

Financial Supermarkets?

Cross-Selling in Retail Banking

Shengmao Cao

Lulu Wang *

May 16, 2026

Abstract

We study how consumers' tendency to hold multiple financial products from the same bank shapes competition and policy in U.S. retail banking. We document pervasive cross-holding across deposits, credit cards, auto loans, mortgages, and brokerage services, especially at smaller banks. Combining an original consumer survey on search and switching with an equilibrium model of bank choice under incomplete consideration, we identify three forces behind cross-holding: correlated preferences across banks, consideration spillovers from banks' cross-selling efforts, and utility complementarities between same-bank products. Together, the consideration spillovers and the utility complementarities account for a substantial share of bank franchise value, with reward credit cards in particular functioning as a loss-leader product whose relationship value far exceeds its standalone margin. Accounting for these synergies qualitatively changes standard policy conclusions. A narrow-bank entrant captures considerably less of the deposit market than a model without complementarity would predict. Interchange-fee regulation that depresses credit card demand spills over into deposit competition. And a horizontal merger that appears anticompetitive on a deposit-HHI screen can become pro-competitive once the merged bank uses its credit card as a loss leader for a larger depositor base.

*Cao: Northwestern Kellogg (email: shengmao.cao@kellogg.northwestern.edu). Wang: Northwestern Kellogg (email: lulu.wang@kellogg.northwestern.edu). We thank participants at the Midwest IO Fest and IIOC for their helpful feedback. Marziyeh Aghaei, Lucia Bertolotti, Keigo Makino, Rodrigo Surraco, and Tongtong Yao provided excellent research assistance. All errors are our own.

I Introduction

Cross-selling — the promotion of additional financial products to existing customers — is a cornerstone of modern retail banking. The 1998 merger of Travelers Group and Citicorp sought to create a “financial supermarket”, a one-stop shop for all consumer financial needs. By 2014, Chase’s retail banking division was generating approximately \$4 billion in gross revenue through cross-selling, with roughly half arising from synergies between its credit card and consumer banking operations (JPM, 2015). The practice drew widespread criticism after the 2016 Wells Fargo fake account scandal, yet it has remained a defining feature of the retail banking landscape.

Despite this practical significance, academic research and regulatory policy have largely approached each banking product market in isolation. The literature on bank franchise value, for example, focuses primarily on deposit market power (Egan et al., 2017; Drechsler et al., 2021). The literature on credit card regulation examines how caps on interest rates and interchange fees affect consumers’ use of credit (Agarwal et al., 2015; Wang, 2023). Bank merger reviews, following longstanding federal guidelines, similarly assess competitive harm primarily through deposit market concentration alone (U.S. Department of Justice, 1995; U.S. Department of Justice and Federal Trade Commission, 2023). These studies offer valuable insights but share a common blind spot: how regulatory changes or shifts in competitive structure within one product market spill over to others.

In this paper, we examine how consumers build portfolios of financial products and assess the implications of cross-selling for competition and policy in retail banking. First, using three U.S. consumer surveys, we document empirical patterns in consumers’ cross-product holdings — including deposit accounts, credit cards, mortgages, auto loans, and investment accounts — across banks. Second, we develop and estimate a model of how consumers choose deposit accounts and credit cards over time. Third, we use the estimated model to derive new insights on the sources of banks’ franchise value, policy spillovers between deposit and credit card markets, and the competitive effects of bank mergers.

Our analysis leverages three consumer surveys: the Survey of Consumer Finances (SCF), the MRI-Simmons Ultimate Survey of Americans, and an original survey of U.S. adults conducted on the Prolific platform. The first two are large, nationally representative studies that document consumers’ financial product portfolios across banks in detail. Our Prolific survey complements these with data on consumers’ search processes, account-switching histories, and responses to hypothetical choice scenarios.

Our central fact is that many consumers choose banks, not individual products. Across the SCF, MRI-Simmons, and Prolific surveys, roughly 35% of consumers with one checking account and one credit card hold both at the same bank — more than five times the 6.4% benchmark implied by independent product choice. The pattern is especially pronounced at smaller banks. At megabanks like Chase and Bank of America, depositors are on average 6 times more likely to hold the bank’s credit card than non-depositors; at regional banks such as PNC or Citizens, this ratio exceeds 60. We document analogous patterns for auto loans, mortgages, and brokerage accounts, with cross-holding rates of 15%, 18%, and 8% respectively — all well above the corresponding independent-choice benchmarks.

Additional cross-sectional patterns point away from a purely informational explanation of cross-holding. The relationship-banking literature emphasizes how deposit accounts generate information about borrower risk, enabling banks to offer better prices or more favorable underwriting to existing customers (Puri et al., 2017; Agarwal et al., 2018; Parlour et al., 2022). Yet we find similar cross-holding rates among high-credit-score consumers who use credit cards as a payment instrument and among low-credit-score consumers who revolve balances. The information channel applies most naturally to the latter group but cannot account for cross-holding among high-credit-score transactors, who face minimal underwriting barriers at any lender. These patterns therefore point to preference heterogeneity and product complementarity, beyond asymmetric information, as the primary drivers of cross-holding.

While these cross-sectional patterns are striking, they cannot on their own distinguish between the mechanisms driving cross-holding. Consumers may simply have correlated preferences across products — a consumer drawn to Chase’s branch network and reputation may prefer it for both deposits and credit cards, generating cross-holding without any causal connection between the two decisions. Alternatively, cross-holding may reflect causal channels: banks’ cross-selling efforts may funnel existing customers into considering additional products, and consumers may derive direct utility from consolidating their banking products at a single institution. These mechanisms have fundamentally different implications for competition and regulatory policy (Gentzkow, 2007). Under correlated preferences alone, cross-holding is a form of consumer sorting and carries no strategic implications. When causal mechanisms are operative, the strategic landscape changes: banks find it profitable to use loss leaders to generate demand for complementary products, and consumers’ tendency to bundle insulates incumbents against entry and transmits competitive or regulatory shocks from one product market to another.

We designed our survey to gather direct evidence on each of these mechanisms. For each respondent’s most recent primary deposit account and credit card opening, we reconstruct the search process: the prior accounts they held, the banks they considered, exposure to direct-mail advertising, and the factors they found most important in the decision. We complement this with a set of discrete choice experiments in which respondents indicate how they would react to hypothetical price changes for a focal product and — for cross-holders — how they would respond if their bank closed their account for the complementary product. In each scenario, respondents report whether they would remain with their current bank or switch.

The survey responses point to all three mechanisms. Consumers who consider a bank for a deposit account also tend to consider the same bank for a credit card — even when they ultimately chose neither — indicating that consideration probabilities are correlated across products. Consumers who hold a product at a bank are considerably more likely to include that bank in their consideration set for the other product, showing that existing banking relationships shape the search process. Many cross-holders report that they would switch their focal product if their bank closed the complementary account, indicating that consumers derive direct utility from holding both products at the same institution.

We then develop an empirical model of consumer choice over deposit accounts and credit cards in which consumers stochastically receive opportunities to switch one product at a time. When a switching opportunity arrives, the consumer forms a consideration set (Goeree, 2008) and selects the utility-maximizing bank. The model embeds all three mechanisms: correlated bank preferences, consideration spillovers from existing product holdings, and direct utility complementarity between products held at the same institution. Banks set prices to maximize the present discounted value of future profits in a Nash equilibrium, internalizing the cross-product spillovers through which competitive pricing in one market helps attract customers in the other. We model two consumer types: transactors, who pay their credit card balance in full each month and optimize over the rewards rate, and revolvers, who carry balances and optimize over the APR (Drechsler et al., 2025). We estimate demand via a Monte Carlo EM algorithm (Train, 2008), recovering consumer-level heterogeneity in attentiveness, bank preferences, price sensitivity, and taste for cross-holding.

The estimated model lets us reassess bank franchise value and the role of complementarity in regulation and competition. Credit cards prove far more important for franchise value than accounting profits suggest. Although direct credit-card profits account for roughly 40% of the combined deposit-and-credit-card franchise, eliminating

credit cards would cost banks around 60% of their total profits. Most of this gap is driven by rewards credit cards: even though banks roughly break even on these cards, rewards cards help retain high-income depositors at the bank, making them a loss-leader for the higher-margin deposit franchise. For many large banks, rewards-card margins would be one-third higher were it not for this loss-leader strategy.

The same complementarity also generates spillovers across markets in policy counterfactuals. A narrow bank offering only deposits captures a meaningful deposit share, but considerably less than a model without complementarity would predict: consumers who value cross-holding stay with their multi-product bank even at the entrant's higher deposit rate. Entry also influences the credit card market, as incumbents who lose depositors also shed credit share. Similarly, interchange-fee regulation that depresses credit-card demand carries significant deposit-market spillovers, because banks rely on credit card relationships to retain depositors. Chase, whose retail franchise is among the most exposed to cross-selling, loses deposit share even as it raises its deposit rate to compensate.

Finally, we apply our framework to a horizontal merger between Citi and the super-regional bank to show how complementarity can flip the competitive verdict. Without complementarity, joint ownership increases concentration and raises prices — the standard anticompetitive prediction. Once complementarity is restored, the merger becomes profitable for a different reason: each firm's credit card can now serve as a loss leader for the other's deposit base, expanding the cross-selling scope that neither firm could exploit alone. The combined firm makes credit-card rewards substantially more attractive — using the card as a loss leader across its expanded portfolio — while letting deposit rates drift slightly lower to harvest surplus from its captive base. Rivals respond by raising their own deposit rates — the opposite of what a concentration analysis predicts. The full Nash response is invisible to a standard deposit-HHI screen: the competitive action runs through credit-card pricing and rivals' defensive responses on deposits, not through deposit concentration alone.

This paper contributes to several strands of literature. We show that cross-product complementarity is a significant source of bank franchise value and market power, extending recent work on the industrial organization of financial markets. This literature has largely studied individual product markets in isolation — deposit accounts (Egan et al., 2017; Honka et al., 2017), credit cards (Stango and Zinman, 2016; Nelson, 2018; Cuesta and Sepúlveda, 2021; Galenianos and Gavazza, 2022; Wang, 2023; Drechsler et al., 2025), and wealth management (Hortaçsu and Syverson, 2004; Hastings et al., 2017; Brown et al., 2023). We complement this work by developing an equilibrium model of

multi-product bank competition that reveals how cross-selling shapes banks’ strategic incentives and market power.

We provide additional evidence for demand-based theories of economies of scope in retail banking. Most studies on cross-selling in finance have focused on economies of scope in corporate banking (Hellmann et al., 2008; Laux and Walz, 2009; Ivashina and Kovner, 2011; Santikian, 2014; Qi, 2024). In retail banking, Puri et al. (2017); Agarwal et al. (2018); Parlour et al. (2022) argue that information from deposit accounts helps lenders mitigate credit risks, providing one foundation for why depository institutions make loans. We contribute a complementary foundation: cross-holding rates are strikingly similar across consumers with very different credit profiles, pointing to consumer search behavior and preferences as an additional driver. Basten and Juelsrud (2023) and Aguirregabiria et al. (2024) document related scope economies in retail banking. We build on this literature by separately identifying the mechanisms behind cross-holding — correlated preferences, consideration spillovers, and utility complementarity — and tracing their equilibrium implications for pricing, entry, and regulatory policy.

Our paper also contributes to the literature on the use of survey data to identify demand models (Petrin, 2002; Berry et al., 2004; Grieco et al., 2024). We innovate by showing how questions about past product changes can help identify demand even in markets featuring switching costs and choice frictions. Our purpose-built survey collects consumers’ consideration sets, search histories, and stated responses to hypothetical shocks — moments that directly inform the structural mechanisms of interest. Although these survey data carry well-known limitations, this approach offers a valuable complement to revealed-preference methods, uncovering dimensions of consumer financial behavior that aggregate market shares alone cannot identify.

More broadly, our paper relates to the literature on search and bundling. On the search side, we characterize how consumers’ narrow consideration sets and infrequent reoptimization shape bank market power and the returns to cross-selling, connecting our work to the literature on limited consideration (Goeree, 2008; Honka et al., 2017) and consumer inattention in financial markets (Egan et al., 2025). On the bundling side, while explicit bundle discounts are uncommon in U.S. retail banking, cross-selling strategies and product complementarity generate similar patterns of cross-holding. We show that consumers’ decision to “bundle” creates significant entry barriers for single-product firms (Whinston, 1989; Nalebuff, 2004).

The rest of the paper is organized as follows. Section II describes the data. Section III presents stylized facts on cross-holding behavior using three complementary surveys of U.S. consumers. Section IV introduces an equilibrium model of consumer choice and

bank competition in the deposit account and credit card markets. Section V estimates the model. Section VI presents counterfactual exercises. Section VII concludes.

II Data

We combine three complementary datasets to describe and analyze household choices of financial products.

Survey of Consumer Finances (SCF) Our first dataset is the 2022 Survey of Consumer Finances (SCF), the latest wave of the triennial survey conducted by the Federal Reserve Board. The SCF is widely regarded as the most comprehensive source on the finances of U.S. households. We use the 2022 Full Public Data Set, which includes detailed information on the financial positions and demographic characteristics of 4,595 households.

The SCF documents ownership of a broad set of financial products, including deposit accounts, credit cards, auto loans, mortgages, and brokerage services. Although bank anonymity has limited prior retail banking research,¹ we exploit a key structural detail: the SCF assigns consistent institution codes within households across different products. This allows us to identify whether a household holds multiple products at the same institution (e.g., cross-holding of a checking account and a mortgage). One important limitation is that we cannot measure cross-holding patterns at the level of specific banks.

MRI-Simmons USA (MRI) Our second dataset, the MRI-Simmons USA Study (MRI), complements the SCF by identifying the specific financial institutions consumers use. A large-scale, nationally representative survey, MRI provides detailed product-ownership records alongside demographic data. We primarily use the 2021 and 2022 waves (approximately 100,000 consumers); we additionally draw on the 2020 wave (referred to as NHCS, approximately 40,000 consumers), which captures consumer tenure at their financial institutions. MRI’s product coverage is narrower than the SCF’s, focused on deposit accounts, credit cards, and brokerage services, but it overcomes the SCF’s anonymity restriction. 9 banks in the data offer both deposit accounts and credit cards: Chase, Bank of America, Wells Fargo, Citi, Capital One, U.S. Bank, PNC, TD Bank, and Citizens. Among these, Chase, Bank of America, and Wells Fargo also offer brokerage services.

Two important limitations remain in both datasets. First, as they are repeated cross-sections, we observe static product portfolios rather than consumers’ search or switching

¹The survey anonymizes financial institutions, preventing direct identification of which bank a household uses.

behavior over time. Second, they capture consumers’ cross-holding decisions in equilibrium but not banks’ cross-selling strategies that shape these patterns. Our Prolific survey is designed to address both gaps.

Original Prolific Survey To overcome these limitations, we fielded an original survey of approximately 955 U.S. consumers on Prolific. The survey takes around ten minutes, and each participant receives \$3 upon completion. The survey mirrors the MRI structure by collecting data on all deposit and credit card holdings by each consumer, distinguishing between primary and other accounts. As shown in Table 2, our sample tracks aggregate market shares in the deposit and credit card markets reasonably well, although it modestly overrepresents customers of online-focused banks such as American Express, Capital One, and Discover.

The survey introduces two novel features designed to better understand consumer choices. First, we reconstruct consumers’ search and switching history. For each respondent’s most recent primary account opening, we elicit the timing of the decision, their consideration set (the set of banks they looked into or reviewed), prior account choices, and exposure to direct-mail advertising.

Second, we use a stated-preference design to collect information on substitution patterns and complementarities. We ask consumers to identify their ‘ideal’ bank for a specific product if they were to make a choice *de novo* today. Differences between this ideal and their current choice indicate inattention or switching costs. We also elicit their ‘second-best’ choice if their ideal bank were not available. Finally, to measure cross-product complementarities, we present respondents with hypothetical shocks, such as an increase in competitors’ deposit interest rates, a reduction in credit-card rewards, and, for cross-holders, the closing of one account by their bank. These questions resemble the “discrete choice experiment” commonly used to estimate the value of non-wage amenities and labor-supply elasticities (Mas and Pallais, 2017, 2019; Caldwell et al., 2025). The resulting switching probabilities provide key empirical moments needed to estimate the magnitude of demand-side complementarities.

We carefully design our survey instrument to mitigate potential framing and order effects.² For example, to minimize framing, we elicit the respondent’s consideration set before asking about their prior banking relationships. This ordering ensures that consideration-set reports are not contaminated by the respondent’s recall of their prior banks. Similarly, we present the product discontinuation shock prior to the price shocks.

²The survey can be accessed at the following link: https://kellogg.qualtrics.com/jfe/form/SV_eXKObVgGMnzCcGa.

This ensures that responses to the “unbundling shock” are not anchored by the price sensitivity elicited in the subsequent questions.³

Table 1 compares key characteristics across our data sources. Panel A shows that median incomes are broadly comparable across sources, ranging from \$62k in our Prolific sample to \$71k in the SCF; our Prolific sample somewhat overrepresents college-educated consumers, consistent with known selection patterns on online survey platforms. Product ownership rates in Panel B are similar across datasets. Panels C1 and C2 reveal substantial differences in deposit balances and credit card spending between revolvers and transactors, a pattern we exploit in calibrating our supply-side model.

Additional Data To complete the supply-side analysis, we collect price data from several additional sources. We combine financial filings from banks and payment volume data from the Nilson Report to estimate average rewards expense by bank. We use the CFPB Terms of Credit Card Plans (TCCP) survey to obtain average credit card APRs by bank. We use the quarterly Y-14 bank holding company filings to measure average bank deposit rates as the ratio of interest expense on domestic deposits to total domestic deposits.

III Stylized Facts

We document four stylized facts on the prevalence of cross-holding and three potential channels — correlated preferences, consideration spillovers, and utility complementarity — through which cross-holding arises.

