_{jt+1}}{Dep_{jt}} + r_{jt} + c_{jt} - \mu_{jt}}{\sigma_{jt}}\right) \tag{12}$$

All objects in equation (12) are already known or have been estimated, other than σ_{jt}.[34] To recover σ_{jt}, we first convert observed credit default swap spreads into implied risk-neutral probabilities of default and then invert equation (12) for each bank in each year of our sample period.[35]

To calculate counterfactual probabilities of default, we replace the baseline objects V_{jt+1} and r_{jt} with their counterfactual counterparts, and then simply re-evaluate the right-hand-side.

We plot the time series of baseline and counterfactual default probabilities for the three largest banks (JP Morgan, Bank of America, Wells Fargo) in Figure 8.[36] Panel (a) displays the results for JP Morgan. Baseline default probabilities range from near 0% to about 2.5%, whereas counterfactual default probabilities fluctuate around 10%. Excluding the global financial crisis, fluctuations in counterfactual default probabilities are driven primarily by the shape of the yield curve: when future rates are expected to be high, more franchise value is lost in the counterfactual.

7 Conclusion

Bank deposits are often described as sticky, sleepy, and inelastic. We examine the implications of sleepy deposits and their role in creating value and stability in the banking sector. We demonstrate that whether the observed stickiness is driven by depositor sleepiness or inelastic demand (i.e., preference heterogeneity) has significant implications for consumers, banks, and policymakers. Using both our microdata and branch-level deposit data, we show that depositors exhibit high levels of inactivity, which we attribute to depositor sleepiness. This sleepiness appears to vary both across depositors and over time.

The sleepiness we document has critical consequences. For consumers, it implies that depositors are much more responsive to deposit rates than previously believed once sleepiness is accounted for. For banks, depositor inactivity creates substantial value: on average, nearly 60% of a bank's deposit franchise value arises from the fact that most of its depositors remain inactive in any given period.

Depositor sleepiness also raises concerns for policymakers. In particular, it suggests that banks may be slow to pass changes in aggregate rates through to depositors. For example, a bank might delay increasing its deposit rates following a rise in aggregate rates, knowing that most of its depositors are unlikely to notice or respond. These concerns have prompted regulators to consider policies aimed at increasing rate pass-through to depositors. We analyze the effects of such policies, including requiring banks to offer floating-rate deposits with a fixed spread. While these policies could benefit depositors by raising deposit rates by approximately

[34] Recall that the expected lending return, μ_{jt}, is not separately identified from marginal costs c_{jt}. As a result, our empirical estimates of c_{jt} correspond to $c_{jt} - \mu_{jt}$ in this equation.

[35] We use five-year CDS spreads, assuming a recovery rate of 40% and a constant hazard rate.

[36] Our results use the primary commercial bank subsidiary within each holding company's organization.

36 basis points, they also come with significant costs. For instance, following the 2022-2023 monetary tightening, the probability of default for two major money-center banks in the US would have increased to more than 20% had banks been forced to offer floating-rate deposits. Thus, while these policies may enhance consumer outcomes, they also pose risks to financial stability.

References

Ackerberg, D. A. and G. Gowrisankaran (2006). Quantifying Equilibrium Network Externalities in the ACH Banking Industry. *The RAND Journal of Economics 37*(3), 738–761. Publisher: [RAND Corporation, Wiley].

Agarwal, S., R. J. Rosen, and V. Yao (2016). Why do borrowers make mortgage refinancing mistakes? *Management Science 62*(12), 3494–3509. Publisher: INFORMS.

Aguirregabiria, V., R. Clark, and H. Wang (2024). The geographic flow of bank funding and access to credit: Branch networks, local synergies, and competition. *American Economic Review (forthcoming)*.

Allen, J., R. Clark, J.-F. Houde, S. Li, and A. Trubnikova (2024). The role of intermediaries in selection markets: Evidence from mortgage lending. *The Review of Financial Studies*, hhae075. Publisher: Oxford University Press.

Andersen, S., J. Y. Campbell, K. M. Nielsen, and T. Ramadorai (2020). Sources of inaction in household finance: Evidence from the danish mortgage market. *American Economic Review 110*(10), 3184–3230. Publisher: American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Argyle, B., B. C. Iverson, J. D. Kotter, T. Nadauld, and C. Palmer (2025). Sticky Deposits, not Depositors.

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating Dynamic Models of Imperfect Competition. *Econometrica 75*(5), 1331–1370. Publisher: Wiley Online Library.

Bancography (2014, January). Bancology: Market analysis and insights for financial institutions. *Bancography*.

Beggs, A. and P. Klemperer (1992). Multi-period competition with switching costs. *Econometrica: Journal of the Econometric Society*, 651–666. Publisher: JSTOR.

Benetton, M. (2021). Leverage regulation and market structure: A structural model of the UK mortgage market. *The Journal of Finance 76*(6), 2997–3053. Publisher: Wiley Online Library.

Benetton, M., A. Gavazza, and P. Surico (2021). Mortgage pricing and monetary policy. *Working Paper*.

Berger, A. N. and T. H. Hannan (1991). The Rigidity of Prices: Evidence from the Banking Industry. *The American Economic Review 81*(4), 938–945. Publisher: American Economic Association.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 841–890. Publisher: JSTOR.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 242–262. Publisher: JSTOR.

Berry, S. T. and G. Compiani (2023, July). An Instrumental Variable Approach to Dynamic Models. *The Review of Economic Studies 90*(4), 1724–1758.

Blickle, K., J. Li, X. Lu, and Y. Ma (2025, August). The Dynamics of Deposit Flightiness and its Impact on Financial Stability. Working Paper 34128, National Bureau of Economic Research.

Brown, Z. Y., M. Egan, J. Jeon, C. Jin, and A. A. Wu (2024, May). Why Do Index Funds Have Market Power? Quantifying Frictions in the Index Fund Market. *NBER Working Paper 31778*.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru (2018). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics 130*(3), 453–483. Publisher: Elsevier.

Cirelli, F. and A. Olafsson (2025). What Makes Depositors Tick? Bank Data Insights into Households' Liquid Asset Allocation. CEPR Discussion Paper 20612, CEPR Press, Paris & London.

Corbae, D. and P. D'Erasmo (2020). Rising Bank Concentration. *Journal of Economic Dynamics and Control 115*, 103877.

Corbae, D. and P. D'Erasmo (2024). A Quantitative Model of Banking Industry Dynamics.

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

Cox, J. C., J. Jonathan E. Ingersoll, and S. A. Ross (1985). A Theory of the Term Structure of Interest Rates. *Econometrica 53*(2), 385–407. Publisher: The Econometric Society.

DeMarzo, P., A. Krishnamurthy, and S. Nagel (2024). Interest Rate Risk in Banking. *Working Paper*.

Dick, A. A. (2008). Demand estimation and consumer welfare in the banking industry. *Journal of Banking & Finance 32*(8), 1661–1676. Publisher: Elsevier.

Drechsler, I., A. Savov, and P. Schnabl (2017). The Deposits Channel of Monetary Policy. *The Quarterly Journal of Economics 132*(4), 1819–1876. Publisher: Oxford University Press.

Drechsler, I., A. Savov, and P. Schnabl (2021). Banking on Deposits: Maturity Transformation without Interest Rate Risk. *The Journal of Finance 76*(3), 1091–1143. Publisher: Wiley.

Drechsler, I., A. Savov, P. Schnabl, and O. Wang (2023). Deposit Franchise Runs. Working Paper 31138, National Bureau of Economic Research.

Dubé, J.-P., G. J. Hitsch, and P. E. Rossi (2009). Do Switching Costs Make Markets Less Competitive? *Journal of Marketing Research 46*(4), 435–445.

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

Egan, M., S. Lewellen, and A. Sunderam (2022). The cross-section of bank value. *The Review of Financial Studies 35*(5), 2101–2143. Publisher: Oxford University Press.

Einav, L., B. Klopack, and N. Mahoney (2023). Selling Subscriptions. *Working Paper*.

Farrell, J. and C. Shapiro (1988). Dynamic Competition with Switching Costs. *The RAND Journal of Economics 19*(1), 123–137. Publisher: Wiley.

Fleitas, S. (2017, May). Dynamic Competition and Price Regulation When Consumers Have Inertia: Evidence from Medicare Part D. Working Paper, Northwestern University.

Fowlie, M., M. Reguant, and S. P. Ryan (2016, February). Market-Based Emissions Regulation and Industry Dynamics. *Journal of Political Economy 124*(1), 249–302. Publisher: The University of Chicago Press.

Gowrisankaran, G. and J. Krainer (2011). Entry and pricing in a differentiated products industry: evidence from the ATM market. *The RAND Journal of Economics 42*(1), 1–22. Publisher: [RAND Corporation, Wiley].

Gowrisankaran, G. and M. Rysman (2012). Dynamics of Consumer Demand for New Durable Goods. *Journal of Political Economy 120*(6), 1173–1219. Publisher: University of Chicago Press.

Hanson, S. G., A. Shleifer, J. C. Stein, and R. W. Vishny (2015). Banks as patient fixed-income investors. *Journal of Financial Economics 117*(3), 449–469. Publisher: Elsevier.

Honka, E., A. Hortaçsu, and M. A. Vitorino (2017, August). Advertising, Consumer Awareness, and Choice: Evidence from the U.S. Banking Industry. *RAND Journal of Economics 48*(3), 611–646.

Jiang, E. X., G. Matvos, T. Piskorski, and A. Seru (2024). Monetary tightening and US bank fragility in 2023: Mark-to-market losses and uninsured depositor runs? *Journal of Financial Economics 159*, 103899. Publisher: Elsevier.

Leland, H. E. (1994). Corporate Debt Value, Bond Covenants, and Optimal Capital Structure. *The Journal of Finance 49*(4), 1213–1252. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1994.tb02452.x.

Lu, X., Y. Song, and Y. Zeng (2025, March). Tracing the Impact of Payment Convenience on Deposits: Evidence from Depositor Activeness. Research Paper, Jacobs Levy Equity Management Center for Quantitative Financial Research, The Wharton School.

Lu, X. and L. Wu (2025). Banking on Inattention. *mimeo*.

MacKay, A. and M. Remer (2024). Consumer inertia and market power. *Working Paper*.

Matvos, G. and A. Seru (2014, April). Resource Allocation within Firms and Financial Market Dislocation: Evidence from Diversified Conglomerates. *The Review of Financial Studies 27*(4), 1143–1189.

Melnikov, O. (2013). Demand for Differentiated Durable Products: The Case of the U.S. Computer Printer Market. *Economic Inquiry 51*(2), 1277–1298. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1465-7295.2012.00501.x.

Miller, K. (2019, June). Estimating Costs When Consumers Have Inertia: Are Private Medicare Insurers More Efficient? Working Paper, University of Oregon.

Modigliani, F. and M. H. Miller (1958). The Cost of Capital, Corporation Finance and the Theory of Investment. *The American Economic Review 48*(3), 261–297. Publisher: American Economic Association.

Neumark, D. and S. A. Sharpe (1992). Market Structure and the Nature of Price Rigidity: Evidence from the Market for Consumer Deposits. *The Quarterly Journal of Economics 107*(2), 657–680. Publisher: Oxford University Press.

Padilla, A. J. (1995, December). Revisiting Dynamic Duopoly with Consumer Switching Costs. *Journal of Economic Theory 67*(2), 520–530.

Pakes, A. and P. McGuire (1994). Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model. *The RAND Journal of Economics 25*(4), 555–589. Publisher: [RAND Corporation, Wiley].

Pakes, A., J. R. Porter, M. Shepard, and S. Calder-Wang (2021, July). Unobserved Heterogeneity, State Dependence, and Health Plan Choices.

Ryan, S. P. (2012). The Costs of Environmental Regulation in a Concentrated Industry. *Econometrica 80*(3), 1019–1061. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA6750.

