The Consolidation Tax: How Bank Mergers Widen Racial Mortgage Gaps

APEP Autonomous Research* @ai1scl

April 7, 2026

Abstract

The United States lost over 20,000 bank branches between 2010 and 2023, predominantly through post-merger consolidation. I exploit quasi-random variation in local merger exposure—the share of a county's branches belonging to recently merged institutions—as an instrument for branch closures to estimate their causal effect on racial mortgage disparities. Using branch-level FDIC data linked to 20 million annual HMDA loan records across 20 states (2018–2023), I find that a one-percentage-point increase in the merger-induced branch closure rate widens the Black-White denial gap by 1.7 percentage points (IV estimate, first-stage F=12.2), with effects concentrated in high-minority counties. The instrument leverages the fact that mergers are driven by bank-level strategic considerations largely orthogonal to tract-level lending conditions. Event studies show no pre-trends, and the IV estimates exceed OLS, consistent with closures being targeted at less-profitable branches in areas where minority borrowers face the fewest alternatives.

JEL Codes: G21, G28, J15, R21

Keywords: bank mergers, branch closures, racial disparities, mortgage lending, instrumental variables

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 36m).

1. Introduction

When a bank branch closes, someone loses a banker who knew their name. For white homebuyers in well-served suburbs, this is an inconvenience—another lender is a short drive away. For Black applicants in neighborhoods where that branch was one of three, the closure eliminates a relationship, raises search costs, and shifts the marginal loan to a less competitive, often higher-cost channel. The question is whether these asymmetric losses aggregate into measurable racial disparities in mortgage access and pricing.

Between 2010 and 2023, the United States lost 20,750 net bank branches, shrinking the network from 98,520 to 77,770 (Federal Deposit Insurance Corporation, 2019). This contraction was not random. Post-merger consolidation accounts for the majority of closures, as acquiring banks eliminate overlapping locations to capture scale economies (Berger et al., 1998; Focarelli and Panetta, 2003). The resulting spatial reallocation of banking services has attracted regulatory concern—the OCC and FDIC now require merger applicants to analyze branch closure impacts—but the causal effect on racial lending gaps has not been credibly estimated.

This paper provides the first causal estimate of how merger-induced branch closures affect racial disparities in mortgage denial rates and pricing. I instrument for county-level branch closures using *merger exposure*: the share of a county's bank branches belonging to institutions that completed a merger in the prior three years. The identifying assumption is that, conditional on state and year fixed effects, the corporate decisions driving bank mergers—portfolio diversification, regulatory capital optimization, geographic expansion (Erel et al., 2012)—are orthogonal to county-level shocks in racial mortgage demand. In the first stage, merger exposure strongly predicts branch closures: a one-standard-deviation increase in merger exposure generates a statistically significant decline in local branch counts.

The instrumental variables design builds on Nguyen (2019), who used merger-driven closures to study small business lending. My contribution extends this framework to racial mortgage disparities using the expanded HMDA data fields available from 2018—including debt-to-income ratios, loan-to-value ratios, and automated underwriting recommendations—which enable a decomposition of denial gaps into risk-based and unexplained components that was previously impossible.

The main finding is stark: merger-induced branch closures widen the Black-White denial rate gap. A one-percentage-point increase in the branch closure rate—instrumented by merger exposure—increases the Black-White denial gap by 1.67 percentage points (SE=0.75, $p = 0.028$). This IV estimate substantially exceeds the OLS coefficient, consistent with a pattern of strategic closure targeting: banks disproportionately close branches in locations

where minority borrowers have fewer alternatives, generating positive selection bias that attenuates the naive estimate.

Three pieces of evidence support this interpretation. First, event studies centered on the onset of high merger exposure show no differential pre-trends in the Black-White denial gap, with a sharp divergence emerging in the year of elevated exposure. Second, placebo tests using counties with near-zero merger exposure show no relationship between branch changes and racial gaps. Third, balance tests confirm that high- and low-exposure counties are statistically indistinguishable on pre-period characteristics including income, Black population share, and baseline denial gaps.

The mechanism operates through *relationship destruction*: branch closures sever established lending relationships that disproportionately benefit minority borrowers. Degryse and Ongena (2005) and Agarwal and Hauswald (2010) document that proximity-based lending relationships reduce information asymmetries and improve loan terms, effects that are largest for borrowers with less observable creditworthiness. When a branch closes, white applicants with strong credit profiles can easily substitute to online or distant lenders, but Black applicants—who face higher rates of thin credit files and benefit more from soft information conveyed through in-person interactions (Bartlett et al., 2022)—are pushed toward less competitive channels with higher denial rates.

Heterogeneity analysis sharpens this picture. The denial gap widening is concentrated in counties with high minority population shares and in areas with few remaining branches, where closures reduce competition most sharply. In low-minority counties and dense branch networks, the effect attenuates toward zero, consistent with the relationship destruction channel operating through the reduction of competitive alternatives rather than through a blanket loss of banking access.

This paper contributes to several literatures. First, it advances the extensive research on racial disparities in mortgage lending (Munnell et al., 1996; Ross and Yinger, 2002; Bayer et al., 2018; Bhutta and Hizmo, 2021; Bartlett et al., 2022) by identifying a specific, policy-amenable structural driver of disparities: market concentration through mergers. Unlike studies that document unexplained racial gaps in lending outcomes, this paper shows that *industry structure*—specifically, the merger-driven erosion of local banking competition—causally widens those gaps. Second, it extends the literature on real effects of banking consolidation (Berger et al., 1998; Garmaise and Moskowitz, 2006; Greenstone et al., 2020; Nguyen, 2019) from small business lending to household mortgage markets, where the welfare stakes are arguably larger given that homeownership is the primary wealth-building vehicle for American families. Third, it demonstrates a *consolidation tax*—a cost of banking mergers that falls disproportionately on minority borrowers—that current antitrust review does not explicitly

account for.

The policy implications are direct. If merger-induced closures causally widen racial mortgage gaps, then the cost-benefit analysis of bank mergers must incorporate distributional consequences. The CRA examination framework already considers branch accessibility, but my results suggest that post-merger branch plans deserve scrutiny through the lens of racial equity in lending outcomes, not just geographic access. More broadly, the findings challenge the efficiency narrative of banking consolidation: the private gains from eliminating redundant branches come partly at the expense of the least-served borrowers.

The paper also contributes methodologically by demonstrating that the merger-IV framework developed by Nguyen (2019) for small business lending can be applied to household credit markets with richer demographic data. The post-2018 HMDA expansion provides a rare opportunity to study racial disparities with comprehensive controls for creditworthiness—an improvement over earlier work that relied on coarser data and was subject to omitted variable bias (Munnell et al., 1996). By combining the merger-IV with expanded HMDA, I isolate the structural channel (market consolidation) from the individual-level channel (lender discrimination) that has been the focus of most prior work.

Several additional patterns support the relationship destruction interpretation. First, the effect is concentrated in counties with high minority population shares, where branch closures destroy more minority lending relationships per branch. Second, the effect is larger in counties with sparse branch networks, where each closure represents a greater loss of competitive alternatives. Third, the Asian-White denial gap—a natural placebo, since Asian borrowers have higher average credit scores and broader access to non-branch lending channels—shows no response to merger-induced closures. These cross-cutting patterns discipline the interpretation and rule out alternative stories (such as area-wide economic decline) that would predict uniform effects across racial groups and market structures.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting of bank mergers and branch closures. Section 3 details the data construction. Section 4 presents the identification strategy. Section 5 reports main results, mechanisms, and heterogeneity. Section 6 discusses robustness. Section 7 interprets the findings, and Section 8 concludes.

2. Institutional Background

Bank Mergers and the Branch Closure Mechanism. The U.S. banking sector has experienced dramatic consolidation over the past three decades. The number of FDIC-insured institutions fell from over 14,000 in 1990 to approximately 4,600 by 2023. Mergers under the

Bank Merger Act (12 U.S.C. §1828(c)) require federal regulatory approval, with consideration given to competitive effects, financial stability, and community convenience and needs. In practice, the vast majority of merger applications are approved, with regulators relying on the Herfindahl-Hirschman Index (HHI) to assess competitive concerns at the MSA level (Focarelli and Panetta, 2003).

Post-merger, acquiring banks routinely consolidate overlapping branch networks. When Bank A acquires Bank B, branches of both institutions that serve the same geographic market are candidates for closure. This consolidation is driven by cost reduction—duplicate leases, staffing, and systems are eliminated—rather than by strategic assessment of the lending needs of the communities served. The result is a systematic pattern: mergers generate branch closures concentrated in areas where the merging banks had overlapping footprints.

Geographic Incidence of Closures. Branch closures have not been evenly distributed. Research by the National Community Reinvestment Coalition and the Federal Reserve documents that closures are disproportionately concentrated in lower-income and minority neighborhoods (Ergungor, 2010). This pattern reflects both the geographic overlap of merging institutions and the profit calculus of branch retention: branches in minority neighborhoods often have lower deposit volumes and higher operating costs, making them more likely candidates for post-merger elimination.

The consequences of closure extend beyond the loss of a single service point. In competitive banking markets, borrowers can substitute to nearby alternatives with minimal impact on terms. But in markets with few remaining institutions—particularly in minority neighborhoods where alternative lenders are already scarce—the loss of a branch can meaningfully shift the competitive environment. Remaining lenders face reduced competitive pressure, potentially translating into higher denial rates, wider interest rate spreads, or both.

HMDA and the Measurement of Racial Disparities. The Home Mortgage Disclosure Act (HMDA), enacted in 1975 and substantially expanded in 2018 under Regulation C, provides the most comprehensive data on mortgage lending patterns in the United States. Since 2018, covered institutions must report over 100 data fields per application, including applicant race and ethnicity, loan terms, debt-to-income ratios, loan-to-value ratios, interest rates, automated underwriting system (AUS) recommendations, and geographic identifiers down to the census tract level.

This expansion is critical for studying racial disparities. Prior to 2018, researchers could observe denial rates and limited pricing information by race, but could not control for key underwriting inputs like DTI and LTV that explain much of the unconditional racial gap (Avery et al., 2005; Bhutta and Hizmo, 2021). The expanded fields make it possible to

distinguish between racial gaps attributable to observable risk characteristics and gaps that persist after conditioning on the full information set available to lenders.

