Financial Supermarkets? Cross-Selling in Retail Banking

Shengmao Cao Lulu Wang *

November 2, 2026

Abstract

Many consumers choose banks, not individual financial products. We study how complementarity across financial products shapes competition and regulation in retail banking. Across three U.S. consumer surveys, consumers are three to five times more likely to hold credit cards, auto loans, mortgages, and brokerage accounts at their deposit bank than independent product choice implies. Using the 2019 transfer of Walmart's credit-card portfolio to Capital One, we show that winning a consumer's card business helps win their deposits, providing quasi-experimental evidence of complementarity. To separate the mechanisms behind cross-holding, we field an original survey that records consumers' bank search and switching histories and runs discrete choice experiments, and we use it to estimate an equilibrium model of joint deposit and credit-card choice. We find that complementarity accounts for about three-quarters of excess cross-holding, with the remaining quarter reflecting correlated preferences. Complementarity changes how banks price and how policy transmits across markets. Rewards credit cards are loss leaders: they earn 7 percent of banks' accounting profits but generate 22 percent of franchise value by attracting and retaining depositors. Interchange-fee regulation that compresses rewards raises deposit rates and costs banks depositors, so about a third of the profit loss falls on their deposit business. A bank merger that creates complementarity brings consumers little gain, because the merged bank captures most of the synergy by worsening its prices. Analyses that treat each product market in isolation miss these effects.

*Cao: Northwestern Kellogg (email: shengmao.cao@kellogg.northwestern.edu). Wang: Northwestern Kellogg (email: lulu.wang@kellogg.northwestern.edu). We thank our discussants Christoph Herpfer, Stefan Lewellen, and Xu Lu, and participants at the Midwest IO Fest, IIOC, the Georgia Tech–Atlanta Fed Household Finance Conference, the Stanford SITE conference, and the Holden Conference at Indiana University for their helpful feedback. We gratefully acknowledge funding from the Asset Management Practicum at Northwestern University through the Simpson Querrey Research Fund. Marziyeh Aghaei, Lucia Bertoletti, 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 reported more than $4 billion in annual synergies from cross-selling products to retail consumers, with roughly half arising from its credit-card and consumer-banking operations (JPM, 2015). Although the practice drew widespread criticism after the 2016 Wells Fargo fake-accounts scandal, it remains a defining feature of the retail banking landscape.

Despite this practical significance, academic research and regulatory policy have largely examined one banking product market at a time. The literature on bank franchise value, for example, focuses primarily on deposit market power (Egan et al., 2017; Drechsler et al., 2021), while work on credit-card markets studies how interest-rate caps and interchange-fee regulation affect credit-card use (Agarwal et al., 2015; Wang, 2026). Bank merger reviews have historically assessed competitive harm through deposit market concentration alone (U.S. Department of Justice, 1995; U.S. Department of Justice and Federal Trade Commission, 2023). By design, these analyses do not capture how a regulatory change or a shift in competitive structure in one product market spills over into others. When such spillovers are large, single-market analysis has major blind spots and can yield miscalibrated policy prescriptions.

In this paper, we examine how consumers build portfolios of financial products and assess how cross-product complementarities shape competition and the effects of policy in retail banking. First, using three U.S. consumer surveys, we document empirical patterns of cross-holding across banks, covering deposit accounts, credit cards, mortgages, auto loans, and brokerage services. Second, we develop and estimate a model of deposit-account and credit-card choice that embeds the mechanisms through which consumers come to consolidate their products at a single bank. Third, we use the estimated model to quantify the drivers of cross-holding and to derive new insights into the sources of banks' franchise value, policy spillovers between deposit and credit-card markets, and the competitive effects of bank mergers.

Our analysis draws on three consumer surveys: the Survey of Consumer Finances (SCF), the MRI-Simmons USA Study, and an original survey of U.S. adults conducted on the Prolific platform. The SCF and MRI are large, nationally representative studies of consumers' financial product portfolios; the Prolific survey complements them with data on 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, 34% of consumers with one checking account and one credit card hold both at the same bank; the analogous shares are 15% for auto loans, 14% for mortgages, and 5% for brokerage accounts. These rates are three to five times those implied by independent product choice. Cross-holding matters most for smaller banks. Depositors at megabanks such as Chase and Bank of America are on average 9 times more likely to hold the bank's credit card than non-depositors, while at regional banks such as PNC or Citizens this ratio exceeds 60. Most of the regional banks' credit-card customers also hold a deposit account there, whereas megabanks attract many standalone cardholders.

The pattern of who cross-holds suggests that relationship banking, the standard explanation for why consumers borrow where they bank, is only part of the story. The relationship-banking literature shows that deposit relationships 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 low-credit-score revolvers, who stand to gain most from this channel, cross-hold at nearly the same rate as high-credit-score transactors, who pay their balances in full, face minimal underwriting barriers, and thus have little to gain. Asymmetric information may contribute to cross-holding, but it cannot be the whole story; the patterns point toward preference heterogeneity, product complementarity, or both.

Consumers may cross-hold because of correlated preferences or complementarity. Under correlated preferences, a consumer drawn to Chase's branch network and reputation may prefer it for both deposits and credit cards, generating cross-holding even though neither choice affects the other. Under complementarity, holding one product at a bank makes a consumer more likely to choose that bank for the next one. The two mechanisms imply fundamentally different incentives for banks and regulators (Gentzkow, 2007). If preferences are merely correlated, cross-holding is a form of consumer sorting. Winning a customer's first product does not help a bank win the second, and shocks to one market stay in that market. If the products are true complements, winning the first raises demand for the second. Loss-leader pricing becomes profitable, and competitive or regulatory shocks transmit across product markets.

We establish that complementarity is at work using a natural experiment, the transfer of a large credit-card portfolio from one bank to another. In October 2019, Walmart moved its consumer credit-card program from Synchrony to Capital One, and millions of cardholders acquired a Capital One credit-card relationship for reasons unrelated to

their own bank preferences. In the three years following the transfer, Capital One's deposit share among Walmart cardholders rises by about 5 percentage points relative to consumers without a Walmart card and to Target cardholders, with no differential movement at placebo banks. Scaled by the first stage, these estimates imply that Capital One converts roughly 8 of every 100 newly acquired credit-card holders into deposit customers.

Complementarity itself can operate through two channels with very different welfare implications. Holding one product at a bank may raise the chance that a consumer considers that bank when shopping for the other. This consideration channel resembles steering and need not make the consumer better off. Alternatively, consumers may derive direct utility from consolidating their financial products at a single institution. This utility channel makes consumers better off. The Walmart experiment, however, identifies only the combined effect of the two channels.

We designed the Prolific survey to gather direct evidence on correlated preferences, consideration complementarity, and utility complementarity. 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, and the factors they found most important in the decision. We complement this with discrete choice experiments in which respondents report whether they would keep a focal product at their current bank or switch after a hypothetical cut in the savings rate, a cut in credit-card rewards, or, for revolvers, an APR increase. In a separate scenario, respondents report how they would respond if their other product had just moved to a different bank for reasons unrelated to the focal product.

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 choose neither, indicating that consideration is correlated across products. Consumers who already hold a product at a bank are even more likely to include that bank in their consideration set for the other product, showing that existing banking relationships themselves shape the search process. Finally, about 30% of consumers would move the focal product to follow the other account, roughly the same share that a median-to-large own-price shock would move. This implies that some consumers derive substantial 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 only occasionally 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 complementarity from existing product holdings, and direct utility complementarity between same-bank products. Existing models of dynamic deposit competition attribute infrequent switching to inattention alone (Egan et al., 2025). Our model adds switching costs and persistent unobserved heterogeneity in bank preferences, and these three forces together explain why consumers so rarely switch banks. Banks set prices in a Nash equilibrium to maximize the present discounted value of future profits, internalizing the cross-product spillovers by which a price cut in one market raises demand in the other.

We estimate the model by maximum likelihood using a Monte Carlo EM algorithm (Train, 2008). The survey instruments separately target the three mechanisms and the frictions that slow switching. Consideration sets identify correlated consideration and consideration complementarity. Consumers' willingness to follow the other product when it moves, benchmarked against their own-price responses, identifies utility complementarity. The *de novo* ideal-bank question, together with responses to price shocks, separates persistent preferences, switching costs, and inattention. The estimated model also fits moments we did not target. Out of sample, it reproduces the deposit-share gain in the Walmart experiment, and its implied margins are consistent with industry estimates.

The estimated model delivers new insights into why consumers cross-hold, the sources of bank franchise value, and the effects of pricing regulation and mergers. Complementarity accounts for about three-quarters of excess cross-holding, mostly through utility rather than consideration. Strong utility complementarity and limited consideration spillovers suggest that cross-selling is not merely steering, because consumers value holding both products at one bank once they have them. Correlated preferences account for the remaining quarter. Attributing all cross-holding to complementarity would therefore overstate how many customers a bank can expect to convert to a second product by winning the first.

Complementarity drives banks to price rewards credit cards as loss leaders to attract depositors. U.S. banks issue credit cards whose rewards are generous relative to the interchange income they receive (Drechsler et al., 2025). Existing explanations of this generosity stay within the card market: consumers respond strongly to rewards (Wang, 2026), and naive cardholders end up paying interest that funds the rewards earned by others (Agarwal et al., 2026). Our estimates show that much of this generosity originates outside the card market, because rewards cards are loss leaders for deposit accounts. They generate about 7% of banks' accounting profits, yet eliminating them would cost banks around 22% of franchise value, because they attract and retain depositors with

large balances. At the megabanks, rewards-card margins are about half the markups the cards would earn as standalone products.

Complementarity also transmits regulation across markets, and in a direction that may be unexpected. Mukharlyamov and Sarin (2025) show that reductions in debit interchange fees led banks to raise checking-account fees, so one might expect a cut in credit-card interchange fees to worsen deposit terms as well. Our model predicts the opposite. When we simulate a one-percentage-point increase in the cost of transactor cards, the kind of shock interchange-fee regulation imposes, banks pass most of the cost through to rewards, cutting them by 81 to 93 basis points. Absent complementarity, the adjustment would stop there. With it, weaker card relationships put the deposit franchise at risk, and banks compensate depositors by raising deposit rates by up to 7 basis points (SE 3.0). The difference arises because deposits and credit cards are not formally bundled. A debit card comes with the checking account, but a consumer whose credit card has become less attractive can move the deposit account elsewhere, which drives the bank to offer better deposit terms. Even with higher deposit rates, lower card rewards reduce deposit market shares. About a third of banks' profit loss from the regulation therefore falls on their deposit business rather than on the card market alone, rising to two-fifths at the banks that offer checking.

A cost shock to revolving credit moves deposit rates in the opposite direction. When we raise the cost of lending to revolvers by 10 percentage points, banks pass almost the entire increase through to APRs and cut deposit rates by about 1 basis point. Lower margins on revolver cards reduce the value of attracting revolvers through the deposit market, though the deposit-rate response is small because revolvers hold modest deposit balances.

Finally, we apply our framework to a horizontal merger between Citi and a super-regional bank and ask whether complementarity can serve as an efficiency defense. Absent complementarity, the merger is a conventional consolidation that results in worse prices and lower market share for the merging parties. With complementarity, the merger creates a synergy, because a Citi card can now be paired with a super-regional deposit account and vice versa. We estimate that at unchanged prices, this synergy increases the merged bank's deposit share by 10% (SE 1.0%) and its credit-card share by 15% (SE 1.6%). However, the merged bank captures the value of consolidation by cutting deposit rates by 8 basis points (SE 4.1) and raising APRs for credit-card revolvers by 17 basis points (SE 5.9). Credit-card rewards remain essentially unchanged. Consumers value the new opportunity to consolidate at one bank, and that value gives the merged bank more pricing power. After prices adjust, the merged bank's deposit share

and revolver share are unchanged, and only its transactor share rises, by an imprecisely estimated 23%. The synergy thus accrues mostly to the merged bank rather than to consumers. The lesson for merger review is that an efficiency defense based on complementarity requires evidence that the synergy will be passed on to consumers, and that pass-through can differ across product markets.

Our paper contributes to three strands of literature. First, we add to the industrial organization literature on financial markets by providing an equilibrium framework for multi-product bank competition and showing that demand-side complementarity is a significant source of bank franchise value. This literature has developed rich models of individual product markets, including deposit accounts (Egan et al., 2017; Honka et al., 2017; Egan et al., 2025), credit cards (Stango and Zinman, 2016; Nelson, 2025; Cuesta and Sepúlveda, 2021; Galenianos and Gavazza, 2022; Wang, 2026; Drechsler et al., 2025), and wealth management (Hortaçsu and Syverson, 2004; Hastings et al., 2017; Brown et al., 2023), with less attention to how banks compete across them. Within credit cards, existing models study either revolvers who borrow (Nelson, 2025; Galenianos and Gavazza, 2022) or transactors who pay in full and earn rewards (Wang, 2026). We model both segments, and the distinction is central to our core insights: transactors hold large deposits and respond to rewards, which makes rewards cards a loss leader for deposits. We show that competition across markets changes how one should think about the sources of bank franchise value, the incidence of regulation, and the evaluation of mergers. On mergers in particular, the classic insight is that combining complementary products lowers prices, because the merged firm internalizes cross-demand externalities (Economides and Salop, 1992; Song et al., 2017). We show that a bank merger can instead create complementarity, but that the merged bank captures most of the synergy rather than passing it through.

Second, we provide new evidence for demand-based theories of economies of scope in banking, most of which concern corporate banking (Hellmann et al., 2008; Laux and Walz, 2009; Ivashina and Kovner, 2011; Santikian, 2014; Qi, 2024). In retail banking, Aguirregabiria et al. (2024) show how banks' asset and liability sides complement each other through their branch networks, Allen et al. (2014, 2019) document the incumbent bank's substantial advantage in the mortgage market, and Basten and Juelsrud (2023) provides population-wide administrative evidence of deposit-to-loan conversion. We add novel quasi-experimental evidence from the Walmart portfolio transfer, separately identify the three demand-side mechanisms behind cross-holding, and derive their equilibrium implications using a model of joint deposit and credit-card competition.

Third, we contribute to the literature on survey-based demand identification (Petrin,

2002; Berry et al., 2004). Standard consumer-finance data sets do not distinguish the mechanisms behind cross-holding, so we design an original survey that collects consideration sets, search histories, and stated responses to hypothetical shocks, each targeted at a distinct mechanism. We follow best practices for discrete choice experiments in labor economics (Mas and Pallais, 2017, 2019; Caldwell et al., 2025) and tailor the hypothetical shocks to each respondent's balances and spending (Hastings and Tejeda-Ashton, 2008). These survey-based moments uncover important dimensions of consumer behavior that observational data alone cannot identify. The resulting estimates imply bank margins consistent with industry estimates and reproduce the deposit-share gain from the Walmart portfolio transfer, a quasi-experiment the estimation does not target.

Section II describes the data. Section III documents stylized facts on cross-holding and presents quasi-experimental and survey evidence on the mechanisms behind it. Section IV develops an equilibrium model of consumer choice and bank competition in the deposit-account and credit-card markets. Section V describes the estimation strategy, reports the estimates, and validates the model against untargeted external benchmarks. Section VI decomposes observed cross-holding, analyzes bank franchise value, and studies interchange-fee regulation and a bank merger. Section VII concludes.

II Data

Our analysis draws on three consumer surveys. The SCF measures cross-holding across the broadest set of financial products, the MRI-Simmons survey identifies the specific banks consumers use, and our original Prolific survey reconstructs the search processes, switching histories, and stated choices that separate the mechanisms behind cross-holding. We supplement the surveys with bank-level price data for the supply-side analysis.

Survey of Consumer Finances (SCF) Our first dataset is the Survey of Consumer Finances (SCF), a triennial survey conducted by the Federal Reserve Board and widely regarded as the most comprehensive source on the finances of U.S. households. We pool the 2019, 2022, and 2025 waves of the Full Public Data Set, which together cover 14,732 households, and weight each wave equally.[1] The SCF has the broadest product coverage of our three surveys, spanning deposit accounts, credit cards, auto loans, mortgages, brokerage services, and more. It anonymizes the financial institutions each household uses, which has limited its use in prior retail banking research. We exploit a feature of

[1]The 2025 wave no longer records which institution holds a household's credit card or auto loan, so our cross-holding measures for these two products pool the 2019 and 2022 waves only. We report account balances in 2025 dollars, deflated by the CPI-U.

its design: the survey assigns consistent institution codes across products within each household. These codes reveal whether a household holds multiple products at the same institution (for example, a checking account and a mortgage at bank A), even though the identity of bank A is never disclosed.

MRI-Simmons USA (MRI) Our second dataset, the MRI-Simmons USA Study (MRI), fills the SCF's main gap by naming the financial institutions consumers use. MRI is a large-scale, nationally representative survey of product ownership and demographics. We primarily use the 2021 and 2022 waves, covering approximately 100,000 consumers, and draw on the 2020 wave (approximately 40,000 consumers) for its data on consumer tenure at financial institutions.[2] The Walmart analysis in Section III additionally uses the 2016–2019 waves, which cover the years before the card-portfolio transfer. MRI covers fewer products than the SCF (deposit accounts, credit cards, and brokerage services), but every holding is linked to a named bank. Nine 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.

Both datasets share two limitations. First, as repeated cross-sections, they record static product portfolios rather than consumers' search and switching behavior over time. Second, they show who cross-holds but not why; equilibrium holdings alone cannot separate the mechanisms that generate them. Our Prolific survey is designed to fill both gaps.

Original Prolific Survey Our third dataset is an original survey of approximately 2012 U.S. consumers, fielded on the Prolific platform.[3] We fielded two waves, of 983 and 1,029 respondents. Both waves collect holdings, consideration sets, and switching histories, and we pool them for the descriptive evidence in Section III and for the structural histories in Section V. Estimation uses stated-preference responses from the second wave, whose instruments are comparable across respondents. The survey mirrors the MRI structure, collecting each consumer's deposit and credit-card holdings and distinguishing primary from other accounts.

The survey introduces two design features aimed at uncovering the mechanisms behind cross-holding. First, we reconstruct consumers' search and switching histories. For

[2]The 2020 wave is referred to as NHCS.

[3]The survey is available at `https://kellogg.qualtrics.com/jfe/form/SV_eXK0bVgGMnzCcGa`. It takes about ten minutes to complete, and each participant receives $5 upon completion.

each respondent's most recent primary account opening, we elicit the timing of the decision, the banks they considered, the prior accounts they held, and the factors they weighed most heavily.

Second, a stated-preference module measures price sensitivity and cross-product complementarities. We present a series of hypothetical shocks to the respondent's current holdings, each expressed in personalized dollar terms based on their reported balances and spending (Hastings and Tejeda-Ashton, 2008): a cut in their savings-account interest rate, a cut in their credit-card rewards, an increase in their credit-card APR (asked of revolvers, who carry a balance), and an unbundling shock in which we ask what they would do with one product if the other product had just moved to a different bank for unrelated reasons.[4] These questions resemble the "discrete choice experiment" design used to estimate the value of non-wage amenities and labor-supply elasticities (Mas and Pallais, 2017, 2019; Caldwell et al., 2025), and the resulting switching probabilities provide key empirical moments for estimating demand-side complementarities.

The Prolific sample lines up well with the other data sources and with market aggregates. Table 1 shows that mean ages, median incomes, and product ownership rates are broadly comparable across our three surveys, with the Prolific sample modestly overrepresenting college-educated consumers, consistent with known selection patterns on online survey platforms. Panels C1 and C2 show large differences in deposit balances and credit-card spending between revolvers (who carry a balance from month to month) and transactors (who pay in full), a pattern we exploit to calibrate the supply-side model. Appendix Table A.1 shows that the Prolific sample also 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.

Additional Data The supply-side analysis requires the three prices banks set (deposit rates, credit-card reward rates, and credit-card APRs), at which we evaluate each bank's first-order conditions. We construct deposit rates from WRDS bank holding company filings, dividing annual interest expense on domestic deposits by year-end domestic deposits. We construct bank-level reward rates by dividing rewards expenses in bank financial filings by Nilson Report payment volumes. We obtain average credit-card APRs

[4]The survey design mitigates framing and order effects. We elicit consideration sets before asking about prior banking relationships, so that recall of past accounts does not contaminate consideration reports, and we randomize the order of the unbundling and price shocks, so that any anchoring between the two can be detected and netted out by comparing responses across the two orderings.

by bank from the CFPB Terms of Credit Card Plans (TCCP) survey.

III Stylized Facts

We document five stylized facts about cross-holding in retail banking. Together, these facts provide the identifying variation for the structural model in Section IV.

Fact 1. *Cross-holding is pervasive, with consumers approximately* 3–5 *times more likely to hold products at the same bank than independent product choice implies. Smaller, regional banks rely on cross-holding the most.*

We measure the aggregate cross-holding rate, the fraction of consumers who hold a checking account and a second financial product at the same bank, across four products in the SCF, MRI, and Prolific data. Figure 1 plots the rate for checking accounts paired with each of the other products. Approximately one-third of consumers hold their checking account and credit card at the same bank; the analogous rates for auto loans, mortgages, and brokerage accounts are 15, 14, and 5 percent, respectively. The raw rates differ widely across products, yet in every category consumers are approximately 3–5 times more likely to hold both products at the same bank than would be expected if they chose products independently. We construct this benchmark by comparing the joint probability of holding both products at the same bank with the product of the corresponding marginal probabilities, computed at banks that operate in both markets.[5] Appendix Table A.2 reports the estimated ratios for each product and dataset, along with 95% confidence intervals; all lie in the range of 3–5.

The pattern of who cross-holds suggests that relationship banking, the standard explanation for why consumers borrow where they bank, is only part of the story. In this literature, deposit relationships generate information about borrower risk, so cross-holding should be concentrated among consumers for whom this information is most valuable in underwriting (Puri et al., 2017; Agarwal et al., 2018; Parlour et al., 2022). Instead, Appendix Figure A.4 shows nearly identical cross-holding rates across income, education, and credit-score levels, and between credit-card transactors and revolvers, in all three datasets. High-credit-score transactors, who face minimal underwriting barriers and thus have little to gain from this channel, cross-hold as much as the low-score revolvers who stand to gain the most. Asymmetric information may thus contribute to

[5]This comparison requires knowing each product's bank, so we compute it only for the MRI and Prolific datasets; unlike the aggregate rates, it includes consumers with multiple accounts. In Prolific, we elicit both the primary deposit account and the primary credit card directly. MRI identifies a consumer's main credit card but not their main deposit account, so when a consumer has multiple deposit banks we treat each banking relationship as equally likely to be primary.

cross-holding, but it cannot be the whole story; the patterns point toward preference heterogeneity, product complementarity, or both.

