

Consolidation-Driven De-Branching

Evidence from Failed-Bank Resolutions

Sam Deegan^{a,*}

^a*University College Dublin, School of Economics, Belfield, Dublin 4, Dublin, D04 C1P1*

Abstract

Three forces drive bank de-branching: technology, fragility and consolidation. All three act at once, so each is hidden. This paper isolates consolidation in 1,261 FDIC resolutions that sold a failed bank whole to one acquirer. Bids are sealed and the branches reopen within days, so neither bank can know who will win. Each deal hands that acquirer counties where both banks had branches and counties where only one did. Within-deal comparison differences out technology and fragility. In the 146 deals that carry the contrast, first-year branch growth where both banks operated runs 9.1 log points below the counties the acquirer entered for the first time. Over four years the loss is about 4.6 branches per county, nine-tenths of what the failed bank brought, with no comparable contraction where only the acquirer operated.

Keywords: Branch networks, Bank mergers, Bank branch closures, Bank consolidation, JEL classification: G21, G28, G33, G34

1. Introduction

American banks closed a fifth of the network between 2009 and 2024 ([Federal Deposit Insurance Corporation, 2024b](#)). Three forces drive this de-branching: technology, fragility and consolidation. Technology has changed the returns to a physical presence. Fragility pushes banks to retrench from their weaker markets. Consolidation removes duplicate coverage ([Keil and Ongena, 2024](#)). The contribution of each remains unclear, because all three drive closures at once.

Where a bank locates its branches has consequences for the households and firms around it. Branch presence raises account ownership and, through it, savings among low-income households ([C  lerier and Matray, 2019](#)); lending to small and opaque borrowers rests on soft information carried by durable ties to a lender ([Petersen and Rajan, 1994](#)), and closures cut credit to local small firms ([Nguyen, 2019](#)). Closures from any of the three channels impose the same costs. What drives these contractions, and to what extent, is an open empirical question.

*Corresponding author

Email address: sam.deegan@ucdconnect.ie (Sam Deegan)

This paper isolates the consolidation channel in 1,261 FDIC purchase-and-assumption resolutions. A failed bank is auctioned by sealed bid and its network transfers to the winner within days, so neither bank can position for the outcome. The resolution shocks the depth and breadth of the acquiring network, leaving the acquirer with new and duplicate coverage and a consolidation decision to make.

Each deal-county falls into one of three types, fixed by which networks operated there before the resolution. In overlap counties, the acquirer already ran branches and now runs the failed bank’s as well, so it can serve the acquired customers from offices it already owns. In expansion counties it ran none, so the offices it takes over are the only ones it has there, and closing them means leaving the market. In pre-existing counties it ran branches and receives none, so the resolution should not move its branch count; those counties are the design’s placebo. The same acquirer takes all three on the same date. Consolidation predicts contraction in overlap counties alone.

Four confounds stand between a merger and a causal estimate of consolidation: technology, fragility, timing and selection. Each is neutralised by a feature of the resolution rather than by a control. Technology is a property of the institution, and a single acquirer stands behind every county in a deal on a single date, so it is common to all of them. Fragility is a property of the failed bank’s markets. A bank fails because its borrowers cannot repay, so overlap and expansion counties both sit inside the footprint that failed and both carry the conditions that closed it. Fragility differences out; duplication does not. The timing problem does not arise: failure and sale are announced together, and the branches reopen within days, so there is no window in which to reposition and none in which divestiture can strip the most duplicated offices. Selection survives only over who wins. Bidders are unaware of the decision weights and cannot observe competing bids (Allen et al., 2024), and which markets overlap is fixed before the auction. A deal fixed effect absorbs the winner’s identity, and confines the comparison to deals holding an overlap county alongside a county of another type: 146 of the 1,261, carrying 256 of the 369 overlap deal-counties.

Branch growth in overlap counties falls by 9.1 log points in the first post-resolution year relative to the counties the acquirer entered in the same deal, and remains below them through the fourth. Over those four years, the loss is about 4.6 branches per county, nine-tenths of what the failed bank brought. The counties the acquirer held alone receive no branches and show no sustained contraction. Estimates there are indistinguishable from zero in three of the four post-resolution years, and the exception, -4.6 log points in the third, does not survive corrective weighting. The acquirer absorbs the failed network without a general retrenchment.

This paper makes two contributions. The first is a measure of the consolidation channel net of the acquirer’s own condition and of the failed bank’s distress. Existing estimates set the merged bank against other banks, Nguyen (2019) against other large banks in the same county and Vij (2021) against the runner-up bidder, so both carry differences in institutional technology and financial condition alongside consolidation. Holding the acquirer, the failed bank and the date fixed removes technology and condition together, and leaves the treatment as whether the acquirer was already there. The second is that the design contains an untreated arm inside the same transaction. Nguyen (2019) observes buyer-only tracts and benchmarks them against tracts served by other banks; here pre-

existing counties sit inside the same transaction as the treated ones, so no institution varies across the contrast. Acquisition keeps the failed bank’s offices open where the acquirer was new to the county; where it was already there, the network sheds about as many branches as the failed bank brought. The acquirer’s prior footprint predicts which counties keep their branches.

Vij (2021) comes closest. He compares counties where the auction winner overlapped the failed bank against counties where the runner-up did, and finds more of the failed bank’s branches closed in the former. That contrast spans two banks, so differences in technology and financial condition remain. This one holds the institution fixed and adds the counties the failed bank held alone, which the same resolution treats and which have nothing to consolidate. Akkus et al. (2016) show why the choice of target matters: expected savings from overlapping networks are part of merger value, so the markets where networks overlap are the markets the acquirer chose. A resolution removes that choice.

The paper proceeds as follows. Section 2 sets out the channels and the identification problems they pose; Section 3 describes the auction and the panel; Section 4 the specification and its assumptions; Section 5 the results.

2. Channels and Identification

Keil and Ongena (2024) identify three channels of de-branching: technology, fragility and consolidation. First, technology changes the returns to a physical presence (Petersen and Rajan, 2002; Fuster et al., 2019). Online banking reaches customers without an office and lowers those returns; interstate branching deregulation had raised them by opening new markets. Second, fragility leads banks to retreat from weaker markets and hold on to profitable ones. Banks under funding pressure contract credit (Iyer et al., 2014), so each office carries less lending, and in US data, small banks close branches when they are fragile (Keil and Ongena, 2024). Third, consolidation works through redundancy. A bank can hold more offices in a market than it will support, and closures follow where two networks lie close together (Benson et al., 2024) or where coverage runs ahead of demand (DeLong, 2001). That redundancy also accumulates slowly, through gradual expansion or a legacy footprint the bank has outgrown. Because consolidation works on redundancy, its prediction is sharp: contraction concentrates in the markets the two networks held in common and nowhere else in the deal.

Branch closures are often endogenous to local conditions, and accumulated redundancy unwinds gradually in the absence of a shock. Bank mergers are attractive: they change a network’s depth and breadth on a single date. Keil and Ongena (2024) report that a branch of a merged bank is 21 per cent more likely to close, 215 per cent where the merging networks overlapped and 120 per cent for the acquisition targets, and state that they can neither detect causal effects nor rule out reverse causality. Timing is closed by the process; technology and fragility by the within-deal comparison; selection by both.

Comparing across banks matters for technology. An acquirer recasts its target in its own image. The consolidated bank’s small-business lending converges on the acquirer’s pre-merger share (Peek and Rosengren, 1998), so the thinning that follows may reflect

the adoption of the acquirer’s practices as much as the duplication of coverage. [Keil and Ongena \(2024\)](#) find internal information-processing technology to be more strongly associated with closure than customer-facing technology, so the variation that matters lies at the institutional level. A comparison that allows the institution to vary, therefore, conflates the acquirer’s technology with the duplication it inherited. Because that capability is chosen and run by the institution, it applies a single value across the counties of a single deal. It moves overlap, expansion and pre-existing counties together, and the deal–event–time effect absorbs it. It cannot generate a contrast between them.

