

# Branch Closures and Spatial Competition in Retail Banking

Alvaro Contreras

Boston University

[alvarocm@bu.edu](mailto:alvarocm@bu.edu)

October 21, 2026

[Latest version of this paper](#)

## Abstract

Across retail markets, firms are retreating from physical locations, generating cost savings from closures but leaving consumers who value proximity with fewer alternatives and longer travel distances. Mergers may exacerbate this trade-off because of the additional closures generated by the overlapping structure of the combined location networks. I study this trade-off in U.S. retail banking. I combine branch-level deposits, bank balance sheets, and census-tract demographics to estimate a model of retail deposit demand and dynamic branch closures by banks. In the model, banks choose which branches to close permanently by evaluating the future profitability of their entire branch network, which depends on branch-level deposit intake and operating costs. Through the demand for deposits, the model captures heterogeneity in households' valuation of proximity and measures how the consequences of branch closures vary across households. I find that proximity matters more for households in older neighborhoods, and closures are more likely for branches that are small, recently acquired or located near offices of the same bank. Counterfactuals show that mergers can accelerate network contraction, reducing depositor access and surplus. Policies aimed at protecting branches from closing can raise depositor surplus, but the costs of maintaining additional branches exceed consumer gains.

**JEL Codes:** L13, L22, G21, G34, R12

**Keywords:** retail networks, deposit markets, branch closures, spatial competition, bank consolidation, local market power, depositor welfare.

# 1 Introduction

In retail markets, technological change is reshaping the way consumers use firms’ physical locations. As more transactions move online, firms increasingly need to reconsider which stores and branches to keep open, leading to the closure of locations that have become less valuable or redundant. In industries such as grocery, pharmacy, clothing, and banking, the consequences of these changes depend on consumers’ reliance on in-person access and their ability to substitute toward digital alternatives. On the one hand, reducing the number of physical locations may generate opportunities to realize efficiencies and reduce operating costs. On the other hand, it can also reduce the alternatives available to consumers in local markets, with distributional consequences. This trade-off is particularly salient in industries characterized by frequent mergers, which are often justified by the cost synergies that may arise due to the elimination of overlapping physical networks, but where concerns about local market concentration and market power can also arise.

I study this trade-off in the U.S. banking industry, where the number of physical branches has contracted substantially, even as proximity remains relevant for depositors. Between 2011 and 2025, the number of commercial bank branches fell by 16% (from 82,564 to 69,167), despite office location remaining the most commonly cited reason households gave for choosing the institution for their main checking account, at 42%.<sup>1</sup> The resulting contraction of branch networks has raised concerns about reduced access to in-person banking services, particularly in underserved communities, and the emergence of *banking deserts*.<sup>2</sup> However, the branch decline may also be a consequence of depositors caring less about proximity than they did before the rise of digital banking. I combine branch-level deposits, bank balance-sheet information, and census-tract demographics to study three questions. How do merger-driven consolidation and changes in the value of proximity for depositors shape the evolution of local branch networks? What are the welfare consequences of the resulting network configuration? And can policies preserve valuable local access while allowing banks to rationalize redundant locations?

To answer these questions, I contribute to the literature that describes the classic merger trade-off between cost efficiencies and increased market power ([Williamson, 1968](#); [Farrell and Shapiro, 1990](#); [Ashenfelter et al., 2015](#)). In general, classic merger analysis focuses on how mergers might reduce marginal costs through synergies and how firms adjust their prices

---

<sup>1</sup>2022 Survey of Consumer Finances.

<sup>2</sup>See, for example, the U.S. Senate Committee on Banking, Housing, and Urban Affairs, “[Brown Pushes Bank Regulators to Protect Consumers and Communities When Reviewing Mergers](#),” June 17, 2024, and “[Brown Presses OCC to Address Bank Mergers and Branch Closures in Low-Income, Rural, and Minority Communities](#),” March 2, 2023.

accordingly, affecting consumers but also generating efficiency gains from lower costs per good sold. In the setting of this paper, I focus on how mergers can reduce recurring operating costs through additional branch closure decisions made by banks, given the overlapping structure of the combined branch network, and how these closures affect depositors through remaining branch availability and distance. In this setting, efficiencies related to cost savings can arise from banks internalizing the cannibalization among branches in the combined network and minimizing the recurring operating costs tied to physical locations. Additionally, even though branch closures might be efficient from the point of view of cost savings, they might also reflect increased market power exerted through the removal of branches near depositors. Even in the absence of changes in prices, closing a newly acquired branch that is close to an existing branch of the acquiring bank implies that the bank can still retain economic benefits from attracting deposits in the area, but generate longer distances for depositors who value proximity.

To capture these economic mechanisms, I build a model of retail deposit demand and dynamic branch networks. On the demand side, heterogeneous households allocate deposits across nearby branches and an outside option that represents other types of investments. Household deposit allocations depend on the distance elasticity of deposit demand, which interacts with local demographics to represent how older households may have stronger preferences for closer in-person banking, and varies over time to capture how the value of proximity has changed in recent years. On the banks' side, every period their profits combine the revenue generated from deposit funding, the recurring cost of operating branches, and a temporary cost for branches obtained through mergers. Since closure decisions are irreversible, every period banks choose which branches to close as part of a dynamic spatial portfolio. When a branch closes, the demand system allows deposits to be reallocated across other offices of the same bank, competing banks, and the outside option.

In this model, mergers are exogenous and they are represented through changes in the ownership of the branch locations that are involved in merger transactions. After a merger, banks re-optimize and branches may be more likely to close due to additional operating costs for managing newly acquired branches and the internalization of business stealing effects through the demand of the combined branch network. However, if the acquiring bank has an overall higher profitability than the previous owner and its bank-specific demand effect is higher (reflecting stronger brand effects, for example), then the realized profits from the transferred branch might be higher and closure less likely than before the merger.

I then turn to the data and identify three patterns that the model is able to capture. First, branch deposits are larger when total income from households is concentrated closer to the office. Second, closures are more common among smaller branches, branches located

near other offices of the same bank, branches serving areas with older populations, and recently acquired branches. Among the offices obtained through mergers, the presence of an existing branch of the new owner within five miles is associated with a higher probability of exit within three years. Third, the negative relationship between household-level checking-account balances and distance to the household’s main institution has weakened between 2016 and 2022, suggesting that the reliance on in-person access has fallen in this period. These facts establish that distance matters for deposit demand, that overlap from branches owned by the same bank predicts branch closures, and that the value of proximity differs across households and over time.

I estimate the model sequentially. The demand estimation step extends the spatial-retail demand approach of [Ellickson et al. \(2020\)](#) to a banking setting, and estimates the parameters that best match predicted branch deposits with observed data. The demand estimates confirm that proximity remains economically relevant. The distance coefficient is negative throughout the sample, indicating the disutility of facing larger distances between households and branches. Sensitivity to distance is substantially stronger in older neighborhoods, indicating relevant distributional consequences of closures for this demographic group. The estimates imply that the welfare compensation for doubling the distance to the branch from a household that belongs to the older quartile is equivalent to 44 basis points, almost twice the cost for a household of the younger quartile (23 basis points). However, even though the time-varying estimates for the elasticity of distance are consistent with a decline in the value of branch proximity in recent years, the magnitude of this change is not statistically significant.

The supply estimation step follows the dynamic discrete choice literature ([Bajari et al., 2007](#); [Pakes et al., 2007](#); [Aguirregabiria and Mira, 2007](#)) to relate the observed closure behavior by banks to the model’s dynamic profit payoff components. The supply estimates show that foregone deposit revenue discourages closure, while operating costs and recent acquisition exposure create incentives to close locations. The estimated per-period operating cost of maintaining a branch open is equivalent to \$4.0 million, which is reasonably close to the upper end of similar estimates for the industry in the literature ([Aguirregabiria et al., 2016](#)). Considering that the median branch for the banks with the largest branch networks generated an annual revenue of \$1.7 million per branch in terms of the interest rate spreads on deposit funding during the 2011-2024 period, the cost savings estimates from branch closure are sizable.

I use the estimated model to solve counterfactual dynamic branch-closure games. Since solving the model for large branch networks that cover the entire geography of the United States is not computationally feasible, I perform these counterfactuals in an illustrative small

county to keep the solution tractable. I select Haywood County, North Carolina, which had 17 branches in 2011, a representative size compared to counties that experienced branch declines between 2011 and 2025. More importantly, there were six banks with large branch networks participating in the county and four branches were involved in merger transactions between large banks. These facts allow Haywood to be an ideal location to decompose the effects that the model can illustrate, particularly related to branch contraction, mergers, and other secular changes.

In the first set of counterfactuals, I decompose the effect of mergers by considering three scenarios for the initial ownership configuration of branches, starting from 2011. The baseline case imposes mergers as observed between 2011 and 2024, and can be considered a benchmark that captures how banks optimize these larger networks. In the second case, I keep branch ownership from 2011 constant, so no mergers are imposed, which would remove the business stealing internalization and cost-savings found in the merger-baseline benchmark. In the third case, I assign all branches to one bank, which serves as an upper bound on local consolidation in the county. Maintaining the 2011 ownership configuration raises the number of branches by 6.1% above the baseline, and increases discounted consumer surplus by 1.1%. The single-bank case produces a smaller response, reducing the number of surviving branches by 0.5% below the baseline, and lowering discounted consumer surplus by about 1.0%.

The second group of counterfactuals aims to inform the debate over whether bank merger guidelines should be revised to incorporate the evaluation of post-merger branch closures and their effects on vulnerable households. To do this, I impose various forms of branch protection policies that set relevant branches to be kept open across the simulation, effectively removing them from the closure action space of the bank owner. I evaluate three simple illustrative branch protection policies that differ in strength of protection. First, a five-mile access policy that focuses on census tracts that only have one branch location in their choice set. Second, a merger safeguard policy that prevents closure of branches that were transferred due to mergers. Third, an extreme policy that prevents every branch closure in the period. The results from these three counterfactuals indicate that preserving additional branches increases consumer surplus, but the bank-side cost of maintaining additional locations is an order of magnitude greater than the consumer gains, making restrictions costly in terms of aggregate welfare.

This paper contributes to the multi-location firm literature by providing an application where firms contract their physical networks, and where the model is able to quantify the trade-off between operating cost savings and reduced consumer access for heterogeneous households. In this setting, the key mechanism captured by the model is how closure decisions at individual locations are evaluated in the context of their effect on the profitability of the

entire bank network across space. In terms of policy, this paper provides a framework that can be used for prospective merger analysis and may be extended to other geographically distributed industries. Starting from a pre-merger network, the model is able to identify locations whose incentives to remain open change due to mergers and evaluate the access consequences of their potential closure.

## Related Literature

This paper contributes to four strands of the literature: spatial retail competition, mergers and endogenous network adjustments, banking competition, and access to financial services.

First, this paper builds on the empirical industrial organization literature on spatial competition and multi-location retail networks. A significant part of this work explores endogenous entry decisions (Seim, 2006) and market structure implications of the expansion of large retail chains and their competition with independent retailers (Jia, 2008). Along these lines, Holmes (2011) evaluates the expansion of the store network of Walmart, and how it balances economies of density and sales cannibalization. Ellickson et al. (2020) use a flexible demand model of spatial competition with overlapping consumer choice sets to estimate rich substitution patterns for retail stores, by leveraging demographic data. Compared to this strand of the literature, I focus on the contraction of physical presence of large chains, and how firms internalize both cannibalization due to overlapping markets and operating costs of running various locations. In particular, I study dynamic branch-level closure decisions within existing networks, which has an impact on the future profitability of remaining branches nearby.<sup>3</sup>

Second, I contribute to the study of mergers and endogenous market structure through network adjustments. Many studies have highlighted that mergers may affect competition through nonprice channels, such as product availability and other characteristics. Berry and Waldfogel (2001) find that firm consolidation in radio broadcasting reduces station entry and increases programming variety. Fan (2013) and Jeziorski (2014) include applications where mergers generate changes in product characteristics, variety for consumers, and consumer welfare. Other papers related to banking explore the consequences of mergers on depositors. Granja and Paixão (2026) show that uniform pricing across branch networks can materially alter the rate consequences of mergers and the welfare effects of branch divestitures. Benson et al. (2024) find that the geographic proximity of merging branch networks predicts anticompetitive effects that can be missed by conventional HHI-based screening. Closely related to this paper, Pakes et al. (2024) quantify the welfare effects of bank mergers through observed

---

<sup>3</sup>Similar to my setting, Nishida (2015) uses a model of competing convenience store chains to study the trade-off between business stealing and cost savings. However, the paper uses a static equilibrium while my paper focuses on dynamic equilibrium.

changes in deposit rates, branch availability, and bank quality. I contribute to this strand of the literature by showing how branch closures, generated by network adjustments after a merger, produce increased distances between depositors and branches. Consequently, my framework is able to jointly evaluate the cost savings resulting from the endogenous responses to mergers and the loss in consumer welfare obtained from reduced branch availability and longer distances.

Third, this paper relates to research on deposit demand and spatial competition in the banking industry. [Dick \(2008\)](#), [Egan et al. \(2017\)](#), [Wang et al. \(2022\)](#), and [Egan et al. \(2022\)](#) use differentiated product models based on [Berry \(1994\)](#) and [Berry et al. \(1995\)](#) to study deposit competition, consumer welfare, financial fragility, monetary-policy transmission, and the sources of bank value. Other papers add a spatial dimension to this setting. [Ho and Ishii \(2011\)](#) estimate a model of local retail banking competition, while [Kutzbach et al. \(2025\)](#) construct neighborhood level measures of deposit competition with a branch choice model similar to [Ellickson et al. \(2020\)](#). Other papers have focused on studying banking through geographically connected networks, with applications related to understanding the incentives for branch expansion, interactions across local markets, and the geographic allocation of deposits and loans ([Kuehn, 2018](#); [Aguirregabiria et al., 2016, 2025](#); [Ji et al., 2023](#); [Oberfield et al., 2024](#); [Morelli et al., 2025](#)). I contribute to this literature by proposing a framework to study deposit competition across local markets, in which depositors value proximity to branches and where banks consolidate their branch networks to maximize discounted profits in a dynamic equilibrium.

Finally, this paper contributes to the literature on access to financial services. Studies of deposit demand underscore that households choose deposit institutions taking into account various frictions, such as search, switching, awareness, and inattention ([Ho, 2015](#); [Honka et al., 2017](#); [Yankov, 2024](#); [Egan et al., 2025](#)). Other studies focus on the impact of branch availability and closures on local credit, financial inclusion and other outcomes ([Nguyen, 2019](#); [Bonfim et al., 2021](#); [C  lerier and Matray, 2019](#); [Ji et al., 2023](#); [Rysman et al., 2026](#)).<sup>4</sup> More specifically, about branch access, [Sakong and Zentefis \(2025\)](#) use geolocation data to separate access from demand in explaining disparities in branch use. [Jiang et al. \(2026\)](#) study how digital disruption influences banks' branching decisions and results in heterogeneous welfare consequences for depositors.<sup>5</sup> I contribute to this literature by focusing on distance and access in demand, and quantifying the consumer welfare implications of endogenous branch closure responses across locations. In particular, this paper evaluates policy counterfactuals

---

<sup>4</sup>For evidence on the importance of distance for financial relationships in credit markets see [Degryse and Ongena \(2005\)](#), [Alessandrini et al. \(2009\)](#), and [Herpfer et al. \(2023\)](#).

<sup>5</sup>Other studies on digital technologies and their interaction with physical banking networks include [Kim \(2023\)](#), [Keil and Ongena \(2024\)](#), and [Narayanan et al. \(2025\)](#).

related to protecting branches from closures, in a dynamic equilibrium model, allowing for the quantification of welfare implications in terms of depositors and banks.

The remainder of the paper is organized as follows. Section 2 describes the institutional background. Section 3 presents the structural model of spatial deposit demand and dynamic branch closure decisions. Section 4 describes the data and documents key features of the data that the model is able to capture. Section 5 shows the estimation procedures and Section 6 the counterfactuals. Section 7 concludes.

## 2 Institutional Background

### 2.1 Aggregate branch-network contraction

The U.S. retail banking network expanded through the mid-2000s and began to contract after the Great Financial Crisis. According to the FDIC, the end-of-year number of commercial-bank branches fell from 82,564 in 2011 to 69,167 in 2025, a decline of about 16%. Figure 1, panel (a), shows this trend and additionally plots the number of commercial banks over the same period, which fell from 6,275 to 3,806, a decline of about 39%.

The contraction is explained by fewer openings and an increased pace of closures. Figure 1, panel (b) shows that the branch opening rate fell sharply around 2010 and remained around the 1–1.5% range, while the closure rate trended upward and peaked during the COVID-19 pandemic. The branch counts fell in 67.8% of mainland U.S. counties between 2011 and 2025, which indicates that the decline is geographically widespread.<sup>6</sup>

### 2.2 Bank mergers and network adjustment

Contraction also coincides with continuing consolidation in bank ownership. The 20 banks with the largest branch networks in 2024 accounted for almost half of all closures between 2011 and 2024.<sup>7</sup> Merger events create an additional reason to adjust a network because they can generate local duplication when the new owner already operates nearby offices. During 2011–2024, about 10.2% of branches that closed did so within three years after changing owners.

Figure 2 illustrates three patterns that show how large bank network trajectories have behaved in this context. Bank of America contracted gradually, without a large acquisition expansion during the period. In contrast, PNC and Truist experienced discrete increases in

---

<sup>6</sup>Appendix Figure 11 includes a map with the branch count variation at the county level.

<sup>7</sup>Appendix Table 13 reports the corresponding breakdown by network size.




Figure 1: Banking network contraction, 2001–2025

*Notes:* Panel (a) plots the number of FDIC-insured commercial banks and bank branches. Panel (b) plots annual branch opening and closure rates. Openings are offices active in year  $t$  but not  $t - 1$ . Closures are offices present in year  $t - 1$  but not  $t$ . Both rates divide the corresponding flow by the branch stock in  $t - 1$ . Appendix A describes the branch identifier corrections and sample restrictions.

*Source:* FDIC Summary of Deposits and FDIC historical institution data.

their networks following acquisitions, followed by reductions in branch counts. PNC Bank branch network expanded due to a merger with RBC Bank (2012) and BBVA USA (2022), while Truist Bank (formerly BB&T) experienced mergers with Susquehanna and National Penn (both in 2016), and SunTrust (2020). In these examples, closure processes after a merger take between two to three years to materialize.

## 2.3 Household branch use in a digital banking environment

Branch contraction occurred during a period of rapid digital adoption. Between 2013 and 2023, the share reporting mobile banking as their primary access method among banked households that accessed an account in the preceding 12 months rose from 5.5% to 48.3%. Over the same period, the share relying primarily on a bank teller fell from 33.3% to 15.1%.<sup>8</sup>

The primary-access measure nevertheless understates the continued reach of branches, and the substantial heterogeneity among demographics. Figure 3 shows that digital banking most often complements rather than replaces branch use. In 2023, 55.2% of banked households used both teller and digital channels, compared with 28.9% that used only digital channels and 12.7% that used only tellers. Reliance on bank tellers was especially high among groups that are likely to face higher costs of substituting to digital channels, such as older demographic groups. A teller was the primary access method for 30.4% of households aged 65 or older and 4.6% among those aged 25–34. This statistic was also higher outside

<sup>8</sup>Appendix Figure 12 reports the full primary-method series.




Figure 2: Illustrative consolidation patterns among large branch networks

*Notes:* Each panel plots the annual number of active branches owned by the bank from 2011 through 2025. Merger episodes are detailed in vertical gray dotted lines in the year that the acquired branches were incorporated. *Source:* FDIC Summary of Deposits.

metropolitan areas (24.0%), compared to within metropolitan areas (13.8%).

Physical proximity also remains important when households choose where to bank. In the 2022 Survey of Consumer Finances (SCF), 42.1% of families with a checking account cited office location as the most important reason for choosing the institution that held their main checking account. Office location was the most common response in every SCF wave from 2007 through 2022, with only a modest decline over the period.<sup>9</sup>

These facts point to coexistence between digital and physical banking. Digital channels dominate routine account management, but many households continue to use tellers and to value nearby offices.

## 2.4 Branch closure regulation and merger review

At the federal level, an individual branch closure generally requires advance notice but not specific regulatory approval. Section 42 of the Federal Deposit Insurance Act requires an insured depository institution to notify its federal regulator and affected customers at least 90 days before a proposed closure, and to provide reasons and supporting information. However, even if the regulator has the option to consult the community authorities when an interstate bank closes an office in a low or moderate income area, the bank maintains the authority to

<sup>9</sup>Appendix Table 14 reports the full distribution of responses.




Figure 3: Household account access by teller and digital channels in the past 12 months, 2013–2023

*Notes:* "Digital only" combines the survey's mobile-phone and computer/tablet account-access categories. Bars report mutually exclusive and exhaustive channel-use options among banked households that accessed an account in the preceding 12 months. In 2013–2017, the survey measured teller and digital use through a single mark-all-that-apply question. Beginning in 2019, the survey asked separate yes/no questions for each channel. *Source:* FDIC National Survey of Unbanked and Underbanked Households.

close once it satisfies the notice requirements.<sup>10</sup>

More recently, branch closures have received greater attention in pre-merger evaluation. The Bank Merger Act requires approval from regulators for merger transactions and directs authorities to consider competitive effects, financial and managerial resources, and the needs of the communities served. In 2024, the FDIC explicitly made projected branch plan changes part of the review for merger transactions under its jurisdiction.<sup>11</sup> Specifically, the 2024 policy stated that the *"FDIC will evaluate all projected or anticipated branch expansion, closings, or consolidations for the first three years following consummation of the merger"*. The FDIC rescinded the 2024 guidelines in 2025 and reinstated its prior policy.<sup>12</sup>

### 3 Model

This section develops a dynamic model of retail branch networks. In the model, households allocate a proportion of their income to deposits, which are distributed among nearby branches and an outside option that represents other savings alternatives. Meanwhile, banks collect deposits from all branches and use them to generate revenues, but they also need to pay recurring operating costs for each open branch. In every period, banks decide which existing branches to close. A closure decision removes that branch as an alternative for nearby households to allocate their deposits, forcing them to choose among the remaining branches in their consideration set and the outside option. This may benefit the bank if it still has other branches among the remaining options, allowing them to maintain some market share in the area, while removing the fixed operating cost at the same time. On the contrary, closing a branch might not be beneficial if the revenues generated by its deposits are larger than these fixed costs.