Smaller banks depend on cross-holding the most. Figure 2 plots each bank's credit-card market share separately among its depositors and non-depositors in the full MRI sample. Depositors at "megabanks" like Chase and Bank of America are on average 9 times more likely than non-depositors to hold the bank's credit card (5× for Chase, 13× for Bank of America); at regional banks such as PNC or Citizens, the ratio exceeds 60. Most of these smaller banks' credit-card customers also hold a deposit account there, whereas megabanks attract many standalone cardholders. The pattern is symmetric: deposit-account market shares are similarly elevated among a bank's credit-card holders, especially at smaller banks. This "relationship multiplier" is central to customer acquisition at smaller banks.

Fact 2. *Quasi-experimental evidence from the 2019 Walmart card-portfolio transfer establishes complementarity: winning a consumer's credit-card relationship raises the bank's probability of also winning their deposit account.*

Cross-sectional patterns of cross-holding cannot distinguish complementarity from correlated preferences, and naturally occurring switching cannot either, since consumers who choose a new card themselves do so partly out of preference for the issuing bank. We therefore exploit a natural experiment. In October 2019, Walmart moved its consumer credit-card program, the store card and the co-brand Mastercard, from Synchrony to Capital One. Millions of cardholders thus acquired a Capital One credit-card relationship for reasons unrelated to their own bank preferences. Under the assumption that the transfer did not change these consumers' intrinsic preferences over banks, any differential post-2019 rise in Capital One's deposit share among Walmart cardholders reflects complementarity rather than sorting.

We track Walmart cardholders in the MRI surveys from 2016 to 2022, comparing them against all consumers without a Walmart card and, as a tighter store-card control, Target cardholders. Because MRI is a repeated cross-section, we follow the group rather than individual consumers, tracking how the probability of holding a Capital One deposit account conditional on holding a Walmart card changes over time. Figure 3 shows that the share of Walmart cardholders with a Capital One deposit account rises from about 11% before the transfer to roughly 21% by 2021–22, while rising far less among both comparison groups, as shown in panel (a). The event-study estimates in panel (b) show a flat pre-trend and a deposit effect of about 5 percentage points two to three years after the transfer. Since every Walmart card is Capital One-issued after the transfer,

all Walmart cardholders hold a Capital One credit card from 2020 onward. Panel (c) shows that around 30% of them already owned a Capital One card before 2019, so the transfer mechanically raised Capital One's card share among Walmart cardholders by 70 percentage points (62 percentage points relative to the comparison groups).

Appendix Table A.3 quantifies these effects in a difference-in-differences design. The average deposit effect, the complementarity estimate, is 0.048 in the full sample and 0.041 with Target controls, or 4–5 percentage points. The effect is specific to Capital One. Walmart cardholders do not pull away from the comparison groups at four placebo banks, and a triple difference that uses deposit holdings at the placebo banks as an additional control group yields similar estimates.[6] Against the first stage of 0.62, these estimates imply that Capital One converts 8 out of every 100 newly acquired credit-card holders into depositors.

The Walmart transfer thus establishes that complementarity is real and economically large. It identifies, however, only the combined effect of consideration complementarity and utility complementarity. The next two facts draw on our Prolific survey, which we designed to decompose this effect into its two components.

Fact 3. *Holding a product at a bank makes consumers significantly more likely to consider that bank when searching for the other product, on top of a strong correlation in consideration across products.*

Banks actively cross-sell to existing customers through targeted marketing, generating consideration complementarity that helps explain the cross-holding patterns documented above. To measure this channel, 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 holdings, we quantify the impact of current relationships on the search process.

Table 2 shows clear evidence of consideration complementarity. Holding one product at a bank raises the probability of considering that bank for the other product by 6 to 11 percentage points. Effects of this size matter for final choices because consideration sets are narrow. Consumers considered an average of only 2.3 banks for deposit accounts and 2.6 for credit cards (Appendix Figure A.3).

These estimates condition on a strong baseline correlation in consideration across product markets. Consumers who consider a bank for one product are approximately 25 percentage points more likely to consider it for the other, even among those who

[6]See Appendix Table A.4 for the triple difference, which absorbs person×year, bank×year, and bank×Walmart fixed effects and yields deposit estimates of 0.064 (full sample) and 0.050 (Target control).

chose neither product from that bank. This co-consideration likely reflects shared brand awareness. The same advertising or reputation that brings a bank to mind for deposits also brings it to mind for credit cards. Because the banks consumers consider overlap heavily with the banks already in their portfolios, conditioning on this correlation is what separates consideration complementarity from shared brand awareness.

Fact 4. *An unbundling shock that moves one product to another bank triggers follow rates comparable to median-to-large own-price shocks, indicating substantial utility complementarity between same-bank deposit accounts and credit cards.*

We measure utility complementarity by benchmarking consumers' willingness to follow a moved product against their responses to own-price shocks. The stated-preference module (Section II) first identifies a second-choice bank (Berry et al., 2004) for each product in two steps. We ask consumers to name the bank they would choose if they were to open an account *de novo*, without any cost of switching. For those whose ideal is their current bank, we then ask for an explicit second choice if that bank were no longer an option.[7] We then phrase the remaining stated-preference questions as an option between staying and moving to their second choice. We randomize the magnitude of own-price shocks (savings-rate cuts and reward cuts of 0.1, 0.5, or 1 percentage point, and APR increases of 0.5, 2, or 5 percentage points), and separately ask whether consumers would follow if their other product had just moved to the different bank for reasons unrelated to the focal product. Because each scenario prompts an active decision, we interpret them as how a consumer would react if they were awake and actively making a decision.

Own-price switching rises with the magnitude of the shock (Figure 4). The fraction who would switch the shocked product rises from 19% to 28% to 40% under small, median, and large savings-rate cuts; from 15% to 25% to 50% under reward cuts; and from 21% to 43% to 56% under APR increases. This dose response addresses a central critique of stated-preference surveys in the contingent-valuation literature, where willingness to pay often fails to scale with the size of the stimulus (Hausman, 2012).

The unbundling question reveals utility complementarity. If cross-holding reflected only correlated preferences or consideration complementarity, a consumer facing the same two options at the same prices would make the same choice whether or not the other product had moved. Yet about 30% of consumers would follow the moved product, as many as a median-to-large own-price shock moves (the dashed line in Figure 4). Some consumers therefore derive substantial utility from holding their financial products at a single institution.

[7] An ideal that differs from the current bank is no contradiction: switching costs and infrequent reconsideration can keep consumers at a bank they would not choose starting fresh.

Motivations for Cross-Holding Comparing the reasons cross-holders and other consumers give for their account choices shows that cross-holding is motivated by convenience and existing relationships rather than price discounts or forward-looking relationship building (Appendix Figure A.11). Across both survey waves, respondents reported why they chose the bank of their primary credit card and their primary checking or savings account from a common core of reasons covering access and convenience, interest rates, fees, and rewards, and brand trust. Cross-holders are 7 to 13 percentage points more likely than other consumers to cite access and convenience, such as mobile banking and branch or ATM proximity. Conversely, cross-holders are 7 to 15 percentage points less likely to cite interest rates, fees, or rewards when choosing their credit card, and no more likely to cite interest rates or fees when choosing their bank account (-8 to $+4$ percentage points). If multi-product price discounts drove cross-holding, cross-holders would cite interest rates and rewards more often, not less.

Each wave also included targeted items to compare static and dynamic relationship motives, and the levels among cross-holders show that static motives are much larger than dynamic ones. On static motives, 38% to 45% of cross-holders report choosing their credit card or bank account because they already had another account at the bank, and 33% to 65% cite an easy application or account setup. By contrast, only 15% to 17% of cross-holders report the dynamic motive of opening their bank account to apply for a credit card at that bank.

Fact 5. *Consumers switch banks infrequently, and a majority are not at their self-described ideal bank. Non-switching reflects genuine satisfaction for some consumers and inattention or switching costs for others.*

Consumers keep their banking products for a long time. Appendix Figure A.6 plots the distributions of tenure at the current bank from our Prolific survey.[8] The median tenure is 8 years for a primary deposit account and 6 years for a primary credit card, and 46% and 37% of consumers, respectively, have held the same product for at least a decade. Only 4.5% of consumers opened their primary deposit account and credit card simultaneously, so consumers typically acquire these products sequentially rather than as a bundle.

To understand why switching is so rare, we combine the *de novo* ideal-bank elicitation with the stated-preference responses from Fact 4, which together separate three explanations that observed tenure alone cannot distinguish. First, some consumers are

[8]We 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.

genuinely satisfied. Asked which bank they would choose *de novo* today, 42% name their current bank for their deposit account and 38% do so for their credit card. These shares far exceed any single bank's market share, suggesting persistent bank-specific preferences in addition to the transitory taste shocks emphasized in models of deposit competition (Egan et al., 2025). Second, some consumers stay because they are inattentive. The smallest price shocks would move 15–20% of consumers (Figure 4), an order of magnitude more than the 2–3% who switched either product in the past year. When the survey prompts an active choice, these consumers are ready to move; in the field, they remain "sleepy." Third, some consumers stay because of switching costs. Their *de novo* ideal already differs from their current bank, yet they state they would not move when a price shock makes it even worse.

IV Model

A dynamic model of consumer choices over deposits and credit cards separates correlated preferences from complementarity. On the demand side, consumers first open accounts and then receive opportunities to switch, one product at a time. On the supply side, banks set deposit rates, reward rates, and APRs in a Nash equilibrium to maximize profits across all product lines, internalizing that winning a customer in one market raises demand in the other. Separating correlated preferences from complementarity matters because only complementarity generates loss-leader pricing incentives and transmits shocks across product markets.

IV.A Demand

The model follows each consumer's holdings of two banking products in continuous time. Let i index consumers, j banks, $k \in \{1, 2\}$ product categories ($k = 1$: deposits; $k = 2$: credit cards), and t time. Let $\mathcal{J} = \{1, 2, \ldots, J\}$ denote the set of banks that offer both products. Each consumer belongs to one of two observed segments, $\sigma_i \in \{\mathrm{T}, \mathrm{R}\}$. Transactors repay their card balance in full each month, and revolvers carry a balance. Consumer i's state is the pair $x_i = (x_{i1}, x_{i2})$. For either product, $x_{ik} \in \{-1\} \cup \{0\} \cup \mathcal{J}$. The state -1 denotes a product never adopted, 0 an account at an unlisted provider, and $j \in \mathcal{J}$ an account at a listed bank. Consumers start at $(-1, -1)$, and neither coordinate returns to -1.

Consideration and Choices Motivated by Fact 5, we assume consumers are inattentive and do not continuously optimize their choices. If a consumer does not have product k, opportunities to adopt it arrive at rate ν_k. Otherwise, opportunities to switch arrive

at rate λ_k. Because of continuous time, the consumer can only switch one product at a time. When an opportunity for product k arrives, the other product's choice is fixed, so $x_{i,-k}$ enters as a state variable in both the consideration and utility functions.

When an opportunity for product k arrives at time t, the consumer forms a consideration set $C_{ikt} \subset \mathcal{J}$ (Honka et al., 2017). The product-specific outside option is always available, so the consumer chooses from $C_{ikt} \cup \{0\}$. Conditional on the state x_i, each bank enters the consideration set independently, so the probability of forming a given set is:

$$P_{ik}(C_{ikt}|x_i) = \prod_{j \in C_{ikt}} \phi_{ijk}(x_i) \cdot \prod_{j \in \mathcal{J} \setminus C_{ikt}} (1 - \phi_{ijk}(x_i)), \tag{1}$$

where $\phi_{ijk}(x_i)$ is the probability that bank j enters the consideration set for product k. Consumers who currently hold product k at a bank automatically include that current bank; for all other banks, consideration follows a logistic form:

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

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

$$z_{ijk}(x_i) = \eta^c_{ik} + v^c_{ijk} + \chi^c_k \mathbb{1}(x_{i,-k} = j). \tag{2}$$

The first term, η^c_{ik}, is consumer i's search propensity for product k, which shifts the consideration probability of all banks. The second term, v^c_{ijk}, is a persistent, consumer-specific visibility of bank j; its correlation across the two products generates the baseline co-consideration documented in Fact 3. The third term is consideration complementarity: holding the other product at bank j raises the probability of considering j for product k, the relationship boost also documented in Fact 3.

Given a consideration set, consumer i's utility from alternative $j \in C_{ikt} \cup \{0\}$ is

$$u_{ijkt}(x_i) = \delta_{ijk}(x_i) + \varepsilon_{ijkt}, \tag{3}$$

where ε_{ijkt} is an idiosyncratic match value drawn from a Type I Extreme Value distribution. For inside alternatives $j \in \mathcal{J}$, the deterministic utility is

$$\delta_{ijk}(x_i) = v^u_{ijk} + \alpha_k p_{jk} + \chi^u_{ik} \mathbb{1}\{x_{i,-k} = j\} + \kappa_{ik} \mathbb{1}\{x_{ik} = j\} - \kappa^0_{ik} \mathbb{1}\{x_{ik} = 0\},$$

and the outside option is normalized to $\delta_{i0k}(x_i) = 0$.

Persistent bank preferences v^u_{ijk} can be correlated across products, generating cross-holding through sorting. The price coefficient α_k measures sensitivity to consumer terms. For deposits, p_{j1} is the deposit rate. For cards, p_{j2} is the reward rate for transactors and the negative of the APR for revolvers, so higher p_{jk} is always preferred. Utility complementarity χ^u_{ik} raises the value of holding both products at the same bank.

Switching costs favor the incumbent provider. A listed incumbent receives the bonus κ_{ik}. For an unlisted incumbent, subtracting κ^0_{ik} from every listed alternative gives the outside option an equivalent advantage while preserving its zero normalization. The outside option aggregates a continuum of unlisted banks. The cost attaches to the consumer's incumbent within that continuum, and κ^0_{ik} denotes its representation in the aggregate choice problem. Neither term applies before adoption.

We allow consumer-level heterogeneity in the search propensities η^c_{ik}, the bank-preference intercepts v^c_{ijk}, v^u_{ijk}, utility complementarities χ^u_{ik}, and switching costs κ_{ik}. In what follows, θ_i collects consumer i's continuous latent parameters, drawn from a segment-specific distribution $\theta_i \mid \sigma_i \sim F_{\sigma_i}$. We defer the parametric form of F_σ, and which remaining parameters vary across consumers, to Section V.A.

The conditional probability that consumer i chooses $j \in C_{ik} \cup \{0\}$ given consideration set C_{ik} is:

$$s_{ijk}(x_i, C_{ik}) = \mathbb{1}\{j \in C_{ik} \cup \{0\}\} \frac{\exp\left(\delta_{ijk}(x_i)\right)}{\sum_{r \in C_{ik} \cup \{0\}} \exp\left(\delta_{irk}(x_i)\right)}. \tag{4}$$

The unconditional choice probability is the weighted average across all possible consideration sets:

$$m_{ijk}(x_i) = \sum_{S \in 2^{\mathcal{J}}} P_{ik}(S|x_i) s_{ijk}(x_i, S). \tag{5}$$

Continuous-Time Generator Matrix Both complementarity channels make a bank's success in one market depend on its presence in the other. We collect these dynamics in a continuous-time Markov process (Arcidiacono et al., 2016), which simplifies the high-dimensional transition space because consumers switch one product at a time.

The economically central object is the post-adoption generator Q_i on $(\{0\} \cup \mathcal{J})^2$, which governs holdings once both products have left -1. Write $Q_i = Q_{i1} + Q_{i2}$. A product-k move cannot change x_{-k}, and both products use reconsideration clocks λ_{ik}:

$$q^{x \to y}_{ik} = \begin{cases} \lambda_{ik}\, m_{iy_k k}(x) & x_{-k} = y_{-k},\ x_k \neq y_k, \\ -\sum_{y' \neq x} q^{x \to y'}_{ik} & x = y, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

The full generator G_i adds the states in which at least one product has not been adopted. Q_i is $(J+1)^2 \times (J+1)^2$; G_i is $(J+2)^2 \times (J+2)^2$, adding $2J+3$ pre-adoption states. Ordering the pre-adoption states before the post-adoption states gives

$$G_i = \begin{pmatrix} A_i & B_i \\ 0 & Q_i \end{pmatrix}. \tag{7}$$

The rates in A_i and B_i follow the same rule as Equation 6, with the first-opening clock replacing the reconsideration clock whenever the moving product is still at -1. For $x \neq y$,

$$g_{ik}^{x \to y} = \begin{cases} \nu_{ik}\, m_{iy_k k}(x) & x_k = -1,\ y_k \in \{0\} \cup \mathcal{J},\ y_{-k} = x_{-k}, \\ \lambda_{ik}\, m_{iy_k k}(x) & x_k, y_k \in \{0\} \cup \mathcal{J},\ x_k \neq y_k,\ y_{-k} = x_{-k}, \\ 0 & \text{otherwise}, \end{cases} \tag{8}$$

and each diagonal entry is minus the sum of its row. A_i collects the rates among pre-adoption states, B_i the rates from pre-adoption states into $(\{0\} \cup \mathcal{J})^2$, and the zero block records that neither product returns to -1. Neither switching cost applies at -1, because the consumer has no incumbent provider.

IV.B Supply

Banks' pricing incentives depend on both complementarity and inattention. Complementarity generates cross-product spillovers. A competitive price on one product attracts customers who may later adopt the other, so the marginal value of a new customer exceeds the flow profit from any single product. Inattention makes demand adjust only slowly to price changes, granting banks temporary market power over their "sleepy" existing customers.

Banks set prices to maximize the steady-state present discounted value (PDV) of joint profits across product lines, taking competitors' prices as given. Each bank sets a three-dimensional price vector $p_j = (p_j^{\text{dep}}, p_j^{\text{rew}}, p_j^{\text{apr}})$ consisting of a deposit rate, a reward rate, and an APR. Because transactors respond to the reward rate and revolvers to the APR, these three instruments map into the prices and balances that enter a consumer's problem as

$$p_{j1} = p_j^{\text{dep}}, \qquad p_{j2}(\sigma) = \begin{cases} p_j^{\text{rew}} & \sigma = \text{T} \\ -p_j^{\text{apr}} & \sigma = \text{R} \end{cases}, \qquad b_2(\sigma) = \begin{cases} b_2^{\text{spend}} & \sigma = \text{T} \\ b_2^{\text{balance}} & \sigma = \text{R}, \end{cases} \tag{9}$$

where $b_1(\sigma)$ is the checking balance of a consumer in segment σ. Banks post a common

deposit rate and offer one card with both a reward rate and an APR. Transactors respond to rewards and revolvers to the APR, so banks need not observe σ to set these prices.

Let $\bar{s}_i(\theta, \sigma)$ denote consumer i's initial distribution over post-adoption banking states, on the grid $\{0\} \cup \mathcal{J}$ for both products; in equilibrium it is the steady-state distribution under Q_i, defined below. Because pricing starts from this distribution, only Q_i enters the supply side, and the pre-adoption blocks (A_i, B_i) play no role. A consumer belongs to segment σ with probability π_σ, where $\pi_{\mathrm{T}} = 1 - \pi_{\mathrm{R}}$. Within segment, the consumer draws a type from F_σ. Bank j's total profit averages the discounted flow payoff over both segments and, within each segment, over types:

$$\Pi_j(p_j, p_{-j}) = \sum_\sigma \pi_\sigma \, \mathbb{E}_{\theta|\sigma} \left[\int_0^\infty e^{-rt} \, \bar{s}_i(\theta, \sigma) \, e^{Q_i(p_j, p_{-j}; \theta, \sigma) t} \, h_j(\theta, \sigma) \, \mathrm{d}t \right]. \tag{10}$$

For each segment and type, the integrand is the PDV of flow profits for a consumer who starts from $\bar{s}_i(\theta, \sigma)$ and whose holdings evolve according to $Q_i(p_j, p_{-j}; \theta, \sigma)$.

The flow profit vector is $h_j(\theta, \sigma) = \sum_{k=1}^{2} O_{jk} \, b_k(\sigma) \left(-p_{jk}(\sigma) - c_{jk}(\sigma) \right)$, where O_{jk} selects the states in which the consumer holds product k at bank j and $c_{jk}(\sigma)$ is bank j's marginal cost of serving product k to segment σ.[9]

Each consumer's inner integral admits a closed-form solution via the *resolvent* of their generator matrix. At discount rate r, the resolvent is[10]

$$\mathcal{R}(p_j, p_{-j}; \theta, \sigma) \equiv \left(rI - Q_i(p_j, p_{-j}; \theta, \sigma) \right)^{-1} = \int_0^\infty e^{-rt} \, e^{Q_i(p_j, p_{-j}; \theta, \sigma) t} \, \mathrm{d}t. \tag{11}$$

With this identity, the profit function simplifies to

$$\Pi_j(p_j, p_{-j}) = \sum_\sigma \pi_\sigma \, \mathbb{E}_{\theta|\sigma} \left[\bar{s}_i(\theta, \sigma) \, \mathcal{R}(p_j, p_{-j}; \theta, \sigma) \, h_j(\theta, \sigma) \right]. \tag{12}$$

The resolvent $\mathcal{R}(\cdot; \theta, \sigma)$ converts the instantaneous transition rates in Q_i into the discounted expected time the consumer spends in each banking state, so each consumer's profit contribution is their initial holdings, times discounted state occupancies, times flow margins. Equation 12 averages these lifetime values within each segment and then across segments, holding $\bar{s}_i(\theta, \sigma)$ fixed at its initial value.

[9] Because higher p_{jk} is preferred by the consumer, $-p_{jk}(\sigma) - c_{jk}(\sigma)$ is the bank's per-unit margin: for revolvers it equals $\mathrm{APR}_j - c_{j2}(\mathrm{R})$, with $c_{j2}(\mathrm{R})$ the cost of funds from overhead and default risk, while for deposits and for transactor cards $c_{jk}(\sigma)$ is net of the investment return on deposits and of interchange revenue, respectively. We defer the measurement of $b_k(\sigma)$ and π_σ to Section V.

[10] The resolvent is well-defined because each Q_i is a generator matrix with non-positive real eigenvalues and $r > 0$.

Equilibrium Concept We define a steady-state equilibrium as a price vector $p^\star$ such that (1) each bank's prices solve $\max_{p_j} \Pi_j(p_j, p^\star_{-j})$ and (2) each $\bar{s}_i(\theta, \sigma)$ is a left zero-eigenvector of the post-adoption generator Q_i at $p^\star$:

$$\bar{s}_i(\theta, \sigma)\, Q_i(p^\star; \theta, \sigma) = 0 \quad \text{for every } (\theta, \sigma). \tag{13}$$

Our supply-side model is "essentially static," following Brown et al. (2023): rather than resetting prices period by period as the state evolves, banks set long-run optimal prices, recognizing that demand responds slowly because consumers are inattentive. We thus study a steady state in which no bank wants to deviate from the equilibrium prices even though consumers are slow to respond.