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

Whited, T. M., Y. Wu, and K. Xiao (2021). Low interest rates and risk incentives for banks with market power. *Journal of Monetary Economics 121*, 155–174.

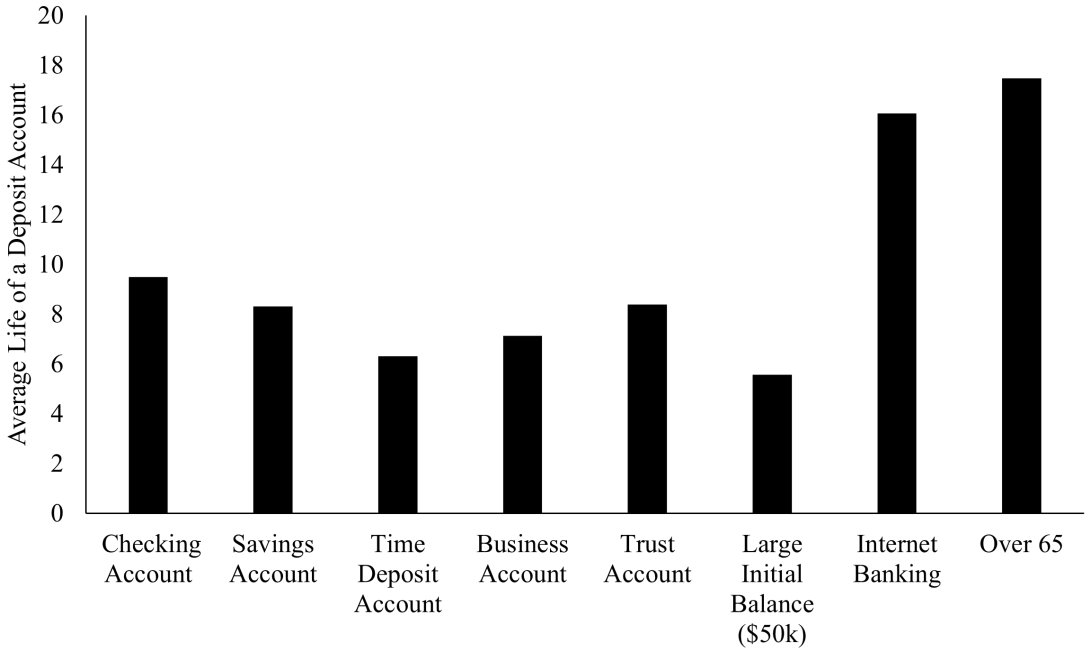
Whited, T. M., Y. Wu, and K. Xiao (2022). Will central bank digital currency disintermediate banks? *Working Paper*.

Xiao, K. (2020). Monetary transmission through shadow banks. *The Review of Financial Studies 33*(6), 2379–2420. Publisher: Oxford University Press.

Zhang, D. (2022). Closing costs, refinancing, and inefficiencies in the mortgage market. *Working paper*.

Tables and Figures

Figure 1: Average Life of a Deposit Account

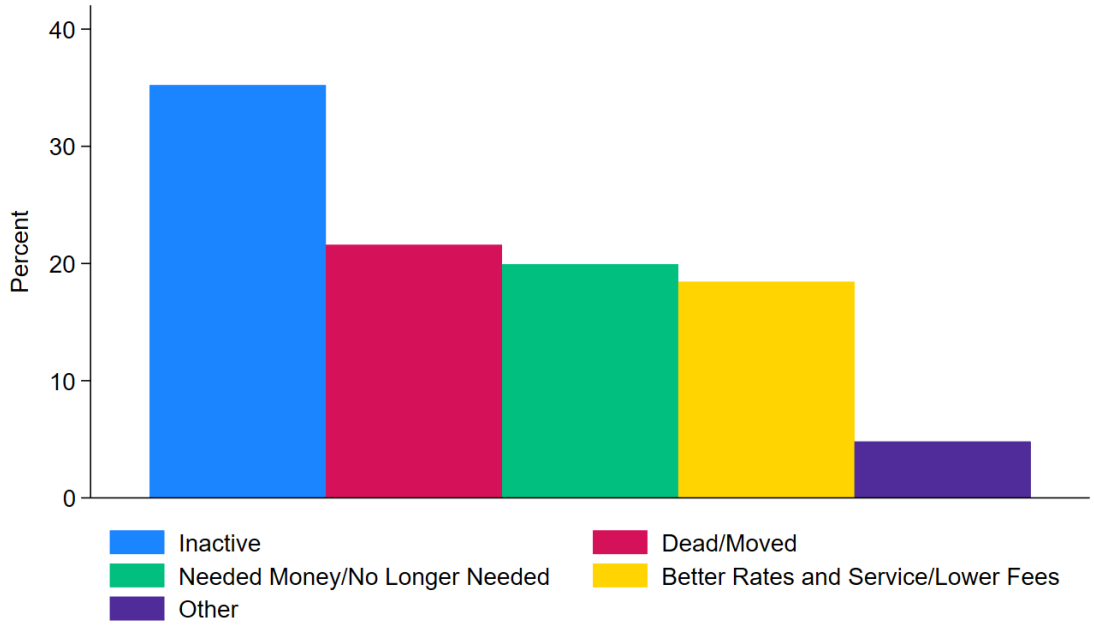


Notes: Figure 1 plots the implied average life of several different types of deposit accounts, in years, using our main account-level microdata. This dataset covers about 12 million accounts from 89 banks and credit unions. Average life is calculated using the rate of account turnover (i.e., the fraction of existing accounts in a given year that are newly-opened) and the following formula:

$$\text{Average life}_i = \frac{1}{\text{Turnover Rate}_i}$$

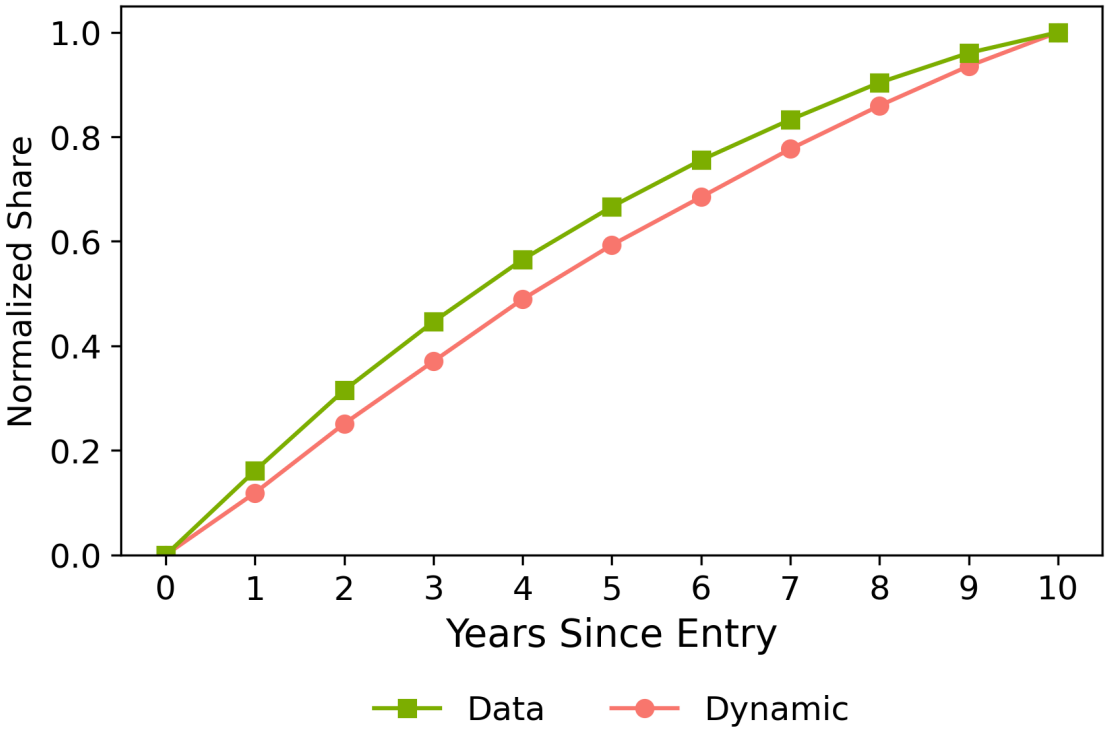
where i denotes account type. The turnover rate, and by extension average life, is calculated in units of years.

Figure 2: Reasons for Closing a Deposit Account



Notes: Figure 2 displays a histogram of the reasons depositors report for closing their accounts. This data is derived from our main account-level microdata, which covers about 12 million accounts from 89 banks and credit unions. We focus on closures where the depositor provided a reason, excluding those resulting from internal transfers within the bank (i.e., cases where the depositor closed an account to open another within the same bank). In 25% of closures, no reason is provided, and internal transfers account for 18% of all closures.

Figure 3: Sleepiness versus Unobserved Preference Heterogeneity: Evidence from Market Entry



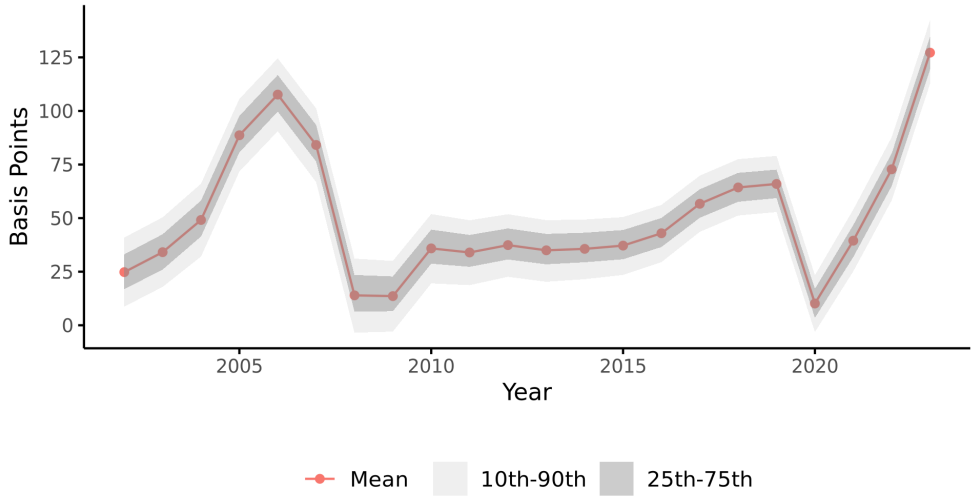
Notes: Figure 3 plots the progression of market shares for entering firms as a function of the number of years since entry:

$$\text{normalized share}_t = \frac{\text{share}_t - \text{share}_0}{\text{share}_{10} - \text{share}_0}$$

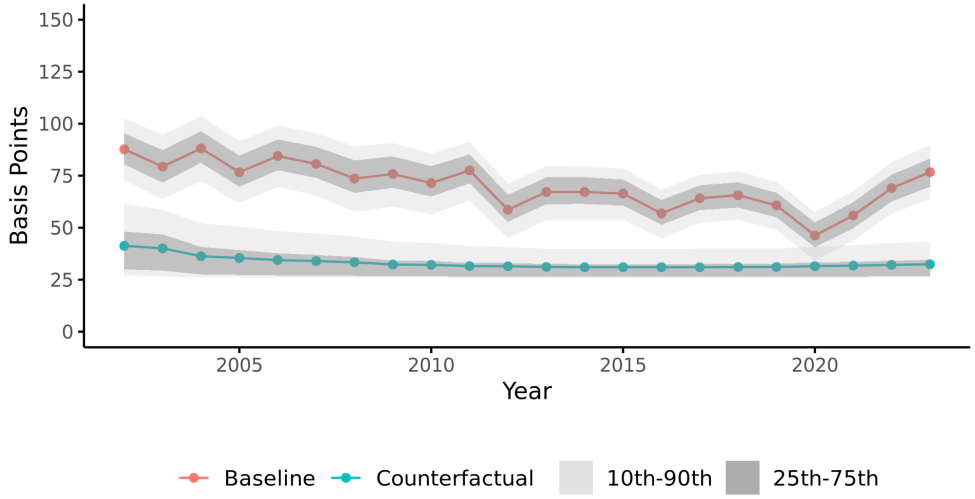
We difference out market shares in year 0 to account for the fact that banks entering new counties according to the Summary of Deposits immediately report large market shares for largely mechanical reasons (i.e., deposits previously collected in nearby areas without any of the entrant bank's branches are immediately booked at the new branch). The green line plots medians from Summary of Deposits data and the red line plots medians from forward-simulations of a 3-bank Markov Perfect equilibrium. See Appendix Figure A1 for an alternative normalization of market shares including the year-0 shares; see Appendix B for additional details on our computation of Markov-Perfect equilibria.

