Consolidation Driven De-Branching

Sam Deegan[a,*]

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

Abstract

I isolate the consolidation channel of bank de-branching using FDIC Purchase-and-Assumption resolutions, in which failed banks are transferred to solvent acquirers through regulator-driven auctions completed within days. A stacked Difference-in-Differences design compares branch adjustments across overlap, expansion, and acquirer-only counties within the same-resolution event. Branch growth in overlap counties falls by 12.4 percentage points in the first post-resolution year and remains depressed through year three, with a cumulative decline of approximately 26%. No comparable immediate contraction occurs elsewhere. The effect does not scale with overlap intensity, indicating that geographic duplication, rather than its degree, drives the result. Because expansion counties absorb target-side distress, the estimated decline isolates consolidation from financial fragility and technological substitution.

Keywords: Bank Consolidation, Branch Networks, De-Branching, Resolution Policy

1. Introduction

Since 2009, U.S. banks have closed more than one-third of their branches. This contraction, commonly described as de-branching, has no single explanation. Existing accounts emphasise three broad forces: digital banking may substitute for physical presence, financial distress may force banks to retrench, and mergers may prompt rationalisation of overlapping networks (Keil and Ongena, 2024). Yet the relative contribution of each channel remains unclear, because branch closures typically occur in settings where all three operate simultaneously.

*Corresponding author

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

This paper isolates the consolidation channel by exploiting a shock to bank branch networks: FDIC Purchase-and-Assumption resolutions, in which failed banks are transferred to solvent acquirers through regulator-driven auctions completed within days. Mergers are the natural shock for studying consolidation because they combine previously separate branch networks, generating geographic overlap that may prompt rationalisation. In voluntary mergers, however, two features contaminate the overlap variation. First, the acquirer selected its target, so overlap geography is endogenous to the acquirer's expansion strategy; post-merger closures in overlap markets may reflect pre-existing plans rather than consolidation per se. Second, voluntary transactions unfold over months of due diligence and regulatory approval, giving acquirers scope to reposition their networks before and during the merger.

P&A resolutions remove both sources of contamination. The acquirer inherits the target's entire branch footprint, with no opportunity to pre-position its own network, and the resulting geographic overlap is a mechanical consequence of which bidder wins the sealed auction rather than a margin the acquirer optimises. The empirical design compares branch adjustment across three types of counties within the same resolution event. Overlap counties, those in which both the acquirer and the failed bank operated before resolution, capture inherited geographic duplication. Expansion counties, those in which only the failed bank operated, serve as the baseline and absorb any contraction attributable to target-side distress. Pre-existing counties, served only by the acquirer, provide a diagnostic for acquirer-wide retrenchment. Because all comparisons are made within deals, the acquiring institution, the failed bank, and the regulatory process are held fixed.

Branch growth falls sharply in overlap counties following resolution. In the preferred specification, the cumulative decline over the first three post-resolution years is approximately 0.30 log points, or roughly 26 per cent, relative to expansion counties. The effect is immediate, concentrated in the first post-resolution year, and persists through year three. Counties served only by the acquirer show no comparable contraction. The consolidation effect does not scale with the measured intensity of inherited overlap or pre-existing market concentration, indicating that the binary presence of overlap captures the relevant variation.

These findings make two contributions. First, the paper provides clean causal identification of the consolidation channel of de-branching, isolated from both distress and technological substitution. The P&A setting breaks the link between overlap geography and acquirer strategy that confounds

identification in voluntary M&A, permitting the consolidation channel to be identified in isolation. Second, the paper establishes that the consolidation effect operates on the extensive margin: the presence of inherited overlap, rather than its measured intensity, is associated with post-resolution contraction. While power is limited by the small number of overlapping counties, the absence of a dose-response relationship indicates that the binary overlap classification captures the relevant variation, rather than discarding information in continuous measures of duplication.

The remainder of the paper proceeds as follows. Section 2 reviews related research on relationship lending, bank consolidation, and resolution design. Section 3 describes the institutional setting, data construction, and measurement. Section 4 outlines the empirical strategy and identifies assumptions. Section 5 presents the results and robustness analyses. Section 6 concludes.

2. Related Literature

Branch proximity remains central to credit access for small firms and informationally opaque households because relationship lending relies on soft information that is difficult to transmit across distance (Petersen and Rajan, 1994; Berger and Udell, 1995; Degryse and Ongena, 2005; Agarwal and Hauswald, 2010). Consistent with this view, branch closures have been shown to reduce local credit availability and economic activity, particularly in markets with limited alternatives (Nguyen, 2019; Morgan et al., 2016). These findings establish that the spatial footprint of banking has real consequences, but they offer limited evidence on why that footprint contracts. Nguyen (2019) is closest to the present paper in documenting branch closures following mergers, and demonstrates that these closures reduce local credit supply. The present paper addresses a logically prior question: whether consolidation, isolated from the confounding forces present in voluntary M&A, actually drives branch closures.