Comparing across banks also matters for financial health. Thinly capitalised banks are more often acquired ([Wheelock and Wilson, 2000](#)), so the target may already be in the process of retrenching by the time the deal closes. Absorbing an institution is costly; acquirers often lose value on announcement ([Moeller et al., 2005](#)). Both are properties of the banks rather than the markets, and both add post-merger closures that duplication did not cause. Such a comparison also confounds the target’s distress with the duplication.

Fragility works the other way. It is a property of the failed bank’s markets rather than of the institution, so it reaches overlap and expansion counties alike and differences out of their comparison; it does not reach pre-existing counties, which lie outside that footprint.

Timing is a separate problem. A negotiated merger can take months from announcement to completion, and the networks change during that window. Direct evidence on what the parties do is thin, because closures are rarely measured before a deal closes. Banks divest branches as a condition of antitrust approval ([Pilloff, 2003](#)), and the probability of a required divestiture jumps at the concentration threshold ([Liebersohn, 2024](#); [Benson et al., 2024](#)). Those divestitures fall in the overlapping markets, so the network data record is already stripped of its most duplicated offices. Negotiated timing removes the duplication from the record before the outcome is measured.

Selection is the last. An acquirer chooses its target, and expected savings from overlapping networks are part of what it pays ([Akkus et al., 2016](#)), so the markets where two networks overlap are the markets the acquirer selected into. A comparison across mergers conditions on that choice.

A resolution meets them all at once. One acquirer takes the entire failed network on a single date, so its markets can be compared with one another rather than with other banks, and no window exists in which either party could reposition.

3. Institutional Setting, Data and Measurement

3.1. The Resolution Process

A bank fails when its chartering authority closes it and appoints the FDIC as receiver. The FDIC does not wait for that day to look for a buyer. Once a failure looks likely, it assembles a list of bidders qualified on capital and supervisory ratings, gives them access to the failing bank’s books under confidentiality agreements, and sets a bid deadline. Bids are sealed. The FDIC opens them, applies the least-cost test, and sells the bank to the bidder whose offer costs the insurance fund less than paying out the deposits. Bidders know the structure of the scoring rule but not the weights the FDIC applies within it, and they observe neither the number of rival bidders nor their characteristics. A bidder,

therefore, cannot forecast the outcome from what it holds at the deadline (Allen et al., 2024). Failure and sale are announced together, almost always on a Friday, and the acquirer opens the branches under its own name on the next business day (Granja et al., 2017).

In a purchase-and-assumption resolution, the acquirer assumes the failed bank’s deposits and buys some or all of its assets. Depositors keep their accounts and the offices stay open under new ownership. The alternative is a deposit payout, in which the FDIC pays insured depositors directly and the offices close. The sample is confined to whole-bank purchase-and-assumptions, which transfer substantially all assets to a single acquirer, so the failed network passes intact rather than being divided among buyers.

3.2. Dataset Construction

Three sources are combined: the FDIC Failures and Assistance Transactions database (Federal Deposit Insurance Corporation, 2024a), the Federal Reserve’s National Information Center relationship files (Board of Governors of the Federal Reserve System, 2024), and the FDIC Summary of Deposits (Federal Deposit Insurance Corporation, 2024b). From them, a panel is constructed at the deal–county–year level: for every county either institution operated in, the branches each held in every filing on both sides of the resolution. Comparisons are drawn between counties within a single deal, so each county enters the panel once per deal: a county in two resolutions appears twice, once for each acquirer and date. The rest of this section builds the branch panel, defines a deal, cuts the sample and classifies its counties.

Branch locations are derived from the Summary of Deposits, which records every deposit-taking office of an FDIC-insured institution, along with its county, as of June 30 each year. The public series begins in 1994; earlier filings come from the FDIC’s historical files, which report the same office-level fields on the same calendar. Branches are counted by bank–county–year, so each institution’s presence in every county is observed in every filing.

The Summary of Deposits records institutions. Pairing them into resolutions takes the two other sources. A deal is a transaction in which a solvent acquirer assumes the deposits and substantially all assets of a failed institution. The failed institution and its failure date are taken from the failure records; the acquirer and the transfer date are taken from the relationship files. The two records date the same event differently, so a deal is retained only where the two dates fall within seven days of each other, and every failed network in the sample passes to its acquirer within a week. The two institutions’ branches are then counted together by deal–county–year, so their combined presence in each county is observed on both sides of the resolution. Event time follows the filing calendar: $\tau = 0$ is the last filing before the resolution, so the first post-resolution observation falls between one and twelve months after the transfer.

The sample is drawn from the universe of FDIC resolutions via a sequence of restrictions, as traced in Section A.1. The FDIC records 4,111 failure and assistance transactions, of which 2,667 are whole-bank purchase-and-assumption resolutions. Failures whose network was split across more than one acquirer are dropped, because the records identify the acquirers but not the allocation of branches among them; each surviving acquirer–target

pair defines a unique deal, 2,270 in all. The seven-day window removes 475 of them. Matching both institutions to the Summary of Deposits, then requiring branch presence, a stable classification and a complete window, leaves 1,261 deals in 2,243 counties. The panel spans Summary of Deposits filings from 1981 to 2023. The estimation panel is 160,623 deal–county–year observations on 17,847 deal–counties in the contiguous United States.

The panel is built with six annual filings on either side of the resolution. The baseline event study estimates over the inner four; [Section A.3](#) re-estimates over wider windows. A wider window requires a longer balanced panel, so the set of deals narrows as the horizon widens.

Overlap counties are those where both institutions operated branches; expansion counties, those where only the failed institution did; and pre-existing counties, those where only the acquirer did. Each deal–county is classified from its pre-resolution filings and held fixed thereafter. Deal–counties whose classification changes across those filings are dropped, as are those not observed in every filing of the window; the window requirement removes 264, the balance requirement 1,351, and the restriction to the contiguous United States 173. Of the 17,847 that remain, 369 are overlap, 1,556 are expansion, and 15,922 are pre-existing. Overlap deal–counties contribute a stable ninety per cent of their cells to the growth-rate estimation at every event time.

3.3. Variable Description

[Table 1](#) defines the variables used in the analysis, in three groups. Each deal–county is assigned a county type from the two institutions’ presence over event time $\tau = -6$ to $\tau = 0$, measured relative to the resolution. The types are overlap, expansion and pre-existing. The classification is held fixed thereafter. County type is categorical, not a single indicator: it enters the specification as two dummies, $\text{Overlap}_{c,d}$ and $\text{PreExisting}_{c,d}$, each equal to one for a deal–county of that type and zero otherwise, with expansion the omitted category; both are interacted with event time, and $\tau = 0$ is the omitted reference. Branches are counted separately for the acquiring and failed institutions and summed to obtain deal–county branches, whose year-over-year log difference is the outcome. A closure matched by an opening in the same county registers as no change, and branches operated by other banks lie outside the measure.¹ Overlap share and county deposit concentration complete the set; both are measured at $\tau = 0$ and interacted with the overlap indicator in [Section 5.2](#). Overlap share, defined only for overlap deal–counties, is $\min(A, T)/(A + T)$, where A and T are the two institutions’ branch counts in the county: 50 per cent where the networks are the same size, near zero where one is much larger, so it measures symmetry rather than distance. Deposit concentration is a Herfindahl–Hirschman index of deposit shares across all banks in the county, on the 0–10,000 scale.