Figure 4: Markup Estimates

(a) Current



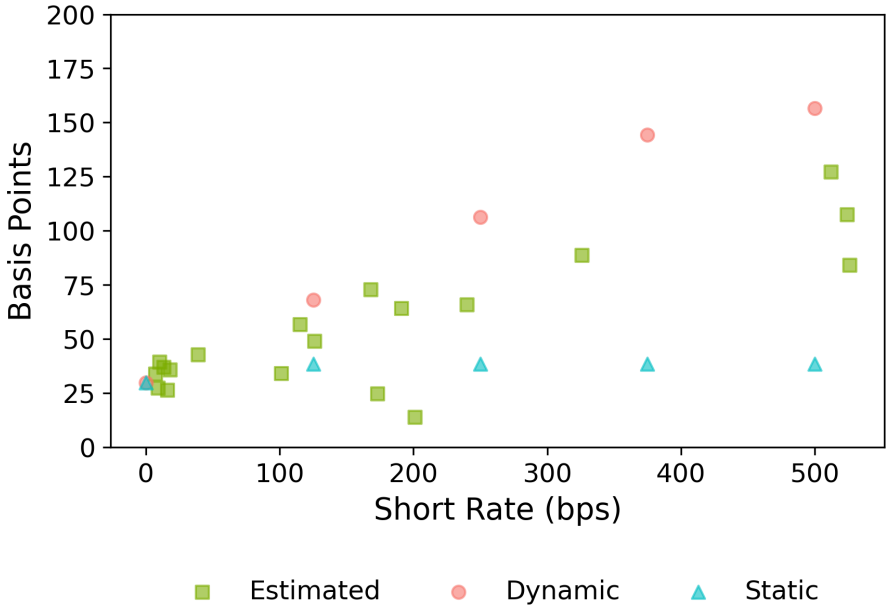
(b) Expected



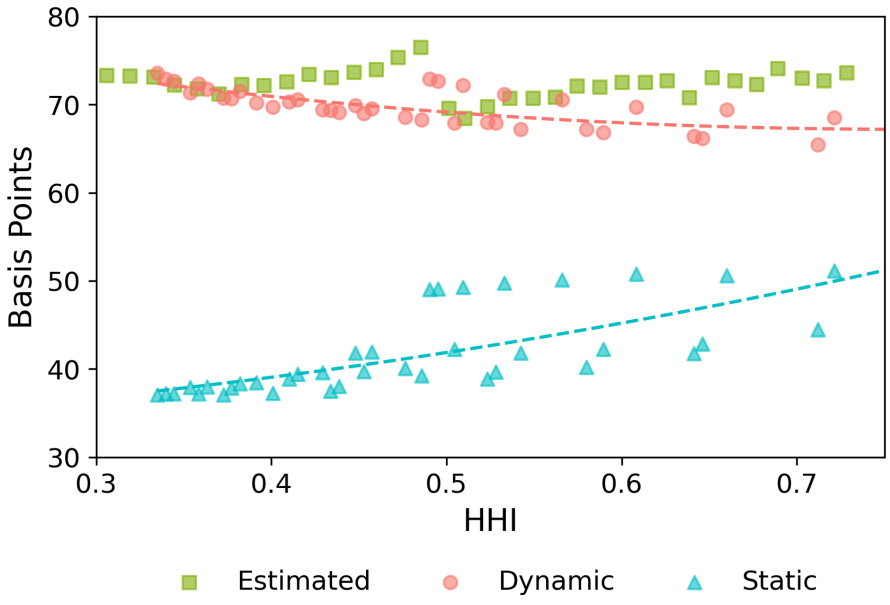
Notes: Panel (a) of Figure 4 displays the yearly bank-county level distributions of current markups under our baseline estimates. Panel (b) of Figure 4 displays the yearly bank-county level distributions of expected markups under our baseline estimates, as well as markups under the counterfactual policy constraining banks to fix the spread they charge below the short rate. "Expected markups" are defined as the constant markup that a bank would need to charge, holding expectations of current and future deposit volumes fixed, in order to achieve the same franchise value that we estimate in the forward simulation. We winsorize expected markups at the 1% level to reduce the influence of extreme outliers.

Figure 5: Properties of Estimated Markups

(a) Markups and Short Rate



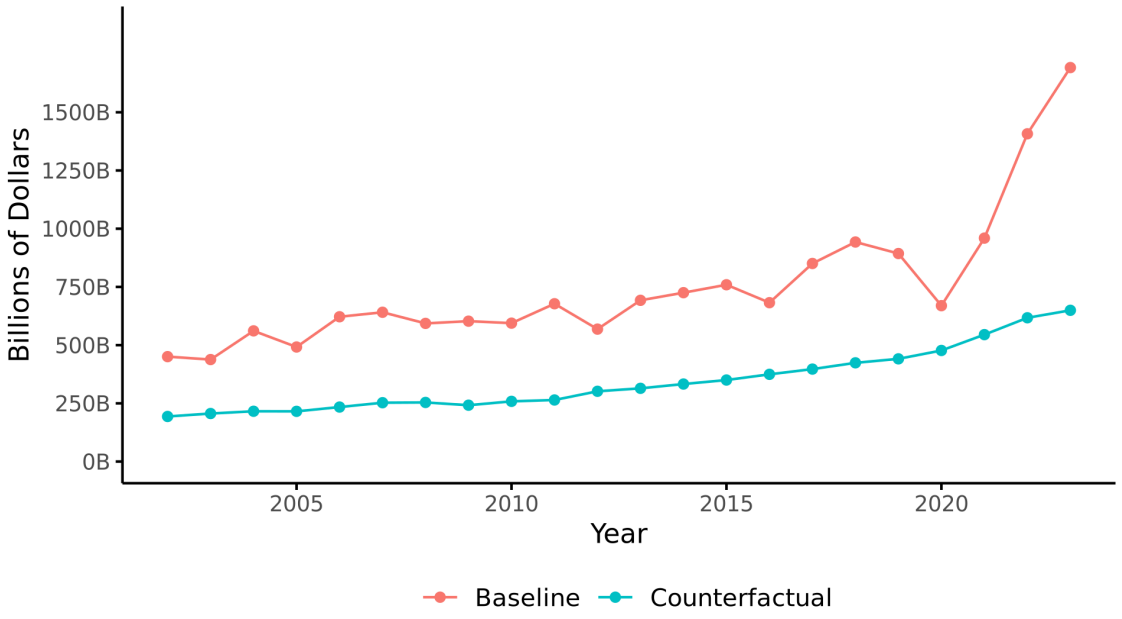
(b) Markups and Market Concentration



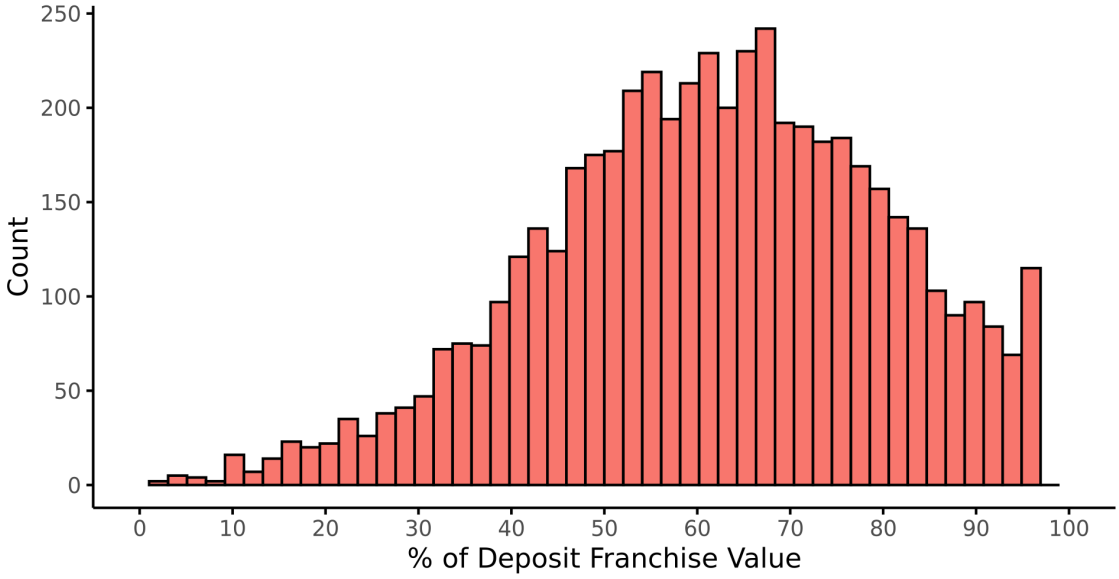
Notes: Panels (a) of Figure 5 plots markups from our two-step estimates (green line) against those obtained from a 3-bank Markov Perfect equilibrium (red line) and from a 3-bank Bertrand-Nash equilibrium without depositor sleepiness (teal line). Panel (b) of Figure 5 plots market-level markups from our two-step estimates (green line) against average market-level markups from a series of 3-bank Markov Perfect equilibria (red line) and a series of Bertrand-Nash equilibria without depositor sleepiness (teal line). See Appendix B for additional details on our computation of Markov Perfect and Bertrand-Nash equilibria.

Figure 6: Counterfactual Franchise Value

(a) Total Deposit Franchise Value in the Banking Sector



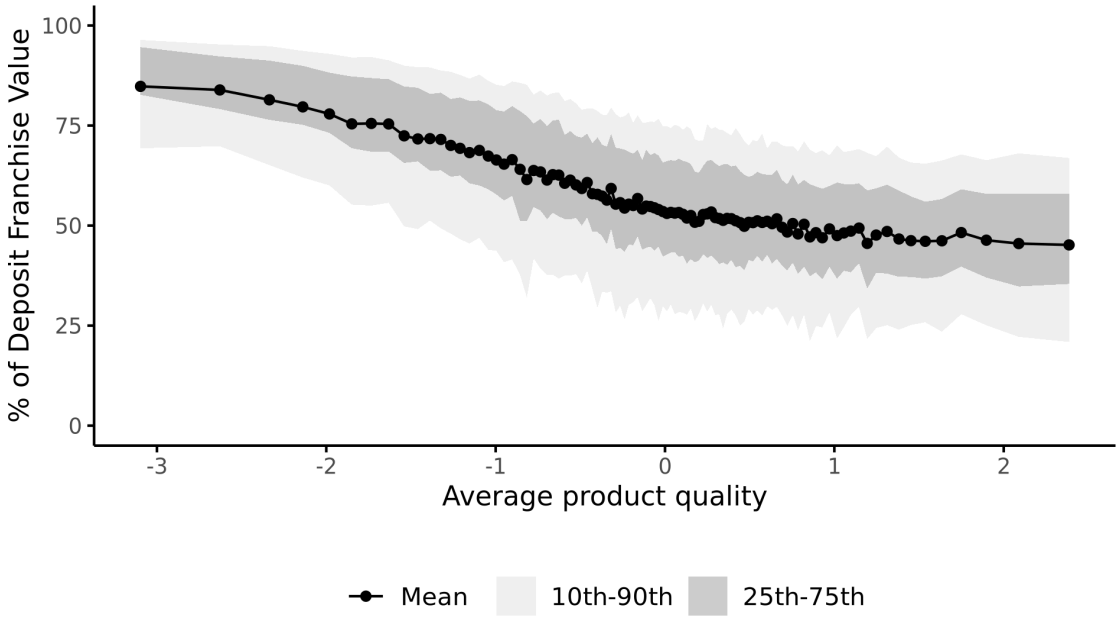
(b) Share of Deposit Franchise Value Lost in Counterfactual – 2010