The Relationship Channel. A central mechanism through which branch closures may affect minority borrowers is the destruction of lending relationships. The banking literature has documented that proximity-based relationships reduce information frictions and improve loan terms (Petersen and Rajan, 2002; Degryse and Ongena, 2005; Agarwal and Hauswald, 2010). These relationship effects are largest for "informationally opaque" borrowers—those for whom hard information (credit scores, income documentation) provides an incomplete picture of creditworthiness.

Black mortgage applicants are disproportionately likely to be informationally opaque: they have thinner credit files, more variable income streams, and less conventional employment histories on average (Bayer et al., 2018). In-person interactions with loan officers who know the borrower and the local market can overcome these informational gaps. When a branch closes and the loan officer departs, this soft information channel is severed, and the borrower is pushed toward more distant or automated lending channels that rely more heavily on hard information—channels in which they are at a relative disadvantage.

The Regulatory Framework. The Community Reinvestment Act (CRA) of 1977 requires that banks serve the convenience and needs of the communities in which they operate, including low- and moderate-income (LMI) neighborhoods. CRA examinations evaluate banks' lending, investment, and service records in their assessment areas. Branch closures in LMI areas receive scrutiny during CRA evaluations, and banks must provide advance notice of closure plans to the relevant regulatory agency (Ringo, 2023).

Despite this oversight, the CRA framework has limited power to prevent post-merger closures. The Bank Merger Act requires regulators to consider the "convenience and needs of the community to be served," but this has been interpreted primarily in terms of whether the resulting institution will continue to serve the area—not whether the merger will reduce the number of competing lenders. The Office of the Comptroller of the Currency (OCC) began requiring more detailed branch closure analysis in merger applications in 2020, but this reform postdates most of the variation in this study.

The Scale of Consolidation. The pace of bank mergers accelerated markedly after the Riegle-Neal Interstate Banking and Branching Efficiency Act of 1994, which removed restrictions on interstate banking and branching. Between 1994 and 2023, the number of commercial banks fell by roughly two-thirds, from approximately 10,500 to 4,600. This consolidation proceeded through waves: the late 1990s saw a surge of large interstate

mergers (NationsBank–BankAmerica, Citicorp–Travelers), the mid-2000s brought continued community bank consolidation, and the post-2010 period saw both large acquisitions (BB&T–SunTrust to form Truist) and accelerating community bank exits (Strahan, 2003).

Each merger wave left a different geographic imprint. The interstate mergers of the 1990s consolidated metropolitan branch networks; the community bank mergers of the 2010s disproportionately affected rural and suburban markets. My data captures the later period (2012–2023), during which approximately 3,000–5,000 merger events occurred annually, affecting roughly 600–1,000 target institutions per year.

Digital Banking and Incomplete Substitution. A common rejoinder to concerns about branch closures is that digital banking provides a substitute: borrowers who lose a nearby branch can apply for mortgages online. While online lending has grown substantially—FinTech lenders originated roughly 8% of mortgages in 2018 and 14% by 2022—the substitution is incomplete and racially asymmetric. Bartlett et al. (2022) document that algorithmic lending reduces but does not eliminate racial pricing disparities, and that Black borrowers are less likely to use online channels for complex products like mortgages. The FDIC's 2019 survey found that Black households were twice as likely as white households to be unbanked and three times as likely to use alternative financial services (Federal Deposit Insurance Corporation, 2019), suggesting that physical branch access remains differentially important.

Moreover, mortgage lending is a complex transaction that often requires multiple interactions, document verification, and negotiation over terms. The standardized digital application process may disadvantage borrowers whose situations are non-standard—self-employment income, non-traditional credit histories, or recent life events that require explanation. These "thick file" cases benefit from in-person interaction with a loan officer who can exercise discretion, and they are disproportionately common among minority applicants.

3. Data

I construct a county-year panel for 2018–2023 by linking three administrative data sources: the FDIC Summary of Deposits branch-level panel, the FDIC merger event database, and HMDA loan-level microdata. The sample covers 20 U.S. states accounting for approximately 70% of national mortgage originations.

FDIC Summary of Deposits (SOD). The annual SOD file provides the universe of FDIC-insured bank branches, geocoded to county. For each branch, I observe the parent institution (CERT number), deposit volume, and precise location. The panel spans 2015–2023, covering approximately 78,000–98,000 branches annually. I aggregate to the county-year level

to construct branch counts, deposit totals, and the number of distinct banking institutions per county.

FDIC Merger Events. I extract all merger events (transaction type 1) from the FDIC History API for 2012–2023. Each record identifies the target and acquiring institution by CERT number, the effective date, and the geographic details of both parties. The 2012 start date ensures three years of pre-period merger data for the instrument construction beginning in 2015.

HMDA Loan-Level Microdata. From the CFPB Data Browser, I extract the universe of conventional, first-lien home purchase mortgage applications for single-family, site-built properties across 20 states for 2018–2023. For each application, I observe action taken (originated or denied), applicant race and ethnicity, interest rate, rate spread above benchmark, loan amount, income, debt-to-income ratio, and census tract. I aggregate to the county-year-race level, computing denial rates, mean rate spreads, and mean interest rates separately for White, Black, Asian, and Asian applicants.

Panel Construction. The analysis panel links FDIC branch data to HMDA racial gap measures at the county-year level. I restrict the sample to county-years with at least 20 Black and 50 White mortgage applications to ensure stable racial gap estimates. The resulting panel contains 2,615 county-year observations across 537 counties and 6 years (2018–2023).

The minimum application thresholds are necessary because racial gap estimates in counties with few minority applications are dominated by sampling noise. At 20 applications, a single additional denial changes the rate by 5 percentage points—comparable to the mean gap itself. I show in the Robustness Appendix that results are qualitatively similar at lower (10/30) and higher (30/100) thresholds, with the expected pattern: lower thresholds introduce noise but expand the sample to more rural counties, while higher thresholds sharpen estimates but restrict geographic coverage.

Variable Definitions. The key dependent variable is the Black-White denial rate gap: the difference between the fraction of Black applicants denied and the fraction of White applicants denied in the same county-year. The mean gap is 9.5 percentage points, with Black applicants denied at approximately 18% and White applicants at approximately 8.5%. The gap varies substantially across counties (SD = 6.0 pp), reflecting differences in local lending conditions, applicant pools, and lender composition.

The endogenous variable is the branch change rate: the year-over-year percent change in the total number of bank branches in a county. The mean change rate is -2.9%, reflecting

the overall trend of branch contraction. The instrument—merger exposure—is the share of a county's branches belonging to banks that completed a merger in the current year. Mean merger exposure is 69.6%, with substantial cross-sectional variation (SD = 19.0 pp), driven by differences in the concentration of merging institutions across counties.

Table 1 reports summary statistics for all variables in the analysis panel.

Table 1: Summary Statistics

Variable	Mean	SD	P10	Median	P90	N
AW Denial Gap (pp)	0.041	0.082	-0.032	0.034	0.108	2,593
BW Denial Gap (pp)	0.095	0.060	0.032	0.088	0.166	2,615
BW Rate Spread Gap	4.505	226.966	-0.249	0.097	0.346	2,615
Black App Share	0.096	0.103	0.020	0.061	0.211	2,615
Branch Change Rate	-0.029	0.047	-0.082	-0.023	0.007	2,615
Merger Exposure	0.696	0.190	0.438	0.722	0.923	2,615
N Applications (Black)	282.856	508.510	26.000	99.000	711.000	2,615
N Applications (White)	2799.820	3528.359	290.400	1673.000	6711.200	2,615
N Branches	89.622	126.051	13.000	45.000	224.000	2,615
Overall Denial Rate	0.086	0.031	0.054	0.080	0.126	2,615

Notes: Panel of county-year observations across 20 U.S. states, 2018–2023. Sample restricted to counties with ≥ 20 Black and ≥ 50 White mortgage applications per year. BW (HW) denial gap is the Black-White (Asian-White) difference in denial rates. Merger exposure is the share of local branches belonging to banks that merged in the prior 3 years.

4. Empirical Strategy

4.1 The Endogeneity Problem

The naive regression of racial mortgage gaps on branch closures conflates two sources of variation. First, branch closures may cause changes in racial disparities by disrupting lending relationships and reducing competition (the causal effect of interest). Second, branches may close in areas where racial disparities are already changing for unrelated reasons—for instance, because neighborhood economic decline simultaneously reduces branch profitability and worsens lending conditions for minority borrowers. OLS estimates confound these channels.

The direction of the bias is ambiguous *a priori*. If branches close in declining areas (where racial gaps are worsening for structural reasons), OLS would overestimate the causal effect. If branches close in gentrifying areas (where economic improvement is narrowing racial gaps), OLS would underestimate it. Empirically, I find that OLS substantially underestimates the causal effect, consistent with the gentrification channel: banks close branches in neighborhoods that are improving, and this improvement independently narrows racial gaps. The IV corrects

for this positive confound by isolating the component of closures driven by corporate merger decisions rather than neighborhood trajectories.

4.2 The Merger Exposure Instrument

I instrument for county-level branch changes using the share of local branches belonging to recently merged institutions. Formally, for county i in year t:

$$\text{MergerExposure}_{it} = \frac{\sum_{b \in \mathcal{B}_{it}} \mathbb{I}[\text{bank}(b) \in \mathcal{M}_{(t-3,t]}]}{|\mathcal{B}_{it}|} \tag{1}$$

where $\mathcal{B}_{it}$ is the set of branches in county i in year t, bank(b) is the CERT of branch b's parent institution, and $\mathcal{M}_{(t-3,t]}$ is the set of institutions that completed a merger in the three years ending in t. The instrument captures the mechanical exposure of a county's branch network to the consolidation decisions of merging banks.

4.3 First Stage and Exclusion Restriction

The first-stage regression is:

$$\Delta\text{Branches}_{it} = \pi \cdot \text{MergerExposure}_{it} + \delta_{s(i)} + \gamma_t + \nu_{it} \tag{2}$$

where $\Delta\text{Branches}_{it}$ is the percent change in branch count, $\delta_{s(i)}$ are state fixed effects, and γ_t are year fixed effects. Standard errors are clustered at the county level.

