

Failure in the Margins

Local Exposure to Non-local Bank Distress

Sam Deegan^{a,*}

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

Abstract

Estimates of the effect of US bank failures on their home counties are mixed, although variation in exposure peaks there and reverse causality should bias them away from zero and make any effect detectable. Two properties of branch networks generate offsetting biases of known sign. Branch networks are contiguous; omitting neighbours' exposure inflates the direct coefficient by two-fifths and discards a larger indirect effect. Plausibly exogenous failures appear roughly twice as damaging as the locally concentrated failures that dominate exposure, so averaging over all failures understates the effect. Instrumenting with exposure from failures at banks holding at most 5 per cent of their footprint locally, own exposure at its median non-zero value lowers annual income growth by 0.22 percentage points, and neighbour exposure at its own median by a further 0.10.

JEL: E44, G21, G28

Keywords: Bank Failures, Credit Disruptions, Branch Networks, Regional Economic Growth, Financial Stability

1. Introduction

Ample evidence exists for the negative effects of bank failure on economic growth through reduced firm investment and activity (Khwaja and Mian, 2008; Ivashina, 2009; Chodorow-Reich, 2014). Further, this is corroborated by aggregate evidence (Bernanke, 1983; Bernanke and Blinder, 1992; Schularick and Taylor, 2012). However, evidence at finer levels of geographic gradation is significantly more mixed (Gilbert and Kochin, 1989; Clair et al., 1994; Driscoll, 2004; Ashcraft, 2006; Greenstone et al., 2020; Calomiris et al., 1986; Ashcraft, 2005; Ghosh, 2017; Contreras et al., 2022; Huber, 2018; Bonfim et al., 2021;

*Corresponding author

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

Preprint submitted to XXXX

September 9, 2026

[Ge and Qiu, 2007](#)). That is difficult to square with the direction of the biases. Variation in exposure increases at increasingly smaller units. And failure is often conditional on borrowers no longer meeting repayments, so failures are triggered by the local weakness they are then credited with. Both should make an effect easier to find, not harder.

This paper assembles a county-year panel of U.S. bank failures from 1981 to 2023, mapping each failure onto counties through the failing bank's pre-failure branch footprint. Exposure is the deposits of failing banks as a share of the county's total, a proxy for the bank's weight in the local credit market, and personal income growth is regressed on it. Two things follow from the branch network. First, branch networks are contiguous and their catchments do not adhere to county lines, so a design measuring only a county's own exposure loads the effect of its neighbours' exposure onto its own, inflating the coefficient on own exposure and dropping the response to neighbours' exposure altogether. Second, I exploit an asymmetry in importance between a bank and its markets to construct an instrument for exposure. A bank may represent a large share of a county's banking services while that county is a negligible share of that bank's deposits, so the county's contribution to the failure is marginal.

At benchmark exposure, with each margin priced at its own median non-zero value, a county loses 0.32 percentage points of annual income growth in the failure year against a median growth rate of 2.34, of which 0.22 comes through its own exposure and 0.10 through its neighbours'. Omitting the neighbour margin inflates the direct effect by two-fifths, from 0.22 to 0.32, attributing the whole of it to the county itself. Instrumenting moves the direct effect further from zero rather than toward it, from 0.11 to 0.22, which is not what simultaneity predicts, since a failure that local decline helped cause should carry the inflated estimate rather than the deflated one. Split by crisis window, the indirect effect is a precise zero outside systemic episodes and 0.30 percentage points inside, while the direct effect does not detectably differ across the two. Those are also the episodes the geographically diversified failures arrive in, so calendar and composition cannot be separated here.

Designs in this literature differ in how they break the simultaneity of failure and local decline. [Ashcraft \(2005\)](#) pairs a national panel of realised failures with an experiment on healthy banks closed when their holding company failed, and in that experiment the treatment lands on a block of adjacent counties in a single state. [Kandrac \(2014\)](#) matches exposed counties to unexposed ones on observables, with geography entering the exposure measure but not the pairing. Others exploit shocks originating outside the local economy ([Peek and Rosengren, 2000](#); [Huber, 2018](#)), and this paper is closest to those, though the variation here is organisational, coming from where the failing bank held its balance sheet

rather than from where the shock began. Each maps bank distress onto counties through branch and deposit geography, as this paper does, but none carries the contiguity of that mapping into the estimating equation, so neighbours' exposure appears in the errors if at all and never in the conditional mean.

The paper makes three contributions. Firstly, it estimates the interference branch networks create rather than assuming it away. Firms enter markets adjacent to those they already serve, so a footprint is contiguous by construction and one failure treats a block of neighbouring counties at once, and a design treating those counties as independent overstates the direct effect while dropping the indirect one. Both errors follow in sign from that geometry before any estimate is run. Secondly, it records both sides of the relationship between a bank and a county. Where a shock reaches a local economy through a firm's multi-market footprint, the firm's weight in the market and the market's weight in the firm are distinct quantities, and only the second supplies exogeneity. Thirdly, it shows the indirect effect is not a constant. Residents who bank or work across the county line lose that lender or employer when a neighbouring bank fails, and outside a crisis the lenders and employers in their own county take them up. In a crisis there is no spare capacity. A pooled spillover averages the two states and describes neither, and the cost of omitting it falls on designs estimated in systemic episodes.

The remainder of the paper proceeds as follows. Section 3 describes the sources and constructs county exposure from pre-failure branch footprints. Section 4 sets out the estimating equation and the instrument. Section 5 reports the baseline estimates, traces the response over eight years, and then asks which failures supply the identifying variation, and whether it anticipates local decline. Section 6 asks when the neighbour effect exists and decomposes the response across income components, and Section 7 concludes.

2. Related Literature

Evidence on the local real effects of bank failures at the county and state level remains unsettled. Early studies find modest or short-lived effects (Gilbert and Kochin, 1989; Clair et al., 1994; Driscoll, 2004; Ashcraft, 2006; Greenstone et al., 2020), whereas others document substantial and sometimes persistent declines (Calomiris et al., 1986; Ashcraft, 2005; Ghosh, 2017; Contreras et al., 2022). Non-US studies show that bank-specific distress propagates to local activity (Huber, 2018; Bonfim et al., 2021; Ge and Qiu, 2007), but operate in institutional settings distinct from the U.S. The divergence is usually attributed to differences in samples, exposure definitions and identifying variation. Two of its sources are structural. Branch networks are contiguous, so a design treating counties as independent misallocates the effect in a direction the geometry fixes before any estimate

is run. And the neighbour effect those designs omit is not a constant, so studies estimated on different episodes are not estimating the same object.