In the model, closure decisions are irreversible and there is no branch entry or reopening. As a result, banks consider the impact of their closure decision across the future profitability of their branch networks. These assumptions imply that the model is able to reproduce aggregate branch contraction paths, which is consistent with the behavior observed between 2011 and 2025 in the data. However, individually, a bank can end up with more branches than its initial portfolio due to exogenous mergers with other banks. More importantly, in the model, a bank's closure actions have strategic consequences, since they can alter household choice sets, rival branches' deposits, and the profitability of all surviving branches.

In this model, mergers are exogenous events that change the ownership of certain trans-

---

<sup>10</sup>Federal Deposit Insurance Act, Section 42, 12 U.S.C. §1831r-1. State-level requirements may differ.

<sup>11</sup>FDIC, *Final Statement of Policy on Bank Merger Transactions*, 89 Fed. Reg. 79,125 (Sept. 27, 2024).

<sup>12</sup>FDIC, *Statement of Policy on Bank Merger Transactions*, 90 Fed. Reg. 29,413 (July 3, 2025).

ferred branches. The model treats these events as unanticipated, since banks do not foresee if their network will change due to a merger in the next periods. Different merger pairings can be simulated by altering the assignments of branch locations to specific banks. The model can also be used to infer alternative closure patterns in the absence of mergers, by maintaining the branch ownership to the original banks. A merger entails three main consequences in the model. First, the attractiveness of the transferred branches may change due to the value of the new brand to the depositor. Second, the new owner faces temporary additional costs from managing the newly transferred branches. Third, the merged bank has a larger branch network, and the optimization of its closure plans might change due to the internalization of business stealing between branches.

The demand system includes a time-varying sensitivity to distance. This captures the fact that depositors' sensitivity to distance might change over time which is consistent with changes that may arise from technological developments, such as online banking, and other changes in depositor behavior. Other secular changes that occur during the period are captured by time-varying demographic characteristics and demand shifters that affect the attractiveness of physical branches relative to the outside option.

### 3.1 Environment and timing

Time is discrete and measured in years,  $t = 0, 1, \dots$ . Let  $\mathcal{B}$  denote the set of banks,  $\mathcal{L}$  the set of potential branch locations, and  $\mathcal{C}$  the set of census tracts. Branches, banks, and tracts are indexed by  $\ell$ ,  $b$ , and  $c$ , respectively. The indicator  $o_{\ell t} \in \{0, 1\}$  records whether branch  $\ell$  is open at the beginning of period  $t$ , and  $b_{\ell t}$  identifies its bank owner.

Banks begin the period with exogenously defined branch locations, and manage closure decisions based on their evolving branch portfolio.

The beginning-of-period state is given by the following vectors:

$$s_t = (\mathbf{o}_t, \mathbf{b}_t, \mathbf{x}_t, \mathbf{m}_t, \mathbf{q}_t, \mathbf{z}_t) \in \mathcal{S}, \quad (1)$$

where  $\mathbf{o}_t$  and  $\mathbf{b}_t$  are vectors that collect branch status and ownership for each branch  $\ell$ . The vector  $\mathbf{x}_t$  contains tract demographics for each census tract  $c$ , while  $\mathbf{m}_t$  contains bank deposit spreads at the bank level  $b$ . The vector  $\mathbf{q}_t$  records branch acquisition histories for each branch  $\ell$ . Finally,  $\mathbf{z}_t$  contains aggregate demand and payoff shifters, such as time fixed effects and the time variant disutility of distance. These variables enter the demand and bank payoff functions described in the following subsections. Every period, these state variables evolve following individual transition processes.

Banks choose branch closures but do not choose deposit rates, since bank deposit spreads are predetermined within the period and evolve according to an estimated exogenous process.<sup>13</sup>

Let

$$\mathcal{O}_b(s_t) = \{\ell \in \mathcal{L} : o_{\ell t} = 1, b_{\ell t} = b\} \quad (2)$$

denote the network that bank  $b$  inherits at the beginning of the period. Bank  $b$  chooses a set of branches to close,

$$a_{bt} \in \mathcal{A}_b(s_t) \equiv 2^{\mathcal{O}_b(s_t)}, \quad (3)$$

where  $\ell \in a_{bt}$  means that branch  $\ell$  closes in period  $t$ . In this equation, we can see how the number of feasible closure plans grows exponentially with the number of branches that a bank has.

Let  $\mathbf{a}_t = (a_{bt})_{b \in \mathcal{B}}$  be the joint action comprising all banks' branch-closure decisions. Closure is irreversible: a closed branch cannot reopen. The branches available after the action are

$$\mathcal{O}(s_t, \mathbf{a}_t) = \{\ell : o_{\ell t} = 1\} \setminus \bigcup_{b \in \mathcal{B}} a_{bt}, \quad \mathcal{O}_b(s_t, a_{bt}) = \{\ell \in \mathcal{O}(s_t, \mathbf{a}_t) : b_{\ell t} = b\}. \quad (4)$$

From these equations we can see that  $\mathcal{O}(s_t, \mathbf{a}_t)$  is the resulting total network obtained by removing, from all the open branches, the set of all branches closed by the banks. As such,  $\mathcal{O}_b(s_t, \mathbf{a}_t)$  refers to the resulting network of branches owned by bank  $b$ .

Deposits are determined jointly by the state variables and closure decisions, distributing deposits across the resulting network.

Within a period, events occur in the following order. First, banks observe the state  $s_t$ . Second, banks choose their closure plans  $a_{bt}$  simultaneously and the selected branches close immediately. Third, households allocate deposits across the resulting branch network and the outside option, and banks receive current flow payoffs. Finally, the post-action network carries into the next period, acquisition histories advance, and the exogenous state variables follow their transition processes. This sequence ensures that current revenues generated by managing deposits and all current operating-payoff components are computed on the same post-closure network.

---

<sup>13</sup>This is consistent with a uniform interest rate policy determined at the national level, while the branch network game is decided locally.

### 3.2 Deposit Demand

The deposit demand model below is adapted from [Ellickson et al. \(2020\)](#) for the retail banking industry in a panel setting. Depositors are located in census tracts. Conditional on  $(s_t, \mathbf{a}_t)$ , the choice set for tract  $c$  is

$$\mathcal{J}_{ct}(s_t, \mathbf{a}_t) = \{\ell \in \mathcal{O}(s_t, \mathbf{a}_t) : d_{c\ell} \leq \bar{r}\} \cup \{0\}, \quad (5)$$

where  $d_{c\ell}$  is the distance between the tract centroid and branch  $\ell$ , and  $\bar{r}$  is the choice-set radius, set at  $\bar{r} = 5$  miles. The outside option, which represents the nonlocal deposit choices, is denoted by  $\{0\}$ .

Figure 4 shows a graphical example of the geographical consideration set for two census tracts. Both census tracts in the figure have three branches in their consideration sets. Note that these consideration sets are overlapping, such that the branch in location  $\ell_2$  serves both census tracts.




Figure 4: Overlapping five-mile consideration sets

The indirect utility of household  $i$  located in census tract  $c$  from choosing a branch  $\ell \in \mathcal{J}_{ct}(s_t, \mathbf{a}_t)$  within time  $t$  is given by:

$$\begin{aligned} u_{clti} &= v_{clt} + \epsilon_{clti}, \\ v_{clt} &= \tau_t \tilde{d}_{cl} + \tau_y \tilde{d}_{cl} \tilde{y}_{ct} + \tau_{age} \tilde{d}_{cl} \widetilde{age}_{ct} + \delta_t + \mu_{b_{it}}. \end{aligned} \quad (6)$$

Here  $\tilde{d}_{cl}$  is demeaned log distance,  $\tilde{y}_{ct}$  is demeaned log household income, and  $\widetilde{age}_{ct}$  is demeaned log age. The coefficient  $\tau_t$  governs average distance sensitivity (which varies at the year  $t$  level), while  $\tau_y$  and  $\tau_{age}$  allow it to vary with tract demographics. This term allows the model to capture changes in behavior by depositors, related to their preference for proximity.

The terms  $\delta_t$  and  $\mu_{b_{\ell t}}$  are a year effect and a bank owner effect, respectively.

The outside option represents deposits not assigned to a branch in the constructed choice set, and therefore includes other saving products available at the census tract. Its utility is

$$\begin{aligned} u_{c0t} &= v_{c0t} + \epsilon_{c0ti}, \\ v_{c0t} &= \lambda_h \tilde{h}_{ct} + \lambda_{h^2} \tilde{h}_{ct}^2 + \lambda_{hy} \tilde{h}_{ct} \tilde{y}_{ct}, \end{aligned} \quad (7)$$

where  $\tilde{h}_{ct}$  is demeaned log number of households. Changes in the outside option are represented by shifts in household density, household density squared, and an interaction term with income. These variables serve as a proxy for the availability of alternative savings options for the census tract located in densely populated areas. The shocks  $\epsilon_{cjt}$  are independently distributed type-I extreme value across households and alternatives.

The resulting allocation probabilities are

$$\begin{aligned} p_{c\ell t}(s_t, \mathbf{a}_t) &= \frac{\exp(v_{c\ell t})}{\sum_{j \in \mathcal{J}_{ct}(s_t, \mathbf{a}_t)} \exp(v_{cjt})}, \\ p_{c0t}(s_t, \mathbf{a}_t) &= \frac{\exp(v_{c0t})}{\sum_{j \in \mathcal{J}_{ct}(s_t, \mathbf{a}_t)} \exp(v_{cjt})}. \end{aligned} \quad (8)$$

Since wealth is not observed at the census tract level, I proxy the total resources that can be allocated for deposits by using household income. Let  $Y_{ct}$  be average household income and  $H_{ct}$  the number of households in tract  $c$ . The tract's deposit mass is

$$M_{ct} = \kappa_D Y_{ct} H_{ct}, \quad \kappa_D > 0, \quad (9)$$

and predicted deposits at branch  $\ell$  are

$$D_{\ell t}(s_t, \mathbf{a}_t) = \sum_{c \in \mathcal{C}: \ell \in \mathcal{J}_{ct}(s_t, \mathbf{a}_t)} M_{ct} p_{c\ell t}(s_t, \mathbf{a}_t). \quad (10)$$

We can see in equation (10) that the deposits collected by branch  $\ell$  are obtained from aggregating all the individual allocations from all census tracts that include branch  $\ell$  in their choice set. Equations (5)–(10) are evaluated after the joint closure decision of all banks. As a result, a branch closure can redirect deposits to the other branches of the same bank, to competing banks, or to the outside option in the same period.

The demand model above is intentionally parsimonious and abstracts from switching costs or depositor inertia. This provides tractability and simplifies computational requirements for the dynamic supply model, while still capturing the spatial substitution patterns that

determine the profitability of branches.

### 3.3 Bank Flow Payoffs

Let  $m_{bt}$  denote bank  $b$ 's annual net interest return per dollar of deposits, implied by the interest rate spread between the loan and deposit franchises.<sup>14</sup> Net interest revenue attributable to deposits is

$$R_{bt}(s_t, \mathbf{a}_t) = m_{bt} \sum_{\ell \in \mathcal{O}_b(s_t, \mathbf{a}_t)} D_{\ell t}(s_t, \mathbf{a}_t). \quad (11)$$

Because both  $\mathcal{O}_b(s_t, \mathbf{a}_t)$  and  $D_{\ell t}(s_t, \mathbf{a}_t)$  refer to the operating network, revenue is measured after deposits have reallocated in response to all banks' closures.

The number of branches operated by bank  $b$  after the closure decisions is

$$N_{bt}(s_t, \mathbf{a}_t) = \sum_{\ell \in \mathcal{O}_b(s_t, \mathbf{a}_t)} 1. \quad (12)$$

This term represents the component of profits that depends on the costs of operating the branch network.

Let  $q_{\ell t}^A$  indicate that branch  $\ell$  is both acquired and within the recent-acquisition window. The number of recently acquired branches that bank  $b$  continues to operate is

$$A_{bt}(s_t, \mathbf{a}_t) = \sum_{\ell \in \mathcal{O}_b(s_t, \mathbf{a}_t)} q_{\ell t}^A. \quad (13)$$

This term represents an additional temporary cost of retaining recently acquired branches. Every open branch contributes to  $N_{bt}$ , whereas only recently acquired open branches contribute to  $A_{bt}$ .

Bank  $b$ 's deterministic flow payoff therefore is

$$\pi_b(s_t, \mathbf{a}_t; \boldsymbol{\theta}^\pi) = \theta_R R_{bt}(s_t, \mathbf{a}_t) - \theta_N N_{bt}(s_t, \mathbf{a}_t) - \theta_A A_{bt}(s_t, \mathbf{a}_t), \quad (14)$$

with  $\boldsymbol{\theta}^\pi = (\theta_R, \theta_N, \theta_A)'$ . The three profit function components value deposit revenue, the recurring cost of operating branches, and the additional temporary cost of recently acquired branches. The model does not include a parameter related to the exit cost (or scrap value) for closing a branch, since it would not be separably identified from the fixed operating

---

<sup>14</sup>This treatment is consistent with evidence that U.S. banks set deposit rates with substantial uniformity across their branch networks (Granja and Paixão, 2026). The model therefore focuses on local branch-network adjustment while taking bank-level deposit spreads as predetermined within the annual closure decision. See A.4 for construction of the loan-deposit spread in the data.

cost.<sup>15</sup> In that sense, the model is consistent with a normalization of the exit cost parameter to zero.

### 3.4 Dynamic Branch Closure Problem

Across every period, the state transition of the branch network is governed by branch closure decisions. For a branch owned by  $b$  at the beginning of the period, its open status is kept if the location is not included in the closure action set:

$$o_{\ell,t+1} = o_{\ell t} \mathbf{1}\{\ell \notin a_{bt}\}. \quad (15)$$

The rest of the state space evolves according to different transition rules. Acquisition history dummies advance deterministically for every additional year that the acquired branch is maintained. Demographics, spreads, and aggregate year fixed effects shifters follow specified Markov processes generated from AR(1) regressions. Let  $\mathcal{P}(ds' \mid s, \mathbf{a})$  denote the resulting transition kernel. Note that, since deposits are recomputed from equations (8)–(10) given the state variables, they do not require a separate transition process.

After observing the state  $s$ , bank  $b$  privately observes an action-specific payoff shock  $\nu_b(a_b)$  for every feasible closure plan. Each bank observes its own vector of shocks before acting but not the contemporaneous shocks of its rivals. Conditional on the state, these shocks are independent across banks, periods, and actions and follow a type-I extreme-value distribution. As a result, the functional form for the realized current payoff is

$$u_b(s, a_b, \mathbf{a}_{-b}, \boldsymbol{\nu}_b) = \pi_b(s, a_b, \mathbf{a}_{-b}; \boldsymbol{\theta}^\pi) + \sigma_\nu \nu_b(a_b), \quad (16)$$

where rival banks' joint action profile is defined by  $\mathbf{a}_{-b}$  and  $\sigma_\nu$  scales the shocks into payoff units.

A Markov strategy  $\varphi_b(a_b \mid s)$  gives the probability that bank  $b$  chooses closure plan  $a_b$  when the industry state is  $s$ . Let  $\mathbf{a}_{-b}$  collect the closure plans chosen by all other banks, and let  $\mathcal{A}_{-b}(s)$  denote the set of all possible combinations of those plans. Conditional on the state, the probability of a particular combination of rival actions is

$$\varphi_{-b}(\mathbf{a}_{-b} \mid s) = \prod_{j \neq b} \varphi_j(a_j \mid s).$$

---

<sup>15</sup>See [Kalouptsi et al. \(2021\)](#) for more detail about identification of cost parameters in dynamic discrete choice models.

Bank  $b$ 's action-specific value averages over these possible actions by the other banks and is

$$Q_b^\varphi(s, a_b) = \sum_{\mathbf{a}_{-b} \in \mathcal{A}_{-b}(s)} \varphi_{-b}(\mathbf{a}_{-b} \mid s) \left[ \pi_b(s, a_b, \mathbf{a}_{-b}; \boldsymbol{\theta}^\pi) + \beta \int_{\mathcal{S}} V_b^\varphi(s') \mathcal{P}(\mathrm{d}s' \mid s, a_b, \mathbf{a}_{-b}) \right], \quad (17)$$

where  $\beta \in (0, 1)$  is the annual discount factor. Before bank  $b$  observes its private action shocks, its expected (ex ante) value is

$$V_b^\varphi(s) = \sigma_\nu \left\{ \gamma_E + \log \left[ \sum_{a_b \in \mathcal{A}_b(s)} \exp \left( \frac{Q_b^\varphi(s, a_b)}{\sigma_\nu} \right) \right] \right\}, \quad (18)$$

where  $\gamma_E$  is the Euler constant, which appears because a standard type-I extreme-value shock has mean  $\gamma_E$ . Because this term is common to all feasible actions, it does not affect their relative choice probabilities. The implied conditional choice probability is

$$\varphi_b(a_b \mid s) = \frac{\exp(Q_b^\varphi(s, a_b)/\sigma_\nu)}{\sum_{\tilde{a}_b \in \mathcal{A}_b(s)} \exp(Q_b^\varphi(s, \tilde{a}_b)/\sigma_\nu)}. \quad (19)$$

Under the log-sum expression in equation (18), a bank with more open branches also has more feasible closure sets and therefore receives a larger expected inclusive value from its action shocks, holding deterministic action values fixed. To identify the profit payoff parameters from choice behavior, I normalize  $\sigma_\nu = 1$ .

### 3.5 Markov-Perfect Equilibrium

A Markov perfect equilibrium is a profile  $\boldsymbol{\varphi}^*$  for which equations (17)–(19) hold for every bank and state when  $\boldsymbol{\varphi} = \boldsymbol{\varphi}^*$ . This fixed point accounts for two types of strategic interaction in the model. On one hand, closures reallocate deposits across surviving branches and may affect rivals' current payoffs. On the other hand, because closures are irreversible, they change the future profitability of the branch network and, consequently, also affect the subsequent decisions and payoffs of every bank.

## 4 Data and Empirical Evidence

This section presents the data and documents three empirical patterns that are characterized in the model in Section 3. First, branch deposits are higher when a larger share of nearby household income is located close to the branch. Second, branch exit is associated with

branch size and overlap within the bank’s own network, with a particularly strong relationship following ownership transfers. Third, household-level checking balances are negatively associated with institution distance, but the magnitude of this relationship is smaller in recent years. These findings are consistent with the economic mechanisms captured by the model, and serve as benchmarks for the interpretation of the structural estimation results in Section 5.

## 4.1 Data and sample construction

**Branch network and bank data.** The main dataset used in this paper is the FDIC Summary of Deposits (SOD), which reports the location, ownership, service type, and deposits of each office of an FDIC-insured institution as of June 30 each year. The main panel uses the 2009–2024 period. A limitation is that reported deposits are office-level accounting totals, which are aggregates from unobserved account-level flows from households and firms to branches. For SOD reporting, the FDIC recommends that banks assign deposits to each office according to their internal record-keeping practices, which may reflect the proximity to the account holder’s address, the office where the account was opened, or the branch where it is most active.

The SOD contains a unique and persistent branch identifier, which is maintained even after ownership changes such as merger events.<sup>16</sup> I focus only on branches that report positive deposits. I also remove investment banks whose deposits are concentrated in one or two reporting offices and exclude New York and New Jersey because their unusually dense branch networks generate exceptionally large local choice sets.<sup>17</sup> Additionally, I merge the SOD panel with annual bank balance-sheet information from the FFIEC Call Reports, which include bank-level loan revenue, loan-to-deposit ratio, and deposit expenses. These are used to construct the bank-level spread on deposits that enters the supply model as a proxy of profit margin on deposits.

**Census tract demographics and locations.** I focus on census-tract level data from the American Community Survey (ACS) five-year samples. A census tract is a geographic unit, for which the ACS reports aggregated household-level data, that is larger than a census block but smaller than a zip code or county, which allows me to treat a census tract as a representative household in the area. The ACS provides annual tract-level measures of the number of households, median household income, and median age. I use median household

---

<sup>16</sup>For a small number of cases, I construct persistent branch identifiers, correcting identifier relabeling when possible, and distinguish an office closure from a surviving office that changes owners.

<sup>17</sup>Appendix A records the branch cleaning and sample rules in greater detail.

income as the representative income measure for a tract because it is less sensitive than the mean to the substantial income outliers within small geographic areas. To calculate the distance from the census-tracts to the branches nearby, I assign the location of the Census Bureau’s population center to each tract, which represents the geographic balance point of its residents. These locations are identified by latitude and longitude coordinates that I can use to calculate the geodesic distance from the census-tract to the branches nearby.

**Household surveys.** I use the Survey of Consumer Finances (SCF) for reduced form evidence related to the size of checking account holdings and the distance from a household’s tract centroid to the nearest branch of that institution, across the 2016, 2019 and 2022 waves.

**Branch–tract matching and outcomes.** I link each census tract  $c$  with physical branch locations  $\ell \in \mathcal{L}$ , in year  $t$ . Let  $o_{\ell t} \in \{0, 1\}$  be an indicator equal to one if branch  $\ell$  is open in year  $t$ , and let  $d_{c\ell}$  denote the distance, in miles, from the population center of tract  $c$  to branch  $\ell$ . For each tract–year pair  $(c, t)$ , I define the local branch choice set as

$$\mathcal{J}_{ct} = \{\ell \in \mathcal{L} : o_{\ell t} = 1, d_{c\ell} \leq 5 \text{ miles}\}.$$

Here,  $\mathcal{J}_{ct}$  contains all branches that are open in year  $t$  and located within five miles of tract  $c$ ’s population center. Each element  $\ell \in \mathcal{J}_{ct}$  represents a choice alternative for depositors in tract  $c$  and implies a tract–branch–year observation in the panel.<sup>18</sup> A five-mile radius provides a common definition of local access across markets while allowing the outside option to absorb more distant relationships or alternative savings. In the 2022 SCF, the median share of total household deposits that were allocated to savings institutions within five miles to the household was 67%.