We measure cross-holding using two complementary approaches. Our *aggregate cross-holding rate* captures the fraction of consumers who hold a checking account and a second financial product at the same bank. To compute this rate, we restrict the sample to consumers with exactly one checking account and exactly one secondary product (e.g., one credit card), so that the definition of cross-holding is unambiguous.⁴ Our *bank-specific conditional share* measures the probability of holding product B at a specific named bank j , conditional on holding product A at the bank. We compute this measure using the full sample, without restrictions on the number of accounts, in both MRI and our

³Any potential anchoring effects on price sensitivity can be assessed by comparing the responses of cross-holders to those of non-cross-holders, who were not presented with the unbundling shock.

⁴In MRI and our Prolific survey, banks outside a set of named institutions are grouped into a residual “Other” category. We exclude consumers for whom both products are assigned to “Other,” because it is not clear whether they are cross-holding at the same bank. This exclusion affects a nontrivial share of consumers in MRI and Prolific, where a meaningful fraction of respondents hold products at banks outside the set of named institutions. The SCF does not face this issue: although it anonymizes bank names, it assigns consistent institution codes within each household, so we can determine whether two products are held at the same institution regardless of which bank it is.

Prolific survey; it cannot be computed from the SCF, where institution identities are anonymous.

Fact 1. *Cross-holding is pervasive across all demographic groups and represents a fivefold increase over a benchmark of independent product choices. Reliance on cross-holding is significantly greater among smaller, regional banks.*

Figure 1 documents a consistent pattern of within-bank bundling across deposit accounts, credit cards, auto loans, mortgages, and brokerage accounts using the aggregate cross-holding rate computed from all three datasets (SCF, MRI, and Prolific). Approximately 35% of consumers with exactly one checking account and one credit card hold both at the same bank, while the analogous rates for auto loans, mortgages, and brokerage accounts are 15, 18, and 8 percent, respectively.⁵ To contextualize the magnitude of cross-holding, we construct a benchmark of independent choices based on aggregate market shares. If consumers selected deposit accounts and credit cards independently, only 6.4% would hold both at the same bank. The observed 35% cross-holding rate is thus more than five times the level implied by independent choices.

The propensity to cross-hold is remarkably stable across demographic groups. Figure 2 shows nearly identical cross-holding rates across different levels of income, education, and credit scores, as well as between credit card transactors and revolvers, replicated across all three datasets. This uniformity helps distinguish our findings from the “information channel” emphasized in prior literature (Puri et al., 2017; Agarwal et al., 2018; Parlour et al., 2022) — the hypothesis that banks use deposit data to better assess credit risks and lend to depositors who would face higher underwriting standards elsewhere. We find that cross-holding is just as prevalent among high-credit-score individuals who face minimal underwriting barriers, and among “transactors” who optimize over publicly posted reward rates. This suggests the decision to bundle reflects preference heterogeneity or product complementarity, rather than asymmetric information alone.

Cross-holding is also associated with lower financial sophistication. Using the MRI sample, we regress an indicator for cross-holding on consumer characteristics (Table 3). Column (1) restricts to consumers with exactly one checking account and one credit card, while Columns (2) and (3) use the full sample of consumers with at least one of each. Consumers in specialized financial roles, such as CFOs and treasurers, and those who hold stocks or tax-advantaged investments are substantially less likely to cross-hold.

⁵The absolute level of the cross-holding rate varies somewhat across datasets because of differences in bank coverage and the share of consumers assigned to unidentified institutions, but the qualitative pattern is consistent.

This pattern reveals heterogeneity in how consumers weigh the convenience of one-stop banking against shopping across providers: financially sophisticated consumers are more likely to mix and match across banks, while others place greater value on consolidating products at a single institution.

Finally, we use the bank-specific conditional shares to document that smaller banks are disproportionately dependent on cross-holding. Figure 3 plots each bank’s credit card market share separately among its depositors and non-depositors, using the full MRI sample. We define the *cross-holding ratio* for bank j as the ratio of these two conditional shares: the probability of holding a credit card at bank j given a deposit account there, divided by the same probability given no deposit account at bank j . For “mega-banks” like Chase and Bank of America, the cross-holding ratio averages roughly $6\times$ ($4\times$ for Chase, $8\times$ for Bank of America). For regional banks such as PNC or Citizens, this ratio exceeds 60. For these smaller banks, credit card adoption is almost entirely driven by the deposit relationship, while megabanks attract a meaningful share of standalone credit card customers. We document a symmetric pattern for deposit account market shares conditional on credit card ownership in Figure 3. These results suggest that the “relationship multiplier” is central to customer acquisition at smaller banks, even though their thin credit card franchises limit the absolute dollars involved compared to those at the megabanks.

Fact 2. *Consumers do not continuously optimize their bank choices and switch banks infrequently.*

Figure 4 shows the distribution of consumer tenures at their current banks, drawn from our Prolific survey.⁶ The median tenure for a primary deposit account is 8 years. About 46% of consumers have been using the same bank for over a decade, while only 2.3% switched their deposit account in the past year. The median tenure for a primary credit card is 6 years, and 37% have maintained the card for at least 10 years. Only 4.5% of consumers opened their primary deposit account and credit card simultaneously, suggesting that consumers typically do not search for or acquire these products at a single point in time.

The long tenures suggest that banks retain consumers for extended periods once they are acquired. To gauge whether this stickiness reflects true satisfaction or underlying frictions, we asked respondents to identify their “ideal” bank for a deposit account or a credit card if they were to make a choice *de novo* today. For deposit accounts, 9.9% of

⁶We ask each respondent to recall when they opened their primary accounts. Because of imperfect recall and rounding, we observe masses at 5, 10, and 15 years in the reported tenure distributions.

respondents prefer a bank different from their current provider, more than four times the 2.3% who actually switched accounts in the past year. We find a similar wedge for credit cards, where 10.7% of consumers identify a different “ideal” issuer, compared with only 2.5% who obtained a new primary credit card in the past year. These large gaps between consumers’ stated preferences and actual choices suggest that infrequent switching partly reflects consumer inattention or switching costs (Egan et al., 2025).

The remaining two facts draw on our Prolific survey, which we designed to recover the search and complementarity mechanisms that cross-sectional shares cannot identify on their own. For each respondent’s most recent primary deposit account and credit card, we record the prior bank, the set of banks they “looked into or reviewed” during the search, exposure to direct-mail advertising, and the factors they viewed as most important in the choice. We then elicit stated responses to a sequence of hypothetical scenarios at their current bank: changes in deposit rates, credit card rewards, and credit card APRs, and, for cross-holders, the closing of the complementary product.

Fact 3. *Consumers are significantly more likely to consider their current bank when searching for a new product.*

Banks’ cross-selling efforts are designed to funnel existing customers into additional product lines, potentially generating the patterns of cross-holding documented above. To identify this mechanism, we use our Prolific survey, in which respondents listed the specific banks they “looked into or reviewed” during their most recent account search. By matching these consideration sets with the consumers’ pre-existing banking relationships, we quantify the impact of current holdings on the search process.

Consideration sets are narrow: consumers considered an average of only 2.0 banks for deposit accounts and 2.5 banks for credit cards (Appendix Figure A.3). Table 4 shows clear evidence of relationship spillovers: owning one product from a bank increases the likelihood of considering that bank for a new product by 6 to 11 percentage points. These relationship spillovers expand the overall consideration set rather than shifting consideration away from rivals. For example, consumers with an existing deposit account at a bank offering credit cards considered an average of 2.4 credit card products, compared to 2.1 among non-holders.

Table 4 also documents the role of correlated brand preferences. Conditional on current holdings, consumers who consider a bank for one product are approximately 25 percentage points more likely to consider it for the other. This strong co-consideration, likely driven by brand-level advertising or general reputation, indicates a high degree of correlation in consumers’ consideration sets across markets.

Fact 4. *Cross-holding consumers exhibit high switching probabilities in response to forced unbundling, indicating substantial perceived complementarities from bundling.*

We use the stated-preference module of our Prolific survey to elicit consumers’ responses to a “discontinuation shock”: we ask cross-holders whether they would switch their primary bank for one product if the bank closed their account for the other. We benchmark these responses against sensitivity to standard price shocks, such as changes in deposit rates, credit card reward rates, and credit card interest rates.

Figure 5 summarizes the results. Approximately 32% of cross-holders would switch their deposit account if their bank closed their credit card. This fraction is comparable to the 40% of transactors and 27% of revolvers who would switch in response to a 50-basis-point increase in competitors’ savings account interest rates. Similarly, 36% of cross-holders state they would switch credit cards if their bank closed their deposit account. This switching probability can be benchmarked against 56% of transactors switching if reward rates fell by half a percentage point, and 72% of revolvers switching if credit card interest rates increased by 5 percentage points.

If cross-holding reflected only prior cross-selling or correlated preferences, losing one product should leave demand for the other unchanged. The large fraction of switchers following the discontinuation shock — around one-third of cross-holders — rules this out. These responses provide direct evidence that the convenience of one-stop service is a first-order driver of cross-holding.

IV Model

A dynamic model of consumer choice over deposits and credit cards separates the three mechanisms behind cross-holding: correlated preferences, increased consideration, and true utility benefits. The three mechanisms have sharply different implications for competition: correlated preferences create neither loss-leader pricing incentives nor entry barriers, increased consideration generates loss-leader pricing but not entry barriers, and χ^u generates both. Separating them is therefore essential for the counterfactuals.

IV.A Demand

Consumer i ’s state is the pair $x_i = (x_{i1}, x_{i2})$ of current bank choices across the deposit ($k = 1$) and credit card ($k = 2$) markets. Let j denote a bank, $k \in \{1, 2\}$ a product category, and t time.

Consideration and Choices Motivated by Fact 2, we assume that consumers are usually inattentive and do not continuously optimize their choices. Instead, opportunities

for consumer i to reconsider product k arrive stochastically according to a Poisson process with arrival rate λ_{ik} . We emphasize that λ_{ik} captures the rate at which already-adopted product choices are revisited; the timing of a consumer's *initial* adoption of product k is taken as given and is not structurally modeled. These reconsideration opportunities arrive independently across products, consistent with our survey evidence that consumers typically search for and acquire these products sequentially rather than simultaneously. This independent arrival structure implies that when an opportunity to switch one product arises, the consumer's choice of the other product remains fixed, acting as a state variable that influences both the set of banks considered and the utility derived from the choice.

When a switching opportunity for product k arrives at time t , the consumer forms a consideration set $C_{ikt} \subset \mathcal{J}_k$ (Honka et al., 2017). The probability that consumer i considers a specific set of banks C_{ikt} , given their current state x_i , is the product of individual consideration probabilities:

$$P_{ik}(C_{ikt}|x_i) = \prod_{j \in C_{ikt}} \phi_{ijk}(x_i) \cdot \prod_{j \notin C_{ikt}} (1 - \phi_{ijk}(x_i)), \quad (1)$$

where $\phi_{ijk}(x_i)$ is the probability that bank j is included in the consideration set for product k , conditional on x_i . Consumers automatically include their current bank in the consideration set. For other banks, the probability of consideration follows a logistic form:

$$\phi_{ijk}(x_i) = \begin{cases} 1 & x_{ik} = j \\ \frac{1}{1 + \exp(-z_{ijk}(x_i))} & x_{ik} \neq j, \end{cases}$$

where the latent consideration index $z_{ijk}(x_i)$ is defined as:

$$z_{ijk}(x_i) = v_{ijk}^c + \chi_k^c \mathbb{1}(x_{i,-k} = j). \quad (2)$$

This index $z_{ijk}(x_i)$ captures two distinct forces. The first term, v_{ijk}^c , is a consumer-specific visibility of bank j 's product k . The second term, $\chi_k^c \mathbb{1}(x_{i,-k} = j)$, captures cross-selling effects. If a consumer holds product $-k$ at bank j , this increases the probability of considering bank j for product k (Fact 3). Consistent with our survey evidence, we assume cross-selling expands consideration sets rather than crowding out rivals.

Given a consideration set, consumer i chooses bank $j \in C_{ikt}$ to maximize utility. The utility for bank j 's product k (given state x_i) is defined as:

$$u_{ijkt}(x_i) = \underbrace{v_{ijk}^u + \alpha_{ik}p_{jk} + \chi_{ik}^u \mathbb{1}(x_{i,-k} = j)}_{\delta_{ijk}(x_i)} + \varepsilon_{ijkt}, \quad (3)$$

where v_{ijk}^u represents the consumer-specific quality of bank j 's product k , and $\alpha_{ik}p_{jk}$ captures price sensitivity. In this setting, p_{jk} is signed so that consumers prefer higher values of p_{jk} . This is natural for the deposit interest rate for a deposit account, or the reward rate for credit card transactors. For credit card borrowers, we define price as $-\text{APR}$. Appendix Figure A.6 reports the observed deposit rates, reward rates, and APRs we target across banks. Under this parameterization, $\alpha_{ik} > 0$.⁷ The term $\chi_{ik}^u \mathbb{1}(x_{i,-k} = j)$ measures the perceived complementarity from holding both products at the same bank, as documented in Fact 4.⁸ The idiosyncratic match value ε_{ijkt} is assumed to follow a Type 1 Extreme Value distribution.

Unobserved Heterogeneity We allow consumer-level heterogeneity in the bank-preference intercepts v_{ij}^c, v_{ij}^u , the price sensitivities α_{ik} , the utility complementarities χ_{ik}^u , and the arrival rates λ_{ik} , collected in a consumer-specific type vector θ_i . The consideration complementarity χ_k^c is treated as homogeneous, since the data identify its mean (Fact 3) but not a separate consumer-level variance. We defer the parametric form of the type distribution to Section V.A.

The conditional probability that consumer i chooses bank j from a consideration set C_{ik} is:

$$s_{ijk}(x_i, C_{ik}) = \mathbb{1}\{j \in C_{ik}\} \times \frac{\exp(\delta_{ijk}(x_i))}{1 + \sum_{r \in C_{ik} \setminus \{0\}} \exp(\delta_{irk}(x_i))}. \quad (4)$$

The unconditional choice probability, $m_{ijk}(x_i)$ is the weighted average across all possible consideration sets:

$$m_{ijk}(x_i) = \sum_{C_{ik} \in 2^{\mathcal{J}_k}} P_{ik}(C_{ik}|x_i) s_{ijk}(x_i, C_{ik}) \quad (5)$$

Steady-State Market Shares The choice processes described above generate dynamic cross-product linkages through spillovers across product lines. A bank's success in one

⁷The model describes a specific type of credit card consumer. We separately estimate the model for transactors, who pay their balance in full each month and focus on the reward rate, and revolvers, who carry a balance and are more concerned about interest rates.

⁸This term may reflect genuine complementarity, such as the convenience of a unified mobile app, or misperceptions induced by banks' cross-selling efforts. In the latter case, this term affects decision utility without conferring real benefits, thereby introducing choice distortions (Allcott, 2013).

market depends on its presence in the other, via both cross-selling effects and product complementarities. We characterize each consumer's long-run holdings using a continuous-time Markov framework (Arcidiacono et al., 2016). This approach simplifies the high-dimensional transition space by assuming that consumers reconsider one product at a time (Fact 2).

For each product k , let $Q_{ik}(\theta_i)$ be consumer i 's transition intensity matrix, whose (x, y) th element $q_{ik}^{x \rightarrow y}$ gives the rate at which consumer i moves from state x to state y on product k . Given the consumer-specific arrival rate λ_{ik} and conditional choice probabilities $m_{ijk}(x)$, the transition rates are:

$$q_{ik}^{x \rightarrow y} = \begin{cases} 0, & \text{if } x_{-k} \neq y_{-k}, \\ \lambda_{ik} m_{iy_k k}(x), & \text{if } x_k \neq y_k \text{ and } x_{-k} = y_{-k}, \\ -\lambda_{ik} (1 - m_{ix_k k}(x)), & \text{if } x = y. \end{cases} \quad (6)$$

The first case reflects our assumption that the switching opportunities arise for only one product at a time. When the opportunity to switch product k arrives, there are no transitions between states that differ on the other product $-k$. The second case describes the transition rate from bank x_k to y_k . The diagonal elements $q_{ik}^{x \rightarrow x}$ represent the total rate of transitioning out of the current state, which equals the negative of the sum of transition probabilities to all other states.

Under the assumption of independent Poisson arrivals of switching opportunities, the aggregate transition intensity matrix for consumer i is the sum of product-specific intensities: $Q_i(\theta_i) = Q_{i1}(\theta_i) + Q_{i2}(\theta_i)$. Consumer i 's steady-state holding distribution $\bar{s}_i(\theta_i)$ is the unique row vector satisfying

$$\bar{s}_i(\theta_i) Q_i(\theta_i) = 0, \quad \sum_x \bar{s}_i(x; \theta_i) = 1. \quad (7)$$

To obtain aggregate market shares, average over the population distribution of types: $\bar{s}(\bar{\theta}, \Sigma_\theta) = \mathbb{E}_\theta[\bar{s}_i(\theta)]$.

We can also derive the transition probabilities over finite intervals from the model. For any time interval t , consumer i 's transition matrix is the matrix exponential

$$S_i(t) = \exp(t \cdot Q_i) \equiv \sum_{k=0}^{\infty} \frac{(t Q_i)^k}{k!}. \quad (8)$$

In Appendix B, we solve a simple example with two symmetric banks and cross-

holding complementarity. We show that it is straightforward to compute the choice probabilities, the transition intensity matrix, and the equilibrium market shares under this framework.

IV.B Supply

In our setting, banks' pricing decisions are distinctive in two ways. First, pricing decisions generate significant spillovers across product lines. A competitive price for one product, such as a high deposit rate, attracts new customers who may subsequently adopt the other product. As a result, the marginal value of a new customer exceeds the flow profits from any single product. Second, consumer inattention implies that demand responds only slowly to price changes. This consumer inertia grants banks temporary market power, allowing them to widen margins before existing "sleepy" customers respond.