A central institutional feature of U.S. failures is that they typically do not eliminate banking infrastructure. A failed bank is auctioned and reopens the following Monday under a competitor’s branding, its insured depositors made whole. Over the twentieth century, depositor protection and resolution evolved from slow liquidation toward rapid transfer of deposits and assets to solvent acquirers ([Friedman and Schwartz, 1963](#); [Ashcraft, 2005](#); [Balla et al., 2019](#); [Demirgüç-Kunt et al., 2015](#)), aimed at limiting runs, containing fire sales and preserving continuity of intermediation ([Bernanke, 1983](#); [Diamond and Dybvig, 1983](#); [Cowan and Salotti, 2015](#)). Branches often continue under new ownership; what does not survive intact is balance-sheet capacity, organisational capital and the lending relationships they carried.

Relationship-intensive lending remains geographically concentrated ([Petersen and Rajan, 1994](#); [Berger and Udell, 2002](#); [Berger et al., 2005](#)), physical branches play a central role in screening and monitoring borrowers ([Petersen and Rajan, 2002](#); [Degryse and Ongena, 2005](#)), and distance impedes the production of soft information ([Agarwal and Hauswald, 2010](#); [Granja et al., 2022](#)). Evidence from branch openings, closures and consolidation shows that changes in local banking presence affect credit supply and activity over narrow spatial radii ([Jayaratne and Strahan, 1996](#); [Amberg and Becker, 2024](#); [Célerier and Matray, 2019](#); [Nguyen, 2019](#); [Ranish et al., 2024](#)). If failures disrupt branch-based intermediation, their effects should be detectable in the counties the failing institution served. Relative to state aggregates, county outcomes offer sharper variation in exposure, since a branch footprint can be large relative to a county economy.

Analyses based on realised failures, including [Ashcraft \(2005\)](#), risk confounding failure effects with contemporaneous local deterioration. He notes that healthy-bank closures may understate the effects of typical failures, since ordinary failures follow poor underwriting and fall when local activity is already depressed, so the cleanest variation in that design comes from a small and unusual sample. [Kandrac \(2014\)](#)’s matched controls are selected without reference to geography, so a matched control may adjoin the treated county, or adjoin a different failure, and neither possibility is addressed. [Peek and Rosengren \(2000\)](#) and [Huber \(2018\)](#) identify from a contraction in lending by banks that survive; here the institution ceases to exist, so relationships are terminated rather than curtailed.

3. Data and Measurement

The analysis uses an unbalanced panel of U.S. counties from 1981 to 2023. Bank failures are recorded at the institution level while outcomes are observed at the county level,

so the panel has to be built rather than merged. I link each failed institution to its pre-failure branch footprint and aggregate exposure to the county-year level, matching 10,601 failed branches to counties over the period. Three restrictions then cut the merged frame of 130,563 county-year observations. County-years without a boundary match drop first, leaving 130,294 across 3,100 counties. County-years missing the outcome, the exposure measure, the instrument or a county centroid drop next, and since the outcome is a year-over-year change that removes 1981 entirely, taking 7,303 observations and 92 counties and leaving 122,991 county-years across 3,008 counties. One further county is absorbed by the state-year effects, so the baseline runs on 122,949 county-years and 3,007 counties.

Bank failures are identified using the FDIC Bank Failures and Assistance Database, which reports failure dates and institution identifiers but contains no geographic detail. I recover geographic exposure by matching each failed institution to its last reported branch network in the FDIC Summary of Deposits (SOD) through the institution's RSSD identifier. The match places 2,024 failing institutions in the branch data, yielding 3,121 failing bank-county cells. Branch coverage has three vintages. The Summary of Deposits reports branch locations and deposits from 1994. For 1987 to 1993 the branch panel is assembled from historical FDIC documentation, and for the years before that from reconstructed records following [Berger and Bouwman \(2009\)](#). All branch-level information is collapsed to county-year observations using pre-failure branch locations.

County economic outcomes are drawn from the Bureau of Economic Analysis. Personal income is measured at the place of residence and includes labour earnings, proprietors' income, capital income, and transfers. Outcomes are calendar-year totals. Because branch footprints and failures are observed as of June 30, income in year t is interpreted as the outcome realised over the calendar year containing the observed exposure. Nominal values are deflated using the GDP deflator, and outcomes are expressed as year-over-year changes in logs, scaled by 100, in real terms.

Geographic information is taken from Census TIGER/Line county shapefiles. When a shapefile is unavailable for a given year, the most recent prior boundary is used, preserving consistent spatial matching over time.

Together, these data provide county-level exposure to bank failures across the Savings and Loan Crisis, the Global Financial Crisis, and the 2023 banking episode. Relative to existing work, what the construction adds is the systematic linkage of institution-level failures to pre-failure branch footprints over four decades, which is what makes spatially explicit panel estimation possible.

3.1. Variable Description

Table 1 and Table 2 summarise the variables used in the analysis.

Panel A of Table 1 characterises local banking structure using branch counts, deposit and branch growth, and market concentration measured by the Herfindahl–Hirschman Index (HHI). These variables describe the scale and evolution of county banking markets. None enters the baseline specification, which carries no covariates beyond the fixed effects.

Panel B defines bank-failure exposure. An indicator captures whether a county experiences a failure in a given year. Intensity measures scale failed branches or deposits by a county base, using pre-failure branch footprints. The primary exposure variable is the Failed Deposit Ratio, which scales deposits in failed banks by the county deposit base, that is the sum of all banks’ branch deposits in the county, at the same pre-failure snapshot; the Failed Deposit-to-Income Ratio, which scales instead by lagged county personal income, is the primary robustness check. The instrument is exposure at failures for which the county accounts for no more than 5 per cent of the institution’s branch deposits across all counties, in markets marginal to the failing bank.

Table 2 reports income outcomes. Panel C decomposes personal income by place of residence into net earnings, transfers, and capital income. Panel D reports income by place of work, including wages, wage supplements, and proprietors’ income. These measures allow the analysis to distinguish resident-based income adjustments from production-side responses.