A branch is defined as closed between  $t$  and  $t + 1$  if it is open in  $t$  and its persistent identifier is not observed in  $t + 1$ . A branch that remains open under a new owner is instead identified as an ownership transfer. Most large ownership transfers arise from mergers or acquisitions.<sup>19</sup>

The combined branch-tract panel covers the period 2009–2024 and includes 43 million observations. Since I do not observe individual level information on bank accounts and location of their branches, this dataset allows me to relate the geography of household

---

<sup>18</sup>These choice alternatives are not observed household choices or observed tract-to-branch deposit flows. The structural demand model adds an outside option, indexed by 0, that captures deposits not assigned to a branch in the constructed local set  $\mathcal{J}_{ct}$ .

<sup>19</sup>There are a small proportion of cases of branch sales which do not involve mergers. In the main dataset, direct branch transfers in which the selling bank remains active in the following year represent 15.6% of transfer events and 8.2% of transferred branches in the 2011–2024 period.

income to total branch deposits and capture the behavior of branch exit across the recent years. Appendix Table 12 reports the corresponding distributions separately at the tract–branch-pair, tract, and branch levels.

## 4.2 Branch deposits are spatially determined

Branch deposits are related not only to the amount of income located near an office, but also to its spatial distribution. For each branch, I divide the area within five miles of the branch into five concentric one-mile distance bands, denoted by the 0–1, 1–2, 2–3, 3–4, and 4–5 miles ranges respectively. The band indexed by  $r$  contains census tracts whose population centers are located at least  $r$  miles and less than  $r + 1$  miles from branch  $\ell$ . I define the income-mass proxy for each band as:

$$M_{\ell t}^r = \sum_{c: r \leq d_{c\ell} < r+1} Y_{ct} H_{ct}, \quad r = 0, \dots, 4,$$

where  $Y_{ct}$  is median household income in census tract  $c$ ,  $H_{ct}$  is the number of households, and  $d_{c\ell}$  is the distance in miles from the tract centroid to branch  $\ell$ . Let  $M_{\ell t} = \sum_{r=0}^4 M_{\ell t}^r$  denote total income mass within five miles and let  $s_{\ell t}^r = M_{\ell t}^r / M_{\ell t}$  denote the share located in distance band  $r$ . I estimate

$$\log D_{\ell t} = \beta \log M_{\ell t} + \sum_{r=0}^3 \delta_r s_{\ell t}^r + \gamma_{z(\ell),t} + \mu_{b(\ell),t} + \varepsilon_{\ell t}, \quad (20)$$

where  $D_{\ell t}$  denotes branch deposits,  $\gamma_{z(\ell),t}$  are ZIP-by-year fixed effects, and  $\mu_{b(\ell),t}$  are bank owner-by-year fixed effects. The coefficients related to the bands' shares measure the relationship between the location of income closer to the branch and the 4–5-mile distance band (the omitted outer band category).

Table 1 shows that branch deposits are higher when a greater share of nearby household income is located closer to the branch. Holding total income mass within five miles as fixed, branches for which the 0–1-mile distance band accounts for 10 percentage points more of that mass than the 4–5-mile band have approximately 10.8% higher deposits. This percentage decreases to 6.9% for the 1–2-mile band and 2.4% for the 2–3-mile band, indicating spatial decay when moving to the outer bands. The estimate for the 3–4-mile band is small and not significant.

Table 1: Branch deposits and the spatial distribution of nearby income

|                                  | Income only         | Income + shares     |
|----------------------------------|---------------------|---------------------|
| Log income mass, 0–5 miles       | 0.211***<br>(0.018) | 0.310***<br>(0.020) |
| Share of mass within 0–1 miles   | —                   | 1.024***<br>(0.053) |
| Share of mass within 1–2 miles   | —                   | 0.663***<br>(0.045) |
| Share of mass within 2–3 miles   | —                   | 0.236***<br>(0.042) |
| Share of mass within 3–4 miles   | —                   | 0.020<br>(0.038)    |
| Observations                     | 1,121,418           | 1,121,418           |
| Branches                         | 100,359             | 100,359             |
| ZIP $\times$ year fixed effects  | Yes                 | Yes                 |
| Bank $\times$ year fixed effects | Yes                 | Yes                 |
| Within $R^2$                     | 0.0015              | 0.0072              |

*Notes:* The dependent variable is log branch deposits. Income mass represents tract median household income multiplied by the number of households, summed within five miles of the branch. Share variables are fractions of total income mass within each band. The omitted share is the 4–5 mile band. The sample covers 2009–2024 and excludes branches that correspond to bank headquarters, as well as branches from the top 0.5% of branch deposits within year. Standard errors, in parentheses, are clustered by branch: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

### 4.3 Branch closures and local competition

The model in Section 3 predicts closure behavior as the result of evaluating the value of keeping a branch open on the profit payoff for the bank, which includes deposits intake and costs incurred by that branch. This exercise serves as the empirical counterpart of the closure action in the model. I examine how branch size, local branch competition, demographics, and recent ownership changes are associated with exit. Let  $E_{\ell,t+1}$  equal one if branch  $\ell$  is present in year  $t$  but absent in year  $t + 1$ . The fixed-effect specification is given by:

$$\Pr(E_{\ell,t+1} = 1 \mid X_{\ell t}) = \Phi\left(X'_{\ell t}\theta + \gamma_{z_{(\ell)},t} + \mu_{b_{(\ell)},t}\right), \quad (21)$$

where  $\Phi(\cdot)$  is the standard normal cumulative distribution function. The covariates include log branch deposits, the log number of own-bank and rival branches within five miles, local age and income, and indicators for the first three years after an ownership transfer.<sup>20</sup> Table 2 reports the implied average marginal effects. Column (1) includes separate bank, ZIP, and year fixed effects. Column (2) replaces the additive bank and year effects with bank-by-year effects while retaining ZIP fixed effects. Column (3) additionally replaces ZIP fixed effects

<sup>20</sup>Local demographics are averaged across nearby census tracts using household counts and inverse-distance weights.

with ZIP-by-year effects. Compared to the first two columns, Column (3) retains a smaller share of the original sample (17.9%) due to a lack of variation in the outcome, but shows higher sensitivities because the exit rate is higher in this subsample (15.3%).

Table 2: Average marginal effects for one-year branch exit

|                                     | (1)                | (2)                | (3)                |
|-------------------------------------|--------------------|--------------------|--------------------|
| Log deposits                        | -2.18***<br>(0.03) | -2.80***<br>(0.04) | -7.71***<br>(0.13) |
| Log(1 + rival branches)             | 0.22***<br>(0.05)  | 0.25***<br>(0.08)  | -0.64***<br>(0.23) |
| Log(1 + own branches)               | 1.23***<br>(0.04)  | 1.76***<br>(0.06)  | 4.89***<br>(0.17)  |
| Log weighted age                    | 1.59***<br>(0.53)  | 2.22***<br>(0.67)  | 11.73***<br>(2.62) |
| Log weighted income                 | -0.02<br>(0.24)    | -0.15<br>(0.31)    | -2.99***<br>(1.14) |
| Ownership transfer: year 1          | 3.14***<br>(0.16)  | 3.34***<br>(0.24)  | 7.59***<br>(0.69)  |
| Ownership transfer: year 2          | 1.73***<br>(0.15)  | 2.12***<br>(0.22)  | 5.48***<br>(0.72)  |
| Ownership transfer: year 3          | 1.01***<br>(0.15)  | 1.90***<br>(0.24)  | 5.57***<br>(0.81)  |
| Bank fixed effects                  | Yes                | No                 | No                 |
| ZIP fixed effects                   | Yes                | Yes                | No                 |
| Year fixed effects                  | Yes                | No                 | No                 |
| ZIP $\times$ year fixed effects     | No                 | No                 | Yes                |
| Bank $\times$ year fixed effects    | No                 | Yes                | Yes                |
| Total branch-years                  | 1,158,005          | 1,158,005          | 1,158,005          |
| Estimation branch-years             | 912,322            | 693,655            | 206,976            |
| Retained share (percent)            | 78.8               | 59.9               | 17.9               |
| Total sample exit rate (percent)    | 2.99               | 2.99               | 2.99               |
| Retained sample exit rate (percent) | 3.77               | 4.87               | 15.28              |

*Notes:* The first line in each cell reports the average marginal effect in percentage points. For continuous regressors, such as income or age, the marginal effects are average derivatives, while discrete regressors such as ownership transfer effects are average changes from zero to one. All columns use the same sample for years 2009–2023. Appendix Table 15 reports the underlying probit coefficients. Branch-clustered standard errors appear in parentheses: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Across the three specifications, branches with more deposits are less likely to exit, which is consistent with banks maintaining branches that possess larger deposit intake. The presence of rival branches is positively related to exit in columns (1) and (2), but turns negative in column (3). In contrast, the number of own bank branches nearby is positively related to exit across the three columns, indicating efforts from the bank to reduce local cannibalization. Interestingly, local age of households nearby is also positively associated with exit. Even though survey data indicates that older populations value branch proximity the most, banks seem more likely to close branches that are near this demographic group. Household income is statistically significant only in column (3), where it indicates that if the households nearby

are richer, the exit probability is lower. This effect is also related to the potential deposit intake. Finally, across all columns, the merger dummies show positive marginal effects, indicating the higher probability of branch closure after the branch has been transferred to a new bank.

**Branch exit and new owner network.** To inspect more closely the merger effects on branch closure, I focus on whether a transferred branch is more likely to exit when the acquiring bank already operated another office within five miles before the transfer.

Let  $e$  index an ownership-transfer event occurring in year  $m$ , and let  $\ell$  index the branches transferred in that event. For each post-transfer horizon  $h \in \{1, 2, 3\}$ , I estimate the linear probability model

$$\text{Exit}_{\ell e}^{(h)} = \beta_h \text{Overlap}_{\ell e} + X_{\ell e}^{\text{pre}} \gamma_h + \alpha_e + \lambda_{z(\ell), m} + \varepsilon_{\ell e}^{(h)}, \quad (22)$$

where  $\text{Exit}_{\ell e}^{(h)}$  equals one if branch  $\ell$  exits within  $h$  years after the transfer. The variable  $\text{Overlap}_{\ell e}$  is equal to one if the acquiring bank operated another branch within five miles of  $\ell$  in the previous year ( $m - 1$ ). This implies that all overlap measures are calculated before the ownership change was introduced.

The vector  $X_{\ell e}^{\text{pre}}$  includes pre-transfer branch deposits, the log numbers of nearby branches belonging to the acquired bank and to rival banks, and local age and income. The transfer event fixed effect  $\alpha_e$  compares branches transferred between the same acquired and acquiring banks in the same year. The term  $\lambda_{z(\ell), m}$  is a ZIP-by-transfer-year fixed effect that absorbs local conditions common to branches in the same ZIP and year.

Table 3 shows that post-transfer branch exits are more likely in locations where the new owner already operated a nearby branch. The presence of another branch already operated by the new owner within five miles is associated with a 7.7 percentage point higher probability of exit within one year and a 12.9 percentage point higher probability within three years. This pattern is consistent with network rationalization following merger events.

## 4.4 Checking balances and institution distance

I complement the branch-level evidence with household-level data from the 2016, 2019, and 2022 Survey of Consumer Finances (SCF). For these waves, the SCF reports a constructed distance from the household’s census-tract centroid to the nearest branch of the institution associated with each checking account. This allows me to analyze the relationship between

Table 3: Pre-transfer overlap with the new owner’s network and subsequent exit

|                                          | Exit within     |                    |                    |
|------------------------------------------|-----------------|--------------------|--------------------|
|                                          | 1 year          | 2 years            | 3 years            |
| Existing new-owner branch within 5 miles | 7.73*<br>(4.43) | 11.05***<br>(4.08) | 12.90***<br>(4.45) |
| Pre-transfer controls                    | Yes             | Yes                | Yes                |
| Transfer-event fixed effects             | Yes             | Yes                | Yes                |
| ZIP $\times$ transfer-year fixed effects | Yes             | Yes                | Yes                |
| Mean dependent exit rate                 | 6.2             | 10.6               | 14.4               |
| Observations                             | 21,424          | 21,424             | 21,424             |
| Transfer events                          | 2,933           | 2,933              | 2,933              |
| Identifying events                       | 329             | 329                | 329                |

*Notes:* Coefficients and standard errors, reported in parentheses, are in percentage points. The sample covers ownership transfers from 2010 to 2021 and uses the first observed ownership change for each branch. The outcomes are indicators for exit within one, two, or three years. Branches that are transferred again or disappear and subsequently reappear during the three-year horizon are excluded. Overlap is measured before the transfer period and equals one when the acquiring bank operated another branch within five miles of the transferred branch in year  $m - 1$ . Pre-transfer controls include log branch deposits, acquired bank network density, rival density, and local age and income. Standard errors are clustered by ownership-transfer event:\*  $p < 0.10$ , \*\*  $p < 0.05$ , and \*\*\*  $p < 0.01$ .

checking balances and distance to a banking institution.<sup>21</sup>

Let  $i$  index households,  $j$  index an account or household–institution relationship, and  $d_{ijt}$  denote the distance to the corresponding institution. I estimate

$$\log B_{ijt} = \sum_{s \in \{2016, 2019, 2022\}} \beta_s \mathbf{1}\{t = s\} \log d_{ijt} + X'_{ijt} \gamma + \alpha_t + \varepsilon_{ijt}, \quad (23)$$

where  $B_{ijt}$  is the checking balance. The coefficient  $\beta_s$  is the elasticity of checking balances with respect to institution distance in SCF wave  $s$ . For example, an elasticity of  $-0.03$  implies that doubling the distance to the institution is associated with approximately a 2% lower checking balance.

I report three specifications that exploit the fact that households may hold more than one checking account, potentially across multiple institutions. The main-account specification uses the household’s primary checking account and includes demographic, account, and relationship controls. The all-accounts specification expands the data to the account level and

<sup>21</sup>I restrict the sample to positive checking balances with valid institution links and distance information. For observations reported as less than one-half mile, distance is assigned a value of 0.25 miles. Estimates use the SCF final family weights and combine uncertainty across the five implicates and 999 replicate weights.

includes secondary checking accounts. The household–institution specification aggregates all of a household’s checking balances held at the same institution when multiple members of a household have accounts at the same bank. The last two specifications include household-by-wave fixed effects, which reduces the sample to households with more than 1 checking account (50% of total households in 2022).

Table 4: Elasticity of checking balances with respect to institution distance, 2016–2022

|                                       | Main account only    | All accounts          | Consolidated accounts per institution |
|---------------------------------------|----------------------|-----------------------|---------------------------------------|
| 2016                                  | -0.026***<br>(0.009) | -0.053***<br>(0.013)  | -0.030**<br>(0.013)                   |
| 2019                                  | -0.005<br>(0.011)    | -0.065***<br>(0.015)  | -0.019<br>(0.015)                     |
| 2022                                  | -0.010<br>(0.014)    | -0.028<br>(0.019)     | -0.008<br>(0.017)                     |
| Unit of observation                   | Household            | Checking-account slot | Household–institution                 |
| Wave fixed effects                    | Yes                  | No                    | No                                    |
| Household $\times$ wave fixed effects | No                   | Yes                   | Yes                                   |
| Main-institution indicators           | No                   | No                    | Yes                                   |
| Additional covariates                 | Yes                  | No                    | No                                    |
| Observations in 2016                  | 5,684                | 7,616                 | 4,095                                 |
| Observations in 2019                  | 5,356                | 7,411                 | 4,001                                 |
| Observations in 2022                  | 4,236                | 6,510                 | 3,588                                 |
| Total estimation observations         | 15,276               | 21,537                | 11,684                                |

*Notes:* Entries report coefficients on year-specific  $\log(\text{distance})$  terms from weighted log OLS regressions estimated on positive checking balances. The dependent variable is the log positive checking balance for the main account, checking-account slot, or household–institution relationship, respectively. The main-account specification includes age, age squared, the inverse hyperbolic sine of household income, the number of checking accounts, the number of financial institutions, an indicator for direct deposit, and log relationship length, along with wave fixed effects. The all-accounts specification expands checking-account slots 1–6 and includes household-by-wave fixed effects. The household–institution specification aggregates checking balances across slots 1–6 within institution, includes household-by-wave fixed effects, and adds wave-specific indicators for whether the institution holds the household’s main checking account. Standard errors, in parentheses. Significance: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table 4 shows that checking balances are negatively associated with institution distance in all three 2016 specifications. This indicates that longer distances to chosen deposit institutions were negatively associated with checking balances. By 2022, the magnitudes of these coefficients are lower than those of 2016 for all specifications, and none of their values are statistically significant. The 2019 estimates show mixed results and only remain significant in the “all accounts” column for 2019.

The implied economic magnitudes are modest, but they indicate lower elasticities for recent years. In the main-account results, moving from one to five miles is associated with approximately 4.2% lower balances in 2016, compared with 1.5% in 2022. For the household–institution specification these magnitudes are respectively 4.7% and 1.3%.

## 5 Estimation

Estimation is implemented sequentially in two linked stages. The demand stage recovers the parameters governing the allocation of tract-level deposit mass across nearby branches and the outside option. The supply stage takes the estimated demand system as given and then estimates the structural payoff parameters that are consistent with banks' closure decisions.

For estimation, I define the strategic banks as the 20 banks with the largest branch networks in 2024. These represent almost 50% of all closures that occurred between 2011 and 2024. This allows the estimated model to remain tractable, while focusing on the agents with larger network decisions. The rest of banks are considered nonstrategic, so their branches remain part of household choice sets and enter the competitive environment, but their closure behavior is exogenous.

### 5.1 Demand Estimation

#### 5.1.1 Estimator and normalizations

The demand estimator implements the deposit allocation model in Section 3.2. For each census-tract  $c$  and year  $t$ , the choice set contains every open branch within five miles of the tract's center of population and an outside option. Given a set of parameter values, the model produces individual choice probabilities to these alternatives and aggregates the resulting tract-level allocations to obtain predicted deposits for each branch within a year. The parameters are estimated by matching these predicted totals to the branch deposits reported in the Summary of Deposits.

Let  $\theta^D$  collect

$$\theta^D = (\{\tau_t\}_{t=2009}^{2024}, \tau_y, \tau_{\text{age}}, \lambda_h, \lambda_{h^2}, \lambda_{hy}, \kappa_D, \{\delta_t\}_{t=2010}^{2024}, \{\mu_b\}_{b \in \mathcal{B}_{20}}).$$

The utility and deposit-mass equations are given by equations (6)–(10). Distance, income, age, and household counts enter in logs and are demeaned using the estimation sample. I normalize the year effect to zero,  $\delta_{2009} = 0$ . Owner effects are estimated for the strategic banks, denoted by  $\mathcal{B}_{20}$ . Branches belonging to all other banks remain in the relevant choice sets, but their common owner effect is normalized to zero.

The scale parameter  $\kappa_D$  converts the tract income-mass proxy,  $Y_{ct}H_{ct}$ , into potential deposit mass allocated across nearby branches and the outside option. For a candidate

parameter vector, predicted deposits are

$$\hat{D}_{\ell t}(\boldsymbol{\theta}^D) = \sum_{c:\ell \in \mathcal{J}_{ct}} \kappa_D Y_{ct} H_{ct} p_{c\ell t}(\boldsymbol{\theta}^D). \quad (24)$$

Observed deposits and the product  $Y_{ct}H_{ct}$  are expressed in millions of dollars. The parameter vector is estimated by nonlinear least squares using one residual per branch-year:

$$\hat{\boldsymbol{\theta}}^D = \arg \min_{\boldsymbol{\theta}^D} \sum_{\ell, t} \left[ \log D_{\ell t}^{\text{obs}} - \log \hat{D}_{\ell t}(\boldsymbol{\theta}^D) \right]^2. \quad (25)$$

### 5.1.2 Identification

This subsection describes the variation in the data and additional assumptions needed to identify the parameters in the demand model. Since deposit flows from households to branches are not observed, identification instead comes from how observed branch deposits vary with the spatial distribution of potential deposit mass from census tracts nearby, and the competing alternatives around each branch.

Following [Ellickson et al. \(2020\)](#), I assume that  $\epsilon_{cjt_i}$  and the measurement error of the nonlinear least squares equation are independent of the location of branches, as well as depositors' residential locations, incomes and other characteristics. More specifically, depositors take branch locations as given, and their perceptions about branch quality and interest rate pricing are formulated at the bank chain level, not at the branch. As a result, I control for the endogeneity of chain policies using the bank owner fixed effects. Evidence in the literature, such as [Granja and Paixão \(2026\)](#) favors the idea that banks set uniform pricing policies for deposit rates.

The  $\kappa_D$  parameter represents the fraction of overall income allocated to potential deposits by households. This parameter is identified from the variation of the number of branches across similar markets, and observing the change in total deposits across branches. For example, two branches exposed to similar total income within five miles may face different distances to that income and different sets of competing branches. This ratio is identified separately from the outside option in the model because changes in the choice set may lead depositors to substitute deposits from the outside option to the observed branches, while the  $1 - \kappa_D$  proportion is not substitutable with branch deposits.

The rest of the parameters affect the utility of the branch choice. Identification comes from the cross-sectional and time-series variation provided by the tract-branch distance and demographic characteristics which govern the resulting shares of deposits that households within a branch exposure set have. Therefore, varying each of these observable variables

alters the share of deposits that would be assigned to that branch, relative to others in the choice set.

It is important to note that, even though I interpret changes in the  $\tau_t$  parameter as absorbing changes in the disutility of distance across time, these may be conflated with unobserved bank-level changes in bookkeeping policies for deposit assignment across branches. On one hand, these changes may capture the impact of digitalization and online banking. On the other hand, other types of bank level policies may generate deposit reassignments that may be independent from consumer preferences.