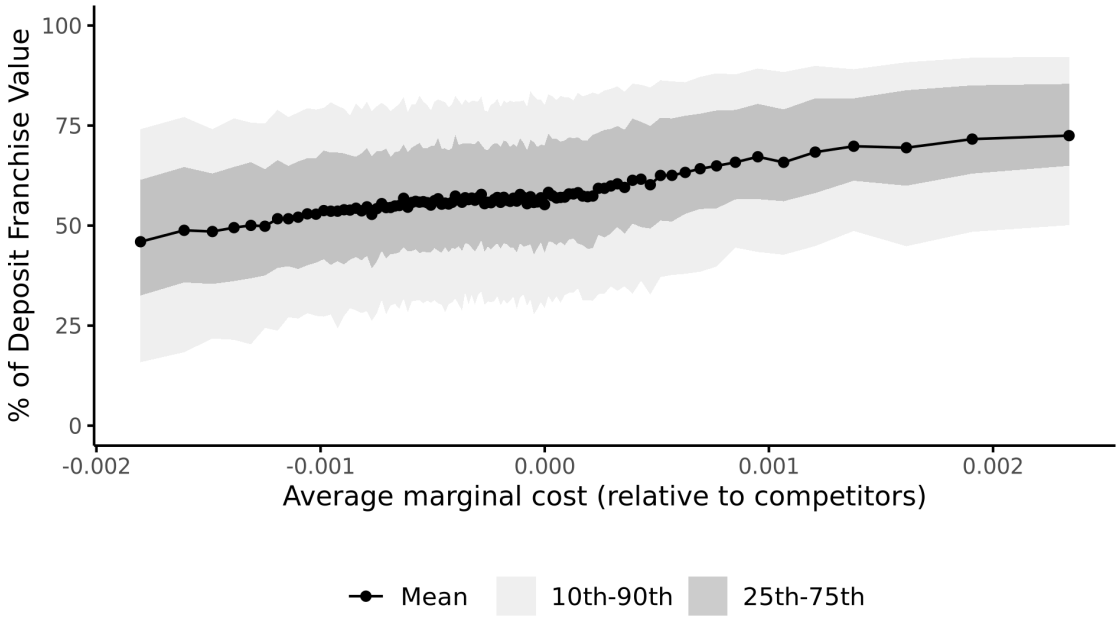
Notes: Panel (a) of Figure 6 displays total deposit franchise value in the U.S. banking sector under our baseline estimates and in our simulated static equilibria. To facilitate comparison with other aggregate quantities, we scale the yearly totals by the fraction of total deposits that are accounted for by observations in our sample. Panel (b) of Figure 6 displays the cross-sectional bank-level distribution of the fraction of deposit franchise value lost under the counterfactual policy (i.e., one less the ratio of counterfactual to baseline value) using data from 2010.

Figure 7: Bank-Level Heterogeneity in Policy Exposure

(a) Deposit Product Quality



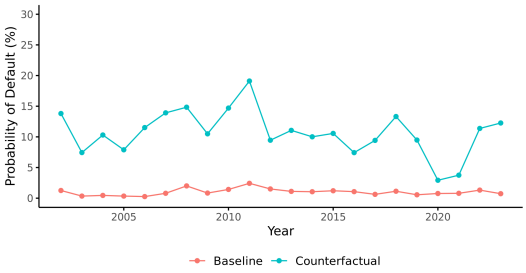
(b) Marginal Cost



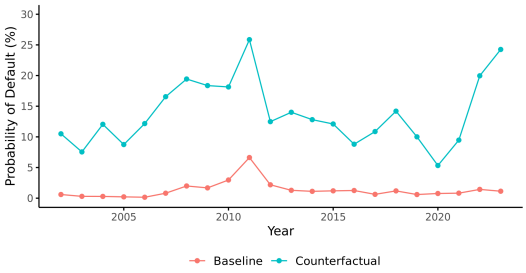
Notes: Panel (a) of Figure 7 displays a bank-level bin-scatter plot of the fraction of deposit franchise value lost under the counterfactual policy against product quality δ_j. Panel (b) of Figure 7 displays a bank-level bin-scatter plot of the fraction of the same statistic against marginal costs, demeaned by the average marginal cost in each of the bank's markets.

Figure 8: Counterfactual Probabilities of Default – Top 3 Banks

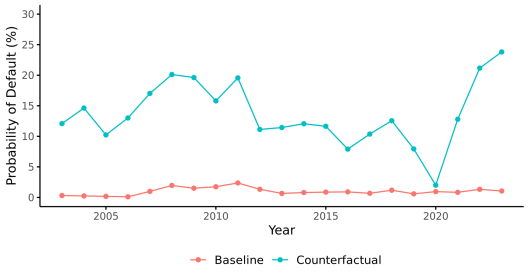
(a) JP Morgan



(b) Bank of America



(c) Wells Fargo



Notes: The panels of Figure 8 display baseline and counterfactual estimates of risk-neutral probabilities of default for the top 3 banks: JP Morgan, Bank of America, Wells Fargo. We compute baseline risk-neutral one-year default probabilities using five-year credit default swap spreads (assuming constant hazard rates). To obtain counterfactual default probabilities, we use the baseline five-year CDS spreads to estimate the risk-neutral probability of default. We then use our baseline estimated franchise value to calculate asset volatility. Finally, we reduce the bank's franchise value by the estimated difference in expected franchise value between the baseline and counterfactual scenarios and recalculate the probability of default while holding asset volatility constant. We assume a 40% recovery rate in the event of default.

Table 1: Summary Statistics

Variable	Observations	Mean	Median	SD
Bank-County Level:				
Deposit Rate (bps)	341,395	34	15	50
Deposit Spread (bps)	341,395	97	17	145
Deposit Market Share	341,395	0.16	0.10	0.18
County Deposits ($M)	341,395	405	84	4046
Bank Level:				
Wages ($1000s)	95,693	69	64	24
$100 \times \frac{\text{Fixed Expenses}}{\text{Total Assets}}$	95,693	0.39	0.36	0.18
Total Deposits ($M)	95,693	1874	171	30712

Notes: Table 1 summarizes the key variables in our data sample. As discussed in the text, we limit our sample to bank-county pairs appearing in at least 5 years of data and for which we have all required information in both the current and previous years. Throughout the entirety of the paper, including here, we winsorize our wage and fixed expense measures at the 2.5% level to reduce the influence of extreme outliers.

Table 2: Account Turnover

	Main sample (89 Banks)				Extended Sample (903 Banks)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Internet Banking	-0.007**	-0.028***	-0.043***	-0.052***				
	(0.003)	(0.003)	(0.003)	(0.003)				
Business Account	0.064***	0.042***	0.035***	0.041***	0.006	0.003	0.013***	0.013***
	(0.007)	(0.007)	(0.007)	(0.006)	(0.010)	(0.010)	(0.001)	(0.001)
Trust Account	0.055***	0.022***	0.014**	0.013**				
	(0.006)	(0.006)	(0.006)	(0.005)				
Over 65	-0.040***	-0.048***	-0.048***	-0.048***	-0.063***	-0.065***	-0.045***	-0.045***
	(0.003)	(0.003)	(0.002)	(0.002)	(0.014)	(0.012)	(0.001)	(0.001)
Large Initial Balance (50k+)	0.094***	0.086***	0.074***	0.072***				
	(0.006)	(0.006)	(0.005)	(0.005)				
ln(Bank Size (# Accounts))	0.006***	-0.004***	0.036***		0.026	0.022*	-0.022	
	(0.001)	(0.001)	(0.007)		(0.016)	(0.013)	(0.015)	
Savings Account	0.014***	0.010***	0.015***	0.013***	-0.034*	-0.031**	0.000	-0.000
	(0.003)	(0.003)	(0.003)	(0.002)	(0.020)	(0.016)	(0.001)	(0.001)
Time Deposit Account	0.056***	0.053***	0.053***	0.051***	0.059***	0.051***	0.073***	0.074***
	(0.004)	(0.004)	(0.003)	(0.003)	(0.009)	(0.011)	(0.005)	(0.004)
Observations	55,041	55,041	55,041	55,022	132,683	132,683	132,683	132,459
R-squared	0.133	0.281	0.431	0.616	0.184	0.226	0.879	0.958
Year FE		X	X			X	X	
Bank FE			X				X	
Bank-Year FE				X				X

Notes: Table 2 displays the regression results corresponding to equation (1). For columns (1)-(4), we classify bank accounts along five key attributes: account type (checking, savings, or time deposits), holder type (person, business, or trust), online banking setup, initial balance above $50,000, and whether the account-holder is over age 65 (for personal accounts). This results in 72 unique account profiles (3×3×2×2×2). Observations are at the account profile-by-bank-by-year level. The dependent variable is account turnover, which is defined as the number of new accounts with that profile divided by the number of existing accounts with that profile. For columns (5)-(8), we classify accounts by account type, holder type, and whether the account-holder is over age 65 (for personal accounts). This results in 12 unique account profiles (3x2x2). The regression is weighted by number of existing accounts, and robust standard errors are in parentheses. *** p<0.01, ** p<0.05, * p<0.10.

Table 3: Sleepiness Estimates

	Microdata only		SoD only		Combined	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Upsilon_{1,\text{cons}}$	0.868***	0.877***	0.977***	0.979***	0.904***	0.945***
	(0.008)	(0.010)	(0.001)	(0.002)	(0.007)	(0.006)
Υ_{1,R_t^F}		-0.416*		-0.105***		-0.355**
		(0.249)		(0.041)		(0.169)
$\Upsilon_{2,\text{over 65}}$		0.0526***				0.017***
		(0.005)				(0.005)
$\Upsilon_{2,\text{over \$50k}}$		-0.094***				-0.130***
		(0.006)				(0.007)
$\Upsilon_{2,\text{internet banking}}$		0.036***				0.002
		(0.007)				(0.007)
Observations	8,878	8,878	341,395	341,395	350,273	350,273
Within R-squared	0.966	0.966	0.963	0.963	0.962	0.966
Bank FE			X	X	X	X
County-Year FE			X	X	X	X

Notes: Table 3 reports the results of panel regressions of depositor sleepiness. Observations are either at the bank-county level (SoD) or the bank-depositor-cell level (using our main sample of microdata). In the case of the SoD sample, all branches of a given bank (as defined by FDIC certificate number) in a given county-year are aggregated into a single observation. In all columns including SoD data, we control for both deposit spreads and a 1st-order polynomial of the residuals obtained from equation (9): a spread-setting control function that uses salaries and fixed expenses as cost shifters. The F-statistic testing joint nullity of the cost shifters in the control function is 12.8. Standard errors are clustered by bank. *** p<0.01, ** p<0.05, * p<0.10.

Table 4: Demand Estimates

	Total Market (No Inertia)		Active Market	
	OLS	CF	OLS	CF
	(1)	(2)	(3)	(4)
ρ_{jkt}	-3.17***	-290.4***	-16.2***	-388.7***
	(1.14)	(15.3)	(2.05)	(29.3)
Observations	294,809	294,809	294,809	294,809
Rate Elasticity of Demand	0.004	0.392	0.022	0.525
R^2	0.733	0.734	0.625	0.626
Bank fixed effects	✓	✓	✓	✓
County-Year fixed effects	✓	✓	✓	✓

Notes: Table 4 reports OLS and control function estimates obtained from a Berry (1994) demand system, where j indexes banks, k indexes counties, and t indexes years. Observations are at the bank-county level. In particular, all branches of a given bank (as defined by FDIC certificate number) in a given county-year are aggregated into a single observation. In columns (3) and (4), we control for a 1st-order polynomial of the residuals obtained from (9): a spread-setting control function that uses salaries and fixed expenses as cost shifters. Standard errors are clustered by bank-year. The F-statistic testing joint nullity of the cost shifters in the control function is 12.8. The rate elasticity of demand is calculated assuming a market share of 10% and an offered interest rate at the sample median (about 15 bps). *** p<0.01, ** p<0.05, * p<0.10.