Forward-looking banks set prices to maximize the present discounted value (PDV) of future profits across both product lines, taking the competitors' prices as given. Because consumers are heterogeneous, the bank's profit is the expectation of the discounted flow payoff across the population distribution of types:

$$\max_{p_j} \Pi_j(p_j, p_{-j}) = \mathbb{E}_\theta \left[\int_0^\infty e^{-rt} \bar{s}_i(\theta) e^{Q_i(p_j, p_{-j}; \theta)t} h_j(\theta) dt \right]. \quad (9)$$

For each type θ , the inner expression is the PDV of flow profits from a consumer whose initial distribution is the type-specific steady state $\bar{s}_i(\theta)$ and whose holdings evolve according to $Q_i(p_j, p_{-j}; \theta)$. The flow profit vector is $h_j(\theta) = \sum_{k=1}^2 O_{jk} b_k \pi_{jk}$, where O_{jk} identifies the states in which a consumer holds bank j 's product k , π_{jk} is the per-unit profit margin, and b_k is the average account balance.⁹ The balances b_k do not depend on θ because consumer types θ are drawn within a given transactor/revolver segment, and we calibrate balances at the segment level.

Each type's inner integral admits a closed-form solution via the *resolvent* of its generator matrix. At discount rate r , the type-specific resolvent is

$$\mathcal{R}(p_j, p_{-j}; \theta) \equiv (rI - Q_i(p_j, p_{-j}; \theta))^{-1} = \int_0^\infty e^{-rt} e^{Q_i(p_j, p_{-j}; \theta)t} dt, \quad (10)$$

⁹We calibrate each account balance parameter using the simple average of median values from independent survey sources (SCF, NHCS, and MRI), separately by consumer type. Appendix C provides details of the calibration procedure. The margin π_{jk} is $r_{jk} - p_{jk}$ for deposit accounts and credit card transactors, where r_{jk} represents investment return on deposits and the interchange fee revenue, respectively. For credit card revolvers, the margin is $(-p_{jk}) - c_{jk} = \text{APR} - c_{jk}$, consistent with the demand-side convention $p_{jk} = -\text{APR}$, where c_{jk} is the bank's cost of funds from overhead and default risk.

which is well-defined because each Q_i is a generator matrix with non-positive real eigenvalues and $r > 0$. Applying this identity, the profit function simplifies to an expectation over types:

$$\Pi_j(p_j, p_{-j}) = \mathbb{E}_\theta [\bar{s}_i(\theta) \mathcal{R}(p_j, p_{-j}; \theta) h_j(\theta)] . \quad (11)$$

For any type θ , the resolvent $\mathcal{R}(\cdot; \theta) = (rI - Q_i(\theta))^{-1}$ converts the instantaneous transition rates in Q_i into the discounted expected time each consumer spends in each banking state. The population profit in Equation 11 averages these type-specific lifetime values.

Our supply-side model is “essentially static”, following Brown et al. (2023). While banks internalize that market shares evolve slowly as consumers respond to price changes with a lag, they do not set prices period by period based on the current state. Instead, banks set long-run optimal prices, recognizing that demand is slow to respond due to consumer inattention. Lower arrival rates of switching opportunities imply slower adjustments, which grant banks greater market power over their existing “sleepy” customer base.

First-Order Conditions The resolvent representation yields clean first-order conditions for pricing. Differentiating Equation 11 with respect to p_{jk} requires the derivative of each type’s resolvent. Applying the matrix inverse identity $\partial A^{-1} = -A^{-1}(\partial A)A^{-1}$ to $\mathcal{R}(\cdot; \theta) = (rI - Q_i(\theta))^{-1}$ gives:¹⁰

$$\frac{\partial \mathcal{R}(\theta)}{\partial p_{jk}} = \mathcal{R}(\theta) \frac{\partial Q_i(\theta)}{\partial p_{jk}} \mathcal{R}(\theta). \quad (12)$$

Since $\partial h_j / \partial p_{jk} = -O_{jk} b_k$ for all products (this holds for revolvers because $\pi_{jk} = -p_{jk} - c_{jk}$ under our sign convention), the first-order condition $\partial \Pi_j / \partial p_{jk} = 0$ takes an expectation over types on both sides:

$$\underbrace{\mathbb{E}_\theta \left[\bar{s}_i(\theta) \mathcal{R}(\theta) \frac{\partial Q_i(\theta)}{\partial p_{jk}} \mathcal{R}(\theta) h_j(\theta) \right]}_{\text{Marginal benefit of gained customers}} = \underbrace{\mathbb{E}_\theta [\bar{s}_i(\theta) \mathcal{R}(\theta) O_{jk} b_k]}_{\text{Infra-marginal cost}} . \quad (13)$$

Banks balance the long-run benefit of attracting customers against the cost of lower margins on the existing customer base. Recall that p_{jk} is defined in consumer-favorable units (deposit rate, rewards rate, or negative APR), so raising p_{jk} makes the bank more attractive but reduces per-customer profits. The left-hand side of Equation 13 is the marginal benefit. A more generous price alters the transition rates, drawing consumers

¹⁰The two negative signs (one from the inverse identity, one from $\partial(rI - Q_i) = -\partial Q_i$) cancel out.

into bank j 's products and raising the discounted lifetime flow of profits across both product lines. Because h_j aggregates profits from both products, banks internalize that gaining a deposit customer today also raises the future probability of earning credit card margins from that same individual, and the cross-product value amplifies the incentive to compete on gateway products. The right-hand side is the infra-marginal cost. The bank sacrifices b_k dollars of margin on every existing holder of product k , and because switching is stochastic and infrequent, there are many such infra-marginal consumers.

In the estimation, we evaluate the outer expectation $\mathbb{E}_\theta[\cdot]$ by a Monte Carlo average over M draws from the estimated type distribution, $M^{-1} \sum_{m=1}^M \bar{s}_i(\theta_m) \mathcal{R}(\theta_m) (\cdots)$. Section V details this computation.

In the empirical implementation, we impose separate credit card FOCs for transactors (reward-rate) and revolvers (APR), because the two segments respond to different price dimensions. The deposit rate FOC pools across segments, since deposit pricing enters both groups' utility symmetrically. Section V gives the computational details.

V Estimation and Results

Estimation translates the survey facts of Section III into quantitative estimates of all three cross-holding mechanisms: correlated preferences, increased consideration, and true utility complementary. Consistent with the reduced-form facts, the resulting estimates suggest that all three mechanisms are active. We then invert the demand model to recover banks' costs. The resulting estimates of margins are similar to past work, and our model also matches untargeted moments on how bank tenures vary across consumers.

V.A Parametric Restrictions on Heterogeneity

The bank-product preferences are jointly normal in levels: for each bank j ,

$$(v_{ij1}^c, v_{ij2}^c, v_{ij1}^u, v_{ij2}^u) \sim \mathcal{N}(\bar{v}_j, \Sigma_v),$$

where Σ_v is shared across banks and allows correlations within a bank across products and across the consideration and utility margins. Because price sensitivities, complementarities, and arrival rates are positive on theoretical grounds, we model them in logs. For each product k , $(\log \alpha_{ik}, \log \chi_{ik}^u)$ is jointly normal, with within-product correlation between price sensitivity and the utility benefit of cross-holding. The two product- k blocks $(\log \alpha_{i1}, \log \chi_{i1}^u)$ and $(\log \alpha_{i2}, \log \chi_{i2}^u)$ are independent of each other, because the Prolific survey randomizes each respondent into hypotheticals for only one product, so the cross-product joint distribution is not identified. Arrival rates are jointly normal across products in logs, $(\log \lambda_{i1}, \log \lambda_{i2}) \sim \mathcal{N}(\log \bar{\lambda}, \Sigma_\lambda)$.

This heterogeneity plays two roles in the estimation. First, the bank-preference covariance Σ_v accommodates correlated tastes for a bank's products. A consumer who particularly values Chase's checking account is also more likely to value its credit card, so Σ_v helps match the level of cross-holding observed in the data without attributing all of it to the causal cross-selling channels χ^c and χ^u . Second, the dispersion in (α_i, χ_i^u) generates a non-degenerate distribution of preferences required to rationalize the hypothetical shock responses described below. Without heterogeneity, the mixture over price-dominant and unbundling-dominant types would assign zero probability to some observed response patterns.

V.B Intuition for Identification

The three cross-holding mechanisms map directly onto distinct features of the survey data. Correlated bank preferences Σ_v are pinned down by joint co-consideration of the same bank's products in the Prolific consideration sets. The consideration spillover χ^c is identified by how often a consumer's existing bank shows up in the consideration set for the other product, after netting out the correlated-preference component. The utility complementarity χ^u is identified by switching responses to the discontinuation shock, which holds bank identity fixed and varies only the bundle, severing the link to Σ_v and χ^c .

Identification of the Population Means The parameter $\bar{\chi}_k^c$ (consideration complementarity) is identified from the cross-holding and consideration-set patterns reported in Fact 3; $\bar{\chi}_k^u$ (utility complementarity) is identified from unbundling-shock responses in Fact 4. The mean price sensitivity $\bar{\alpha}_k$ is identified from stated responses to hypothetical price shocks. The visibility intercepts $\bar{v}_j^c$ are identified by how often each bank's product appears in the consideration sets reported in the Prolific survey. The utility intercepts $\bar{v}_j^u$ are identified by the probabilities that each product is chosen conditional on being considered. Finally, $\bar{\lambda}_k$ is identified by combining the average observed tenure with the average would-switch rate from the Prolific survey. Tenure measures how long consumers stay, and the would-switch rate isolates the fraction of wake-ups that trigger a switch; together they pin down $\bar{\lambda}_k$.

Identification of Unobserved Heterogeneity Conditional on population means, the survey identifies preference distribution across consumers because we observe multiple decisions per respondent.

First, the diagonal elements of Σ_v are identified by how often consumers would switch banks when awakened. The Prolific survey elicits each respondent's ideal bank

alongside her current holdings, so we observe the share of awakened consumers who rank a bank other than their current one at the top. If Σ_v were large, almost every awakened consumer would already be at her preferred bank and the would-switch rate would be small. Substantial dispersion in bank preferences is the only way to reconcile a low would-switch rate with the observed steady-state holdings.

Second, given Σ_v , the empirical tenure distribution identifies arrival-rate dispersion around $\bar{\lambda}_k$. The model's tenure distribution is a mixture of exponentials where the wake-up rate and post-wake-up switch probability determine the shape. Once we pin down switch probability from Σ_v , the shape of the tenure distribution directly reveals λ_{ik} and its heterogeneity.

Third, for each product k , the joint distribution of $(\log \alpha_{ik}, \log \chi_{ik}^u)$ is identified from each respondent's stated responses to the hypothetical price and unbundling shocks, together with the rate at which consumers switch to the outside option. A consumer who abandons her ideal bank under a small price shock reveals a large α_{ik} . One who abandons it under unbundling reveals a large χ_{ik}^u . Because each respondent answers both types of shock questions for the same product, the joint distribution of her stated responses traces out the within-product correlation between α_{ik} and χ_{ik}^u , for instance whether price-sensitive consumers tend to place lower or higher value on cross-holding.

Sample selection complicates the unbundling shock. By design, we pose the unbundling question only to consumers who currently cross-hold both products at the same bank. Cross-holders are exactly the consumers with high realized χ_i^u , so their stated responses overstate the population mean of χ_i^u if treated naively. Heterogeneity is what corrects this. The cross-holders' stated responses identify the conditional distribution $f(\chi_i^u \mid \text{cross-hold})$, while the cross-holding rates and past-switching transitions in Prolific identify the population distribution of types that produces the observed selection into cross-holding. Combining the two pins down the unconditional distribution of χ_i^u , recovering both its mean and its variance without inheriting the cross-holder upward bias.

V.C Likelihood Function

We estimate the population mean and covariance of the heterogeneity distribution by maximum likelihood, integrating over the consumer-specific parameters θ_i :

$$\max_{\bar{\theta}, \Sigma_{\theta}} \sum_{i=1}^N \log \int L(X_i \mid \theta) \phi(\theta; \bar{\theta}, \Sigma_{\theta}) d\theta,$$

where $\phi(\cdot; \bar{\theta}, \Sigma_{\theta})$ is the multivariate normal density specified in Section V.A and $L(X_i | \theta_i)$ is the complete-data likelihood for consumer i evaluated at her type θ_i . Section V.D describes the Monte Carlo EM algorithm we use to solve this optimization problem.

The complete-data likelihood combines three blocks of evidence about each Prolific respondent. The static-choices block models the consideration sets and bank choices reported at her most recent product switches. The stated-preference block models her responses to hypothetical price and unbundling shocks. The dynamic block wraps the static choices in arrival rates and matrix-exponential terms that govern how the consumer transitions between states.

Static choices (Prolific historical decisions) The Prolific survey reconstructs each respondent’s most recent deposit-account opening and credit-card opening: the bank she chose, the set of banks she “looked into or reviewed” in the search, and her holdings at the time of each switch. For each switch k , we model the joint probability of her observed consideration set C_k and chosen bank x_k^* given the pre-switch state $x(t_k^-)$:

$$D_k(\theta_i) = \underbrace{P(C_k | x(t_k^-), \theta_i)}_{\text{consideration-set probability}} \cdot \underbrace{P(x_k^* | C_k, x(t_k^-), \theta_i)}_{\text{logit choice given consideration}}. \quad (14)$$

The consideration-set probability is the product of bank-specific inclusion probabilities $\phi_{ijk}(x_i)$ defined in Section IV; the choice probability is the conditional logit over banks in the consideration set. The static block contributes

$$L^{\text{stat}}(\theta_i) = D_1(\theta_i) \cdot D_2(\theta_i),$$

two snapshots of the consumer’s preferences at her two most recent decision moments. Each snapshot conditions on a different pre-switch state; that variation identifies how cross-holding shifts the consideration set and the conditional choice (Facts 3–4).

Stated responses (Prolific hypothetical questions) For each Prolific respondent we elicit the ideal bank V , the second-choice bank A , and the stated bank choice under two hypothetical shocks at the ideal bank: an unfavorable price change (W) and the loss of the complementary product (R). We interpret these answers as arising from a single draw of latent Gumbel preference shocks at the time of the question, so that the four answers form a consistent ranking of banks under successively more severe shocks. The contribution to the likelihood is then the joint probability of the four responses,

marginalized over the unobserved consideration set C :

$$L^{\text{hyp}}(\theta_i) = \sum_{C \in 2^{\mathcal{J}}} P(V, W, R, A \mid C, \theta_i) P(C \mid x^*, \theta_i). \quad (15)$$

The conditional probability $P(V, W, R, A \mid C, \theta_i)$ is the difference of choice probabilities under the four progressively-degraded utility states (baseline, the unbundling shock, the price shock, full removal of the focal product). Whether a particular consumer falls into a regime in which the price shock dominates the unbundling shock depends on the realized $(\alpha_{ik}, \chi_{ik}^u)$ ratio, so the population-marginalized likelihood is a mixture across the two regimes weighted by the heterogeneity distribution. The mixture assigns positive probability to every observed response pattern, including patterns that would otherwise look “inconsistent” under any single (α, χ^u) pair. Appendix D writes out the regime-conditional formulas.

Dynamics (Prolific past switching) The static block describes the choice events themselves; the dynamic block adds the timing layer that turns those events into a coherent stochastic process for each consumer. Specifically, the Prolific survey records the elapsed time since each switch (t_1 since the deposit switch, t_2 since the credit-card switch, with $t_1 > t_2$), and we model this timing using the continuous-time Markov framework from Section IV. The dynamic contribution is

$$L^{\text{dyn}}(\theta_i) = \underbrace{\lambda_{i1}}_{\text{deposit-switch arrival rate}} \times \underbrace{M}_{\text{intervening tenure}} \times \underbrace{\lambda_{i2}}_{\text{credit-switch arrival rate}} \times \underbrace{F}_{\text{final stability}}. \quad (16)$$

Each λ_{ik} is the Poisson rate at which the consumer reconsiders product k (the wake-up rate), M is the matrix-exponential probability of the consumer’s state evolution during the interval (t_2, t_1) between the two switches, and F is the matrix-exponential probability of remaining in the current state from the most recent switch through the survey date. Closed-form expressions for M and F , the case adjustments for first-time adoption (where the relevant friction differs from reconsideration), and the closed-form integrals that handle censored switch times beyond the survey’s lookback window are derived in Appendix D.

Non-cardholders The stated-preference and dynamic blocks both require a credit-card holding. The stated-preference module asks consumers to evaluate price and unbundling shocks at a bank where they currently hold both products, and the dynamic block needs

a credit-card switch event to time. Prolific respondents who do not currently hold a credit card therefore enter only through a simplified version of the static block: a single choice probability that the consumer holds the outside option on credit cards conditional on her checking-account holding, marginalized over the unobserved credit-card consideration set. Concretely, the non-cardholder static contribution is

$$L^{\text{stat,nc}}(\theta_i) = \sum_{C \in 2^{\mathcal{J}}} P(x_2^* = 0 \mid C, x_1^*, \theta_i) P(C \mid x_1^*, \theta_i),$$

which disciplines how attractive the credit-card outside option must be relative to the inside options. Non-cardholders contribute neither L^{hyp} nor L^{dyn} , because the survey did not elicit hypothetical responses or switching histories from this group.

Combined likelihood For a cardholder, the complete-data log-likelihood factors into the three blocks:

$$\log L^{\text{card}}(X_i \mid \theta_i) = \log L^{\text{stat}}(\theta_i) + \log L^{\text{hyp}}(\theta_i) + \log L^{\text{dyn}}(\theta_i).$$

Combining cardholders and non-cardholders, the full log-likelihood for observation i is

$$\log L(X_i \mid \theta_i) = \mathbb{1}\{i \in \text{cardholders}\} \log L^{\text{card}}(\theta_i) + \mathbb{1}\{i \in \text{non-cardholders}\} \log L^{\text{stat,nc}}(\theta_i). \quad (17)$$

The three cardholder blocks operate on different timescales but share the consumer-level θ_i , which is what allows information from each block to flow into the others through the population heterogeneity distribution.