The resolvent representation of the profit function yields the first-order conditions for pricing. The bank optimizes over its three instruments p_j^ℓ, $\ell \in \{\text{dep}, \text{rew}, \text{apr}\}$, each of which affects only $\mathcal{R}(\theta, \sigma)$ and $h_j(\theta, \sigma)$. Applying the matrix inverse identity $\partial M^{-1} = -M^{-1}(\partial M)M^{-1}$ gives:

$$\frac{\partial \mathcal{R}(\theta, \sigma)}{\partial p_j^\ell} = \mathcal{R}(\theta, \sigma)\, \frac{\partial Q_i(\theta, \sigma)}{\partial p_j^\ell}\, \mathcal{R}(\theta, \sigma). \tag{14}$$

The first-order condition for an interior optimum in instrument ℓ is then

$$\begin{aligned} 0 = &\sum_\sigma \pi_\sigma \mathbb{E}_{\theta|\sigma}\left[\bar{s}_i\, \mathcal{R}\, \frac{\partial Q_i}{\partial p_j^\ell}\, \mathcal{R}\, h_j\right] \\ &+ \sum_\sigma \pi_\sigma \mathbb{E}_{\theta|\sigma}\left[\bar{s}_i\, \mathcal{R}\, \frac{\partial h_j}{\partial p_j^\ell}\right], \end{aligned} \tag{15}$$

where the arguments (θ, σ) are suppressed. The first term is the demand effect and the second is the direct margin effect. Equation 9 makes the reward-rate condition depend only on transactors and the APR condition only on revolvers. The deposit condition averages over both segments, weighted by π_σ. A more generous offer means a higher deposit or reward rate, or a lower APR. Such an offer attracts customers but reduces margins on existing balances.[11] Because h_j aggregates profits from deposits and credit cards jointly, the bank internalizes that winning a credit-card customer today raises the probability of also earning deposit margins from that same consumer. On the cost side, the more generous price is paid on the discounted stream of balances $b_k(\sigma)$ held by the

[11] For a positive APR derivative in Equation 15, the demand effect is negative and the direct margin effect is positive.

bank's product-k customers. Because switching opportunities arrive stochastically, this inframarginal base is slow to respond to less favorable terms, strengthening the bank's incentive to harvest it.

Complementarity pushes equilibrium margins below the markups banks would charge on each product alone. Index product lines by $\ell \in \{\text{dep}, \text{rew}, \text{apr}\}$ and express every price in the consumer-preferred direction, so $p_{j,\text{apr}} = -p_j^{\text{apr}}$ while deposit and reward prices retain their posted signs. Let $H_{j\ell}$ denote the PDV of balances held on line ℓ at bank j,[12] and let $\zeta_{j,\ell m} = H_{j\ell}^{-1}\, \partial H_{jm}/\partial p_{j\ell}$ be the semi-elasticity of line-m balances with respect to the line-ℓ price. With margins $\mu_{j\ell} = -p_{j\ell} - c_{j\ell}$, the first-order conditions in Equation 15 give

$$\mu_{j\ell} \;=\; \underbrace{\frac{1}{\zeta_{j,\ell\ell}}}_{\text{own-product markup } \mu_{j\ell}^{\text{own}}} \;-\; \sum_{m \neq \ell} \frac{\zeta_{j,\ell m}}{\zeta_{j,\ell\ell}}\, \mu_{jm}. \tag{16}$$

Absent complementarity, $\zeta_{j,\ell m} = 0$ for $m \neq \ell$ and the margin reduces to the standard inverse-semi-elasticity markup $\mu_{j\ell} = 1/\zeta_{j,\ell\ell}$. With complementarity, a better price on line ℓ also raises balances on the bank's other lines, so $\zeta_{j,\ell m} > 0$, and the cross-selling adjustment $\mu_{j\ell}^{\text{cross}} = -\sum_{m \neq \ell} (\zeta_{j,\ell m}/\zeta_{j,\ell\ell})\mu_{jm}$ is negative whenever those lines are profitable. The bank then prices line ℓ below its own-product markup, which is the loss-leader logic we quantify in Section VI.

IV.C Discussion of Key Assumptions

Extensive and Intensive Margins We assume that post-adoption choices determine where consumers hold each product rather than whether they hold it, and that balances $b_k(\sigma)$ are fixed within each segment. In principle, access to credit-card liquidity could reduce a consumer's need for precautionary deposit balances. However, this margin does not affect bank pricing incentives in our setting because a bank's prices do not shift consumers between having a credit card and not having one; the outside option pools unlisted card issuers rather than non-adoption. A bank's card pricing only affects which bank issues the consumer's card, leaving the consumer's overall access to credit and precautionary deposit balances unchanged.

One Bank per Product We assume that each consumer holds each product at a single bank, interpreting x_{ik} as the consumer's primary bank for product k and $b_k(\sigma)$ as the balance on that primary account. Although many consumers hold accounts at multiple

[12] $H_{j\ell}$ applies the discounted state occupancies $\bar{s}_i\mathcal{R}$ of Equation 12 to the product-ℓ balance and averages over types (Appendix Equation E.2).

institutions, they concentrate most of their checking balances and card spending in a single primary account (Appendix B.3). Card borrowing is concentrated in the same way, since in the CFPB Consumer Credit Panel most borrowers carry more than half of their card balances on a single card (Nelson, 2025, Online Appendix A.8). Tracking the primary account therefore captures the bulk of consumers' balances and banks' revenues.

Interpreting Correlated Preferences We model correlated bank preferences (v^c_{ijk}, v^u_{ijk}) as exogenous consumer primitives, even though in practice some of this correlation may come from bank investments such as brand-wide advertising. This assumption does not affect our pricing, franchise-value, or policy counterfactuals because correlated preferences do not make winning one product raise demand for the other. Brand exposure raises a consumer's valuation of both products regardless of current holdings, so it generates neither loss-leader pricing incentives nor cross-market pass-through of cost shocks, both of which depend on the structural complementarities (χ^c_k, χ^u_{ik}).

Utility Complementarity versus Relationship Pricing We assume that banks post uniform prices p_{jk} that do not condition on a consumer's other holdings, interpreting χ^u_{ik} as the convenience of consolidating accounts. Institutional practice and the survey evidence in Section III support this specification. Aside from Bank of America and U.S. Bank, the major banks in our study post uniform card rewards and offer no advertised deposit-rate bonuses for holding a credit card, and cross-holders in our survey cite convenience and existing relationships rather than interest rates or rewards as the reason for their choices.

Consideration, Switching Costs, and Search We model consideration probabilities in Equation 1 as a reduced-form function of current holdings rather than deriving them from optimal sequential search. In an endogenous search model, consumers who consolidate products at a single bank would search less and therefore switch less often. Our model captures this retention effect directly through utility complementarity and switching costs: a consumer who holds a card at bank j is more likely to consider j for deposits ($\chi^c_1 > 0$) and to move deposits there ($\chi^u_{i1} > 0$), and once both products are at j, switching either product elsewhere forfeits χ^u_{ik} on top of the switching cost κ_{ik}.

V Estimation and Results

We estimate demand by maximum likelihood optimized via a Monte Carlo EM algorithm. The survey separately identifies correlated bank preferences, consideration

complementarity, and utility complementarity. The estimated model reproduces the deposit-share gain from the Walmart credit-card transfer studied in Section III, and its product-level profit margins are consistent with external evidence.

V.A Parametrization of Consumer Preferences

The consumer-level parameters from Section IV are the search propensities η_{ik}^c, the bank-preference intercepts v_{ijk}^c and v_{ijk}^u, the utility complementarities χ_{ik}^u, and the switching costs κ_{ik}. Price sensitivity α_k, consideration complementarity χ_k^c, and the clocks (ν_k, λ_k) are homogeneous within segment. All of these objects are estimated separately for transactors and revolvers.

The bank-preference intercepts are jointly normal. For each bank j, the consideration pair (v_{ij1}^c, v_{ij2}^c) and the utility pair (v_{ij1}^u, v_{ij2}^u) are bivariate normal with bank-specific means and with covariance matrices Σ_v^c and Σ_v^u shared across banks. The within-bank correlations $\rho^c \equiv \rho(v_{ij1}^c, v_{ij2}^c)$ and $\rho^u \equiv \rho(v_{ij1}^u, v_{ij2}^u)$ are the correlated-preferences mechanism. A consumer who considers a bank for one product tends to consider it for the other (ρ^c), and a consumer who values one of a bank's products tends to value the other (ρ^u).

The search propensities $(\eta_{i1}^c, \eta_{i2}^c)$ are bivariate normal with mean zero and their own variances and correlation, independent of the other coefficients. Their mean is zero because the bank means of v_{ijk}^c absorb any common level.

The remaining random coefficients are jointly normal but floored at zero. Utility complementarity and switching costs thus may covary within and across products. Formally, let

$$(\chi_{i1}^{u*}, \kappa_{i1}^*, \chi_{i2}^{u*}, \kappa_{i2}^*) \sim \mathcal{N}\big((\overline{\chi}_1^u, \overline{\kappa}_1, \overline{\chi}_2^u, \overline{\kappa}_2), \Sigma_\theta^{\chi^u,\kappa}\big)$$

be jointly normal with mean vector $\overline{\theta}^{\chi^u,\kappa}$ and covariance matrix $\Sigma_\theta^{\chi^u,\kappa}$. Utility complementarity and switching costs are then the truncated-at-zero versions of the latent draws, $\chi_{ik}^u = \max\{\chi_{ik}^{u*}, 0\}$ and $\kappa_{ik} = \max\{\kappa_{ik}^*, 0\}$. This terminology emphasizes the economically active coefficient: the underlying normal is retained for estimation, while its negative tail maps to a point mass at zero in the kernel.[13]

Heterogeneity is concentrated in search propensities, bank preferences, utility complementarity, and switching costs. Appendix C.8 describes how the homogeneous parameters are estimated within the MC-EM algorithm.

The outside option aggregates unlisted providers, and Section IV lets an unlisted

[13] An unrestricted normal would permit negative switching costs. A lognormal would rule out zero switching costs and zero utility complementarity. Applying a zero lower bound to a normal combines a point mass with a continuous positive component (Train, 2009).

incumbent carry its own switching cost κ_{ik}^0. We set $\kappa_{i1}^0 = 0$ for deposits and $\kappa_{i2}^0 = \kappa_{i2}$ for cards. A consumer with an existing card at an unlisted issuer therefore faces the same cost of moving to a listed bank as any other cardholder, while a consumer whose deposit account sits at an unlisted bank can move to a listed bank without that cost. We treat deposits and credit cards asymmetrically because the market share of unlisted banks is much larger than the market share of any listed bank in the deposit market, whereas it is similar to the market share of a typical listed bank in the credit card market. Thus applying the same switching cost convention in the deposit market would result in implausibly long tenures at unlisted banks for depositors.

V.B Intuition for Identification

The reduced-form facts from the Prolific survey help to separate the roles of correlated preferences and complementarity in driving observed cross-holding. They also allow us to separately identify inattention, switching costs, and persistent preferences in explaining why consumers rarely switch banks.

Correlated Consideration and Preferences The baseline co-consideration in Fact 3 identifies ρ^c directly. Consumers tend to include the same bank in both consideration sets even when they hold neither, reflecting common brand exposure that does not depend on any existing relationship. Given ρ^c and the two complementarity parameters identified next, the utility correlation ρ^u picks up the excess cross-holding in Fact 1 that complementarity and co-consideration alone cannot explain.

Search Propensity Variation in consideration-set size across consumers identifies the dispersion of η_{ik}^c, since bank-specific intercepts v_{ijk}^c are independent across banks. The covariance in consideration-set size across the two products identifies $\rho(\eta_{i1}^c, \eta_{i2}^c)$. Allowing for correlated search propensities ensures that consumers who shop widely in both markets do not inflate the estimated within-bank visibility correlation ρ^c.

Consideration Complementarity The relationship boost in Fact 3 identifies χ_k^c as the residual over baseline co-consideration. Holding a bank's product raises the probability that a consumer considers the bank for the other product, over and above the co-consideration that ρ^c absorbs. This boost reflects the existing relationship itself, targeted cross-selling and the salience of one's own bank, rather than shared brand awareness.

Utility Complementarity Responses to the unbundling shock in Fact 4 identify the mean utility complementarity $\overline{\chi}_k^u$. When a consumer's other product moves to a dif-

ferent bank for reasons unrelated to the focal product, following it reveals a taste for consolidation that cannot be attributed to correlated preferences or prior consideration (Gentzkow, 2007). The dispersion of χ^u_{ik} is identified by comparing the unbundling and own-price responses within respondent, both elicited in Fact 4. Section V.C writes those comparisons, together with the *de novo* ideal, as interval restrictions on a common Gumbel draw.

Wake-Up Rates, Switching Costs, and Persistent Preferences Fact 5 shows that consumers switch infrequently for three reasons: they may like their current bank, they may have high switching costs, or they may be inattentive. Existing models of dynamic deposit demand only allow for inattention (Egan et al., 2025; Lu and Wu, 2026). The *de novo* ideal-bank elicitation and the stated-preference price shocks in Facts 4 and 5 separate these mechanisms. Consumers who name their current bank as ideal *de novo* identify persistent preferences. Consumers whose ideal already differs from their current bank, yet who would not switch even under a small price shock, identify high switching costs. Given preference persistence and switching costs, observed tenures in Fact 5 identify the wake-up rates λ_k.

Other Key Parameters Price sensitivities α_k, common to all consumers within a segment, are identified by the dose response to own-price shocks of randomized magnitudes in Fact 4. Bank-specific mean utilities are identified by steady-state market shares.

V.C Likelihood and Estimator

We estimate the demand model by maximum likelihood, integrating over 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 of Section V.A and $L(X_i \mid \theta)$ is the complete-data likelihood at type θ.

The likelihood combines each respondent's account histories with their stated responses,

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

The retrospective block scores which banks were considered and chosen at each recorded opening, and when those openings occurred. The stated-preference block scores the joint pattern of the *de novo* ideal, the second-choice challenger, the hypothetical price shock, and the unbundling question. In the main text we focus on the case where the data

are uncensored and consumers remember their account histories perfectly. Appendix C integrates over censored times, incomplete recall, and unknown event order.

Retrospective record The data records, for each product, the most recent opening, with the bank chosen, the banks looked into, the holdings at the time, and how long ago it happened. Thus the likelihood is the product of the likelihood of the initial state before the older switch, the choice at that time, any intermediate transitions, the likelihood of the most recent switch, and the likelihood of not switching from then on. We assume that all consumers start at age 16 with neither a deposit account nor a credit card. We also condition on the fact that all consumers have a deposit account.

For example, suppose the consumer opened a deposit account t_1 years ago at the age of $E + 16$ at bank x_1^* and a credit card t_2 years ago at bank x_2^*. Before acquiring the deposit account, the consumer's previous deposit account was at bank x_1^- and their previous credit card was at bank x_2^-.[14] Taking the deposit opening to be the older event ($t_1 > t_2$), and writing $x(t_1^-) = (x_1^-, x_2^-)$, $x(t_1^+) = (x_1^*, x_2^-)$, and $x^* = (x_1^*, x_2^*)$, the retrospective contribution is

$$
\begin{aligned}
L^{\text{retro}}(\theta_i) = \frac{1}{1 - e^{-\nu_1 E}} & && \text{deposit ownership} \\
& \times \left[\exp\big(G_i(E - t_1)\big)\right]_{(-1,-1),\, x(t_1^-)} && \text{state before older switch} \\
& \times \lambda_1 \, P\big(C_1 \mid x(t_1^-)\big) \, s_{i x_1^* 1}\big(x(t_1^-), C_1\big) && \text{choice at that time} \\
& \times \left[\exp\big((t_1 - t_2)(\widetilde{Q}_{i1} + Q_{i2})\big)\right]_{x(t_1^+),\, x(t_2^-)} && \text{intermediate transitions} \\
& \times \lambda_2 \, P\big(C_2 \mid x(t_2^-)\big) \, s_{i x_2^* 2}\big(x(t_2^-), C_2\big) && \text{most recent switch} \\
& \times \left[\exp\big(t_2(\widetilde{Q}_{i1} + \widetilde{Q}_{i2})\big)\right]_{x^*,\, x^*} && \text{no switch thereafter}
\end{aligned}
\tag{17}
$$

where $\left[M\right]_{x,x'}$ is the (x, x') entry of the matrix M, and exp denotes the matrix exponential.

Stated responses The stated-preference block scores the joint pattern of answers from Facts 4 and 5. Each product module supplies a second-choice bank a_{ik}, a prospective consideration set C_{ik}, and answers to three hypothetical questions $h \in \{\text{ideal}, \text{price}, \text{unb}\}$. Let Y_{ik} be the set of questions to which the respondent answers yes, where yes means naming a_{ik} as the *de novo* ideal, switching to a_{ik} after an unfavorable price change, or switching to a_{ik} when the non-focal product moves there. We interpret these answers as the choices the respondent would make if a reconsideration opportunity arrived today,

[14] If the consumer did not have an account, we can set $x_1^- = -1$.

under a single draw of Gumbel shocks held fixed across questions. The contribution is therefore the mass of shock realizations that generate the whole pattern, rather than the product of independent choice probabilities.

Suppressing the product subscript, reduce the incumbent's systematic utility by d and write the challenger's resulting share

$$s_a(d; C, \theta_i) = \frac{e^{\delta_{ia}(x_i)}}{e^{\delta_{ix}(x_i)-d} + \sum_{j \in (C \cup \{0\}) \setminus \{x\}} e^{\delta_{ij}(x_i)}}, \tag{18}$$

where $\delta_{ij}(x_i)$ denotes $\delta_{ijk}(x_i)$ from Section IV, which depends on θ_i. Then $s_a(-\infty; C, \theta_i) = 0$, while $s_a(\infty; C, \theta_i)$ is the probability that a is the best non-current alternative.

Each question lowers the incumbent's utility by a known amount d_h. The ideal question removes the switching-cost advantage, so $d_{\text{ideal}} = \kappa$. The price question sets $d_{\text{price}} = \alpha \Delta p$, where Δp is the deterioration in consumer terms. The unbundling question sets d_{unb} to $2\chi^u$ when the other product moves from the same listed incumbent to a, since the incumbent loses the consolidation benefit and the challenger gains it, and to χ^u when it moves from a third provider. Because the Gumbel draw is common across questions, the respondent answers yes to question h exactly when d_h exceeds a single latent cutoff, the smallest reduction at which a beats the incumbent. The answers therefore bracket that cutoff, and $s_a(d; C, \theta_i)$ is the probability that it lies below d. The contribution is the probability of the bracket, times the probability of the elicited consideration set,

$$\begin{aligned} L_{ik}^{\text{hyp}}(\theta_i) = \Big[s_a\big(\min_{h \in Y_{ik}} d_h; C, \theta_i\big) - s_a\big(\max_{h \notin Y_{ik}} d_h; C, \theta_i\big) \Big]_+ & \quad \text{joint response cell} \\ \times P(C \mid x_i, \theta_i) & \quad \text{consideration today}, \end{aligned} \tag{19}$$

with $\min \emptyset = \infty$ and $\max \emptyset = -\infty$. For example, a respondent who would switch after the price change but not after the unbundling shock contributes only draws with $d_{\text{unb}} < \alpha \Delta p$.

The Prolific sample identifies consumer-level heterogeneity but is too small to pin down each bank's mean utility. In a post-EM step, we update the mean bank utilities $\overline{v}_{jk}^u$ to match the MRI-Simmons cross-section. Holding the remaining parameters at their EM estimates, we maximize a weighted steady-state likelihood on roughly 100,000 MRI-Simmons respondents, integrating over the EM draws of θ_i.

Standard Errors We compute standard errors using a Bayesian bootstrap with 40 replicates. Each replicate draws an independent exponential weight for every Prolific and MRI-Simmons respondent and re-runs the MC-EM estimator, the post-EM refinement,

marginal-cost recovery, and all counterfactual equilibria on the weighted sample.

Supply-Side Cost Recovery Given demand, we recover bank-product marginal costs by inverting Equation 15 at observed prices, taking $\bar{s}_i(\theta, \sigma)$ as the left zero-eigenvector of Q_i. Because each bank's flow payoff is affine in its own costs, the first-order conditions remain linear in costs under heterogeneity. Stacking the three conditions, pooled checking, transactor credit, and revolver credit, gives a 3×3 linear system per bank. Appendix E presents the derivation.

We set the discount rate to $r = 0.02$, matching long-run real Treasury yields. We set π_{R} to the average, across our four survey sources, of the share of credit-card holders who carry a balance, $\pi_{\mathrm{R}} = 41.5\%$ (Table 1). We calibrate $b_k(\sigma)$ as the simple average of median values from the SCF, NHCS, and MRI, separately by segment. We use medians rather than means because reported balances are subject to measurement error, and a few misreported large balances can move a mean substantially. For revolvers, we use deposits of $\$5,472$ and credit-card balances of $\$3,850$. For transactors, we use deposits of $\$25,552$ and annual spending of $\$15,063$ (Appendix B).

V.D Estimation Results

Our estimates suggest that both correlated preferences and complementarity contribute to observed cross-holding. They also show that inattention, switching costs, and persistent unobserved heterogeneity all help explain why consumers rarely switch banks. Table 3 reports population means by segment, and Table 4 reports the associated dispersions and correlations. We report utility complementarity and switching costs as equivalent changes in the product's own price. Each is the population mean of the active coefficient, $E[\max\{\chi^{u*}_{ik}, 0\}]$ or $E[\max\{\kappa^{*}_{ik}, 0\}]$, divided by the segment's price coefficient α_k, so that a value of one means one percentage point of price. The own price is the deposit rate for checking, the rewards rate for transactor cards, and the APR for revolver cards. The standard deviations in Table 4 are scaled the same way.

Correlated Consideration and Preferences Consideration intercepts are highly correlated within bank ($\rho^c = 0.97$ for transactors and 0.97 for revolvers). A consumer who looks into a bank for deposits is very likely to look into it for cards as well, even absent an existing relationship. Utility intercepts are less strongly correlated ($\rho^u = 0.50$ and 0.42). Thus even in the absence of complementarity, we should expect consumers to cross-hold. Search propensities also vary widely across consumers, with standard deviations of 1.0 for checking and 0.81 for cards among transactors, and 0.84 and 0.77 among

revolvers. Consumers who shop widely for one product also shop widely for the other ($\rho(\eta_1^c, \eta_2^c) = 0.91$ for transactors and 0.84 for revolvers).