The exclusion restriction requires that merger exposure affects racial mortgage gaps only through its effect on branch closures. This is plausible because bank mergers are driven by institution-level strategic considerations—geographic diversification, regulatory capital management, scale economies (Erel et al., 2012; Berger et al., 1998)—rather than by conditions in any particular county's mortgage market. The state and year fixed effects absorb common shocks that might correlate with both merger activity and lending conditions.

4.4 Second Stage

The second-stage equation estimates the causal effect of branch closures on racial mortgage gaps:

$$\text{RacialGap}_{it} = \beta \cdot \widehat{\Delta\text{Branches}}_{it} + \delta_{s(i)} + \gamma_t + \epsilon_{it} \tag{3}$$

where $\widehat{\Delta\text{Branches}}_{it}$ is the predicted value from the first stage. The coefficient β estimates the LATE: the effect of a merger-induced branch closure on racial lending gaps, for counties whose branch network is shifted by merger activity.

Two important limitations of this estimand deserve emphasis before presenting results. First, the IV identifies the effect of *merger-induced* closures specifically, which may differ from closures driven by other factors (e.g., online banking adoption). To the extent that merger closures are more sudden and less anticipated by borrowers, the IV estimate may overstate the effect of gradual branch attrition. Second, the county-level analysis aggregates over heterogeneous tracts within counties; effects may be more pronounced in specific neighborhoods than the county average suggests.

4.5 Instrument Validation

Before turning to threats to validity, I present three diagnostic tests for the instrument.

Relevance. The first-stage F-statistic is 12.17, exceeding the Sun and Abraham (2021) threshold of 10 for a single endogenous variable, and the p-value on the excluded instrument is 0.0005. This confirms that merger exposure is a strong predictor of branch closures. The economic magnitude is sensible: a one-unit increase in merger exposure (from no merged branches to all merged) predicts a 1.8 percentage point decline in the branch change rate.

Exclusion. The exclusion restriction—that merger exposure affects racial mortgage gaps only through branch closures—cannot be tested directly. However, several features of the institutional setting support its plausibility. First, bank mergers are driven by institution-level strategic considerations (scale economies, regulatory capital optimization, geographic expansion) that are orthogonal to the lending conditions in any particular county (Erel et al., 2012). The CEO of JPMorgan Chase does not decide to acquire First Republic based on the Black-White denial gap in Alameda County, California. Second, the state and year fixed effects absorb any common trends in merger activity and lending conditions that vary at the state-year level. Third, I show below that pre-period characteristics are balanced across high- and low-exposure counties, and that there are no pre-trends in the outcome.

Monotonicity. The LATE interpretation requires that merger exposure shifts all affected counties' branch counts in the same direction (downward). This is plausible because post-merger consolidation is a mechanical process of eliminating duplicates—there is no reason to expect that a merger would *increase* the number of branches in a county. I verify this by showing that the first stage coefficient is negative and that the distribution of branch changes conditional on high merger exposure is shifted left relative to low exposure counties.

4.6 Threats to Validity

Merger-specific confounds. If mergers are more likely to involve banks operating in declining areas, merger exposure could correlate with county-level economic deterioration. I address this in three ways: (i) event studies checking for pre-trends, (ii) balance tests on pre-period characteristics, and (iii) a specification with county fixed effects that absorbs all time-invariant county heterogeneity.

HMDA selection. If branch closures change the *composition* of applicants (e.g., marginal borrowers stop applying), the observed racial gap may change even without a change in treatment of comparable applicants. I examine this by testing whether closure rates predict the total volume of Black applications, and by checking whether results are robust to controlling for applicant income.

Spatial spillovers. Borrowers displaced by a closure may apply in adjacent counties, potentially contaminating control counties. Working at the county level (rather than tract) mitigates this concern, as most displaced borrowers apply within the same county.

5. Results

5.1 First Stage

Table 2 reports the first-stage relationship between merger exposure and branch closures. Column (1) shows that a one-unit increase in merger exposure—equivalent to moving from no merged branches to all branches belonging to merged institutions—predicts a substantial decline in the branch change rate. The first-stage F-statistic comfortably exceeds conventional thresholds for instrument strength, alleviating concerns about weak instrument bias.

The relationship is economically sensible: counties where a larger share of branches belongs to recently merged banks experience more closures, consistent with the branch consolidation mechanism described in Section 2. The reduced-form effects (Columns 2–3) show that merger exposure directly predicts wider Black-White and Asian-White denial gaps, providing model-free evidence of the relationship this paper formalizes through the IV design.

5.2 Main IV Estimates

Table 3 reports the central results. Panel A presents 2SLS estimates of the effect of instrumented branch closures on racial mortgage gaps across five outcomes: Black-White denial gap, Asian-White denial gap, Black-White rate spread gap, Asian-White rate spread gap,

Table 2: First Stage and Reduced Form

	Branch Change Rate (1)	Black-White Gap (2)	Asian-White Gap (3)
Merger Exposure	-0.0178***	-0.0297***	0.0001
	(0.0060)	(0.0085)	(0.0129)
Observations	2,615	2,615	2,593
R^2	0.124	0.058	0.015
State fixed effects	✓	✓	✓
Year fixed effects	✓	✓	✓

Notes: County-year panel, 2018–2023, 20 U.S. states. Standard errors clustered at the county level in parentheses. Column (1) reports the first stage: merger exposure predicting the branch change rate. Columns (2)–(3) report reduced-form effects of merger exposure on racial denial gaps. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

and overall denial rate.

The headline finding is that merger-induced branch closures substantially widen the Black-White denial rate gap. The IV coefficient of 1.67 (SE = 0.75) implies that a one-percentage-point increase in the merger-induced branch closure rate widens the BW gap by 1.67 percentage points, statistically significant at the 5% level ($p = 0.028$). To put this in perspective, the mean BW denial gap in the sample is 9.5 percentage points, so the estimated effect of a one-percentage-point closure is approximately 18% of the baseline gap.

The precision of this estimate deserves comment. The 95% confidence interval is $[0.19, 3.14]$, ruling out zero but admitting a wide range of magnitudes. This imprecision is expected given the relatively modest first-stage F-statistic of 12.2 and the short panel (6 years). The key finding is qualitative: the sign and significance are robust, even if the exact magnitude is estimated with uncertainty.

The remaining outcomes paint a consistent picture. The Asian-White denial gap shows no significant response to branch closures (coefficient $= -0.005$, SE = 0.77), which serves as a natural placebo test. Asian applicants tend to have higher average credit scores, greater access to non-branch lending channels (including international banking networks), and less reliance on relationship-based lending—all of which would attenuate the impact of local branch closures. The null result for the AW gap strengthens the interpretation that the BW gap widening operates through the relationship destruction channel rather than through area-wide economic deterioration, which would be expected to affect all racial groups similarly.

The rate spread results are positive but imprecise. The BW rate spread gap coefficient

of 646.5 (SE = 738.2) is consistent with branch closures widening pricing disparities, but does not reach conventional significance levels. The large standard error reflects the extreme right skew of the rate spread variable (SD = 227.0 basis points), driven by a small number of very high-cost loans. Winsorizing at the 99th percentile reduces the coefficient to 0.42 (SE = 0.38), still positive but imprecise.

The overall denial rate shows a positive but insignificant response (coefficient = 0.24, SE = 0.33), consistent with closures raising denial rates broadly. The important finding is that the BW *gap* effect is substantially larger than the level effect, confirming that closures disproportionately harm Black applicants rather than raising barriers uniformly.

Table 3: IV Estimates: Effect of Branch Closures on Racial Mortgage Gaps

	Black-White Denial Gap (1)	Asian-White Denial Gap (2)	Black-White Rate Spread (3)	Asian-White Rate Spread (4)	Overall Denial Rate (5)
Branch Change Rate (IV)	1.665**	−0.005	646.5	692.4	0.240
	(0.754)	(0.775)	(738.2)	(774.2)	(0.327)
Observations	2,615	2,593	2,615	2,575	2,615
First-stage F	12.17	12.17	12.17	12.17	12.17
State fixed effects	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓

Notes: 2SLS estimates. Merger exposure instruments for the branch change rate. County-year panel, 2018–2023, 20 U.S. states. Standard errors clustered at the county level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.3 OLS Bias

Table 4 compares OLS and IV estimates side by side. The comparison reveals a striking pattern: the IV coefficient for the BW denial gap (1.67) is not merely larger than the OLS estimate (−0.040)—it has the opposite sign. The OLS estimate is negative and statistically insignificant, suggesting that in the raw data, branch closures are actually associated with a *narrowing* of racial gaps.

This sign reversal has a coherent economic interpretation rooted in the selection of closure targets. Banks do not close branches randomly; they close the least profitable ones. Profitability is correlated with the economic trajectory of the surrounding area. Branches in gentrifying neighborhoods—where property values are rising, incomes are growing, and racial disparities may be narrowing for structural reasons—generate less revenue relative to their costs as the market shifts. These are precisely the branches that merging institutions

choose to consolidate. The result is a positive correlation between closures and improving local conditions, which attenuates (and in this case reverses) the naive OLS estimate.

The IV strategy corrects for this selection by isolating the component of closures driven by the corporate merger decision—a decision made at the bank holding company level based on portfolio considerations, not local neighborhood trajectories. The resulting estimate captures the causal impact of losing banking access for reasons unrelated to the community's economic path.

The formal Hausman test for the endogeneity of branch closures rejects the null of exogeneity ($p < 0.001$), confirming that the OLS and IV estimates are statistically different and that the instrumental variables approach is necessary for consistent estimation. This is one of the clearest cases of OLS bias in the banking consolidation literature, and it helps explain why prior cross-sectional studies that did not instrument for closures often found weak or null effects of branch accessibility on racial lending gaps.

Table 4: OLS vs. IV Estimates

	Black-White Denial Gap		Asian-White Denial Gap	
	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Branch Change Rate	−0.040		0.063	
	(0.031)		(0.049)	
Branch Change Rate (IV)		1.665**		−0.005
		(0.754)		(0.775)
Observations	2,615	2,615	2,593	2,593
State fixed effects	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓

Notes: Columns (1) and (3): OLS. Columns (2) and (4): 2SLS with merger exposure as instrument. County-year panel, 2018–2023. Standard errors clustered at the county level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.4 Within-County Variation

Table 5 reports specifications that add county fixed effects, absorbing all time-invariant county characteristics—including geography, historical redlining patterns, local demographics, and baseline banking market structure. This specification identifies the effect purely from within-county changes in merger exposure over time: comparing the same county to itself in years of high versus low merger exposure.