V.D Algorithm

We maximize the marginal likelihood by Monte Carlo EM, following Train (2008) with one practical modification. The E-step forms a posterior over each consumer’s type θ_i by drawing from a proposal distribution and reweighting by importance ratios; the M-step updates the population mean and covariance $(\bar{\theta}, \Sigma_{\theta})$ as weighted first and second moments of the simulated draws. We depart from Train (2008) in one detail: rather than drawing all types from the population prior, we draw each consumer’s types from her own posterior summary from the previous iteration. This observation-specific proposal sharply reduces the variance of the importance weights, since each draw is already concentrated in the region that the previous iteration judged likely for that consumer, and lets the E-step converge with a manageable number of draws per observation. Appendix D.8 writes out the full pseudocode along with the parametric block structure im-

posed on Σ_θ , the phased estimation schedule (static moments first, dynamic transitions added later), the adaptive proposal covariance, and the effective-sample-size criterion that triggers re-simulation.

Calibrating bank intercepts to aggregate shares The Prolific sample identifies consumer-level heterogeneity well but is too small to pin down each bank’s mean utility precisely. After the MC-EM iterations converge, we refine the bank-specific intercepts $\bar{v}_{jk}^u$ to match aggregate market shares observed in the larger MRI-Simmons cross-section. Holding the heterogeneity, complementarity, and switching-rate parameters fixed at their EM estimates, we iterate $\bar{v}_{jk}^u$ until the model-implied steady-state share for each bank-product cell, averaged over the estimated type distribution, equals the corresponding MRI share. The refinement adjusts only the bank-specific levels; the consumer-level parameters identified from Prolific are unchanged.

V.E Estimation Results

The estimates reveal a striking asymmetry in cross-product complementarity: revolvers value holding both products at the same bank at a price-equivalent of 183 basis points of APR, while transactors value it at only 24 basis points of rewards. Combined with the thin margins on transactor credit cards, this gap supports the gateway-product rationale, in which rewards cards subsidize deposit acquisition even when they generate little direct profit. We estimate the model separately for credit card transactors and revolvers; Table 5 presents the population means of the structural parameters, Table 7 the bank-specific preference intercepts, and Table 6 the estimated standard deviations and correlations of Σ_θ .¹¹

The arrival rate λ_{ik} conflates genuine inattention with switching costs, since a consumer who reconsiders but faces high fixed costs looks observationally identical to one who never reconsiders. We interpret λ_{ik} as a reduced-form measure of the total friction limiting reoptimization. The counterfactual implications differ across the two interpretations. A policy that reduces switching costs, such as portable account numbers, would raise effective λ_{ik} under the switching-cost reading but not under pure inattention.

Price Sensitivity The price sensitivity parameters $\bar{\alpha}_k$ represent the semi-elasticity of choice with respect to a one-percentage-point change in rates. For checking accounts, transactors and revolvers exhibit elevated sensitivities ($\bar{\alpha}_1^{\text{transactor}} = 2.25$, $\bar{\alpha}_1^{\text{revolver}} = 1.68$), implying that a one-percentage-point increase in the deposit rate raises the probability

¹¹Standard errors are omitted from the structural tables; parametric bootstrap inference is under development.

of choosing a bank by roughly 1.7–2.3%, conditional on consideration. Our estimates lie in the upper range of the empirical literature on deposit demand, which spans semi-elasticities of roughly 0.8–2.3. Reported values include 1.12 in d’Avernas et al. (2023), 1.39 (insured) and 2.26 (uninsured) in Koont (2023), 0.97 in Wang et al. (2022), 0.78 in Dick (2008), and approximately 2.10 in Honka et al. (2017). Even within a single consumer segment, the estimated heterogeneity in α_{ik} is substantial. The log-scale standard deviation of α_{i1} is 1.04 for transactors and 1.40 for revolvers, so the population distribution of deposit-rate responsiveness is broad relative to its mean.

The two segments differ sharply on credit cards. Transactors, who pay in full and optimize over rewards, are highly responsive ($\bar{\alpha}_2^{\text{transactor}} = 2.71$). Revolvers, who borrow and face higher switching barriers (including credit-risk-based underwriting at alternative issuers), are much less sensitive to APR ($\bar{\alpha}_2^{\text{revolver}} = 0.49$), a roughly fivefold difference. The estimated variance of α_{i2} is small for both segments, so the difference is mean-level rather than compositional selection within a broad distribution.

Cross-Selling Effects Cross-selling shifts consumer demand through two channels. First, we find substantial consideration complementarity. For transactors and revolvers alike, an existing credit card relationship significantly increases the likelihood of considering the same bank for a checking account ($\bar{\chi}_1^{c,\text{transactor}} = 1.23$, $\bar{\chi}_1^{c,\text{revolver}} = 0.85$). The impact of the deposit relationship on credit card consideration is stronger, with $\bar{\chi}_2^{c,\text{transactor}} = 1.53$ and $\bar{\chi}_2^{c,\text{revolver}} = 1.02$. The boost is meaningful but insufficient on its own to elevate the low-visibility super-regional bank to the consideration level of a credit-card monoline. For transactors, adding the +1.53 boost to the super-regional bank’s baseline credit card consideration index of -3.42 still falls short of Capital One’s baseline of -1.19 , and revolvers face a similar pattern. Even where the mean boost does not fully close the gap, cross-selling expands the super-regional bank’s effective reach.¹²

Second, consumers derive substantial benefits from consolidating both products at the same bank. For checking accounts, the utility boost from cross-holding is equivalent to a 35–43 basis-point increase in the deposit rate, computed as $\bar{\chi}_1^u / \bar{\alpha}_1$. For credit cards, the price-equivalent value is 24 basis points of rewards for transactors and a substantially larger 183 basis points of APR for revolvers. The revolver gap likely reflects that revolvers, who may face higher underwriting barriers elsewhere or possess lower finan-

¹²Our estimates are conservative with respect to the “Other” bank aggregate, which pools all banks outside the set of named institutions. By construction, $\chi^c = \chi^u = 0$ for the Other aggregate, since cross-holding at two different unidentified banks is indistinguishable from cross-holding at the same bank. This assumption biases the estimated complementarity parameters downward for the named banks, as the model attributes some of their cross-holding advantage to the baseline rather than to the structural parameters.

cial sophistication, perceive much higher relative utility from opening a credit card at their primary deposit bank. These implied benefits are averages over the estimated log-normal distribution of χ_{ik}^u , which has meaningful dispersion ($\sigma(\log \chi_1^u) = 1.00$ for transactor checking, $\sigma(\log \chi_2^u) = 0.37$ for revolver credit), so a nontrivial share of consumers either derive little complementarity benefit or place a very high value on consolidation.

Correlated Preferences An important question is whether the cross-holding mechanisms reflect genuine causal complementarity or correlated preferences, since consumers who like a bank may both consider and choose its products more. The estimated bank-preference covariance Σ_ν in Table 6 points to very strong correlated preferences. The within-bank correlation between checking and credit consideration, $\rho(\nu_1^c, \nu_2^c)$, is 0.96 for both transactors and revolvers, and the analogous utility correlation $\rho(\nu_1^u, \nu_2^u)$ is 0.65 for transactors and 0.15 for revolvers.¹³ These correlations absorb most of the raw association between considering and choosing a bank’s products. Identification of χ_k^u is protected from this concern by the discontinuation shock (Fact 4), which holds bank identity fixed and varies only the product bundle offered. A consumer who would switch her deposit account upon losing her credit card at the same bank reveals a complementarity that cannot be explained by correlated tastes alone. Identification of χ_k^c is harder to protect, but because we treat χ_k^c as a homogeneous parameter and let ν^c absorb the consumer-level dispersion in consideration tastes, the joint consideration-and-choice patterns in the Prolific survey allow the MC-EM estimator to identify $\bar{\chi}_k^c$ separately from the bank-preference covariance Σ_ν .

Switching Frictions Consumers reconsider their products infrequently. The estimated average arrival rates imply reconsideration horizons of about 4.5 to 3.3 years for checking accounts ($\bar{\lambda}_1^{\text{transactor}} = 0.22$, $\bar{\lambda}_1^{\text{revolver}} = 0.30$) and 4.8 to 3.6 years for credit cards ($\bar{\lambda}_2^{\text{transactor}} = 0.21$, $\bar{\lambda}_2^{\text{revolver}} = 0.28$). The model attributes some of the low observed switching rate to low- λ_{ik} types, so the population mean $\bar{\lambda}$ can sit above the average reconsideration horizon implied by aggregate switching counts while still matching the population-average switching hazard. The estimated log-scale standard deviation of λ_{ik} is small (0.10–0.16), so most of the dispersion in observed switching frequencies is driven by heterogeneity in α_i and χ_i^u rather than in λ_i .

What an analyst ignoring complementarity would conclude Table 5 compares full and restricted ($\chi^c = \chi^u = 0$) specifications. When complementarity is shut down,

¹³The standard deviations of the ν block are economically large, ranging from 1.33 to 2.01 in log-odds units across transactor and revolver specifications.

the model must attribute observed cross-holding entirely to correlated bank preferences ν_{ijk}^u , so to match the same cross-holding levels it inflates the bank-specific utility intercepts. Higher intercepts reduce the implied responsiveness to price and, through the supply-side FOCs, lower the recovered marginal cost of attracting deposit customers. The restricted specification is not a nuisance benchmark. Section VI shows that it misses deposit–credit spillovers in every counterfactual by construction, and the magnitude of the miss represents the policy-relevant bias from modeling retail banking markets as two independent products.

V.F Goodness of Fit

Demand and Tenure Distributions To assess the model’s performance, we compare predicted market shares and cross-holding patterns against observed data at the bank level. Figure 6 shows that our model closely replicates the market shares for both checking accounts and credit cards, capturing the dominance of national players and the long tail of smaller banks. Figure 7 evaluates the model’s performance in matching cross-holding rates. The strong within-bank correlations in bank-specific preferences ($\rho(\nu_1^c, \nu_2^c) = 0.96$) allow the model to match both the cross-bank variation in cross-holding intensity and the aggregate level of cross-holding.

Figure 8 shows that the model reproduces the empirical tenure distributions from our survey and the NHCS, correctly matching the heavy right tail of long-tenured consumers. Although heterogeneity in λ_{ik} is small (Table 6), heterogeneity in α_i and χ_i^u generates a distribution of type-specific hazards that produces the observed tenure dispersion. Appendix D.6 describes how we derive these model-implied tenure distributions from the estimated demand parameters by averaging phase-type distributions over the type distribution.

Appendix Figure A.5 checks that the model reproduces tenure patterns across the life cycle rather than only in aggregate. We simulate a panel of synthetic consumers under the estimated model, assigning each a first-account arrival age drawn uniformly from a fixed window and then running the estimated switching dynamics forward to their current age. The model matches the monotone age gradient in survey tenure for both products, indicating that our estimated switching dynamics are consistent with the observed life-cycle patterns even though λ_{ik} is estimated only from the behavior of existing account holders.

V.G Supply-Side Cost Recovery and Margin Validation

The counterfactuals in Section VI require marginal cost estimates so that banks can re-optimize prices in each new equilibrium. Given the demand-side parameter esti-

mates, we recover marginal costs by inverting the supply-side first-order conditions (Equation 13) from banks' profit maximization problem. Even under consumer heterogeneity, the first-order conditions remain *linear* in costs, so cost recovery reduces to solving a small linear system for each bank.

The linearity survives because for every consumer type θ , bank j 's flow payoff vector is affine in its own costs: $h_j(\theta) = h_j^0(\theta) - c_{j1} O_{j1} b_1 - c_{j2} O_{j2} b_2$, where $h_j^0(\theta)$ collects the price- and return-dependent terms and c_{jk} are the unknown marginal costs. Each type-specific resolvent $\mathcal{R}(\theta) = (rI - Q_i(\theta))^{-1}$ and its price derivatives depend only on demand parameters and observed prices, not on costs. Taking the expectation over types in Equation 13 preserves linearity in (c_{j1}, c_{j2}) , so the FOCs aggregate to a linear system whose coefficients are expectations over θ . We evaluate these expectations by a Monte Carlo average over M draws from the estimated type distribution, reusing the $\{(\bar{s}_i(\theta_m), \mathcal{R}(\theta_m))\}_{m=1}^M$ already computed on the demand side. Because each bank's flow payoffs involve only its own costs, the system decouples across banks. In the pooled model with transactors and revolvers, each bank has three cost parameters (checking c_{check} , transactor credit c_{trans} , and revolver credit c_{rev}) and three first-order conditions: a deposit-weighted pooled checking FOC and two segment-specific credit card FOCs. This yields an independent 3×3 linear system per bank, which we solve in closed form.

Calibrated Parameters The first-order conditions require several parameters that we calibrate rather than estimate. The discount rate is set to $r = 0.02$, matching long-run real Treasury yields. The balance parameters b_k , which measure the typical dollar amounts consumers hold in checking deposits, credit card outstanding balances, and annual credit card spending, are calibrated as the simple average of median values from independent survey sources (SCF, NHCS, and MRI), separately by consumer type. For revolvers, the averaged values yield deposits of \$4,620 and credit card balances of \$3,430. For transactors, the averaged values yield deposits of \$22,736 and spending of \$14,247 per year, with zero outstanding balances by definition. Together, these calibrated values fully determine the mapping from estimated demand parameters and observed prices to implied marginal costs. Appendix C provides further details on the balance calibration.

Implied Margins and the Loss-Leader Mechanism Table 8 summarizes the implied margins by product and segment, providing a supply-side validation of the demand estimates against external benchmarks. The median checking account margin of 0.83 percentage points aligns with the deposit spreads documented by Drechsler et al. (2017)

and the markups estimated by Wang et al. (2022).

Credit card margins differ sharply by consumer type. The median transactor margin is a thin 0.45 percentage points. This is consistent with evidence from the Federal Reserve’s credit card profitability reports, which find that interchange revenue is largely offset by rewards for the transactors.¹⁴ In contrast, the median revolver margin is 4.3 percentage points, driven by interest income on outstanding balances. This is consistent with Nelson (2018) and Drechsler et al. (2025), who find substantial markups in credit card lending and identify interest income as the primary driver of issuer profitability.

VI Counterfactuals

Product complementarity reshapes how we assess competition and regulation in retail banking. We begin by quantifying the contribution of complementarity to bank franchise values. We then run three exercises (narrow-bank entry, interchange-fee regulation, and a horizontal merger) and compare predictions under the full model to a restricted specification that shuts down complementarity ($\chi = 0$). The gap between the two measures what a single-market analyst misses by treating deposits and credit cards as independent.

VI.A Complementarity and Bank Franchise Values

Figure 9 presents two perspectives on how checking accounts, transactor credit cards, and revolver credit cards contribute to banks’ profits. Panel A shows a standard accounting decomposition of each bank’s equilibrium profits across the three product lines. For deposit-oriented banks, such as Chase, Bank of America, and Wells Fargo, checking accounts generate the majority of profits. By contrast, Capital One and Citi derive a larger share of their profits from revolving credit card balances, reflecting their credit-heavy business models. The contribution from transactor credit cards is usually small, owing to their thin margins.

Panel B measures the marginal contribution of each product line by simulating its removal and recording the resulting decline in total franchise value. For every bank, these contributions sum to well above 100%. The excess arises because removing any line destroys not only its direct profits but also the demand spillovers it generates for complementary products — precisely the complementarity that Panel A’s accounting cannot capture.

Comparing the two panels underscores that the strategic importance of credit cards far exceeds what their accounting shares suggest. Aggregating across the named banks,

¹⁴Board of Governors of the Federal Reserve System, “Profitability of Credit Card Operations of Depository Institutions,” 2023.

credit cards account for roughly 40% of accounting profits in Panel A but contribute close to 60% of franchise value in Panel B once their cross-product pull on deposits is included. While transactor credit cards usually represent a minority of accounting profits, their marginal contributions to total profits are much larger. This divergence confirms that transactor cards function as a loss leader. They draw consumers into the bank and boost demand for the higher-margin checking accounts.

For megabanks such as Chase and Bank of America, transactor credit cards serve as the lead product. Generous rewards raise credit card consideration and acquisition, and cardholders then adopt the bank's higher-margin checking account. For the super-regional bank (which aggregates PNC, Citizens, and similar institutions), the pattern inverts. Standalone credit card adoption is rare, and the deposit account initiates the relationship instead. Holding a deposit then raises the probability that the same consumer considers and adopts the bank's credit card. Panel B therefore shows the super-regional bank's excess value concentrated on the deposit side, consistent with the 60× cross-holding ratio documented in Section III.

VI.B Margin Decomposition

A second decomposition recovers the same complementarity from banks' equilibrium margins rather than their profits. The supply-side first-order conditions in Equation 13 are linear in costs, so once we substitute the flow payoff $h_j(\theta) = h_j^0(\theta) - \sum_m c_{jm} O_{jm} b_m$ and use p_{jk} in consumer-favorable units (so that $\mu_{jk} \equiv -p_{jk} - c_{jk}$ is the margin), the FOCs rearrange into a bank-specific linear system in margins:

$$A_j \mu_j = e_j, \tag{18}$$

where

$$\begin{aligned} A_j[k, m] &= \mathbb{E}_\theta \left[\bar{s}_i(\theta) \mathcal{R}(\theta) \frac{\partial Q_i(\theta)}{\partial p_{jk}} \mathcal{R}(\theta) O_{jm} b_m \right], \\ e_j[k] &= \mathbb{E}_\theta \left[\bar{s}_i(\theta) \mathcal{R}(\theta) O_{jk} b_k \right]. \end{aligned}$$

The diagonal coefficient $A_j[k, k]$ shows how a price cut on product k changes product- k balances. The off-diagonal $A_j[k, m]$ captures the cross-selling spillover into product m 's balances. The right-hand side $e_j[k]$ is the infra-marginal cost of a price cut: the bank sacrifices b_k dollars on every existing holder of product k .

We split each margin into two pieces:

$$\mu_{jk} = \underbrace{\frac{e_j[k]}{A_j[k, k]}}_{\text{own-product markup } \mu_{jk}^{\text{own}}} + \underbrace{-\frac{1}{A_j[k, k]} \sum_{m \neq k} A_j[k, m] \mu_{jm}}_{\text{cross-selling adjustment } \mu_{jk}^{\text{cross}}}. \quad (19)$$

The own-product markup μ_{jk}^{own} is the margin a single-product optimizer would charge. The cross-selling adjustment μ_{jk}^{cross} collects the off-diagonal terms that a multi-product bank internalizes, each weighted by the complementary product's own margin μ_{jm} . Under our sign convention, a larger μ_{jk} is a more generous price, so a negative cross-selling adjustment means the bank prices more aggressively than a single-product rival, harvesting downstream profits on the complementary product.