Table 5: Policy Function Estimates

	(1)
Constant	-0.001***
	(0.0003)
R_t	0.909***
	(0.005)
R_{t-1}	-0.116***
	(0.008)
δ_{jkt}	0.001***
	(3.26×10^{-5})
Log(Deposit Share)$_{jkt-1}$	-0.0010***
	(3.19×10^{-5})
$R_{t-1} \times$ Log(Deposit Share)$_{jkt-1}$	0.001
	(0.001)
$\sum_{\ell \neq j}$ Log(Deposit Share)$_{\ell kt-1}$	9.88×10^{-6}***
	(1.08×10^{-6})
$\frac{\text{Wages}_{jt}}{\text{\# Employees}_{jt-1}}$	3.12×10^{-5}***
	(1.96×10^{-6})
$\sum_{\ell \neq j} \frac{\text{Wages}_{\ell t}}{\text{\# Employees}_{\ell t-1}}$	1.37×10^{-6}***
	(9.02×10^{-8})
$\frac{\text{Fixed Expenses}_{jt}}{\text{Assets}_{jt-1}}$	-0.023
	(0.025)
$\sum_{\ell \neq j} \frac{\text{Fixed Expenses}_{\ell t}}{\text{Assets}_{\ell t-1}}$	-0.035***
	(0.001)
fraction over 65$_k$	0.002***
	(0.0002)
fraction over \$50k HH income$_k$	-0.002***
	(0.0002)
fraction w/internet$_k$	-0.004***
	(0.0003)
Observations	341,329
R^2	0.942

Notes: Table 5 reports OLS estimates of the spread-setting policy function, where j indexes banks, k indexes counties, and t indexes years. Observations are at the bank-county level. In particular, all branches of a given bank (as defined by FDIC certificate number) in a given county-year are aggregated into a single observation. Unobserved product qualities δ are taken from the demand system estimated in column (4) of Table 4. Standard errors are clustered by bank-year. *** p<0.01, ** p<0.05, * p<0.10.

Table 6: Marginal Cost Estimates

ζ	0.695
	(0.030)
ω	-0.008
	(0.003)
γ_1	5.139 $\times 10^{-6}$
	(1.190 $\times 10^{-6}$)
γ_2	0.013
	(0.325)

Notes: Table 6 reports the results of the two-step estimator based on Bajari et al. (2007). We construct 60 alternative policies: 20 perturb the responsiveness of deposit spreads with respect to the short rate; 20 perturb the responsiveness of deposit spreads with respect to lagged log market share; 20 multiply spreads by a constant randomly drawn around 1. For each alternative policy, independent draws are taken for one randomly-selected bank in each county-year – 3,278,400 total inequalities are used in estimation. Each expectation is constructed by averaging over 1,000 simulated paths of: (i) interest rates, generated from estimated yield curves as described in Section 5.3.3; and (ii) policy shocks, drawn from the distribution of policy function residuals (censored at the 5th and 95th percentiles). Standard errors are calculated via bootstrap: we resample, with replacement, 500 times – each time constructing datasets of the same length as the total number of inequalities. We then calculate standard errors as the standard deviation of the vector of estimates.

Costs are parameterized as:

$$c_{jkt} = \zeta \text{Fed Funds}_t + \omega + \gamma_1 \frac{\text{Salaries}_{jt}}{\text{\# Employees}_{jt-1}} + \gamma_2 \frac{\text{Fixed Expenses}_{jt}}{\text{Assets}_{jt-1}}$$

Appendix

A Details of marginal cost estimation

We begin by rewriting the bank value function as an expectation that is linear in the marginal cost parameters:

$$\psi_{\sigma,j,k} \coloneqq \begin{pmatrix} \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t \rho_{jks} Dep_{jks}\Big] \\ \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t R_s Dep_{jks}\Big] \\ \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t Dep_{jks}\Big] \end{pmatrix}$$

$$V_{\sigma,j,k} \equiv \psi_{\sigma,j,k}^T \cdot \begin{pmatrix} 1 \\ -\zeta \\ -\omega - \gamma' \mathbf{Z}_{jt} \end{pmatrix}$$

For any alternative spread-setting policy, $\tilde{\sigma}$, it should be the case that a bank is better off at the original estimated policy function, σ, than at the alternative one:

$$g(\sigma, \tilde{\sigma}) = (\psi_\sigma - \psi_{\tilde{\sigma}})^T \cdot \begin{pmatrix} 1 \\ -\zeta \\ -\omega - \gamma' \mathbf{Z}_{jt} \end{pmatrix} \geq 0$$

where $\psi_{\tilde{\sigma}}$ denotes the analogous vector of expectations under the perturbed policy.

To operationalize this idea, we construct sixty such alternative policies. Twenty perturb the responsiveness of deposit spreads to the short rate, twenty perturb responsiveness to lagged market share, and twenty are multiplicative deviations. All perturbations are normal draws centered around the estimated policy function. For each alternative, we perturb the policy function of a single randomly-selected bank in each county-year, holding the policy function of its competitors fixed at the original estimated policy. Each policy alternative thus generates one inequality per county-year, namely that the selected bank is better off at the estimated policy function than at the perturbed one.

We compute expectations over 1000 paths of: (i) interest rates drawn from instantaneous forward rates, as described in Section 5.3.3 ; and (ii) private shocks drawn from the empirical distribution of policy function residuals. We set the discount rate $\beta = 0.9$ and forward-simulate out to fifty periods in the future. We additionally assume that the only source of growth in deposit market size is from interest accumulation.

Following Bajari et al. (2007), we estimate marginal costs by minimizing the sum of squared violations over all inequalities:

$$(\omega, \zeta, \gamma) \coloneqq \arg\min \sum (\min\{g(\sigma, \tilde{\sigma}), 0\})^2$$

While all cost parameters are jointly identified, it is helpful to consider the variation driving the estimates of each parameter. Consider a bank's decision of whether or not to deviate in some way from the equilibrium policy function. Doing so changes both the expected path of spreads as well as the expected path of deposit volumes. To fix ideas, let us suppose the bank is considering a deviation that tends to increase spreads. Holding the path of deposit volumes fixed, this will increase franchise value. But through the active demand system, it will also reduce deposit volumes and thus franchise value. Both of these effects are fully quantifiable without any knowledge of the marginal cost parameters. That said, the change in deposit volumes also changes the present value of costs, and this effect is fully characterized by the marginal cost parameters. The question therefore becomes whether the net benefit of the deviation (i.e., the sum of the two "cost-independent" effects) outweighs the cost effect. The BBL estimator penalizes candidate vectors for which this is the case, and the penalty scales in the size of potential profits from deviating.

Importantly, changes in deposit volumes differentially impact each of the four cost terms at different points in the state space, thereby allowing us to separately identify each parameter. A given volume shock: (i) has the same effect on the unweighted path of volumes across all states; (ii) impacts the rate-weighted path of volumes much more in a high-rate environment than a low rate environment; and (iii) impacts the cost-shifter-weighted paths of volumes much more when costs are higher. Thus, we pin down ζ by selecting inequalities from different years of our sample (i.e., as rates vary), whereas we pin down the γ parameters by sampling inequalities widely from both the cross-section and the time-series.

While this identification argument holds for any set of selected policy deviations, we choose our alternative policies in part to make the separate identification of the key cost parameters more transparent. For example, perturbations to the federal funds pass-through provide an additional dimension of differential variation in the rate-weighted path of volumes, since they strengthen the correlation between the path of rates and the path of volumes.

B Computation of Markov Perfect and Bertrand-Nash Equilibria

B.1 Markov Perfect Equilibrium

We work with markets consisting of three banks and set sleepiness, ϕ, to $.935$ (the average estimated level in our data). We set rate sensitivity, α, to the estimated level of -388.7. We set ζ, the coefficient on the short rate in the marginal cost equation, to the estimated $.695$. We set "base" marginal costs to -30 bps for all three banks, which is about the 75th percentile of our estimated distribution.[37] Mirroring the empirical model, we add mean-zero, normally-distributed shocks to banks' chosen spreads, with the standard deviation set to the level in the estimated policy function (about 21 bps). Finally, we parameterize the short rate transition process using a discretized Cox et al. (1985) process with mean reversion, κ, of .11, a long-term average rate, θ, of .03, and a volatility, σ, of .08.

We solve for Markov Perfect equilibrium using value function iteration, as in Pakes and McGuire (1994). As is well-known, with strategic interactions between firms, this algorithm is not guaranteed to converge. We have found that a few modifications to the algorithm tend to smooth out best response correspondences, thereby reducing cycling behavior and facilitating convergence: (i) interpolating the value function off of a discretized grid with flexible splines, as in Fowlie et al. (2016); and (ii) imposing a small quantity adjustment cost. There is also no guarantee that any Markov Perfect equilibrium is unique, but in practice we have found little to no evidence of multiple equilibria.

For panel (a) of Figure 5, we endow all three banks in the market with equivalent non-price deposit product qualities δ and solve for a symmetric equilibrium. We then plot the policy function, at different levels of the short rate, for a bank with 10% market share when one other bank has 10% share and the third bank has 80%.

For panel (b) of Figure 5, we perform the following procedure for each red dot on the graph:

- pick the vector of lagged market shares at which we want to start (i.e., to target some value of HHI)
- solve for the product quality vector $\boldsymbol{\delta}$ that generates these market shares in static equilibrium (i.e., when all depositors are active)
- solve for the Markov Perfect Equilibrium of the dynamic game via value function iteration
- forward-simulate both the static and dynamic games over 100 paths of interest rates and policy shocks, each 100 years in length, recording the average market-level expected markup across all runs.

To obtain starting values of lagged market shares for the dynamic game, we assume that in the first period everyone is active, and we solve for the Nash equilibrium of the first period

[37] None of our results are sensitive to the exact choice of marginal costs.

stage game holding the continuation values fixed at the solution to the infinite-horizon Markov Perfect equilibrium. Then, for each simulation run, we begin at these lagged market shares and run a burn-in of 50 periods, at which point we then record the results of the forward simulation. All dots begin with a starting aggregate short rate of 5%.

B.2 Bertrand-Nash Equilibrium

We again work with markets with three banks, but we now set sleepiness, ϕ, to 0. All other parameters are as described in the prior section. For panel (a) of Figure 5, we endow all three banks in the market with equivalent non-price deposit product qualities δ and solve for a symmetric Markov Perfect equilibrium using the same value function iteration as above, though with $\phi = 0$ this is isomorphic to solving the usual nonlinear fixed point problem characterizing equilibrium Bertrand-Nash markups.

For panel (a) of Figure 5, we endow all three banks in the market with equivalent non-price deposit product qualities δ and solve for a symmetric equilibrium. We then plot the policy function, at different levels of the short rate, for a bank with 10% market share when one other bank has 10% share and the third bank has 80%.

For panel (b) of Figure 5, we perform the following procedure for each teal dot on the graph:

- pick the vector of lagged market shares at which we want to start (i.e., to target some value of HHI)
- solve for the product quality vector $\boldsymbol{\delta}$ that generates these market shares in static equilibrium (i.e., when all depositors are active)
- solve for the Bertrand-Nash equilibrium of the static game
- forward-simulate the static game over 100 paths of interest rates and policy shocks, each 100 years in length, recording the average market-level expected markup across all runs.