¹The log difference is undefined where the count reaches zero at either endpoint: 25,073 of the 160,623 cells, 19,203 with no branches at t and 5,870 with branches at t but none at $t - 1$.

Table 1: Variable Definitions: Deal-County Branch Panel

Variable	Description	Source
Panel A: Event Timing and Classification		
Event time (τ)	Index of Summary of Deposits filings relative to the resolution. $\tau = 0$ is the final pre-resolution filing and the omitted reference; negative and positive values index earlier and later filings. The panel is built over $\tau = -6$ to $+6$; the baseline estimates over $\tau = -4$ to $+4$.	FDIC SOD
Calendar Year	Calendar year associated with each deal-county- τ observation.	FDIC SOD
County Type	Categorical classification of a deal-county from the pre-resolution filings, held fixed thereafter: <i>overlap</i> where both the acquiring and the failed institution operated branches, <i>expansion</i> where only the failed institution operated, and <i>pre-existing</i> where only the acquirer operated. Enters as two indicators, Overlap and PreExisting, each one for a deal-county of that type and zero otherwise; expansion is the omitted category.	FDIC SOD/FRB
Panel B: Deal-County Branch Counts		
Acquirer-County Branches	Number of branches operated by the acquiring institution in county c in year t .	FDIC SOD/FRB
Target-County Branches	Number of branches operated by the failed institution in county c in year t .	FDIC SOD/FRB
Deal-County Branches	Total branches in county c associated with the deal in year t : the sum of acquirer and target branches.	FDIC SOD/FRB
Deal-County Branch Growth	Year-over-year log difference in deal-county branches. This is the outcome variable throughout.	FDIC SOD/FRB
Panel C: Exposure, measured at $\tau = 0$		
Overlap Share (%)	$\min(A, T)/(A + T)$, where A and T are the acquirer's and the failed institution's county branch counts at $\tau = 0$. Bounded at 50 per cent, where the two networks are the same size. Measures the symmetry of duplication, not distance. Defined only for overlap deal-counties.	FDIC SOD/FRB
Deposit Concentration (HHI)	Herfindahl-Hirschman Index of county deposit concentration, constructed from deposit shares across all banks operating in the county.	FDIC SOD

Notes: The unit of observation is a deal-county-year cell; a deal is a matched acquirer-target pair from an FDIC whole-bank Purchase-and-Assumption resolution. Log-difference growth rates are $\log(x_t) - \log(x_{t-1})$ and are undefined where the count reaches zero at either endpoint. Deposit concentration is on the 0–10,000 scale; overlap share is in percentage points. Sources are the FDIC Summary of Deposits (SOD) and the Federal Reserve Board (FRB) National Information Center relationship files.

3.4. Descriptive Statistics

Table 2 reports descriptive statistics. Three features of the sample bear on the design. First, local banking markets are concentrated: deposit concentration averages 2,207, above the 1,800 threshold at which a market is treated as highly concentrated in merger screening, and county bank and branch counts are right-skewed, with means well above their medians. Second, the panel is dominated by pre-existing counties: of 17,847 deal-counties, 369 (2.1 per cent) are overlap, 1,556 (8.7 per cent) expansion and 15,922 (89.2 per

Table 2: Descriptive Statistics: Deal–County Branch Panel

Variable	Mean	SD	Min	Median	Max	N
Panel A: Market Characteristics						
Calendar Year	2005.8	8.5	1981	2008	2023	160,623
County Bank Count	16.8	20.9	1	11.0	289	160,623
County Branch Count	77.5	162.4	1	25.0	1,811	160,623
Panel B: Deal–County Branch Counts						
Acquirer–County Branches	4.09	8.50	0	2.00	260	160,623
Target–County Branches	0.18	1.90	0	0.00	195	160,623
Deal–County Branches	4.27	8.81	0	2.00	260	160,623
Deal–County Branch Growth	0.014	0.172	−2.674	0.000	3.135	135,550
Panel C: Exposure						
Overlap Share (%)	30.7	14.7	1.4	33.3	50.0	3,051
Deposit Concentration (HHI)	2,207	1,404	382	1,826	10,000	160,614

Notes: The unit of observation is a deal–county–year cell; the baseline panel is balanced by construction, with 17,847 deal–counties at nine event times giving 160,623 cells. Statistics are unweighted. County bank and branch counts cover all institutions in the county, not only the two parties to the deal. Counts below 160,623 arise where a growth rate is undefined at a zero endpoint, where a concentration measure is missing, or — for overlap share — where the variable is defined only for overlap deal–counties. Exposure measures are defined at $\tau = 0$ and held fixed, so their counts are deal–county–year cells; the 3,051 overlap-share cells correspond to 339 overlap deal–counties. Variables are defined in [Table 1](#).

cent) pre-existing, so the acquirer’s footprint reaches far beyond the failed bank’s. Third, the outcome is dispersed relative to its mean: 0.172 versus 0.014, and that dispersion sets the precision of the estimates.

[Table 3](#) reports deal–county characteristics by county type at $\tau = 0$. The branch rows differ by construction: the acquirer holds none in expansion counties and the failed institution holds none in pre-existing counties. The market rows do not. Overlap counties are systematically larger, hosting 44.6 banks and 300.7 branches on average, compared with 22.5 banks and 114.5 branches in expansion counties, with correspondingly lower deposit concentration (1,821 against 2,302). Acquirer and target branches are more likely to co-locate in larger, more competitive markets. The consolidation effect is therefore identified in those markets. Whether it extends to smaller, more concentrated ones, where inherited overlap is rare, is untested.

Table 3: Deal–County Characteristics by County Type at $\tau = 0$

	Expansion	Overlap	Pre-Existing
Panel A: Deal-Network Exposure (differs by construction)			
Acquirer–County Branches	0.00	12.26	4.44
Target–County Branches	2.83	5.08	0.00
Deal–County Branches	2.83	17.34	4.44
Overlap Share (%)	—	30.7	—
Panel B: Market Characteristics			
County Bank Count	22.5	44.6***	15.5***
County Branch Count	114.5	300.7***	69.7***
Deposit Concentration (HHI)	2,302	1,821***	2,206
Deal–Counties	1,556	369	15,922

Notes: The unit of observation is a deal–county, measured at $\tau = 0$. Panel A differs across columns by construction — the acquirer operates no branches in expansion counties and the failed institution none in pre-existing counties — so no tests are reported. Panel B reports differences relative to the expansion baseline, with standard errors clustered by deal. County bank and branch counts cover all institutions in the county, not only the two parties to the deal. Deposit concentration is measured at the final pre-resolution filing and held fixed. Overlap share is defined for 339 of the 369 overlap deal–counties. Variables are defined in Table 1. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

4. Empirical Strategy

4.1. Identification

Each deal–county falls into one of three types, fixed by which networks operated there over the filings before the resolution. Table 4 sets out the classification and what each type is for; Figure 1 shows the three in a single resolution. Consolidation predicts contraction only in overlap counties.

Absent mismanagement or fraud, a bank fails because its borrowers are unable to repay their loans, so the failure is endogenous to the markets in which it lent. Those markets are fragile, and banks de-branch in response to fragility whether or not a resolution follows. Contraction inside the failed bank’s footprint is therefore consistent with fragility alone. Overlap and expansion counties both lie within that footprint and both carry the

Table 4: Deal–County Classification

At $\tau = 0$	Acquirer operates branches	Acquirer operates none
Failed bank operates branches	<i>Overlap</i> (369). Duplicated coverage passes to one owner. Contraction predicted.	<i>Expansion</i> (1,556). Inherited offices are the acquirer’s only presence; closing them means exit. Baseline.
Failed bank operates none	<i>Pre-existing</i> (15,922). No branches received. Placebo.	Outside the deal.