Figure 10 plots this decomposition for each bank and product. Across essentially all bank-product pairs, the own-product markup is positive, and the cross-selling adjustment is negative. Equilibrium margins are therefore systematically below their single-product counterparts, reflecting the implicit subsidy that complementarity generates. The adjustment is largest on transactor credit cards at megabanks, where rewards rates exceed what a standalone issuer would offer. For these banks, the own-product markup on rewards cards is roughly one-third higher than the equilibrium margin. For the super-regional bank, the adjustment concentrates on deposits, mirroring the franchise-value pattern in Figure 9 Panel B.

VI.C Narrow Bank Entry

A long-standing puzzle in retail banking is the relative scarcity of entry in the deposit market despite persistently high profit margins (Drechsler et al., 2017). One potential mechanism is that product complementarity creates significant entry barriers for single-product entrants (Nalebuff, 2004).

We simulate the entry of a narrow bank that offers a deposit account at a 3% interest rate, equivalent to the federal funds rate and well above the near-zero rates offered by most incumbents. We set the entrant's consideration probability at the median of existing banks and allow incumbents to re-optimize all prices to the new Nash equilibrium. Under our full model, the narrow bank captures 15.3% of the deposit market, while an analyst who ignored complementarity would predict 22%. Complementarity insulates incumbents and helps them retain consumers who value cross-holding even at less competitive deposit rates.

Figure 11 Panel A shows the decomposition across incumbents. Multi-product banks lose between 15% and 28% of their checking share under the full model, compared with 24% to 39% when complementarities are shut down. The gap is most pronounced for

banks with strong credit-card franchises. American Express’s checking franchise loses only 11% under the full model versus 50% under $\chi = 0$.

Narrow bank entry also spills over into the credit card market. Deposit-heavy banks shed 1–5% of their credit franchise as their deposit base shrinks. Panel B shows that they also cut credit-card rewards in response: with a smaller deposit base, the loss-leader strategy would yield lower returns. Credit-dominant banks gain modestly in credit-card share as their deposit-heavy rivals pull back on rewards.

VI.D Regulatory Spillovers

The recent reintroduction of the Credit Card Competition Act (2026) aims to reduce interchange fees by requiring large banks to enable at least two unaffiliated networks for transaction routing. While the primary intent is to lower merchant costs, product complementarity implies that the regulation may spill well beyond the credit card market. To assess these cross-market spillovers, we simulate a 1 percentage-point increase in transactor credit-card costs, mimicking the impact of interchange-fee regulation.

Figure 12 shows the impact of the policy. Banks pass through the cost increase by reducing rewards rates by 0.78 to 0.93 percentage points. These reductions spill over into the deposit market: banks with large credit-card franchises lose deposit share as weaker rewards erode cross-holding ties. Chase, whose retail franchise is the most exposed across both products, loses roughly 4% of its pre-shock checking franchise and raises its checking rate by 3.2 basis points to stem the outflow. Capital One and Citi take sharper relative hits, roughly 5% and 6% of their pre-shock deposit shares. Bank of America sheds about 2% of its deposit franchise and responds by raising deposit rates by 3.5 basis points. Across banks, the correlation between deposit-share responses and credit-rewards changes is 0.90.

VI.E Horizontal Merger

We simulate a horizontal merger between Citi and the super-regional bank, comparing price responses under the full model against a restricted specification with $\chi = 0$. Both brands remain active. The combined firm owns all four products and solves a joint first-order condition over them, internalizing flow payoffs and cross-product complementarity across the wider portfolio. Rivals re-optimize in the new Nash equilibrium.

Under $\chi = 0$, the merger delivers the standard result: internalizing competition between the two brands raises prices, as a conventional concentration analysis would predict (Figure 13, orange markers).

Complementarity inverts the pricing logic. The combined firm exploits cross-selling scope that neither partner could access alone, using each brand’s credit card as a loss

leader for the other’s deposit base. Citi lowers its deposit rate by 9.4 basis points — harvesting surplus from its captive depositors — while raising credit-card rewards by 5.9 basis points to funnel new cardholders into the higher-margin super-regional deposit franchise. The super-regional plays the symmetric role: it nudges its deposit rate up slightly while cutting credit rewards by 3.6 basis points, harvesting deposits already attracted by Citi’s stronger card (Figure 13, blue markers).

Rival banks respond defensively, raising deposit rates — Bank of America by 1.1 basis points, Chase by 0.4 basis points — the opposite of what a standard concentration analysis would predict. The full Nash response runs through credit-card pricing and rivals’ deposit adjustments, not through deposit concentration; a standard deposit-HHI screen misses it entirely.

When merging parties have meaningful cross-product complementarities, joint pricing creates cross-selling gains that two separate firms could not realize — and that deposit concentration metrics cannot detect. Merger analysis in retail banking should thus incorporate cross-holding diagnostics, rather than relying on deposit concentration alone.

VII Conclusion

This paper shows that cross-selling and perceived complementarity, not correlated preferences alone, drive the widespread pattern of consumers holding multiple financial products at the same bank. Across three U.S. consumer surveys, roughly 35% of consumers with one checking account and one credit card hold both at the same bank, more than five times the 6.4% benchmark implied by independent product choice; the relationship multiplier averages $6\times$ at megabanks and exceeds $60\times$ at regional banks. Banks exploit infrequent consumer reoptimization by using one product as a gateway to capture demand for the other. A dynamic model of joint deposit and credit card choices, estimated on survey and cross-sectional data, lets us separately identify these mechanisms and trace their equilibrium implications. Cross-product complementarity creates entry barriers for narrow banks, generates regulatory spillovers from credit card policy to deposit markets, and amplifies the bundling synergies in horizontal mergers.

These mechanisms reshape standard measures of bank performance. Credit cards account for roughly 40% of bank franchise profits in standard accounting but close to 60% once their cross-product pull on deposits is included, and rewards-card margins at large banks would be about a third higher absent the cross-product subsidy. In our merger counterfactual, ignoring complementarity makes consolidation appear unprofitable for the combining firms; restoring complementarity flips the sign and yields a meaningful

combined-profit gain.

These findings bear directly on ongoing policy debates. Cross-product complementarity limits the competitive threat from narrow bank entry, since even at substantially higher deposit rates, a narrow bank cannot replicate the bundling value that retains cross-holding consumers. Interchange fee regulation, by eroding credit card rewards, weakens cross-product ties and transmits into deposit market competition, an indirect channel single-market analysis misses. For merger review, the longstanding focus on deposit-market HHI in federal bank merger guidelines (U.S. Department of Justice, 1995; U.S. Department of Justice and Federal Trade Commission, 2023) understates consolidation when the merging parties hold complementary franchises, and cross-holding data are needed to score the full diversion—which runs through the credit-card market and rivals’ defensive pricing on deposits, not through deposit concentration alone.

Several limitations suggest avenues for future work. Most notably, our model takes the timing of initial product adoption as given. The arrival rate λ_k governs how often already-adopted choices are reconsidered. However, the extensive-margin decision of whether and when a consumer first acquires a deposit account or credit card lies outside the model. A richer framework that endogenizes this adoption timing would allow a more complete welfare analysis of policies that affect first-time product uptake.


References

- Agarwal, Sumit, Souphala Chomsisengphet, Chunlin Liu, Changcheng Song, and Nicholas S Souleles**, “Benefits of Relationship Banking: Evidence from Consumer Credit Markets,” *Journal of Monetary Economics*, 2018, 96, 16–32.
- , – , **Neale Mahoney, and Johannes Stroebel**, “Regulating Consumer Financial Products: Evidence from Credit Cards *,” *The Quarterly Journal of Economics*, February 2015, 130 (1), 111–164.
- Aguirregabiria, Victor, Robert Clark, and Hui Wang**, “The geographic flow of bank funding and access to credit: Branch networks, local synergies and competition,” *arXiv preprint arXiv:2407.03517*, 2024.
- Allcott, Hunt**, “The Welfare Effects of Misperceived Product Costs: Data and Calibrations from the Automobile Market,” *American Economic Journal: Economic Policy*, 2013, 5 (3), 30–66.
- Arcidiacono, Peter, Patrick Bayer, Jason R Blevins, and Paul B Ellickson**, “Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition,” *The Review of Economic Studies*, 2016, 83 (3), 889–931.
- Basten, Christoph and Ragnar Juelsrud**, “Cross-Selling in Bank-Household Relationships: Mechanisms and Implications for Pricing,” *The Review of Financial Studies*, 2023, p. hhad062.
- Berry, Steven, James Levinsohn, and Ariel Pakes**, “Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market,” *Journal of Political Economy*, February 2004, 112 (1), 68–105.
- Brown, Zach Y, Mark L Egan, Jihye Jeon, Chuqing Jin, and Alex A Wu**, “Why Do Index Funds Have Market Power? Quantifying Frictions in the Index Fund Market,” Technical Report, National Bureau of Economic Research 2023.
- Caldwell, Sydnee, Ingrid Haegeler, and Jörg Heining**, “Why Workers Stay: Pay, Beliefs, and Attachment,” 2025.
- Cuesta, José Ignacio and Alberto Sepúlveda**, “Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access,” *Available at SSRN 3282910*, 2021.
- Drechsler, Itamar, Alexi Savov, and Philipp Schnabl**, “The deposits channel of monetary policy,” *The Quarterly Journal of Economics*, 2017, 132 (4), 1819–1876.
- , – , and – , “Banking on Deposits: Maturity Transformation without Interest Rate Risk,” *The Journal of Finance*, 2021, 76 (3), 1091–1143.
- , **Hyeyoon Jung, Weiyu Peng, Dominik Supera, and Guanyu Zhou**, “Credit Card Banking,” *Working Paper*, March 2025.

- Egan, Mark, Ali Hortaçsu, and Gregor Matvos**, “Deposit Competition and Financial Fragility: Evidence from the Us Banking Sector,” *American Economic Review*, 2017, 107 (1), 169–216.
- Egan, Mark L., Ali Hortaçsu, Nathan A. Kaplan, Adi Sunderam, and Vincent Yao**, “Dynamic Competition for Sleepy Deposits,” September 2025.
- Galenianos, Manolis and Alessandro Gavazza**, “Regulatory Interventions in Consumer Financial Markets: The Case of Credit Cards,” *Journal of the European Economic Association*, 2022, 20 (5), 1897–1932.
- Gentzkow, Matthew**, “Valuing New Goods in a Model with Complementarity: Online Newspapers,” *American Economic Review*, 2007, 97 (3), 713–744.
- Goeree, Michelle Sovinsky**, “Limited Information and Advertising in the U.S. Personal Computer Industry,” *Econometrica*, 2008, 76 (5), 1017–1074.
- Grieco, Paul LE, Charles Murry, and Ali Yurukoglu**, “The evolution of market power in the us automobile industry,” *The Quarterly Journal of Economics*, 2024, 139 (2), 1201–1253.
- Hastings, Justine, Ali Hortaçsu, and Chad Syverson**, “Sales Force and Competition in Financial Product Markets: The Case of Mexico’s Social Security Privatization,” *Econometrica : journal of the Econometric Society*, 2017, 85 (6), 1723–1761.
- Hellmann, Thomas, Laura Lindsey, and Manju Puri**, “Building Relationships Early: Banks in Venture Capital,” *The Review of Financial Studies*, 2008, 21 (2), 513–541.
- Honka, Elisabeth, Ali Hortaçsu, and Maria Ana Vitorino**, “Advertising, Consumer Awareness, and Choice: Evidence from the US Banking Industry,” *The RAND Journal of Economics*, 2017, 48 (3), 611–646.
- Hortaçsu, Ali and Chad Syverson**, “Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds,” *The Quarterly journal of economics*, 2004, 119 (2), 403–456.
- Ivashina, Victoria and Anna Kovner**, “The Private Equity Advantage: Leveraged Buy-out Firms and Relationship Banking,” *The Review of Financial Studies*, 2011, 24 (7), 2462–2498.
- JPM**, “JPM 2015 Investor Day.Pdf,” February 2015.
- Laux, Christian and Uwe Walz**, “Cross-Selling Lending and Underwriting: Scope Economies and Incentives,” *Review of Finance*, 2009, 13 (2), 341–367.
- Mas, Alexandre and Amanda Pallais**, “Valuing alternative work arrangements,” *American Economic Review*, 2017, 107 (12), 3722–3759.
- and —, “Labor supply and the value of non-work time: Experimental estimates from the field,” *American Economic Review: Insights*, 2019, 1 (1), 111–126.


- Nalebuff, Barry**, “Bundling as an Entry Barrier,” *The Quarterly Journal of Economics*, 2004, 119 (1), 159–187.
- Nelson, Scott**, “Private Information and Price Regulation in the Us Credit Card Market,” *Unpublished working paper*, 2018.
- Parlour, Christine A, Uday Rajan, and Haoxiang Zhu**, “When FinTech Competes for Payment Flows,” *The Review of Financial Studies*, November 2022, 35 (11), 4985–5024.
- Petrin, Amil**, “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, 2002, 110 (4).
- Puri, Manju, Jörg Rocholl, and Sascha Steffen**, “What Do a Million Observations Have to Say about Loan Defaults? Opening the Black Box of Relationships,” *Journal of Financial Intermediation*, July 2017, 31, 1–15.
- Qi, Yingjie**, “Big Broad Banks: How Does Cross-Selling Affect Lending?,” *Review of Finance*, 2024, 28 (2), 551–592.
- Santikian, Lori**, “The Ties That Bind: Bank Relationships and Small Business Lending,” *Journal of financial intermediation*, 2014, 23 (2), 177–213.
- Stango, Victor and Jonathan Zinman**, “Borrowing High versus Borrowing Higher: Price Dispersion and Shopping Behavior in the US Credit Card Market,” *The Review of Financial Studies*, 2016, 29 (4), 979–1006.
- Train, Kenneth E.**, “EM Algorithms for Nonparametric Estimation of Mixing Distributions,” *Journal of Choice Modelling*, 2008, 1 (1), 40–69.
- U.S. Department of Justice**, “Bank Merger Competitive Review – Introduction and Overview,” 1995.
- U.S. Department of Justice and Federal Trade Commission**, “Horizontal Merger Guidelines,” 2023.
- Wang, Lulu**, “Regulating Competing Payment Networks,” URL: https://luluywang.github.io/PaperRepository/payment_jmp.pdf (last accessed on 10 November 2023), 2023.
- Wang, Yifei, Toni M. Whited, Yufeng Wu, and Kairong Xiao**, “Bank Market Power and Monetary Policy Transmission: Evidence from a Structural Estimation,” *The Journal of Finance*, 2022, 77 (4), 2093–2141.
- Whinston, Michael D**, “Tying, Foreclosure, and Exclusion,” 1989.

Figure 1: Cross-holding of Checking Accounts and Other Products



Notes: The fraction shown reflects the share of customers holding additional products at their primary bank. For each product, we calculate the fraction among consumers holding both a checking account and that product.

Figure 2: Comparing Cross-holding of Disadvantaged versus Advantaged Groups




Notes: The figure compares cross-holding between disadvantaged and advantaged consumer groups. We define the disadvantaged group as those with low income, low education, low credit scores, or high credit card balances (i.e., carrying a balance sometimes or often). Specifically, the lower education group consists of consumers with a high school degree or below, and low credit scores are those below 720. Income is segmented using the sample median. We proxy revolvers using the balance carrier group (those who carry a balance sometimes or often) in contrast to transactors, who spend to collect rewards.

Figure 3: Cross-holding of Checking Accounts and Credit Cards by Bank




Notes: This figure shows the prevalence of cross-holding between checking accounts and credit cards. The top panel shows each bank's checking account market share among two groups of consumers: those who hold a credit card at the bank (blue dots) and those who do not (orange dots). The bottom panel shows the credit card market shares among consumers who do and do not have a checking account at the bank. Market shares reflect both primary and other accounts. The sample includes respondents in the 2021 and 2022 MRI data.

Figure 4: Consumer Tenure at Current Banks




Notes: This figure shows the distribution of consumer tenure at their current bank for credit account and deposit account, separately. 0 means the consumer obtained their current product less than a year ago, while 20 means more than twenty years ago.

Figure 5: Stated Responses to Hypothetical Shocks




Notes: This figure shows the fraction of consumers who state they would switch to a different bank or significantly reduce the usage of their current product under various hypothetical shocks. Transactors are individuals who never or rarely carry a balance on their credit cards, while revolvers are individuals who carry a balance at least occasionally. The “crosshold” group consists of consumers who currently hold their primary deposit account and credit card at the same bank.

Figure 6: Goodness of Fit: Market Shares




Notes: These figures compare observed and model-predicted market shares for checking accounts (left) and credit cards (right). Each point represents a bank–segment pair (transactor and revolver shown separately, distinguished by color and shape), with the horizontal axis showing observed shares and the vertical axis showing predicted shares. Points along the 45-degree line indicate perfect fit.

Figure 7: Goodness of Fit: Cross-holding Rates by Bank and Consumer Type




Notes: This figure compares observed and model-predicted cross-holding rates across banks for transactors and revolvers. The figure shows the probability of holding a credit card at a bank conditional on having a checking account there.

Figure 8: Goodness of Fit: Tenure Distributions




Notes: These figures compare model-predicted and observed tenure distributions for checking accounts (left) and credit cards (right). Tenures are binned into five categories: less than 1 year, 1–3 years, 3–5 years, 5–10 years, and 10+ years. The model predictions are computed from the backward recurrence time distribution of the estimated CTMC (see Appendix D.6). Observed tenures come from the Prolific survey, with NHCS survey data shown as an additional benchmark for checking accounts.

Figure 9: Profit Decomposition and Franchise Values



Notes: This figure decomposes bank franchise values across product lines. Panel A shows the accounting decomposition: each bank's total profit is split into contributions from checking accounts, credit card transactors, and credit card revolvers. Panel B shows the marginal contribution of each product line, computed by removing that product and measuring the resulting loss in the bank's total franchise value. The dashed vertical line marks 100%. Marginal contributions sum to more than 100%, reflecting cross-product complementarities: removing any single product line reduces profits from the remaining lines as well.