Spatial-lag (“SLX”) variants of key variables are constructed as row-standardised averages over the counties that share a border or a corner with the county in question, under queen contiguity.

3.2. Measuring Failure Exposure

County exposure is built from branch-level deposits, which proxy the geographic distribution of a bank’s economic weight. Deposits come from the FDIC Summary of Deposits. The primary measure is the Failed Deposit Ratio,

$$x_{c,t} = \frac{\sum_{b \in \mathcal{F}_t} \text{Deposits}_{cb}^{\text{pre}}}{\text{CountyDeposits}_c^{\text{pre}}}, \quad (1)$$

where $\mathcal{F}_t$ is the set of banks failing in year t and both terms are read at the pre-failure snapshot. A bank does not appear in the Summary of Deposits once it has failed, so its deposits can only be observed beforehand, and the county base is taken at the matching vintage. The ratio therefore compares two same-year deposit amounts and carries none

Table 1: County Banking Infrastructure and Failure Exposure Measures

Variable	Description	Source
Panel A: Bank Infrastructure		
Branch Count	Number of bank branches operating in the county in year t .	FDIC
Branch Growth (%)	Year-over-year log change ($\times 100$) in the number of branches located in the county.	FDIC
Deposit Growth (%)	Year-over-year log change ($\times 100$) in total deposits held in branches located in the county.	FDIC
Deposit Concentration (HHI)	Herfindahl–Hirschman Index of deposit concentration across banks in the county.	FDIC
Panel B: Bank Failure Indicators		
Failure Year (Dummy)	Indicator equal to 1 if at least one bank operating a branch in the county fails in year t (measured using pre-failure branch footprints).	FDIC
Failed Branch Ratio (%)	Failed branches as a percentage of total county branches, measured using branch counts in the year prior to failure.	FDIC
Failed Deposit Ratio (%) $[x_{c,t}]$	Deposits in failed banks as a percentage of the county deposit base, the sum of all banks' branch deposits in the county, both measured at the pre-failure snapshot. This is the paper's exposure measure.	FDIC
Exog. Failed Deposit Ratio $[z_{c,t}]$	The part of the ratio above arising at failed banks for which the county accounts for no more than 5% of the bank's branch deposits across all counties — exposure in markets marginal to the failing bank. This is the instrument.	FDIC
Failed Deposit-to-Income Ratio (%)	Deposits in failed bank branches prior to failure scaled by lagged county personal income, following Ashcraft (2005) and Kandrac (2014) . Reported as the primary robustness check.	FDIC/BEA

Note: Exposure ratios are measured at the pre-failure SOD snapshot, so the deposit stock and the county base both predate the failure they size; the failures themselves occur in year t . “SLX” denotes a spatial-lag variant of a variable, constructed as the row-standardised average of that variable across queen-contiguous neighbouring counties.

Table 2: County Income Measures: Residence- and Workplace-Based Components

Variable	Description	Source
Panel C: County Income by Place of Residence		
Personal Income Growth (%)	Year-over-year log change ($\times 100$) in total personal income, defined as the sum of net earnings by residence, capital income, and transfer receipts.	BEA
Net Earnings Growth (%)	Year-over-year log change ($\times 100$) in net earnings by place of residence, defined as earnings by place of work minus social insurance contributions plus the residence adjustment.	BEA
Transfer Growth (%)	Year-over-year log change ($\times 100$) in personal current transfer receipts, including Social Security, Medicare and Medicaid benefits, unemployment insurance, and other government and business transfers.	BEA
Capital Income Growth (%)	Year-over-year log change ($\times 100$) in capital income, defined as the sum of dividends, personal interest income, and rental income of persons.	BEA
Panel D: County Income by Place of Work		
Earnings Growth (%)	Year-over-year log change ($\times 100$) in earnings by place of work, including wages and salaries, supplements to wages and salaries, and proprietors' income earned within the county of employment.	BEA
Wage Growth (%)	Year-over-year log change ($\times 100$) in wage and salary disbursements earned within the county of work.	BEA
Wage Supplement Growth (%)	Year-over-year log change ($\times 100$) in employer pension and insurance contributions and other non-wage compensation allocated to the county of work.	BEA
Proprietor Income Growth (%)	Year-over-year log change ($\times 100$) in proprietors' income (farm and nonfarm) earned by businesses operating in the county of work.	BEA

Note: All series are annual county aggregates from the BEA Local Area Personal Income accounts, deflated to constant dollars and entered as $100 \times \Delta \log$. Panel C measures income where recipients live and Panel D where it is earned, and the Mechanism section reports both.

of the deposit flight the failure itself induces. The event is contemporaneous while the measurement is predetermined. That vintage is also the one the instrument is built on, so exogenous exposure is a component of total exposure in every county-year. Deposit-scaled exposure of this kind is used by [Bernanke \(1983\)](#), [Ashcraft \(2005\)](#) and [Kandrac \(2014\)](#). The Failed Deposit-to-Income Ratio, which replaces the denominator with lagged county personal income, is reported throughout as the primary robustness check.¹

Exposure is spatially clustered. Figure 1 maps the mean Failed Deposit Ratio for 2010 to 2019, the decade carrying the bulk of the financial-crisis failures. Exposure falls in contiguous blocks of counties rather than in isolated ones, so neighbouring counties are frequently exposed together and cannot be treated as independent observational units. Maps for the remaining decades, drawn on the same colour scale, are reported in Section [Appendix B](#).

3.3. Descriptive Statistics

Table 3 reports summary statistics for the variables entering the baseline specification. Failure exposure is rare and heavily right-skewed: the median county-year records no exposure at all, the mean Failed Deposit Ratio is 0.20 points against a standard deviation of 2.65, and failures occur in roughly 2 per cent of county-years. The instrument is rarer still by construction, averaging 0.04 points, since the marginal-market threshold admits only those county-years in which the failing bank held at most 5 per cent of its deposits locally. Deposit growth is volatile, with a standard deviation of 20 points against a median of 1.5, so the tails of the banking-structure series are wide even where the outcome is not. No series is trimmed or winsorised at any point in the pipeline. Spatially lagged variables show lower dispersion than their county-level counterparts. The neighbour average of the Failed Deposit Ratio has a standard deviation of 1.49 against the county-level 2.65, consistent with spatial aggregation smoothing idiosyncratic local variation. Panel C reports the alternative scalings used in Table [A.6](#), and Panel F the queen-contiguity neighbour counts: the average county has 5.84 neighbours and the average exposed county 5.83, so the neighbour margin is not drawn from an unusually connected part of the map.