Consideration Complementarity Consideration complementarity is modest for checking and larger for cards. We translate the estimates into percentage-point lifts in the probability that a bank enters the consideration set for one product, given an existing relationship in the other.[15] A credit card raises the probability that a consumer considers the same bank for checking by about 4 percentage points for transactors and 5 for revolvers, and both lifts are more than twice their standard errors. A checking account raises card consideration by about 12 percentage points for transactors and 8 for revolvers, in line with the reduced-form lifts in Section III.

In the data, many consumers report considering the same bank for both products even when they do not have a product at that bank. Our model attributes this pattern mostly to correlated search and preferences across products.

Utility Complementarity Utility complementarity varies across consumers, with a substantial share placing no value on consolidation. Between 34% and 45% of consumers have $\chi_{ik}^u = 0$, depending on product and segment. Averaged over all consumers, including those at zero, holding the other product at the same bank is worth 52 basis points of deposit rate for transactors and 65 for revolvers. On cards it is worth 40 basis points of rewards for transactors and 176 basis points of APR for revolvers. The dispersion is large relative to these means. Among transactors, for example, the value depositors place on holding their credit card at the same bank has a standard deviation of around 80 basis points of deposit rate. The value cardholders place on holding their deposit account at the same bank has a standard deviation of around 70 basis points of rewards. A consumer who values consolidation on one product tends to value it on the other ($\rho(\chi_1^u, \chi_2^u) = 0.63$ for transactors and 0.63 for revolvers). Thus banks can attract a substantial share of consumers not by offering better prices but by offering a wide array of products.

Wake-Up Rates, Switching Costs, and Persistent Preferences The estimates attribute infrequent switching to all three mechanisms. Checking reconsideration rates are $\lambda_1 = 0.28$ for transactors and 0.23 for revolvers, a horizon of about four years; transactors reconsider cards about every year and a half ($\lambda_2 = 0.70$) and revolvers about every two

[15] Because consideration is logistic in the intercepts, a common χ_k^c implies a bank-specific probability lift $\Lambda(\overline{v}_{jk}^c + \chi_k^c) - \Lambda(\overline{v}_{jk}^c)$. The reported figure is the unweighted average of these lifts across banks with a finite mean consideration intercept, evaluated at the estimated bank means.

years ($\lambda_2 = 0.47$). These wake-up rates are fast relative to observed tenures, so when consumers reconsider, they often do not switch. Average switching costs are 134 basis points of deposit rate for transactors and 213 for revolvers on checking, and 100 basis points of rewards for transactors and 509 basis points of APR for revolvers on cards, the larger figure reflecting the underwriting frictions revolvers face when applying elsewhere. Unlike utility complementarity, switching costs are almost always active: 11% of consumers have no checking switching cost and at most 3% have no card switching cost. Revolvers' switching costs are also more strongly correlated across products ($\rho(\kappa_1, \kappa_2) = 0.76$, against 0.15 for transactors). Persistent preferences account for the rest. Utility intercepts are dispersed across consumers ($\text{sd}(v^u_{ijk}) = 0.30$–$0.54$), so many consumers are at a bank they genuinely prefer. Switching costs and persistent preferences thus explain why consumers so often stay put when they have the chance to move.

Price Sensitivity Deposit-rate sensitivity is 1.16 for transactors and 0.78 for revolvers. Conditional on consideration, those values sit at or just below the 0.8–2.3 range in the deposit-demand literature.[16] Card sensitivities are 2.57 for rewards (transactors) and 0.35 for APRs (revolvers). The revolver coefficient matches the structural price coefficients in Nelson (2025), which range from about 0.19 to 0.53 across the FICO distribution: like ours, that model embeds APR sensitivity in a dynamic discrete-choice problem where switching costs make card relationships persistent. The rewards coefficient is lower than the network-level reward semi-elasticities in Wang (2026), where consumers substitute among Visa, Mastercard, and AmEx rather than among banks' differentiated card products.

V.E Goodness of Fit

Fit to Targeted Moments Market shares, the stated switching responses, and reported tenures each enter the likelihood, directly or through the post-EM refinement, so the fit reported here checks that the estimator recovers them rather than testing the model out of sample. The estimated model reproduces bank-level same-bank holding even though the complementarity parameters are common across banks (Figure 5).[17] It also matches the stated switching responses at each shock, including the higher propensity of existing cross-holders to follow the other product (Appendix Figure A.8). The model matches the cross-sectional distribution of tenure, including its heavy right tail (Figure 6). It also

[16] 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).

[17] Market shares are matched largely by construction, because the post-EM refinement fits bank-level mean utilities to the MRI cross-section. Appendix Figure A.7 plots the share fit.

matches how average deposit-account and credit-card tenure vary with age in Prolific (Appendix Figure A.9).

External Validation The model reproduces the untargeted deposit gain from the Walmart portfolio transfer in Section III. With the Capital One card standing in for the Walmart card, which previously came from Synchrony and had no checking account, each consumer type starts at its steady state without Capital One complementarity (Appendix D). At the transfer date, we restore complementarity and follow the treated group of Capital One cardholders and the comparison group of all other consumers for three years, reweighting types to match pre-transfer Capital One deposit shares. Three years after the transfer, the difference-in-differences effect on Capital One deposit holding is 4.9 percentage points in the model, compared with 5.4 percentage points (SE 1.1) in the 2022 MRI-Simmons wave.

Recovered margins are consistent with industry estimates, providing a supply-side check on the demand estimates. The median checking margin of 1.2 percentage points aligns with the deposit spreads in Drechsler et al. (2017) and the markups in Wang et al. (2022). The median transactor margin is 0.3 percentage points, consistent with the finding in Drechsler et al. (2025) that interchange revenue is largely offset by rewards. The median revolver margin is 3.1 percentage points, driven by interest on outstanding balances, consistent with Nelson (2025) and Drechsler et al. (2025).

VI Counterfactuals

Treating deposit and credit-card demand as independent misses how cross-selling shapes competition and policy. In the policy exercises, we compare the full model with a restricted specification that sets $\chi^c = \chi^u = 0$ and removes every correlation between a consumer's checking parameters and card parameters, including bank intercepts, search propensities, and switching costs. In this restricted model, checking and credit-card choices are independent.

VI.A Decomposing Observed Cross-Holding

Observed cross-holding in the full model far exceeds the rate under independent product choice (Table 5). The outcome is the unconditional steady-state probability of holding checking and a named credit card at the same bank. The full model predicts a pooled rate of 24.3% (SE 0.8). Without complementarity and without correlated preferences, checking and card choices are independent and the rate falls to 5.4% (SE 0.2). Correlated preferences capture every covariance between a checking parameter and a

card parameter, including bank intercepts, search propensities, and switching costs.[18] Because the channels interact, we split the 18.9 percentage-point excess with Shapley values, which average each channel's contribution over every order of removal.

Complementarity explains about three-quarters of excess cross-holding, mostly through utility. Utility complementarity accounts for 62% (SE 4) and consideration complementarity for 13% (SE 2). Correlated preferences account for the remaining 24% (SE 3). The split is similar for transactors and revolvers, though correlated preferences account for a somewhat larger share among revolvers (28% versus 22%). Consumers who like a bank's branches or reputation may choose it for both products even if holding one does not affect their choice of the other. Attributing all excess cross-holding to cross-selling therefore overstates what a bank gains by winning a customer's first product. Strong utility complementarity and limited consideration spillovers also suggest that cross-selling is not merely steering, because consumers value consolidating their financial products at one institution.

Correlated preferences also explain nearly half of the long-run relationship effect documented by Basten and Juelsrud (2023), who find that a consumer's early-life bank predicts later borrowing.[19] In the model counterpart of their main specification, a bank's depositors are more likely than depositors at other listed banks to open a card at the bank within 14 years (Appendix D). The gap between the treated group and the comparison group is 15.3 percentage points in the full model. Removing correlated preferences while holding all other parameters fixed lowers this gap to 8.5 percentage points. Attributing the full relationship effect to complementarity would therefore overstate the role of complementarity by a factor of about two.

VI.B Loss-Leader Pricing and Bank Franchise Value

Complementarity gives banks an incentive to offer generous card rewards to attract depositors. U.S. banks offer rewards that absorb much of their credit-card interchange income (Drechsler et al., 2025). Existing work explains the resulting thin margins within the card market: consumers respond strongly to rewards (Wang, 2026), and naive cardholders end up paying interest that funds the rewards earned by others (Agarwal et al., 2026). Our estimates point to a third source, outside the card market: rewards cards are loss leaders for deposit accounts. We first show that rewards cards contribute far more

[18] Removing correlated preferences leaves the joint distribution of each product's own parameters unchanged. The restricted model in the policy exercises removes these correlations and also sets $\chi^c = \chi^u = 0$.

[19] Their evidence concerns deposits and household loans rather than credit cards, so our exercise does not decompose their estimate directly, though we expect similar economic forces to operate in both settings.

to bank franchise value than to accounting profits, and then that banks price them below what a standalone issuer would charge.

Figure 7 compares two measures of what each product is worth to a bank. Panel A reports accounting profits, which attribute to each product the margins booked on it. For deposit-oriented banks, such as Chase, Bank of America, and Wells Fargo, deposit accounts generate the majority of profits. Capital One and Citi instead derive just over half of their profits from credit cards, reflecting their credit-heavy business models. Transactor credit cards usually contribute a small share because of their thin standalone margins.

Panel B measures each product's marginal contribution by the decline in total franchise value when the bank stops offering it. Product removal destroys cross-selling profits as well as the profits booked to that product, so for every bank these contributions sum to more than 100%. Their profit-weighted average is 122% of accounting profit across the named banks, ranging from 114% at the super-regional bank to 130% at Chase. The excess over 100% is the cross-selling value that accounting profits miss. Because these removal exercises overlap, their contributions are not additive shares of franchise value.

The gap between the two measures is largest for transactor credit cards. Across the named banks combined, they account for about 7% of accounting profits but contribute about 22% of franchise value, a gap of 14.4 percentage points (SE 1.8). At the deposit-oriented banks, removing rewards cards would cost roughly three to nine times their accounting contribution. Transactors' larger deposit balances explain why they are worth attracting through card rewards. Median total deposits are \$23.7k for transactors versus \$4.1k for revolvers (Table 1, Panels C1–C2). Winning a transactor's deposit account can therefore make generous card rewards profitable for the bank as a whole. For revolvers, the smaller deposit balance weakens this incentive, while banks earn interest on the \$3.9k revolving balance.

Cross-selling lowers rewards-card margins most at banks with deposit franchises. Figure 8 plots the margin decomposition of Equation 16, splitting each bank-product margin into the own-product markup and the cross-selling adjustment. Across essentially all bank-product pairs, the negative cross-selling adjustment lowers equilibrium margins below their own-product markups. The adjustment is concentrated almost entirely on transactor credit cards. Deposit pricing sits close to what a standalone deposit bank would choose, with only a small cross-selling wedge. The reward-card wedge, in contrast, is substantial at every bank with a deposit franchise, and largest at the deposit-oriented banks, where the cross-selling adjustment erases roughly half of the

own-product markup. American Express and Discover, which have no deposit franchise to cross-sell into, price with essentially no wedge. Rewards cards are thus acquisition and retention tools for the higher-margin deposit franchise, and their thin margins reflect the value of winning depositors as well as competition within the card market.

VI.C Regulatory Spillovers

Credit-card interchange regulation spills over into deposit competition because consumers value holding both products at the same bank. The Credit Card Competition Act, reintroduced in 2026, would reduce interchange fees by requiring large banks to enable at least two unaffiliated networks for transaction routing. Mukharlyamov and Sarin (2025) show that reductions in debit interchange fees caused banks to raise checking-account fees. Since many consumers also hold their credit card at their deposit bank, one might expect banks to recover lost credit-card interchange revenue through worse deposit terms. Our model predicts the opposite: banks raise deposit rates to retain customers whose card relationships become less attractive.

We simulate a one-percentage-point increase in the cost of transactor cards, the kind of shock imposed by a reduction in interchange fees, and allow banks to reoptimize prices.[20] Figure 9 reports the price responses. Without complementarity, the adjustment remains within the transactor card market. Banks reduce rewards by roughly 86–102 basis points, while deposit rates and revolver APRs remain unchanged. Correlated preferences would not transmit the shock either. Customers may like both products at the same bank, but changing one product's terms does not change their demand for the other.

With complementarity, lower rewards weaken the deposit relationship as well as demand for the card. Banks still cut rewards substantially, by about 81–93 basis points, but they also compete more directly for depositors. Relative to the restricted model, deposit rates rise by 7.0 basis points at Chase (SE 3.0), 3.9 at Citi (SE 2.5), and 3.7 at Capital One (SE 2.0). The mechanism depends on consumers being able to choose the products separately: a less attractive card reduces the value of keeping deposits at the same bank, and because consumers can move those deposits elsewhere, raising the deposit rate helps the bank retain them. The resulting improvement in deposit terms also makes the bank more attractive to revolvers, allowing banks to raise APRs slightly, by up to 3 basis points.

[20] Whereas the Credit Card Competition Act exempts smaller issuers and could induce some consumers to substitute toward debit cards, we raise transactor card costs at all listed banks (including American Express and Discover) and hold the value of the outside option fixed. As discussed in Section IV.C, the results therefore capture reallocation across banks rather than substitution out of credit cards.

Card-centric banks lose the most deposit share after the interchange shock. Citi's deposit share falls by about 9% and Capital One's by 6%, compared with a 2.3% decline at Chase and changes of less than 1% in either direction at Bank of America and Wells Fargo.[21] The larger losses at card-centric banks reflect the acquisition channel: their deposit franchises depend more heavily on the appeal of their credit cards, while the megabanks' deposit franchises have broad standalone appeal. A fee regulation aimed narrowly at the card market thus propagates, through complementarity, into the deposit market, with consequences a single-market analysis would miss.

Banks bear part of the regulatory burden in their deposit business rather than in the card market alone. At fixed prices, the interchange shock falls entirely on transactor cards and reduces industry profit by 35.1%. After banks cut rewards and reoptimize prices, industry profit falls by 13.0% (SE 2.2). Lost deposit share and higher deposit rates generate 4.2 percentage points of that decline in checking accounts (SE 1.3). Checking carries about a third of the loss, though this share is imprecisely estimated (SE 15 percentage points). Among the six banks that offer checking, deposit profits fall by 4.5 percentage points out of an 11.0 percentage-point total profit decline, or roughly two-fifths of the loss (SE 19 percentage points). Without complementarity, total industry profit falls by a similar 12.9% (SE 2.4), but none of the loss falls on the deposit franchise.

VI.D A Cost Shock to Revolving Credit

A cost shock to the borrowing side of the credit-card market moves deposit rates in the opposite direction. We represent regulation that makes lending to revolvers more costly, such as tighter underwriting rules or limits on interest charges, as a 10 percentage-point rise in the marginal cost of revolver credit at every bank. Banks pass the cost increase through almost one-for-one and raise APRs by 9.9 to 10.4 percentage points. Higher APRs reduce revolver card shares by 3.6 to 14.5 percentage points and lower industry profit by 15.9% (SE 2.7).

Deposit rates fall by 0.9 to 1.1 basis points at every bank that offers checking (SEs 0.3 to 1.3), whereas they remain unchanged without complementarity. Higher APRs shrink revolver card demand, so winning a revolver's deposit account yields a smaller expected cross-selling payoff in the card market and weakens banks' incentive to compete for deposits. The deposit-rate decline is small because revolvers hold small deposit balances and carry a small cross-selling wedge in baseline card pricing (Figure 8). Even so, the reversal in sign relative to the interchange shock shows that cross-market spillovers depend on which product the shock hits.

[21] These are percentage changes in market shares, not percentage-point changes.

VI.E Horizontal Merger

We study a merger between Citi and the model's super-regional bank to assess whether complementarity can support an efficiency defense. The combination pairs a national card franchise with a regional deposit base, allowing consumers to hold a Citi card and regional deposits within the same banking group. Both brands remain active and retain separate prices, while the combined firm chooses prices jointly and rivals reoptimize in the new Nash equilibrium. Reported combined price changes use baseline market shares as weights. Figure 10 compares the resulting equilibrium with the no-complementarity benchmark.

Absent complementarity, the merger produces the conventional price effects of horizontal consolidation. The combined bank cuts its average deposit rate by 4 basis points and raises its average APR by 11 basis points. Its deposit and credit-card shares fall by 4% and 4%, respectively.

With complementarity, the merger creates a synergy that the merged bank largely captures through larger price changes. At unchanged prices, the expanded opportunities to cross-hold would raise the merged bank's deposit share by 10.3% (SE 1.0%) and its credit-card share by 14.8% (SE 1.6%). Customers who can combine a Citi card with regional deposits are harder to lose, and the merged bank exploits this attachment by cutting its average deposit rate by 8 basis points (SE 4.1) and raising its average APR by 17 basis points (SE 5.9), while leaving credit card rewards essentially unchanged. After prices adjust, the merged bank's deposit share and revolver share are unchanged, and only its transactor share rises, by an imprecisely estimated 23%. This counterfactual shows that complementarities may not be passed through to consumers, and that the effect can differ by market.

VII Conclusion

Many consumers choose banks rather than individual financial products, and that choice reshapes competition and regulation in retail banking. We combine quasi-experimental evidence from the Walmart portfolio transfer with an original survey and an equilibrium model of joint deposit and credit-card choice to separate the mechanisms behind cross-holding. Complementarity accounts for about three-quarters of the excess over independent choice, mostly through utility rather than consideration. Correlated preferences account for the remaining quarter. Complementarity, together with infrequent reconsideration and large switching costs, makes customer relationships durable and valuable. Rewards credit cards are therefore loss leaders, which banks use to acquire deposit customers rather than as standalone profit centers.

The policy lesson is that complementarity ties retail banking's product markets into a single competitive system. Accounting profitability understates what rewards cards are worth to a bank; regulation aimed narrowly at the card market spills into deposit competition, in a direction that single-market evidence would not predict; and mergers create synergies that the merged bank captures rather than passes through, so prices can worsen even as deposit concentration measures look benign. Competition analysis in retail banking should treat the banking relationship, not the individual product, as the unit of analysis.

Two extensions are natural. First, we document pervasive cross-holding across a wide range of product lines — mortgages, auto loans, and brokerage accounts among them — while our model deliberately focuses on two products to isolate the core economic forces. The framework extends tractably to additional products and could capture the full scope of the financial supermarket. Second, our counterfactuals trace prices, market shares, and profits, leaving welfare unquantified. The model contains the ingredients for welfare analysis — it separates complementarity from switching costs and inattention — but welfare accounting requires a stance on how to value choices made under these frictions, and we leave that to future work.

References

Agarwal, Sumit, Andrea F. Presbitero, André F. Silva, and Carlo Wix, "Who Pays For Your Rewards? Redistribution in the Credit Card Market," CEPR Discussion Paper 17733, Centre for Economic Policy Research May 2026. New draft.

_ , 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.

Allen, Jason, Robert Clark, and Jean-François Houde, "The effect of mergers in search markets: Evidence from the Canadian mortgage industry," *American Economic Review*, 2014, *104* (10), 3365–3396.

_ , _ , and _ , "Search frictions and market power in negotiated-price markets," *Journal of Political Economy*, 2019, *127* (4), 1550–1598.

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 Haegele, 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.

d'Avernas, Adrien, Andrea L Eisfeldt, Can Huang, Richard Stanton, and Nancy Wallace, "The deposit business at large vs. small banks," Technical Report, National Bureau of Economic Research 2023.

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

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.

Economides, Nicholas and Steven C Salop, "Competition and integration among complements, and network market structure," *The Journal of Industrial Economics*, 1992, pp. 105–123.

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.

Greene, Claire and Joanna Stavins, "2020 Diary of Consumer Payment Choice," Technical Report 21-2 May 2021.

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.

Hastings, Justine S. and Lydia Tejeda-Ashton, "Financial Literacy, Information, and Demand Elasticity: Survey and Experimental Evidence from Mexico," Technical Report w14538, National Bureau of Economic Research 2008.

Hausman, Jerry, "Contingent Valuation: From Dubious to Hopeless," *Journal of Economic Perspectives*, 2012, *26* (4), 43–56.

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 Buyout Firms and Relationship Banking," *The Review of Financial Studies*, 2011, *24* (7), 2462–2498.

JPM, "JPM 2015 Investor Day.Pdf," February 2015.

Koont, Naz, "The digital banking revolution: Effects on competition and stability," *Available at SSRN 4624751*, 2023.

Laux, Christian and Uwe Walz, "Cross-Selling Lending and Underwriting: Scope Economies and Incentives," *Review of Finance*, 2009, *13* (2), 341–367.

Lu, Xu and Lingxuan Wu, "Banking on Inattention," January 2026.

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.

Mukharlyamov, Vladimir and Natasha Sarin, "Price Regulation in Two-Sided Markets: Empirical Evidence from Debit Cards," *Journal of Financial Economics*, April 2025, *Forthcoming*.

Nelson, Scott T., "Private Information and Price Regulation in the US Credit Card Market," *Econometrica*, 2025, *93* (4), 1371–1410.

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.

Song, Minjae, Sean Nicholson, and Claudio Lucarelli, "Mergers with interfirm bundling: a case of pharmaceutical cocktails," *The RAND journal of Economics*, 2017, *48* (3), 810–834.

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.

_ , *Discrete Choice Methods with Simulation*, 2nd ed., Cambridge: Cambridge University Press, 2009.