Keil and Ongena (2024) organise these forces into three channels: technological substitution, financial fragility, and consolidation. The first two have received substantial empirical attention. Digital adoption reduces demand for in-person transactions (Petersen and Rajan, 2002; Fuster et al., 2019), and balance-sheet distress induces banks to retrench from weaker markets (Iyer et al., 2014). The consolidation channel is less well identified. Because mergers are voluntary and anticipated, post-merger branch closures conflate targeted rationalisation with pre-emptive repositioning. Disentangling the two requires variation in consolidation that is orthogonal to the strategic forces that shape branch networks — variation that voluntary M&A, by

construction, cannot provide. Mergers can generate efficiency gains by eliminating redundant capacity (DeLong, 2001), but they may simultaneously erode relationship-lending capacity for borrowers with high switching costs (Beck et al., 2018; Degryse and Ongena, 2004), making the net effect on local banking an empirical question that existing designs cannot cleanly answer.

The P&A setting studied in this paper resolves this identification problem by generating variation in inherited overlap that is not shaped by acquirer strategy, permitting consolidation to be isolated from the confounding forces described above. Bank failures are not random events; they cluster in recessions and correlate with local economic conditions(Peek and Rosengren, 2000; Ashcraft, 2005; Calomiris and Mason, 2003). A growing literature studies the FDIC resolution process itself, documenting that geographically proximate banks are more likely to bid (Granja et al., 2017), that auction design shapes acquisition outcomes (Allen et al., 2024), and that acquirer identity affects post-resolution performance (Vij). These findings imply that the identity of the winning bidder is not independent of local conditions, which would threaten identification strategies that compare across deals. This paper sidesteps the problem by exploiting a different source of variation: within-deal cross-county differences in inherited overlap. The paper thereby complements Dinger et al. (2024), which studies the balance-sheet and firm-level effects of distressed mergers, by focusing on the spatial reorganisation of branch networks.

3. Data and Measurement

3.1. Data Collection

I construct a panel dataset of deal-county-year observations covering all FDIC Purchase-and-Assumption resolutions in the contiguous United States from 1981 to 2023. The starting point is the FDIC Bank Failures and Assistance Database, restricted to P&A cases in which a solvent acquirer assumes all assets and liabilities. Acquiring institutions are identified through National Information Centre transformation files and linked to failed banks using FDIC-provided concordances. Transactions involving multiple acquirers are excluded. Each acquirer-target pair defines a unique deal.

To align failure events with annual branch data, I assign failures occurring before June 30 to the current calendar year and those occurring after June 30 to the subsequent year, corresponding to the first Summary of Deposits observation for the consolidated network. Event time is measured relative to this transformation year. I then merge deal identifiers to the FDIC Summary of Deposits, which reports the universe of U.S. bank branches and their

county locations as of June 30 each year. Branches are aggregated to the deal-county-year level, yielding 93,282 observations.

Counties are classified ex ante using the full pre-resolution window. A county is classified as overlap if both the acquiring and failed institutions operated branches in at least one pre-resolution year; as expansion if only the failed institution operated; and as pre-existing if only the acquirer operated. Events with unstable classifications, where counties switch categories before failure, are excluded, as are events with incomplete event windows. The primary outcome variable is the year-over-year log difference in the total number of branches operated by the combined acquirer-target network in a given deal-county.

3.2. Variable Description

Table 1 defines the variables used in the analysis. The primary outcome is the year-over-year log difference in the number of deal-county branches[1]. Overlap intensity and deposit concentration are measured at $\tau = -1$ and held fixed across event time to avoid post-treatment endogeneity.

[1]Because the log difference is undefined when branch counts reach zero, counties in which all deal-associated branches close exit the sample in that period. This attrition is asymmetric across county types. Overlap counties are mechanically protected: by construction, both institutions operate at least one branch before resolution, ensuring a floor of two. Expansion counties, where only the target operated, are more exposed. If a single-branch expansion county loses its branch, it exits the sample, removing the most severely contracting control observations. This would bias *against* finding a consolidation effect, making the estimates conservative.

Table 1: Variable Definitions: Deal–County Branch Networks and Market Structure

Variable	Description	Source
Panel A: Event Timing and Deal Characteristics		
Relative Year (τ)	Event time measured relative to the transformation year of the deal. $\tau = 0$ denotes the omitted reference period in event-study specifications; negative (positive) values correspond to pre- (post-) transformation years.	FDIC SOD/FRB
Calendar Year	Calendar year associated with each deal–county–event-time observation.	FDIC SOD
Panel B: County Banking Structure		
County Bank Count	Number of distinct banking institutions operating at least one branch in the county in year t.	FDIC SOD
County Bank Count Growth (Log-Difference)	Year-over-year log difference in the number of distinct banks operating in the county.	FDIC SOD
County Branch Count	Total number of bank branches operating in the county in year t.	FDIC SOD
County Branch Growth (Log-Difference)	Year-over-year log difference in the total number of branches located in the county.	FDIC SOD
Deposit Concentration (HHI, fixed at $\tau = -1$)	Herfindahl–Hirschman Index of county-level deposit concentration, constructed from deposit shares across banks and measured at $\tau = -1$. Held fixed across event time to avoid post-treatment endogeneity.	FDIC SOD
Deposit Concentration (HHI, time-varying)	Herfindahl–Hirschman Index of county-level deposit concentration, constructed from contemporaneous deposit shares across banks.	FDIC SOD
Panel C: Deal–County Outcomes and Exposure		
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 target institution in county c in year t.	FDIC SOD/FRB
Deal–County Branches	Total number of branches in county c associated with the deal in year t, equal to the sum of acquirer and target branches.	FDIC SOD/FRB
Deal–County Branch Growth (Log-Difference)	Year-over-year log difference in the number of deal-associated branches in a given county. This is the primary outcome variable in the event-study analysis.	FDIC SOD/FRB
Overlap Share (%, fixed at $\tau = -1$)	Percentage of deal–county branches that are overlapping between acquirer and target, measured at $\tau = -1$ and held fixed across event time.	FDIC SOD/FRB
Overlap Share (%, time-varying)	Percentage of deal–county branches that are overlapping between acquirer and target, measured contemporaneously.	FDIC SOD/FRB