¹The income-scaled measure is reported for comparability with the literature rather than preferred. Placing lagged county income in the denominator puts the outcome variable on both sides of the regression, so measurement error in income is shared between the regressor and the dependent variable. The original motivation for that denominator was to scale the shock by the size of the local economy, which is appropriate where depositor losses are the first-order channel. Over this sample insured balances are largely made whole and the failed institution's operations are usually preserved through resolution, so what a failure withdraws is intermediation capacity rather than deposits. The deposit denominator measures that capacity directly and compares two same-year deposit amounts drawn from one source, where the income-scaled ratio combines two series and so carries two sources of measurement error.




Figure 1: Mean Failed Deposit Ratio by County, 2010 to 2019

Notes: The map reports the spatial distribution of the treatment variable, $x_{c,t}$, averaged over the decade. County exposure is the deposits of banks failing in each year scaled by the county deposit base and constructed from pre-failure branch footprints. The decade from 2010 to 2019 is shown as an illustration because it carries the bulk of the financial-crisis failures. The crisis wave begins in the preceding decade, so no single decade map spans it. The colour scale is common to this map and to the decade maps in Section [Appendix B](#). Counties with no deposit base recorded in the Summary of Deposits are shaded as no data.

Table 3: Descriptive statistics: variables entering the baseline specification

Variable	Mean	Median	SD	Min	Max	Count
Panel A: Outcome						
Personal Income Growth (%)	2.22	2.34	5.54	−90.3	82.8	122,991
Personal Income Growth (SLX, %)	2.28	2.42	4.23	−40.4	37.8	119,971
Panel B: Failure exposure						
Failed Deposit Ratio (x)	0.20	0.00	2.65	0.0	100.0	122,991
Exog. Failed Deposit Ratio (z)	0.04	0.00	1.04	0.0	100.0	122,991
Failed Deposit Ratio SLX (Wx)	0.20	0.00	1.49	0.0	100.0	122,991
Exog. Failed Deposit Ratio SLX (Wz)	0.05	0.00	0.68	0.0	65.9	122,991
Panel C: Alternative exposure scalings						
Failed Branch Ratio	0.21	0.00	2.50	0.0	100.0	122,991
Exog. Failed Deposit-to-Income Ratio	0.02	0.00	0.40	0.0	49.6	122,991
Panel D: Failure incidence						
Failure Year (Dummy)	0.02	0.00	0.13	0.0	1.0	122,991
Panel E: County banking structure (not entered in the baseline)						
Deposit Concentration (HHI)	3,529.40	2,843.72	2,307.47	271.4	10,000.0	122,991
County Deposit Growth (%)	2.29	1.52	20.12	−805.1	647.7	122,979
Branch Count Growth (%)	0.80	0.00	15.45	−277.3	277.3	122,991
Panel F: Neighbourhood structure						
Neighbours per county, all counties	5.84	6	—	0	14	3,109
Neighbours per county, counties with exposure	5.83	6	—	1	13	1,143

Notes: County-year descriptives for the variables entering the baseline specification. All dollar-denominated series are deflated to constant 2024 Q1 dollars before any transformation, so growth rates are real; they are $100 \times \Delta \log$ year-over-year changes, so a value of -100 is a fall of 63 per cent and the bounds are wider than a percentage change can be; no series is trimmed or winsorised. SLX variables are row-standardised spatial averages over queen-contiguous neighbouring counties. Panel F counts queen-contiguity neighbours per county rather than county-years; the full set is the 3,109 counties in the weights matrix, which is built from the county shapefile rather than restricted to the estimation sample, and four of them are isolates, none exposed. Counts differ across rows because the estimation sample conditions jointly on the outcome, the treatment and the instrument.

4. Empirical Strategy

4.1. Identification

Identification rests on an asymmetry between a bank’s economic weight in a county and that county’s weight in the bank. The two need not move together, and failures where the second is small are failures a county cannot plausibly have caused but is still exposed to, once shocks common to the bank and the county are held fixed. Exposure is otherwise not randomly assigned. A bank fails when the markets carrying most of its balance sheet deteriorate, so endogeneity concentrates in counties that made up a large part of the failing institution. The instrument exploits the variation from failures where the county’s share of the institution was small.

Those common shocks are absorbed by the fixed-effects structure. Identification operates within county and within state-year, so the estimates rest on deviations from a county’s own trajectory and from the contemporaneous experience of its state. Branch networks also span administrative boundaries, so branches of a failing bank frequently lie in adjacent counties and a failure delivers indirect exposure to neighbours alongside direct exposure to the county itself. That is interference: one county’s exposure entering another’s outcome. The specification carries both margins and estimates it rather than assuming it away, and the estimated indirect effect is the coefficient on the SLX term.

Counties are not independent observational units. County growth rates are volatile, especially in small and rural areas, which weakens power and amplifies sensitivity to specification choices; more importantly, failure episodes are geographically concentrated (Figure 1), so neighbouring counties often share lenders and markets. Both contamination of nominal controls and interference among treated areas are then plausible, implying that bias in standard exposed–unexposed contrasts is a priori ambiguous. Partial treatment of nearby controls through shared banks and integrated labour markets compresses treated–control differences in clustered settings, consistent with formal results on interference (Miguel and Kremer, 2004; Duflo and Saez, 2003).

4.2. Econometric Specification

4.2.1. The Estimating Equation

Two quantities have to be estimated, the effect of a county’s own exposure and the effect of exposure among its neighbours, at the failure year and at horizons after it. A spatial lag of X specification delivers the first pair and local projections (Jordà, 2005) deliver the second. Let $\Delta y_{c,t} = 100 [\log Y_{c,t} - \log Y_{c,t-1}]$ denote annual real personal income growth in county c . The baseline, horizon-zero specification is

$$\Delta y_{c,t} = \beta x_{c,t} + \theta Wx_{c,t} + \alpha_c + \lambda_{s(c)t} + \varepsilon_{c,t}, \quad (2)$$

where W is a row-standardised queen-contiguity weights matrix, so $Wx_{c,t}$ is the average failure ratio among county c 's adjacent counties. The SLX term is a second treatment margin rather than a control, since a county with no failure of its own may be partially treated through its neighbours' exposure. Because the outcome carries no spatial lag, β is the direct impact of own exposure and θ the indirect impact of neighbour exposure, each priced at its own margin's exposure, and $\beta + \theta$ the current-year total impact of a one-point rise in the failure ratio in a county and in each of its neighbours simultaneously. Direct and indirect impacts in the sense of [LeSage and Pace \(2009\)](#) therefore coincide with the coefficients, and there is no feedback component to separate.