All dots begin with a starting aggregate short rate of 5%.

C Background on FCA Policy and Equivalence to Static Counterfactual

C.1 Background on the FCA Policy

In July of 2023, the FCA published a 14-point plan with the stated goal of increasing the deposit rates offered to consumers.[38] Under the power of "Consumer Duty", a set of newly-enacted consumer protection regulations, the FCA began to require banks with deposit rates substantially lower than the "base rate" (the aggregate short-term interest rate) to justify how their offered rates provided consumers with a "fair value." If the FCA deemed such justifications to be insufficient, it would have recourse to levy fines onto the banks or place restrictions on their business operations. The plan also specified that the FCA would review the timing of bank deposit rate changes after any adjustment in the base rate. In September of 2024, the FCA published another statement announcing that they would investigate evidence of asymmetric deposit betas: banks that passed on rate reductions more quickly and fully than rate increases.[39] Jointly, these statements express a clear desire on the part of the FCA to restrict the relationship between aggregate rates and bank deposit rates. We therefore use our structural estimates to simulate a counterfactual policy in which banks are forced to fully pass aggregate rate fluctuations through to their deposit rates.

C.2 Equivalence to Static Counterfactual

In this section, we demonstrate that constraining banks to fully pass through aggregate interest rate changes to their deposit rates is in fact isomorphic to eliminating depositor sleepiness (i.e., setting $\phi = 0$).

To see this, suppose that at time 0 all depositors are active. Each bank seeks to maximize franchise value by choosing the constant spread it will charge between the federal funds rate and its offered deposit rate. The bank's problem is:

$$\max_{\rho} M_0 s_0(\rho - c(R_0^F)) + \sum_{t=1}^{\infty} \beta^t \mathbb{E}[M_t s_t(\rho - c(R_t^F))] \tag{A1}$$

where s_t refers total deposit market share, $c(\cdot)$ refers to marginal cost and deviates from its initial level solely due to fluctuations in R_t^F, and M_t refers to market size at time t and grows

[38] Financial Conduct Authority. "FCA sets out 14-point action plan on cash savings." Available at: https://www.fca.org.uk/news/press-releases/action-plan-cash-savings. Accessed December 17, 2024

[39] Financial Conduct Authority. "FCA update on cash savings – September 2024." Available at: https://www.fca.org.uk/publications/multi-firm-reviews/fca-update-cash-savings-september-2024. Accessed December 17, 2024.

at a deterministic rate g.[40] We can rewrite equation (A1) as:

$$\max_{\rho} M_0 s_0(\rho - c(R_0)) + \sum_{t=1}^{\infty} \beta^t \mathbb{E}[\Big(\phi_t M_{t-1}(1+g)s_{t-1} + (1-\phi_t)M_t s_t^{active}\Big)(\rho - c(R_t^F))]$$

Because all depositors are active in period 0, we clearly have $s_0 = s_0^{active}$. But since active depositors have the same preferences in every period, it follows that $s_0 = s_0^{active} = s_t^{active} = s_t, \ \forall t$. Noting also that $(1+g)M_0 = M_1$, the first term in the infinite sum thus collapses to $\beta\mathbb{E}[M_1 s_0(\rho - c(R_1)]$, the second term to $\beta^2\mathbb{E}[M_2 s_0(\rho - c(R_2)]$, and so forth, leaving us with:

$$\max_{\rho} s_0 \sum_{t=0}^{\infty} \beta^t \mathbb{E}[M_t(\rho - c(R_t^F))]$$

The bank's first-order condition is:

$$\frac{\partial s}{\partial \rho}\Big(\sum_{t=0}^{\infty} \beta^t \mathbb{E}[M_t(\rho - c(R_t^F))]\Big) + s_0\Big(\sum_{t=0}^{\infty} \beta^t M_t\Big) = 0 \tag{A2}$$

This pricing equation is precisely the present-value formulation of the static Bertrand-Nash first-order condition. We solve for outcomes under the FCA policy in each market in our data by finding the vector of spreads and shares under which equation (A2) holds for all banks in the market.

[40] In our forward simulation, market size actually grows in accordance with total interest accumulation and the sleepy deposit stock grows in accordance with bank-specific interest accumulation – both, in turn, depend on offered deposit rates in the previous period. This force is quantitatively unimportant, and we therefore assume exogenous market size growth in our counterfactual for the sake of tractability. That said, to further ensure comparability between our baseline and counterfactual results, we assume that market size grows each year by the average offered deposit rate in our sample (about 34 bps).

D Franchise Value Decomposition

The proof laid out in the previous section is constructive in the sense that it provides a simple algorithm to compute outcomes under the counterfactual FCA policy. To perform a true decomposition of deposit franchise value into a "sleepy" component and an "inelastic" component, however, we need to retain a zero lower bound on deposit rates. This complicates the required procedure in two, related ways. First, in some markets, some banks may not be able to charge their Bertrand-Nash markup without breaching this constraint. Solving for equilibrium then consists of both determining which banks are constrained in equilibrium and, given the set of constrained banks, solving for spreads and market shares that satisfy the first-order conditions for the unconstrained banks. That is, we search for a vector of deposit spread ρ^* that, for all banks j in the market, satisfies:

$$\rho_j^*(R^F) = \begin{cases} c_j - \frac{1}{\alpha(1-s_j(\boldsymbol{\rho^*},\boldsymbol{\delta}))}, & \text{if } R^F - (c_j - \frac{1}{\alpha(1-s_j(\boldsymbol{\rho^*},\boldsymbol{\delta}))}) \geq 0 \\ 0, & \text{otherwise} \end{cases}$$

Second, consequently, equilibrium spreads and market shares are now a function of the short rate. Thus, unlike in the FCA computation, we need to solve for the equilibrium in each market *for each level of the short rate*.

Solving for the equilibrium of the more than 54,000 markets in our data for each of 51 periods across 1000 simulations (i.e., for each draw of the short rate), is computationally infeasible. Instead, we discretize the range of short rates into an equally-spaced grid from 0 to .05, with buckets of 25 bps in width. We then solve for equilibria in all markets at each point of the discretized grid. Finally, we perform the forward simulation over the same 1000 paths of drawn short rates used in our baseline two-step estimation by assigning each drawn short rate to a discretized outcome using the floor function (e.g., a short rate draw of 80 bps is assigned to the equilibrium outcomes calculated for a short rate of 75 bps) and then averaging the resulting discounted sum of profits:

$$\mathbb{E}[\text{Static Franchise Value w/ZLB}_{jkt}] = \sum_{sim=1}^{1000} \sum_{t=0}^{50} \beta^t (1+g)^t M_t s_j^*(R_{sim,t}^F, \boldsymbol{\delta})(\rho_j^*(R_{sim,t}^F) - c_j(R_{sim,t}^F))$$

As in the FCA computation, we assume that the size of the market grows deterministically at the average offered deposit rate in the sample (34 bps).

E Extension: Price Discrimination between New and Incumbent Depositors

Recall that in our baseline set-up, banks must choose a single deposit spread each period to offer to both their "incumbent" depositors (i.e., depositors who banked with them in the prior period) and their "new" depositors. In this extension, we allow banks to offer separate spreads to their incumbent and new depositors, and we estimate separate marginal cost parameters for servicing new versus incumbent deposits.

Importantly, because banks cannot distinguish which of their incumbent depositors are active versus asleep, the fundamental invest-versus-harvest tradeoff when setting the incumbent spread substantially resembles our initial formulation. The primary difference in the incumbent spread decision relative to baseline is that the propensity of running off awake incumbent depositors with high spreads will be larger if competitors choose to offer low spreads to their new depositors in equilibrium. With regard to the choice of spread offered to new depositors, offering a very low "new spread" does not come with a cost of offering better terms to many asleep depositors, though it still has the benefit of carrying over a larger deposit base into the following period.

Under this set-up, bank j's problem becomes:

$$V(\mathbf{S}_t, \boldsymbol{\sigma}^i, \boldsymbol{\sigma}^n) = \mathbb{E}_{\boldsymbol{\chi}_{jt}}[Dep^i(\mathbf{S}_t, \sigma_j^i, \boldsymbol{\sigma}^n)(\sigma_j^i - c^i(\mathbf{S}_t, \chi_{jt}^i)) \,|\, \mathbf{S}_t] + \mathbb{E}_{\boldsymbol{\chi}_{jt}}[Dep^n(\mathbf{S}_t, \boldsymbol{\sigma}^i, \boldsymbol{\sigma}^n)(\sigma_j^n - c^n(\mathbf{S}_t, \chi_{jt}^n)) \,|\, \mathbf{S}_t] + \beta \mathbb{E}_{\mathbf{S}_{t+1}, \boldsymbol{\chi}_{t+1}}[V(\mathbf{S}_{t+1}, \boldsymbol{\sigma}^i, \boldsymbol{\sigma}^n) \,|\, \boldsymbol{\sigma}^i, \boldsymbol{\sigma}^n, \mathbf{S}_t]$$

where the i and n superscripts denote, respectively, "incumbent" and "new."

$$\psi_{\sigma,j,k} \coloneqq \begin{pmatrix} \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t \rho_{jks}^i Dep_{jks}^i\Big] \\ \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t R_s Dep_{jks}^i\Big] \\ \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t Dep_{jks}^i\Big] \\ \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t \rho_{jks}^n Dep_{jks}^n\Big] \\ \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t R_s Dep_{jks}^n\Big] \\ \mathbb{E}\Big[\sum_{s=t}^{\infty} \beta^t Dep_{jks}^n\Big] \end{pmatrix}$$

$$V_{\sigma,j,k} \equiv \psi_{\sigma,j,k}^T \begin{pmatrix} 1 \\ -\zeta^i \\ -\omega^i - \gamma^{i\prime} \mathbf{Z}_{jt} \\ 1 \\ -\zeta^n \\ -\omega^n - \gamma^{n\prime} \mathbf{Z}_{jt} \end{pmatrix}$$

We create alternative polices as in the baseline case, but we now perturb "incumbent" spreads in half of the runs and "new" spreads in the other half. For any alternative *pair* of spread-setting policies, $(\tilde{\sigma}^i, \tilde{\sigma}^n)$, it should be the case that a bank is better off at the original estimated policy functions, (σ^i, σ^n), than at the alternative ones:

$$g((\sigma^i, \sigma^n), (\tilde{\sigma}^i, \tilde{\sigma}^n)) = (\psi_\sigma - \psi_{\tilde{\sigma}})^T \begin{pmatrix} 1 \\ -\zeta^i \\ -\omega^i - \gamma^{i\prime}\mathbf{Z}_{jt} \\ 1 \\ -\zeta^n \\ -\omega^n - \gamma^{n\prime}\mathbf{Z}_{jt} \end{pmatrix} \geq 0$$