Notes: Counts are deal–counties in the estimation panel. The row differences out the fragility that closed the bank, because both treated types lie inside its footprint. The column differences out the acquirer’s technology and financial condition, because one acquirer takes every county in the deal on one date.




Figure 1: Deal-county classification in one resolution. The target, First National Bank, Edinburg, Texas, was closed on 13 September 2013 holding \$3.1 billion in assets and \$2.3 billion in deposits, and its offices reopened as branches of the acquirer, PlainsCapital Bank, Dallas, Texas. Branch counts are the two institutions' own offices as of the final pre-resolution filing on 30 June 2013 and exclude all other banks. Overlap counties, hatched, are the four in which both operated: Dallas, Travis, Williamson and Bexar, the metropolitan markets the acquirer already served, holding 8 target and 13 acquirer branches. Expansion counties, in mid grey, are the eleven in which only the target operated, running along the Rio Grande and the Gulf Coast, holding 44 target branches and none of the acquirer's. Pre-existing counties, in light grey, are the four in which only the acquirer operated, holding 15 acquirer branches. Tarrant County is not shaded because its classification varies within the window, and it is therefore dropped. Counties outside the deal are unshaded, and state boundaries and abbreviations are drawn for orientation. This deal is one of 1,261 and its composition is not that of the panel, in which 369 deal-counties are overlap, 1,556 expansion and 15,922 pre-existing.

conditions that led to the bank’s closure. What separates them is duplication. The same acquirer takes both on the same date, so comparing them holds the deal fixed. The shared fragility falls out, and the duplication does not. Pre-existing counties lie outside the failed bank’s markets and receive none of its branches, so changes in their branch count reflect the acquirer’s own strategy. The acquirer might cut branches because absorbing the failed bank strained its balance sheet, or because it was thinning its network across the board in response to a new technology. Overlap against expansion nets out the failure. Pre-existing against expansion tests whether branches move where there is nothing to consolidate. Both contrasts are reported.

Technology here refers to the information-processing and delivery capability with which a bank serves a customer: what its systems do inside an office and what they do outside one (Petersen and Rajan, 2002; Keil and Ongena, 2024). It sets how much business an office must carry to be worth holding, and so how densely the bank branches; the internal processing side is more strongly associated with closure than the customer-facing side (Keil and Ongena, 2024). That capability is bought and run at the level of the institution rather than the market, so it applies a single value across the counties of a single deal. That density is a property of the bank rather than of any market. After the resolution, every branch in the deal is the acquirer’s, so the same density governs overlap, expansion and pre-existing counties alike (Peek and Rosengren, 1998). A change in it moves the three together, and the overlap-against-expansion comparison differences it out. The resolution changes the number of offices the acquirer holds where it already has a presence. Pre-existing counties take no branches in the deal, so the resolution predicts no movement there. Section 5.1 reports that arm.

4.2. Empirical Specification

The design is a stacked difference-in-differences (Cengiz et al., 2019). Each resolution defines a stack: the deal’s counties, observed over the event window around it. Every comparison is drawn inside one stack, between county types, so a county treated in one resolution never serves as a control for a county in another.² Treatment varies across counties within a stack, and every county in a stack shares one event date, so the forbidden comparisons of staggered two-way fixed-effects estimators do not arise (Goodman-Bacon, 2021; Callaway and Sant’Anna, 2021; Sun and Abraham, 2021). What remains is how the stacks are aggregated, and Section A.2 reports the estimates using the corrective weights from Wing et al. (2024).

The outcome $y_{c,d,\tau}$ is the year-over-year log difference in deal–county branches, the two institutions’ offices in the county, summed in every filing on both sides of the resolution. Summing them from the start keeps the count continuous through the transfer. Before it, the deal–county network is two banks measured together. After it, the same offices are under one management. Movement at $\tau = +1$ is therefore a change in the number of offices. The construction also gives expansion counties a pre-period. The acquirer holds nothing there until the resolution, so that an acquirer-only count would be undefined before it. Expansion counties are the omitted baseline. The estimating equation is

²Dependence across stacks is discussed below.

$$y_{c,d,\tau} = \sum_{\substack{k=-4 \\ k \neq 0}}^4 \beta_k^O \text{Overlap}_{c,d} \cdot \mathbf{1}_{\{\tau=k\}} + \sum_{\substack{k=-4 \\ k \neq 0}}^4 \beta_k^{PE} \text{PreExisting}_{c,d} \cdot \mathbf{1}_{\{\tau=k\}} \\ + \gamma_{c,d} + \delta_{d,\tau} + \lambda_{s(c),t} + \varepsilon_{c,d,\tau},$$

where $\tau = 0$ is the final pre-resolution filing and the omitted reference, t the observation's calendar year, and $s(c)$ the state containing county c . Deal-county fixed effects $\gamma_{c,d}$ absorb each cell's average growth rate over the window, and with it the county-type main effects. Deal-event-time fixed effects $\delta_{d,\tau}$ restrict every comparison to one deal at one point in event time. State-year fixed effects $\lambda_{s(c),t}$ control for regional conditions, and bind only for the 14.6 per cent of deals whose counties lie in more than one state. In a single-state deal, state and calendar year are both fixed within a deal-event-time cell, so $\delta_{d,\tau}$ absorbs $\lambda_{s(c),t}$.

β_k^O is the gap in branch growth between overlap counties and expansion counties at event time k , measured against the same gap at $\tau = 0$. The coefficients β_k^{PE} do the same for pre-existing counties.

Because $\delta_{d,\tau}$ confines every comparison to one deal at one event time, a deal contributes to β_k^O only if it holds an overlap county alongside a county of another type. That requirement is met by 146 deals, holding 256 of the 369 overlap deal-counties. Of the 135,550 cells with a defined growth rate, 3,666 are singletons under the preferred fixed-effect structure and drop out. The remaining 131,884 form the common sample on which every column is estimated; the preferred column loses a further fifteen.

Standard errors are clustered two-way, by deal and by county (Cameron et al., 2011). The panel repeats counties across deals as well as deals across counties, and neither margin alone exhausts the dependence.

4.3. Assumptions and Validation

Reading the two contrasts as the effect of inherited overlap requires three assumptions.

First, absent the resolution, branch growth would have followed the same path in all three county types. The outcome is a growth rate, so the assumption is parallel growth rather than parallel levels. The threat is size: overlap counties hold six times the branches of expansion counties, and deal-county fixed effects on a differenced outcome remove each cell's mean growth rate and nothing else. If larger markets grow faster, the overlap arm gains, so the size effect understates the contraction.

Second, the $\tau = 0$ filing must precede any repositioning motivated by the failure. Every coefficient is measured against $\tau = 0$, so movement before that filing shifts the reference and carries into all of them. The auction turns on weights that the FDIC does not disclose and on rivals that a bidder does not observe (Allen et al., 2024), so neither institution could position itself against the acquirer's identity. That argument reaches the identity of the winner. What either institution did as the failure approached is a separate question, and Section A.5 tests it.

Third, the treatment must not spill over. Consolidation in overlap counties must not displace branching into the deal's other counties, and a county consolidated in one

resolution must not carry that thinning into a later stack. Displacement into expansion counties would raise the baseline and inflate the measured contraction. Deal-county fixed effects give a recurring county its own unit in each deal, so earlier thinning enters only through what is left to consolidate. [Section 5.1](#) reports the estimate against each baseline.