### 5.1.3 Demand estimates and economic magnitudes

Table 5 reports the branch utility, outside-option, and market-size parameters. The distance coefficient is negative in every year, indicating that deposits are less likely to be allocated to more distant branches. Figure 5 shows that the annual point estimates become gradually less negative through 2021 before returning toward their earlier magnitude.<sup>22</sup>

Table 5: Demand Parameter Estimates

| Coefficient                       | Variable                                | Estimate   | Std. error |
|-----------------------------------|-----------------------------------------|------------|------------|
| $\bar{\tau}$                      | Distance, average 2009–2024             | -0.367***  | (0.026)    |
| $\tau_y$                          | Distance $\times$ income                | 0.042      | (0.062)    |
| $\tau_{age}$                      | Distance $\times$ age                   | -0.551***  | (0.137)    |
| $\lambda_h$                       | Outside option: density                 | 5.109***   | (0.159)    |
| $\lambda_{h^2}$                   | Outside option: density squared         | -2.777***  | (0.130)    |
| $\lambda_{hy}$                    | Outside option: density $\times$ income | 6.380***   | (0.307)    |
| $\kappa_D$                        | Market-size scaling                     | 0.764***   | (0.019)    |
| Year fixed effects                |                                         | Yes        |            |
| Top-20 bank fixed effects         |                                         | Yes        |            |
| Estimation window                 |                                         | 2009–2024  |            |
| Candidate tract-branch-year pairs |                                         | 43,720,184 |            |
| Branch-year NLS observations      |                                         | 1,224,332  |            |

*Notes:* The table reports the baseline demand estimates that exclude branch deposit rates. The dependent variable in the nonlinear least-squares objective is branch-year log deposits. Distance, income, age, and household count enter in logs and are demeaned across the candidate-pair estimation sample. Density is shorthand for demeaned log households. The first row reports  $\bar{\tau} = 16^{-1} \sum_{t=2009}^{2024} \tau_t$ . Its standard error is  $\{w' \widehat{\text{Var}}(\hat{\tau}) w\}^{1/2}$ , where  $w$  places equal weight on the 16 annual coefficients. Standard errors use the selected covariance matrix, two-way clustered by branch and bank certificate. Year and Top-20 bank fixed effects are included but omitted from the parameter panel.

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

<sup>22</sup>This time-series variation is imprecisely estimated. A joint Wald test does not reject equality of the 16 annual distance coefficients ( $p = 0.125$ ).




Figure 5: Time-varying demand estimates

*Notes:* Panel (a) plots the annual distance coefficients  $\tau_t$ . Panel (b) plots the common inside-option year effects  $\delta_t$ , with  $\delta_{2009}$  normalized to zero. Vertical bars show 95% confidence intervals based on the two-way clustered covariance matrix.

The estimates also reflect the empirical heterogeneity across demographic preferences for branches. The interaction between distance and income is small and statistically imprecise, with an estimate of 0.042. By contrast, the distance-age interaction is  $-0.551$  and statistically significant, implying that tracts with older residents are more sensitive to branch distance.

The outside option parameters show that household density around the census tract has a positive linear coefficient and a negative quadratic coefficient. This concavity indicates that the larger the population density around the branch is, the preference for the outside option becomes less positive. Also, the income term shows that if the population nearby has a higher income, the preference for the outside option is higher. These estimates are consistent with stronger outside option for branches in richer and more urban areas, where more alternatives for savings are available. The estimated scale parameter,  $\hat{\kappa}_D = 0.764$ , represents the share of tract income mass that goes into potential deposit mass.

Table 6: Demand heterogeneity and branch-deposit fit

| Quartile                                | Mean income<br>or age | Effective<br>distance slope | Aggregate<br>distance elasticity |
|-----------------------------------------|-----------------------|-----------------------------|----------------------------------|
| <i>Panel A: Income quartiles</i>        |                       |                             |                                  |
| Q1 (lowest)                             | \$31,814              | -0.375                      | -0.332                           |
| Q2                                      | \$49,576              | -0.390                      | -0.344                           |
| Q3                                      | \$67,956              | -0.387                      | -0.351                           |
| Q4 (highest)                            | \$112,516             | -0.385                      | -0.365                           |
| <i>Panel B: Age quartiles</i>           |                       |                             |                                  |
| Q1 (youngest)                           | 29.5                  | -0.257                      | -0.244                           |
| Q2                                      | 36.1                  | -0.344                      | -0.322                           |
| Q3                                      | 40.9                  | -0.409                      | -0.373                           |
| Q4 (oldest)                             | 48.8                  | -0.495                      | -0.443                           |
| <i>Panel C: Branch-year deposit fit</i> |                       |                             |                                  |
| Branch-year NLS observations            |                       | 1,224,332                   |                                  |
| $R^2$ , log deposits                    |                       | 0.109                       |                                  |
| RMSE, log deposits                      |                       | 1.199                       |                                  |
| Median absolute log residual            |                       | 0.636                       |                                  |

*Notes:* Income and age quartiles are formed from their pooled, unweighted tract-year distributions over 2009–2024. The second column reports the unweighted mean characteristic within each quartile. Effective distance slopes are defined in equation (26). Aggregate distance elasticities average the own-distance allocation elasticities in equation (27) across the branches available to a census tract and then aggregate across tract-years in the quartile using predicted deposits. The slope and elasticity columns are predicted-deposit-weighted means. The pseudo- $R^2$  is one minus the ratio of squared log-deposit residuals to the total sum of squares of observed log deposits.

The distance parameter slope can be summarized in the following equation:

$$\tau_{ct}^d = \tau_t + \tau_y \tilde{y}_{ct} + \tau_{age} \tilde{age}_{ct}. \quad (26)$$

For a branch, the corresponding tract-level own-distance elasticity of its allocation probability is

$$\frac{\partial \log p_{c\ell t}}{\partial \log d_{c\ell}} = \tau_{ct}^d (1 - p_{c\ell t}). \quad (27)$$

Branch-level deposit elasticities aggregate this expression across the tracts contributing to a branch.

Table 6 translates the distance coefficients into effective slopes and elasticities. Across the unconditional income quartiles, effective distance slopes are similar, ranging from  $-0.390$  to  $-0.375$ . The changes across the age distribution are substantially stronger, indicating that deposit allocation in older neighborhoods is more sensitive to branch distance. The aggregate distance elasticity increases in magnitude from  $-0.244$  for tracts in the youngest age quartile

to  $-0.443$  for those in the oldest quartile.

Table 6, Panel C, reports branch-level fit. The model explains 10.9% of the variation in log reported branch deposits. This modest fit is expected given the model’s deliberately spatial interpretation of SOD deposits. In the data, reported deposit balances may also reflect non-local accounts, institutional customers, and within-bank allocation conventions. In contrast, the role of the demand estimation is to capture the spatial component of deposits that is explained by local household demand.

## 5.2 Supply Estimation

### 5.2.1 Overview of the supply estimator

I estimate the supply parameters in two steps. In the first step, I estimate reduced-form policy functions to link the binary branch closure decision with observable state variables. These policies are used to approximate the equilibrium Conditional Choice Probabilities (CCP) in the model, without solving the dynamic game during estimation. Additionally, I estimate the transition processes governing the state variables.

In the second step, I use the reduced-form CCPs to forward simulate the closure actions for all the branches under two different initial conditions: a path in which a branch is closed and another in which that branch is open. After simulating for 20 periods, I calculate discounted change in the branch owner’s revenue, operating costs and recent-acquisition exposure payoffs caused by the current closure decision. I repeat this process for all the branches within each year that are observed in the data, since the option to close is possible whenever a branch is observed. I then estimate the structural payoff parameters using the discounted change terms to rationalize the actual closure decision in the data.

Methodologically, the procedure combines forward policy evaluation in the spirit of [Bajari et al. \(2007\)](#) with a CCP-based second-stage likelihood as in [Pakes et al. \(2007\)](#) and [Aguirregabiria and Mira \(2007\)](#). Unlike a standard CCP inversion, I obtain the relevant value differences directly by forward simulation from observed states.

It is important to note that the full bank-level policy in Section 3 assigns probabilities to complete closure plans, not individual binary closure actions. Estimating this object nationally is unfeasible because the number of possible plans grows exponentially with the size of a bank’s network. I therefore use branch-level closure probabilities to approximate the model’s continuation behavior in the forward simulations. Additionally, although the forward simulation focuses on a specific focal closure at a time, its simulated payoff consequences are calculated for the bank considering its entire network, bringing the estimation closer to the

model’s bank-level decision problem.

### 5.2.2 First-stage policies and state transitions

Let  $e_{\ell t} = 1$  if branch  $\ell$  is open in year  $t$  but absent in  $t + 1$ , and let  $e_{\ell t} = 0$  otherwise. If a branch survives under a new owner, I consider it an ownership transfer. For branches owned by the strategic banks, the estimated closure policy is

$$\hat{p}_{\ell t} \equiv \mathbb{P}(e_{\ell t} = 1 \mid \mathbf{w}_{\ell t}) = \Lambda(\mathbf{w}_{\ell t}' \boldsymbol{\gamma}^P + \eta_{z(\ell)}^P + \eta_{b_{\ell t}, t}^P), \quad (28)$$

where  $\Lambda(\cdot)$  is the logistic cumulative distribution function,  $\mathbf{w}_{\ell t}$  contains reported branch deposits, nearby same-bank and rival branch counts, local age, and indicators for the first three years following an ownership transfer. The specification includes ZIP and bank-by-year fixed effects. The fitted probabilities determine subsequent closure behavior along the forward-simulated paths.

During the forward simulation, the reported deposits used in the reduced-form policy function are replaced by recomputed branch deposits from the estimated demand system under the simulated network. This ensures that the estimated demand system is the source of endogenous variation in branch deposits across the simulated network paths. In particular, this assumption allows the closure actions to redistribute deposits across the remaining network, further reinforcing deposit levels of the surviving neighbor branches in the next periods.

The first-stage panel covers 2009–2024. Since a branch closure in year  $t$  can only be identified if its status in  $t + 1$  is non-open, closure outcomes are available only until 2023.

I estimate a separate reduced-form policy for non-Top-20 bank branches. This auxiliary policy allows non-strategic competitors to evolve in the national forward simulations used for estimation.<sup>23</sup> Beyond the observed sample, the strategic bank-year closure intercepts are extrapolated using fitted bank-specific AR(1) processes. The simulations also use the estimated demand system, tract demographic processes, and bank-year deposit-spread processes. Appendix C reports the first-stage estimates and transition equations.

The strategic-bank policy sample contains 447,892 branch-year observations, of which 369,597 satisfy the fixed-effect support requirements. Figure 6 compares observed closure rates with predicted probabilities obtained from the reduced-form logit policy function. The estimated policy closely follows the closure rates by year and post acquisition history patterns. Additionally, it captures the strong decline in closure probabilities across the deposit

---

<sup>23</sup>The county counterfactual games instead hold nonstrategic branches fixed open.

distribution for branches. However, fit is weaker in several intermediate deposit deciles.




Figure 6: First-stage closure-policy calibration

*Notes:* The figure uses the 369,597 Top-20 branch-year observations retained by the first-stage logit after fixed effect support restrictions. Each panel compares the observed branch closure rate with the mean predicted probability obtained from the fixed effects sample. Panel (a) groups observations by year. Panel (b) uses pooled deciles of branch deposits. Panel (c) classifies a branch as not recently acquired or as being in the first, second, or third year after an ownership transfer. Vertical bars show 95% confidence intervals for the observed closure rates.

### 5.2.3 Forward-simulated payoff differences

The structural sample contains strategic banks' branch-year decision nodes  $i = (\ell, b, t)$  from 2011 through 2023, where  $b$  denotes the owner of branch  $\ell$  at the beginning of  $t$ .

For each decision node, the simulation begins from the observed national branch and tract state. I generate one path in which the branch corresponding to the decision node is forced to remain open in period  $t$  and another in which it is forced to close (independently of the observed outcome). This intervention applies only in the specific period where the decision node is observed. In the open path, the intervened branch may close from period  $t + 1$  onward by following the estimated closure policy, which depends on the evolution of the state variables and a random draw. In the close path, the intervened branch remains closed because exit is irreversible. All other branches follow their estimated closure policies, which govern if these close or remain open in the next periods. By comparing the two paths we can measure the effect of the current branch decision on the entire discounted payoff of its owner, including deposit reallocation across the bank's surviving offices and subsequent changes in the branch network.

For payoff component  $k \in \{R, N, A\}$  and focal branch action  $a \in \{0, 1\}$ , define

$$F_i^k(a) = \widehat{\mathbb{E}} \left[ \sum_{h=0}^{T_{\text{sim}}-1} \beta^h f_b^k(s_{t+h}, a_{t+h}) \middle| s_t = s_i, e_{\ell t} = a, \widehat{\mathbf{p}}, \widehat{\mathcal{P}} \right], \quad a \in \{0, 1\}, \quad (29)$$

where  $\widehat{\mathbf{p}}$  denotes the estimated branch policies,  $\widehat{\mathcal{P}}$  the estimated transition processes and  $a = 1$  denotes closing the focal branch. I set  $T_{\text{sim}} = 20$  and  $\beta = 0.95$  for the simulation.

The flow primitives are signed to match the payoff equation:

$$f_b^R = R_b, \quad f_b^N = -N_b, \quad f_b^A = -A_b, \quad (30)$$

The structural regressor is the close-minus-open value difference of each component

$$\Delta F_i^k = F_i^k(1) - F_i^k(0). \quad (31)$$

The two paths use common random numbers for the exogenous transition processes so that  $\Delta F_i^k$  isolates the effect of the focal intervention action. This process allows us to simulate the otherwise unobserved changes in the discounted streams of the profit function components, which, according to the model, the bank evaluates before implementing a closure action.

The revenue difference  $\Delta F_i^R$  captures the revenue lost at the branch net of deposits redirected elsewhere in the bank's network. The operating cost difference  $\Delta F_i^N$  measures the discounted number of open branches eliminated by the closure. The acquisition term  $\Delta F_i^A$  similarly measures the reduction in discounted exposure to recently acquired branches. The forward simulation also includes the continuation correction implied by the private shock of the model.<sup>24</sup> For every decision node, each expectation is approximated using 50 paired simulations with different random draws per simulation for the exogenous processes.

#### 5.2.4 Second-stage estimation and identification

The second stage relates observed branch closures to the simulated close-minus-open value differences. For branch-year observation  $i$ , under the branch-level binary logit specification,

$$\mathbb{P}(e_i = 1 \mid \Delta \mathbf{F}_i, \chi_{b(i),t}, \chi_{z(i)}) = \Lambda \left( \Delta F_i^\nu + \theta_R \Delta F_i^R + \theta_N \Delta F_i^N + \theta_A \Delta F_i^A + \chi_{b(i),t} + \chi_{z(i)} \right). \quad (32)$$

The CCP continuation offset  $\Delta F_i^\nu$  enters with coefficient normalized to one. The remaining coefficients correspond directly to the three payoff parameters in equation 14. Similarly

---

<sup>24</sup>The construction of this term is specified in Appendix C.4.

to the first-stage reduced form policy for strategic banks, the specification includes bank-by-year and ZIP fixed effects.

Identification of the supply parameters depends on branch differences generated by how closure varies with the bank’s discounted profit payoff component, conditional on the ZIP code and bank-by-year fixed effects. Revenue consequences vary with the removed branch deposits, bank level spreads on deposits, and the amount of demand that is redirected to surviving offices. The operating cost component varies with the discounted number of branches across subsequent years that are eliminated by the closure. Acquisition exposure varies with the timing of exogenous ownership transfers (which occur as observed in the data). Bank-by-year fixed effects absorb common changes in a bank’s closure propensity over time, and ZIP code fixed effects absorb persistent local differences.

The sample used for the forward simulation contains 391,600 branch-year observations (decision nodes). The fixed-effect logit uses 317,403 observations after removing decisions without the within-group support required by the fixed effects. The observed closure rate is 3.54% in the full sample and 4.33% in the retained estimation sample.

### 5.2.5 Supply estimates and economic fit

Table 7: Supply payoff parameter estimates

| Coefficient                       | Discounted close-minus-open component                           | Estimate  | Std. error |
|-----------------------------------|-----------------------------------------------------------------|-----------|------------|
| 1 (fixed)                         | Future private-shock value, $\Delta F^\nu$                      | 1.000     | —          |
| $\theta_R$                        | Deposit revenue, $\Delta F^R$                                   | 39.963*** | (1.945)    |
| $\theta_N$                        | Fixed open-branch cost, $\Delta F^N = \Delta(-N)$               | 0.159***  | (0.008)    |
| $\theta_A$                        | Negative recent-acquisition exposure, $\Delta F^A = \Delta(-A)$ | 0.214***  | (0.038)    |
| Bank-by-year fixed effects        |                                                                 | Yes       |            |
| ZIP-code fixed effects            |                                                                 | Yes       |            |
| Common closure adjustment cost    |                                                                 | Excluded  |            |
| Fixed open-branch cost            |                                                                 | Included  |            |
| Full-sample branch-year decisions |                                                                 | 391,600   |            |
| Fixed-effect-logit observations   |                                                                 | 317,403   |            |
| Full-sample closure rate          |                                                                 | 3.54%     |            |
| Estimation-sample closure rate    |                                                                 | 4.33%     |            |

*Notes:* The dependent variable equals one when the physical branch office disappears before the next year and zero otherwise. An office that survives under a new owner is not coded as closed. The deterministic regressors are simulated close-minus-open differences in 20-period discounted payoff components. Revenue is the bank's deposit spread multiplied by deposits measured in billions of dollars, so  $\Delta F^R$  is measured in discounted spread-times-billion-dollar years. The fixed open-branch cost regressor is the discounted close-minus-open difference in the negative number of open branches,  $\Delta F^N = \Delta(-N)$ . The acquisition regressor contains the negative primitive  $\Delta F^A = \Delta(-A)$ . The CCP-implied future private-shock difference  $\Delta F^\nu$  enters as an offset with coefficient fixed at one under the normalized branch-logit shock scale. The specification excludes a separate common closure adjustment cost. The fixed-effect estimator retains 317,403 of 391,600 decisions. Standard errors are clustered by branch.

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

Table 7 shows the structural payoff estimates. The revenue coefficient is positive, which is consistent with the fact that closing a branch lowers discounted deposit revenue. In other words, the revenue component discourages closure. The operating cost coefficient enters negatively in the profit function, so closing a branch eliminates its current and future operating costs and consequently raises the value of a closure decision. Finally, the merger acquisition coefficient implies an additional temporary cost of maintaining recently acquired branches open.

Since the revenue component is measured in monetary units, the other payoff coefficients can be expressed relative to  $\widehat{\theta}_R$  to obtain implicit monetary cost measures. The ratio  $\widehat{\theta}_N/\widehat{\theta}_R$  implies a revenue-equivalent recurring operating cost of approximately \$4.0 million per branch per year, while  $\widehat{\theta}_A/\widehat{\theta}_R$  implies an additional revenue-equivalent cost of approximately \$5.4 million per acquired branch during the temporary acquisition period.

To interpret these coefficients in terms of economic magnitudes, Table 8 reports their contributions to the closure index and their changes in predicted closure probabilities.

Table 8: Supply payoff contributions and marginal effects

| Payoff component     | Median index contribution | One-SD AME (p.p.) |
|----------------------|---------------------------|-------------------|
| Deposit revenue      | -0.740                    | 2.336             |
| Fixed operating cost | 1.699                     | 1.470             |
| Recent acquisition   | 0.000                     | 0.238             |

*Notes:* The second column reports the unconditional median of each deterministic contribution to the logit closure index in the 317,403-observation fixed-effect estimation sample. Each contribution is  $\hat{\theta}_k \Delta F_i^k$ . The final column is the average marginal effect, in percentage points, of a one-standard-deviation increase in the corresponding value difference. The recent-acquisition regressor is nonzero for 6.93 percent of the sample, so its median index contribution is zero.

Table 8 shows that the most important structural components are deposit revenue and the recurring operating cost. A one-standard-deviation increase in their respective values changes the predicted closure probability by approximately 2.3 and 1.5 percentage points, respectively. The acquisition component effect is concentrated among recently transferred branches and therefore has a smaller average effect.

Figure 7 compares observed and predicted closure rates across predicted closure rate deciles. The model ranks closure probability reasonably well, with an in-sample area under the ROC curve (AUC) of 0.735. However, the model overpredicts closure in the highest closure probability decile.




Figure 7: Structural closure-probability calibration

*Notes:* The figure uses the 317,403 observations retained by the actual-outcome fixed-effect logit. Points are equal-count deciles of the fitted closure probability, numbered from lowest to highest risk. Vertical bars are 95% confidence intervals for the observed closure rates. The dashed line is the 45-degree line. The fitted probabilities include bank-by-year and ZIP-code fixed effects.

The estimation described above uses the full national sample to recover the demand parameters  $\theta^D$  and the structural payoff parameters  $\theta^\pi$ . However, for the counterfactuals, I instead focus on a selected local banking market at the county level, where I use the estimated demand and bank payoff parameters to solve dynamic closure games. By restricting the counterfactual analysis to one county, the branch network state space is sufficiently small to recover bank-level policies over complete closure plans, as described in Section 3. Otherwise, the state space grows exponentially for larger geographic areas such as commuting zones or states, which may result in settings where the solution is computationally infeasible. Each counterfactual changes conditions of the model (ownership over branches, demand or payoff components, or feasible closure sets), so I resolve the corresponding equilibrium for each scenario.

## 6 Counterfactuals

This section uses the estimated model to study dynamic counterfactuals in which a local banking network evolves under different conditions. I implement all counterfactual scenarios in Haywood County, North Carolina, which serves as an illustrative case.

In the first group of counterfactuals, I compare three initial market structure configu-

rations defined by branch ownership and the degree of concentration in the county. In the second group of counterfactuals, I implement policies that aim to protect branches based on census tract availability and mergers. This last set of policies seeks to quantify trade-offs related to shifting the pre-merger guidelines towards more protective policies for branch closure plans.