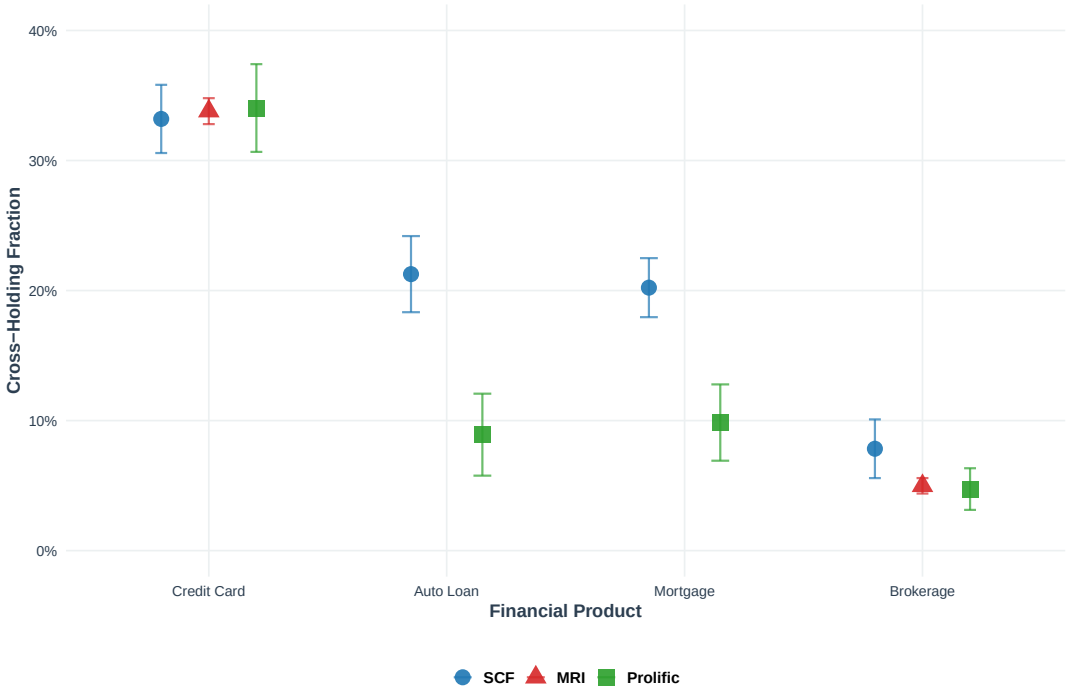
U.S. Department of Justice, "Bank Merger Competitive Review – Introduction and Overview," 1995.

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*, 2026.

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.

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



Notes: For each product, the bar shows the fraction of consumers who hold that product at the same bank as their checking account, among consumers who hold both a checking account and the product. To harmonize the measure across datasets, the sample is restricted to consumers with exactly one checking account and exactly one of the second product (e.g., one credit card), so that cross-holding status is unambiguous; the datasets are inconsistent in reporting primary accounts for consumers who hold multiple accounts. This restriction applies only to these aggregate rates; the bank-level comparisons and independence benchmarks in the text include consumers with multiple accounts. The 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.

Figure 2: Cross-holding of Checking Accounts and Credit Cards by Bank



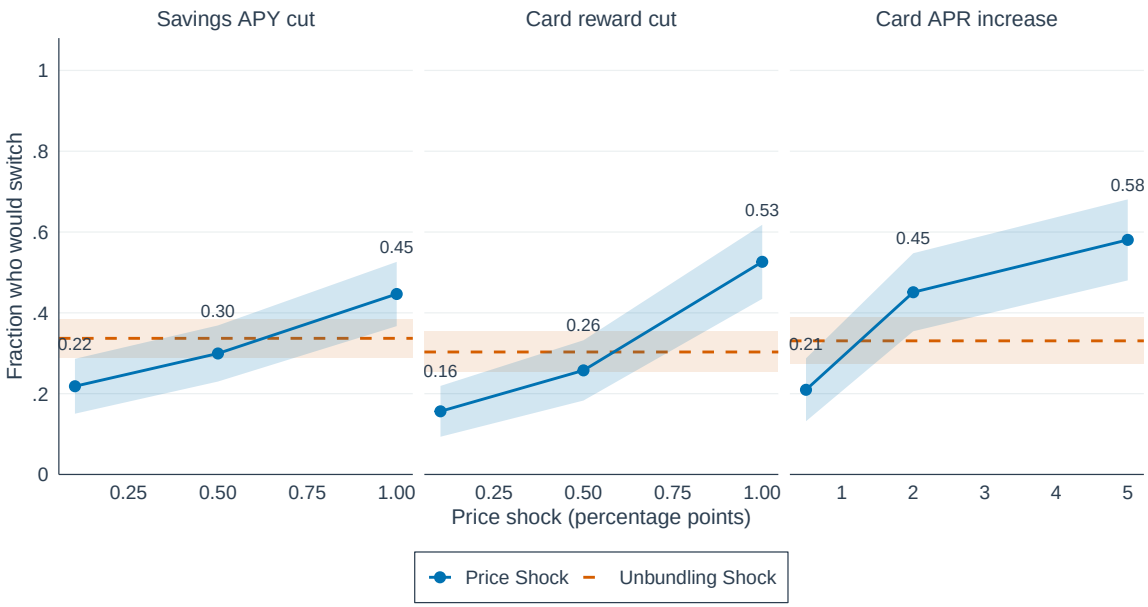
Notes: This figure shows the prevalence of cross-holding between checking accounts and credit cards. The top panel shows each bank's checking account market share among two groups of consumers: those who hold a credit card at the bank (blue dots) and those who do not (orange dots). The bottom panel shows 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 3: Effects of the Walmart portfolio transfer on Capital One market shares



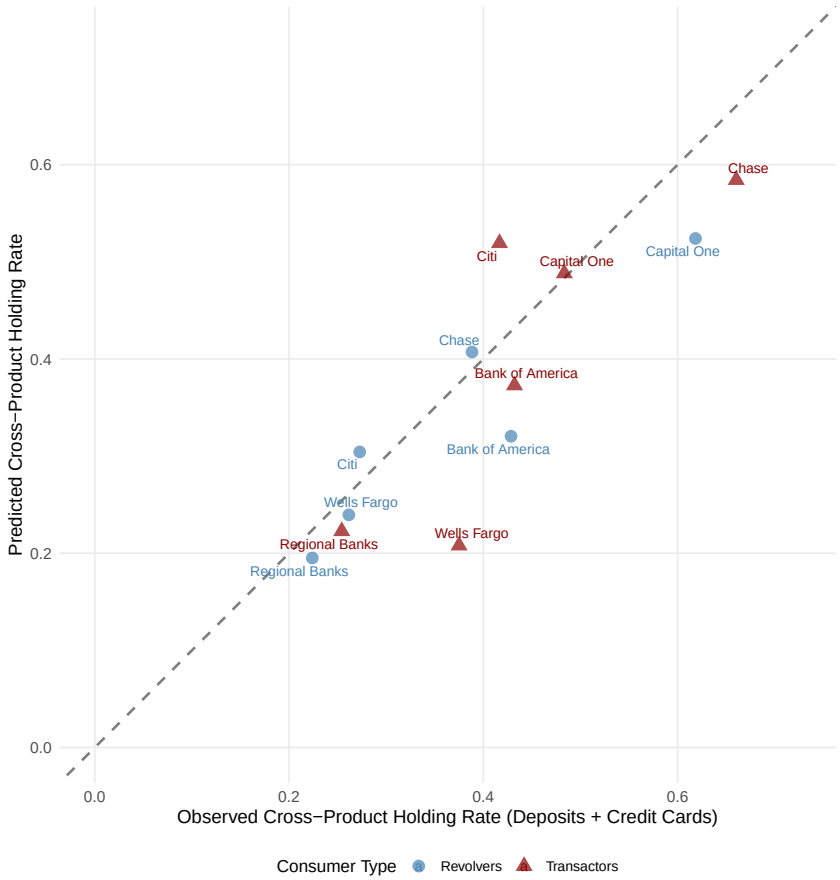
Notes: Left column: raw Capital One market shares by group, 2016–2022; the dashed vertical line marks the 2019 transition. Right column: event-study coefficients (Walmart×year interactions relative to 2018) with 95% confidence intervals, full sample. The credit-card outcome is issuer-coded: every Walmart cardholder is coded as holding a Capital One credit card from 2020 onward, since all Walmart cards are Capital One-issued after the transfer; self-reported Capital One card ownership is not used for these cells.

Figure 4: Stated Responses to Hypothetical Shocks



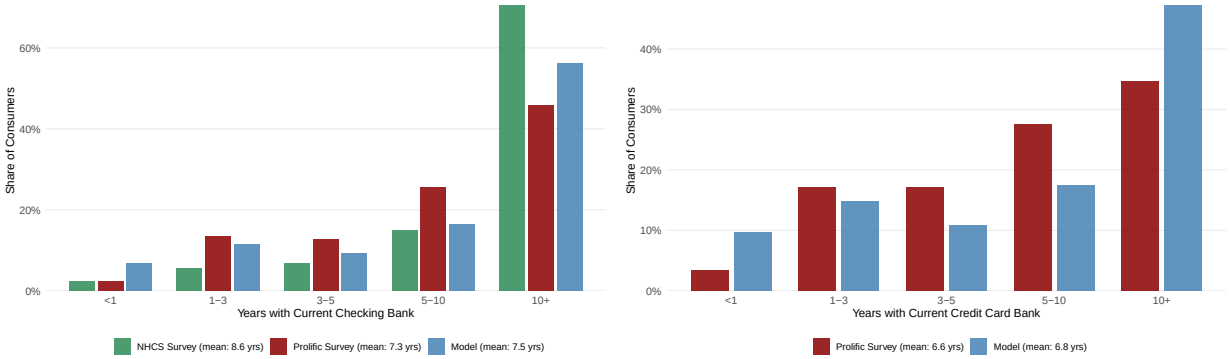
Notes: This figure summarizes stated responses to hypothetical shocks. Each facet plots the fraction who state they would switch their current product against the magnitude of a randomized own-price shock (solid line), identifying price sensitivity. The dashed *Unbundling Shock* line shows the fraction who state they would move the focal product if the other product had just switched banks for unrelated reasons, identifying utility complementarity from cross-holding. Shaded bands give 95% confidence intervals.

Figure 5: Goodness of Fit: Cross-holding Rates by Bank and Consumer Type



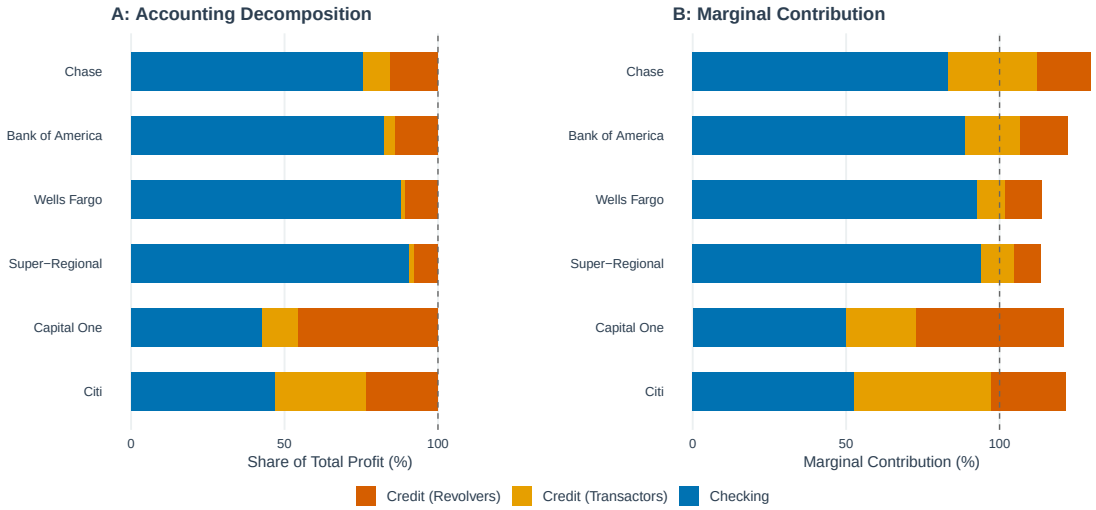
Notes: This figure compares observed and model-predicted cross-holding rates across banks for transactors and revolvers. Observed rates come from the Prolific survey. The figure shows the probability of holding a credit card at a bank conditional on having a checking account there.

Figure 6: Goodness of Fit: Tenure Distributions



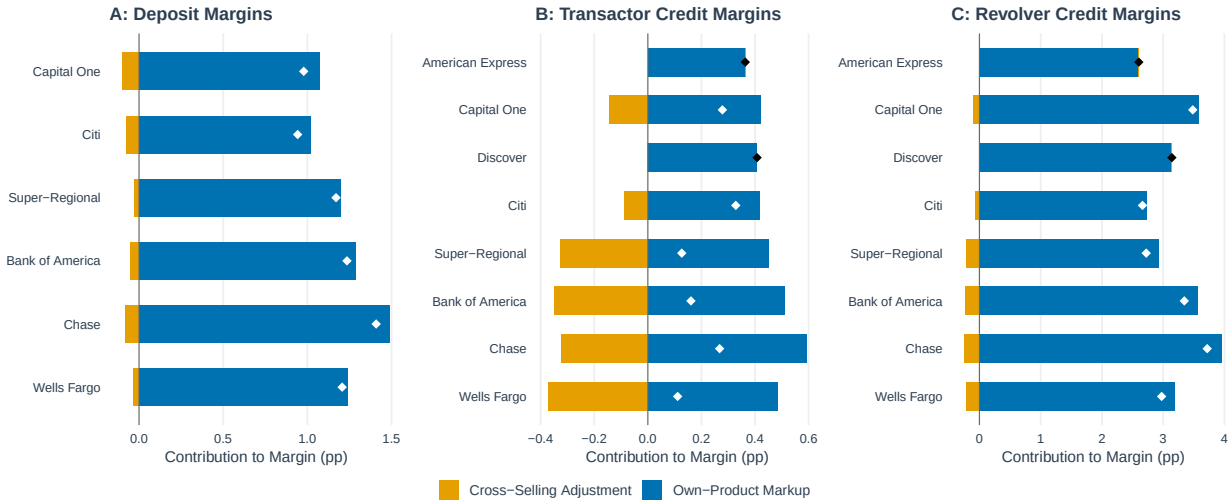
Notes: These figures compare model-predicted and observed tenure distributions for checking accounts (left) and credit cards (right). Tenures are binned into five categories: less than 1 year, 1–3 years, 3–5 years, 5–10 years, and 10+ years. The model predictions use the cross-sectional tenure distribution under Q_i. Observed tenures come from the Prolific survey, with NHCS survey data shown as an additional benchmark for checking accounts.

Figure 7: 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 8: FOC Margin Decomposition: Own-Product Markup vs. Cross-Selling Adjustment



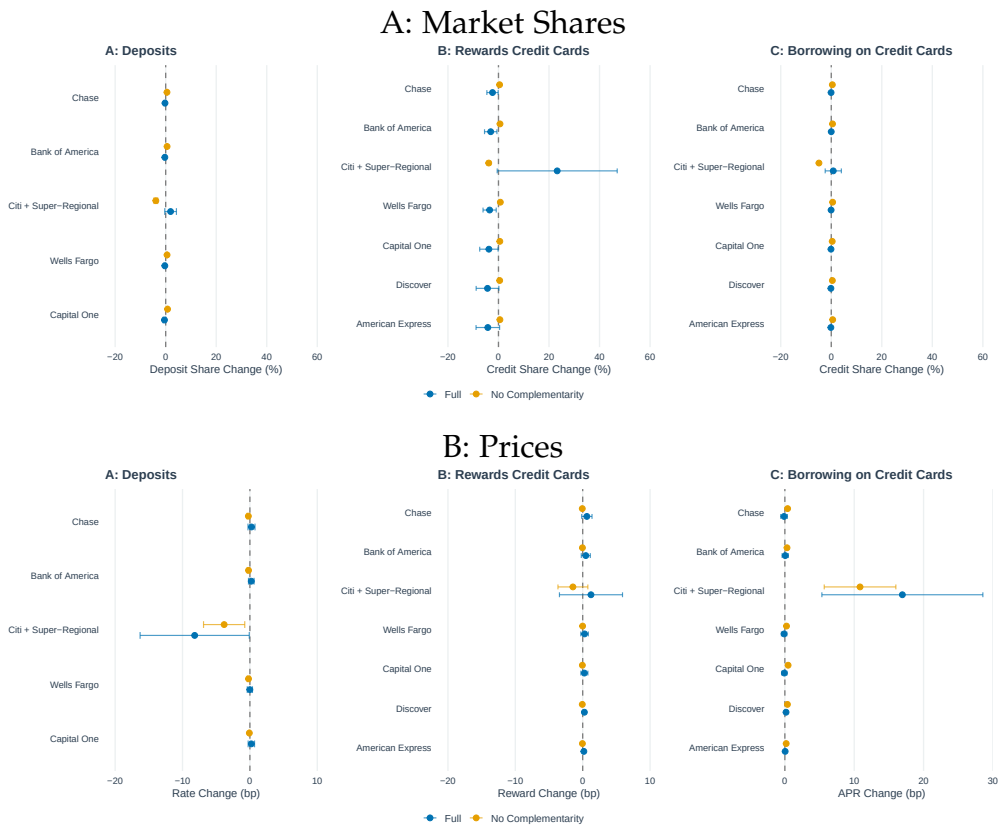
Notes: This figure decomposes each bank's equilibrium arithmetic margin (percentage points) using the multiproduct inverse-semi-elasticity rule in Equation 16. The own-product markup $\mu_{j\ell}^{\text{own}} = 1/\zeta_{j,\ell\ell}$ is the inverse of the dynamic own-price semi-elasticity of discounted balances; the cross-selling adjustment $\mu_{j\ell}^{\text{cross}} = -\sum_{m\neq\ell}(\zeta_{j,\ell m}/\zeta_{j,\ell\ell})\,\mu_{jm}$ uses the dynamic cross-price semi-elasticities (Equations E.2–E.5). Panels A–C show deposits, transactor credit, and revolver credit, respectively. Diamonds mark the total margin $\mu_{j\ell}$. A negative cross-selling adjustment lowers the margin on line ℓ because attracting customers raises profits on complementary lines.

Figure 9: 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 10: 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. The combined firm owns all four products (checking and credit at each brand) and solves a joint first-order condition over them. Cross-product complementarity χ pools across the firm, so a consumer earns the complementarity bonus 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, $\chi = 0$ (orange), which removes complementarity and every correlation between a consumer's checking and card parameters. Under the restricted model, joint pricing alone yields no synergy, and the combined firm loses both deposit and credit share to rivals. Under the full model, the combined firm lowers its average deposit rate and leaves average rewards essentially unchanged.

Table 1: Summary Statistics

	Data Source			
Characteristic	Prolific	MRI	NHCS	SCF
Panel A: Demographics				
Mean Age (years)	45.4	47.9	48.8	49.9
Median Income (in thousand dollars)	61.2	79.7	83.4	77.3
Bachelor's Degree or Above (%)	55.1	35.9	34.7	39.2
Credit Score >= 720 (%)	65.7	78.1		
Carrying Credit Card Balance (%)	45.4	36.6	43.4	40.5
Panel B: Product Holdings (%)				
Deposit Account (%)	99.2	94.4	94.1	94.2
Credit Card (%)	85.9	76.0	73.1	77.8
Auto Loan (%)	22.4	24.3	31.4	33.8
Mortgage (%)	29.4	41.3	32.0	40.6
Brokerage Account (%)	31.6	32.5	24.5	23.2
Panel C1: Account Balances - Revolvers (thousands of 2025 dollars)				
Median Total Deposits (k$)	1.3		6.8	4.1
Median Credit Card Balance (k$)	1.0			3.9
Median Credit Card Spending (k$/yr)	2.9	4.5		4.0
Panel C2: Account Balances - Transactors (thousands of 2025 dollars)				
Median Total Deposits (k$)	10.0		27.4	23.7
Median Credit Card Spending (k$/yr)	8.2	12.1		18.0
Observations	2,012	100,884	39,716	14,732

Notes: This table reports summary statistics across our main data sources. Prolific refers to our original survey of approximately 2,000 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 2019, 2022, and 2025 Surveys of Consumer Finances, pooled with equal wave weights. Dollar amounts are in 2025 dollars. Panel A compares demographic characteristics. Carrying a credit card balance is the share among credit-card holders. 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 generalize.

Table 2: Factors for Banking Product Considerations

	Deposit OLS	Deposit Bank FE	Credit OLS	Credit Bank FE
Credit Used	0.080***	0.072***		
	(0.018)	(0.018)		
Credit Considered	0.247***	0.230***		
	(0.010)	(0.011)		
Deposit Used			0.206***	0.202***
			(0.021)	(0.020)
Deposit Considered			0.279***	0.241***
			(0.013)	(0.013)
Constant	0.092***	0.095***	0.142***	0.147***
	(0.003)	(0.003)	(0.003)	(0.003)
Observations	22132	22132	22132	22132

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

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 3: Structural Parameter Estimates, Population Means

Parameter	Transactors	Revolvers
Price Sensitivity		
Checking (α_1)	1.16 (0.22)	0.78 (0.13)
Credit (α_2)	2.57 (0.71)	0.35 (0.08)
Cross-Selling Effects		
Consideration lift: Checking (χ_1^c, pp)	4.39 (1.67)	4.52 (1.34)
Consideration lift: Credit (χ_2^c, pp)	11.97 (2.35)	8.38 (2.21)
Utility: Checking ($E[\max(\chi_1^u, 0)]$, pp)	0.52 (0.20)	0.65 (0.18)
Utility: Credit ($E[\max(\chi_2^u, 0)]$, pp)	0.40 (0.14)	1.76 (0.24)
Switching Costs		
Checking ($E[\max(\kappa_1, 0)]$, pp)	1.34 (0.50)	2.13 (0.53)
Credit ($E[\max(\kappa_2, 0)]$, pp)	1.00 (0.32)	5.09 (0.58)
Arrival Rates		
Checking (λ_1)	0.28 (0.02)	0.23 (0.02)
Credit (λ_2)	0.70 (0.23)	0.47 (0.08)
Adoption Rates (from age 16)		
Checking (ν_1)	0.11 (0.01)	0.09 (0.01)
Credit (ν_2)	0.08 (0.00)	0.07 (0.00)

Notes: Reports population means from the MC-EM estimator. Utility complementarity (χ^u) and switching costs (κ) enter utility as the positive part of a normal latent, generating a point mass at zero; the table reports the population mean of that positive part divided by the segment's price coefficient α_k, in percentage points of the product's own price (deposit rate for checking, rewards rate for transactor cards, APR for revolver cards). Price sensitivity (α) is homogeneous within segment. Consideration complementarity (χ^c) is the percentage-point lift in the probability that a consumer considers a bank for one product given an existing relationship with that bank in the other product. Estimation uses the Prolific survey. After the MC-EM step, bank mean utilities are refined on the MRI-Simmons cross-section (approximately 100,000 observations) via a weighted steady-state likelihood (see Section V). Bootstrap standard errors from 40 replicates are in parentheses.