Figure 10: FOC Margin Decomposition: Own-Product Markup vs. Cross-Selling Adjustment



Notes: This figure decomposes each bank's equilibrium margin (percentage points) into two components. The own-product markup ($\mu_{jk}^{\text{own}} = e_j[k] / A_j[k, k]$) is the margin a single-product optimizer would set; the cross-selling adjustment ($\mu_{jk}^{\text{cross}} = -A_j[k, k]^{-1} \sum_{m \neq k} A_j[k, m] \mu_{jm}$) is the margin shading a multi-product bank applies to internalize complementary-product profits (see Equation 19). Panels A–C show deposits, transactor credit, and revolver credit, respectively. Diamonds mark the total margin μ_{jk} . A negative cross-selling adjustment indicates that the bank prices product k more generously than a standalone competitor would, in order to harvest profits on the complementary product line.

Figure 11: Counterfactual: Narrow Bank Entry



Notes: This figure shows equilibrium outcomes under the entry of a narrow bank that offers only a deposit account. Panel A shows market shares, Panel B shows prices, and Panel C shows profits under varying levels of cross-product complementarity.

Figure 12: Counterfactual: Interchange Fee Regulation



Notes: This figure shows equilibrium outcomes under interchange fee regulation that reduces credit card reward rates. Panel A shows market shares, Panel B shows prices, and Panel C shows profits. The results illustrate regulatory spillovers from the credit card market to the deposit market.

Figure 13: Counterfactual: Horizontal Merger (Citi + Super-Regional Bank)



Notes: This figure shows equilibrium outcomes under a horizontal merger between Citi and the super-regional bank. Both brands remain active in the choice set; what changes is that the combined firm now owns all four products (checking and credit at each brand) and solves a joint first-order condition over them, while cross-selling complementarity χ pools across the firm so the bonus fires for any pair of products held within the combined entity. Panel A shows post-merger changes in credit card and deposit market shares by bank; Panel B shows changes in credit card reward rates and deposit rates. Results are reported for the full model (blue) and the restricted model that shuts down cross-product complementarity, $\chi = 0$ (orange). Under the restricted model, joint pricing alone yields no synergy: the combined firm loses both deposit and credit share to rivals and profits decline. Under the full model, the broader span of complementarity within the firm makes the merger profitable: Citi lowers its deposit rate to extract surplus while raising credit-card rewards to acquire customers, and combined profits rise by 6.6%.

Table 1: Summary Statistics

Characteristic	Data Source			
	Prolific	MRI	NHCS	SCF
<i>Panel A: Demographics</i>				
Median Income (in thousand dollars)	62.0	67.1	67.1	70.5
Bachelor's Degree or Above (%)	53.9	35.9	34.7	40.1
Credit Score ≥ 720 (%)	77.8	78.1		
Carrying Credit Card Balance (%)	43.3	36.6	43.4	39.3
<i>Panel B: Product Holdings (%)</i>				
Deposit Account (%)	98.8	94.4	94.1	94.6
Credit Card (%)	86.0	76.0	73.1	80.0
Auto Loan (%)	22.5	24.3	31.4	34.1
Mortgage (%)	30.3	41.3	32.0	40.9
Brokerage Account (%)	32.6	32.5	18.9	26.0
<i>Panel C1: Account Balances - Revolvers (in thousands)</i>				
Median Total Deposits (k\$)	1.1		5.5	3.7
Median Credit Card Balance (k\$)	0.9			3.4
Median Credit Card Spending (k\$/yr)	2.8	4.1		3.8
<i>Panel C2: Account Balances - Transactors (in thousands)</i>				
Median Total Deposits (k\$)	9.0		22.0	23.5
Median Credit Card Spending (k\$/yr)	8.4	10.2		18.4
Observations	955	100,884	39,716	4,595

Notes: This table reports summary statistics across our main data sources. Prolific refers to our original survey of approximately 1,000 U.S. consumers. MRI refers to the 2021–2022 waves of the MRI-Simmons USA Study. NHCS refers to the 2020 wave of the MRI-Simmons USA Study, which additionally captures consumer tenure at their financial institutions. SCF refers to the 2022 Survey of Consumer Finances. Panel A compares demographic characteristics. Panel B compares product ownership rates. Panels C1 and C2 report median account balances for revolvers (consumers who carry credit card balances) and transactors (consumers who pay their balance in full each month), respectively. Cells are blank where the data source does not report the variable. Prolific respondents hold lower median balances than SCF and NHCS because the Prolific panel skews younger and lower-income; the structural price sensitivities estimated on this sample nonetheless match external benchmarks (see Section V), suggesting that the demand parameters identified from this sample generalize.

Table 2: Bank Market Shares in Deposit Accounts and Credit Cards

Bank	Deposit (MRI)	Deposit (Prolific)	Deposit (SOD)	Credit (MRI)	Credit (Prolific)	Credit (Nilson)
Chase	13.6	10.7	11.7	17.4	22.4	17.0
Bank of America	11.1	7.7	11.0	10.6	9.5	8.3
Wells Fargo	10.6	7.7	8.3	6.0	4.5	3.0
Capital One	5.2	10.3	2.2	14.2	24.0	14.1
PNC	3.0	3.4	2.3	1.4	0.5	0.5
Citi	2.9	1.2	4.1	8.6	7.0	12.1
US Bank	2.8	2.9	2.5	2.2	2.1	3.3
TD Bank	1.8	1.4	2.2	0.8	0.8	1.2
Citizens Bank	1.3	1.3	1.5	0.6	0.1	0.2
American Express	-	3.5	0.5	7.5	8.1	9.0
Discover	-	3.3	0.4	4.9	11.6	7.6
Others	47.6	46.4	53.2	25.7	9.4	23.7
Total	100.0	100.0	100.0	100.0	100.0	100.0

Notes: This table shows market shares for deposit accounts and credit cards in the MRI sample and our Prolific Survey sample. Market shares in MRI are weighted by the provided population weights. The "Others" category includes credit unions, internet banks, mutual funds, and various other national or regional banks. For MRI deposit accounts, the "Others" category also includes American Express and Discover. When a consumer holds multiple accounts for a single product, each account is treated as a separate observation in the market share calculations. For comparison, Column 3 shows deposit account market shares based on the Summary of Deposit Data, and Column 6 shows credit card market shares from the Nilson Report.

Table 3: Cross-holding and Consumer Characteristics

	1 Checking + 1 CC (1)	Main CC at Checking (1)	CC at Each Checking (1)
Age / 10	-0.10*** (0.01)	-0.10*** (0.01)	-0.12*** (0.01)
College Degree	0.00 (0.03)	0.02 (0.01)	0.04*** (0.01)
Log Income	-0.08*** (0.02)	-0.08*** (0.01)	-0.11*** (0.01)
CFO or Treasurer	-0.09 (0.15)	-0.32* (0.15)	-0.35** (0.13)
Online Banking	0.06+ (0.03)	0.03 (0.02)	0.00 (0.06)
Invests in Stocks	-0.20*** (0.03)	-0.15*** (0.04)	-0.09 (0.11)
Money Market	-0.09+ (0.05)	-0.08** (0.03)	-0.02 (0.05)
Tax-Advantaged Accounts	-0.19*** (0.03)	-0.15*** (0.03)	-0.20*** (0.05)
FICO > 780	-0.02 (0.04)	-0.03 (0.02)	0.01 (0.01)
Factor: Interest Rates	-0.13*** (0.03)	-0.11*** (0.02)	-0.08** (0.03)
Factor: Reward Program	0.66*** (0.03)	0.61*** (0.04)	0.59*** (0.02)
Num.Obs.	28858	50100	50053
FE: num_cc^num_checking		X	X

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

+ p \$<\$ 0.1, * p \$<\$ 0.05, ** p \$<\$ 0.01, *** p \$<\$ 0.001

Notes: This table shows results of logistic regressions estimating the relationships between consumer characteristics and the propensity to hold checking accounts and credit cards at the same bank. Column (1) restricts the sample to consumers with exactly one checking account and one credit card; the dependent variable equals one if both products are from the same bank. Columns (2) and (3) use the full sample of consumers with at least one checking account and one credit card. In Column (2), the dependent variable indicates whether the consumer has a checking account at the same bank as the primary credit card. In Column (3), the dependent variable indicates whether all pairs of checking accounts and credit cards are held at the same banks. Both columns flexibly control for the number of checking accounts and credit cards. *** implies significance at 0.001 level, ** 0.01, * 0.05, + 0.1.

Table 4: Factors for Banking Product Considerations

Outcome	(1) Deposit	(2) Deposit	(3) Credit	(4) Credit
Credit Used	0.060*** (0.027)	0.063*** (0.025)		
Credit Considered	0.242*** (0.015)	0.235*** (0.014)		
Deposit Used			0.061** (0.026)	0.108*** (0.024)
Deposit Considered			0.257*** (0.017)	0.257*** (0.017)
Constant	0.117*** (0.003)	0.117*** (0.003)	0.137*** (0.004)	0.135*** (0.004)
Observations	11,796	11,796	11,796	11,796
<i>Controlling for:</i>				
Bank FE		✓		✓

Notes: This table shows factors that predict whether a consumer considers a bank when choosing a deposit account (Columns 1 and 2) or a credit card (Columns 3 and 4). *Credit Used* indicates whether the consumer was using the bank's credit card when choosing their deposit account. *Deposit Used* indicates whether the consumer was using the bank's deposit account when choosing a credit card. *** implies significance at 0.01 level, ** 0.05, * 0.1.

Table 5: Structural Parameter Estimates: Transactors vs. Revolvers

Parameter	Transactors	Revolvers
Price Sensitivity		
Checking (α_1)	2.25	1.68
Credit (α_2)	2.71	0.49
Cross-Selling Effects		
Consideration – Checking (χ_1^c)	1.23	0.85
Consideration – Credit (χ_2^c)	1.53	1.02
Utility – Checking (χ_1^u)	0.78	0.73
Utility – Credit (χ_2^u)	0.64	0.90
Implied Benefits		
Checking (pp deposit rate)	0.35	0.43
Credit (pp cashback/APR)	0.24	1.83
Arrival Rates		
Checking (λ_1)	0.22	0.30
Credit (λ_2)	0.21	0.28

Notes: Reports population means θ of the structural parameters from the MC-EM estimator. N reports the Prolific sample sizes; the full likelihood also incorporates approximately 100,000 MRI-Simmons observations via a weighted summand on the steady-state likelihood (see Section V). Standard errors are omitted; parametric bootstrap inference is under development.

Table 6: Estimated Heterogeneity: Standard Deviations and Correlations of Σ_θ

	Transactors	Revolvers
Bank Preference Std. Deviations (Σ_v)		
v_1^c (Check. Consid.)	1.65	1.59
v_2^c (Credit Consid.)	1.56	1.70
v_1^u (Check. Utility)	1.33	1.65
v_2^u (Credit Utility)	1.51	2.01
Bank Preference Correlations (Σ_v)		
$\rho(v_1^c, v_2^c)$	0.96	0.96
$\rho(v_1^u, v_2^u)$	0.65	0.15
Structural Parameter Std. Deviations (Σ_θ, log scale)		
$\log \alpha_1$ (Checking)	1.04	1.40
$\log \alpha_2$ (Credit)	0.11	0.08
$\log \chi_1^u$	1.00	0.57
$\log \chi_2^u$	0.09	0.37
$\log \lambda_1$ (Checking)	0.11	0.16
$\log \lambda_2$ (Credit)	0.10	0.14

Notes: Reports standard deviations and within-bank correlations of the estimated heterogeneity distribution. The Σ_v block captures bank-preference heterogeneity in levels, with strong within-bank correlation across consideration and utility for the same product (near-unit $\rho(v_1^c, v_2^c)$); cross-correlations between consideration (v^c) and utility (v^u) blocks are restricted to zero and omitted from the table. The $(\log \alpha_k, \log \chi_k^u)$ blocks are estimated separately by product (the cross-product joint distribution is not identified by the survey design) and exhibit substantive dispersion; $(\log \lambda_1, \log \lambda_2)$ is jointly normal across products with limited within-segment dispersion. The consideration complementarity χ^c is held homogeneous within consumer segment ($\bar{\chi}_k^c$ identified by Fact 3; variance not identified).

Table 7: Bank Preference Parameters: Transactors and Revolvers

Bank	v_1^c (Check. Consid.)	v_2^c (Credit Consid.)	v_1^u (Check. Utility)	v_2^u (Credit Utility)
Panel A: Transactors				
American Express	—	-1.69	—	-5.84
Bank of America	-2.07	-1.84	-2.88	-4.65
Capital One	-2.46	-1.19	-8.85	-3.83
Chase	-2.05	-0.61	-3.27	-4.42
Citi	-3.71	-1.84	-9.29	-5.81
Regional Banks	-2.71	-3.42	-4.18	-2.30
Discover	—	-1.44	—	-4.90
US Bank	-2.51	-2.77	-2.07	-3.46
Panel B: Revolvers				
American Express	—	-1.85	—	10.20
Bank of America	-2.10	-2.17	-1.68	14.63
Capital One	-2.59	-0.31	-6.12	13.26
Chase	-1.91	-0.81	-1.91	14.50
Citi	-3.69	-2.18	-7.76	10.30
Regional Banks	-2.25	-3.29	-3.56	10.76
Discover	—	-1.12	—	11.42
US Bank	-2.64	-3.34	-1.14	11.02

Table 8: Implied Margins and Costs from Supply-Side First-Order Conditions


Statistic	Value
Margins (Median)	
Median Checking Margin (pp)	0.83
Median Transactor Credit Margin (pp)	0.45
Median Revolver Credit Margin (pp)	4.27
Checking Cost Range	
Min Checking Cost	-0.0369
Max Checking Cost	-0.0202
Transactor Credit Cost Range	
Min Transactor Credit Cost	-0.0278
Max Transactor Credit Cost	-0.0135
Revolver Credit Cost Range	
Min Revolver Credit Cost	0.1549
Max Revolver Credit Cost	0.2226

Notes: This table reports implied margins and marginal costs recovered from the supply-side first-order conditions of the EM specification. Margins are the difference between per-unit revenue and marginal cost for each product, in percentage points, averaged across consumer types. Costs represent the marginal cost of serving a customer per dollar of balance, net of interchange and investment-return revenue; negative values indicate that these net revenues exceed direct servicing costs.

Appendices. For Online Publication Only

A Additional Figures and Tables

Figure A.1: Crossholding of Checking Accounts and Credit Cards by Bank, Prolific Data




Notes: This figure shows the deposit account market shares of Chase, Wells Fargo, and Bank of America among two groups of consumers: those who have a brokerage service account at the bank (orange dots) and those who do not (green dots). The sample includes all consumers in the 2021 and 2022 MRI data.

Figure A.2: Cross-holding at Big and Small Banks




Notes: This figure shows the prevalence of cross-holding across banks. The top panel plots, for each bank, the ratio of credit card market shares among consumers who do and do not have a checking account at the bank. The bottom panel plots the analogous ratio for checking account market shares conditional on credit card ownership. Banks are sorted in descending order by their aggregate market shares in checking accounts.

Figure A.3: Sizes of Consideration Sets




Notes: This figure shows the distributions of the sizes of the consideration sets for deposit accounts and credit cards in our Prolific Survey Sample. The consideration set includes each consumer’s bank of choice, and other banks they “looked into or reviewed” when choosing the product.

Figure A.4: Cross-holding by Consumer Subgroups: Detailed View




Notes: This figure provides a detailed breakdown of cross-holding rates by consumer subgroups. Each panel shows the rate of credit card cross-holding separately for different demographic characteristics: income, education, credit score, and credit card balance. Within each panel, we compare disadvantaged and advantaged groups as defined in the main text.

Figure A.5: Goodness of Fit: Tenure by Age



Notes: This figure compares model-predicted and empirical average tenure (time since the most recent bank switch) for checking accounts and credit cards, broken down by consumer age. Solid lines show Prolific survey averages; dashed lines show the model-implied analog from a forward simulation in which each synthetic consumer is assigned a first-account arrival age drawn uniformly from a fixed window and then experiences switching dynamics under the estimated CTMC until their current age.

Figure A.6: Observed Prices by Bank



Notes: This figure reports the observed prices used as the supply-side targets in the model: deposit interest rates, credit card reward rates, and credit card APRs by bank. Deposit rates are constructed from Y-14 quarterly bank holding company filings as the ratio of interest expense on domestic deposits to the total amount of domestic deposits. Reward expense per dollar of credit card spending is constructed from bank financial filings combined with payment volume data from the Nilson Report. APRs are from the CFPB Terms of Credit Card Plans (TCCP) survey.

B An Example of Consumer Demand Model

Consider a simple example with two banks and the outside option for each product. There are nine possible states:

$$x \in \{(0,0), (0,1), (0,2), (1,0), (1,1), (1,2), (2,0), (2,1), (2,2)\}$$

Assume for simplicity that $\delta_{jk} = 0$ and $p_{jk} = 0$ for all products. Let $\chi = \log 2$, so that consumers are twice as likely to choose the credit card from the bank where they have the deposit account and vice versa. We can derive the Markov transition intensity matrix for the deposit account as:

$$Q_1 = \lambda_1 \times \begin{pmatrix} -2/3 & & & 1/3 & & & 1/3 & & \\ & -3/4 & & & 1/2 & & & 1/4 & \\ & & -3/4 & & & 1/4 & & & 1/2 \\ 1/3 & & & -2/3 & & & 1/3 & & \\ & 1/4 & & & -1/2 & & & 1/4 & \\ & & 1/4 & & & -3/4 & & & 1/2 \\ 1/3 & & & 1/3 & & & -2/3 & & \\ & 1/4 & & & 1/2 & & & -3/4 & \\ & & 1/4 & & & 1/4 & & & -1/2 \end{pmatrix}$$

This matrix illustrates key model features. The first row shows transitions from state $(0,0)$. Only 2 off-diagonal elements appear: consumers can move only to states holding a credit card at bank 0, since the matrix captures deposit-switching opportunities only. Given equal product qualities, the transition rate to each of those states is $\lambda_1/3$, so the total flow out of $(0,0)$ is $-2\lambda_1/3$.