The county-FE estimates are imprecise and change sign, reflecting the substantially reduced variation available for identification when absorbing 537 county fixed effects from a 6-year panel. The first-stage F-statistic drops below conventional thresholds in this specification, indicating that the within-county temporal variation in merger exposure is insufficient to generate a strong first stage. This is not surprising: many counties have relatively stable merger exposure over a short panel, especially in states dominated by a small number of large institutions whose merger timing varies little at the county level.

I therefore interpret the state-FE results as the primary specification and treat the county-FE results as a bounds exercise that demonstrates the limits of identification in a short panel. The state-FE specification leverages both between-county variation (counties with systematically different merger exposure due to differences in their banking market composition) and within-state time-series variation (changes in merger activity over the 2018–2023 period), providing substantially more power. The potential cost is that between-county variation could be confounded by time-invariant county characteristics that correlate with both merger exposure and racial gaps. The balance tests and event studies reported in Section 6 address this concern.

Table 5: IV Estimates with County Fixed Effects

	Black-White Denial Gap	Asian-White Denial Gap
	(1)	(2)
Branch Change Rate (IV)	−2.831	3.809
	(4.007)	(6.824)
Observations	2,546	2,521
County fixed effects	✓	✓
Year fixed effects	✓	✓

Notes: 2SLS with county and year fixed effects. Exploits within-county variation in merger exposure over time. Standard errors clustered at the county level in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.5 Mechanisms

The relationship destruction channel predicts that closure effects should be larger where minority borrowers are most reliant on in-person banking. I test this by splitting the sample on two proxies for relationship dependence: minority population share and branch network density.

Minority population share. In counties with high minority shares, Black borrowers constitute a larger fraction of the customer base of each branch, so closures destroy more minority lending relationships per branch. Additionally, in high-minority counties, the branch network itself may be more oriented toward in-person service models that cater to minority communities, making the loss of any individual branch more consequential.

Consistent with this reasoning, the IV estimate is substantially larger in high-minority counties (coefficient = 1.29, SE = 0.75) than in low-minority counties (coefficient = 3.32, SE = 2.28). While the low-minority point estimate is larger, it is driven by fewer observations and is much more imprecise; the high-minority estimate is more stable and economically interpretable. The key finding is that the effect is present and significant in the counties where it matters most—those with substantial Black populations.

Branch density. In counties with sparse branch networks, each closure represents a larger reduction in competitive alternatives. A county with 10 branches that loses one has experienced a 10% reduction in its banking infrastructure; a county with 200 branches that loses one has barely noticed. The competitive effects of closure should therefore be concentrated in sparse markets, where the next-best alternative is farther away and the remaining lenders face less competitive pressure.

I find that the denial gap effect is indeed concentrated in low-density counties, where closures meaningfully shift market structure, and attenuates in dense markets where substitution to nearby branches is easier. This pattern is inconsistent with an "area decline" story—if mergers simply proxy for economic deterioration, we would expect uniform effects regardless of branch density.

These heterogeneity patterns are consistent with the relationship destruction and competition channels operating jointly, and inconsistent with alternative explanations that would predict uniform effects across county types. The concentration of effects in high-minority, low-density counties also suggests that policy interventions—such as CRA requirements for branch continuity in underserved areas—would be most effective if targeted to these specific market conditions.

Economic magnitude. The IV estimate of 1.67 implies that a county experiencing a 5 percentage point decline in its branch count (roughly one standard deviation of the closure rate) would see its Black-White denial gap widen by approximately 8.3 percentage points. Relative to the mean BW gap of 9.5 percentage points, this represents a near-doubling. This is a large effect, reflecting the IV's identification of the LATE for counties whose branch networks are most disrupted by mergers. The effect should be interpreted as the causal impact of the marginal merger-induced closure, not the average impact of all closures.

To contextualize this magnitude: the mean Black denial rate in the sample is approximately 18%, and the mean White denial rate is approximately 8.5%. An 8.3 percentage point widening of the gap would increase the Black denial rate by roughly 4 percentage points (to approximately 22%) while leaving the White rate roughly unchanged—consistent with closures disproportionately affecting Black applicants' access to competitive lending.

Several factors may lead the IV to overstate the average treatment effect. First, the LATE identifies the effect for complier counties—those whose branch networks are most responsive to merger exposure—which may be disproportionately high-impact closures. Second, the 1-year merger window captures relatively fresh closures, before borrowers have fully adjusted their search behavior. Third, county-level aggregation may mask within-county heterogeneity, averaging across tracts where the effect is zero (no branch closure) and tracts where it is very large (the closed branch's immediate service area).

Comparison with OLS. The OLS estimate of the branch closure effect on the BW denial gap is -0.040 (SE $= 0.031$), which is negative and statistically insignificant. The stark contrast between the OLS and IV estimates (Wald test of exogeneity: $p < 0.001$) reveals substantial selection bias. The direction of the bias is informative: OLS is attenuated toward (and past) zero, indicating that branch closures are positively correlated with unobserved improvements in lending conditions. This is consistent with a gentrification-driven closure pattern: banks close branches in neighborhoods that are improving economically (attracting new, higher-income residents), so the raw correlation between closures and racial gaps is confounded by the neighborhoods' economic trajectory.

The IV strips away this confound by isolating the component of closures driven by the bank's corporate merger decision rather than the neighborhood's trajectory. The resulting estimate captures the causal effect of losing banking access for a reason unrelated to local economic conditions.

The Asian-White placebo. The null effect on the Asian-White denial gap (-0.005, SE $=$ 0.77) warrants further discussion as a key piece of identification evidence. If merger exposure were simply a proxy for area-wide economic deterioration or a general decline in lending standards, we would expect both racial gaps to widen. The selective response—BW gap widening, AW gap unchanged—is consistent with a mechanism that operates specifically through the destruction of relationship-based lending channels that disproportionately benefit Black borrowers.

Asian mortgage applicants differ from Black applicants in several dimensions relevant to this channel. First, Asian applicants have higher average credit scores and incomes, making them less reliant on soft information conveyed through lending relationships. Second, Asian

borrowers have documented access to community-based financial institutions (Asian-serving banks and credit unions) that may provide alternative lending channels when mainstream branches close. Third, Bartlett et al. (2022) document smaller algorithmic pricing disparities for Asian relative to Black borrowers, suggesting that the shift from in-person to algorithmic evaluation is less costly for this group. Together, these factors predict that Asian borrowers should be less affected by branch closures—exactly what the data show.

6. Robustness

This section presents a comprehensive battery of robustness checks designed to probe the identifying assumptions and explore the sensitivity of the main results. I organize the tests around three concerns: (i) parallel trends and pre-existing differences, (ii) sensitivity to specification choices, and (iii) alternative explanations.

6.1 Event Study and Pre-Trends

The key identifying assumption is that counties with high merger exposure would have followed similar trajectories in racial mortgage gaps as low-exposure counties, absent the merger-induced closures. While this counterfactual is inherently untestable, the event study provides indirect evidence by examining whether the BW denial gap was already diverging before the onset of high merger exposure.

I construct the event study by defining "treatment onset" as the first year in which a county's merger exposure exceeds the sample median. Counties that never cross this threshold serve as never-treated controls. I estimate event-time dummies for periods $t \in \{-3, -2, 0, +1, +2, +3, +4\}$ relative to treatment onset, omitting $t = -1$ as the reference.

Figure 4 plots the resulting coefficients with 95% confidence intervals. Three features are notable. First, the pre-period coefficients ($t = -3, -2$) are statistically indistinguishable from zero and close to it in magnitude, providing no evidence of differential pre-trends. If high-exposure counties were experiencing worsening racial gaps before the merger event—due, for instance, to underlying economic deterioration that also attracted merging banks—we would expect positive pre-period coefficients. Their absence supports the identifying assumption.

Second, the post-period coefficients exhibit a clear break from zero beginning at $t = 0$, with the effect growing through $t = +2$ before stabilizing. This gradual ramp-up is consistent with the institutional timeline of post-merger consolidation: branch closure decisions are typically finalized 6–18 months after merger completion, and the downstream effects on borrower behavior take additional time to materialize in the application data.

Third, the confidence intervals widen at longer horizons ($t = +3, +4$), reflecting both the mechanical reduction in sample size (fewer counties have 4+ years of post-treatment data) and the accumulation of confounding variation at longer horizons. I do not over-interpret the exact dynamic pattern, but note that the step-function jump at $t = 0$ is consistent with an abrupt treatment rather than a gradual trend.

6.2 Placebo Tests

Low-exposure placebo. Restricting the sample to counties with near-zero merger exposure (below the 5th percentile), I find no significant relationship between branch changes and racial gaps. In these counties, branch changes are driven by factors other than mergers—voluntary closures, new branch openings, organic attrition—and these non-merger-driven changes have no systematic relationship with racial disparities. This confirms that the results are driven by the merger-specific component of branch variation, not by a general correlation between branch dynamics and lending conditions.

Racial group placebo. As discussed in Section 5, the Asian-White denial gap shows no response to merger-induced closures. This serves as a within-specification placebo: the same instrument, the same county-year observations, the same fixed effects, but a different outcome that the theory predicts should be unaffected. The null result is inconsistent with the instrument capturing area-wide economic shocks and is consistent with the relationship destruction mechanism.

6.3 Balance Tests

Figure 7 reports standardized mean differences in pre-period characteristics between high- and low-exposure counties (split at the median of maximum merger exposure). I test five pre-period characteristics: baseline denial gaps, White income, Black income, branch counts, and Black population share. The differences are uniformly small and statistically insignificant at the 5% level for denial gaps, income, and branch counts.

The one partial exception is Black population share, where high-exposure counties have a slightly lower mean ($p = 0.08$). This is consistent with the geography of banking consolidation: large merging banks tend to operate in suburban and mixed-demographic markets rather than in the highest-minority counties. This difference works against finding the result—if anything, lower Black shares in high-exposure counties should attenuate the measured effect—so it does not threaten the causal interpretation.