Notes: The unit of observation is a deal–county–event-time cell. Log-difference growth rates are defined as $\log(x_t) - \log(x_{t-1})$. Fixed overlap intensity and fixed deposit concentration are measured at $\tau = -1$ and held constant across event time to avoid post-treatment endogeneity. Deposit concentration is measured using the Herfindahl–Hirschman Index (HHI) on the 0–10,000 scale.

3.3. Descriptive Statistics

Table 2 reports descriptive statistics. Three features of the sample are relevant for the empirical design. First, local banking markets are concentrated: the Herfindahl-Hirschman Index averages approximately 5,600, and the distribution of county-level bank and branch counts is highly right-skewed, driven by large urban markets. Second, pre-resolution overlap is common but not

universal, averaging 38 per cent where defined, with substantial cross-county variation. Third, deal-county branch growth exhibits wide dispersion, indicating sharp local adjustments around resolution events that provide the variation the within-deal design exploits.

Table 2: Descriptive Statistics: Deal–County Branch Networks and Market Structure

Variable	Mean	Median	SD	Min	Max	Count
Panel A: Event Timing and Deal Characteristics						
Relative Year (τ)	0.00	0.00	2.00	-3	3	93,282
Calendar Year	2007	2009	7.12	1981	2023	93,282
Panel B: County Banking Structure						
County Bank Count	16.1	11.0	18.6	1	289	93,282
County Bank Count Growth (Log-Difference)	0.0021	0.000	0.101	-2.40	1.61	92,024
County Branch Count	74.1	24.0	160.0	1	1,811	93,282
County Branch Growth (Log-Difference)	0.00299	0.000	0.0813	-2.08	1.79	92,024
Deposit Concentration (HHI, fixed at $\tau = -1$)	5,638	5,022	2,902	466	10,000	71,722
Deposit Concentration (HHI, time-varying)	5,675	5,081	2,838	397	10,000	73,443
Panel C: Deal–County Outcomes and Exposure						
Acquirer–County Branches	4.87	2.00	9.04	0	243	93,282
Target–County Branches	0.19	0.00	2.04	0	195	93,282
Deal–County Branches	5.05	2.00	9.32	1	243	93,282
Deal–County Branch Growth (Log-Difference)	0.0109	0.000	0.164	-2.67	3.14	92,024
Overlap Share (%, fixed at $\tau = -1$)	37.7	37.5	19.6	1.18	90.5	1,939
Overlap Share (%, time-varying)	38.0	38.5	19.3	1.18	90.9	1,108

Notes: The unit of observation is a deal–county–event-time cell. Relative Year τ is defined relative to the transformation year, with $\tau = 0$ denoting the omitted reference period in event-study specifications. Overlap intensity and fixed HHI are measured at $\tau = -1$ and held constant to avoid post-treatment endogeneity. Deposit concentration is measured using the Herfindahl–Hirschman Index (HHI) on the 0–10,000 scale. Differences in observation counts reflect missing overlap or concentration measures in some deal–county cells.

Table 3 compares pre-treatment means across county types at $\tau = -1$. Overlap counties are systematically larger markets: they host roughly twice as many banks and nearly three times as many branches as expansion counties, with correspondingly lower deposit concentration (HHI of 3,763 versus 5,022). These differences reflect the fact that both institutions are more likely to co-locate in larger, more competitive markets. Pre-existing counties, by contrast, resemble expansion counties in scale but are somewhat more concentrated.

Table 3: Pre-Treatment Mean Comparisons by Deal–County Treatment Category ($\tau = -1$)

	Expansion	**Overlap**	**Pre-Existing**
Acquirer–County Branches	0.000	10.390***	4.776***
Target–County Branches	3.763	4.567	0.000***
Deal–County Branches	3.763	14.957***	4.776
Overlap Share (%, branches)	NA	32.030	NA
County Bank Count	20.889	43.635***	15.376***
County Branch Count	104.641	295.440***	67.518***
Deposit Concentration	5,022	3,763***	5,694**

Notes: All variables measured at $\tau = -1$. Stars denote significance of differences relative to the Expansion baseline using deal-clustered standard errors. Deposit HHI is the Herfindahl–Hirschman Index on the 0–10,000 scale. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Some differences are mechanical given treatment definitions (e.g. acquirer presence in Expansion). The economically informative differences are market size and concentration.