The second row describes transitions out of the $(0,1)$ state. Here, the consumer already has a credit card at bank 1. With $\chi = \log 2$, they are twice as likely to choose the deposit account from the same bank. As a result, the transition rate to the $(1,1)$ state is $\lambda_1/2$, twice the transition rate to the $(2,1)$ state. The total flow out of the $(0,1)$ state then becomes $-3\lambda_1/4$. The third has a similar intuition as the second row, and the fourth row has a similar intuition as the first row.

The fifth row brings out another economic force at work. This row represents transitions from $(1,1)$, where the consumer is already cross-holding at bank 1. With $\chi_1 = \chi_2 = \log 2$, the consumer prefers keeping the deposit account at bank 1. Therefore, the transitions into $(0,1)$ and $(2,1)$ are reduced. This generates the empirical pattern that

consumers who cross-hold are less likely to consider switching and more likely to be at their bank for a longer period of time.

Similarly, the transition matrix for Q_2 is

$$Q_2 = \lambda_2 \times \begin{pmatrix} -2/3 & 1/3 & 1/3 & & & & & & \\ 1/3 & -2/3 & 1/3 & & & & & & \\ 1/3 & 1/3 & -2/3 & & & & & & \\ & & & -3/4 & 1/2 & 1/4 & & & \\ & & & 1/4 & -1/2 & 1/4 & & & \\ & & & 1/4 & 1/2 & -3/4 & & & \\ & & & & & & -3/4 & 1/4 & 1/2 \\ & & & & & & 1/4 & -3/4 & 1/2 \\ & & & & & & 1/4 & 1/4 & -1/2 \end{pmatrix}$$

and the transition intensity matrix for the entire system is $Q = Q_1 + Q_2$.

Now suppose $\lambda_1 = 0.1$ and $\lambda_2 = 0.5$. This reflects the idea that people switch deposit accounts rarely but switch credit cards quickly. The full transition matrix then looks like

$$Q = \begin{pmatrix} -0.4 & 0.17 & 0.17 & 0.03 & & & & 0.03 & \\ 0.17 & -0.41 & 0.17 & & 0.05 & & & & 0.02 \\ 0.17 & 0.17 & -0.41 & & & 0.02 & & & 0.05 \\ 0.03 & & & -0.44 & 0.25 & 0.12 & 0.03 & & \\ & 0.02 & & 0.12 & -0.3 & 0.12 & & 0.02 & \\ & & 0.02 & 0.12 & 0.25 & -0.45 & & & 0.05 \\ 0.03 & & & 0.03 & & & -0.44 & 0.12 & 0.25 \\ & 0.02 & & & 0.05 & & 0.12 & -0.45 & 0.25 \\ & & 0.02 & & & 0.02 & 0.12 & 0.12 & -0.3 \end{pmatrix}$$

Given this process, an approximate steady-state vector of market shares $\mathbf{s}$ is a left eigenvector of the matrix Q with eigenvalue 0, normalized to have a sum of 1:

$$\mathbf{s} = (0.09 \ 0.09 \ 0.09 \ 0.09 \ 0.18 \ 0.09 \ 0.09 \ 0.09 \ 0.18)$$

At the steady state, consumers are more likely to cross-hold both products at the same bank.

The matrix exponential for one time period is $S(1)$ and equals

$$\exp(Q) = \begin{pmatrix} 0.69 & 0.12 & 0.12 & 0.02 & 0.01 & 0.0 & 0.02 & 0.0 & 0.01 \\ 0.12 & 0.69 & 0.12 & 0.0 & 0.04 & 0.0 & 0.0 & 0.02 & 0.01 \\ 0.12 & 0.12 & 0.69 & 0.0 & 0.01 & 0.02 & 0.0 & 0.0 & 0.04 \\ 0.02 & 0.0 & 0.0 & 0.66 & 0.19 & 0.09 & 0.02 & 0.0 & 0.01 \\ 0.0 & 0.02 & 0.0 & 0.09 & 0.76 & 0.09 & 0.0 & 0.02 & 0.0 \\ 0.0 & 0.0 & 0.02 & 0.09 & 0.19 & 0.65 & 0.0 & 0.0 & 0.04 \\ 0.02 & 0.0 & 0.0 & 0.02 & 0.01 & 0.0 & 0.66 & 0.09 & 0.19 \\ 0.0 & 0.02 & 0.0 & 0.0 & 0.04 & 0.0 & 0.09 & 0.65 & 0.19 \\ 0.0 & 0.0 & 0.02 & 0.0 & 0.0 & 0.02 & 0.09 & 0.09 & 0.76 \end{pmatrix}$$

Note that all entries of the matrix exponential are strictly positive (entries displayed as 0.0 reflect rounding to one decimal place), even though the transition intensity matrix Q has many zeros. In the model, it is possible to transition from, say, $(1,1)$ to $(2,2)$ in a short time window if you receive both a deposit account and credit card switching shock in the time window. The matrix exponential encodes this possibility through the process of multiplying the matrix.

C Data and Balance Parameter Calibration

C.1 Balance Parameters in the Model

The supply-side first-order conditions require balance parameters b_k measuring the typical dollar amounts in checking deposits (b_1), credit card outstanding balances (b_2^{balance}), and annual credit card spending (b_2^{spend}). Banks' profit calculations depend on these parameters: for checking, revenue comes from the net interest margin (the difference between rates earned on deposits and rates paid to depositors); for credit cards, from interest on balances and interchange fees on transaction volume.

C.2 Data Sources and Calibration

We calibrate each balance parameter as the simple average of median values from independent survey sources, separately by consumer type (revolvers, who carry a credit card balance, and transactors, who pay in full each month). Three surveys contribute to the calibration. The Federal Reserve's 2022 Survey of Consumer Finances (SCF) provides household-level medians for all three parameters: deposits, credit card outstanding balances, and credit card spending. The National Household Credit Survey (NHCS) provides weighted median deposit balances for a nationally representative sample of 16,489 consumers. The MRI-Simmons USA Study provides median annual credit card spending derived from survey spending buckets across approximately 100,000 consumers. For each parameter, we average across all sources that measure it: deposit balances average the SCF and NHCS medians, credit card spending averages the SCF and MRI medians, and credit card outstanding balances use the SCF median alone (the only source that reports this variable by consumer type). Table C.1 reports the source-specific median values side by side.

The resulting calibrated values are as follows. For revolvers, the averaged values yield deposits of \$4,620, credit card outstanding balances of \$3,430, and credit card spending of \$3,936 per year. For transactors, the averaged values yield deposits of \$22,759 and credit card spending of \$14,295 per year, with zero outstanding balances by definition. The 4.9-fold difference in deposits between transactors and revolvers reflects meaningful heterogeneity in the financial profiles of these consumer groups: transactors tend to be more affluent and hold substantially higher deposit balances.

Multi-Source Comparison Table C.1 compares the source-specific median values across the SCF, MRI/NHCS, and two additional reference sources (Prolific and Nilson Report). The SCF and NHCS deposit medians are broadly consistent: \$3,740 versus \$5,500 for

revolvers and \$23,520 versus \$21,998 for transactors. The SCF and MRI credit card spending medians agree closely for revolvers (\$3,816 versus \$4,056) but diverge more for transactors (\$18,384 versus \$10,206), reflecting differences in survey design and spending measurement. Averaging across sources mitigates idiosyncratic measurement error in any single survey.

Table C.1: Balance Parameters: Multi-Source Comparison (SCF, MRI/NHCS, Prolific, Nilson)

Parameter	SCF	MRI/NHCS	Prolific	Nilson
<i>Panel A: Revolvers (Balance Carriers)</i>				
Deposits (\$)	3,740	5,500	1,100	—
CC Outstanding (\$)	3,430	—	880	6,657*
CC Spending (\$/yr)	3,816	4,056	2,820	\$0*
<i>Panel B: Transactors (Non-Carriers)</i>				
Deposits (\$)	23,520	21,998	9,000	—
CC Outstanding (\$)	\$0	—	—	\$0*
CC Spending (\$/yr)	18,384	10,206	8,400	25,602*

Notes: SCF values are survey medians from the Survey of Consumer Finances. MRI/NHCS: deposit medians from the National Household Credit Survey (NHCS); CC spending medians from MRI Simmons; neither source covers CC outstanding. Prolific values are survey medians from the Prolific Academic panel. *Nilson Report (2021–2022) values are allocated under the strong assumption that revolvers carry 100% of CC outstanding and transactors account for 100% of CC spending; Nilson does not include deposit data. Population shares: 43.3% revolvers, 56.7% transactors.

Table C.2: Checking Deposit Balances from NHCS: Mean versus Median

Consumer Type	Mean Balance	Median Balance	Ratio (Mean/Median)
Transactors	\$87,860	\$21,998	3.99
Revolvers	\$32,769	\$5,500	5.96
Pooled	\$61,297	\$10,999	5.57

Notes: This table compares mean and median checking deposit balances from the NHCS for transactors (consumers who pay credit card balances in full) and revolvers (consumers who carry balances). The pooled estimate combines both groups. The large mean-to-median ratios reflect the highly right-skewed distribution of deposit balances.

Rationale for Using Medians Deposit and credit card balances are highly right-skewed (mean-to-median ratios of 4 to 6; Table C.2), so medians better reflect the typical account holder’s balance than an average inflated by outliers. In unreported results, using means

leaves the qualitative cross-selling findings unchanged, though implied marginal cost levels shift as expected.

D Likelihood Construction Details

The main text decomposes the Prolific cardholder likelihood into three blocks: a static block $L^{\text{stat}} = D_1 \cdot D_2$ that captures the consideration set and bank choice at each historical decision moment (Equation 14), a stated-preference block L^{hyp} for the hypothetical shocks (Equation 15), and a dynamic block $L^{\text{dyn}} = \lambda_{i1} \cdot M \cdot \lambda_{i2} \cdot F$ that adds the timing layer (Equation 16). For brevity, parts of this appendix collect the static and dynamic blocks under the shorthand $L^s \equiv L^{\text{stat}} \cdot L^{\text{dyn}}$ since both depend on the timing parameterization (t_1, t_2) of the past switching history. This appendix gives the closed-form expressions for D_k , M , and F , enumerates the four first-account configurations and explains why specific terms drop in each, derives the closed-form integrals for censored switch times, formalizes the non-cardholder contribution, and lays out the regime-conditional formulas for the hypothetical-shock likelihood.

D.1 Closed-Form Expressions for the Static and Dynamic Terms

For any intensity matrix A , let $\tilde{A}$ denote the modified matrix with all off-diagonal elements zeroed out (i.e., no switches), and let $[A]_{x,y}$ denote the matrix element corresponding to a transition from state x to state y . We use these to write transitions that respect what we know happened during a given interval.

Static choice term D_k . The static block multiplies two switch-event probabilities $D_1 \cdot D_2$, where for each switch k ,

$$D_k = P(x_k^* \mid C_k, x(t_k^-), \theta_i) \cdot P(C_k \mid x(t_k^-), \theta_i), \quad (\text{D.1})$$

C_k is the observed consideration set, x_k^* is the observed bank choice, $P(x_k^* \mid C_k, x(t_k^-), \theta_i)$ is the logit choice probability conditional on the consideration set and the pre-switch state, and $P(C_k \mid x(t_k^-), \theta_i)$ is the probability of drawing that consideration set.

Intervening tenure M . The intervening tenure M is the matrix-exponential probability of the consumer's state evolution during the interval (t_2, t_1) between the two observed switches. Because the consumer's deposit account is fixed to the post-deposit-switch bank during this interval (no further deposit switches have occurred since t_1), we suppress deposit transitions in the generator while retaining the full credit-card

dynamics:

$$M = \left[\exp \left((\tilde{Q}_1 + Q_2) (t_1 - t_2) \right) \right]_{x(t_1^+), x(t_2^-)} . \quad (\text{D.2})$$

The matrix element corresponds to the transition from the post-deposit-switch state $x(t_1^+)$ to the pre-credit-card-switch state $x(t_2^-)$.

Final stability F . The final stability F is the matrix-exponential probability of remaining in the consumer's currently observed state x^* from the most recent switch through the survey date. Because both observed switches have already been accounted for, we suppress both products' transitions:

$$F = \left[\exp \left((\tilde{Q}_1 + \tilde{Q}_2) t_2 \right) \right]_{x^*, x^*} . \quad (\text{D.3})$$

The diagonal element gives the probability of remaining in x^* without any additional switches.

D.2 First-Account Cases

The likelihood in Equation 16 applies when neither the deposit nor the credit card is the consumer's first account of that type. When one or both accounts are first-time adoptions, certain terms drop from the likelihood. We focus throughout on the case where the consumer has held her current deposit account longer than her current credit card ($t_1 > t_2$); the symmetric case follows by relabeling products. Table D.3 summarizes the four configurations, and we explain each below.

Table D.3: Likelihood Terms by First-Account Status

Case	Deposit	CC	λ_1	D_1	M	λ_2	D_2	F
1	Not First	Not First	✓	✓	✓	✓	✓	✓
2	First	Not First		✓	✓	✓	✓	✓
3	Not First	First	✓	✓			✓	✓
4	First	First		✓			✓	✓

Two types of terms can drop from the likelihood, each for a distinct economic reason:

First, arrival rates drop for first accounts. The arrival rate λ_k governs the Poisson process of *reconsideration* opportunities, the rate at which an existing account holder reconsiders their current product. As discussed in the main text, λ_k conflates genuine inattention with the fixed costs of switching, and both frictions contribute to low observed switching rates. For first-time adoption, the relevant frictions may be quite different: the fixed costs of opening a new account may be much smaller than those of switching an

existing one, so consumers may consider acquiring a new product far more frequently than λ_k would imply. Since λ_k conflates these frictions, we condition on the adoption having occurred and retain only the choice component D_k , which captures bank selection and consideration set formation at the time of adoption. In Case 2, λ_1 drops because the deposit is a first adoption; in Case 3, λ_2 drops because the credit card is a first adoption; in Case 4, both drop.

Second, the intervening tenure term drops when the more recently switched product is first. The term M (Equation D.2) is the probability the consumer's state follows the specified dynamics between the deposit switch at t_1 and the credit card switch at t_2 . During this interval, the consumer holds the deposit account at the chosen bank while potentially switching credit cards under the generator $\tilde{Q}_1 + Q_2$. This term is informative only when the consumer actually held a credit card throughout the intervening period. When the credit card is a first account (Cases 3 and 4), the consumer did not hold any credit card during $[t_1, t_2]$, so the credit card state is the outside option ($x_2 = 0$) throughout with certainty. Because the consumer's state path during the intervening period is trivially determined, M contributes no additional information and drops from the likelihood.

Third, the final stability term and choice probabilities always appear. The final stability term F (Equation D.3) captures the probability of no further switches between the most recent event and the survey date, and is always relevant regardless of whether accounts are first adoptions. The choice probabilities D_k , which capture which bank was selected and what consideration set was formed, are likewise always observed and always enter the likelihood.

D.3 Censored Switch Times

When one or both account switches occurred more than $T_{\max}$ years ago, the exact timing is censored and we must integrate over all possible switch times. Because the CTMC yields matrix exponential terms, these integrals admit closed-form solutions.

Deposit switch censored ($t_1 > T_{\max}$, $t_2 \leq T_{\max}$). When the deposit switch is censored but the credit card switch is observed, we integrate the likelihood over all possible deposit switch times beyond the lookback window:

$$L^s(t_1, t_2 | \theta) = \int_{T_{\max}}^{\infty} L(u, t_2 | \theta) du$$

The key step is integrating the matrix exponential in the intervening-tenure term M (Equation D.2) over the censored horizon. Using the identity

$$\int_{T_{\max}}^{\infty} \exp(A u) du = -A^{-1} \exp(A T_{\max})$$

for a matrix A with strictly negative real eigenvalues, M is replaced by:

$$M_{\text{dep}}^{\text{cens}} = - \left[(\tilde{Q}_1 + Q_2)^{-1} \exp((\tilde{Q}_1 + Q_2)(T_{\max} - t_2)) \right]_{x(t_1^+) \rightarrow x(t_2^-)}$$

The matrix inverse $(\tilde{Q}_1 + Q_2)^{-1}$ arises from integrating the matrix exponential over the semi-infinite censored horizon. All other likelihood terms, including which arrival-rate and choice terms appear, depend on first-account status exactly as in Table D.3.

Both switches censored ($t_1 > T_{\max}$, $t_2 > T_{\max}$). When both switches are censored, we perform a double integral over all possible switch times:

$$L^s(t_1 > T_{\max}, t_2 > T_{\max} | \theta) = \int_{T_{\max}}^{\infty} \int_r^{\infty} L(u, r | \theta) du dr$$

Now both M and F change. The intervening tenure becomes:

$$M_{\text{both}}^{\text{cens}} = - \left[(\tilde{Q}_1 + Q_2)^{-1} \right]_{x(t_1^+) \rightarrow x(t_2^-)}$$

This arises from $M_{\text{dep}}^{\text{cens}}$ when the upper integration limit for the credit card switch time also extends to infinity: the inner integral over the deposit switch time u , evaluated at its lower limit $u = r$ where $\exp(\tilde{Q}_1(u - r)) = I$, leaves only the matrix inverse from the outer integral over r . The final stability becomes:

$$F_{\text{both}}^{\text{cens}} = \left(-[\tilde{Q}_1 + \tilde{Q}_2]_{x^*, x^*} \right)^{-1} \exp \left([\tilde{Q}_1 + \tilde{Q}_2]_{x^*, x^*} \cdot T_{\max} \right)$$

which integrates the scalar sojourn probability over $[T_{\max}, \infty)$. As with the single-censored case, the first-account adaptations from Table D.3 apply unchanged: the arrival-rate and choice terms adapt based on which accounts are first accounts, while M and F are replaced by their censored counterparts.