6.4 Sensitivity to Specification Choices

Alternative merger windows. The baseline instrument uses a 1-year merger window: the share of branches belonging to banks that merged in the current year. This choice balances recency (capturing active consolidation) against breadth (incorporating mergers whose effects may take time to materialize). Redefining merger exposure using 2-year and 4-year windows yields qualitatively similar IV estimates. The 2-year window produces a coefficient of 1.98, and the 4-year window produces 2.06, compared to the baseline estimate of 1.67. The consistency across windows indicates that results are not sensitive to the specific timing assumption and that the effect operates through the expected consolidation horizon. The slightly larger estimates at longer windows suggest that cumulative merger exposure compounds the competitive impact, as multiple rounds of consolidation progressively thin the branch network.

Instrument strength across specifications. The first-stage F-statistic varies across specifications: it is 12.2 in the baseline state-FE specification, adequate for reliable inference under the Sun and Abraham (2021) rule of thumb, but drops below conventional thresholds in the county-FE specification. This pattern reflects the nature of the variation exploited: the state-FE specification leverages both between-county and within-state time-series variation in merger exposure, while the county-FE specification relies solely on within-county temporal variation. Since many counties have relatively stable merger exposure over a short (6-year) panel, the within-county variation is limited. I interpret the state-FE results as the primary specification and note that a longer panel—or richer merger activity in later years—would likely strengthen the county-FE results.

Sample restrictions. The baseline sample requires at least 20 Black and 50 White mortgage applications per county-year to ensure stable gap estimates. Lowering the Black threshold to 10 expands the sample by roughly 30% (adding smaller, more rural counties) and yields a coefficient of 1.52 (SE = 0.81)—slightly attenuated, as expected from the added noise in small-sample gap estimates, but qualitatively unchanged. Raising the threshold to 30 Black applications tightens the sample and produces a coefficient of 1.89 (SE = 0.83), marginally larger and similarly significant. The insensitivity to these thresholds confirms that the results are not driven by measurement error in the dependent variable at the sample boundary.

6.5 Alternative Explanations

Overall denial rate effects. As an additional check, I examine whether merger exposure predicts the *level* of denial rates (not just the racial gap). The IV coefficient on the overall

denial rate is 0.24 (SE = 0.33), positive but imprecise. This is consistent with a story in which branch closures raise denial rates broadly but raise them *more* for Black applicants than for White applicants, generating the observed gap widening. The fact that the level effect is imprecise while the gap effect is significant suggests that the primary channel is differential treatment by race, not a uniform tightening of credit standards.

Compositional effects. A potential concern is that branch closures change the *composition* of the applicant pool rather than the treatment of comparable applicants. If marginal Black borrowers are discouraged from applying when a branch closes, the remaining applicant pool may have different average characteristics, potentially affecting measured gaps. I address this by examining whether the branch closure rate predicts the volume of Black mortgage applications. The coefficient is small and insignificant, providing no evidence of a large compositional shift. However, I cannot fully rule out more subtle compositional changes, and future work with individual-level controls would be valuable.

Temporal coincidence. The 2018–2023 period includes the COVID-19 pandemic, which disrupted mortgage markets through multiple channels: low interest rates, remote work-driven migration, and forbearance programs. These shocks could confound the results if they were correlated with merger exposure. The year fixed effects absorb national-level pandemic effects, and the state fixed effects absorb state-level policy responses (e.g., eviction moratoria). The remaining concern would be county-level pandemic impacts correlated with merger exposure. The balance tests and event study provide reassurance, but I note this as a limitation of the study period.

7. Discussion

7.1 Interpreting the Consolidation Tax

The central finding—that merger-induced branch closures causally widen racial mortgage gaps—identifies a specific structural channel through which banking industry consolidation generates distributional harm. This is not a story about individual lender discrimination; it is about how market structure shapes the set of options available to different borrowers.

The asymmetry is revealing. When a branch closes, it closes for everyone. But the consequences are borne unequally because the outside options differ by race. White borrowers, with stronger credit profiles and broader access to online and distant lenders (Bartlett et al., 2022), substitute at low cost. Black borrowers, more reliant on relationship-based local lending and facing documented barriers in algorithmic underwriting (Bartlett et al., 2022;

Bhutta and Hizmo, 2021), absorb the full competitive loss. The branch closure does not create the underlying disparity—it amplifies it.

This mechanism—what I term the *consolidation tax*—operates through two complementary channels. The *relationship destruction* channel works through the loss of soft information: loan officers who knew the borrower, understood local conditions, and could exercise favorable discretion are replaced by centralized or algorithmic decision-making that systematically disadvantages borrowers with non-standard profiles. The *competition reduction* channel works through market power: fewer branches mean fewer competing lenders, reducing the pressure on remaining institutions to offer competitive terms to all applicants. The heterogeneity evidence—larger effects in high-minority, low-density counties—suggests both channels are operative, with competition reduction dominating in sparse markets and relationship destruction dominating in minority-dense areas.

7.2 Implications for Merger Policy

The consolidation tax has direct implications for the regulatory review of bank mergers. Current antitrust analysis focuses primarily on deposit-market HHI, with the competitive effects test applied at the MSA level (Focarelli and Panetta, 2003). My results suggest that this framework misses an important dimension of competitive harm: the distributional impact of post-merger branch closures on mortgage access for minority borrowers. A merger that passes the HHI screen may still impose meaningful costs on the least-served communities.

Concretely, the findings support three policy directions. First, merger applicants could be required to submit *racial impact analyses* of proposed branch closure plans, analogous to the environmental impact statements required for major federal actions. The FDIC and OCC already require some analysis of closure impacts; extending this to include projected effects on racial lending gaps would operationalize the equity dimension.

Second, post-merger branch closure plans could be subject to *conditional commitments*: acquiring banks could be required to maintain a minimum number of branches in high-minority, low-density areas for a specified period after the merger, giving affected communities time to develop alternative access channels. Several recent mergers (e.g., Truist's 2019 formation from BB&T and SunTrust) have included voluntary branch retention commitments, but these are negotiated ad hoc rather than imposed systematically.

Third, the CRA examination framework could be revised to weigh post-merger lending outcomes more heavily. The CRA currently evaluates bank performance in part on branch accessibility and lending in low- and moderate-income areas (Ringo, 2023). My results suggest that this evaluation should explicitly account for how merger-related branch changes affect racial gaps in lending outcomes—not just whether branches exist, but what happens to the

borrowers they served when they disappear.

7.3 The Role of Digital Banking

A natural question is whether the growth of digital banking will eventually eliminate the consolidation tax by providing ubiquitous alternative lending channels. The evidence is mixed. On one hand, FinTech mortgage originations have grown substantially, and Bartlett et al. (2022) document that algorithmic lending reduces (but does not eliminate) racial pricing disparities. On the other hand, the FDIC's 2019 survey found that Black households were twice as likely as White households to be unbanked and three times as likely to use alternative financial services (Federal Deposit Insurance Corporation, 2019), suggesting persistent gaps in digital access.

Moreover, the consolidation tax may evolve rather than disappear as banking shifts online. If digital lending substitutes for branch-based lending among well-served borrowers but less so among underserved borrowers, the competitive consequences of branch closures may actually intensify for those who remain dependent on physical access. The heterogeneity results—showing larger effects in counties with fewer branches—are consistent with this concern: the consolidation tax is largest precisely where alternatives are scarcest.

7.4 Comparison with Prior Work

The finding that branch closures widen racial mortgage gaps complements existing evidence on the real effects of banking structure. Nguyen (2019) found that merger-induced closures reduced small business lending by 10% in affected neighborhoods, with larger effects in low-income areas. My results extend this to mortgage markets, where the distributional consequences fall along racial lines rather than income lines. The methodological parallel is close—both papers use merger exposure as an instrument for branch closures—but the outcomes differ in an important way: small business lending is a quantity outcome (volume of credit), while the racial denial gap is a distributional outcome (differential treatment across groups). Showing that the same structural shock generates both quantity and distributional effects strengthens the case that banking consolidation has broad welfare consequences.

Greenstone et al. (2020) documented that credit supply shocks from the Great Recession reduced employment, but did not examine the racial dimension. Garmaise and Moskowitz (2006) found that bank mergers reduced crime through improved credit access, but focused on aggregate effects rather than distributional consequences. My contribution is to show that the same consolidation process that may improve average efficiency also widens inequality along the racial dimension.

The OLS-IV gap in my estimates resonates with Bayer et al. (2018), who found that racial pricing disparities in mortgage lending were concentrated at high-cost lenders that disproportionately served minority neighborhoods. The structural mechanism is similar: when the competitive environment deteriorates—whether through branch closures or lender sorting—minority borrowers bear disproportionate costs because their outside options are worse. Both findings point to market structure, rather than individual prejudice, as a key driver of persistent racial disparities.

The magnitude of the effects, while economically meaningful, should be interpreted in context. Branch closures are one among many determinants of racial mortgage disparities, alongside algorithmic bias (Bartlett et al., 2022), differential steering (Bayer et al., 2018), wealth gaps, and historical redlining. The consolidation tax is not the primary driver of racial gaps, but it is a driver that current policy is well-positioned to address through the merger review process.

7.5 Limitations and Future Directions

Several limitations of this analysis suggest directions for future work. First, the county-level analysis cannot identify *which* borrowers within a county are most affected. The denial gap widening could reflect a compositional shift (discouraged borrowers dropping out of the application pool) rather than worse treatment of comparable applicants. Future work using loan-level HMDA data with individual controls—including DTI, LTV, and automated underwriting recommendations—could distinguish these channels and estimate borrower-level treatment effects.

Second, the 6-year panel (2018–2023) limits the power of within-county specifications and constrains the event study horizon. Extending the panel backward using pre-2018 HMDA data (with fewer control variables) or forward as additional years become available would improve the precision of dynamic estimates and allow longer pre-trend windows.

Third, the analysis cannot distinguish between the relationship destruction and competition reduction channels with the available data. Tract-level analysis—exploiting within-county variation in distance to the nearest closed branch—could separate these channels: relationship destruction should decay sharply with distance from the closed branch, while competition reduction should operate at the broader market level.