These level differences do not threaten identification. The within-deal design controls for time-invariant county characteristics using deal-county fixed effects, and the parallel-trends evidence in Section 5.1 confirms that branch growth evolved similarly across county types before resolution. The relevant concern is not whether overlap and expansion counties are identical in levels, which they are not, but whether they would have followed parallel trajectories absent acquisition. The table does, however, clarify the external validity of the estimates: the consolidation effect is identified from larger, more competitive markets where co-location of acquirer and target is most common. The design does not speak to whether similar effects would obtain in smaller, more concentrated markets where inherited overlap is rare.

4. Empirical Strategy

4.1. Identification Challenges

The identification strategy rests on a simple comparison. When a bank fails and is acquired through a P&A resolution, the acquirer inherits branches in counties where it may or may not already operate. If the acquirer was already

present, the county now has overlapping networks from two institutions under common ownership. If the acquirer is absent, the county simply gains a new owner of the existing branches. Consolidation predicts contraction in the first type of county but not the second. Target-side distress, by contrast, affects both equally, because both lie within the failed bank's footprint and are exposed to the same failure-related shocks. Any contraction common to the failed bank's markets, whether from depositor flight, loan-book deterioration, or local economic weakness, is absorbed by the expansion baseline. Acquirer-wide retrenchment would appear in counties the acquirer already served, regardless of whether the target was present. Comparing these three county types within the same deal, therefore, separates consolidation from the two main confounding forces.

An alternative approach would compare outcomes across deals using variation in the identity of the winning bidder — for instance, contrasting counties that became overlap markets under the realised acquirer with those that would have been overlap markets under the runner-up. Such a design answers a different question: how auction outcomes shape branch networks across counterfactual acquirers. The present design instead holds the acquirer fixed and exploits geographic variation within a single ownership structure, which is the relevant comparison for isolating consolidation-driven rationalisation from institution-specific confounders. The two approaches are complements rather than substitutes.

The outcome variable, $y_{c,d,\tau}$, is the year-over-year log difference in the total number of branches operated by the combined acquirer-target network in county c at event time τ. Expansion counties form the omitted baseline. Formally, I estimate:

$$y_{c,d,t,\tau} = \sum_{\tau=-3,\,\tau\neq 0}^{3} \beta_\tau^{O}(\text{Overlap}_c \cdot \mathbf{I}_\tau) + \sum_{\tau=-3,\,\tau\neq 0}^{3} \beta_\tau^{\text{PE}}(\text{PreExisting}_c \cdot \mathbf{I}_\tau) + \gamma_{c,d} + \delta_{d,\tau} + \lambda_{s(c),t} + \varepsilon_{c,d,t,\tau}.$$

County fixed effects ($\gamma_{c,d}$) absorb time-invariant local characteristics within each transaction. Deal-event-time fixed effects $\delta_{d,\tau}$ ensure that all comparisons across county types occur at the same point in event time within the same deal, absorbing deal-specific shocks to branch adjustment that are common across markets. State-year fixed effects ($\lambda_{s(c),t}$) control for regional macroeconomic and regulatory conditions. The coefficients β_τ^{O} capture differential branch adjustments in overlap counties relative to the expansion baseline at each event time, controlling for contraction common to all markets

within the failed bank's footprint. The coefficients β_{τ}^{PE} provide a falsification test: if the post-resolution contraction reflects consolidation rather than acquirer-wide retrenchment, pre-existing counties should show no systematic decline.

4.2. Measuring Overlap

The baseline specification treats inherited overlap as a binary exposure. To assess whether consolidation scales with the intensity of inherited duplication, I augment the specification by interacting the overlap indicator with an ex ante measure of overlap intensity defined in the final pre-resolution year ($\tau = -1$):

$$O_{c,d}^{\text{snap}} = \frac{\min(\text{Target}_{c,d}, \text{Acquirer}_{c,d})}{\text{Acquirer}_{c,d} + \text{Target}_{c,d}} \times 100$$

Where $Target_{c,d}$ and $Acquirer_{c,d}$ denote the number of branches operated by the failed and acquiring institutions, respectively, in county c before the resolution of deal d. The measure captures bilateral duplication: it is maximised when the two institutions operate equal numbers of branches and approaches zero when one dominates the county. The measure is held constant over time to avoid post-treatment endogeneity. Because this intensity measure is identified only from overlap counties, which constitute a small subset of the panel, it is interpreted as a diagnostic rather than a structural parameter.

4.3. Assumptions and Validation

The design rests on two identifying assumptions. First, absent acquisition, branch growth in overlap and expansion counties would have followed parallel trends. Both county types lie within the failed bank's pre-resolution footprint and share exposure to the same local economic conditions associated with failure, making differential pre-resolution dynamics unlikely. Table 3 shows that overlap and expansion counties differ in market size and concentration, but these level differences are absorbed by deal-county fixed effects. The parallel-trends assumption is formally tested in Section 5 using pre-event coefficient estimates and visual inspection of pre-trends.