## 6.1 Haywood County and counterfactual design

To maintain the dynamic game solution computationally tractable, I implement counterfactuals on Haywood County, North Carolina, as an illustrative case. Among the counties that registered branch declines between 2011-2024, Haywood was selected because it had 17 branches at the beginning of the period (same as the median of this group) and, more importantly, it had both a relatively large number of strategic banks (six) and four cases of branches changing ownership due to mergers. By the end of the sample period, only 9 branches remained open, which represents a decline of 47%.<sup>25</sup> These characteristics make Haywood an ideal county to decompose the effects that the structural model can illustrate, particularly related to mergers, branch contraction and other secular changes.

Figure 8 shows the 2011 and 2024 branch networks in Haywood County. Six strategic banks were present in the county in 2011: First Citizens, Truist, Wells Fargo, PNC, TD Bank and Woodforest. Of these, only the first three remained active in the county by 2024 due to a combination of consolidation by mergers and branch exits. During this period, an RBC branch became part of PNC, a SunTrust branch became part of Truist and two nonstrategic bank branches became part of First Citizens.

---

<sup>25</sup>One branch from HomeTrust, a nonstrategic bank, entered the county during this period. Since this will affect all scenarios, its effect will not affect the counterfactual comparison




Figure 8: Observed branch-network contraction in Haywood County, 2011–2024  
*Notes:* These maps show branch locations. Colors show institutions owning each branch according to their final 2024 owner (owner family), so that the panels are comparable. In Panel A, crosses identify locations that closed by 2024 and black rings identify branch transfers due to mergers. Panel B marks the new 2024 HomeTrust location with a star. The polygons are the 13 census tracts that belong to the county.

Each counterfactual starts from the 17 branches that were open in 2011. I assign 12 of these locations to six strategic banks, as observed in the data, so that they can endogenously choose which to close in equilibrium. The other five branches are considered nonstrategic locations and are held open across scenarios.

For each counterfactual, I solve a new stationary Markov Perfect Equilibrium. Then, I simulate the model for 20 years using the stationary equilibrium, to match the period considered in the supply estimation. The total number of branches that remain open by year 20 is obtained from expectations over the equilibrium distribution of closure plans.<sup>26</sup>

## 6.2 Welfare measurement

In order to quantify consumer welfare in monetary units, and make it comparable to supply side costs, the quantification exercise requires a price sensitivity parameter in utility. Since the demand model estimated above did not include interest rates on deposits, I obtain an

<sup>26</sup>I use the Haywood counterfactuals as a comparison across equilibria that share a common initial network, not as a forecast of the exact 2024 set of surviving branches. Appendix Table 21 reports the branch level fit.

interest sensitivity parameter separately.

I follow [Ho et al. \(2026\)](#) and calibrate the marginal utility of a one percentage point increase in the annual deposit rate using an external elasticity target. I match the average own-rate elasticity of 1.61, estimated by [Pakes et al. \(2024\)](#), using deposits data from 2015 and 12-month CD rates.

Consider the utility of census tract  $c$  of choosing branch  $\ell$  as  $u_{c\ell} = \hat{u}_{c\ell} + \alpha^{pp} r_{\ell}^{pp}$ , where  $\hat{u}_{c\ell}$  is the utility estimated using the demand system in Section 3.2, and  $r_{\ell}^{pp}$  is the 12-month CD rate for the respective branch in my sample. By matching the implied own elasticity of deposits, I obtain  $\alpha^{pp}$  equal to 9.008 utility units per percentage point.<sup>27</sup>

For every counterfactual scenario  $z$  and the branch network  $\omega$  that is reached after 20 periods, define the tract level inclusive value:

$$\mathcal{I}_{ct}^z(\omega) = \log \left[ \exp\{v_{c0t}\} + \sum_{\ell \in \mathcal{J}_{ct}^z(\omega)} \exp\{v_{c\ell t}^z\} \right]. \quad (33)$$

The discounted consumer surplus change relative to the baseline  $B$  is

$$\Delta CS_{0:19}^z = \frac{1}{100 \alpha^{pp}} \sum_{t=0}^{19} \beta^t \left\{ \mathbb{E}_z \left[ \sum_c M_{ct} \mathcal{I}_{ct}^z(\omega_t) \right] - \mathbb{E}_B \left[ \sum_c M_{ct} \mathcal{I}_{ct}^B(\omega_t) \right] \right\}. \quad (34)$$

Note that equation (34) is measured in millions of dollars. Since the utility of the outside option is held fixed within each tract, a tract-specific location normalization cancels from the change in inclusive value.

The estimated supply parameters imply a monetary value for the bank profit payoff components by dividing the cost coefficients by  $\hat{\theta}_R$ , the revenue coefficient. I multiply each ratio by 1000 to convert them from billions to millions of dollars. For the counterfactual scenario  $z$ , I define the discounted producer surplus change as

$$\Delta PS_{0:19}^z = \frac{1000}{\hat{\theta}_R} \sum_{t=0}^{19} \beta^t \left\{ \mathbb{E}_z \left[ \sum_b \pi_b(s_t, a_t; \hat{\theta}^\pi) \right] - \mathbb{E}_B \left[ \sum_b \pi_b(s_t, a_t; \hat{\theta}^\pi) \right] \right\}, \quad (35)$$

where  $\pi_b$  contains the three structural components in equation (14): deposit revenue, the recurring operating cost, and the temporary post-acquisition cost. The corresponding welfare change is

$$\Delta W_{0:19}^z = \Delta CS_{0:19}^z + \Delta PS_{0:19}^z. \quad (36)$$

---

<sup>27</sup>Appendix E.2 provides the analytical derivation, validation, and alternative calibrations.

### 6.3 Decomposition of merger effects

To decompose the effect of mergers on the equilibrium, I solve the model for different cases of branch ownership, effectively introducing distinct market structure configurations: a merger baseline, no mergers, and single bank owner. These alter the set of branch locations  $\ell$  that belong to a bank  $\mathcal{O}_b$  inherited network.

In the "baseline" case, branches that have been involved in merger transactions in the data are designated as being part of the initial branch network of the acquiring bank. This allows the acquiring bank to evaluate closure plans that include these branches, which makes this scenario reflective of the actual ownership changes observed in the data. In the model, I assign the four transferred branches to three strategic banks (one to PNC, one to Truist, and two to First Citizens). By the end of period 20, the model predicts 6.92 branches remaining and 92.2% of households have a branch within five miles.

To construct the "no mergers" scenario, I maintain the ownership of all branches to their 2011 bank owners, which effectively implies that none of the mergers observed between 2011 and 2024 were implemented in the model. This assumption generates changes on several components of the state space compared to the baseline equilibrium. First, strategic banks have smaller branch networks, which prevents them from internalizing business stealing effects from nearby competing branches. Second, the absence of branch transfers implies no additional temporary post-acquisition costs. Third, the specific bank fixed effect in the demand maintains the initial owner which produces a lower demand attractiveness than those generated by the fixed effects for the large acquiring banks in the baseline (PNC, Truist, and First Citizens). Fourth, the interest rate spread on deposits correspond to the initial owners, which may have a lower profitability than the large acquiring banks. The first two effects listed above induce the equilibrium to show fewer closures, while the last two effects imply that banks have a lower profitability from the branch franchise overall, motivating more closures than the baseline.

In Table 9, Panel A compares the merger baseline outcomes with the no merger scenario. When branch ownership maintains the 2011 configuration (no mergers) the number of branches remaining by period 20 is 6.1% higher than the baseline. The higher number of available branches reduces the outside-option deposit share by 1.14 percentage points and increases discounted consumer surplus by \$0.27 million, about 1.1% greater than the baseline. This result shows that mergers may increase the pace of branch closures and these may reduce the consumer surplus that depositors obtain.

An additional scenario shows that the merger baseline practically does not change if the temporary post-acquisition costs from operating transferred branches disappear. The num-

ber of remaining branches increases only by 0.3%, indicating negligible gains for consumer surplus. In Panel B, we can decompose the changes from the merger baseline to the no merger scenario in terms of the branch reassignment effect and the temporary post-acquisition cost. Across the transition path from period 0 to 20, the reassignment effect dominates and is strongest around Year 10. By the year 20, the reassignment effect produces 0.40 less closures compared to the baseline, which represents 94% of the final difference between the no merger configuration and the baseline.

Table 9: Ownership and Consolidation in Haywood County

*Panel A. Main ownership outcomes*

| Scenario                             | Terminal expected<br>branches | Change in<br>branches (%) | Change in outside<br>share (pp) | 20-year consumer-<br>surplus change (\$m) |
|--------------------------------------|-------------------------------|---------------------------|---------------------------------|-------------------------------------------|
| Merger baseline                      | 6.92                          | 0.0                       | 0.00                            | 0.00                                      |
| No mergers (2011 ownership)          | 7.34                          | 6.1                       | -1.14                           | 0.27                                      |
| Mergers, no post-acquisition<br>cost | 6.95                          | 0.3                       | -0.06                           | 0.08                                      |
| Single bank                          | 6.89                          | -0.5                      | -0.62                           | -0.25                                     |

*Panel B. Sequential timing decomposition*

| Transition                   | Year 0 | Year 5 | Year 10 | Year 20 |
|------------------------------|--------|--------|---------|---------|
| Branch-reassignment effect   | -0.06  | -0.33  | -0.42   | -0.40   |
| Post-acquisition-cost effect | -0.18  | -0.14  | -0.07   | -0.02   |

Notes: Branch counts represent expected number of operating branches resulting from the equilibrium closure plans. Panel A changes are relative to merger baseline, in which branches involved in mergers are assigned to their final bank acquirer. Panel B changes are relative to the No mergers (2011 ownership) starting scenario, so negative values indicate fewer expected branches.

The final scenario presents a case in which all the 12 strategic branches in the county are assigned to a single bank (First Citizens). This is a stronger case for market consolidation than the baseline because 9 of these branches are treated as newly acquired for the single bank owner. Results in table 9, Panel A show that the number of terminal branches falls to 6.89 expected branches, 0.5% below the baseline, and discounted consumer surplus falls by \$0.25 million. This shows that a concentrated market structure induces a faster pace of branch closures, but not by a strong margin. In fact, the positive effects on demand from having a large chain operating the branch network (which comes from overall profitability and demand attractiveness), may partially offset the higher closure rate implied by reassignment and post-acquisition costs.<sup>28</sup>

In summary, these scenarios show that mergers increase branch contraction and lower consumer surplus in Haywood County.

<sup>28</sup>Appendix Figure 14 plots the full equilibrium paths.

## 6.4 Targeted access policies and the cost of preserving branches

In this subsection, I implement counterfactual scenarios by imposing various forms of branch protection policies that require particular branches to remain open across the simulation, effectively removing them from the closure action space of the bank owner. The aim of these policies is to inform the debate over whether bank merger guidelines should be revised to incorporate the evaluation of post-merger branch closures and their effects on vulnerable households.

- The *five-mile access policy* protects branch locations that are needed to preserve five-mile coverage for all census tracts. In Haywood, the rule is binding and protects a former RBC location that was transferred to PNC. Since this branch is a scarce option for a vulnerable census tract, this scenario illustrates how a proximity-based rule can be used to ensure that access is maintained.
- The *transferred-branch safeguard* focuses instead on branches that were transferred to another bank due to a merger. In Haywood, the rule protects the four branches involved in merger transactions.
- The *no closure* benchmark is an illustrative case of the upper bound welfare cost of policies that seek to prevent closures from occurring. For this scenario, I assume that strategic banks do not close their branches across the period, and measure the changes in discounted operating costs and consumer surplus.

Table 10, Panel A shows the equilibrium solutions for each of the policies explained above by period 20. Compared to the baseline, the five mile access policy increases the expected number of branches by 13.6% and increases five mile access to 100% by period 20. In the same direction, the transferred-branch safeguard and the no closure benchmark increase the number of branches, respectively, by 49% and 145% compared to the no intervention baseline. These two policies also raise the 5 mile access coverage to 100% and the 2.5 mile access coverage from 75.6% to 84%.

Panel B of Table 10 shows that the depositor surplus benefits are exceeded by the additional operating costs of keeping these branches open. Since closures are efficient outcomes for banks that evaluate benefits and costs over time of maintaining branches, these policies impose additional operating costs. More importantly, the benefits perceived by depositors in terms of consumer welfare are small compared to operating costs for all the policies. Figure 9 maps the spatial distribution of changes in depositor surplus.

Table 10: Access Policies and Modeled Welfare in Haywood County

*Panel A. Period-20 network and access*


| Scenario                     | Expected branches | Change in branches (%) | Access $\leq$ 2.5 mi (%) | Access $\leq$ 5 mi (%) | Outside share (%) |
|------------------------------|-------------------|------------------------|--------------------------|------------------------|-------------------|
| Baseline                     | 6.92              | 0.0                    | 75.62                    | 92.23                  | 61.66             |
| Five-mile access             | 7.86              | 13.6                   | 75.63                    | 100.00                 | 58.95             |
| Transferred-branch safeguard | 10.31             | 49.0                   | 84.44                    | 100.00                 | 54.13             |
| No-closure benchmark         | 17.00             | 145.6                  | 84.44                    | 100.00                 | 43.49             |

*Panel B. Discounted outcomes*

| Scenario                     | Additional open branch-years | Consumer surplus (\$m) | Bank producer flow (\$m) | Total welfare (\$m) |
|------------------------------|------------------------------|------------------------|--------------------------|---------------------|
| Baseline                     | 0.00                         | 0.00                   | 0.00                     | 0.00                |
| Five-mile access             | 8.06                         | 0.59                   | −22.88                   | −22.29              |
| Transferred-branch safeguard | 28.16                        | 1.67                   | −89.65                   | −87.99              |
| No-closure benchmark         | 81.96                        | 4.52                   | −262.15                  | −257.63             |

Notes: Panel A evaluates all outcomes on the period-20 equilibrium network. Panel B discounts flow outcomes over periods 0–19 using  $\beta = 0.95$ . The five-mile policy protects one strategic office. The transferred-branch safeguard protects all four transferred branches involved in merger transactions in the data. The no-closure benchmark protects all 12 strategic branches.

**(a) Five-mile access floor**



**(b) No acquired-office closures**




Figure 9: Spatial incidence of consumer-surplus changes under targeted access policies  
*Notes:* Each tract is shaded by its discounted consumer surplus change relative to the baseline, divided by the number of households in period 0. Diamonds mark the protected branches.

Figure 10 displays the decomposition of total welfare and the producer surplus. Both the five-mile access policy and the merger safeguard for acquired offices show negative effects on welfare, but the second policy has almost 4 times the cost of the first. In terms of producer surplus decomposition, the main component that explains these losses is the recurring operating cost per period, which far exceeds the increase in revenues from attracting deposits from the protected branches.




Figure 10: Welfare decomposition of targeted access policies

*Notes:* The figure reports 20-year discounted changes relative to the baseline, in millions of dollars. Modeled welfare is the sum of consumer surplus and producer surplus.

These welfare results capture the magnitude of the trade-off between cost efficiencies generated by branch closures and the consumer welfare losses from reduced access. Across the policies tested in Haywood County, the cost savings dominate the consumer losses. This example illustrates how the model can be implemented to evaluate alternative merger scenarios in different locations.

These results suggest that merger guidelines in the banking sector would benefit from analyzing the impact on depositors that are exposed to a prospective merger. This would help identify key demographic groups and locations that are more vulnerable to branch closures.

However, preventing closures from occurring is costly, especially considering that these are efficient responses to the evaluation of the profitability of the branch network altogether. Therefore, results in this county illustrate that not all closures are inefficient in terms of welfare, and regulators should balance the additional costs from maintaining a larger branch network against the benefits of preserving depositor access.

## 7 Conclusion

This paper studies the contraction of physical retail networks through the lens of the U.S. banking industry. Branch closures may reduce banks' operating costs, but depositors might be affected due to the lower access to in-person banking services. I study this trade-off by designing and estimating a structural model of endogenous branch closures within a bank's spatial portfolio, which makes it possible to ask whether branch contraction mainly reflects the loss of local access or the removal of costly and overlapping branches nearby.

The empirical results show that branch deposits are determined by local demographics and the degree of competition with other branches nearby. Branch deposits are positively associated with household income located close to a branch. Additionally, the probability that a branch closes is higher when a branch has fewer deposits or is located near other branches of the same bank. The demand estimation shows that depositors value proximity to branches, and particularly in neighborhoods with older populations. The supply estimates indicate that the operating costs of keeping a branch open are substantial in comparison to the deposit revenue it generates.

The counterfactual results illustrate how the trade-off between higher cost savings and lower depositor access can be quantified in a specific county. In Haywood County, maintaining the 2011 ownership (therefore removing the effect of mergers) increases the number of branches by the end of the period and raises consumer surplus, compared to a counterfactual in which branches are reassigned to new owners, based on actual merger events observed in the data. Most of the difference between these two counterfactuals is explained by the branch reassignment effect, which arises when banks can consider larger branch networks to choose which combination allows them to optimize their deposit intake across space and reduce total operating costs by closing overlapping locations in the process.

Welfare exercises serve to illustrate how policies that seek to prevent branch closures may increase depositor welfare, but the additional operating costs of keeping these branches open far exceed any gains. This favors the interpretation that closures are efficient outcomes for banks, even though there may be distributional consequences for specific demographic

groups. These counterfactuals also demonstrate how the model can be used to inform the evaluation of prospective mergers in the industry.

## References

- Aguirregabiria, V., Clark, R., and Wang, H. (2016). Diversification of geographic risk in retail bank networks: evidence from bank expansion after the Riegle-Neal Act. *The RAND Journal of Economics*, 47(3):529–572.
- Aguirregabiria, V., Clark, R., and Wang, H. (2025). The geographic flow of bank funding and access to credit: Branch networks, synergies, and local competition. *American Economic Review*, 115(6):1818–1856.
- Aguirregabiria, V. and Mira, P. (2007). Sequential estimation of dynamic discrete games. *Econometrica*, 75(1):1–53.
- Alessandrini, P., Presbitero, A. F., and Zazzaro, A. (2009). Bank size or distance: what hampers innovation adoption by SMEs? *Journal of Economic Geography*, 10(6):845–881.
- Ashenfelter, O. C., Hosken, D. S., and Weinberg, M. C. (2015). Efficiencies brewed: Pricing and consolidation in the US beer industry. *The RAND Journal of Economics*, 46(2):328–361.
- Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.
- Benson, D., Blattner, S., Grundl, S., Kim, Y. S., and Onishi, K. (2024). Concentration and geographic proximity in antitrust policy: Evidence from bank mergers. *American Economic Journal: Microeconomics*, 16(3):107–133.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841.
- Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2):242.
- Berry, S. T. and Waldfogel, J. (2001). Do mergers increase product variety? evidence from radio broadcasting. *The Quarterly Journal of Economics*, 116(3):1009–1025.
- Bonfim, D., Nogueira, G., and Ongena, S. (2021). “Sorry, we’re closed” bank branch closures, loan pricing, and information asymmetries. *Review of Finance*, 25(4):1211–1259.
- Célerier, C. and Matray, A. (2019). Bank-branch supply, financial inclusion, and wealth accumulation. *The Review of Financial Studies*, 32(12):4767–4809.

- Degryse, H. and Ongena, S. (2005). Distance, lending relationships, and competition. *The Journal of Finance*, 60(1):231–266.
- Dick, A. A. (2008). Demand estimation and consumer welfare in the banking industry. *Journal of Banking & Finance*, 32(8):1661–1676.
- Egan, M., Hortaçsu, A., and Matvos, G. (2017). Deposit competition and financial fragility: Evidence from the US banking sector. *American Economic Review*, 107(1):169–216.
- Egan, M., Lewellen, S., and Sunderam, A. (2022). The cross-section of bank value. *The Review of Financial Studies*, 35(5):2101–2143.
- Egan, M. L., Hortaçsu, A., Kaplan, N. A., Sunderam, A., and Yao, V. (2025). Dynamic competition for sleepy deposits. NBER Working Paper 34267, National Bureau of Economic Research.
- Ellickson, P. B., Grieco, P. L., and Khvastunov, O. (2020). Measuring competition in spatial retail. *The RAND Journal of Economics*, 51(1):189–232.
- Fan, Y. (2013). Ownership consolidation and product characteristics: A study of the us daily newspaper market. *American Economic Review*, 103(5):1598–1628.
- Farrell, J. and Shapiro, C. (1990). Horizontal mergers: An equilibrium analysis. *American Economic Review*, 80(1):107–126.
- Granja, J. a. and Paixão, N. (2026). Bank consolidation and uniform pricing. *Journal of Financial Economics*, 176:104204.
- Herpfer, C., Mjøs, A., and Schmidt, C. (2023). The causal impact of distance on bank lending. *Management Science*, 69(2):723–740.
- Ho, C.-Y. (2015). Switching cost and deposit demand in China. *International Economic Review*, 56(3):723–749.
- Ho, C.-Y., Rysman, M., and Wang, Y. (2026). Demand for performance goods: Import quotas in the chinese movie market. *Unpublished manuscript, Boston University*.
- Ho, K. and Ishii, J. (2011). Location and competition in retail banking. *International Journal of Industrial Organization*, 29(5):537–546.
- Holmes, T. J. (2011). The diffusion of Wal-Mart and economies of density. *Econometrica*, 79(1):253–302.