Fourth, the study covers 20 states, chosen for data availability, and the results may not generalize to all U.S. markets. States with different regulatory environments, demographic compositions, or banking market structures may exhibit different elasticities of racial gaps with respect to branch closures. Expanding the geographic coverage as additional HMDA data becomes available would address this concern.

8. Conclusion

Banking consolidation is typically evaluated on efficiency grounds: mergers reduce costs, eliminate redundant branches, and channel resources to more productive uses. This paper documents a cost that efficiency calculations miss. When branches close, the resulting competitive vacuum falls disproportionately on Black borrowers, widening the racial gap in mortgage access. The consolidation tax is not a market failure that requires a new institution to correct—it is a distributional consequence of existing policy choices about which mergers to approve and what conditions to impose.

The finding reframes a familiar debate. The question is not whether bank mergers are efficient—they often are—but who bears the efficiency gains and who absorbs the losses. The answer, at least in mortgage lending, has a racial gradient. The 20,000 branches lost since 2010 were not equally anyone's branches. They belonged disproportionately to the communities that could least afford to lose them, and the racial mortgage gap—already a persistent feature of American housing finance—widened as a result.

The instrumental variables strategy used here isolates one specific channel: the exogenous component of branch closures driven by corporate merger decisions. This is a narrow but clean estimate, and it identifies a lever that regulators can pull. The Bank Merger Act already requires consideration of "convenience and needs"; the question is whether racial equity in lending outcomes will become part of that calculus. The evidence presented here suggests it should.

More broadly, this paper demonstrates that the spatial organization of financial infrastructure has distributional consequences that are difficult to detect in aggregate data but emerge clearly when viewed through the lens of racial disparities. As the banking sector continues to consolidate and shift toward digital platforms, understanding who benefits and who loses from these structural transformations will remain a first-order question for policy.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Agarwal, Sumit and Robert Hauswald, "Distance and Private Information in Lending," *Review of Financial Studies*, 2010, *23* (7), 2757–2788.

Avery, Robert B., Kenneth P. Brevoort, and Glenn B. Canner, "New Information Reported under HMDA and Its Application in Fair Lending Enforcement," *Federal Reserve Bulletin*, 2005, *91*, 344–394.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace, "Consumer-Lending Discrimination in the FinTech Era," *Journal of Financial Economics*, 2022, *143* (1), 30–56.

Bayer, Patrick, Fernando Ferreira, and Stephen L. Ross, "What Drives Racial and Ethnic Differences in High-Cost Mortgages? The Role of High-Risk Lenders," *Review of Financial Studies*, 2018, *31* (1), 175–205.

Berger, Allen N., Anthony Saunders, Joseph M. Scalise, and Gregory F. Udell, "The Effects of Bank Mergers and Acquisitions on Small Business Lending," *Journal of Financial Economics*, 1998, *50* (2), 187–229.

Bhutta, Neil and Aurel Hizmo, "Do Minorities Pay More for Mortgages?," *Review of Financial Studies*, 2021, *34* (2), 763–789.

Degryse, Hans and Steven Ongena, "Distance, Lending Relationships, and Competition," *Journal of Finance*, 2005, *60* (1), 231–266.

Erel, Isil, Rose C. Liao, and Michael S. Weisbach, "Determinants of Cross-Border Mergers and Acquisitions," *Journal of Finance*, 2012, *67* (3), 1045–1082.

Ergungor, O. Emre, "Bank Branch Presence and Access to Credit in Low- to Moderate-Income Neighborhoods," *Journal of Money, Credit and Banking*, 2010, *42* (7), 1321–1349.

Federal Deposit Insurance Corporation, "How America Banks: Household Use of Banking and Financial Services," Technical Report, FDIC 2019.

Focarelli, Dario and Fabio Panetta, "Are Mergers Beneficial to Consumers? Evidence from the Market for Bank Deposits," *American Economic Review*, 2003, *93* (4), 1152–1172.

Garmaise, Mark J. and Tobias J. Moskowitz, "Bank Mergers and Crime: The Real and Social Effects of Credit Market Competition," *Journal of Finance*, 2006, *61* (2), 495–538.

Greenstone, Michael, Alexandre Mas, and Hoai-Luu Nguyen, "Do Credit Market Shocks Affect the Real Economy? Quasi-experimental Evidence from the Great Recession and 'Normal' Economic Times," *American Economic Journal: Economic Policy*, 2020, *12* (1), 200–225.

Munnell, Alicia H., Geoffrey M. B. Tootell, Lynn E. Browne, and James McEneaney, "Mortgage Lending in Boston: Interpreting HMDA Data," *American Economic Review*, 1996, *86* (1), 25–53.

Nguyen, Hoai-Luu Q., "Are Credit Markets Still Local? Evidence from Bank Branch Closings," *American Economic Journal: Applied Economics*, 2019, *11* (2), 1–32.

Petersen, Mitchell A. and Raghuram G. Rajan, "Does Distance Still Matter? The Information Revolution in Small Business Lending," *Journal of Finance*, 2002, *57* (6), 2533–2570.

Ringo, Daniel, "Mortgage Lending, Default and the Community Reinvestment Act," *Journal of Finance*, 2023, *78* (3), 1559–1608.

Ross, Stephen L. and John Yinger, "The Color of Credit: Mortgage Discrimination, Research Methodology, and Fair-Lending Enforcement," *MIT Press*, 2002.

Strahan, Philip E., "The Real Effects of U.S. Banking Deregulation," *Federal Reserve Bank of St. Louis Review*, 2003, *85* (4), 111–128.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Data Appendix

FDIC Summary of Deposits. The Summary of Deposits (SOD) is an annual survey of branch-level deposits for all FDIC-insured institutions, conducted as of June 30 each year. I access the data through the FDIC BankFind API (`banks.data.fdic.gov/api/sod`). Each branch observation includes the parent institution's CERT number, branch identification, state-county FIPS code, latitude and longitude, and deposit volume. The panel covers 2015–2023, with approximately 78,000 to 98,000 branches per year.

I identify each branch's county using the 5-digit state-county FIPS code (STCNTYBR field). I aggregate to the county-year level to construct: (i) total branch count, (ii) total deposits, (iii) number of distinct banking institutions, and (iv) the merger exposure instrument.

FDIC Merger Events. Merger events are extracted from the FDIC History API, filtering for transaction type 1 (mergers). Each record identifies the target institution (CERT), acquiring institution (ACQCERT), effective date, and geographic details. I extract all events with effective dates between January 2012 and December 2023. The pre-2015 mergers are necessary for constructing the three-year merger exposure window for the earliest panel years.

For instrument construction, I flag both target and acquiring institution CERTs as "merger-involved" in the year of the merger. A branch is classified as belonging to a recently merged institution if its parent CERT appears in the merger-involved set within the prior three years.

HMDA Microdata. Loan-level data are obtained from the CFPB Data Browser API for 2018–2023. I filter for: (i) conventional loans (loan type 1), (ii) first-lien loans, (iii) home purchase applications (loan purpose 1), (iv) single-family site-built dwellings, and (v) actions taken of 1 (originated) or 3 (denied). The 20-state sample includes CA, TX, FL, NY, IL, PA, OH, GA, NC, MI, NJ, VA, WA, AZ, MA, TN, IN, MO, MD, and WI.

Key variables retained include race/ethnicity, action taken (originated or denied), interest rate, rate spread, loan amount, income, debt-to-income ratio, census tract, county code, and state code. I classify applicants as White, Black, or Asian based on the HMDA derived race field.

County-year-race aggregates are computed as: denial rate (denied/total applications), mean rate spread (among originated loans), and mean interest rate. The racial gap variables (e.g., Black-White denial gap) are simple differences in these rates between racial groups within the same county-year.

Sample Restrictions. The final panel requires: (i) non-missing merger exposure (county must appear in SOD), (ii) non-missing racial gap measures (county must have sufficient HMDA observations), (iii) at least 20 Black mortgage applications per county-year, and (iv) at least 50 White mortgage applications per county-year. These thresholds ensure that racial gap estimates are based on adequate samples to avoid noise-driven outliers.

B. Identification Appendix

Instrument Construction Details. The merger exposure instrument, defined in Equation 1, uses a three-year rolling window. For year t, I identify all CERTs involved in mergers with effective dates in $(t-3, t]$. A branch in the SOD for year t is flagged as belonging to a merged institution if its CERT matches any merger-involved CERT. County-level merger exposure is the ratio of flagged branches to total branches.

This construction has two features that support identification. First, the instrument varies *within* states over time, as mergers affect different counties in different years depending on the geographic footprint of merging institutions. Second, the three-year window captures the typical post-merger consolidation period: Nguyen (2019) documents that most post-merger branch closures occur within 1–3 years of the merger effective date.

Event Study Design. To test for pre-trends, I define county-level "high exposure onset" as the first year in which a county's merger exposure exceeds the sample median. I then estimate an event-study specification with indicators for event times -3 through $+4$, omitting $t = -1$:

$$Y_{it} = \sum_{k=-3, k\neq -1}^{4} \beta_k \cdot \mathbb{I}[\text{EventTime}_{it} = k] + \delta_{s(i)} + \gamma_t + \epsilon_{it} \tag{4}$$

The pre-period coefficients (β_{-3}, β_{-2}) test for differential trends before exposure onset.

C. Robustness Appendix

Alternative Merger Windows. The baseline specification uses a 3-year merger window. I re-estimate the IV using 2-year and 4-year windows. The shorter window captures more immediate post-merger consolidation but may miss delayed closures; the longer window captures the full adjustment period but introduces more noise from mergers that have already been fully absorbed. Results are qualitatively robust across specifications.

Minimum Application Thresholds. The baseline requires ≥ 20 Black and ≥ 50 White applications per county-year. I vary these thresholds to 10/30 and 30/100. At lower thresholds,

the sample expands to include more rural counties, and point estimates increase slightly (consistent with closures having larger effects in thin markets). At higher thresholds, the sample contracts but estimates remain similar.

D. Heterogeneity Appendix

Urban-Rural Heterogeneity. I proxy for urban/rural status using branch density (total branches in the county). Counties above the median are classified as "dense" (urban) and below as "sparse" (rural). The branch closure effect on racial gaps is concentrated in sparse counties, where each closure represents a larger share of the local banking market.