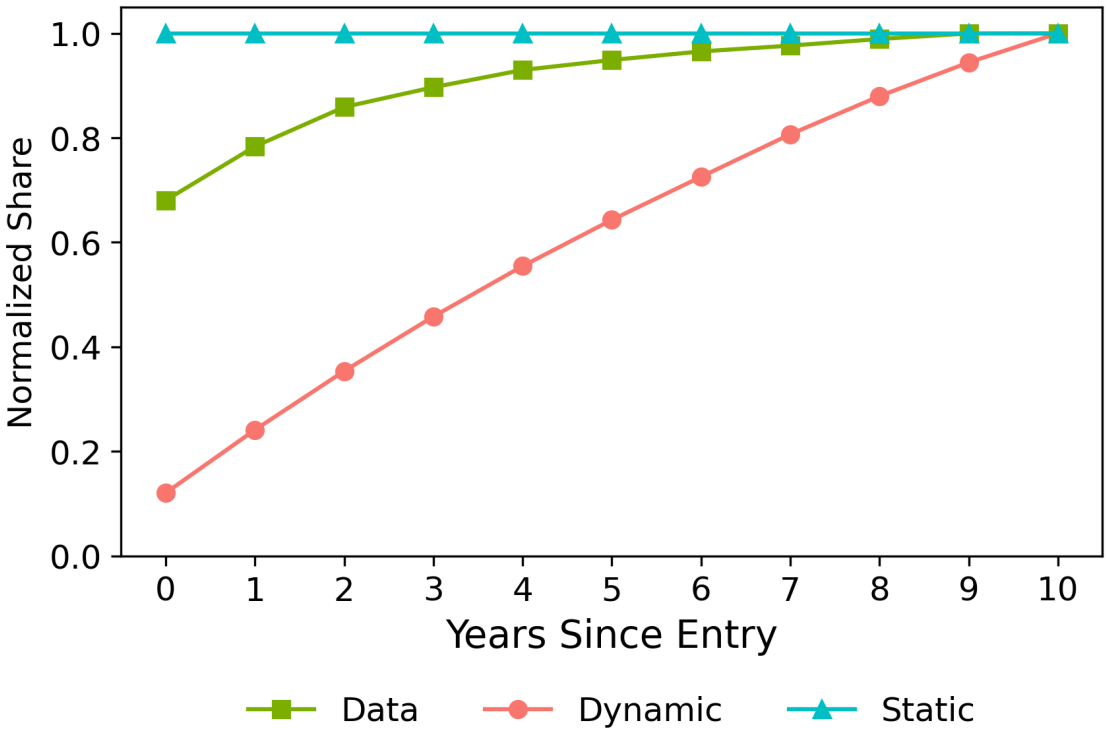
where $\psi_{\tilde{\sigma}}$ denotes the analogous vector of expectations under the perturbed policy. With inequalities in-hand, we choose the vector of cost parameters that minimizes the sum of squared violations from optimality, only now we estimate separate parameters for incumbent and new deposits:

$$(\omega^i, \zeta^i, \gamma^{\mathbf{i}}, \omega^n, \zeta^n, \gamma^{\mathbf{n}}) := \arg\min \sum (\min\{g((\sigma^i, \sigma^n), (\tilde{\sigma}^i, \tilde{\sigma}^n)), 0\})^2$$

We report estimates in Appendix Table A2. In Figure A4, we additionally plot the time-varying distributions of current markups charged to new and incumbent depositors.

Appendix Tables and Figures

Figure A1: Sleepiness versus Persistent Preference Heterogeneity (alternative normalization)

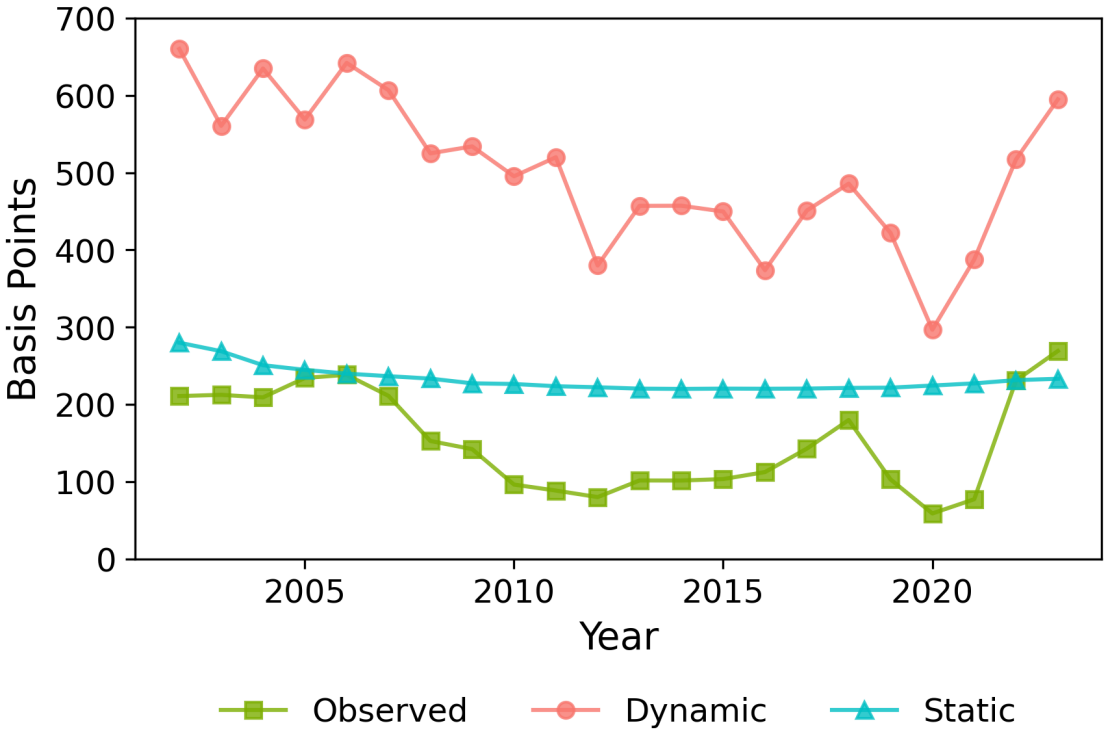


Notes: Figure A1 plots the progression of market shares for entering firms as a function of the number of years since entry:

$$\text{normalized share}_t = \frac{\text{share}_t}{\text{share}_{10}}$$

The green line plots medians from Summary of Deposits data, the red line plots medians from forward-simulations of a 3-bank Markov Perfect equilibrium, and the teal line reflects the fact that static market shares are time-invariant. See Appendix B for additional details on our computation of Markov Perfect equilibria.

Figure A2: Estimated Franchise Values versus Core Deposit Intangible Valuations

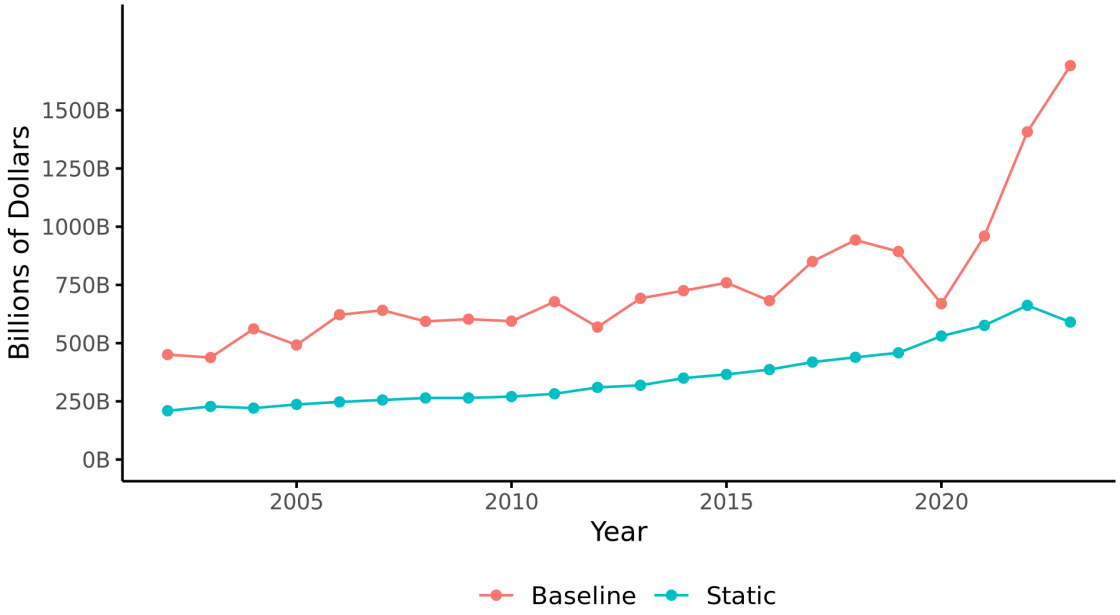


Notes: Figure A2 plots the average Core Deposit Intangible (CDI) announced in realized mergers (green line) against the time-series averages of as close an analog as we can construct from our two-step estimates (red line) and static simulations (teal line). The Core Deposit Intangible, a part of the purchase price allocation process in mergers between banks, is calculated by projecting the present value of savings a bank realizes by offering below-market interest rates on core deposits, rather than raising other forms of more expensive debt, over the deposits' "useful life" (almost always taken to be ten years).[41] We construct analogous statistics using our two-step estimates and static simulations by calculating franchise values only using cash flows over the next ten years. There remain two important differences between our construction of franchise values and the CDI. First, our analysis does not incorporate the fixed costs of running the deposit franchise. This means that our estimates will be biased upward relative to the CDI. Second, if Modigliani and Miller (1958) fails, our estimates will incorporate the value of the lending franchise. In contrast, even if Modigliani and Miller (1958) does fail, banks are instructed to calculate CDI using the gap between offered deposit rates and the rate they would pay on their next-cheapest form of debt (usually the FHLB advance rate). This, again, biases our estimates upwards relative to reported CDI's.

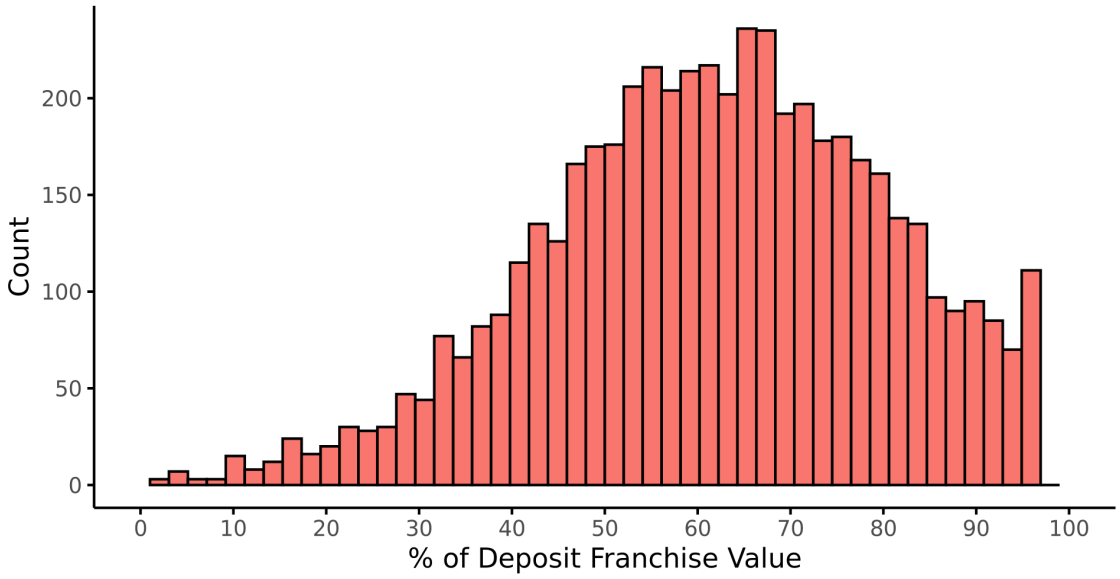
[41] See, for example, https://www.bdo.com/insights/industries/core-of-the-core-deposit-intangible-valuation-and-trends. Accessed 08/05/2025.