Second, the resolution must not be anticipated in a way that permits pre-positioning of branch networks. While bank failures may be partly predictable from balance-sheet deterioration (Cole and White, 2012), the identity of the acquiring institution is determined through a sealed-bid auction and is not known in advance. Moreover, resolutions are completed within

days, leaving no scope for acquirers to adjust their branch footprint before inheriting the target's network. Spillovers are assessed by verifying the absence of systematic post-resolution effects in pre-existing markets.

Finally, the stacked within-deal design avoids the weighting and contamination problems associated with two-way fixed-effects estimators under staggered treatment timing. Identification is achieved solely through within-deal cross-county variation, ensuring that estimates reflect consolidation-driven adjustment rather than differential exposure across transactions.

5. Results

5.1. Descriptive Patterns

Figure 1 plots mean deal-county branch growth by market type over the pre-resolution window. Overlap counties exhibit higher average branch growth than expansion counties throughout the pre-period, reflecting systematic differences in market composition. The gap between the two series, however, remains broadly stable: there is no evidence of differential acceleration or trend breaks as the resolution date approaches. Pre-existing counties track expansion counties closely. Together, these patterns support the parallel-trends assumption and provide no evidence of anticipatory adjustment prior to failure.

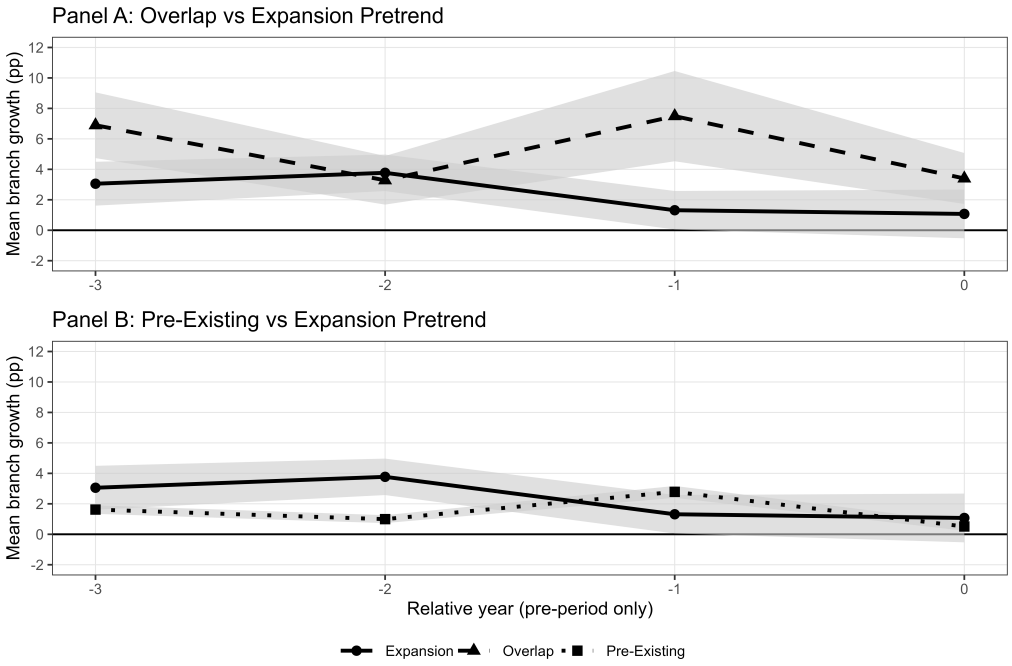

Figure 1: Parallel Trends (Visual Inspection)

5.2. Baseline Estimates

Figure 2 plots event-study estimates from the preferred specification (column 3 of Table 4), expressed relative to expansion counties with $\tau = 0$ omitted. Panel A reports the overlap coefficients. In the pre-resolution period, all three coefficients are small and statistically indistinguishable from zero, confirming the absence of differential pre-trends. In the first post-resolution year, branch growth in overlap counties falls by 12.4 percentage points relative to expansion counties ($p < 0.01$). The contraction persists at 7.6 percentage points in year two ($p < 0.05$) and 9.5 percentage points in year three ($p < 0.01$). The cumulative decline over the three post-resolution years is approximately 0.30 log points, or roughly 26 per cent.

Panel B reports the corresponding estimates for pre-existing counties. Coefficients are small and insignificant in the first two post-resolution years, consistent with the absence of acquirer-wide retrenchment. At $\tau = 3$, however, pre-existing counties show a significant decline of 5.9 percentage points. The critical observation is that pre-existing counties show a precisely zero effect in the first two post-resolution years, the period when the overlap contraction is largest. The delayed appearance and smaller magnitude of the pre-existing decline are inconsistent with consolidation, which should bite

immediately upon inheriting a duplicate network, and more consistent with a gradual post-acquisition portfolio review that is distinct from the overlap-driven response. The overlap effect at the same horizon remains roughly twice as large.