County fixed effects α_c absorb time-invariant heterogeneity. State-year fixed effects $\lambda_{s(c)t}$, where $s(c)$ is the state containing county c , absorb shocks common to a state in a year, so the estimates rest on within-county variation net of what that county's state was doing in the same year.

The specification carries no covariates beyond the fixed effects and no lags of the outcome. Bank health indicators such as non-performing loans or capital ratios are jointly determined with distress and local conditions, so conditioning on them would absorb the credit-contraction channel. They are also not county-level objects. Any such control has to be aggregated from bank balance sheets to the county through the branch network, so the weighting scheme becomes a modelling choice inside the control itself, and the specification would then condition on a constructed variable whose construction is not identified. Lagged growth is excluded for a separate reason. Under a serially clustered treatment it is itself partly an outcome of earlier exposure, so it is a bad control. Adding it changes little in any case, and the lag ladder is reported in [Section Appendix F.2](#).

4.2.2. Instrumenting Both Margins

Both $x_{c,t}$ and $Wx_{c,t}$ are endogenous, for three reasons. Failures are episodic and clustered in periods of macroeconomic stress, so exposure rises precisely when conditions deteriorate. Within episodes, banks sort into regions on business models, risk tolerance and expected profitability, so exposed counties can differ systematically in underlying income trajectories. And neighbour failures select on neighbours' conditions, which co-move with own conditions.

I therefore estimate equation~(2) by two-stage least squares with two excluded instruments. The first is exogenous failure exposure, which exploits the variation from failures at banks holding at most 5 per cent of their deposits in the county. These are the failures

for which the exposure asymmetry is sharpest, and hence for which the county's capacity to have influenced the insolvency is negligible. Formally,

$$z_{c,t} = \frac{\sum_{b \in \mathcal{F}_t} \mathbf{1}\{s_{cb} \leq 0.05\} \text{Deposits}_{cb}^{\text{pre}}}{\text{CountyDeposits}_c^{\text{pre}}}, \quad s_{cb} = \frac{\text{Deposits}_{cb}^{\text{pre}}}{\sum_{c'} \text{Deposits}_{c'b}^{\text{pre}}}, \quad (3)$$

where s_{cb} is the county's share of the failing bank's deposit base and the threshold binds on it, while the ratio itself is scaled by the county's own deposits exactly as the treatment is. The second instrument is the neighbour average, $Wz_{c,t} = \sum_j w_{cj} z_{j,t}$, with the threshold applied county by county before averaging, so that $z_{c,t}$ and $Wz_{c,t}$ are components of $x_{c,t}$ and $Wx_{c,t}$ by construction.

The first stage is the two-equation system

$$\begin{aligned} x_{c,t} &= \pi_{11} z_{c,t} + \pi_{12} Wz_{c,t} + \alpha_c^{(1)} + \lambda_{s(c)t}^{(1)} + u_{c,t}, \\ Wx_{c,t} &= \pi_{21} z_{c,t} + \pi_{22} Wz_{c,t} + \alpha_c^{(2)} + \lambda_{s(c)t}^{(2)} + v_{c,t}. \end{aligned} \quad (4)$$

Figure 2 reports what the threshold admits. Exposure is non-zero in 2,217 county-years and the threshold binds in 806 of them, and the two are not drawn from the same years. The threshold excludes almost the whole Savings and Loan wave from the instrument, where failing institutions were concentrated in single counties, and retains most of the wave that begins in 2008, where they were not.

With two endogenous regressors and two excluded instruments the model is exactly identified, so no overidentification test is available. Relevance is structural rather than estimated, since each instrument is a component of its endogenous regressor and loads on it at close to one (Table C.8). Anderson-Rubin sets, which are valid whatever the first stage delivers, accompany the 2SLS intervals at the baseline and are reported in Section Appendix D.

The identifying content of the design is the exclusion restriction that exogenous failure exposure affects county income growth only through realised failure exposure, own and neighbouring. The SLX term is instrumented rather than excluded because a design that claims interference has already claimed that neighbours' exposure enters the outcome equation, so $Wx_{c,t}$ is endogenous and receives its own excluded instrument.

Where effects differ across failures, two-stage least squares recovers a local average treatment effect: a weighted average across failures in which each failure's weight rises




Figure 2: Exogenous Failure Exposure by Year

Notes: Each bar is the count of counties carrying failure exposure in that year. The base band is the headline threshold at 5 per cent, and the lighter bands above it are what widening the threshold to 10, 15 and 20 per cent adds to the band below, so the bands are increments rather than levels. The top band is the exposure no threshold reaches.

with the exposure the instrument moves, over the population Figure 2 describes. What that population is, and what it implies for the estimand, is taken up in Section 5.3.

4.2.3. Dynamics and Inference

The local projections re-estimate equation~(2) at each horizon. The dependent variable becomes $\Delta y_{c,t+h}$ for $h = 0, \dots, 8$ with the right-hand side unchanged at t , and the $h = 0$ projection, estimated on the baseline sample with the same instruments, is the baseline regression. Cumulative responses come from cumulative local projections.

Panel fixed-effects local projections carry Nickell bias even where the regressors contain no lagged dependent variables, because the predictive specification is itself dynamic, and the bias affects inference as well as the point estimates (Mei et al., 2025). The correction is a split-panel jackknife, and applying it here moves the two margins differently (Section Appendix F.2). The direct effect at impact rises in magnitude and stays inside the uncorrected interval, so the estimate the paper stands on survives the correction. The indirect effect roughly doubles. That asymmetry is a further reason this paper reports θ alongside β rather than in place of it, and it means the uncorrected estimates understate the damage on both margins, in the direction Mei et al. (2025) report across the crisis literature.