Table 4: Estimated Heterogeneity: Standard Deviations and Correlations

	Transactors	Revolvers
Panel A: Consideration		
Standard deviations		
$\mathrm{sd}(v_1^c)$	1.1 (0.071)	1.1 (0.068)
$\mathrm{sd}(v_2^c)$	1.2 (0.066)	1.3 (0.070)
$\mathrm{sd}(\eta_1^c)$	1.0 (0.054)	0.84 (0.078)
$\mathrm{sd}(\eta_2^c)$	0.81 (0.059)	0.77 (0.085)
Correlations		
$\rho(v_1^c, v_2^c)$	0.97 (0.0060)	0.97 (0.0074)
$\rho(\eta_1^c, \eta_2^c)$	0.91 (0.083)	0.84 (0.13)
Panel B: Preference		
Bank intercept standard deviations		
$\mathrm{sd}(v_1^u)$	0.47 (0.056)	0.40 (0.051)
$\mathrm{sd}(v_2^u)$	0.54 (0.10)	0.30 (0.037)
Parameter standard deviations, pp of price		
$\mathrm{sd}(\chi_1^u)$	0.81 (0.32)	1.4 (0.50)
$\mathrm{sd}(\chi_2^u)$	0.72 (0.38)	3.6 (0.92)
$\mathrm{sd}(\kappa_1)$	1.0 (0.44)	1.6 (0.56)
$\mathrm{sd}(\kappa_2)$	0.36 (0.21)	2.7 (0.52)
Correlations		
$\rho(v_1^u, v_2^u)$	0.50 (0.22)	0.42 (0.13)
$\rho(\chi_1^u, \chi_2^u)$	0.63 (0.16)	0.63 (0.20)
$\rho(\kappa_1, \kappa_2)$	0.15 (0.14)	0.76 (0.14)
Share at zero		
$P(\chi_1^u = 0)$	0.34 (0.064)	0.45 (0.074)
$P(\chi_2^u = 0)$	0.39 (0.059)	0.43 (0.077)
$P(\kappa_1 = 0)$	0.11 (0.031)	0.10 (0.036)
$P(\kappa_2 = 0)$	0.0031 (0.0085)	0.030 (0.020)

Notes: Reports latent standard deviations, within-bank correlations, and the share of consumers at zero, with bootstrap standard errors from 40 replicates in parentheses. Panel A reports consideration heterogeneity and Panel B reports preference heterogeneity. Bank intercept standard deviations are in utility units. Standard deviations of utility complementarity (χ^u) and switching costs (κ) are the latent standard deviation divided by the segment's price coefficient α_k, in percentage points of the product's own price, as in Table 3. Consideration and utility bank intercepts have separate covariance blocks. Utility complementarity and switching costs may correlate within and across products. Price sensitivity, consideration complementarity, and arrival-rate dispersions are fixed at 0.05; their covariance rows are fixed. The displayed structural standard deviations come from the parameter export for the reported fit.

Table 5: Decomposition of Cross-Holding Channels

	Transactors	Revolvers	Pooled
P(both products at same bank), full model (%)	26.3 (1.3)	22.0 (1.1)	24.3 (0.8)
P(both products at same bank), independence (%)	5.2 (0.3)	5.7 (0.2)	5.4 (0.2)
Excess cross-holding (pp)	21.1 (1.4)	16.3 (1.2)	18.9 (0.9)
Shapley share of excess cross-holding (%)			
Cross-product correlation	22 (4)	28 (4)	24 (3)
Ex-post utility complementarity (χ^u)	65 (5)	57 (4)	62 (4)
Information complementarity (χ^c)	13 (3)	14 (2)	13 (2)

Notes: The first two rows report the steady-state probability that a consumer holds checking and a card at the same listed bank, measured over the full simulated population. Independence removes all three channels, so checking and card choices are independent. Excess cross-holding is the difference between the two rows. The lower panel splits the excess into Shapley shares, which average each channel's marginal contribution over all six orders of removal. Cross-product correlation sets to zero every covariance between a checking parameter and a card parameter, in both the bank intercepts and the other consumer parameters. Bootstrap standard errors from 40 replicates are in parentheses.

Table 6: Implied Margins from Supply-Side First-Order Conditions

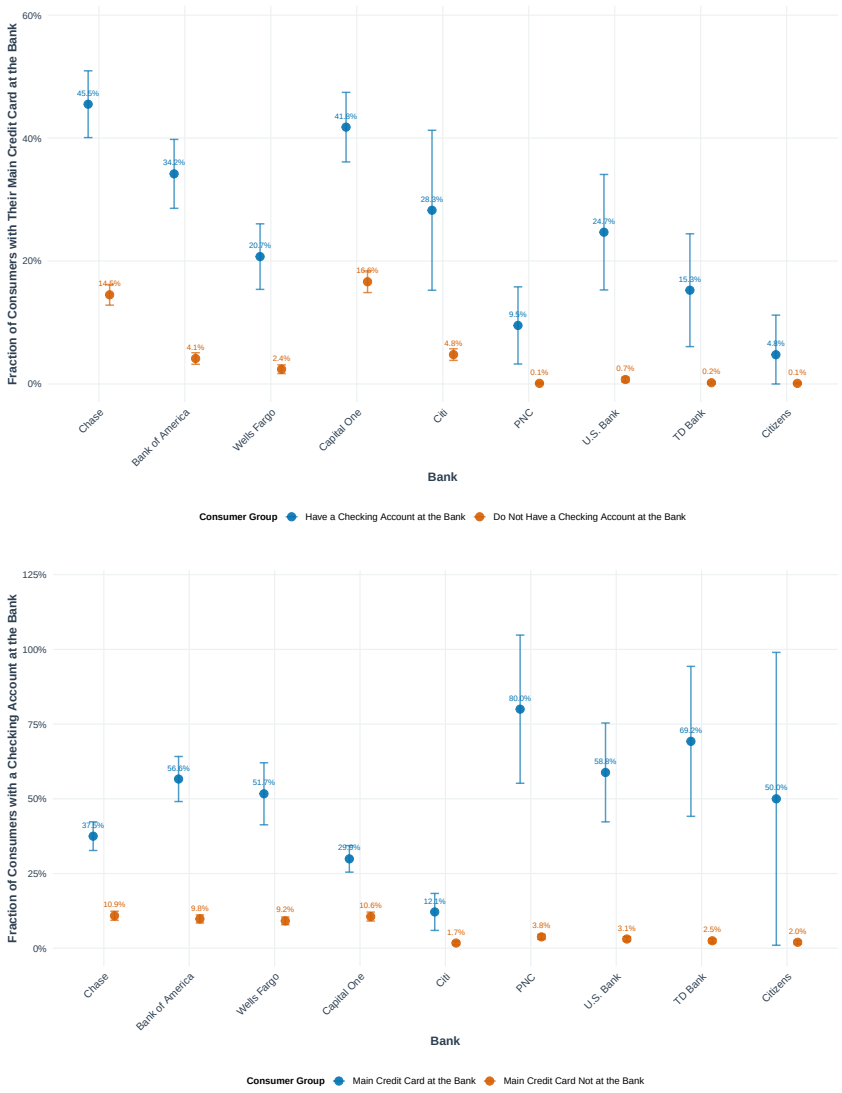
Statistic	Value
Median margin (pp)	
Checking	1.19 (0.44)
Credit (Trans)	0.27 (0.26)
Credit (Rev)	3.06 (0.70)

Notes: This table reports implied margins 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. Bootstrap standard errors from 40 replicates are in parentheses.

Appendices. For Online Publication Only

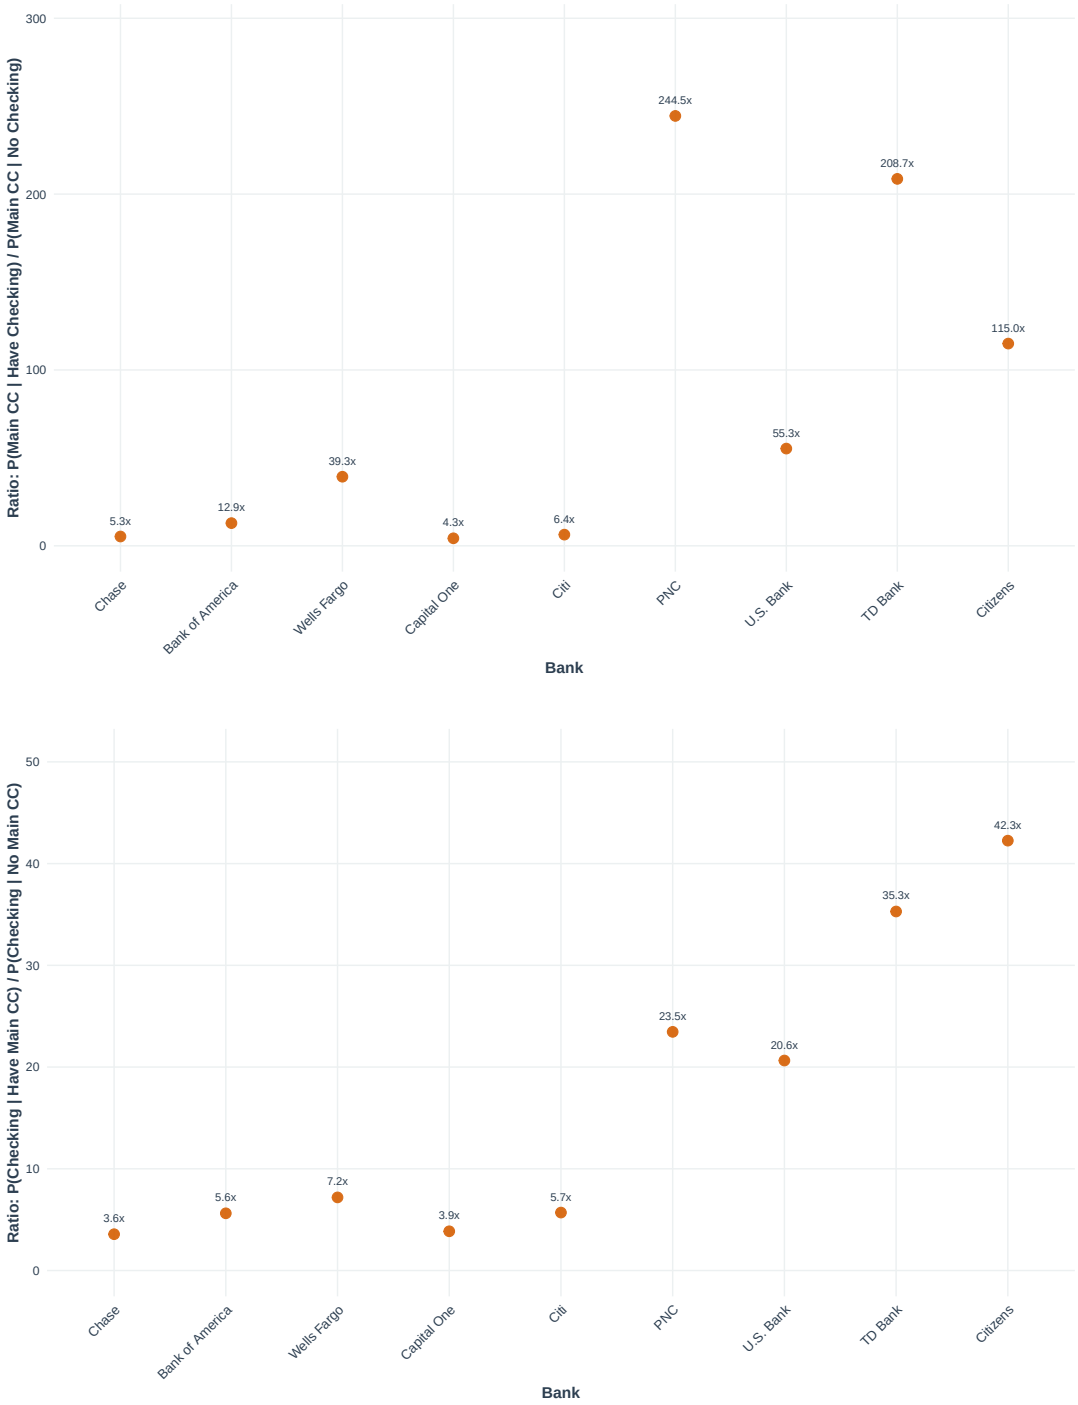
A Additional Figures and Tables

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



Notes: The top panel reports the fraction of Prolific respondents whose main credit card is at each bank, separately for respondents who do and do not have a checking account at that bank. The bottom panel reports the analogous fraction whose checking account is at each bank, separately for respondents whose main credit card is and is not at that bank. The Prolific plots use the respondent's primary credit-card bank, matching the one-card-per-consumer definition used for the MRI comparison.

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



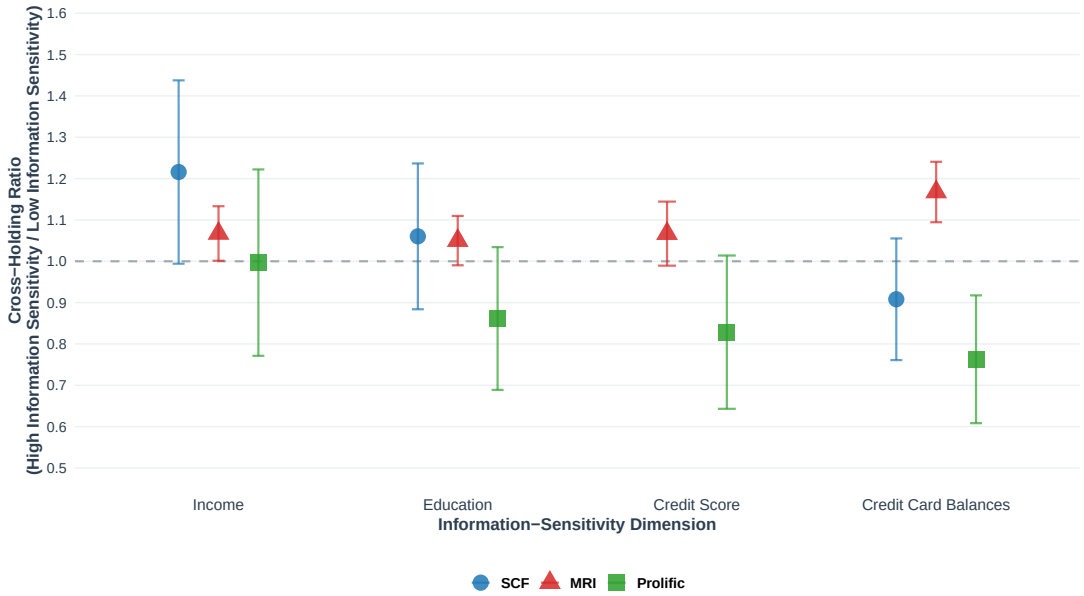
Notes: This figure shows the prevalence of cross-holding across banks. The top panel plots, for each bank, the ratio of credit card market shares among consumers who do and do not have a checking account at the bank. The bottom panel plots the analogous ratio for checking account market shares conditional 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



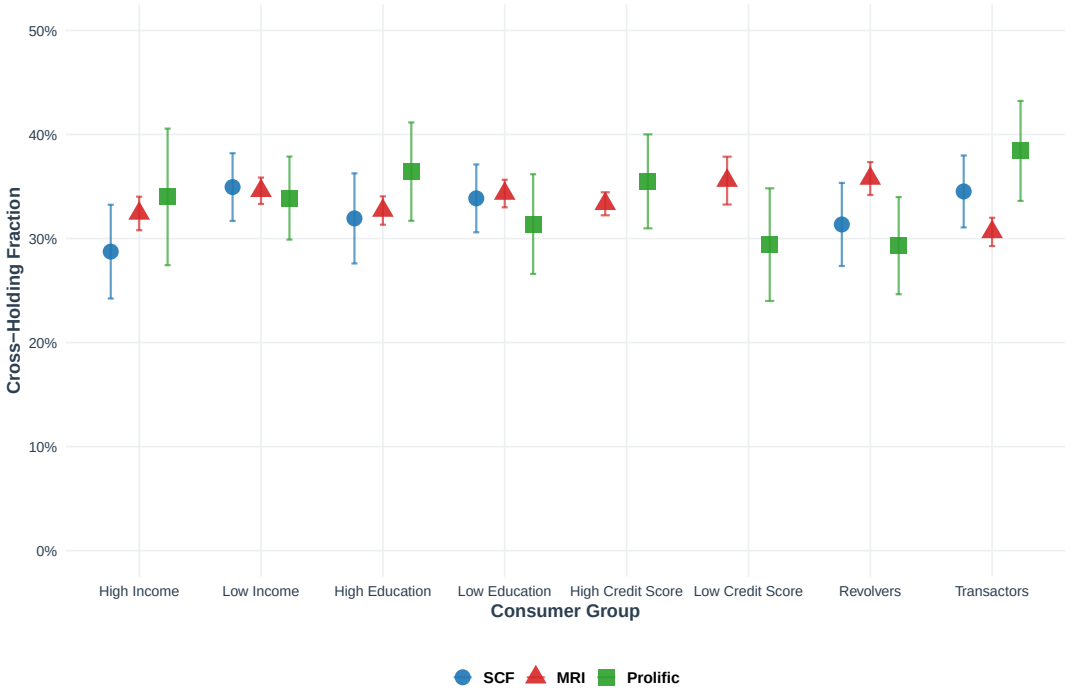
Notes: This figure shows the distributions of the sizes of the consideration sets for deposit accounts and credit cards in our Prolific Survey Sample. The consideration set includes each consumer's bank of choice, and other banks they "looked into or reviewed" when choosing the product.

Figure A.4: Comparing Cross-holding of Disadvantaged versus Advantaged Groups



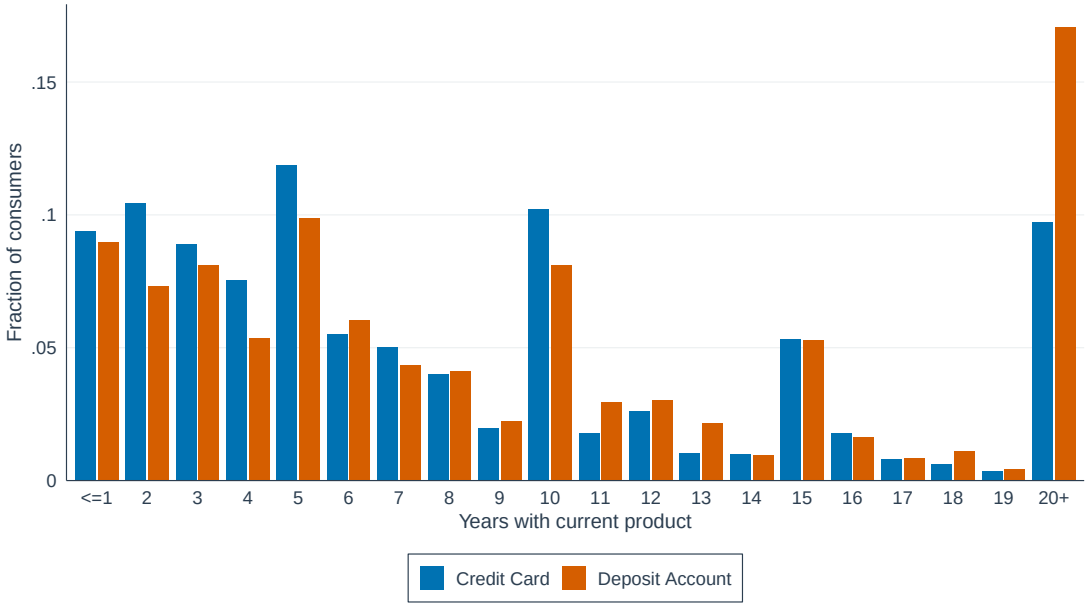
Notes: The figure compares cross-holding between disadvantaged and advantaged consumer groups, computed for consumers with exactly one checking account and one credit card. We compare groups defined by income, education, credit score, and whether consumers carry a credit card balance. Specifically, the 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 by consumers who carry a balance sometimes or often. Transactors pay their card balances in full.