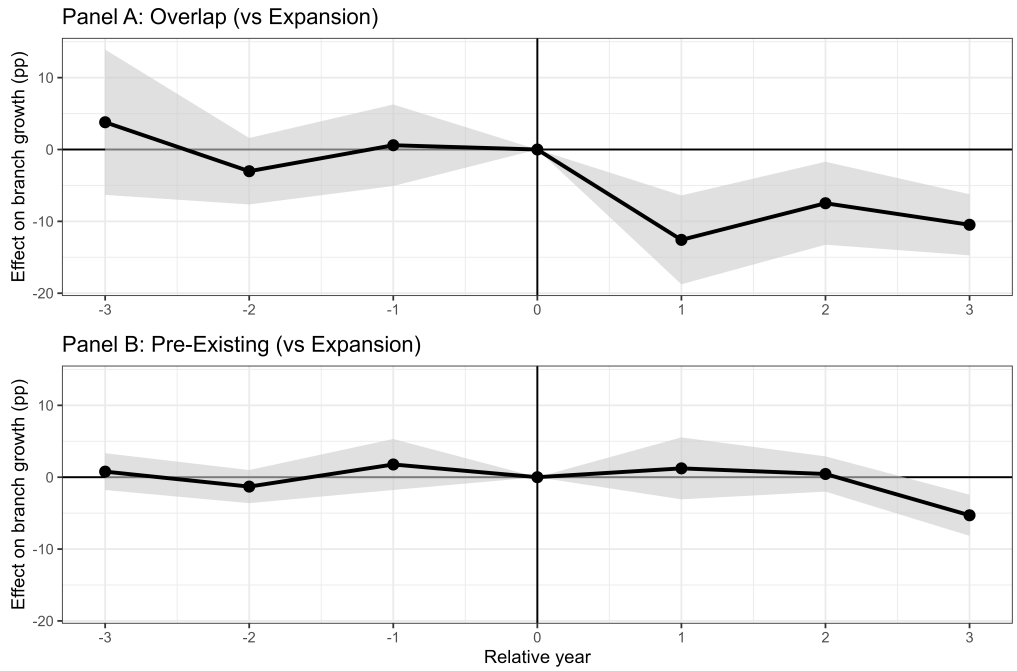

Figure 2: Baseline Stacked DiD

Table 4 traces these estimates across three specifications. Column 1 reports unconditional differences without fixed effects. Column 2 adds deal-county and state-year fixed effects, absorbing time-invariant local characteristics and regional macroeconomic conditions. Column 3, the preferred specification, adds deal-event-time fixed effects, ensuring that all comparisons across county types occur within the same deal at the same point in event time. The overlap coefficients increase as controls are added: the $\tau = +1$ effect roughly doubles from column 1 to column 3, while the pre-resolution coefficients approach zero. This pattern indicates that the earlier specifications understated the consolidation effect because deal-level time shocks common to all county types within a resolution were adding noise. The deal-event-time fixed effects absorb these common shocks, sharpening the within-deal contrast between overlap and expansion counties. The pre-existing coefficients are attenuated across the first two post-resolution years in all specifications, with the $\tau = +3$ effect appearing only in the most saturated column.

Table 4: Dynamic Effects of Consolidation on Branch Growth

	(1)	(2)	(3)
	Branch Growth$_{dct}$		
Panel A: Overlap Markets			
Relative Year$_{-3}$	0.059^{**}	0.024	0.025
	(0.025)	(0.023)	(0.039)
Relative Year$_{-2}$	0.023^{**}	-0.021	-0.030
	(0.011)	(0.017)	(0.023)
Relative Year$_{-1}$	0.065^{***}	0.034	0.024
	(0.018)	(0.023)	(0.030)
Relative Year$_{+1}$	-0.052^{***}	-0.066^{**}	-0.124^{***}
	(0.018)	(0.026)	(0.033)
Relative Year$_{+2}$	-0.046^{**}	-0.062^{***}	-0.076^{**}
	(0.018)	(0.023)	(0.031)
Relative Year$_{+3}$	-0.017	-0.058^{**}	-0.095^{***}
	(0.011)	(0.023)	(0.022)
Panel B: Pre-Existing Markets			
Relative Year$_{-3}$	0.007	0.003	-0.014
	(0.007)	(0.015)	(0.013)
Relative Year$_{-2}$	0.000	-0.013	-0.020
	(0.003)	(0.013)	(0.014)
Relative Year$_{-1}$	0.018	0.046^{**}	0.028
	(0.014)	(0.022)	(0.021)
Relative Year$_{+1}$	0.004	0.032	0.006
	(0.004)	(0.020)	(0.023)
Relative Year$_{+2}$	-0.010^{***}	0.002	-0.010
	(0.003)	(0.014)	(0.015)
Relative Year$_{+3}$	-0.010^{**}	-0.021	-0.059^{***}
	(0.004)	(0.019)	(0.017)
Panel C: Fixed Effects and Errors			
Deal–County FE		✓	✓
State–Year FE		✓	✓
Deal–Event-Time FE			✓
Standard Errors		Clustered – Deal	
Panel D: Model Summary			
Observations	92,024	92,024	89,984
R^2	0.004	0.182	0.332
RMSE	0.163	0.148	0.130