Standard errors are Conley-type spatial HAC estimates (Conley, 1999) that allow arbitrary within-county correlation across years and spatial correlation across counties within 300 miles. The within-county component covers the serial correlation of order h that the projections induce. Alternative cutoffs, clustering schemes and a placebo assessment of empirical size are in Section Appendix D. Alternative weight structures, including distance bands and a bank-network operator, are in Section Appendix E.

4.3. Remaining Threats and Mitigation

The exclusion restriction is untestable, and it faces three threats. First, failure may follow local decline rather than cause it: deposits fall as county income falls, and banks fail after their deposits fall, so exposure lands in counties whose income was already falling. Second, the same shock can hit the bank and the county, and state-year effects remove only its state average: a bank’s markets share an economic base, so in a farm or oil bust the counties in its footprint fall relative to the rest of their state, failure or no failure.² Third, the spatial structure may be wrong: the specification reads a county’s

²Sub-state region-by-year effects are not estimable here: the instrument is non-zero in 0.66 per cent of county-years, so cells that fine would absorb the identifying variation. The SLX term does part of that work in the conditional mean instead, since a shock shared across a bank’s markets reaches the county’s neighbours as well as the county and enters through Wx .

indirect exposure off its adjacent counties, and if a county is instead reached through the failed bank’s branch network, or from beyond one ring, the SLX term mismeasures that exposure and the split between direct and indirect effects is misdrawn. The first is taken up by the lead and cascade tests in Section 5.4, the second by the fixed-effects saturation in Section Appendix E, and the third by the alternative weight operators reported there, including a bank-network variant.³

5. Results

5.1. Baseline Estimates

A one-point rise in the Failed Deposit Ratio lowers annual county income growth by 0.042 percentage points in the exposed county and by 0.072 through a simultaneous rise in exposure at all of its neighbours. Table 4 reports the ladder from ordinary least squares without fixed effects to the two-stage specification carrying the SLX term. Panel B prices each coefficient at the median non-zero value of its own regressor: 5.36 points for x and 1.41 for Wx . Every scaled figure in the paper uses those two medians unless its note says otherwise.

At benchmark exposure the direct effect is 0.22 percentage points of annual income growth, against a median county-year growth rate of 2.34. That is under a tenth of one year’s growth, and it is the number the design is built to produce: the failures it isolates are those of banks for which the county was one market among many, not the failures of the county’s own bank. The indirect effect is larger per point of exposure, but the two are not commensurable as event responses, since the SLX term prices a rise in exposure at every one of a county’s neighbours.

Three comparisons in the table carry the paper’s argument. The first is the fixed-effect structure. Replacing year effects with state-year effects halves the ordinary least squares estimate, from 0.051 to 0.025. Failures cluster regionally and income growth is correlated within state, so a coarser specification conflates common regional and failure exposure. Everything from this point on is identified within a state-year cell.

The second is spatial. Column (5) is the specification closest to the prior literature, instrumented and without an SLX term, and it returns 0.059 against column (6)’s 0.042. A design treating the county as an isolated unit loads part of the surrounding exposure

³One measurement concern is separate from exclusion and is addressed in Appendix Appendix A: deposits measured just before failure understate exposure where depositors have already fled. The measurement-timing ladder bounds this directly (Figure A.11): moving the deposit weight back one snapshot at a time moves the direct effect only between -0.22 and -0.24 , so the residual bias is small rather than merely signed.

Table 4: Baseline estimates: OLS and 2SLS, with and without the SLX term

Estimator	Personal Income Growth _{c,t} (%)					
	OLS				2SLS	
	Out (1)	Out (2)	Out (3)	In (4)	Out (5)	In (6)
Panel A: Coefficient estimates						
Failed Deposit Ratio, $x_{c,t}$	-0.042*** (0.011)	-0.051*** (0.007)	-0.025*** (0.005)	-0.021*** (0.005)	-0.059*** (0.020)	-0.042** (0.016)
Failed Deposit Ratio SLX, $Wx_{c,t}$				-0.037*** (0.006)		-0.072** (0.029)
Constant	2.22*** (0.116)					
Panel B: Effects at benchmark exposure (pp)						
Direct (β)	-0.23 (0.06)	-0.27 (0.04)	-0.13 (0.03)	-0.11 (0.03)	-0.32 (0.11)	-0.22 (0.09)
Indirect (θ)				-0.05 (0.01)		-0.10 (0.04)
Total				-0.17 (0.03)		-0.32 (0.11)
Panel C: Model summary and identification						
Observations	122,991	122,991	122,949	122,949	122,949	122,949
RMSE	5.535	4.887	4.267	4.267	4.268	4.267
Wald, x					3,462.7	4,835.7
Wald, Wx						118.8
Panel D: Fixed effects and errors						
County FE		✓	✓	✓	✓	✓
Year FE		✓				
State $\times$ Year FE			✓	✓	✓	✓
Standard errors	Conley (1999), 300-mile cutoff					

Notes: Dependent variable is annual county real personal income growth ($100 \times \Delta \log$). Exposure $x_{c,t}$ is failed deposits from banks failing in year t , scaled by the county deposit base; both are read at the same pre-failure snapshot, the vintage the instrument is built on, which makes $z \leq x$ hold county-year by county-year. W denotes queen-contiguity neighbour averages. Columns (5)–(6) instrument x (and Wx in (6)) with failures at banks holding at most 5% of their deposits in the county. The panel covers 3,007 counties over 42 outcome years (1982–2023, the first differenced year being 1982). Exposure is non-zero in 2,218 county-years in the wider frame of columns (1)–(2) and in 2,217 in the estimation sample of columns (3)–(6); the instrument is non-zero in 806. Panel A reports Conley (1999) spatial standard errors in parentheses. Panel B prices each coefficient at the median non-zero value of its own regressor: 5.36 points for x , over the 2,218 county-years in which x is non-zero, and 1.41 points for Wx , over the 7,354 in which Wx is non-zero. Total is both margins at their own medians together, with the covariance between the coefficients in its standard error. Standard errors in parentheses. Panel B is computed from unrounded coefficients, so the components need not sum exactly to the reported total. Mean outcome 2.22%. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

onto its own and overstates the direct county-level effect by two-fifths. The same design drops the indirect effect entirely, so it overstates what a county’s own exposure costs it and leaves what its neighbours’ exposure costs it unestimated. Both errors follow from the same fact, that branch networks are contiguous, and their signs are known before the estimates are run.