Granted these three, the comparison identifies the effect of inherited overlap. Which mergers that effect extends to is a separate question. Holding the winner fixed removes selection across bidders, not selection into the auction. The estimand is the effect of inherited overlap on branch growth. It applies to banks that bid on a failed network and win, not to banks merging at large. A bidder pays more where it expects savings from rationalising duplication ([Akkus et al., 2016](#)), so the estimate is plausibly an upper bound on what a randomly assigned merger of the same two networks would produce. Nearby institutions bid disproportionately. An acquirer that already operates in the failed bank’s counties can serve the acquired deposits from its existing offices, so the duplicated network is worth more to it than to an entrant ([Akkus et al., 2016](#); [Abedifar et al., 2024](#)). Overlap is therefore correlated with winning. Within a deal, it is not: once the winner is fixed, which of its counties overlap was settled before the auction. The estimate applies to the acquirers such auctions select.

5. Results

5.1. Branch Growth in Overlap Counties

Branch growth in overlap counties falls by 9.1 log points in the first post-resolution year relative to expansion counties in the same deal (SE 3.1). The contraction continues at 7.4 log points in year two, 9.0 in year three and 5.3 in year four. [Figure 2](#) plots both paths. Panel A, the overlap series, drops at $\tau = +1$ and stays below the baseline; Panel B, the pre-existing series, tracks zero at every horizon but the third. Overlap and expansion counties both lie inside the failed network and both pass to the acquirer at resolution, so the comparison nets out the fragility that put the target into receivership and any policy the acquirer applies to branches it has just absorbed. The difference between them is that the acquirer already operated in the overlap counties.

The pre-existing coefficient cannot be distinguished from zero in every post-resolution year but the third. Wherever that holds, the overlap coefficient is the consolidation effect. A joint test of the four pre-resolution overlap coefficients does not reject ($p = 0.19$). A pre-test that fails to reject is weak evidence on its own ([Roth, 2022](#)), so we report what it could have detected. At eighty per cent power the design detects a linear differential trend in growth of 1.7 log points per year, which would carry 1.7 log points into the first post-resolution coefficient against an estimated 9.1; the implementation of [Roth \(2022\)](#) gives 1.5. No pre-resolution coefficient in either panel is significant at five per cent.³ The third-year exception, -4.6 log points, falls to 1.5 and cannot be distinguished from zero under corrective stacked weights, while the overlap effect rises. The acquirer absorbs the failed network without retrenching elsewhere. The FDIC invites bids from

³The same test on the failed bank’s own pre-resolution path does not reject ($p = 0.44$). The acquirer’s own path rejects ($p = 0.04$) ([Roth, 2022](#); [Rambachan and Roth, 2023](#)); [Section A.5](#) bounds what that implies for the outcome.




Figure 2: Event study: overlap and pre-existing counties against expansion. Coefficients on the year-over-year log difference in deal-county branches by event time τ , $\tau = 0$ omitted. Panel A, overlap against expansion counties in the same deal. Panel B, pre-existing against the same baseline. Deal-county, state-year and deal-event-time fixed effects. Errors two-way clustered by deal and county. 95 per cent bands.

institutions it judges able to absorb the failure, so healthy acquirers are what the auction produces. Measured against the counties the acquirer already held, where the resolution delivered no branches at all, the first-year effect is 10.8 log points (SE 2.3). That comparison rests on an untreated arm inside the same transaction, and it is larger than the estimate against expansion counties. Displacement into expansion counties would push the expansion-baseline estimate above this one, and it is the smaller of the two.

The outcome counts the deal network, the branches the two institutions operated in the

county between them. The average overlap deal-county holds 17.3 at $\tau = 0$, of which 5.1 come from the failed institution (Table 3). The first-year fall of 9.1 log points is 8.7 per cent of that network, about 1.5 branches. Cumulating the four post-resolution coefficients gives a shortfall of 26 per cent by the fourth year, about 4.6 branches per county.⁴ Against the 5.1 branches the failed institution operated, that is a contraction equal to nine-tenths of what the failed bank brought, though the design cannot say whose offices closed. The outcome is net of openings, so these figures measure the size of the contraction rather than the fate of particular branches. Whether a customer of the merged bank can switch to another bank’s office is outside the scope of the design.

Table 5 builds the specification one set of effects at a time. Column (1) imposes none: the raw difference in branch growth across county types at each event time. The first-year coefficient should be stable across the columns, because the identifying variation lies within deal and event time, which none of the added effects absorb; the pre-resolution coefficients should attenuate as the comparison narrows to one deal at one point in event time.

The columns remove, in turn, persistent differences in growth between deal-counties, regional conditions in a given year, and everything common to a deal at one point in event time. The first-year estimate survives all three, so the contraction is not a product of the specification. What the removals fix is the pre-period. Four coefficients are significant in column (1) and none in column (4) at five per cent. Column (3) is the instructive failure. Removing regional calendar time without holding the deal fixed strips the design’s variation, and years two and four attenuate. In the preferred specification the post-resolution average is -7.7 log points, 95 per cent interval $[-11.4, -3.9]$.⁵ Widening the event window leaves the first-year estimate at 9.0 log points at ± 5 and 9.6 at ± 6 , while the balanced-panel requirement reduces the deal count from 1,261 to 1,046.

The closure effects reported elsewhere are much larger than the one estimated here, because the designs measure different quantities. Measured de-branching around a merger has two parts: rationalisation inside the merged network, and reallocation of coverage between the merged bank and its rivals. A comparison that sets a merged bank’s branches against branches owned by other banks contains both. Holding the acquirer and the target fixed within a deal isolates the first. Quantifying the second term requires a common metric that these estimates do not share.

⁴The overlap coefficient partly measures the acquirer’s own network, 12.3 of the 17.3 branches in the average overlap county. Netting it out gives $\beta_k^O - 0.71 \beta_k^{PE}$. The correction is under 1.5 log points except in year three, where it takes the contraction from 9.0 to 5.7 and the four-year cumulation from 26 to 23 per cent. It assumes the acquirer behaves alike in both arms.

⁵The interval rests on 146 identifying deals. Varying the clustering barely moves it, and randomisation inference, which does not use the asymptotics, rejects the sharp null at $p = 0.002$ (Young, 2019). Simultaneous bands cover all eight coefficients at once (Montiel Olea and Plagborg-Møller, 2019). Excluding the largest deals in turn leaves the path unchanged (Figure A.4), and reweighting stacks for unequal composition (Wing et al., 2024) raises the post-resolution average from 8.3 to 12.6 log points (SE 5.2) on the reweighting sample, with the pre-existing effect insignificant at 1.5. Table A.2 collects the fixed-effect ladder, the weighting, the comparison arm and the error structures; window width and leave-one-deal-out are reported in Figure A.3 and Figure A.4.

Table 5: Dynamic Effects of Consolidation on Branch Growth

	(1)	(2)	(3)	(4)
	Deal-County Branch Growth			
Panel A: Overlap counties				
$\tau = -4$	-0.028 (0.022)	-0.035 (0.022)	-0.004 (0.030)	-0.015 (0.022)
$\tau = -3$	-0.000 (0.035)	-0.005 (0.036)	0.011 (0.045)	0.020 (0.039)
$\tau = -2$	-0.049** (0.023)	-0.054** (0.023)	-0.009 (0.028)	-0.024 (0.021)
$\tau = -1$	0.037 (0.039)	0.038 (0.039)	0.059 (0.048)	0.033 (0.040)
$\tau = +1$	-0.092*** (0.033)	-0.092*** (0.034)	-0.093** (0.038)	-0.091*** (0.031)
$\tau = +2$	-0.082*** (0.029)	-0.078*** (0.029)	-0.045 (0.034)	-0.074** (0.031)
$\tau = +3$	-0.083*** (0.022)	-0.079*** (0.023)	-0.055** (0.022)	-0.090*** (0.021)
$\tau = +4$	-0.056*** (0.021)	-0.052** (0.022)	-0.025 (0.022)	-0.053** (0.021)
Panel B: Pre-existing counties				
$\tau = -4$	-0.038*** (0.012)	-0.046*** (0.012)	-0.025* (0.013)	-0.024* (0.014)
$\tau = -3$	-0.025** (0.013)	-0.033*** (0.012)	-0.027** (0.013)	-0.016 (0.013)
$\tau = -2$	-0.026** (0.011)	-0.032*** (0.010)	-0.025** (0.011)	-0.019 (0.014)
$\tau = -1$	0.027 (0.021)	0.026 (0.021)	0.031 (0.021)	0.023 (0.021)
$\tau = +1$	0.013 (0.022)	0.009 (0.022)	0.019 (0.021)	0.017 (0.020)
$\tau = +2$	0.001 (0.013)	0.001 (0.013)	0.001 (0.013)	-0.015 (0.015)
$\tau = +3$	-0.037** (0.015)	-0.036** (0.016)	-0.031** (0.015)	-0.046*** (0.015)
$\tau = +4$	-0.018 (0.013)	-0.016 (0.014)	-0.013 (0.015)	-0.020 (0.014)
Deal-county FE		✓	✓	✓
State-year FE			✓	✓
Deal-event-time FE				✓
Observations	131,884	131,877	131,876	131,869
R^2	0.004	0.153	0.231	0.338
RMSE	0.168	0.155	0.148	0.137