Notes: The dependent variable is the annual growth rate of branches in deal–county markets. Coefficients report event-time interactions between consolidation exposure and market type, relative to Expansion markets in the final pre-treatment period ($\tau = 0$). Column (1) reports unconditional differences without fixed effects. Column (2) includes deal–county and state–year fixed effects. Column (3) is the preferred specification, adding deal–event-time fixed effects to ensure all comparisons across county types occur within the same deal at the same point in event time. The slight reduction in observations in column (3) reflects singleton groups absorbed by the additional fixed effects. Standard errors are clustered at the deal level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.3. Overlap Estimates and Mechanism

The baseline specification treats inherited overlap as a binary exposure. To assess whether consolidation scales with the degree of geographic duplication, I interact the overlap indicator with the bilateral overlap measure defined in Section 4.2, which captures the symmetric extent of inherited duplication between acquirer and target. Figure 3 reports the resulting interaction coefficients. Panel A interacts overlap with the pre-resolution bilateral overlap share. Panel B interacts overlap with pre-resolution county-level deposit concentration, measured by the Herfindahl-Hirschman Index.

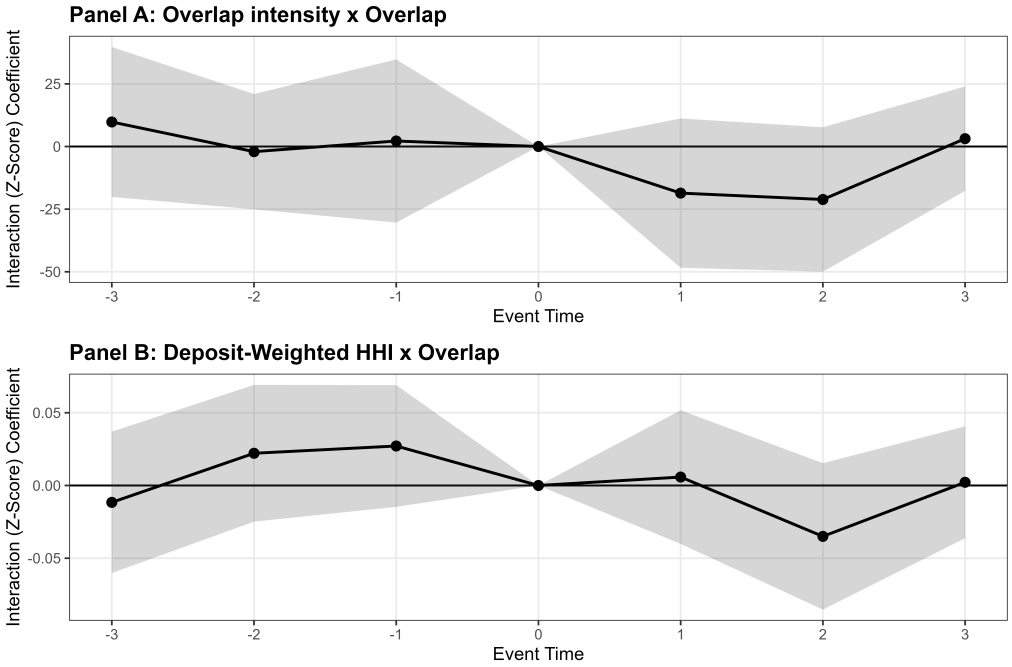

Figure 3: Overlap X Intensity Interaction

The interaction coefficients are imprecisely estimated across both panels and exhibit no stable pattern over event time. There is no evidence that post-resolution branch contraction increases with either the intensity of inherited overlap or pre-existing market concentration. This absence of a dose-response relationship indicates that the consolidation effect does not scale proportionally with marginal variation in overlap across counties.

Statistical power is a natural concern. The interaction specifications are identified exclusively from overlap counties, a small subset of the deal-county panel, and the preferred specification accounts for substantial variation in

deal-county, state-year, and deal-event-time fixed effects. Interaction coefficients are consequently estimated with wide confidence intervals. These results should be interpreted as evidence against strong monotonic scaling rather than as proof that overlap intensity is irrelevant.

The pattern is nonetheless informative. The absence of proportional scaling indicates that the consolidation effect operates on the extensive margin: the presence of overlap is associated with post-resolution contraction, but its measured intensity does not predict additional contraction. This result confirms that the binary overlap classification in the baseline specification captures the relevant variation. The interaction estimates do not separately identify competitive or market-power channels, and are best viewed as diagnostic: they discipline the interpretation of the baseline result without resolving the precise margin through which consolidation operates.

6. Conclusion

Resolution-driven acquisitions generate discrete, geographically targeted reductions in branch capacity. In the preferred specification, the cumulative decline in branch growth in overlap counties over the first three post-resolution years is approximately 0.30 log points relative to expansion counties, with the largest contraction concentrated in the first year. No comparable immediate adjustment occurs in counties served only by the acquirer, confirming that the contraction reflects inherited overlap rather than generalised retrenchment. A delayed decline in acquirer-only markets at the three-year horizon suggests some broader network adjustment may follow, but it is distinct in both timing and scale from the consolidation response.