The third is the instrument. Within a state-year cell realised exposure still selects on the outcome, since a contracting county economy weakens its banks’ borrowers and deposit bases, so columns (5) and (6) instrument both margins. Simultaneity of that kind should attenuate the estimate. Instead it doubles, from 0.021 to 0.042 on the direct margin and from 0.037 to 0.072 on the neighbour margin, the two scaling by almost the same factor. The Wu–Hausman statistic rejects the equality of the two estimators, at 5.06, so the difference is not sampling noise. That gap is not a nuisance to be reported and set aside. What the threshold changes is which failures supply the identifying variation, and Section 5.3 takes it up.

The direct effect is the estimate the paper stands on. It holds at 5 per cent under all three Conley bandwidths and under clustering on county (Table D.9), and it keeps its sign and its significance across seven alternative neighbour operators, ranging from 0.035 to 0.055 (Table E.13). The indirect effect is reported alongside it and carries less weight, since it loses significance under the coarsest clusterings. Its magnitude is stable once each operator is priced at its own benchmark: across the five row-standardised operators the impact effect runs from 0.041 to 0.124 percentage points, and under queen, Delaunay and k -nearest neighbours it lies between 0.102 and 0.124 (Figure E.22). The raw coefficients in Table E.13 are per unit of differently scaled quantities and are not comparable across columns. The two unstandardised operators and the bank-network operator are not commensurable with these on either scaling.

5.2. Dynamics

Figure 3 traces the response over the eight years after failure. Each horizon is a separate estimation of equation (2) with $\Delta y_{c,t+h}$ on the left, so the $h = 0$ estimates reproduce Table 4, column (6), by construction. The failure year and the year after carry the contraction, at 0.22 and 0.23 percentage points at benchmark exposure, both significant at 5 per cent. The response is indistinguishable from zero at $h = 2$ and $h = 3$, and then turns positive, at 0.12 at $h = 5$ and 0.14 at $h = 7$. The growth path does not simply flatten after the shock; it overshoots and then settles.

A growth rate that returns to zero is not a recovery; it is a return to its growth path. The cumulative projection converts the annual path into a level shortfall. It deepens




Figure 3: Dynamic Responses of County Income Growth to Failure Exposure

Notes: The figure reports local-projection responses of annual county real personal income growth to failure exposure over horizons t to $t+8$. The three panels report the direct effect of own-county exposure, the indirect effect of neighbour exposure through the queen-contiguity SLX term, and their total. Each horizon is a separate regression on the specification of Table 4, column (6), so the $h = 0$ estimates reproduce that column by construction. Estimates are scaled to benchmark exposure. OLS and 2SLS arms are shown together; shaded bands denote 95 per cent confidence intervals computed with Conley standard errors at a 300-mile cutoff.

for four years, reaching 0.68 percentage points at $h = 4$ with a 95 per cent interval of $[-1.15, -0.21]$, and then recovers about half of that: income remains 0.29 points below its counterfactual level at $h = 8$, with an interval of $[-0.54, -0.04]$. The recovery is real and it is incomplete. The positive annual responses at $h = 5$ and $h = 7$ are what closes the gap, so the county does grow its way back toward the path, but eight years on it has not reached it and the shortfall is still distinguishable from zero. A county recovers its growth rate long before it recovers its level, and on this evidence it does not fully recover the level inside the window the design can see. The cumulative path is estimated directly at each horizon rather than by summing the annual one, since a summed path carries no valid standard error (Section [Appendix F.1](#)).

Panel fixed-effects local projections carry Nickell bias whether or not the specification includes lagged outcomes ([Mei et al., 2025](#)), and the split-panel jackknife correction moves the direct effect further from zero rather than toward it (Section [Appendix F.2](#)). The dynamics reported here are therefore a lower bound on the corrected path.

5.3. *What Identifies the Estimate*

The 5 per cent threshold changes which failures supply the identifying variation, not the estimand, which remains the effect of total failure exposure. Exposure is non-zero in 2,217 county-years and at least one failure clears the threshold in 806 of them. In the remaining 1,411 every failing bank held more than 5 per cent of its deposits in the county, the configuration in which local conditions and bank distress are hardest to tell apart. The design does not need to know what caused those failures, only which failures the county cannot have caused.

The gap between ordinary least squares and two-stage least squares admits two readings. Either instrumenting removes bias, or it changes the population.

Under the first, the most locally concentrated failures should carry the largest estimates, since a failure that follows local decline is credited with the decline. Figure 4 shows they carry the smallest: averaged across the four bins above the threshold and weighted by treated county-years they give 0.107 percentage points, against a pooled ordinary least squares estimate of 0.11 and an exogenous-bin estimate of 0.223. The gradient has the wrong sign for contamination. The bins do differ in when their exposure falls: 70.8 per cent of the exogenous bin’s exposure lands in crisis years against 44.5 per cent in the most concentrated, declining monotonically in between, so the instrument reweights toward the episodes Section 6.1 identifies the neighbour effect in.

Figure 5 makes the same comparison continuous, re-estimating at every rung. The estimate moves away from zero as the threshold tightens and reaches roughly twice the




Figure 4: The Effect by Bin of the County's Share of the Failing Bank *Notes*: The per-unit effect estimated separately in disjoint bins of the county's share of the failing bank's deposit base, so the bins partition the failures rather than nesting them. The five bins are at most 5 per cent, 5 to 25, 25 to 50, 50 to 75 and above 75, holding 806, 541, 197, 162 and 849 treated county-years. The exogenous bin is the one the headline instrument uses. The estimated effect is concentrated there and is small and unsystematic across the rest, and the share of each bin's exposure falling in crisis years declines monotonically as the county's share of the bank rises. Estimates are scaled to median non-zero exposure, so the vertical axis reads in percentage points of annual income growth for a county at the benchmark.

ordinary least squares value at 5 per cent, with most of the movement at the first rung, where excluding the banks that hold almost all their deposits in one county moves the direct effect from 0.120 to 0.151. Every rung is a valid local average for a different population of failures, so the choice of threshold is an economic question rather than a statistical one. Panel C reports the mix behind the retained exposure, which shifts toward the largest banks as the threshold tightens.⁴

Figure 6 re-estimates on rolling nine-year windows. The direct effect is negative in every one of the twelve centred on 2005 to 2016 and significantly so in ten, and no window between 1997 and 2004 carries enough identifying variation to be plotted.⁵ Panel C shows why, with the counties entering the instrument concentrated in the later wave, because failures in the earlier years are overwhelmingly of the single-county banks the threshold excludes. The episodic identifying variation is therefore a property of the object rather than a defect of the design.