Figure A.5: Cross-holding by Consumer Subgroups: Detailed View



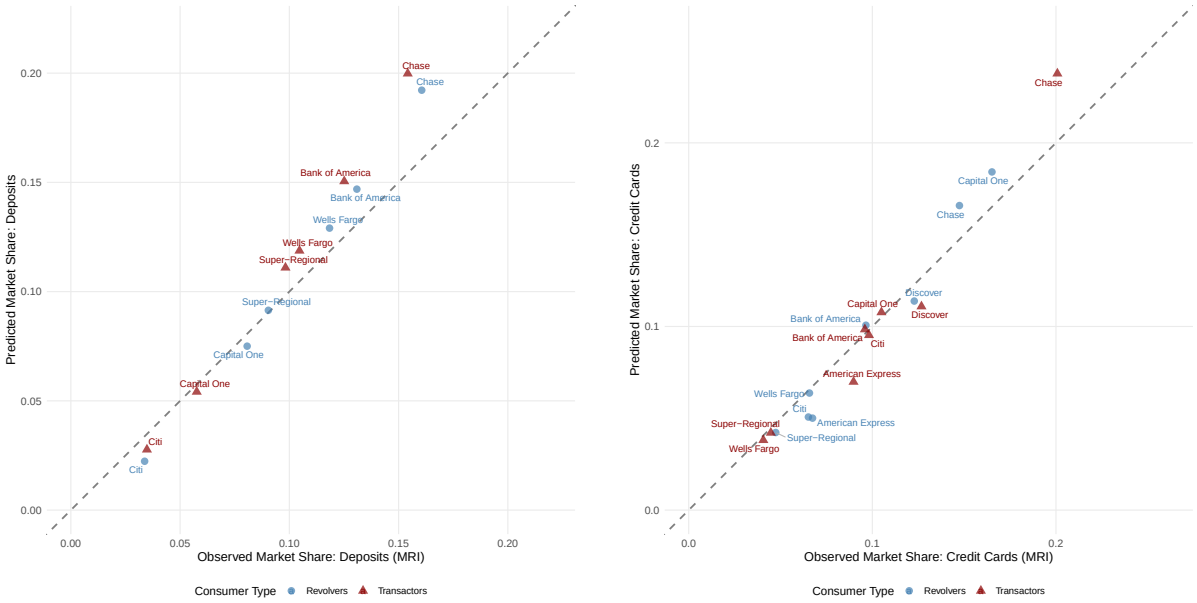
Notes: This figure provides a detailed breakdown of cross-holding rates by consumer subgroups. Each panel shows the rate of credit card cross-holding separately for different demographic characteristics: income, education, credit score, and credit card balance. Within each panel, we compare disadvantaged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

Figure A.6: Consumer Tenure at Current Banks



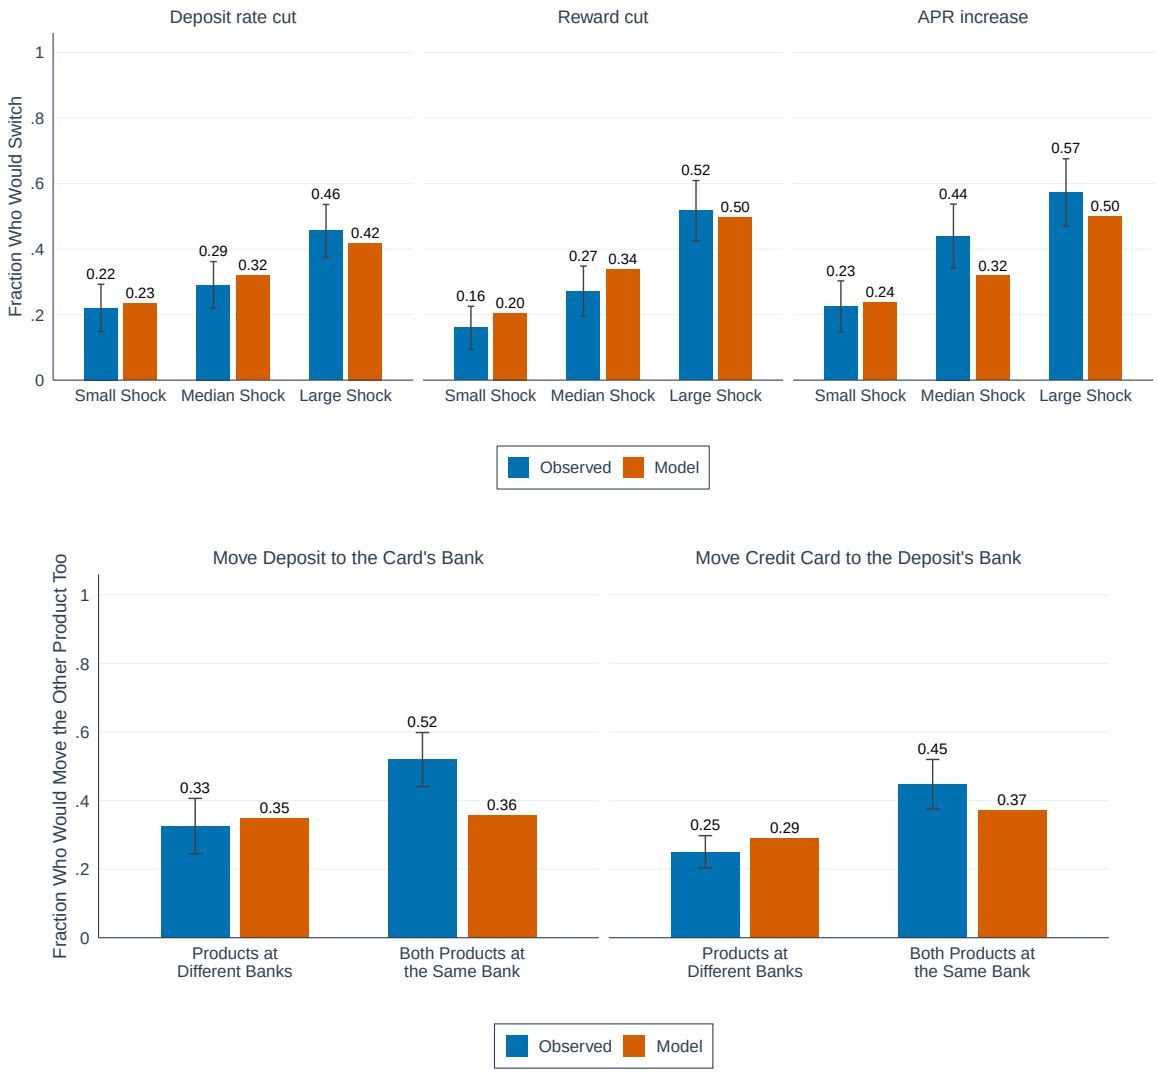
Notes: This figure shows the distribution of consumer tenure at their current bank for credit accounts and deposit accounts, separately. 0 means the consumer obtained their current product less than a year ago, while 20 means more than twenty years ago.

Figure A.7: Goodness of Fit: Market Shares



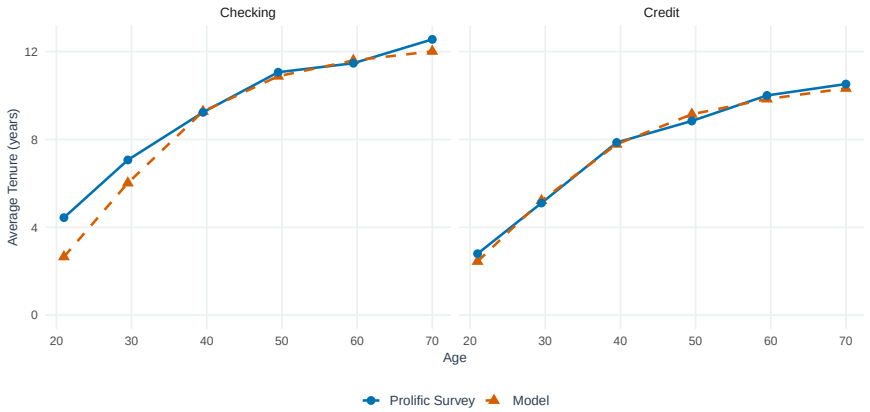
Notes: These figures compare observed and model-predicted market shares for checking accounts (left) and credit cards (right). Each point represents a bank–segment pair (transactor and revolver shown separately, distinguished by color and shape), with the horizontal axis showing observed shares and the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

Figure A.8: Goodness of Fit: Stated Switching Responses



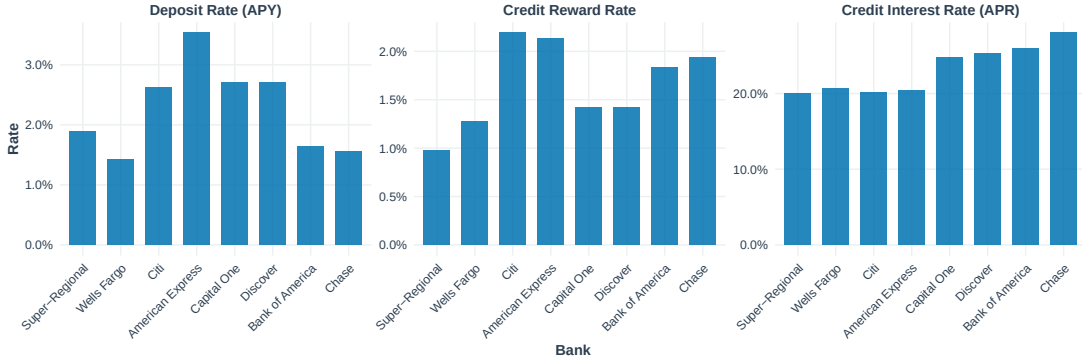
Notes: The top panel compares observed and model-predicted switching responses to own-product price shocks of different magnitudes. The bottom panel compares responses to moving the other product, separately for consumers initially holding the two products at different banks and at the same bank. Model moments simulate long-run holdings and then administer the same hypothetical questions as in the Prolific survey.

Figure A.9: 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 simulate the adoption-augmented transient process G_i: consumers start at age 16 with neither product, adopt through the product-specific ν_k clocks, and then switch under the estimated λ_k wake-up law. Tenure is measured from the most recent adoption or effective switch. Each simulated tenure is rounded to integer years and top-coded at 20 before averaging over holders, matching the Prolific reporting convention.

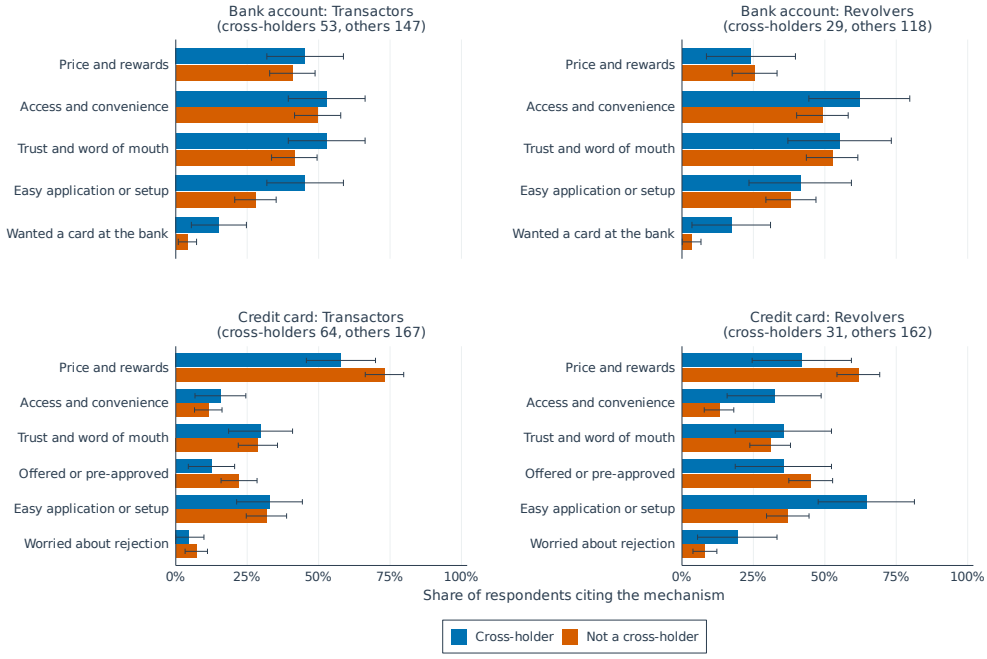
Figure A.10: Observed Prices by Bank



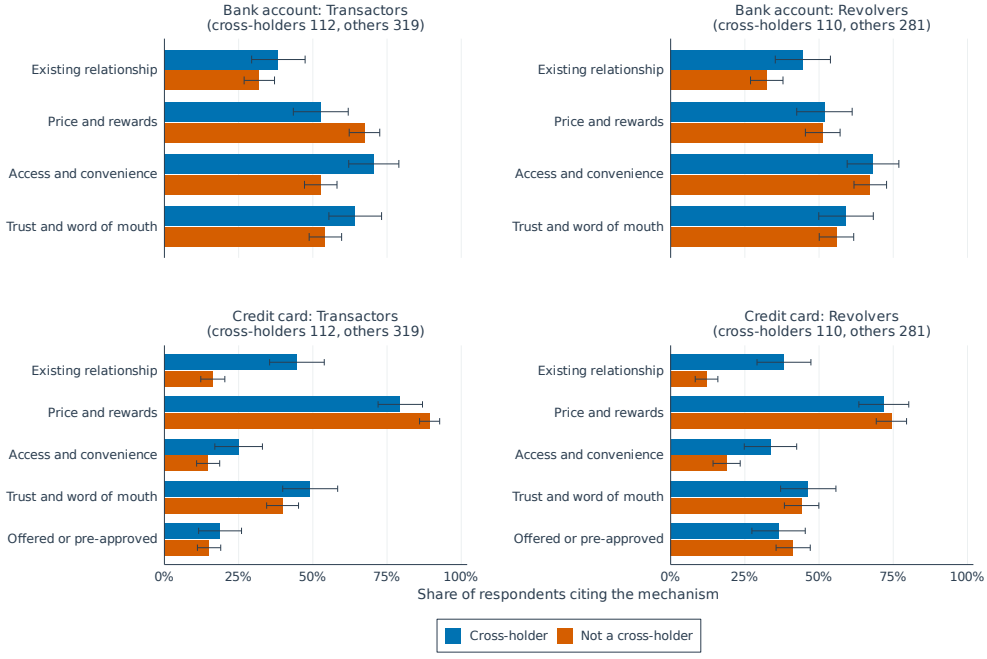
Notes: This figure reports the observed prices used as the supply-side targets in the model: deposit interest rates, credit card reward rates, and credit card APRs by bank. Deposit rates are constructed from WRDS bank holding company filings as annual interest expense on domestic deposits divided by year-end 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.

Figure A.11: Stated Reasons for Choosing a Bank, Cross-Holders and Other Consumers

(a) 2025 wave



(b) 2026 wave



Notes: Share of Prolific cardholders who cite each group of reasons for choosing the bank of their primary credit card or savings account, with 95% confidence intervals. Cross-holders hold their credit card at their checking bank. Both waves asked about interest rates, fees, and rewards, access and convenience, and trust and word of mouth. The 2025 wave (panel a) asked about initial account opening (select all that apply) and added application ease and forward-looking motives. The 2026 wave (panel b) allowed up to three reasons and asked directly about existing banking relationships.

Table A.1: 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	12.0	11.7	16.4	22.7	17.0
Bank of America	11.1	9.5	11.0	10.0	9.6	8.3
Wells Fargo	10.6	7.9	8.3	5.7	5.2	3.0
Capital One	5.2	10.9	2.2	13.4	23.6	14.1
PNC	3.0	3.1	2.3	1.3	0.6	0.5
Citi	2.9	1.3	4.1	8.1	6.2	12.1
US Bank	2.8	3.1	2.5	2.1	2.0	3.3
TD Bank	1.8	1.7	2.2	0.7	0.8	1.2
Citizens Bank	1.3	1.3	1.5	0.6	0.2	0.2
American Express	-	3.1	0.5	7.0	7.9	9.0
Discover	-	2.5	0.4	10.6	10.6	7.6
Others	47.6	43.8	53.2	24.1	10.8	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 (based on active accounts).

Table A.2: Cross-holding Rates and Comparison to an Independent-Choices Benchmark

Product	Source	Cross-Holding Rate	Cross-Holding Ratio
Credit Card	Prolific	34.0% (1.7)	3.71 (0.17)
	MRI	33.8% (0.5)	3.99 (0.03)
	SCF	33.2% (1.3)	—
Auto Loan	Prolific	8.9% (1.6)	3.53 (0.43)
	SCF	21.3% (1.5)	—
Mortgage	Prolific	9.8% (1.5)	3.41 (0.38)
	SCF	20.2% (1.2)	—
Brokerage	Prolific	4.7% (0.8)	4.61 (0.52)
	MRI	5.0% (0.3)	3.41 (0.07)
	SCF	7.8% (1.2)	—

Notes: This table reports aggregate cross-holding rates and the aggregate cross-holding ratio for each product and dataset along with standard errors. Aggregate cross-holding rates are computed as the fraction of consumers who hold a checking account and a second financial product at the same bank among those consumers who hold exactly one checking account and one secondary product. For the MRI dataset aggregate cross-holding rate, if either the primary checking account or the primary secondary product is held at a bank outside the set of named institutions, we count the consumer as not cross-holding. The cross-holding ratio is the ratio of observed cross-holding rate to an independent-choice benchmark on the full dataset that includes consumers who multi-home across banks for a given product line. In the Prolific dataset, we observe primary accounts and use those to compute cross-holding rates. For the MRI dataset, if a consumer multi-homes, we designate a random bank as the primary bank for that product.

Table A.3: DiD Estimates: Capital One Deposit Uptake among Walmart Cardholders

	Full sample		Target control	
	(1) Deposit	(2) Credit card	(3) Deposit	(4) Credit card
Walmart cardholder (Treat)	0.0562***	0.1761***	0.0262***	0.0510***
	(0.0055)	(0.0079)	(0.0080)	(0.0112)
Post (2020+)	0.0409***	0.0655***	0.0403***	0.0424***
	(0.0012)	(0.0017)	(0.0060)	(0.0084)
Walmart × Post (DiD)	0.0483***	0.6240***	0.0408***	0.6685***
	(0.0081)	(0.0080)	(0.0113)	(0.0122)
Implied conversion rate	0.077		0.061	
	(0.013)		(0.017)	
Observations	200,921	200,921	18,335	18,335
R-squared	0.009	0.103	0.010	0.326

Notes: This table shows Difference-in-Differences (DiD) estimates of the impact of the 2019 Walmart-Capital One portfolio transfer on Capital One's deposit market shares among Walmart cardholders. The control group consists of everyone without a Walmart card in columns 1 and 2, and restricts to consumers with a Target card in columns 3 and 4. The credit-card outcome is issuer-coded: every Walmart cardholder is coded as holding a Capital One card from 2020 onward, since all Walmart cards are Capital One-issued after the transfer. The implied conversion rate is the deposit DiD divided by the card DiD, with delta-method standard errors. Linear probability models excluding 2019; robust standard errors in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table A.4: Triple Difference: Specificity of the Capital One Effects

	Full sample		Target control	
	(1) Deposit	(2) Credit card	(3) Deposit	(4) Credit card
Capital One × Walmart × Post (DDD)	0.0635***	0.6397***	0.0503***	0.6673***
	(0.0084)	(0.0090)	(0.0119)	(0.0136)
Implied conversion rate	0.099		0.075	
Person × Year FE	Yes	Yes	Yes	Yes
Bank × Year FE	Yes	Yes	Yes	Yes
Bank × Walmart FE	Yes	Yes	Yes	Yes
Observations	1,004,605	1,004,605	91,675	91,675

Notes: This table reports triple-difference estimates of the impact of the 2019 Walmart–Capital One portfolio transfer on Capital One uptake, adding deposit and credit-card holdings at the four placebo banks (Chase, Bank of America, Citi, and Wells Fargo) as a further control dimension. Columns 1 and 2 use the full sample; columns 3 and 4 restrict the control group to Target cardholders. All specifications absorb person×year, bank×year, and bank×Walmart fixed effects. The credit-card outcome is issuer-coded as in Table A.3; issuer-coding alters Capital One cells only. Standard errors, clustered by person-year, in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table A.5: Multi-Bank Holding and Concentration in the Largest Account

	Bank accounts	Credit cards
Panel A: Hold the product at two or more banks (%)		
MRI 2021–2022	42.1 (0.3)	57.8 (0.3)
Prolific survey	22.1 (0.9)	50.2 (1.2)
SCF 2019, 2022 (institutions)	36.6 (0.6)	46.4 (0.7)
Panel B: Concentration of balances and spending (%)		
Share of checking balances at largest bank	73.6 (0.3)	
Share of credit card transactions on most-used card	78.8 (1.5)	

Notes: Standard errors in parentheses. Panel A reports the share of product holders who hold the product at two or more banks. MRI-Simmons (2021–2022, population weights, $N = 87{,}140$ checking and 81,637 card holders) records every bank used in the past year and counts checking. The Prolific survey ($N = 1{,}995$ and 1,728) counts the primary bank and any other bank the respondent uses for checking or savings. The SCF (2019 and 2022, pooled weights, five implicates, $N = 9{,}773$ and 8,334 households) counts checking institutions and card-issuing institutions. Panel B restricts to consumers with two or more. The checking row uses the SCF (2019, 2022, and 2025, $N = 5{,}590$ households with checking at two or more institutions) and divides the balance at the institution with the largest balance by total checking balances. It excludes households with more than six checking accounts, whose extra accounts carry no institution code. The credit card row uses the Diary of Consumer Payment Choice (2017–2020) and reports the share of diary credit card transactions made on the most-used card. The sample is holders of two or more cards with at least five card transactions in the three-day diary ($N = 298$), because a consumer with one transaction puts all of it on one card by construction. If each of these transactions went to one of the consumer's cards at random, the share would be 60.7%.

B Data and Balance Parameter Calibration

B.1 Balance Parameters in the Model

Balance parameters convert per-dollar margins into profits per consumer. We calibrate deposit balances $b_1(\sigma)$ by segment, annual card spending b_2^{spend} for transactors, and outstanding card balances b_2^{balance} for revolvers.

B.2 Data Sources and Calibration

We calibrate balances from survey medians, separately for revolvers and transactors. The Survey of Consumer Finances (SCF) supplies household deposit balances, outstanding card balances, and card spending. The National Household Credit Survey (NHCS) supplies deposit balances, and MRI-Simmons supplies annual card spending derived from spending brackets. Deposit balances average the SCF and NHCS medians, transactor spending averages the SCF and MRI medians, and revolver outstanding balances use the SCF median. Table B.6 reports each source and the resulting calibration.

For revolvers, the averaged values yield deposits of \$5,472 and credit card outstanding balances of \$3,850.

For transactors, the averaged values yield deposits of \$25,552 and credit card spending of \$15,063 per year, with zero outstanding balances by definition.

Comparison across sources The SCF and NHCS report similar deposit medians for transactors but differ more for revolvers. The revolver medians are \$4,103 and \$6,841, respectively. The transactor medians are \$23,741 and \$27,364. Transactor card spending differs more across sources, with medians of \$18,000 in the SCF and \$12,126 in MRI. We give the two sources equal weight in the spending calibration.

Rationale for Using Medians We use medians to calibrate the balances of a typical consumer in each segment. Deposit means are substantially higher than medians (Table B.7). This calibration does not capture within-segment balance heterogeneity or its association with demand responses.

B.3 Primary and Secondary Accounts

Many consumers hold the same product at more than one bank, but they concentrate their balances and spending in a single primary account. Across our surveys, 46% to 58% of credit-card holders and 22% to 42% of checking holders use two or more banks (Appendix Table A.5). Among multi-homers, however, households keep an average of

Table B.6: Median Balance Parameters by Consumer Type and Source

Parameter	SCF	NHCS	MRI	Prolific	Nilson	Supply
Population shares (supply FOC pooling)						
Revolvers	40.5%	43.4%	36.6%	45.4%	—	41.5%
Transactors	59.5%	56.6%	63.4%	54.6%	—	58.5%
Panel A: Revolvers (Balance Carriers)						
Deposits ($)	4,103	6,841	—	1,275	—	5,472
CC Outstanding ($)	3,850	—	—	971	7,926*	3,850
Panel B: Transactors (Non-Carriers)						
Deposits ($)	23,741	27,364	—	10,000	—	25,552
CC Spending ($/yr)	18,000	—	12,126	8,236	28,287*	15,063

Notes: The Supply column reports the calibrated values for revolvers (Panel A) and transactors (Panel B). Deposit balances average SCF and NHCS medians, transactor spending averages SCF and MRI medians, and revolver outstanding balances use the SCF median. Population shares average the SCF, NHCS, MRI, and Prolific estimates. Prolific dollar balances and Nilson ratios are shown for comparison. A dash indicates an unavailable measure, and asterisks identify Nilson industry ratios.

Table B.7: 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.

74% of their checking balances at their largest depository institution (SCF; Appendix Table A.5), and cardholders with two or more cards make 79% of their card transactions on their most-used card (Diary of Consumer Payment Choice; Greene and Stavins, 2021). Card borrowing is similarly concentrated in the CFPB Consumer Credit Panel, where 80% to over 90% of consumers, depending on credit score, carry more than half of their card balances on a single card (Nelson, 2025, Online Appendix A.8). Our Prolific survey multi-homes at similar rates, with 50% holding cards at two or more banks (versus 46% in the SCF) and 22% holding bank accounts at two or more banks (versus 37% in the SCF). In every dataset, card multi-homing differs by at most five percentage points between revolvers and transactors. Our survey and model therefore track each consumer's primary deposit account and primary credit card (Section IV.C).