- Honka, E., Hortaçsu, A., and Vitorino, M. A. (2017). Advertising, consumer awareness, and choice: Evidence from the US banking industry. *The RAND Journal of Economics*, 48(3):611–646.
- Jeziorski, P. (2014). Effects of mergers in two-sided markets: The us radio industry. *American Economic Journal: Microeconomics*, 6(4):35–73.
- Ji, Y., Teng, S., and Townsend, R. M. (2023). Dynamic bank expansion: Spatial growth, financial access, and inequality. *Journal of Political Economy*, 131(8):2209–2275.
- Jia, P. (2008). What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica*, 76(6):1263–1316.
- Jiang, E. X., Yu, G. Y., and Zhang, J. (2026). Bank competition amid digital disruption: Implications for financial inclusion. *The Journal of Finance*, 81(4):1951–2004.
- Kalouptside, M., Scott, P. T., and Souza-Rodrigues, E. (2021). Identification of counterfactuals in dynamic discrete choice models. *Quantitative Economics*, 12(2):351–403.
- Keil, J. and Ongena, S. (2024). The demise of branch banking—technology, consolidation, bank fragility. *Journal of Banking & Finance*, 158:107038.
- Kim, M. (2023). Does the internet replace brick-and-mortar bank branches? Technical report, Oklahoma State University. Working paper.
- Kuehn, J. (2018). Spillovers from entry: the impact of bank branch network expansion. *The RAND Journal of Economics*, 49(4):964–994.
- Kutzbach, M. J., Wagner, G., and Watson, C. L. (2025). Bank competition for neighborhood deposits. Preliminary and incomplete manuscript, January 2025.
- Morelli, J. M., Moretti, M., and Venkateswaran, V. (2025). Geographic funding risk and market power in deposit markets. Working paper, April 25, 2025.
- Narayanan, R. P., Ratnadiwakara, D., and Strahan, P. E. (2025). The decline of bank branching. NBER Working Paper 33773, National Bureau of Economic Research.
- Nguyen, H.-L. Q. (2019). Are credit markets still local? evidence from bank branch closings. *American Economic Journal: Applied Economics*, 11(1):1–32.
- Nishida, M. (2015). Estimating a model of strategic network choice: The convenience-store industry in okinawa. *Marketing Science*, 34(1):20–38.

- Oberfield, E., Rossi-Hansberg, E., Trachter, N., and Wenning, D. T. (2024). Banks in space. Working Paper 32256, National Bureau of Economic Research.
- Pakes, A., Ostrovsky, M., and Berry, S. (2007). Simple estimators for the parameters of discrete dynamic games (with entry/exit examples). *the RAND Journal of Economics*, 38(2):373–399.
- Pakes, A., Whinston, M. D., and Zheng, F. (2024). The consumer welfare effects of bank mergers. Working paper, October 3, 2024.
- Rysman, M., Townsend, R. M., and Walsh, C. (2026). Bank branching strategies in the 1997 thai financial crisis and local access to credit. *Journal of the European Economic Association*, page jvag007.
- Sakong, J. and Zentefis, A. K. (2025). What drives disparities in bank branch use? a geospatial analysis. *Review of Finance*. Forthcoming. Available at SSRN.
- Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics*, 37(3):619–640.
- Wang, Y., Whited, T. M., Wu, Y., and Xiao, K. (2022). Bank market power and monetary policy transmission: Evidence from a structural estimation. *The Journal of Finance*, 77(4):2093–2141.
- Williamson, O. E. (1968). Economies as an antitrust defense: The welfare tradeoffs. *American Economic Review*, 58(1):18–36.
- Yankov, V. (2024). In search of a risk-free asset: Search costs and sticky deposit rates. *Journal of Money, Credit and Banking*, 56(5):1053–1098.

## A Data construction and additional institutional evidence

This appendix records the data-cleaning and classification choices used to construct the main combined panel, and collects additional descriptive exhibits. The objective is to distinguish the broad FDIC series used to describe the industry from the cleaned branch panel used in the empirical analysis.

### A.1 Summary of Deposits branch panel and sample universes

The raw Summary of Deposits files identify offices, their current bank certificate, service type, address, coordinates, and reported deposits. I construct a persistent physical-office identifier across annual waves. The cleaning first follows the FDIC office identifier and then corrects adjacent-year relabeling when address and geographic information indicate continuity at the same physical location. A surviving office that changes bank certificate is therefore retained as the same physical branch with a new owner.

The paper uses distinct historical and SOD-based branch universes. Table 11 reconciles them. The main headline series is the FDIC December 31 commercial-bank aggregate: it contains 6,275 banks and 82,564 branches in 2011 and 3,806 banks and 69,167 branches in 2025. By contrast, SOD measures offices and deposits as of June 30 and includes a broader set of insured institutions. After imposing positive deposits and a usable persistent office identifier, the cleaned SOD panel contains 7,521 certificates and 94,618 offices in 2011 and 4,429 certificates and 73,876 offices in 2025.

The opening and closure series further excludes any physical office ever classified as a savings association or regulated by the Office of Thrift Supervision. That sample contains 6,790 certificates and 85,473 offices in 2011 and 4,204 certificates and 68,075 offices in 2025. The estimation-ready demand panel is a separate, non-nested SOD universe: it requires positive deposits, complete ACS covariates, a valid branch–tract match, and location outside New York and New Jersey, but it does not impose the ever-SA/OTS restriction. It contains 6,976 certificates, 83,545 branches, 61,338 tracts, and 2,892,806 tract–branch pairs in 2011. Its last available year is 2024, with 4,270 certificates, 66,327 branches, 69,626 tracts, and 2,529,298 pairs.

An opening in year  $t$  is an office observed in  $t$  but not  $t - 1$ , while a closure is an office observed in  $t - 1$  but not  $t$ . The corrected rate series divides both flows by the cleaned

Table 11: FDIC branch universes used in the paper

| Universe                          | Reference date | Year | Institutions | Branches | Tracts | Pairs     |
|-----------------------------------|----------------|------|--------------|----------|--------|-----------|
| Historical commercial-bank series | December 31    | 2011 | 6,275        | 82,564   | —      | —         |
| Historical commercial-bank series | December 31    | 2025 | 3,806        | 69,167   | —      | —         |
| Cleaned positive-deposit SOD      | June 30        | 2011 | 7,521        | 94,618   | —      | —         |
| Cleaned positive-deposit SOD      | June 30        | 2025 | 4,429        | 73,876   | —      | —         |
| Opening/closure rate sample       | June 30        | 2011 | 6,790        | 85,473   | —      | —         |
| Opening/closure rate sample       | June 30        | 2025 | 4,204        | 68,075   | —      | —         |
| Estimation-ready demand panel     | June 30        | 2011 | 6,976        | 83,545   | 61,338 | 2,892,806 |
| Estimation-ready demand panel     | June 30        | 2024 | 4,270        | 66,327   | 69,626 | 2,529,298 |

*Notes:* Institutions in the SOD-based rows are distinct FDIC certificates. The historical series is the FDIC December 31 commercial-bank aggregate. The cleaned SOD sample requires positive reported deposits. The opening/closure sample additionally excludes physical branches ever classified as savings associations or regulated by the Office of Thrift Supervision. The estimation panel is a separate, non-nested SOD universe and ends in 2024. Source: FDIC Annual Data Summaries, FDIC Summary of Deposits, and authors' calculations.

branch stock in  $t - 1$ . Thus,

$$\text{OpenRate}_t = \frac{\#\{\ell : \ell \notin \mathcal{B}_{t-1}, \ell \in \mathcal{B}_t\}}{|\mathcal{B}_{t-1}|}, \quad \text{CloseRate}_t = \frac{\#\{\ell : \ell \in \mathcal{B}_{t-1}, \ell \notin \mathcal{B}_t\}}{|\mathcal{B}_{t-1}|}.$$

An office disappearance caused only by a corrected identifier change is not treated as a closure. For comparison, I also report the former calculations that use a year- $t$  denominator.

The demand sample removes offices with nonpositive reported deposits because the estimator matches log observed and predicted deposits. It also removes investment banks whose aggregate deposits are reported at one or two offices rather than distributed across a conventional retail network.

New York and New Jersey are excluded because their unusually dense branch networks create exceptionally large five-mile choice sets.

## A.2 Ownership histories, transfers, and closures

For each persistent physical office, I compare the owner certificate in adjacent SOD waves. If the office remains active but its certificate changes, I code an ownership transfer rather than a closure. Many such transfers arise from whole-bank mergers or acquisitions. Others reflect sales of selected branches, internal reorganizations, or changes that cannot be classified from the SOD panel alone.

When a bank certificate disappears between  $t$  and  $t + 1$  and at least one of its offices transfers to another certificate, the closure-segment calculations assign its non-transferred offices to the dominant successor certificate. This prevents a disappearing target certificate from mechanically classifying every non-transferred office under the target's pre-merger size

category.

For each disappearing certificate, I count its physical offices that survive under each successor in  $t + 1$  and select the successor receiving the largest number. Ties are resolved deterministically in favor of the numerically lowest successor certificate. Non-transferred closures and the corresponding year- $t$  branch exposure are assigned to that successor; the successor’s network is measured in  $t + 1$ .

The ownership-transfer sample covers SOD transitions ending in 2011–2024. It identifies 3,707 seller–buyer–year events involving 27,935 transferred physical branches. A transfer is classified as direct when the selling certificate remains active in the post-transfer SOD wave. There are 577 such events and 2,299 such branches, equal to 15.6% of events and 8.2% of transferred branches. I match transfers to the FDIC Bank Structure Changes history using the SOD transition year, outgoing certificate, and acquiring certificate. Effective dates from July through December are assigned to the following June SOD wave. This rule matches 2,647 transfer events involving 23,063 branches to an official business-combination record. Unmatched certificate changes remain ownership transfers and are not relabeled as confirmed mergers.

### A.3 Census geography and tract matching

The branch–tract panel combines annual ACS five-year estimates with Census Centers of Population and branch coordinates. Median household income is used as the representative tract income measure because it is less sensitive to within-tract outliers than mean income. Household counts scale the potential deposit mass, and median age measures local demographic composition.

The panel uses the 2000 tract vintage in 2009, the 2010 vintage from 2010 through 2019, and the 2020 vintage from 2020 onward. For each vintage I use the Census Bureau’s corresponding tract-level mean Center of Population. Annual ACS five-year estimates are merged using the tract identifier appropriate for that year. The construction does not impose a single set of tract boundaries across decades: tract splits and mergers are therefore absorbed when the vintage changes in 2010 and 2020 rather than allocated through an areal interpolation.

Connecticut’s replacement of counties with planning regions changes ACS geographic identifiers beginning in 2022. For 2022–2024, I map 2020 Connecticut tract identifiers to the contemporaneous ACS identifiers using unique tract-name labels, with explicit overrides for the two tract-9901 records. Coverage and uniqueness are checked before the merge. No analogous within-vintage crosswalk is needed for other states.

The branch map assigns each persistent office its most recent valid SOD latitude and

longitude observed through 2024 and applies that location to each tract vintage. Missing coordinates, coordinates outside the valid global latitude and longitude ranges, and the (0, 0) placeholder are excluded. The cleaner records the source year of the selected coordinate, the number of distinct reported coordinates for each branch, and the distance between its first and last valid coordinates. Population centers with nonpositive population, invalid coordinates, or duplicate tract identifiers are likewise excluded and reported by vintage.

Distances are WGS84 ellipsoidal geodesic distances. One mile is 1,609.344 meters. A branch is linked to a tract when the distance from the branch coordinate to the tract population center is no more than five miles. In 2024, the positive-deposit branch sample contains 74,400 offices, of which 1,678 have no branch–tract map. Geographic matching initially produces 3,697,568 candidate pairs for 72,722 branches. Excluding New York and New Jersey leaves 2,588,840 pairs and 66,726 branches; the ACS match and complete-covariate requirements leave the final 2,529,298 pairs, 66,327 branches, and 69,626 tracts reported in Table 11. The same construction sequence is applied in every panel year.

## A.4 Call Report matching and deposit spreads

Annual FFIEC Call Report variables are matched to each SOD branch through its bank owner. The resulting bank-year spread is predetermined within the annual closure decision and enters the bank’s deposit-revenue component in the supply model.

For bank certificate  $b$  and year  $t$ , let  $L_{bt}$  and  $D_{bt}$  denote year-end loans and deposits,  $\bar{L}_{bt}$  and  $\bar{D}_{bt}$  their average current- and prior-year balances,  $II_{bt}^L$  interest income on loans, and  $IE_{bt}^D$  domestic deposit interest expense. I construct

$$r_{bt}^L = \frac{II_{bt}^L}{\bar{L}_{bt}}, \quad r_{bt}^D = \frac{IE_{bt}^D}{\bar{D}_{bt}}, \quad \phi_{bt} = \frac{L_{bt}}{D_{bt}}, \quad m_{bt} = \phi_{bt} r_{bt}^L - r_{bt}^D.$$

Loans use RCFD2122, RCON2122, RCFDB528, or RCONB528, plus foreign-office loans RCFN2122. Deposits use RCFD2200 or RCON2200, plus RCFN2200. Loan interest income is RIAD4010. Deposit interest expense uses RIAD4170 when reported; otherwise it is the sum of RIAD4508, RIAD0093, RIADA517, and RIADA518 through 2016, or RIAD4508, RIAD0093, RIADHK03, and RIADHK04 beginning in 2017. Foreign deposit interest RIAD4172 is excluded.

Call Report monetary variables are reported in thousands of dollars, so the rates, leverage ratio, and spread are unit-free decimals. I retain the latest quarterly report in each certificate–year, normally the December report. The average balance is the mean of current and prior year-end balances when the same certificate has a report in the immediately preceding calendar year; it uses the current balance otherwise. Missing income or expense

components remain missing rather than being set to zero, and missing or nonpositive balances produce a missing yield or spread.

The Call Report’s FDIC certificate number is matched directly to the SOD owner certificate. Merger histories are not spliced across predecessor and successor certificates. Consequently, a successor without a consecutive prior-year report uses its current balance, while a disappearing certificate leaves the panel. The resulting 2004–2024 matched panel contains 137,833 certificate–years, of which 135,487 have finite spreads and 2,346 are missing.

## A.5 RateWatch deposit rates and imputation

I use RateWatch to construct branch-level deposit rates for 2009–2023. The interest rate measure is the 12-month CD rate, corresponding to RateWatch product code 12MCD10K. For each RateWatch location and calendar year, I average the available rate observations within the year and match the resulting annual rate to the FDIC branch panel. When multiple RateWatch records remain for the same branch-year after matching, I use their median.

RateWatch does not cover every branch or bank-year. I fill missing rates in three steps. First, when a branch is missing a rate but RateWatch observes another branch of the same bank in that year, I estimate

$$r_{\ell bt}^{RW} = \eta_{bt} + \lambda_{s(\ell)t} + \varepsilon_{\ell bt}, \quad (37)$$

where  $r_{\ell bt}^{RW}$  is the observed 12-month CD rate at branch  $\ell$  of bank  $b$ ,  $\eta_{bt}$  is a bank-by-year effect, and  $\lambda_{s(\ell)t}$  is a state-by-year effect based on the branch’s state. I use the fitted value to impute the missing branch rate. If the relevant state-by-year effect is not supported in the RateWatch sample, I use the median observed rate for that bank-year.

Second, for bank-years with no RateWatch observations, I use the overlap between RateWatch and the Call Reports to relate the 12-month CD rate to the bank’s deposit interest cost. I estimate

$$r_{\ell bt}^{RW} = \beta_r r_{bt}^D + \delta_t + \lambda_{s(\ell)t} + \varepsilon_{\ell bt}, \quad (38)$$

where  $r_{bt}^D$  is the annual deposit interest cost constructed from the Call Reports and described in the preceding subsection. The preferred prediction uses the state-by-year specification. When the corresponding state-by-year effect is unavailable, I use a specification with year effects only.

Finally, for observations that remain unresolved, I assign the median rate among observed and previously imputed branches in the same state and year. If that value is unavailable, I use the corresponding year median. This procedure produces a rate measure at the branch-year level that preserves observed RateWatch information whenever available while using

within-bank geographic variation and bank-level deposit expenses to extend coverage.

These constructed rates are used to calibrate the monetary scale of consumer surplus in Section 6, with the elasticity-matching procedure described in Appendix E.2.

## A.6 Additional descriptive evidence

Table 12 reports summary statistics for the estimation-ready 2024 branch–tract cross-section.

Table 12: Summary statistics for the combined branch–census tract dataset, 2024

| Variable                                  | <i>N</i>  | Mean    | SD        | P10    | P25    | Median | P75     | P90     |
|-------------------------------------------|-----------|---------|-----------|--------|--------|--------|---------|---------|
| <i>Panel A. Branch–census tract pairs</i> |           |         |           |        |        |        |         |         |
| Distance (miles)                          | 2,529,298 | 3.09    | 1.26      | 1.23   | 2.12   | 3.24   | 4.17    | 4.68    |
| <i>Panel B. Census tracts</i>             |           |         |           |        |        |        |         |         |
| Branches available within five miles      | 69,626    | 36.3    | 41.6      | 3      | 8      | 24     | 49      | 82      |
| Banks available within five miles         | 69,626    | 13.6    | 10.3      | 2      | 6      | 12     | 18      | 27      |
| Households                                | 69,626    | 1,585.9 | 661.1     | 830    | 1,121  | 1,508  | 1,957   | 2,424   |
| Median household income (\$)              | 69,626    | 86,239  | 41,213    | 43,515 | 57,770 | 77,535 | 105,238 | 140,350 |
| Median age (years)                        | 69,626    | 40.0    | 8.1       | 30.8   | 34.7   | 39.3   | 44.4    | 49.9    |
| <i>Panel C. Bank branches</i>             |           |         |           |        |        |        |         |         |
| Census tracts exposed to branch           | 66,327    | 38.1    | 49.1      | 2      | 6      | 22     | 51      | 93      |
| Branch deposits (\$000s)                  | 66,327    | 220,232 | 3,645,970 | 18,956 | 40,310 | 78,730 | 147,759 | 260,294 |

*Notes:* The table reports unweighted summary statistics for the estimation-ready 2024 cross-section. Panel A uses each unique census-tract–branch pair; Panel B uses one observation per census tract; and Panel C uses one observation per branch. A branch is available to a tract—and a tract is exposed to a branch—when the distance from the branch to the tract’s population-weighted center is no more than five miles. Distances are geodesic and converted from meters using 1 mile = 1,609.344 meters. Branches and banks are distinct counts within each tract’s choice set. Deposits are reported in thousands of dollars; income is median household income in dollars; age is median age; and households is the number of households. The estimation-ready sample requires positive branch deposits and complete tract covariates and excludes New York and New Jersey. *Source:* Federal Deposit Insurance Corporation, 2024 Summary of Deposits; U.S. Census Bureau, 2024 American Community Survey five-year estimates and 2020 Centers of Population; author’s calculations.

Table 13 compares the fixed set of 20 banks ranked by their 2024 branch counts with all other banks. It reports the number of banks and branches at the 2011 and 2025 endpoints, each group’s 2011 branch share, transition-attributed branch-year exposure, closures, closure share, and closure rate. The table is descriptive: the Top-20 share of closures reflects both its initial exposure and its higher observed closure rate.

Figure 12 reports the primary-access time series referenced in Section 2.3.

Finally, Table 14 reports the full SCF distribution underlying the household-choice discussion in Section 2.3.

Table 13: Branch closures at the 2024 Top-20 banks and other banks

| Group       | Banks |       | Branches |        |            | Branch-years | Closures, 2011–2025 |        |      |
|-------------|-------|-------|----------|--------|------------|--------------|---------------------|--------|------|
|             | 2011  | 2025  | 2011     | 2025   | 2011 share |              | Total               | Share  | Rate |
| 2024 Top 20 | 19    | 20    | 33,483   | 28,526 | 39.2%      | 452,217      | 15,969              | 48.5%  | 3.5% |
| Other banks | 6,771 | 4,184 | 51,990   | 39,549 | 60.8%      | 637,679      | 16,926              | 51.5%  | 2.7% |
| Total       | 6,790 | 4,204 | 85,473   | 68,075 | 100.0%     | 1,089,896    | 32,895              | 100.0% | 3.0% |

*Notes:* The Top-20 group is fixed using branch counts in the full cleaned 2024 FDIC Summary of Deposits panel, matching the structural-sample ranking. The closure sample excludes physical branches ever classified as savings associations or regulated by the Office of Thrift Supervision. Closures are branches observed at least once from 2011 through 2024 and not observed again by 2025. When a certificate disappears between  $t$  and  $t+1$ , its non-transferred closures and its year- $t$  branch exposure are assigned to the successor receiving the largest number of its surviving branches; ties are resolved by the lowest successor certificate. The closure rate is closures divided by transition-attributed branch-years at risk during 2011–2024. Source: FDIC Summary of Deposits.




Figure 11: Change in branch counts by county, 2011–2025

*Notes:* County percentage changes are capped at the 1st and 99th percentiles. Counties with small initial branch counts can exhibit large percentage changes. Source: FDIC Summary of Deposits.




Figure 12: Primary method of bank account access, 2013–2023

*Notes:* The figure plots the primary account-access method among banked households that accessed an account in the preceding 12 months. The residual category includes ATMs, telephone banking, and other reported methods. Source: FDIC National Survey of Unbanked and Underbanked Households.

Table 14: Most important reason for choosing the institution holding the main checking account, 2007–2022

|                                    | 2007 | 2010 | 2013 | 2016 | 2019 | 2022 |
|------------------------------------|------|------|------|------|------|------|
| Location of offices                | 45.9 | 46.0 | 43.7 | 44.0 | 43.3 | 42.1 |
| Many services at one institution   | 16.2 | 16.6 | 18.0 | 15.5 | 16.9 | 14.2 |
| Lowest fees/minimum balance        | 13.7 | 14.2 | 15.1 | 13.8 | 12.6 | 11.4 |
| Recommendation from family/friends | 4.2  | 4.0  | 3.4  | 3.6  | 4.5  | 6.7  |
| Other                              | 20.1 | 19.1 | 19.8 | 23.0 | 22.7 | 25.7 |

*Notes:* The table reports survey-weighted percentages of U.S. families with a checking account that identify each category as the most important reason for choosing the institution holding their main checking account. Location, many services, lowest fees, and recommendation correspond to SCF variable X3530 codes 3, 6, 7, and 1, respectively; Other combines all remaining valid responses. Estimates use the final family weight (X42001), average over the five implicates, and sum to 100 percent within each wave. *Source:* Board of Governors of the Federal Reserve System, Survey of Consumer Finances, 2007–2022 full public-use files; author’s calculations.

## B Reduced-form robustness and additional results

### B.1 Probit coefficients for one-year branch exit

Table 15 reports the probit coefficients underlying the average marginal effects in Table 2.