Notes: The unit of observation is a deal-county-year cell and the dependent variable is the year-over-year log difference in deal-county branches, so coefficients are in log units. Coefficients are event-time interactions with county type, taken relative to expansion counties in the same deal with $\tau = 0$ omitted; standard errors, clustered two-way by deal and by county, are in parentheses. Column (4) is the preferred specification: the deal-event-time effects restrict every comparison to one deal at one event time, and 146 deals contain an overlap county alongside a county of another type. All four columns are estimated on the common sample of 131,884 cells, so differences across columns reflect the controls rather than the sample. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

5.2. Intensity and Market Structure

The baseline asks only whether the acquirer was already in the county. It does not ask how much duplication it inherited, and consolidation implies a dose-response. Where the two networks are evenly matched, there is more to rationalise than where the acquirer holds one office against the target's twenty, so the contraction should deepen with the overlap share, and the interaction carries a negative coefficient if it does. If the acquirer closes a fraction of the branches the two networks hold in common, deal-county branches fall by that fraction times $\min(A, T)/(A + T)$, the bilateral overlap share of [Section 3.3](#), where A and T are the acquirer's and the failed institution's branch counts in the county. Against a mean share of 30.7 per cent, the first-year contraction of 8.7 per cent implies a fraction of about 0.28. [Figure 3](#) interacts the overlap indicator with that share and with county deposit concentration, both standardised and measured at $\tau = 0$.

Panel A does not isolate a scaling effect. Its coefficients are positive at every horizon, before the resolution as well as after, and three of the eight are significantly positive, two in the pre-period and one at $\tau = +3$. A series positive throughout measures a standing difference between high- and low-share counties rather than a response to the resolution. The first-year interaction is 1.5 log points per standard deviation (SE 2.8), and at eighty per cent power the test could only have detected a slope corresponding to closing about half the duplicated branches, well above the 0.28 implied above, so the flat panel is uninformative rather than evidence against a dose. The index measures symmetry rather than distance. A county holding one office from each institution forty miles apart sits at its top. [Benson et al. \(2024\)](#) recommend screening mergers on the proximity of the two networks. This index ignores distance, so the null here does not speak to that screen.

Panel B interacts the overlap indicator with deposit concentration and returns -4.5 log points per standard deviation in the first year (SE 1.6), indistinguishable from zero at every other horizon. One horizon of eight is weak evidence and the paper does not build on it. The sign is the opposite of what the concern in [Section 3.4](#) implies: the contraction is larger where deposits are more concentrated, not smaller.




Figure 3: Event study with overlap-share and concentration interactions. Panel A interacts the overlap indicator with the bilateral overlap share, Panel B with county deposit concentration. Both are measured at $\tau = 0$ and identified from overlap counties only. Otherwise, see [Figure 2](#).

6. Conclusion

This paper isolates the consolidation channel in branch decline using purchase-and-assumption resolutions, comparing markets within a single deal that differ in their exposure to the two pre-merger networks. The same acquirer, on the same date, holds technology and its own condition fixed, and both arms sit within the failed network. A bank can foresee a failure. It cannot foresee which bidder wins, and movement in the acquirer's own branches ahead of the resolution carries at most 2.5 log points into the outcome, against a first-year fall of 9.1.

A bank entering a market for the first time has no branches there to consolidate into, and closing the offices it acquires would mean leaving the market. A bank already present can serve those customers through its existing branches, and it closes surplus offices. Branch growth in overlap counties falls by 9.1 log points in the first post-resolution year. It stays below the expansion baseline for three more years, while the markets the acquirer held alone show no comparable contraction, so the acquirer takes on the failed network without cutting back elsewhere. Over the four years, the loss is about 4.6 branches per overlap county, a contraction equal to nine-tenths of what the failed bank brought, though the design cannot say whose offices closed.

Concentration rises in overlap counties by construction. The market loses an owner, and its deposits sit with one bank from the day of the sale, before a single branch closes. Within four years, the county's deal network has shed nearly as many offices as the failed bank brought. In expansion counties, the acquirer replaces the failed bank, and the market structure remains unchanged. A whole-bank resolution keeps depositors' accounts open across the board. Whether it keeps the offices open depends on whether the acquirer was already in the county. Four to five offices is the scale at which the credit and access effects documented by [C  lerier and Matray \(2019\)](#) and [Nguyen \(2019\)](#) operate, so the counties where the acquirer was already present are the ones exposed to them. Inherited overlap is observable at the bid deadline and predicts whether the county keeps the failed bank's offices. Bidders already price it. The resolution criterion does not.

The design prices the consolidation channel where it operates. It does not weigh that channel against technology and fragility. Both are common to the deal, so the comparison differences them out rather than measuring them, and how much of the national decline each of the three channels accounts for is a question this design cannot answer. The estimate is much smaller than the closure effects reported around mergers, and the two measure different things. It gives the size of the consolidation channel in markets where inherited overlap arises.

The same duplication arises in negotiated mergers, where the acquirer chooses which markets to overlap, has months to plan the closures, and may be in difficulty itself. Those transactions cannot identify the channel. A resolution removes the months, and comparing markets within a single deal nets out the competing channels, technology and fragility. Neither feature would be enough alone.

A. Additional Results and Diagnostics

A.1. Sample Construction

Table A.1 records the restrictions that cut the estimation panel out of the universe of FDIC resolutions, and the magnitude of each. Institutions are matched across three sources through the FDIC’s concordance between its certificate numbers and Federal Reserve RSSD identifiers: the certificate number links the failure record to the concordance, and the RSSD identifier links the concordance to the transformation record that names the acquirer. Pairs whose dates disagree by more than seven days are matches of unrelated events, not resolutions.

Table A.1: Sample Construction: FDIC Resolutions to Estimation Panel

Restriction	Unit	Dropped	Remaining
FDIC failure and assistance transactions	transactions	—	4,111
Retained resolution types	transactions	1,444	2,667
Joined to NIC transformation records	transactions	—	3,458
Failure transformations only	transactions	39	3,419
Resolved non-unique target matches	transactions	1,149	2,270
Within the seven-day transformation window	transactions	475	1,795
Deals with observed branch presence	deals	50	1,745
Deals surviving classification screens	deals	209	1,536
Deal-counties ever observed	deal-counties	—	30,764
With pre-resolution presence	deal-counties	10,495	20,269
With stable pre-resolution class	deal-counties	634	19,635
Observed in every filing of the window	deal-counties	264	19,371
Balanced across the estimation window	deal-counties	1,351	18,020
Contiguous United States	deal-counties	173	17,847

Notes: Counts follow the merge in order. The join to the transformation records raises the transaction count before it falls again, because one failure can match more than one transformation record until the non-unique matches are resolved.