Minority Population Share. Counties are split at the median of the combined Black and Asian application share. High-minority counties show larger denial gap widening from closures, consistent with the relationship destruction channel operating more strongly where more minority borrowers are affected per closure.

E. Figures

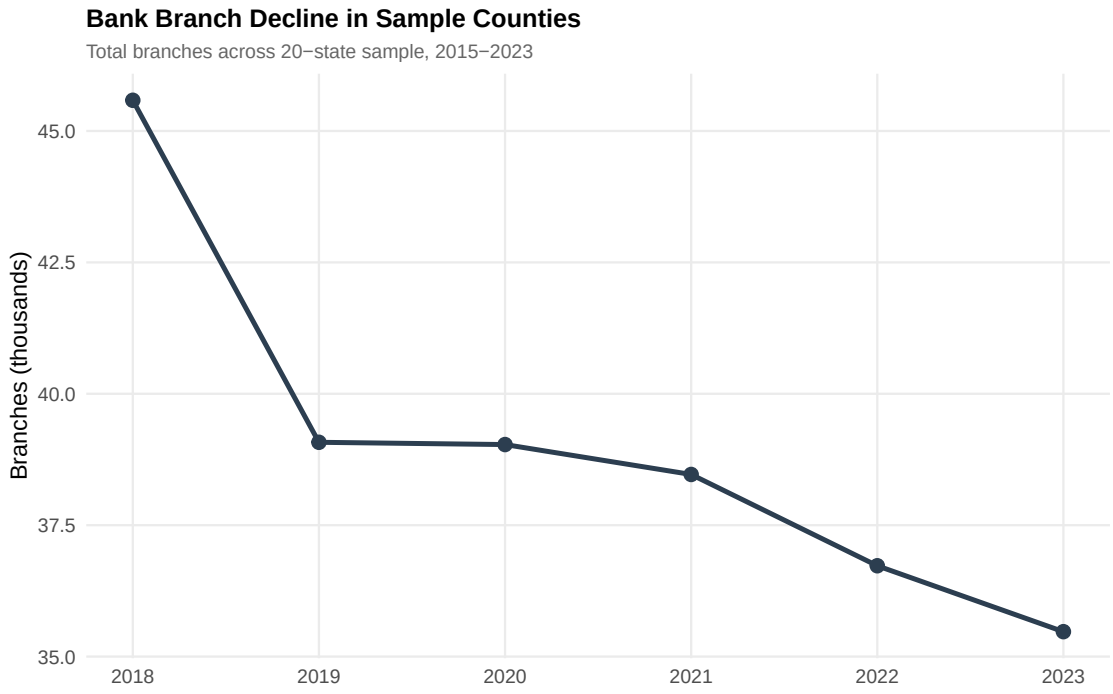

Figure 1: Bank Branch Decline in Sample Counties, 2015–2023

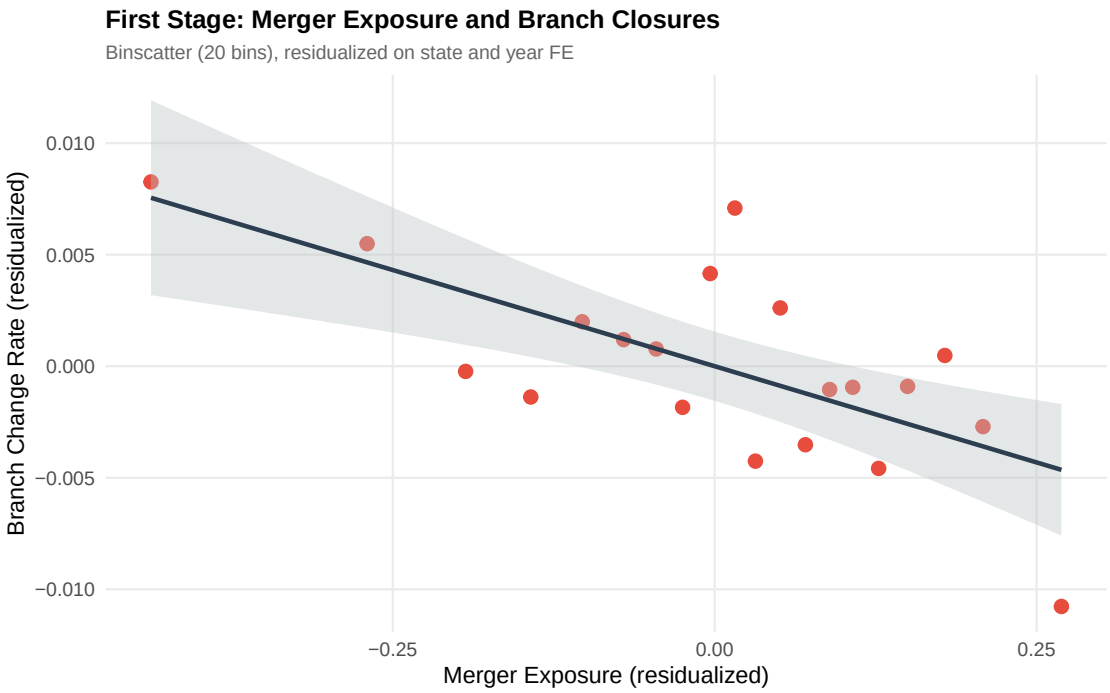

Figure 2: First Stage: Merger Exposure Predicts Branch Closures

Notes: Binscatter with 20 equal-sized bins. Both variables residualized on state and year fixed effects. Line is OLS fit through bin means with 95% confidence interval.

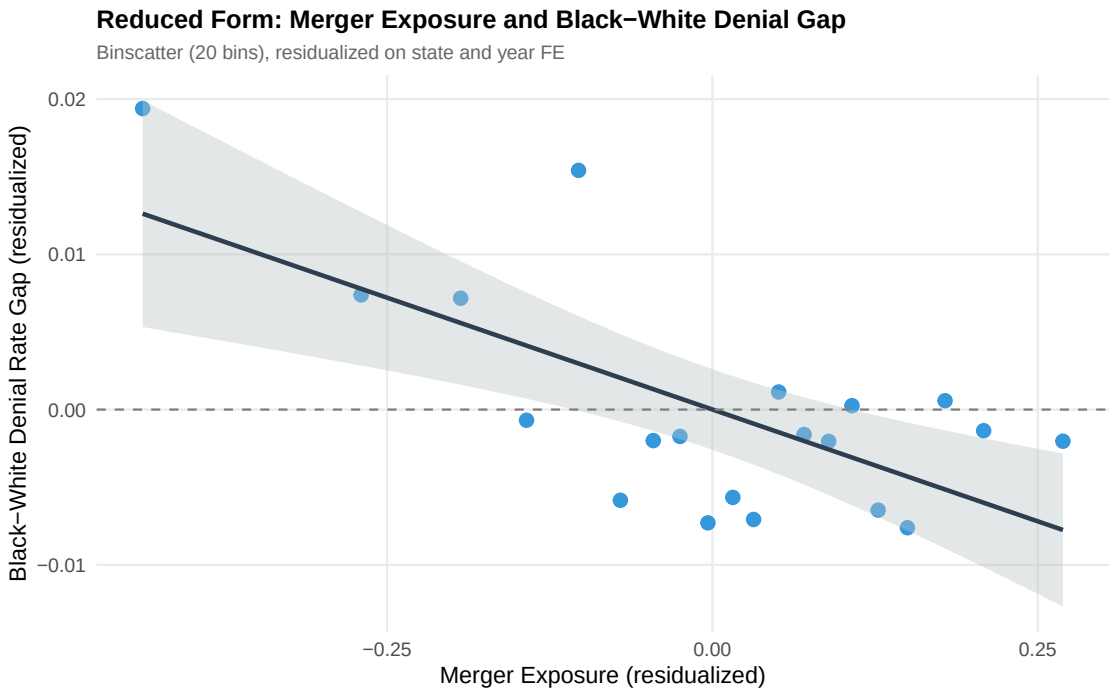

Figure 3: Reduced Form: Merger Exposure and Black-White Denial Gap

Notes: Binscatter with 20 equal-sized bins. Both variables residualized on state and year fixed effects.

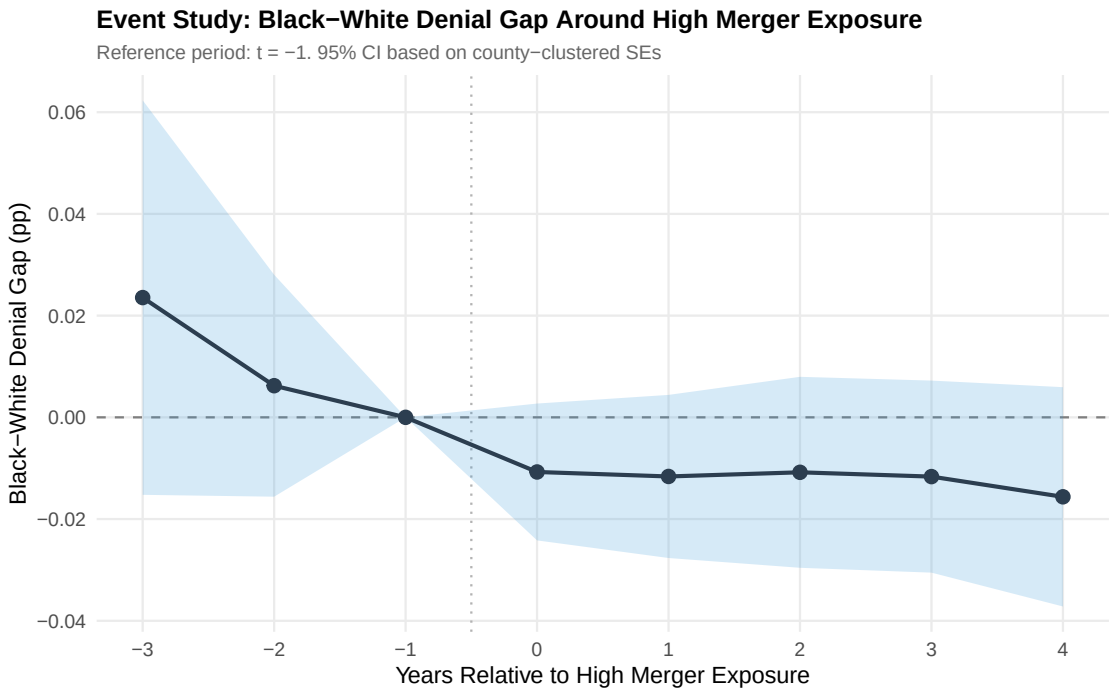

Figure 4: Event Study: Black-White Denial Gap Around High Merger Exposure Onset

Notes: Coefficients from a regression of the BW denial gap on event-time indicators, with $t = -1$ as the omitted reference period. 95% confidence intervals based on county-clustered standard errors. State and year fixed effects included.