Table 15: Probit coefficient estimates for one-year branch exit

|                                     | (1)                  | (2)                  | (3)                  |
|-------------------------------------|----------------------|----------------------|----------------------|
| Log deposits                        | -0.309***<br>(0.004) | -0.358***<br>(0.005) | -0.503***<br>(0.010) |
| Log(1 + rival branches)             | 0.032***<br>(0.007)  | 0.032***<br>(0.010)  | -0.042***<br>(0.015) |
| Log(1 + own branches)               | 0.175***<br>(0.006)  | 0.225***<br>(0.007)  | 0.319***<br>(0.012)  |
| Log weighted age                    | 0.226***<br>(0.076)  | 0.284***<br>(0.086)  | 0.766***<br>(0.171)  |
| Log weighted income                 | -0.003<br>(0.034)    | -0.020<br>(0.040)    | -0.195***<br>(0.074) |
| Ownership transfer: year 1          | 0.347***<br>(0.015)  | 0.345***<br>(0.020)  | 0.431***<br>(0.035)  |
| Ownership transfer: year 2          | 0.210***<br>(0.016)  | 0.233***<br>(0.021)  | 0.321***<br>(0.039)  |
| Ownership transfer: year 3          | 0.131***<br>(0.018)  | 0.211***<br>(0.024)  | 0.326***<br>(0.043)  |
| Bank fixed effects                  | Yes                  | No                   | No                   |
| ZIP fixed effects                   | Yes                  | Yes                  | No                   |
| Year fixed effects                  | Yes                  | No                   | No                   |
| ZIP $\times$ year fixed effects     | No                   | No                   | Yes                  |
| Bank $\times$ year fixed effects    | No                   | Yes                  | Yes                  |
| Total branch-years                  | 1,158,005            | 1,158,005            | 1,158,005            |
| Estimation branch-years             | 912,322              | 693,655              | 206,976              |
| Retained share (percent)            | 78.8                 | 59.9                 | 17.9                 |
| Total sample exit rate (percent)    | 2.99                 | 2.99                 | 2.99                 |
| Retained sample exit rate (percent) | 3.77                 | 4.87                 | 15.28                |

*Notes:* The first line in each cell reports the probit coefficient; branch-clustered standard errors appear in parentheses. The outcome equals one when a branch observed in year  $t$  is absent in year  $t + 1$ . All columns begin from the same all-bank complete-covariate sample for decision years 2009–2023. Estimation counts differ because fixed-effect groups without outcome variation do not contribute to the nonlinear estimator. Bank-by-year and ZIP-by-year effects subsume the corresponding lower-order fixed effects. The corresponding average marginal effects are reported in main-text Table 2. Stars: \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

### B.2 SCF sample construction and inference

The regressions in Section 4.4 use the 2016, 2019, and 2022 full public-use Survey of Consumer Finances. I use the checking-account institution links and the corresponding institution-distance measures reported in the SCF. The distance measure is the distance from the household’s census-tract centroid to the nearest branch of the linked institution. Observa-

tions reported as less than one-half mile are assigned a distance of 0.25 miles. The estimation samples require a valid institution link and distance measure and, for the log-balance specifications in Table 4, a positive checking balance.

The three specifications differ in their unit of observation. The main-account specification uses the household’s primary checking account and the controls reported in Table 4. The all-accounts specification expands checking-account slots 1–6 and includes household-by-wave fixed effects. Each account observation retains the household survey weight. The household–institution specification sums checking balances across accounts linked to the same institution and also includes household-by-wave fixed effects. In this specification, the family weight is divided by the number of qualifying institutions for that household and wave so that a household’s total weight does not increase with its number of banking relationships. Wave-specific indicators identify the institution holding the household’s main checking account.

I use the final SCF family weight, X42001, and account for multiple imputation and the survey sampling design. The SCF contains five implicates for each family. I estimate each specification separately in the five implicates and average the resulting coefficient estimates. Because detailed sample-design information is not released in the public files, sampling uncertainty is estimated using the 999 replicate weights supplied by the Federal Reserve. These replicates are constructed by bootstrap resampling designed to reproduce the main dimensions of the original sampling variation, with analysis weights recomputed for every replicate. For replicate  $r$ , I use the product of the supplied weight  $WT1B_r$  and multiplicity factor  $MM_r$  and re-estimate the regression on the first implicate. The [Federal Reserve Board’s SCF standard-error documentation](#) describes the procedure in more detail.

Let  $B$  denote the covariance matrix across the five implicate estimates and  $U$  the covariance matrix across the 999 replicate estimates. Following the SCF procedure for combining sampling and imputation uncertainty, I compute the reported covariance matrix as

$$\widehat{\text{Var}}(\widehat{\beta}) = U + \frac{6}{5}B.$$

The other SCF statistics reported in the paper use the same final family weights and treatment of the five implicates where applicable. Replicate-weight inference is used when sampling uncertainty is reported.

### B.3 Same-bank deposit reallocation after organic closures

Let  $e$  index an organic closure,  $c(e)$  denote the closing branch, and  $T_e$  denote the first Summary of Deposits wave in which  $c(e)$  is absent. I select  $r(e)$  as the nearest office owned

by the same bank in the final pre-closure wave, requiring that it remain open in wave  $T_e$  and lie within five miles of the closing branch.<sup>29</sup> The resulting sample contains 12,398 closures from 2010 through 2022, with a median closing-branch-recipient distance of 1.62 miles.

I residualize log deposits using branch, bank-by-year, and ZIP-by-year fixed effects in the full branch-year panel. On the resulting residuals, I estimate

$$\hat{u}_{r(e),T_e+h} = \eta_e + \sum_{k \in \mathcal{H}} \beta_k \mathbf{1}\{h = k\} + \varepsilon_{eh}, \quad \mathcal{H} = \{-3, -2, 0, 1, 2, 3\}, \quad (39)$$

where  $h = -1$  is the omitted final pre-closure wave and  $\eta_e$  is a closure-event fixed effect. Because Summary of Deposits balances are branch allocations reported by banks rather than customer-level account flows, the estimates measure changes in reported branch deposits and do not trace individual depositors.




Figure 13: Reported deposits at the nearest same-bank branch around an organic closure

*Notes:* The figure reports exact percentage effects,  $100[\exp(\hat{\beta}_h) - 1]$ , from equation (39). Event time  $h = 0$  is the first wave in which the closing branch is absent, and  $h = -1$  is normalized to zero. Bars denote 95% confidence intervals based on two-way clustering by the focal pre-closure bank and recipient branch. The outcome is residualized using branch, bank-by-year, and ZIP-by-year fixed effects.

Figure 13 shows a discrete increase in reported deposits at the recipient branch when the nearby office disappears. Recipient deposits are 25.3% higher in the first absent wave than in the omitted pre-closure wave, after removing branch, bank-year, and ZIP-year components. The average effect over the following three waves is 19.5%. The two pre-event coefficients are economically small, but they are jointly different from zero in the headline

<sup>29</sup>Organic closures remain absent through  $T_e + 3$  and exclude temporary reporting gaps, unstable ownership histories, inferred ownership transfers, and likely successor or replacement offices. The sub sample also removes overlapping event windows for the same recipient.

sample ( $p = 0.030$ ). Since closure timing is not randomly assigned, I interpret the estimates as a descriptive event-time pattern.<sup>30</sup>

The event-time pattern therefore provides evidence of adjustment within the surviving branch network. When banks remove offices, reported balances rise immediately and persistently at nearby same-bank branches.

### B.3.1 Sample construction and event-time support

The headline sample applies the strict organic-closure and recipient-selection rules. Requiring a closure to remain absent through  $T_e + 3$ , to have usable pre-closure history, and to have no inferred ownership transfer or likely successor leaves 23,766 organic closures from 2010 through 2022. Of these events, 15,965 have an eligible same-bank recipient within five miles. Removing overlapping event windows for the same recipient produces the headline sample of 12,398 closure events and 12,121 recipient branches.

The headline recipient must have positive reported deposits and a nonmissing residualized outcome at  $h = -3, -2, -1, 0$ , but observations after the closing wave are not imputed when the recipient subsequently exits or has an unusable deposit report. The number of events contributing to the estimates therefore falls from 12,398 at  $h = 0$  to 12,082 at  $h = 1$ , 11,760 at  $h = 2$ , and 11,429 at  $h = 3$ . The balanced-event-window robustness instead requires positive deposits, a nonmissing residualized outcome, a physical-branch service code, and stable ownership at every event time from  $h = -3$  through  $h = 3$ . It contains 10,845 events and 10,610 recipient branches at every horizon.

### B.3.2 Detailed estimates and sample robustness

Table 16 shows that the post-closure pattern is stable across the three samples. The immediate effect is 25.3% in the headline sample, 26.3% in the balanced event window, and 24.0% among closure cohorts through 2019. The corresponding averages over  $h = 1, 2, 3$  are 19.5, 20.2, and 19.0%.

The pre-period diagnostics are less uniform. The headline joint test rejects zero pre-event coefficients at the 5% level ( $p = 0.030$ ), even though its largest lead is only 0.48%. The balanced sample does not reject ( $p = 0.138$ ). The closures-through-2019 sample rejects more strongly ( $p = 0.005$ ), driven primarily by a negative 0.82% estimate at  $h = -3$ . These diagnostics reinforce the descriptive interpretation of the event-time pattern. The

---

<sup>30</sup>The scaled comparisons are descriptive since they incorporate reporting reallocations and do not measure the fraction of customers transferred to  $r(e)$ .

Table 16: Same-bank deposit reallocation: detailed event-study estimates

| Event time                    | Headline                | Balanced<br>event window | Closures<br>through 2019 |
|-------------------------------|-------------------------|--------------------------|--------------------------|
| $h = -3$                      | 0.05<br>[−0.63, 0.72]   | 0.08<br>[−0.60, 0.76]    | −0.82<br>[−1.60, −0.04]  |
| $h = -2$                      | 0.48<br>[0.06, 0.90]    | 0.41<br>[−0.04, 0.86]    | 0.20<br>[−0.37, 0.78]    |
| $h = -1$ (reference)          | 0.00                    | 0.00                     | 0.00                     |
| $h = 0$                       | 25.26<br>[23.43, 27.12] | 26.25<br>[24.18, 28.36]  | 24.02<br>[22.15, 25.91]  |
| $h = 1$                       | 22.37<br>[20.47, 24.30] | 23.14<br>[21.03, 25.29]  | 21.40<br>[19.42, 23.41]  |
| $h = 2$                       | 19.40<br>[17.21, 21.63] | 20.19<br>[17.98, 22.44]  | 19.12<br>[16.94, 21.35]  |
| $h = 3$                       | 16.87<br>[14.71, 19.07] | 17.23<br>[14.97, 19.53]  | 16.39<br>[14.02, 18.81]  |
| Average, $h = 1, 2, 3$        | 19.53<br>[17.48, 21.61] | 20.16<br>[17.99, 22.37]  | 18.95<br>[16.81, 21.14]  |
| Joint pre-period $p$ -value   | 0.030                   | 0.138                    | 0.005                    |
| Event-time observations       | 84,863                  | 75,915                   | 55,400                   |
| Events at $h = 0$             | 12,398                  | 10,845                   | 8,094                    |
| Events at $h = 3$             | 11,429                  | 10,845                   | 7,444                    |
| Recipient branches at $h = 0$ | 12,121                  | 10,610                   | 8,048                    |

*Notes:* Coefficient entries are exact percentage effects,  $100[\exp(\hat{\beta}_h) - 1]$ , with 95 percent confidence intervals in brackets. Event time  $h = 0$  is the first Summary of Deposits wave in which the closing branch is absent, and  $h = -1$  is the omitted period. Log branch deposits are residualized in the full branch-year panel using branch, bank-by-year, and ZIP5-by-year fixed effects. The event-study regressions include closure-event fixed effects, with standard errors two-way clustered by the closing branch's pre-closure bank and the recipient branch. The balanced event window requires an eligible recipient observation at all seven event times. The final column restricts closure cohorts to events whose first absent wave occurs by 2019; later event-window observations can occur after 2019. The joint pre-period test is  $\beta_{-3} = \beta_{-2} = 0$ . Deposits are reported branch allocations, and the estimates are descriptive.

final column restricts the closure cohort to events whose first absent wave occurs by 2019; it does not remove later event-window observations that fall after 2019.

### B.3.3 Scaled recipient gains

To compare the reported recipient response with the size of the disappearing office, define the raw scaled gain in the first absent wave as

$$G_{e0}^{\text{raw}} = \frac{D_{r(e),T_e} - D_{r(e),T_e-1}}{D_{c(e),T_e-1}}. \quad (40)$$

I also construct a residual-growth-adjusted analog,

$$G_{e0}^{\text{adj}} = \frac{D_{r(e),T_e-1} [\exp(\hat{u}_{r(e),T_e} - \hat{u}_{r(e),T_e-1}) - 1]}{D_{c(e),T_e-1}}, \quad (41)$$

which converts the change in residualized log deposits into a dollar change using the recipient's final pre-closure deposits.

Table 17: Scale of recipient deposit gains in the first absent wave

|                                                      | Headline | Balanced<br>event window |
|------------------------------------------------------|----------|--------------------------|
| <i>Panel A: Raw scaled recipient gain</i>            |          |                          |
| Nonnegative recipient gain (percent)                 | 89.7     | 90.6                     |
| 10th percentile (percent)                            | -0.3     | 0.6                      |
| Median (percent)                                     | 97.5     | 99.8                     |
| 90th percentile (percent)                            | 189.3    | 192.6                    |
| Aggregate scaled gain (percent)                      | 77.9     | 87.4                     |
| <i>Panel B: Residual-growth-adjusted scaled gain</i> |          |                          |
| Median (percent)                                     | 47.1     | 50.2                     |
| Aggregate scaled gain (percent)                      | 33.3     | 37.6                     |
| Closure events                                       | 12,398   | 10,845                   |

*Notes:* The table reports first-absent-wave scaled gains defined in equations (40) and (41). Event-level quantiles give each closure equal weight. Aggregate scaled gains divide the sum of recipient deposit gains by the sum of final active-wave deposits at the corresponding closing branches. Panel A uses observed deposit changes. Panel B converts the change in residualized log deposits into a dollar change using the recipient's final pre-closure deposits. These measures can exceed 100 percent and do not identify the share of customers transferred from the closing branch.

As shown in Table 17, 89.7% of headline recipients record a nonnegative raw deposit change in the first absent wave. The median raw scaled gain is 97.5%, while the ratio of aggregate recipient gains to aggregate closing-branch deposits is 77.9%. After removing the

estimated branch, bank-by-year, and ZIP5-by-year components, the median and aggregate scaled gains are 47.1% and 33.3%. The balanced event window produces similar values.

Neither measure is an identified customer-transfer share. The raw measure also contains unrelated recipient growth and changes in banks’ internal reporting allocations. The adjusted measure removes the fitted residualization components but remains based on reported branch allocations and inherits the residualization specification. Both measures are therefore scale comparisons, not estimates of the fraction of closing-branch customers who move to  $r(e)$ .

## C Supply Estimation Details

This appendix provides the implementation details underlying the supply estimator in Section 5.2. The national estimation procedure estimates branch-level closure policies and state transitions, uses those objects to evaluate close-versus-open deviations by forward simulation, and then estimates the structural payoff coefficients from the resulting discounted value differences.

### C.1 First-stage closure policies

For the strategic banks, the first-stage policy is the fixed-effect logit in equation (28). Its fitted probability,  $\hat{p}_{\ell t}$ , approximates the probability that branch  $\ell$  closes conditional on the observed state. The covariates include branch deposits, nearby same-bank and rival branch counts, local demographics, and indicators for the first three years following an ownership transfer. ZIP and bank-by-year fixed effects absorb persistent local differences and bank-specific changes in closure propensity over time.

A separate auxiliary policy is estimated for branches outside the Top-20 strategic group. These branches remain part of household choice sets and therefore affect deposit reallocation during the national forward simulations. Allowing them to evolve according to a separate policy prevents the continuation environment faced by the strategic banks from being mechanically fixed at its initial network. In the county counterfactual games, by contrast, nonstrategic branches are held fixed within a scenario.

Table 18: First-stage branch-closure policies

|                                | Top-20 policy           | Non-Top-20 policy |
|--------------------------------|-------------------------|-------------------|
| Log branch deposits            | -0.703                  | -0.686            |
| Log rival branches             | -0.125                  | 0.128             |
| Log same-bank branches         | 0.929                   | 0.316             |
| Log weighted age               | 0.905                   | 0.397             |
| Log weighted income            | —                       | -0.173            |
| Acquisition year 1             | 0.482                   | 0.582             |
| Acquisition year 2             | 0.270                   | 0.377             |
| Acquisition year 3             | 0.149                   | 0.392             |
| Full-sample branch-years       | 447,892                 | 710,113           |
| Estimation observations        | 369,597                 | 448,543           |
| Full-sample physical-exit rate | 3.37%                   | 2.75%             |
| Fixed effects                  | ZIP, bank $\times$ year | Year, ZIP, bank   |
| Years                          | 2009–2023               | 2009–2023         |

*Notes:* The table reports point estimates from fixed-effect logit policy approximations. A physical branch exit equals one if the branch identifier is not observed in the following year; a surviving office that changes ownership is not an exit. Rival and same-bank branch counts are measured locally. Weighted age and income aggregate nearby tract demographics using household and inverse-distance weights. Groups without outcome variation are removed. These nuisance-policy estimates generate CCPs for forward simulation and for the expected future private-shock correction.

Table 18 reports the two policy approximations. Deposits reduce the closure index in both samples, while a larger nearby same-bank network and recent acquisition status increase it. The rival-network coefficient differs across the strategic and auxiliary samples. These coefficients are not interpreted as structural payoff parameters. Their role is to generate the CCPs that govern future branch behavior in the forward simulations.

During a forward simulation, the deposit covariate in the policy function is replaced by deposits recomputed from the estimated demand system under the simulated branch network. As a result, a closure changes future policy inputs through endogenous deposit reallocation. Other policy covariates and state variables evolve according to the transition rules described below.

## C.2 State-transition processes

The forward simulations require laws of motion for tract demographics, bank deposit spreads, and the time-varying components of the estimated closure policies. Table 19 reports the estimated scalar parameters. This subsection records the corresponding equations. For each tract-level demographic variable  $x_{ct} \in \{Y_{ct}, \text{Age}_{ct}, H_{ct}\}$ , I estimate an AR(1) process with

tract fixed effects and a linear time trend:

$$x_{ct} = \alpha_c^x + \rho_x x_{c,t-1} + \delta_x(t - t_0^x) + \varepsilon_{ct}^x, \quad \varepsilon_{ct}^x \sim N(0, \sigma_x^2). \quad (42)$$

The variables are estimated in levels: income in dollars, age in years, and households as a count. In the baseline forward simulations, the demographic innovations are drawn using the estimated residual standard deviations.

I estimate the bank-level deposit spread process as

$$m_{bt} = a_m + \rho_m m_{b,t-1} + \alpha_b^m + \lambda_t^m + \varepsilon_{bt}^m, \quad (43)$$

where  $\alpha_b^m$  and  $\lambda_t^m$  are bank and year effects. Whenever an observed bank-year spread is available within the estimation window, the forward simulations use that observed value directly. The estimated process is used to fill missing values and to extend spreads beyond the observed sample.

The common spread year effect is extrapolated using

$$\lambda_t^m = a_\lambda + \phi_\lambda \lambda_{t-1}^m + \xi_t^m, \quad \xi_t^m \sim N(0, \sigma_\lambda^2). \quad (44)$$

Beyond the observed sample, simulated spreads therefore satisfy

$$m_{bt} = \hat{a}_m + \hat{\rho}_m m_{b,t-1} + \hat{\alpha}_b^m + \lambda_t^m. \quad (45)$$

The baseline simulations allow innovations in the common year component.

The time-varying fixed effects entering the closure policies are observed within the estimation sample and extrapolated only when the simulations extend beyond their observed support. For the auxiliary non-Top-20 policy, I estimate an AR(1) for its year effect,

$$\eta_t^{P,\text{aux}} = a_P + \phi_P \eta_{t-1}^{P,\text{aux}} + \xi_t^P. \quad (46)$$

For the Top-20 policy, which contains bank-by-year fixed effects, I estimate a separate process for each strategic bank:

$$\eta_{b,t}^P = a_b + \phi_b \eta_{b,t-1}^P + \xi_{b,t}^P. \quad (47)$$

The residual standard deviations of these auxiliary AR(1) regressions are reported in Tables 19 and 20. In the forward simulations, however, the observed fixed effects are used whenever available and their out-of-sample paths are extrapolated deterministically, setting the future innovations  $\xi_t^P$  and  $\xi_{b,t}^P$  to zero. These extrapolations are forecasting devices used

only to evaluate continuation behavior beyond the observed data window.

Table 19: Supply first-stage transition processes

| Panel                                                 | Process                   | Lag    | Additional | Innovation SD | Fixed effects   |
|-------------------------------------------------------|---------------------------|--------|------------|---------------|-----------------|
| <i>Panel A: Demographic processes</i>                 |                           |        |            |               |                 |
|                                                       | Income                    | 0.7961 | 683.8346   | 5,924.0892    | Tract           |
|                                                       | Age                       | 0.6462 | 0.0551     | 1.9736        | Tract           |
|                                                       | Households                | 0.8551 | 2.5021     | 71.3548       | Tract           |
| <i>Panel B: Deposit and spread processes</i>          |                           |        |            |               |                 |
|                                                       | Branch deposit residual   | 0.8605 | 0.0266     | 0.3059        | Bank, ZIP, year |
|                                                       | Bank-year spread          | 0.5379 | —          | 0.0032        | Bank, year      |
|                                                       | Common spread year effect | 0.4992 | -0.0048    | —             | —               |
| <i>Panel C: Auxiliary closure-policy time process</i> |                           |        |            |               |                 |
|                                                       | Non-Top-20 year effect    | 0.8421 | 0.2223     | 0.1118        | —               |

*Notes:* The additional coefficient is the linear time trend for the demographic processes, the coefficient on the squared positive lag for the branch deposit residual, and the intercept for the two year-effect processes. Income is measured in dollars, age in years, and households as a count. The main supply simulations set the branch deposit residual to zero; its estimated process is reported for completeness. The bank-year spread process also contains bank effects, while its common year effect follows the separately reported AR(1). The demand system used to recompute deposits is reported in Table 5.