A.2. Specification and Measurement

Table A.2 collects the estimates behind the figures that follow. Figure A.1 plots the four columns of Table 5 on common axes, and Figure A.2 re-estimates the path under corrective stacked weights (Wing et al., 2024).




Figure A.1: Event study across the fixed-effect ladder. One panel per rung, from no fixed effects through deal-event-time, matching columns 1 to 4 of [Table 5](#), on common axes. Otherwise as [Figure 2](#).

Table A.2: Robustness of the Overlap Effect

Specification	Estimate	Exhibit
<i>Panel A: Fixed-effect ladder, $\tau = +1$</i>		
No fixed effects	−0.092 (0.033)	Table 5, col. (1)
Deal-county FE	−0.092 (0.034)	Table 5, col. (2)
+ State-year FE	−0.093 (0.038)	Table 5, col. (3)
+ Deal-event-time FE (preferred)	−0.091 (0.031)	Table 5, col. (4); Figure A.1
<i>Panel B: Outcome, weighting and baseline</i>		
Corrective weights, post-resolution average	−0.126 (0.052)	Figure A.2
Pre-existing counties as baseline, $\tau = +1$	−0.108 (0.023)	Section 5.1
Post-resolution average (preferred)	−0.077 (0.019) [−0.114, −0.039]	Table 5, col. (4)
<i>Panel C: Inference, preferred point estimate</i>		
Heteroskedasticity-robust errors	−0.091 (0.025)	Figure A.5
Deal-clustered errors	−0.091 (0.032)	Figure A.5
County-clustered errors	−0.091 (0.024)	Figure A.5
Two-way, deal and county (preferred)	−0.091 (0.031)	Figure A.5
Simultaneous sup- t band, critical value 2.67	−0.091	Figure A.6
Randomisation inference, post-resolution average	$p = 0.002$	Figure A.7

Notes: Each row varies one feature of the preferred specification, which is deal-county, state-year and deal-event-time fixed effects with $\tau = 0$ omitted and standard errors clustered two-way by deal and by county. Panel A rows report the first-year coefficient on overlap counties relative to expansion counties. Panel B varies the weighting across stacks and the comparison arm, and reports the pooled post-resolution average; the corrective-weighted figure is a mean of the four post-resolution coefficients, not a first-year estimate. Standard errors are in parentheses and 95 per cent intervals in square brackets. Panel C holds the point estimate fixed and varies the error structure only. Window width and leave-one-deal-out re-estimation are reported in [Figure A.3](#) and [Figure A.4](#).




Figure A.2: Event study under corrective stacked weights, so that the pooled coefficient is an equally weighted average of deal-level effects rather than one weighted by each deal's county composition (Wing et al., 2024). Otherwise as Figure 2.

A.3. Sample Composition

Two exercises vary the sample rather than the specification. The first widens the event window ([Figure A.3](#)), which requires a longer balanced panel and so narrows the set of deals. The second drops one deal at a time ([Figure A.4](#)) to test whether any single resolution carries the result.




Figure A.3: Post-resolution coefficients across event-window widths, one series per window. The sample narrows as the horizon widens, because a wider window needs a longer balanced panel, with 1,261 deals at ± 4 , 1,172 at ± 5 , 1,046 at ± 6 . Otherwise as [Figure 2](#).




Figure A.4: First-year coefficient with one deal dropped at a time, one estimate per identifying deal with 95 per cent intervals. The horizontal line marks the full-sample estimate. The spread measures the influence of a single deal, not sampling error.

A.4. Inference

Three exhibits vary the inference rather than the estimate. [Figure A.5](#) re-estimates the intervals under each error structure, [Figure A.6](#) replaces pointwise intervals with simultaneous sup- t bands, and [Figure A.7](#) reports the randomisation distribution behind $p = 0.002$.




Figure A.5: Post-resolution coefficients under heteroskedasticity-robust, deal-clustered, county-clustered and two-way deal-and-county errors. Point estimates are identical by construction. Only the width of the 95 per cent intervals differs.




Figure A.6: Overlap path with simultaneous sup- t bands, which cover all eight coefficients jointly at 95 per cent and are therefore wider than the pointwise intervals plotted elsewhere.




Figure A.7: Permutation distribution of the post-resolution average effect under random reassignment of county types within deals, holding each deal's composition fixed. The vertical line marks the estimate. The p -value is the share of draws beyond it, with one added to the count and to the number of draws.

A.5. Identifying Assumptions

Figure A.8 counts only the branches the acquirer already held in overlap counties, so its pre-resolution coefficients test for repositioning ahead of the failure. No individual coefficient is distinguishable from zero; the joint test rejects ($p = 0.04$). Failure is foreseeable (Cole and White, 2012), so some adjustment ahead of it is expected. The movement is small: the acquirer holds seven-tenths of the branches in the average overlap deal-county, so the largest pre-resolution movement carries at most 2.5 log points into the deal-county outcome, against a first-year fall of 9.1.

Parallel trends can also be relaxed rather than tested. Rambachan and Roth (2023) bound the post-resolution violation by the pre-resolution one and report the largest bound at which the effect stays distinguishable from zero. Under the smoothness restriction, which limits how far the differential trend may bend from one year to the next, the first-year effect survives a bend of 0.03 log units and the second-year effect 0.01. The third and fourth survive at zero: they hold if the differential trend continues in a straight line from the pre-period, but not if it bends. Under the relative-magnitudes restriction, which scales the permitted post-resolution violation to the largest pre-resolution one, the breakdown value is zero. The overlap pre-period does not trend, its four coefficients alternating in sign, but its largest movement between adjacent years is 5.7 log points against a first-year fall of 9.1, so a violation of that size absorbs the estimate.




Figure A.8: The acquirer’s own branches in overlap counties before the resolution. The outcome counts only branches it already held, so the pre-resolution coefficients test for repositioning ahead of the failure. A joint test of the pre-period rejects ($p = 0.04$).