These results speak to a tension in resolution design that has received little empirical attention. P&A auctions are structured to preserve financial stability and minimise losses to the Deposit Insurance Fund. They succeed on both counts, and relative to the alternative of outright liquidation, they almost certainly preserve more branches in aggregate, since the acquirer retains the bulk of the failed bank's network. But P&A does not distribute branch preservation evenly. Communities in which the winning bidder already operates bear a concentrated consolidation penalty that communities in the failed bank's expansion footprint do not. The rationalisation that follows is selective rather than indiscriminate, and the selection is determined by the accident of which bidder wins the auction rather than by local economic conditions. Whether the branches that close were genuinely redundant or whether their closure reduces credit access for locally dependent borrowers is a question the present analysis raises but does not resolve. Answering it re-

quires first establishing that consolidation causes branch closures in a setting where the channel is cleanly identified — which is what this paper provides. Linking the consolidation-driven closures documented here to downstream lending and borrower outcomes is a natural next step, but it requires different data and a different outcome variable, and is beyond the scope of the present analysis.

More broadly, the findings suggest that the spatial consequences of financial restructuring deserve the same empirical attention as its balance-sheet consequences. The banking system that emerges from a wave of resolutions may be better capitalised and more stable, but it is also more geographically concentrated, and the concentration is not randomly distributed. It falls on communities where the auction outcome generates inherited overlap, a selection mechanism that is unrelated to local credit demand.

Appendix

Agarwal, S., Hauswald, R., 2010. Distance and private information in lending. The Review of Financial Studies 23, 2757–2788. URL: https://doi.org/10.1093/rfs/hhq001, doi:10.1093/rfs/hhq001.

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.

Ashcraft, A.B., 2005. Are banks really special? new evidence from the fdic-induced failure of healthy banks. The American Economic Review 95, 1712–1730. doi:10.1257/000282805775014326. dOI: 10.1257/000282805775014326 MAG ID: 3123275304.

Beck, T., Degryse, H., De Haas, R., van Horen, N., 2018. When arm's length is too far: Relationship banking over the credit cycle. Journal of Financial Economics 127, 174–196. URL: https://www.sciencedirect.com/science/article/pii/S0304405X17302957, doi:10.1016/j.jfineco.2017.11.007.

Berger, A.N., Udell, G.F., 1995. Relationship lending and lines of credit in small firm finance. The Journal of Business 68, 351–381. URL: https://www.jstor.org/stable/2353332. publisher: University of Chicago Press.

Calomiris, C.W., Mason, J.R., 2003. Consequences of bank distress during the great depression. The American Economic Review 93, 937–947. doi:10.1257/000282803322157188. dOI: 10.1257/000282803322157188 MAG ID: 2073575615.

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.

Degryse, H., Ongena, S., 2004. The impact of technology and regulation on the geographical scope of banking. Oxford Review of Economic Policy 20, 571–590. URL: https://www.jstor.org/stable/23606936. publisher: Oxford University Press.

Degryse, H., Ongena, S., 2005. Distance, lending relationships, and competition. The Journal of Finance 60, 231–266. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2005.00729.x, doi:10.1111/j.1540-6261.2005.00729.x. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.2005.00729.x.

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.

Dinger, V., Schmidt, C., Theissen, E., 2024. The real effects of distressed bank mergers. Journal of Corporate Finance 89, 102674. URL: https://www.sciencedirect.com/science/article/pii/S0929119924001366, doi:10.1016/j.jcorpfin.2024.102674.

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.

Granja, J., Matvos, G., Seru, A., 2017. Selling failed banks. The Journal of Finance 72, 1723–1784. URL: https://www.jstor.org/stable/26652552. publisher: [American Finance Association, Wiley].

Iyer, R., Peydró, J., Da-Rocha-Lopes, S., Schoar, A., 2014. Interbank liquidity crunch and the firm credit crunch: Evidence from the 2007-2009 crisis. Review of Financial Studies 27, 347–372. doi:10.1093/rfs/hht056. dOI: 10.1093/rfs/hht056 MAG ID: 3080979461.

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.

Morgan, D.P., Pinkovskiy, M.L., Yang, B., 2016. Banking Deserts, Branch Closings, and Soft Information. Technical Report. URL: https://fedinprint.org/item/fednls/87107.

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., 2000. Collateral damage: Effects of the japanese bank crisis on real activity in the united states. The American Economic Review 90, 30–45. doi:10.1257/aer.90.1.30. dOI: 10.1257/aer.90.1.30 MAG ID: 2016839669 S2ID: 3a3cc06a0b40c7a667a05547feafdca63426b3c3.

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://www.jstor.org/stable/2329133, doi:10.2307/2329133. publisher: [American Finance Association, Wiley].

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://www.jstor.org/stable/3094536. publisher: [American Finance Association, Wiley].

Vij, S., . Acquiring failed banks doi:10.2139/ssrn.3234435.