Table 20: Top-20 closure-policy bank-year-effect processes

| Bank certificate | 2011 effect | Intercept | Persistence | Innovation SD |
|------------------|-------------|-----------|-------------|---------------|
| 588              | 0.488       | 0.131     | 0.131       | 0.282         |
| 628              | -1.040      | 0.081     | 0.947       | 0.208         |
| 2,270            | -0.314      | 0.059     | 0.299       | 0.296         |
| 3,510            | -0.063      | 0.131     | 0.710       | 0.264         |
| 3,511            | -0.646      | 0.094     | 0.983       | 0.282         |
| 4,977            | -0.167      | 0.076     | -0.518      | 0.513         |
| 6,384            | -0.641      | 0.061     | 0.786       | 0.319         |
| 6,548            | -0.616      | 0.049     | 0.831       | 0.447         |
| 6,560            | -0.457      | -0.046    | 0.567       | 0.505         |
| 6,672            | -1.132      | -0.039    | 0.664       | 0.519         |
| 7,213            | -0.559      | 0.311     | 0.472       | 0.496         |
| 9,846            | -0.473      | 0.098     | 0.776       | 0.390         |
| 11,063           | 0.269       | 0.303     | 0.419       | 0.276         |
| 12,368           | -0.010      | 0.064     | 0.404       | 0.383         |
| 16,571           | -0.681      | -0.369    | -0.203      | 0.509         |
| 17,534           | -1.345      | -0.020    | 0.590       | 0.447         |
| 18,409           | -0.665      | 0.097     | 0.152       | 0.447         |
| 23,220           | -2.036      | -1.435    | -0.121      | 0.477         |
| 32,541           | -0.365      | -0.237    | 0.036       | 0.408         |
| 57,957           | -0.028      | 0.043     | 0.631       | 0.439         |

*Notes:* For each Top-20 bank certificate, the table reports the AR(1) used to extrapolate the estimated bank-by-year closure-policy effect:  $\eta_{g,t}^P = a_g + \phi_g \eta_{g,t-1}^P + \varepsilon_{g,t}$ . The initial effect is the estimated 2011 bank-by-year fixed effect. Innovation standard deviations are calculated from the fitted time series of effects.

Table 19 reports the lag coefficients, additional trend or nonlinear terms, and innovation

standard deviations associated with equations (42)–(46). Table 20 reports the bank-specific parameters in equation (47).

The annual demand coefficients used in the forward simulations retain their estimated values within the demand-estimation window. Beyond 2024, the annual distance coefficient and the common inside-option year effect are held fixed at their final estimated values. The remaining demand parameters are time invariant, while tract demographics continue to evolve according to equation (42).

### C.3 Observed-state forward simulation

The structural sample consists of strategic-bank branch-year decision nodes  $i = (\ell, b, t)$  from 2011 through 2023. For every node, the simulation restarts from the observed beginning-of-period national state in year  $t$ . I evaluate two paths that differ only in the current action of the focal branch. In the open path, the branch is forced to remain open in period  $t$ . In the close path, it is forced to close. This override applies only to the current decision. If the branch remains open, it follows the estimated closure policy from  $t + 1$  onward. If it closes, irreversibility keeps it closed thereafter. All other branches follow their estimated policies in both paths.

The timing follows Section 3. The current closure decision is applied before period- $t$  deposits and flow payoffs are calculated. The demand system is then recomputed on the post-closure network, allowing deposits to move to other branches of the focal bank, competing banks, or the outside option. The resulting deposit allocations become inputs to subsequent closure policies. This procedure means that a focal branch deviation is evaluated through its consequences for the owner’s full simulated portfolio rather than through the deposits of the focal office alone.

Within an observed-state restart, the tract–branch links and branch-neighbor relationships are held at their restart-year geographic configuration, with simulated closures removing branches from those initial links. Observed subsequent ownership transfers are applied when they occur in the data. After the final observed transfer year, no additional ownership changes are introduced. Exogenous state variables evolve according to Appendix C.2.

For payoff component  $k \in \{R, N, A\}$  and current focal action  $a \in \{0, 1\}$ , Section 5.2 defines the simulated discounted value  $F_i^k(a)$  and its close-minus-open difference  $\Delta F_i^k$ . The implementation uses  $T_{\text{sim}} = 20$ , an annual discount factor  $\beta = 0.95$ , and 50 paired simulation draws. The open and close paths use common random numbers for every node.

## C.4 Private-shock continuation correction

Section 3 specifies type-I extreme-value shocks over complete bank closure plans. The national estimator cannot recover CCPs over that exponentially large action space, so the continuation policy is approximated with branch-level binary logits. The private-shock correction used in estimation is constructed consistently with this branch-level CCP approximation.

Consider a binary closure decision with estimated probability  $p$ . After centering the type-I extreme-value shocks by subtracting their common Euler-constant component, the expected private-shock contribution associated with the binary logit policy is the entropy

$$H(p) = -[p \log p + (1 - p) \log(1 - p)], \quad (48)$$

with the convention  $0 \log 0 = 0$ . This expression follows from averaging the conditional expected logit shocks over the two actions. The omitted Euler constant is common to the actions and therefore cancels from the close-versus-open value difference.

Let  $\mathcal{O}_b^C(s)$  denote the set of open branches owned by bank  $b$  that remain eligible for a strategic closure decision in state  $s$ . Under the branch-level continuation approximation, the expected discounted future private-shock contribution following current focal action  $a$  is

$$F_i^\nu(a) = \hat{\mathbb{E}} \left[ \sum_{h=1}^{T_{\text{sim}}-1} \beta^h \sum_{j \in \mathcal{O}_b^C(s_{t+h})} H(\hat{p}_{j,t+h}) \middle| s_t = s_i, e_{\ell t} = a, \hat{\mathbf{p}}, \hat{\mathcal{P}} \right], \quad (49)$$

and the corresponding close-minus-open continuation correction is

$$\Delta F_i^\nu = F_i^\nu(1) - F_i^\nu(0). \quad (50)$$

The sum begins at  $h = 1$ . The current-period action-specific shock is integrated out by the second-stage binary logit itself, so adding an  $h = 0$  entropy term would count the current choice shock twice. Future entropy terms matter because forcing the focal branch to close changes which branches remain open, their future closure probabilities, and therefore the expected private-shock contribution along the continuation path.

Equation (49) should be interpreted as the estimation counterpart to the complete-plan shock structure in Section 3. The theoretical model assigns one shock to each feasible bank closure plan, while the national estimator instead uses the sum of binary-policy entropy terms because continuation behavior is approximated branch by branch.

## C.5 Second-stage pseudo-likelihood and identifying variation

Conditional on the first-stage policies and transition estimates, the second stage estimates the structural payoff coefficients from the actual observed branch-closure outcomes. For observation  $i$ , define

$$\mathcal{I}_i(\boldsymbol{\theta}^\pi) = \Delta F_i^\nu + \theta_R \Delta F_i^R + \theta_N \Delta F_i^N + \theta_A \Delta F_i^A + \chi_{b(i),t} + \chi_{z(i)}, \quad (51)$$

so that

$$\Pr(e_i = 1 \mid \Delta \mathbf{F}_i, \chi_{b(i),t}, \chi_{z(i)}) = \Lambda(\mathcal{I}_i(\boldsymbol{\theta}^\pi)). \quad (52)$$

This is the same specification as equation (32). The coefficient on  $\Delta F_i^\nu$  is normalized to one. Together with the unit scale of the binary logistic shock, this fixes the payoff scale and identifies the coefficients  $\boldsymbol{\theta}^\pi = (\theta_R, \theta_N, \theta_A)'$  relative to that normalization.

The estimator maximizes the corresponding fixed-effect binary-logit pseudo-likelihood,

$$\widehat{\boldsymbol{\theta}}^\pi = \arg \max_{\boldsymbol{\theta}^\pi, \boldsymbol{\chi}} \sum_i \{e_i \log \Lambda(\mathcal{I}_i) + (1 - e_i) \log [1 - \Lambda(\mathcal{I}_i)]\}. \quad (53)$$

The structural sample contains 391,600 branch-year decision nodes. The fixed-effect estimator retains 317,403 observations after removing groups that do not provide the within-group outcome support required by the bank-by-year and ZIP fixed effects. The observed closure rate is 3.54% in the full structural sample and 4.33% in the retained estimation sample.

Identification comes from differences across branch-year nodes in how a current forced closure changes the owner's discounted payoff, after conditioning on the fixed effects. The revenue regressor varies with the focal branch's deposits, the owner's deposit spread, and the amount of demand recaptured by its surviving network. The operating-cost regressor varies with the number of branch-years removed by the current closure. The acquisition regressor varies with the branch's position in the temporary post-transfer window. The future-shock offset varies because the current action changes the set and probabilities of future branch decisions.

The bank-by-year and ZIP effects in equations (51)–(53) are included to absorb common bank-time closure propensity and persistent local differences that are not assigned to the three structural flow-payoff components. The structural estimates reported in Table 7 are therefore interpreted conditional on these controls. Reported standard errors are clustered by physical branch and condition on the estimated first-stage objects and simulated value-difference regressors.

## D Computational Appendix

This appendix describes how the stationary county games that determine the counterfactual policy functions are solved. The county formulation makes it possible to retain the model’s simultaneous bank decisions, post-closure demand reallocation, and bank-level continuation values without constructing the infeasible national action space.

### D.1 Computational state and action spaces

Fix a county and a counterfactual scenario. Let  $\mathcal{L}^S$  be the set of strategic locations, let  $\mathcal{L}_b^S \subseteq \mathcal{L}^S$  collect the locations assigned to bank  $b$ , and let  $L^S = |\mathcal{L}^S|$ . The endogenous network state is the binary vector  $\omega_t \in \{0, 1\}^{L^S}$ , whose elements indicate which strategic locations are open at the beginning of period  $t$ . Nonstrategic competitors are fixed open within a scenario and therefore do not add dimensions to the state. Ownership, branch and tract locations, and tract–branch links are also fixed within a scenario. Changing ownership across counterfactuals changes the partition of strategic locations across banks and the associated payoff and demand shifters, but not the physical set of locations.

The state space specification holds the common demand shifter, bank spreads, and tract income, age, and household counts at their reference values. The average distance coefficient is fixed at its 2011–2024 mean. These objects therefore affect flow payoffs but are not dimensions of the computational state. Deposits are also not an independent state variable in the county game. For every network reached after the banks act, equations (8)–(10) reallocate the tract-level deposit mass across the open branches and the outside option.

When the post-acquisition channel is active, the network state is augmented by a common acquisition-age state  $r_t \in \mathcal{R} = \{0, 1, 2, 3\}$ . Strategic locations whose scenario owner differs from their 2011 owner begin at age zero. The post-acquisition operating cost applies to those locations at ages zero and one, and the transition is deterministic:

$$r_{t+1} = \min\{r_t + 1, 3\}. \quad (54)$$

State 3 is absorbing. In scenarios without this channel,  $\mathcal{R}$  is a singleton. The computational state space is consequently

$$\tilde{\mathcal{S}} = \{0, 1\}^{L^S} \times \mathcal{R}, \quad |\tilde{\mathcal{S}}| = 2^{L^S} |\mathcal{R}|. \quad (55)$$

Enumerating the full network lattice includes states that may be unreachable from the initial network under a particular policy restriction. Keeping these states nevertheless provides a

common closed state space for the Bellman updates.

At state  $\tilde{s} = (\omega, r)$ , bank  $b$  may choose any subset of its currently open strategic locations. A branch-access rule restricts this set by removing closure plans that contain designated protected locations. If  $a = (a_b)_{b \in \mathcal{B}}$  is the simultaneous joint closure decision, the controlled transition is

$$\omega'_\ell = \omega_\ell \mathbf{1}\{\ell \notin a_b\}, \quad \ell \in \mathcal{L}_b^S. \quad (56)$$

Thus the union of the banks' closure sets determines the next network and no closed location can re-enter. Current deposits and operating payoffs are evaluated on this post-decision network. The time-invariant bank and ZIP closure shifter assigned to a location enters once, in the period in which that location closes.

## D.2 Equilibrium fixed point

For each state, I enumerate the feasible closure plans of every strategic bank and their joint combinations. Conditional on a candidate profile of Markov strategies, a bank's choice-specific value averages over rival closure plans using the product of their conditional choice probabilities. It adds the current post-closure payoff to the discounted expected value of the resulting network and next acquisition-age state. This is the finite-state counterpart of equation (17).

Under the type-I extreme-value action-shock specification in Section 3.4, the best response is a multinomial logit over complete bank closure plans. Given choice-specific values  $Q_b(\tilde{s}, a_b)$ , the ex ante value and policy update are

$$\mathcal{V}_b(\tilde{s}) = \sigma_\nu \log \left[ \sum_{a_b \in \mathcal{A}_b(\tilde{s})} \exp \{Q_b(\tilde{s}, a_b)/\sigma_\nu\} \right], \quad (57)$$

$$\mathcal{T}_b(a_b | \tilde{s}) = \frac{\exp \{Q_b(\tilde{s}, a_b)/\sigma_\nu\}}{\sum_{\tilde{a}_b \in \mathcal{A}_b(\tilde{s})} \exp \{Q_b(\tilde{s}, \tilde{a}_b)/\sigma_\nu\}}. \quad (58)$$

The shock scale is normalized to  $\sigma_\nu = 1$ , matching the normalized scale of the supply estimates. For computation, I subtract the common Euler-constant term from the value function. This normalization does not affect action probabilities or value differences. Adding that constant would shift all value-function levels by the same amount but would not change value differences or policies. I retain the log sum implied by the complete-plan shock specification.

A stationary Markov perfect equilibrium is a fixed point at which the candidate strategies used to average over rival actions equal the logit best responses in equation (58), and the value functions satisfy the associated Bellman equations. The annual discount factor is

$\beta = 0.95$ , the same factor used in the supply-estimation forward simulations. The 20-period reporting window truncates the reported paths, not the infinite-horizon Bellman problem.

### D.3 Iteration and construction of expected paths

I compute the fixed point by simultaneous value and policy iteration. The procedure is:

1. Initialize every bank's value function at zero and assign equal probability to all of its feasible actions at each state.
2. Integrate the current value functions over the acquisition-age transition for every possible next-period network.
3. At each state, evaluate all joint closure plans. For each bank and own action, average the post-closure flow payoff and discounted continuation value over rival actions using the current rival policies.
4. Apply equations (57)–(58) to obtain proposed value functions and policies. Values receive the full update. Policies use damping weight 0.55, so the next iterate is 0.45 times the current policy plus 0.55 times the proposed logit policy.
5. Repeat until the largest absolute change in any value is at most  $5 \times 10^{-6}$  and the largest absolute change in any action probability is at most  $5 \times 10^{-7}$ . The iteration limit is 500.

The log sums and choice probabilities are evaluated after subtracting the largest action value within a choice set, a numerically stable transformation that leaves the mathematical objects unchanged. Because dynamic games can have multiple Markov perfect equilibria, the procedure selects the fixed point reached from the zero-value and uniform-policy initialization.

After finding the equilibrium, I construct expected paths by propagating the entire distribution over network and acquisition-age states. The initial distribution places unit mass on the all-open strategic network and, when relevant, acquisition age zero. At each period, the product of the banks' equilibrium action probabilities maps the beginning-of-period distribution into a post-decision network distribution. The acquisition-age transition then produces the next beginning-of-period distribution. Expected branch counts, deposits, access, and flow payoffs are integrals over the post-decision distribution rather than averages over a finite set of simulated paths. Outcomes are reported for periods 0 through 20. Discounted welfare summaries use flow outcomes from periods 0 through 19 at  $\beta = 0.95$ . Period 20 is retained as a terminal network statistic.

## E Counterfactual Details and Robustness

This appendix provides additional results for the Haywood County exercises in Section 6. Appendix D describes the solution of the county games. The results below focus on calibration, the equilibrium paths behind the ownership comparisons, and the sensitivity of the policy welfare calculations.

### E.1 Local calibration

The county exercises begin from the 17 physical locations observed in 2011. Twelve locations are strategic and five nonstrategic locations remain fixed open in the baseline. Simulation year 13 provides the comparison with the observed 2024 network.

Table 21 reports the corresponding calibration diagnostics. The baseline produces 8.39 expected branches in total at year 13, compared with 9 branches in the data. It predicts 3.39 surviving strategic locations, compared with 6 observed survivors. Replacing the five initial nonstrategic locations with the three nonstrategic locations observed in 2024 gives 6.50 expected branches and leaves the strategic-network fit broadly similar. I use the baseline physical network for the counterfactual comparisons.

Table 21: Haywood Network Calibration Diagnostics

Panel A. Network counts

| Network                                     | Strategic | Nonstrategic | Total |
|---------------------------------------------|-----------|--------------|-------|
| Observed 2011 network                       | 12.00     | 5.00         | 17.00 |
| Observed 2024 network                       | 6.00      | 3.00         | 9.00  |
| Baseline simulated network at year 13       | 3.39      | 5.00         | 8.39  |
| Terminal-nonstrategic robustness at year 13 | 3.50      | 3.00         | 6.50  |

Panel B. Strategic branch-status fit

| Network treatment                           | Expected survivors | Brier score | Deposit RMSE, all locations | Deposit RMSE, observed survivors |
|---------------------------------------------|--------------------|-------------|-----------------------------|----------------------------------|
| Baseline simulated network at year 13       | 3.39               | 0.285       | 76.85                       | 70.24                            |
| Terminal-nonstrategic robustness at year 13 | 3.50               | 0.280       | 76.32                       | 69.45                            |

Notes: Simulation year 13 is the 2024-equivalent horizon. Panel B evaluates the 12 strategic locations present in 2011; the observed number of strategic survivors in 2024 is 6. Deposit RMSE is measured in millions of dollars.

## E.2 Deposit-rate sensitivity and monetary welfare calibration

The baseline demand model identifies how changes in the branch network affect deposit allocation and the value of households' choice sets, but it does not include deposit rates and therefore does not estimate a coefficient that converts utility into monetary units. I calibrate this coefficient by requiring the estimated demand system to reproduce an external branch-level deposit-rate elasticity. This follows the elasticity-matching approach in [Ho et al. \(2026\)](#), who calibrate an otherwise weakly identified price coefficient to an elasticity from outside their estimated demand model.

Let  $p_{c\ell}$  denote the fitted probability that deposits from tract  $c$  are allocated to branch  $\ell$  under the estimated demand model, and let  $M_c$  denote the tract's potential deposit mass. Predicted deposits at branch  $\ell$  are

$$D_\ell = \sum_{c:\ell \in J_c} M_c p_{c\ell}. \quad (59)$$

For the monetary calibration, let  $r_\ell^{pp}$  denote the annual deposit rate at branch  $\ell$ , measured in percentage points, and let  $\alpha^{pp}$  denote the utility effect of a one-percentage-point increase in that rate. Consider adding the term  $\alpha^{pp} r_\ell^{pp}$  to the utility of branch  $\ell$ , holding the baseline demand estimates fixed. Under the logit demand system,

$$\frac{\partial p_{c\ell}}{\partial r_\ell^{pp}} = \alpha^{pp} p_{c\ell} (1 - p_{c\ell}). \quad (60)$$

It follows that the own-rate elasticity of predicted deposits at branch  $\ell$  is

$$\varepsilon_\ell^r = \frac{r_\ell^{pp}}{D_\ell} \frac{\partial D_\ell}{\partial r_\ell^{pp}} = \alpha^{pp} r_\ell^{pp} \frac{\sum_{c:\ell \in J_c} M_c p_{c\ell} (1 - p_{c\ell})}{D_\ell}. \quad (61)$$

I use 2015 because [Pakes et al. \(2024\)](#) estimate deposit demand for that year and report an average branch own-rate elasticity of 1.61 using the annual one-year CD rate. I evaluate equation (61) using the constructed 2015 12-month CD rates described in Appendix A.5 and choose  $\alpha^{pp}$  so that the unweighted average elasticity across eligible branch-years equals this target:

$$\frac{1}{N_{2015}} \sum_{\ell \in \mathcal{L}_{2015}} \varepsilon_\ell^r = 1.61. \quad (62)$$

Because equation (61) is linear in  $\alpha^{pp}$ , the calibration has the closed-form solution

$$\alpha^{pp} = \frac{1.61}{\frac{1}{N_{2015}} \sum_{\ell \in \mathcal{L}_{2015}} r_{\ell}^{pp} \frac{\sum_{c: \ell \in J_c} M_c p_{cl} (1 - p_{cl})}{D_{\ell}}} = 9.008. \quad (63)$$

The calibrated coefficient is used only to express changes in the demand inclusive value in monetary units. The calibration does not change the estimated demand parameters, the supply estimates, or the equilibrium branch-network counterfactuals.

### E.3 Ownership paths

Figure 14 reports the expected network paths behind the ownership experiments in Section 6. The temporary post-acquisition cost has its largest effect early in the path. The difference associated with common portfolio control persists after that cost has largely expired. The common-ownership stress test generates additional contraction throughout the horizon.




Figure 14: Expected branch-network paths in the ownership experiments

*Notes.* Lines report expected post-decision operating branches over simulation years 0–20. Panel (a) compares the 2011 ownership configuration, final ownership without the temporary acquisition cost, and the baseline. Panel (b) compares the baseline with the common-ownership stress test. Both panels use the same vertical scale.

## E.4 No-closure benchmark

Figure 15 reports the producer-side components of the no-closure benchmark together with consumer surplus and modeled welfare. It complements the targeted-policy decomposition in Figure 10 and illustrates the cost of preserving the full initial network.




Figure 15: Welfare decomposition under the no-closure benchmark

*Notes.* The figure reports 20-year discounted changes relative to the baseline in millions of dollars. The no-closure experiment keeps all 17 initial locations open. Consumer surplus uses the central deposit-rate sensitivity calibration. Producer-side components use the revenue-equivalent monetary scale described in Section 6.