Figure A3: Franchise Value Decomposition

(a) Aggregate Banking Sector



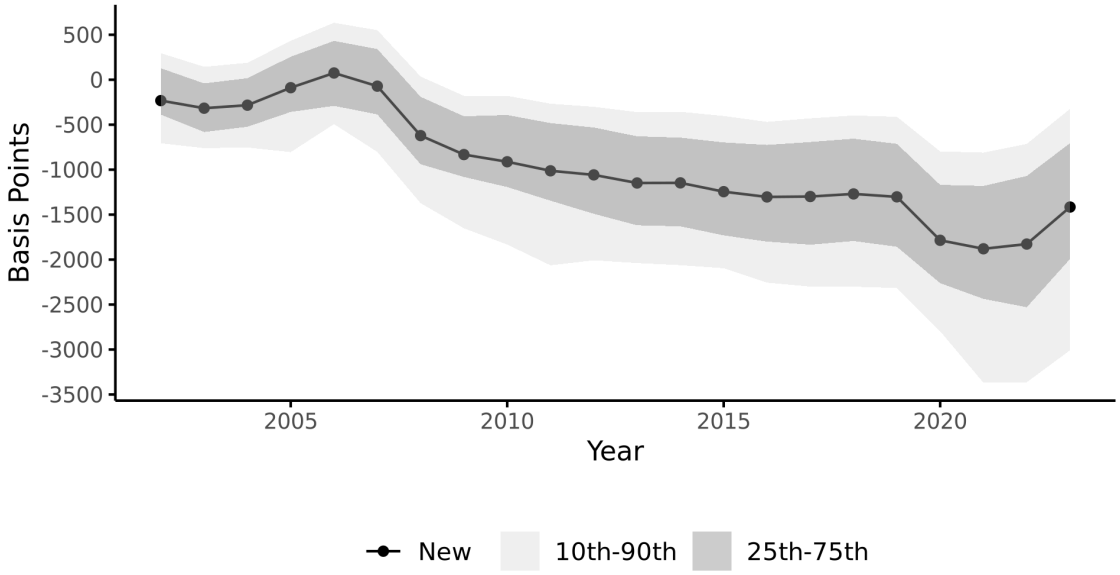
(b) 2010 Cross-Section



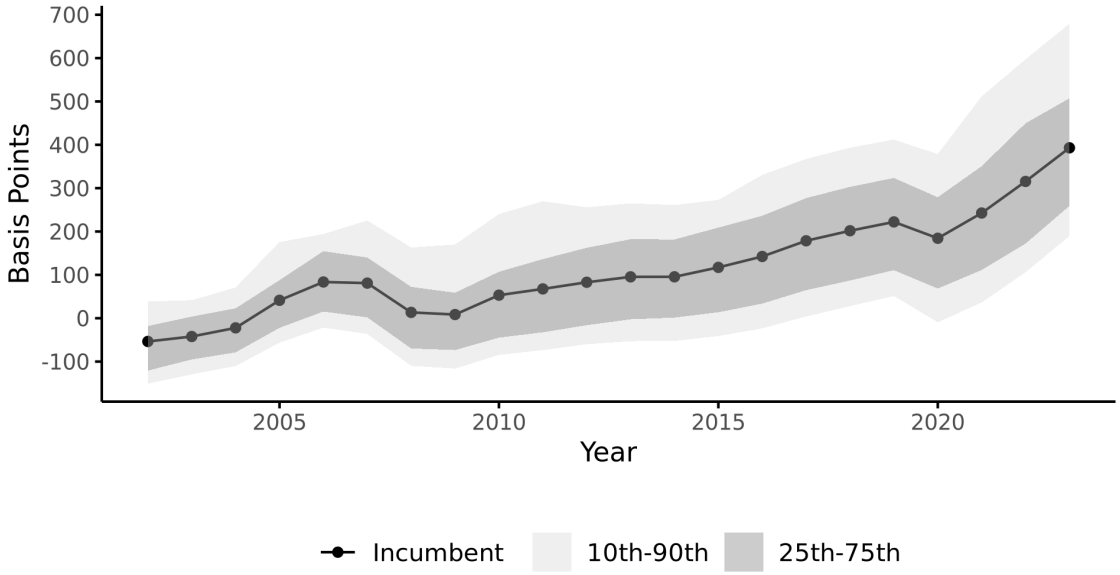
Notes: Panel (a) of Figure A3 displays total deposit franchise value in the U.S. banking sector under our baseline estimates and in our simulated static equilibria eliminating depositor sleepiness. To facilitate comparison with other aggregate quantities, we scale the yearly totals by the fraction of total deposits that are accounted for by observations in our sample. Panel (b) of Figure A3 displays the cross-sectional bank-level distribution of the fraction of deposit franchise value that we attribute to sleepy demand, with the remainder attributed to inelastic demand, using data from 2010.

Figure A4: Markup Estimates under Price Discrimination

(a) New Depositors

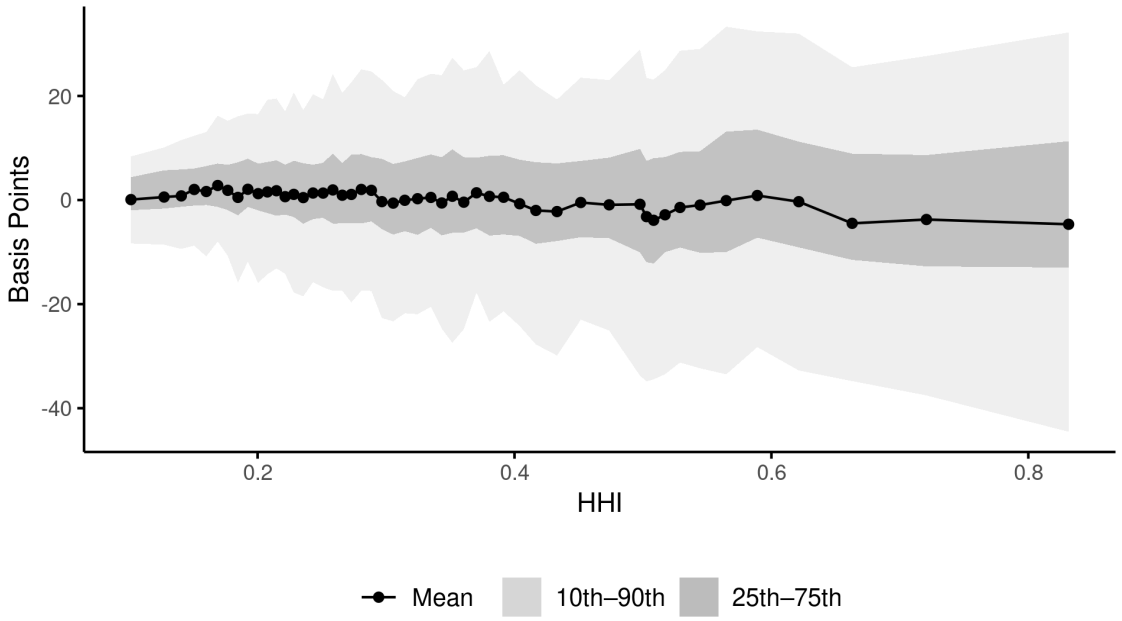


(b) Incumbent Depositors



Notes: Panel (a) of Figure A4 displays the yearly bank-county level distributions of current markups charged to new depositors. Panel (b) of Figure A4 displays the yearly bank-county level distributions of current markups charged to incumbent depositors. We winsorize markups at the 1% level to reduce the influence of extreme outliers.

Figure A5: Deposit Spreads versus Market Concentration



Notes: Figure A5 plots a market-level bin-scatter plot of deposit spreads against the Herfindahl Hirschman Index (HHI). Market-level deposit spreads are demeaned by year.

Table A1: Policy Function for Spreads Offered to New Depositors

	(1)
Constant	0.008***
	(0.002)
R_t	0.197***
	(0.029)
R_{t-1}	0.098***
	(0.031)
δ_{jkt}	0.0003**
	(0.0001)
Log(Deposit Share)$_{jkt-1}$	-0.0003*
	(0.0002)
$R_{t-1} \times$ Log(Deposit Share)$_{jkt-1}$	0.005
	(0.005)
$\sum_{\ell \neq j}$ Log(Deposit Share)$_{\ell kt-1}$	1.03×10^{-5}
	(6.94×10^{-6})
$\frac{\text{Wages}_{jt}}{\text{\# Employees}_{jt-1}}$	-3.88×10^{-6}
	(1.42×10^{-5})
$\sum_{\ell \neq j} \frac{\text{Wages}_{\ell t}}{\text{\# Employees}_{\ell t-1}}$	1.32×10^{-6}**
	(5.83×10^{-7})
$\frac{\text{Fixed Expenses}_{jt}}{\text{Assets}_{jt-1}}$	0.070
	(0.147)
$\sum_{\ell \neq j} \frac{\text{Fixed Expenses}_{\ell t}}{\text{Assets}_{\ell t-1}}$	-0.005
	(0.010)
fraction over 65$_k$	0.001
	(0.002)
fraction over \$50k HH income$_k$	-0.002
	(0.001)
fraction w/internet$_k$	-0.007***
	(0.002)
Observations	24,258
R^2	0.364

Notes: Table A1 reports OLS estimates of the policy function for the "promotional gap": the difference between promotional CD rates of maturity $\leq$ 6 months and regular \$10k 3-month CD rates. We proxy for the spread offered to new depositors as the baseline \$10k money market spread less this promotional gap. Standard errors are clustered by bank-year. *** p<0.01, ** p<0.05, * p<0.10.

Table A2: Marginal Cost Estimates Under Price Discrimination

ζ^i	0.670
	(0.163)
ω^i	0.031
	(0.035)
γ_1^i	-0.001
	(0.000)
γ_2^i	0.264
	(0.226)
ζ^n	-0.796
	(0.337)
ω^n	-0.135
	(0.177)
γ_n^i	0.003
	(0.002)
γ_2^n	0.127
	(0.148)

Notes: Table A2 reports the results of the two-step estimator based on Bajari et al. (2007). We construct 60 alternative policies: 15 perturb the responsiveness of incumbent deposit spreads with respect to the short rate; 15 perturb the responsiveness of new deposit spreads with respect to the short rate; 15 multiply incumbent spreads by a constant randomly drawn around 1; and 15 multiply new spreads by a constant randomly drawn around 1. For each alternative policy, independent draws are taken for one randomly-selected bank in each county-year – 3,278,400 total inequalities are used in estimation. Each expectation is constructed by averaging over 1,000 simulated paths of: (i) interest rates, generated from estimated yield curves as described in Section 5.3.3; and (ii) policy shocks for both new and incumbent spreads, drawn from the distribution of policy function residuals (censored at the 5th and 95th percentiles). Standard errors are calculated via bootstrap: we resample, with replacement, 500 times – each time constructing datasets of the same length as the total number of inequalities. We then calculate standard errors as the standard deviation of the vector of estimates.

Costs are parameterized as:

$$c_{jkt}^i = \zeta^i \text{Fed Funds}_t + \omega^i + \gamma_1^i \frac{\text{Salaries}_{jt}}{\text{\# Employees}_{jt-1}} + \gamma_2^i \frac{\text{Fixed Expenses}_{jt}}{\text{Assets}_{jt-1}}$$

$$c_{jkt}^n = \zeta^n \text{Fed Funds}_t + \omega^n + \gamma_1^n \frac{\text{Salaries}_{jt}}{\text{\# Employees}_{jt-1}} + \gamma_2^n \frac{\text{Fixed Expenses}_{jt}}{\text{Assets}_{jt-1}}$$

Table A3: Forward-Looking Demand Estimates

	(1)	(2)
ρ_{jkt}	-388.7***	-769.6***
	(29.3)	(39.5)
$\mathbb{E}[\sum_{s=t+1}^{\infty}\left(\beta^s u_{ijks}\prod_{m=t+1}^{s}\phi_m\right)]$		0.238***
		(0.0009)
Observations	294,809	294,809
Rate Elasticity of Demand	0.525	1.039
R^2	0.626	0.801
Bank fixed effects	✓	✓
County-Year fixed effects	✓	✓

Notes: Table A3 reports control function estimates obtained from a Berry (1994) demand system. We control for a 1st-order polynomial of the residuals obtained from equation (9): a spread-setting control function that uses salaries and fixed expenses as cost shifters. Standard errors are clustered by bank-year. Column (1) replicates our baseline demand specification, which assumes depositor myopia, as presented in column (4) of Table 4. Column (2) reports our forward-looking estimates, in which we control for the present value of future utility in the event of remaining asleep and continuing to bank with the bank chosen today. The rate elasticity of demand is calculated assuming a market share of 10% and an offered interest rate at the sample median (about 15 bps). *** p<0.01, ** p<0.05, * p<0.10.