References

- Abedifar, P., Abdollahzadeh, M., Tarazi, A., White, L.J., 2024. The sale of failed banks: The importance of their branch networks and of the acquirers' financial strength. *Journal of Financial Stability* 75, 101338. doi:[10.1016/j.jfs.2024.101338](https://doi.org/10.1016/j.jfs.2024.101338).
- Akkus, O., Cookson, J.A., Hortaçsu, A., 2016. The determinants of bank mergers: A revealed preference analysis. *Management Science* 62, 2241–2258. doi:[10.1287/mnsc.2015.2245](https://doi.org/10.1287/mnsc.2015.2245).
- Allen, J., Clark, R., Hickman, B., Richert, E., 2024. Resolving failed banks: Uncertainty, multiple bidding and auction design. *The Review of Economic Studies* 91, 1201–1242. URL: <https://doi.org/10.1093/restud/rdad062>, doi:[10.1093/restud/rdad062](https://doi.org/10.1093/restud/rdad062).
- Benson, D., Blattner, S., Grundl, S., Kim, Y.S., Onishi, K., 2024. Concentration and geographic proximity in antitrust policy: Evidence from bank mergers. *American Economic Journal: Microeconomics* 16, 107–133. URL: <https://www.aeaweb.org/articles?id=10.1257/mic.20220094>, doi:[10.1257/mic.20220094](https://doi.org/10.1257/mic.20220094).
- Board of Governors of the Federal Reserve System, 2024. National information center relationship files. NIC bulk data download, FFIEC. URL: <https://www.ffiec.gov/npw/FinancialReport/DataDownload>. accessed August 2026.
- Callaway, B., Sant'Anna, P.H.C., 2021. Difference-in-differences with multiple time periods. *Journal of Econometrics* 225, 200–230. URL: <https://www.sciencedirect.com/science/article/pii/S0304407620303948>, doi:[10.1016/j.jeconom.2020.12.001](https://doi.org/10.1016/j.jeconom.2020.12.001).
- Cameron, A.C., Gelbach, J.B., Miller, D.L., 2011. Robust inference with multiway clustering. *Journal of Business & Economic Statistics* 29, 238–249. URL: <https://doi.org/10.1198/jbes.2010.07136>, doi:[10.1198/jbes.2010.07136](https://doi.org/10.1198/jbes.2010.07136).
- Célérier, C., Matray, A., 2019. Bank-branch supply, financial inclusion, and wealth accumulation. *The Review of Financial Studies* 32, 4767–4809. URL: <https://doi.org/10.1093/rfs/hhz046>, doi:[10.1093/rfs/hhz046](https://doi.org/10.1093/rfs/hhz046).
- Cengiz, D., Dube, A., Lindner, A., Zipperer, B., 2019. The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics* 134, 1405–1454. URL: <https://doi.org/10.1093/qje/qjz014>, doi:[10.1093/qje/qjz014](https://doi.org/10.1093/qje/qjz014).
- Cole, R.A., White, L.J., 2012. Déjà vu all over again: The causes of U.S. commercial bank failures this time around. *Journal of Financial Services Research* 42, 5–29. URL: <https://doi.org/10.1007/s10693-011-0116-9>, doi:[10.1007/s10693-011-0116-9](https://doi.org/10.1007/s10693-011-0116-9).
- DeLong, G.L., 2001. Stockholder gains from focusing versus diversifying bank mergers. *Journal of Financial Economics* 59, 221–252. URL: <https://www.sciencedirect.com/science/article/pii/S0304405X00000866>, doi:[10.1016/S0304-405X\(00\)00086-6](https://doi.org/10.1016/S0304-405X(00)00086-6).
- Federal Deposit Insurance Corporation, 2024a. Failures and assistance transactions. BankFind Suite, data files. URL: <https://banks.data.fdic.gov/bankfind-suite/failures>. data cover 1934–2023. Accessed August 2026.

- Federal Deposit Insurance Corporation, 2024b. Summary of deposits. BankFind Suite, data files. URL: <https://banks.data.fdic.gov/bankfind-suite/SOD>. data cover 1981–2023. Accessed August 2026.
- Fuster, A., Plosser, M., Schnabl, P., Vickery, J., 2019. The role of technology in mortgage lending. *The Review of Financial Studies* 32, 1854–1899. URL: <https://doi.org/10.1093/rfs/hhz018>, doi:10.1093/rfs/hhz018.
- Goodman-Bacon, A., 2021. Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225, 254–277. URL: <https://www.sciencedirect.com/science/article/pii/S0304407621001445>, doi:10.1016/j.jeconom.2021.03.014.
- Granja, J., Matvos, G., Seru, A., 2017. Selling failed banks. *The Journal of Finance* 72, 1723–1784. URL: <https://doi.org/10.1111/jofi.12512>, doi:10.1111/jofi.12512.
- Iyer, R., Peydró, J.L., Da-Rocha-Lopes, S., Schoar, A., 2014. Interbank liquidity crunch and the firm credit crunch: Evidence from the 2007–2009 crisis. *The Review of Financial Studies* 27, 347–372. doi:10.1093/rfs/hht056.
- Keil, J., Ongena, S., 2024. The demise of branch banking – Technology, consolidation, bank fragility. *Journal of Banking & Finance* 158, 107038. URL: <https://www.sciencedirect.com/science/article/pii/S0378426623002297>, doi:10.1016/j.jbankfin.2023.107038.
- Liebersohn, J., 2024. How does competition affect retail banking? Quasi-experimental evidence from bank mergers. *Journal of Financial Economics* 154, 103797. URL: <https://doi.org/10.1016/j.jfineco.2024.103797>, doi:10.1016/j.jfineco.2024.103797.
- Moeller, S.B., Schlingemann, F.P., Stulz, R.M., 2005. Wealth destruction on a massive scale? A study of acquiring-firm returns in the recent merger wave. *The Journal of Finance* 60, 757–782. doi:10.1111/j.1540-6261.2005.00743.x.
- Montiel Olea, J.L., Plagborg-Møller, M., 2019. Simultaneous confidence bands: Theory, implementation, and an application to SVARs. *Journal of Applied Econometrics* 34, 1–17. URL: <https://doi.org/10.1002/jae.2656>, doi:10.1002/jae.2656.
- Nguyen, H.L.Q., 2019. Are credit markets still local? Evidence from bank branch closings. *American Economic Journal: Applied Economics* 11, 1–32. URL: <https://www.aeaweb.org/articles?id=10.1257/app.20170543>, doi:10.1257/app.20170543.
- Peek, J., Rosengren, E.S., 1998. Bank consolidation and small business lending: It’s not just bank size that matters. *Journal of Banking & Finance* 22, 799–819. URL: <https://www.sciencedirect.com/science/article/pii/S0378426698000120>, doi:10.1016/S0378-4266(98)00012-0.
- Petersen, M.A., Rajan, R.G., 1994. The benefits of lending relationships: Evidence from small business data. *The Journal of Finance* 49, 3–37. URL: <https://doi.org/10.1111/j.1540-6261.1994.tb04418.x>, doi:10.1111/j.1540-6261.1994.tb04418.x.
- Petersen, M.A., Rajan, R.G., 2002. Does distance still matter? The information revolution in small business lending. *The Journal of Finance* 57, 2533–2570. URL: <https://doi.org/10.1111/1540-6261.00505>, doi:10.1111/1540-6261.00505.

- Pilloff, S.J., 2003. What's Happened at Divested Bank Offices? An Empirical Analysis of Antitrust Divestitures in Bank Mergers. Finance and Economics Discussion Series 2002-60. Board of Governors of the Federal Reserve System. URL: <https://www.federalreserve.gov/pubs/feds/2002/200260/200260pap.pdf>, doi:10.2139/ssrn.380343.
- Rambachan, A., Roth, J., 2023. A more credible approach to parallel trends. The Review of Economic Studies 90, 2555–2591. URL: <https://doi.org/10.1093/restud/rdad018>, doi:10.1093/restud/rdad018.
- Roth, J., 2022. Pretest with caution: Event-study estimates after testing for parallel trends. American Economic Review: Insights 4, 305–322. URL: <https://doi.org/10.1257/aeri.20210236>, doi:10.1257/aeri.20210236.
- Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics 225, 175–199. URL: <https://www.sciencedirect.com/science/article/pii/S030440762030378X>, doi:10.1016/j.jeconom.2020.09.006.
- Vij, S., 2021. Acquiring failed banks. URL: <https://ssrn.com/abstract=3234435>, doi:10.2139/ssrn.3234435. working paper, University of Georgia.
- Wheelock, D.C., Wilson, P.W., 2000. Why do banks disappear? The determinants of U.S. bank failures and acquisitions. The Review of Economics and Statistics 82, 127–138. URL: <https://doi.org/10.1162/003465300558560>, doi:10.1162/003465300558560.
- Wing, C., Freedman, S.M., Hollingsworth, A., 2024. Stacked Difference-in-Differences. NBER Working Paper 32054. National Bureau of Economic Research. URL: <https://www.nber.org/papers/w32054>, doi:10.3386/w32054.
- Young, A., 2019. Channeling Fisher: Randomization tests and the statistical insignificance of seemingly significant experimental results. The Quarterly Journal of Economics 134, 557–598. URL: <https://doi.org/10.1093/qje/qjy029>, doi:10.1093/qje/qjy029.