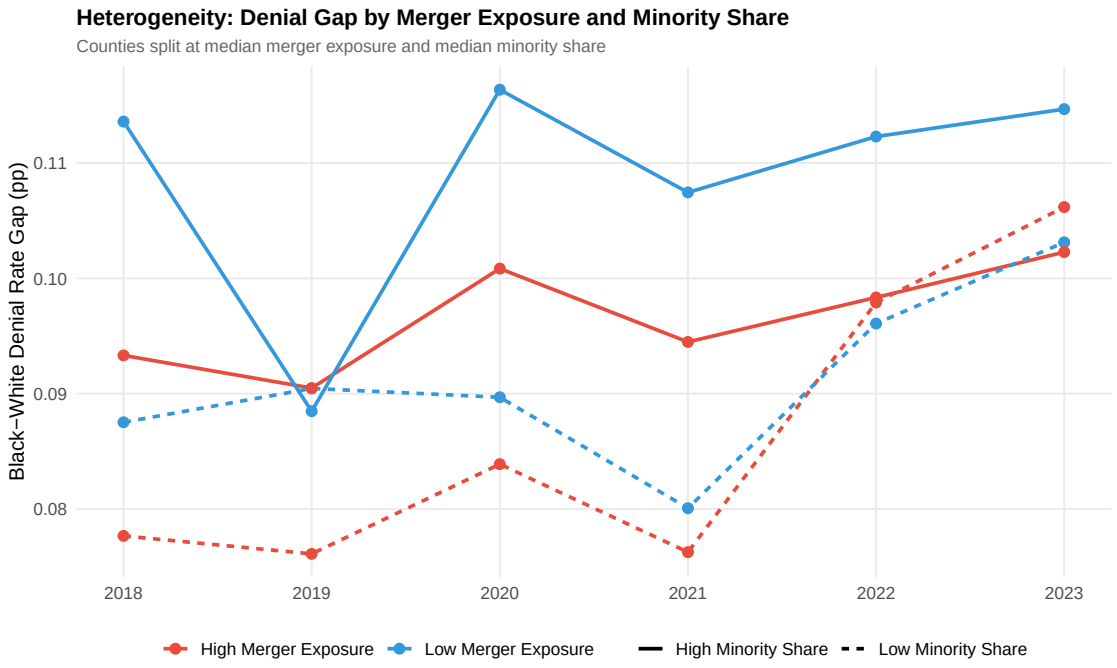

Figure 5: Heterogeneity: Black-White Denial Gap by Merger Exposure and Minority Share

Notes: Counties split at median merger exposure and median minority (Black + Asian) application share. Lines show mean BW denial gap over time for each group.

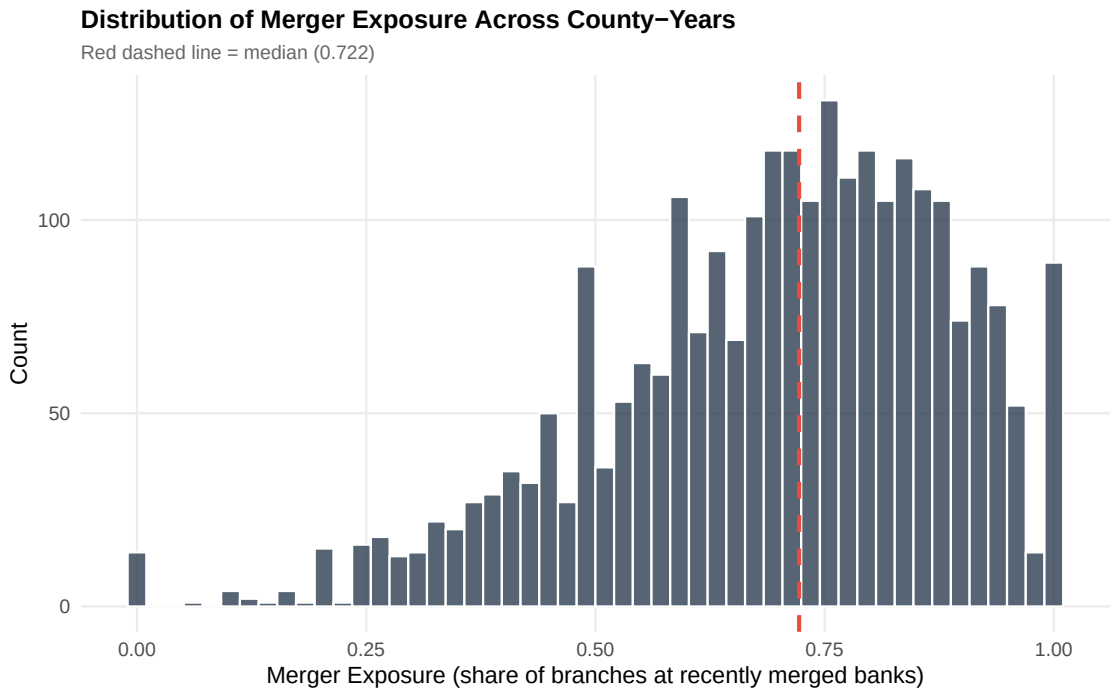

Figure 6: Distribution of Merger Exposure Across County-Years

Notes: Histogram of merger exposure (share of branches at recently merged banks) across all county-years in the sample. Red dashed line indicates the median.

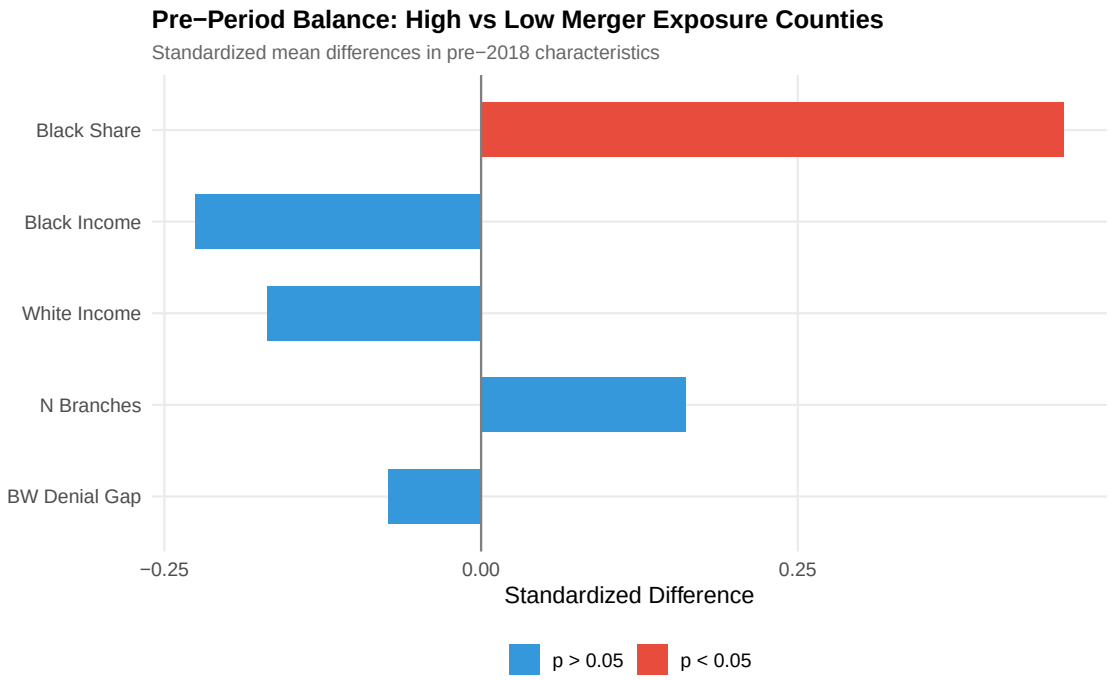

Figure 7: Pre-Period Balance: High vs. Low Merger Exposure Counties

Notes: Standardized mean differences in pre-2018 county characteristics between above-median and below-median merger exposure counties. Blue bars indicate $p > 0.05$; red bars indicate $p < 0.05$.

F. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
BW Denial Gap	1.665	0.754	0.060	27.96	12.66	Large positive
AW Denial Gap	−0.005	0.775	0.082	−0.05	9.42	Moderate neg.
BW Rate Spread	646.5	738.2	227.0	2.85	3.25	Large positive
Overall Denial	0.240	0.327	0.031	7.70	10.50	Large positive
Panel B: Heterogeneous (by minority population share)						
BW Gap (High Min.)	1.288	0.750	0.057	22.53	13.12	Large positive
BW Gap (Low Min.)	3.322	2.281	0.061	54.40	37.36	Large positive

Notes: **Country:** United States. **Research question:** Do merger-induced bank branch closures widen racial disparities in mortgage denial rates and pricing? **Policy mechanism:** Bank mergers approved under the Bank Merger Act lead acquiring institutions to close overlapping branches, reducing physical lending access points and eliminating established banking relationships. **Outcome definition:** Black-White denial rate gap (difference in fraction of home purchase mortgage applications denied between Black and White applicants in the same county-year); rate spread gap (difference in mean APR spread above benchmark). **Treatment:** Continuous; percent change in branch count instrumented by merger exposure (share of local branches at recently merged banks). **Data:** FDIC Summary of Deposits branch panel merged with CFPB HMDA loan-level microdata, 2018–2023, 20 U.S. states, county-year observations. **Method:** Two-stage least squares with merger exposure as instrument for branch changes; state and year fixed effects; standard errors clustered at county level. **Sample:** Counties with ≥ 20 Black and ≥ 50 White mortgage applications per year; 20-state sample covering approximately 70% of U.S. mortgage originations. SDE $= \hat{\beta}/\text{SD}(Y)$ where SD(Y) is the cross-sectional standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).