D.4 Non-Cardholders

Prolific respondents who currently hold a checking account but no credit card are loaded into the estimator through a separate static-only contribution. The survey did

not elicit deposit-side switching dynamics or hypothetical-shock responses from this group, so they enter the likelihood only through the choice probability that they hold the credit-card outside option conditional on their checking-account holding:

$$L^{\text{nc}}(\theta_i) = \sum_{C \in 2^{\mathcal{J}}} P(x_2^* = 0 \mid C, x_1^*, \theta_i) \cdot P(C \mid x_1^*, \theta_i), \quad (\text{D.4})$$

which marginalizes over the unobserved credit-card consideration set C . None of the dynamic terms (λ_{ik} , D_k , M , or any component of F) and none of the hypothetical-shock terms appear; the data needed to evaluate them are not collected for non-cardholders. This contribution disciplines how attractive the credit-card outside option must be relative to the inside options offered by each bank.

D.5 Hypothetical-Shock Likelihood

Section V.C introduces the stated-preference likelihood

$$L^{\text{hyp}}(\theta_i) = \sum_C P(V, W, R, A \mid C, \theta_i) P(C \mid x^*, \theta_i).$$

This appendix derives the regime-conditional formulas for the conditional probability $P(V, W, R, A \mid C, \theta_i)$.

We model each respondent as drawing a single set of Gumbel preference shocks $\{\varepsilon_{ijk}\}$ at the time of the question; she then reports her hypothetical choices across all four scenarios using that draw. Define the consumer's choice probabilities under four progressively-degraded utility states for the focal product: the baseline (s_j), the unbundling shock (s'_j , where the ideal bank V no longer offers the complementary product $-k$), the price shock (s''_j , where the price at V becomes less attractive by Δp), and the complete removal of the focal product at V (s'''_j). Each s_j is a logit choice probability over the latent consideration set C .

Whether the price shock or the unbundling shock dominates the consumer's utility loss at the ideal bank depends on the realized $(\alpha_{ik}, \chi_{ik}^u)$. In the **price-dominant** case ($\alpha_{ik}\Delta p > \chi_{ik}^u \mathbb{1}\{V = x_{-k}^*\}$), any consumer who would switch under the unbundling shock would also switch under the price shock. The conditional probability collapses to

a stack of utility-gap differences:

$$P(V, W, R, A \mid C, \theta_i) = \begin{cases} s_0 & V = 0, \\ s_A''' - s_A'' & V = R = W \neq A, \\ s_A'' - s_A' & V = R \neq W = A, \\ s_A' - s_A & V \neq R = W = A, \\ 0 & \text{otherwise.} \end{cases}$$

Each non-zero case isolates the mass of consumers whose Gumbel draws place them in a particular utility window: just severe enough to trigger a switch under the stronger shock but not the weaker one.

In the **unbundling-dominant** case ($\alpha_{ik}\Delta p < \chi_{ik}^u \mathbb{1}\{V = x_{-k}^*\}$), the roles of W and R swap and the corresponding stack of differences uses the unbundling-side rather than price-side ordering. Because $(\alpha_{ik}, \chi_{ik}^u)$ are heterogeneous, the marginal likelihood mixes over the two regimes through the population distribution of types, which is what gives the model positive likelihood for every observed response pattern (including patterns that any single (α, χ^u) pair would assign zero probability to).

D.6 Model-Implied Tenure Distributions

We derive the model-implied distribution of account tenures from the estimated CTMC parameters. The key object is the *backward recurrence time*: the elapsed time since the consumer last switched to their current bank, as observed in a cross-sectional survey.

D.6.1 Phase-Type Backward Recurrence Time

Consider the tenure of a consumer's deposit account at bank j . Define $\mathcal{S}_1(j) = \{x : x_1 = j\}$ as the set of states in which the consumer holds a deposit account at bank j , varying over credit card banks $x_2 \in \{0, \dots, J\}$. While the consumer remains within $\mathcal{S}_1(j)$, they may switch credit cards but their deposit tenure continues to accumulate.

The sub-generator $Q_{\mathcal{S}} = Q[\mathcal{S}_1(j), \mathcal{S}_1(j)]$ governs the dynamics within $\mathcal{S}_1(j)$. Because $Q = Q_1 + Q_2$, this sub-generator captures both the (zero) within-set deposit transitions and the credit card transitions that occur while the consumer remains at deposit bank j . Exit from $\mathcal{S}_1(j)$, that is, switching the deposit account to a different bank, is encoded in the off-diagonal blocks of Q that are excluded from $Q_{\mathcal{S}}$. The sojourn time in $\mathcal{S}_1(j)$ therefore follows a phase-type distribution $\text{PH}(\alpha, Q_{\mathcal{S}})$, where α is the entry distribution described below.

D.6.2 Entry Distribution

The entry distribution α describes the composition of states within $\mathcal{S}_1(j)$ at the moment a consumer switches *into* deposit bank j . This distribution matters because a consumer's expected tenure at bank j depends on her credit card state x_2 at the time of entry. A consumer who enters as a cross-holder (j, j) faces lower exit rates because of the utility complementarity χ^u , and therefore tends to remain longer than a consumer who enters holding a credit card elsewhere. The entry rate into state $(j, x_2) \in \mathcal{S}_1(j)$ is:

$$\alpha_{x_2} = \sum_{j' \neq j} \bar{s}(j', x_2) \cdot q_1^{(j', x_2) \rightarrow (j, x_2)}$$

where $\bar{s}(j', x_2)$ is the steady-state probability of state (j', x_2) and $q_1^{(j', x_2) \rightarrow (j, x_2)}$ is the deposit transition rate from bank j' to j (holding credit card bank x_2 fixed). Normalizing so that $\sum_{x_2} \alpha_{x_2} = 1$ gives the probability distribution over credit card states at the moment of entry. Credit card tenure at bank j is defined analogously, with $\mathcal{S}_2(j) = \{x : x_2 = j\}$.

D.6.3 Backward Recurrence Distribution

In a stationary CTMC, the backward recurrence time B (the time elapsed since the most recent entry into $\mathcal{S}_k(j)$) has the following distribution:

$$f_B(t) = \frac{\alpha' \exp(Q_S t) \mathbf{1}}{\mu} \quad (\text{D.5})$$

$$P(B > t) = \frac{-\alpha' Q_S^{-1} \exp(Q_S t) \mathbf{1}}{\mu} \quad (\text{D.6})$$

where $\mathbf{1}$ is the vector of ones and $\mu = -\alpha' Q_S^{-1} \mathbf{1}$ is the mean sojourn time in $\mathcal{S}_k(j)$.

Intuitively, $\exp(Q_S t) \mathbf{1}$ gives the probability of remaining within $\mathcal{S}_k(j)$ for at least t years starting from each state. Averaging over the entry distribution α and normalizing by μ converts this into the backward recurrence density observed in a cross-section. The survival function in Equation (D.6) exploits the commutativity $Q_S^{-1} \exp(Q_S t) = \exp(Q_S t) Q_S^{-1}$ to precompute $Q_S^{-1} \mathbf{1}$ once.

D.6.4 Aggregation

We compute per-bank tenure distributions separately for each bank $j = 1, \dots, J$ and each product k . The overall tenure distribution for product k is the share-weighted average across banks:

$$f_B^{\text{overall}}(t) = \sum_{j=1}^J w_j f_B^{(j)}(t), \quad w_j = \frac{\bar{s}_j}{\sum_{j'=1}^J \bar{s}_{j'}}$$

where $\bar{s}_j$ is the marginal steady-state share of bank j for product k (excluding the outside option). We further combine the transactor and revolver subsamples using their sample frequency weights to produce a single predicted tenure distribution for comparison with the survey data. Because the demand parameters θ_i are heterogeneous, each type's tenure distribution is computed separately and then averaged over the Monte Carlo draws from the heterogeneity distribution.

D.7 Parametric Form of the Heterogeneity Distribution

We specify the bank-preference intercepts in levels and the structural primitives in logs, since α, χ^u, λ are positive on theoretical grounds. The heterogeneity distribution has block-diagonal structure:

1. **Bank preferences (levels).** For each bank j , the consumer's preference vector is jointly normal,

$$\begin{pmatrix} v_{ij1}^c \\ v_{ij2}^c \\ v_{ij1}^u \\ v_{ij2}^u \end{pmatrix} \sim \mathcal{N}(\bar{v}_j, \Sigma_v),$$

independent across consumers i and banks j , with Σ_v shared across banks. The off-diagonal entries of Σ_v allow consumers' preferences to correlate within a bank across products and across the consideration and utility margins, which is empirically central: we estimate the consideration–consideration correlation $\rho(v_1^c, v_2^c)$ to be close to one.

2. **Price sensitivity and utility complementarity (logs, per product).** For each product k , $(\log \alpha_{ik}, \log \chi_{ik}^u)$ is jointly normal,

$$\begin{pmatrix} \log \alpha_{ik} \\ \log \chi_{ik}^u \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_{\alpha,k} \\ \mu_{\chi^u,k} \end{pmatrix}, \Sigma_{(\alpha, \chi^u),k}\right),$$

with within-product correlation freely estimated. The two product- k blocks $(\log \alpha_{i1}, \log \chi_{i1}^u)$ and $(\log \alpha_{i2}, \log \chi_{i2}^u)$ are independent of each other. The independence reflects the survey design: each Prolific respondent is randomized into hypothetical-shock questions for only one of the two products, so the within-product joint distribution

is identified but the cross-product joint distribution is not.

3. **Arrival rates (logs).** The pair of log arrival rates is jointly normal,

$$\begin{pmatrix} \log \lambda_{i1} \\ \log \lambda_{i2} \end{pmatrix} \sim \mathcal{N}(\mu_\lambda, \Sigma_\lambda),$$

with their own 2×2 covariance. Tenure distributions for both products in the same consumer's life identify the cross-product correlation in arrival rates.

The consideration complementarity χ_k^c is treated as a homogeneous parameter, with population mean $\bar{\chi}_k^c$ identified by Fact 3 but no consumer-level variance. The consumer-level shifters η_{ik}^c and η_{ik}^u used in earlier specifications are not separately identified from ν^c and ν^u , so we drop them. We use log specifications for α, χ^u, λ because the model imposes positivity on these parameters; the bank-preference intercepts ν have no sign restriction and are modeled in levels.

D.8 MC-EM Algorithm and Implementation Details

This appendix gives the full pseudocode and the practical features of the implementation glossed over in the main text. Let X_i denote the data for observation i , $\phi(\cdot \mid \mu, \Sigma)$ the multivariate normal PDF with mean μ and covariance Σ , and $\ell = 1, 2, \dots, L$ the iteration index.

Algorithm.

1. **Initialize.** Set $\bar{\theta}^{(1)}$ and $\Sigma^{(1)}$ to starting values described in the initialization paragraph below, and set each observation's proposal distribution to the population prior, $\bar{\theta}_i^{(1)} = \bar{\theta}^{(1)}$ and $\Sigma_i^{(1)} = \Sigma^{(1)}$.
2. For iterations $\ell = 1, \dots, L$ or until convergence:
 - (a) **Draw.** For each consumer i , draw R proposals $\theta_{i,r}^{(\ell)} \sim \mathcal{N}(\bar{\theta}_i^{(\ell)}, \Sigma_i^{(\ell)})$. The proposal distribution is observation-specific and is updated each iteration based on the previous iteration's posterior summary for i .
 - (b) **Evaluate.** Compute $L(X_i \mid \theta_{i,r}^{(\ell)})$, the population prior density $p_r^{(\ell)} = \phi(\theta_{i,r}^{(\ell)} \mid \bar{\theta}^{(\ell)}, \Sigma^{(\ell)})$, and the proposal density $q_{i,r}^{(\ell)} = \phi(\theta_{i,r}^{(\ell)} \mid \bar{\theta}_i^{(\ell)}, \Sigma_i^{(\ell)})$.
 - (c) **E-step (importance weights).** Form normalized posterior weights

$$w_{i,r}^{(\ell)} = \frac{L(X_i \mid \theta_{i,r}^{(\ell)}) p_r^{(\ell)} / q_{i,r}^{(\ell)}}{\sum_{t=1}^R L(X_i \mid \theta_{i,t}^{(\ell)}) p_t^{(\ell)} / q_{i,t}^{(\ell)}}.$$

The ratio $p_r^{(\ell)} / q_{i,r}^{(\ell)}$ reweights each draw from the proposal back to the population prior; multiplying by the likelihood yields draws from the posterior $p(\theta_i \mid X_i, \bar{\theta}^{(\ell)}, \Sigma^{(\ell)})$ up to a normalizing constant.

- (d) **M-step (population parameters).** Update $\bar{\theta}^{(\ell+1)}$ and $\Sigma^{(\ell+1)}$ as weighted first and second moments of the simulated types:

$$\begin{aligned}\bar{\theta}^{(\ell+1)} &= \frac{1}{N} \sum_{i=1}^N \sum_{r=1}^R w_{i,r}^{(\ell)} \theta_{i,r}^{(\ell)}, \\ \Sigma^{(\ell+1)} &= \frac{1}{N} \sum_{i=1}^N \sum_{r=1}^R w_{i,r}^{(\ell)} (\theta_{i,r}^{(\ell)} - \bar{\theta}^{(\ell+1)}) (\theta_{i,r}^{(\ell)} - \bar{\theta}^{(\ell+1)})^\top.\end{aligned}$$

- (e) **Update proposal.** Refresh each consumer's proposal using her own posterior moments:

$$\begin{aligned}\bar{\theta}_i^{(\ell+1)} &= \sum_{r=1}^R w_{i,r}^{(\ell)} \theta_{i,r}^{(\ell)}, \\ \Sigma_i^{(\ell+1)} &= \sum_{r=1}^R w_{i,r}^{(\ell)} (\theta_{i,r}^{(\ell)} - \bar{\theta}_i^{(\ell+1)}) (\theta_{i,r}^{(\ell)} - \bar{\theta}_i^{(\ell+1)})^\top.\end{aligned}$$

Step 2(c) is the E-step implemented by importance sampling; step 2(d) is the M-step in closed form, since the weighted log-likelihood of the simulated types has a Gaussian-conjugate maximizer in $(\bar{\theta}, \Sigma_\theta)$. The algorithm departs from Train (2008) in step 2(a): rather than drawing all types from the population prior, we draw each consumer's types from her own posterior summary from the previous iteration, sharply reducing the variance of the importance weights and allowing the E-step to converge with a manageable number of draws per observation.

The remainder of this appendix describes practical features of the implementation needed to run reliably on samples of our size.

Phased estimation schedule. We estimate the model in four phases. The first phase fits only the static moments (the static-choice and hypothetical-choice likelihoods), holding the arrival rates λ_{ik} at an initial calibration and estimating only the preference and complementarity blocks. The second phase adds the dynamic past-switching likelihood with the preference and χ parameters frozen at their phase-one estimates, so that the first iteration can update $\bar{\lambda}$ and σ_λ in isolation. The third phase unfreezes preferences while keeping the complementarity block fixed. The fourth phase estimates all parameters jointly. The phased schedule avoids the instability that arises from simultaneously

updating parameters whose moments compete in the early iterations.

Adaptive proposal covariance. The proposal covariance $\Sigma_i^{(\ell)}$ in step 2(a) of the algorithm is a convex combination of the previous-iteration posterior covariance for observation i and the current-iteration population covariance, with the mixing weight adapted based on the effective sample size of the importance weights. When weights are concentrated on few draws (low ESS), we shrink the proposal toward the population covariance to broaden the search; when weights are close to uniform (high ESS), we rely more on the observation-specific posterior covariance. The mixing weight is also scaled by a diffusion factor greater than one to prevent the proposal from collapsing below the true posterior scale.

Draw regeneration. Each MC-EM iteration can either re-simulate draws from the current proposal or reuse the draws from the previous iteration with updated importance weights. We regenerate draws when any of three conditions holds: the effective sample size $(\sum_r w_{i,r}^{(\ell)})^2 / \sum_r (w_{i,r}^{(\ell)})^2$ drops below a per-consumer threshold, the population parameters $\bar{\theta}^{(\ell)}$ have drifted far enough from their values at the last regeneration, or a fixed number of iterations has elapsed since the last regeneration. Reusing draws across consecutive iterations is exact up to the importance ratio; regeneration is needed only when the proposal has moved far enough that the importance weights degenerate.

Block-diagonal covariance constraint. The heterogeneity covariance is block-diagonal across the four blocks defined in Appendix D.7: the bank-preference block Σ_v (4×4 , shared across banks); the per-product log blocks $\Sigma_{(\alpha, \chi^u), 1}$ and $\Sigma_{(\alpha, \chi^u), 2}$ (2×2 each, independent of each other); and the joint log arrival-rate block Σ_λ (2×2). Cross-block correlations are constrained to zero. After each M-step we project the updated covariance onto this block structure and enforce positive-definiteness by eigenvalue flooring.

Aggregate-share refinement. After the MC-EM iterations converge, we adjust the bank-specific utility intercepts $\bar{v}_{jk}^u$ to match aggregate market shares from the MRI-Simmons cross-section. Holding all other parameters fixed at their EM estimates, we iterate

$$\bar{v}_{jk}^{u,(\ell+1)} = \bar{v}_{jk}^{u,(\ell)} + \log s_{jk}^{\text{MRI}} - \log \hat{s}_{jk}(\bar{v}^{u,(\ell)}),$$

where $\hat{s}_{jk}(\bar{v}^u)$ is the model-implied steady-state share for bank-product cell (j, k) , averaged over the estimated type distribution by Monte Carlo. The contraction converges

quickly because the residual share gap is small after EM, and it leaves the heterogeneity, complementarity, and switching-rate estimates unchanged.

Initialization. We initialize $\bar{\theta}^{(1)}$ at moderate values consistent with the reduced-form moments. In particular, bank-preference intercepts match observed market shares and arrival rates match the average reciprocal tenure observed in the survey. $\Sigma^{(1)}$ is initialized as a block-diagonal matrix with moderate diagonal entries and identity correlation structure within blocks. Convergence is declared when the relative change in the population log-likelihood and in $\bar{\theta}$ falls below prespecified tolerances across five consecutive iterations.