C Likelihood Construction Details

The retrospective likelihood combines event-specific consideration and choice probabilities with the density of the recorded event times. The formulas below apply conditional on consumer type; the estimator then integrates over types.

C.1 First Openings and Generator Blocks

The transition rates of the full generator are given in Equation 8. The outside option denotes an unlisted provider for either product. Restricting the rows and columns to states with at least one coordinate at -1 gives A_i. Restricting those rows to post-adoption destination columns gives B_i. In particular, each pre-adoption diagonal satisfies

$$(A_i)_{xx} = - \sum_{y \neq x : \text{at least one } y_k = -1} (A_i)_{xy} - \sum_{y \in (\{0\} \cup \mathcal{J})^2} (B_i)_{xy}. \tag{C.1}$$

A first opening and a later switch use the same destination contribution D_k in the event factors of Equation 17. The event rate is ν_{ik} for a first opening and λ_{ik} for a switch.

C.2 Record Density

The observation window has length $E = \text{age}_i - 16$ and begins at $(-1, -1)$. Equation 17 combines three propagation or survival factors with two event factors and a deposit-ownership normalizer. The lead factor $\left[\exp\big(G_i(E - t_1)\big)\right]_{(-1,-1),\, x(t_1^-)}$ is the law of the pre-event state under the full generator. Each recorded event contributes its clock times the destination weight $D_k = P(C_k \mid x)\, s_{ix_k^* k}(x, C_k)$. Between the older and newer events, $\widetilde{Q}_{ik}$ zeroes the older product's off-diagonal transitions while retaining its exit rates on the diagonal; the other product's generator is left intact, so unrecorded moves on that product are integrated out. After the newer event, both products are killed. For the deposit-first case in Equation 17, $\left[\exp\big(t_2(\widetilde{Q}_{i1} + \widetilde{Q}_{i2})\big)\right]_{x^*, x^*}$ gives survival at the current holdings. The sample normalizer $1 - e^{-\nu_{i1} E}$ conditions on having a deposit account by the survey date.

An event's reported other-product holding may not distinguish never adopted from an unlisted provider. Where the survey leaves that distinction unresolved, the event operator sums over the compatible augmented states.

C.3 Censored Event Times

Event times are top-coded at $T_{\max} = 20$ years. When the two products share a calendar year and the survey does not name an order, we return to the order mixture in Section C.4. When order is named but the two events share a calendar year, we move

the older event back one year, within the survey's integer-year resolution, so the record has a one-year gap. A tie at the top code is left censored, and a move that reaches the top code censors the older event.

A censored event time lies in $[T_{\max}, E]$. The record density integrates the corresponding event time over that finite window. A block-triangular exponential evaluates the integrated portion of the record density,

$$\exp\left(\begin{pmatrix} G^{\text{pre}} & \mathcal{E} \\ 0 & G^{\text{post}} \end{pmatrix} s\right) = \begin{pmatrix} e^{G^{\text{pre}} s} & \int_0^s e^{G^{\text{pre}}(s-u)}\, \mathcal{E}\, e^{G^{\text{post}} u}\, du \\ 0 & e^{G^{\text{post}} s} \end{pmatrix},$$

where G^{pre} and G^{post} are the generators governing the path before and after the censored event, $\mathcal{E}$ is its event operator, and $s = E - T_{\max}$ is the remaining exposure. The same construction applies to first openings and later switches, using the corresponding event rate. For a censored older event, the initial-state distribution premultiplies this block. The remaining factors propagate holdings from the censoring threshold to the newer event, score that event, and impose survival through the survey date.

If both events are censored, we integrate over $T_{\max} < t_n < t_o \leq E$, where t_o and t_n denote the older and newer lookback times. The upper-right block of a three-block triangular exponential evaluates the ordered integral over the remaining exposure. Final survival over $T_{\max}$ is applied separately. Both cases retain the deposit-ownership normalizer.

C.4 Incomplete Recall

Missing survey items are integrated out of the retrospective or response likelihood.

Forgotten origin. For an unrecalled prior provider, we sum the record density over listed and unlisted banks, excluding the recorded destination. The same rule applies to deposits and credit cards; first openings are handled separately. Reported moves between unlisted providers are recoded as forgotten origins because the model represents all such providers by one state. Each candidate contributes its joint record density, including its pre-event probability.

Unknown order. When both products have the same reported tenure year and no named order, we sum the record densities under the two possible orders. Each order uses the compatible holding of the newer product at the older event and moves the hypothesized older event back one year, as for a named-order tie. A named order is not mixed.

Unrecalled consideration sets. If the respondent does not recall the banks reviewed at a past event, the destination contribution is replaced by the unconditional choice probability

$$D_k(\theta_i) = m_{ix_k^*k}\big(x(t_k^-)\big) = \sum_{S \in 2^{\mathcal{J}}} P_{ik}\big(S \mid x(t_k^-)\big)\, s_{ix_k^*k}\big(x(t_k^-), S\big).$$

An unobserved *prospective* consideration set is integrated out of the stated-preference likelihood as described in Section C.6.

Unasked hypotheticals. A question that was not asked is omitted from the interval restrictions in Section C.6. A response of "would not switch" is retained as an informative no. We exclude the first wave's hypothetical responses because its instruments differ from those in the second wave. Those respondents contribute L^{retro} alone.

C.5 Non-Cardholders

Respondents with no recorded card adoption remain at credit -1. Their checking histories contribute under the same process and deposit-ownership normalizer as other respondents. Remaining without a card contributes survival under the first-card clock, so it is informative about ν_{i2}. A card held at an unlisted issuer instead occupies state 0 and follows reconsideration at rate λ_{i2}. Unasked card hypotheticals contribute no response restriction.

C.6 Hypothetical-Shock Likelihood

Stated-preference responses share one draw of the idiosyncratic utility shocks within respondent and product. This draw is independent of the shocks at earlier account openings. The elicited prospective set is an outcome. The likelihood scores $P(C)$ with the same consideration technology as the retrospective block, evaluated at current holdings, and then scores the named challenger and the response pattern given that set.

Incumbent and challenger. For a fixed product and consumer type, let x be the current provider, x_- the other-product provider, and C the prospective consideration set. A listed incumbent is always in C, and the alternatives are $C \cup \{0\}$. An incumbent at $x = 0$ holds an account at an unlisted provider; the outside slot is that incumbent. The switching-cost advantage is $d_{\text{ideal}} = \kappa$ for both listed and unlisted incumbents. In the hypothetical, an unlisted incumbent is the respondent's own provider rather than the aggregate of unlisted providers, so the full switching cost applies in place of κ^0. Otherwise the utility $\delta_{ijk}(x_i)$ is as specified in Section IV.

The challenger a is the stated ideal when that ideal differs from x, and the stated second choice otherwise. It is the best non-current alternative under the common shock draw. The response likelihood includes only observations with $a \in \mathcal{J}$ because the cross-product hypothetical moves the other account to a listed bank.

Modified shares and response thresholds. Equation 18 defines the modified share $s_a(d; C, \theta_i)$. The unbundling reduction d_{unb} equals $2\chi^u$ when the other product moves from the same listed incumbent to a, χ^u when it moves from a third provider, and zero when it is already at a. An unlisted other-product provider receives no complementarity, so its move to a uses $d_{\text{unb}} = \chi^u$. An inattentive consumer may already prefer the challenger at the survey date, so the incumbent–challenger gap can be negative.

Joint response cells. Write $\text{cell}_a(C, \theta_i) = [s_a(\min_{h \in Y_{ik}} d_h; C, \theta_i) - s_a(\max_{h \notin Y_{ik}} d_h; C, \theta_i)]_+$ for the joint cell in Equation 19. Unasked questions contribute no bound. Summing over response patterns yields $s_a(\infty; C, \theta_i)$, retaining the probability of naming a. Equation 19 multiplies the cell by $P(C \mid x, \theta)$. We condition the cell on a listed challenger, dividing by $P(a \in \mathcal{J} \mid C, \theta)$, without imposing that the incumbent remains optimal under the survey shock draw. For a listed incumbent that probability is the listed mass among non-current alternatives in $C \cup \{0\}$; for an unlisted incumbent it is 1 whenever C is nonempty.

Unobserved prospective sets. When the prospective set is unobserved, the same joint is mixed over consideration sets,

$$L_{ik}^{\text{hyp}} = \frac{\sum_{C \subseteq \mathcal{J}: a \in C} P(C \mid x, x_-, \theta)\, \text{cell}_a(C, \theta)}{\sum_{C \subseteq \mathcal{J}} P(C \mid x, x_-, \theta)\, P(a \in \mathcal{J} \mid C, \theta)}.$$

A listed incumbent is forced into every set. The denominator includes sets without the observed challenger because selection requires any listed challenger.

C.7 Model-Implied Tenure by Age

We compute the model-implied Prolific age profiles by forward simulation. Simulations use the adoption-augmented generator G_i from the retrospective likelihood. Each consumer starts at age 16 with neither product. First openings occur at rates (ν_1, ν_2), followed by reconsideration and choice under the estimated switching process. At each age, tenure is the time since the most recent adoption or effective switch. We round each simulated tenure to integer years and top-code it at 20 before averaging over consumers who hold the product. The diagnostic compares these averages with the corresponding Prolific survey moments.

C.8 MC-EM Algorithm and Estimation Details

The EM algorithm integrates the joint retrospective and response likelihood over normal latent coefficients. A joint maximum-likelihood fit supplies initial means, using several starting regimes for the stated-preference parameters.

Importance-sampling E-step. For consumer i, the proposal density combines a Gaussian posterior approximation with the population prior,

$$q_i^{(g)}(\theta) = (1-\omega)\phi(\theta; \widehat{m}_i, \widehat{C}_i) + \omega\phi(\theta; \overline{\theta}^{(g)}, \Sigma^{(g)}),$$

where $g \leq \ell$ is the iteration at which consumer i's draws were last generated. We set $\omega = 0.2$. A fixed share ω of each consumer's draws comes from the prior component, and every draw is scored under the full mixture density, so the prior-to-proposal ratio is bounded by $1/\omega$. The Gaussian component takes the consumer's posterior mean and covariance from the most recent E-step, masked to the block structure of the population covariance. If any of its correlations exceeds 0.98 in absolute value, all correlations are shrunk proportionally toward zero until the largest equals 0.98, which preserves the variances and positive definiteness. At the second iteration, the Gaussian component pairs the posterior mean with the population covariance.

At the first iteration no posterior approximation exists yet. The Gaussian component then uses the consumer's posterior under the likelihood without its event-time factors, which requires no matrix exponentials. We compute this static posterior by adaptive multiple importance sampling in rounds of 100,000 draws. The first round draws from the population prior. Each later round draws from the same mixture, with its Gaussian centered at the current static posterior mean and carrying the population covariance, and every draw is weighted against the mixture of all rounds so far. Rounds stop once the static effective sample size reaches 200 or after ten rounds. The static posterior only places proposals; the weights always use the full likelihood.

At iteration ℓ, normalized posterior weights are

$$w_{ir}^{(\ell)} \propto L(X_i \mid \theta_{ir}) \frac{\phi(\theta_{ir}; \overline{\theta}^{(\ell)}, \Sigma^{(\ell)})}{q_i^{(g)}(\theta_{ir})}, \qquad \sum_r w_{ir}^{(\ell)} = 1.$$

Between regenerations, each draw and its likelihood are retained, and only the prior density in the numerator is recomputed. The weights therefore remain exact as the population parameters move; movement lowers the effective sample size rather than

biasing the weights.

Draw counts and regeneration. Each consumer starts with 500 draws. All consumers' draws are regenerated at the first two iterations, whenever the symmetric Kullback–Leibler divergence between the current population prior and the prior at the last full regeneration exceeds one nat, and every 100 iterations. Between full regenerations, a consumer whose effective sample size falls below 50 receives new draws from the Gaussian component at their current posterior moments. For a consumer whose effective sample size falls below 10, the Gaussian component is replaced by the static posterior evaluated at the current population prior, because their full-likelihood posterior moments rest on too few effective draws. A redrawn consumer's draw count targets an effective sample size of 150, using the effective sample size per draw of that consumer's most recent redraw. The count is bounded between 150 and 10,000 draws. Fresh draws whose effective sample size exceeds the target are trimmed to the target count, and retained draws are trimmed only above 1.5 times the target. Sizing the draw count from the most recent redraw prevents overshooting as parameter updates erode the effective sample size of retained draws.

The mean of each search propensity η_{ik}^c is normalized to zero. After each M-step, any common shift in η_{ik}^c moves into the bank means of v_{ijk}^c, leaving the likelihood unchanged. In constructing consumer-level proposal distributions, however, we retain each consumer's unshifted posterior mean of η_{ik}^c, because a consumer-level shift in η_{ik}^c and an equal offsetting shift across that consumer's v_{ijk}^c are close substitutes in the likelihood. Moving the proposal along that ridge would center it away from the consumer's posterior, lower the effective sample size, and pull the posterior means of η_{ik}^c toward zero.

Population updates. The Gaussian moment updates are

$$\bar{\theta}^{(\ell+1)} = \frac{1}{N} \sum_i \sum_r w_{ir}^{(\ell)} \theta_{ir},$$
$$\Sigma^{(\ell+1)} = \frac{1}{N} \sum_i \sum_r w_{ir}^{(\ell)} (\theta_{ir} - \bar{\theta}^{(\ell+1)})(\theta_{ir} - \bar{\theta}^{(\ell+1)})^\top.$$

The implementation projects these updates onto the covariance restrictions. Means whose variances are fixed follow the homogeneous-parameter update below rather than the unrestricted moment formula.

Latent distributions and restrictions. Bank consideration and utility intercepts form separate covariance blocks, each allowing within-bank correlation across products. Utility complementarity and switching costs may covary within and across products. The reconsideration rates form a separate block with fixed covariance; their cross-covariances with the utility-side block are restricted to zero.

Homogeneous parameters. Price sensitivity α_k, consideration complementarity χ_k^c, and the arrival rates (ν_k, λ_k) are treated as common within each segment. The algorithm represents them with normal pseudo-latents whose covariances with freely varying coordinates are zero. Their final standard deviation is 0.05, except for the first-opening rates ν_k, whose final standard deviation is 0.01. Their means are estimated by a missing-information Newton step on the observed-data score. Each coordinate's step is capped at half of its prior standard deviation per iteration, and the first iteration takes a plain EM step. Because the Newton signal for a pinned mean scales with its prior standard deviation, the algorithm anneals these standard deviations. They start at 0.2, or 0.05 for ν_k, and hold there for 1,000 iterations. They then fall geometrically to their final values over 500 iterations and remain there. The algorithm runs at most 3,000 iterations and cannot stop before the annealing ends, so its endpoint is a fixed point under the final standard deviations. The imposed dispersion is a computational approximation, not estimated consumer heterogeneity.

Latents are stored as normal draws in levels. At utility evaluation, the price sensitivity, utility complementarity, and switching-cost latents are truncated at zero; the reported coefficients are their common or positive-part values as appropriate. At generator evaluation, both arrival rates have a floor of 0.02 per year. The latent draws remain untruncated in the moment updates. For numerical evaluation, response probabilities are floored at 10^{-20} before taking logarithms. This floor bounds the penalty for impossible response patterns; it is not a behavioral model of response error.

Post-EM refinement. The MRI-Simmons refinement estimates bank mean utilities by maximizing a weighted steady-state likelihood, holding the other parameters at their EM values. The sample includes checking-only consumers and joint holders. The refinement integrates over the estimated heterogeneity using the exported consumer draws.

D Model Counterparts of the Reduced-Form Evidence

Population The Walmart portfolio transfer and Basten and Juelsrud (2023) exercises evaluate the estimated model on the same consumer types as the supply side and the

counterfactuals. A consumer belongs to segment σ with probability π_σ and has type θ_i drawn from F_σ, with types weighted equally within each segment and segments weighted by π_σ. Given θ_i, a consumer's holdings follow the continuous-time Markov chain with generator Q_i at baseline prices. Prices stay at their baseline values throughout both exercises, so banks do not reprice. The population distribution over holdings at any date is the weighted mixture of the type-level distributions. Type i has steady-state distribution $\bar{s}_i(\theta, \sigma)$, and a type starting from distribution s has distribution $s\, e^{Q_i \tau}$ after τ years. We compute these distributions exactly from the matrix exponential, without simulating individual histories.

Conditioning on holdings We compute each group statistic as a ratio of population probabilities because conditioning on holdings selects on consumer types. Consumers whose types favor a bank are more likely to hold its checking account or credit card. For any group defined by holdings, the statistic divides the weighted probability that a consumer is in the group and has the outcome by the weighted probability that a consumer is in the group. An equal-weighted average of the per-type conditional probabilities would remove the selection that correlated preferences create.

Walmart portfolio transfer In the Walmart exercise, the Capital One card stands in for the Walmart card. Before the transfer, Walmart cards were issued by Synchrony, which offers no checking account. Each type therefore starts in the pre-transfer steady state of the estimated model with Capital One's consideration and utility complementarity set to zero, which implies no pre-trend. At the transfer date, Capital One's complementarity returns to its estimated value, and each type's distribution evolves under the estimated generator for 1, 2, and 3 years. At each date, the treated group consists of consumers who hold a Capital One card and the comparison group consists of everyone else, including consumers with no card. Group membership changes over time as consumers switch cards, matching the repeated cross-sections in MRI-Simmons.

We reweight consumer types by entropy balancing because Walmart cardholders in the data differ from Capital One cardholders in the model. We choose type weights closest in Kullback–Leibler divergence to the population weights so that the pre-transfer Capital One deposit share is 11.6% in the treated group and 6.0% in the comparison group, matching the 2016–2018 MRI-Simmons means. Each weight belongs to a type and remains fixed across dates. The treatment effect at horizon τ is the change from before the transfer in the treated group's Capital One deposit share minus the corresponding change in the comparison group. This difference-in-differences matches the Walmart $\times$

Post coefficient from the regression in Section III.

Deposit relationships and later borrowing Following the main specification of Basten and Juelsrud (2023), we compare the probability of opening a card at a bank within τ years between the bank's own depositors (the treated group) and depositors at other listed banks (the comparison group). For each of the six checking banks, we start in steady state and keep consumers with no card at that bank and a checking account at any listed bank. Excluding consumers without a checking account matches the banked-only comparison group in Basten and Juelsrud (2023). We remove from Q_i every state in which the consumer holds a card at that bank. The restricted generator governs paths without a card opening there, so each row sum of its matrix exponential at τ is the probability of not opening a card by year τ. Each bank's relationship effect is the opening probability in the treated group minus that in the comparison group, which we average across the six banks using depositor-count weights. In the specification without correlated preferences, we rebuild the type draws from the same underlying shocks with every covariance between a checking parameter and a card parameter set to zero. All other parameters, including mean utilities, stay at their estimated values.

E Supply-Side Cost Recovery

Given demand estimates and observed prices, each bank's pricing conditions form a linear system in its marginal costs. The system pools deposit demand across transactors and revolvers while keeping rewards and APR incentives separate. We evaluate $\bar{s}_i(\theta, \sigma)$ at the observed-price steady state and hold it fixed when differentiating profits with respect to a bank's prices.

Index product lines by $\ell, m \in \{\text{dep}, \text{rew}, \text{apr}\}$, as in Section V.C. Prices are expressed in the consumer-preferred direction, so $p_{j,\text{apr}} = -p_j^{\text{apr}}$. For these line indices, define

$$
\begin{aligned}
O_{j,\text{dep}} &= O_{j1}, \quad b_{\text{dep}}(\sigma) = b_1(\sigma), \\
O_{j,\text{rew}} &= O_{j2}, \quad b_{\text{rew}}(\sigma) = \mathbf{1}\{\sigma = \text{T}\}\, b_2^{\text{spend}}, \\
O_{j,\text{apr}} &= O_{j2}, \quad b_{\text{apr}}(\sigma) = \mathbf{1}\{\sigma = \text{R}\}\, b_2^{\text{balance}}.
\end{aligned}
$$

Deposits contribute in both segments; each card line contributes only in its corresponding segment. The recovered deposit cost is common across segments, while the two card lines have separate costs.

Discounted balances and their price responses are

$$H_{jm} = \sum_{\sigma} \pi_{\sigma} \, \mathbb{E}_{\theta|\sigma} \left[\bar{s}_i \, \mathcal{R} \, O_{jm} \, b_m(\sigma) \right], \tag{E.2}$$

$$\frac{\partial H_{jm}}{\partial p_{j\ell}} = \sum_{\sigma} \pi_{\sigma} \, \mathbb{E}_{\theta|\sigma} \left[\bar{s}_i \, \mathcal{R} \, \frac{\partial Q_i}{\partial p_{j\ell}} \, \mathcal{R} \, O_{jm} \, b_m(\sigma) \right], \tag{E.3}$$

where the arguments (θ, σ) of $\bar{s}_i$, $\mathcal{R}$, and Q_i are suppressed. A price change alters transition rates and hence the discounted balances the bank expects to retain or acquire. The derivative in Equation E.3 follows from the resolvent identity, holding initial holdings fixed.

With margins $\mu_{jm} = -p_{jm} - c_{jm}$, discounted profits are $\Pi_j = \sum_m H_{jm} \mu_{jm}$. The first-order condition for line ℓ is

$$\sum_{m} \frac{\partial H_{jm}}{\partial p_{j\ell}} \, \mu_{jm} = H_{j\ell}. \tag{E.4}$$

Balances and their derivatives depend on demand parameters and observed prices, but not on marginal costs. Substituting for margins therefore gives a linear system,

$$\sum_{m} \frac{\partial H_{jm}}{\partial p_{j\ell}} \, c_{jm} = -H_{j\ell} - \sum_{m} \frac{\partial H_{jm}}{\partial p_{j\ell}} \, p_{jm}.$$

For banks offering all three lines, the solution recovers costs for deposits, transactor cards, and revolver cards. For card-only issuers, the system includes only the card lines.

Normalize each response by the own-line balance,

$$\zeta_{j,\ell m} \equiv \frac{1}{H_{j\ell}} \frac{\partial H_{jm}}{\partial p_{j\ell}}. \tag{E.5}$$

This normalization yields

$$\sum_{m} \zeta_{j,\ell m} \, \mu_{jm} = 1. \tag{E.6}$$

Isolating $\mu_{j\ell}$ gives the own-product markup and cross-selling adjustment in Equation 16.