Mechanical explanations for the gap fall short. Reverse causality predicts the opposite ordering across bins. Pre-failure deposit flight is bounded by the timing ladder (Section Appendix A). And interference beyond the first contiguity ring cannot explain it, since both columns carry the SLX term and the direct effect holds across eight topologies including the failed bank’s own branch network.

5.4. *Is the Identifying Variation Clean?*

The exclusion restriction requires that exogenous exposure not be preceded by a decline in income. Table G.14 enters exposure at $t+1$ through $t+4$ alongside the contemporaneous term and asks whether failures that have not yet happened predict income growth today, once on realised exposure and once on the instrument, and Figure 7 plots the two against one another. The instrument’s leads are the noisier of the two, with standard errors up to three times the treatment’s.

Realised exposure fails, in the direction the design predicts. On the direct margin its two nearest leads are -0.020 and -0.018 , both significant at 5 per cent, and both carry the sign of the post-failure effect: counties whose banks fail next year are already growing

⁴Banks are ranked by total assets within year. Tightening to ten per cent retains 44 per cent of exposure at the largest quintile against 6 per cent at the fourth and almost nothing below; institutions with no matched quintile survive at 78 per cent, so they are large banks too.

⁵A window is plotted only where it holds 100 treated county-years, and the two-stage arm only where it also holds 64 with a non-zero instrument; two-stage estimates with a first-stage F below 10 are withheld as well. All three floors keep the figure legible rather than making a claim about identification. Twenty windows clear them on the two-stage arm, eight centred on 1988 to 1995 and twelve on 2005 to 2016; the ordinary least squares arm, facing only the treated floor, is drawn on twenty-seven. Windows centred between 1997 and 2004 hold between two and twenty instrumented county-years against 623 in the 2008 window, and the 1996 window clears both support floors only to return a first stage of 0.58.




Figure 5: Effects and Bank Mix as the Threshold Is Tightened

Notes: Panels A and B report the direct and indirect effects as the threshold on the county's share of the failing bank's deposit base is tightened, so that progressively more locally concentrated failures are excluded from the instrument. Both panels are the headline specification of Table 4, column (6), carrying the spatial lag of the treatment, so the two series are the same two objects the baseline table reports. The rung at $\tau = 100\%$ is the unrestricted benchmark: no failure is excluded from the instrument, the instrument reproduces the treatment, and the estimate reproduces OLS. Every rung clears a first-stage F of 10; markers are drawn at decade rungs. Panel C reports the composition of the retained exposure by quintile of failing-bank total assets, ranked within year. Estimates are scaled to median non-zero exposure. Shaded bands denote 95 per cent confidence intervals.




Figure 6: Failure Effects Estimated on Rolling Windows

Notes: Panels A and B re-estimate the direct and indirect effects of Table 4, column (6) on rolling windows, plotting each estimate at the window's midpoint; OLS and 2SLS arms are shown together and estimates are scaled to median non-zero exposure. Windows are plotted where they clear the floors given in the text; the ordinary least squares arm faces only the treated floor, since it uses no instrument. Panel C reports the count of counties with a failure in each window, split into those entering the instrument because the county held at most 5 per cent of the failing bank's deposits and the remaining failures, and is the panel that shows which episodes the instrument is identified off.

more slowly today. The instrument does not show that pattern. None of its four direct leads is distinguishable from zero at 5 per cent, none exceeds 0.031 in absolute value against a contemporaneous coefficient of 0.062, and the two furthest are positive, so counties that later receive exogenous exposure were growing faster three and four years earlier, not more slowly.




Figure 7: Pre-Failure Coefficients on Leads of Exposure *Notes:* The figure plots coefficients on *leads of failure exposure*, exposure in years $t+1$ to $t+4$ entered alongside the contemporaneous term, so each coefficient asks whether exposure that has not yet happened predicts income growth today. All four leads enter one regression on one common sample and are tested jointly; the headline specification itself carries no leads. The panel reports the direct margin under the headline specification, plotting the treatment leads and the instrument leads together so the two can be read against one another. Joint tests over the four leads are reported in the text. The vertical axis reports raw coefficients, per one point of exposure, and is not scaled to the benchmark. Shaded bands denote 95 per cent confidence intervals.

The restriction is better put as two single contrasts than as one four-degree-of-freedom test. The first is the mean of the four leads, which asks whether exposure is preceded by a level shift. On the direct margin it is -0.009 (0.006) on realised exposure and 0.009 (0.009) on the instrument, and neither rejects. The second is the slope across the leads, which asks whether anticipation builds as the failure approaches. It is 0.009 (0.004) on realised exposure, rejecting at 5 per cent ($p = 0.020$), and 0.008 (0.008) on the instrument, which does not ($p = 0.29$). The two are orthogonal by construction, so the pattern is unambiguous: realised exposure shows a lead path that steepens toward the failure, and

the instrument shows neither contrast. The joint Wald rejects on the instrument, but four leads of one series are near-collinear, so it can reject on a linear combination of them that does not resemble a pre-trend. The reciprocal condition number is 0.004 on the instrument's direct arm against 0.043 on the treatment's indirect arm, so the joint statistics are not equally well determined, and on the treatment's own direct arm the statistic cannot be computed at all. Section [Appendix G](#) reports all four arms and the balance of the instrument on pre-period characteristics.

A second question is whether failures cascade locally, an earlier wave creating a later one. Figure 8 projects each failure measure on the other. Endogenous failures do not predict exogenous ones at $t+1$, $t+2$ or $t+4$; at $t+3$ the coefficient is positive and significant at 1 per cent, at 0.0046 (0.0018), which is the sign a cascade would carry. Its size rules it out: a county carrying the benchmark 5.36 points of endogenous exposure acquires 0.025 points of exogenous exposure by $t+3$, which at the estimated coefficient is a thousandth of a percentage point of income growth.




Figure 8: Failure Cascades Between the Exogenous and Endogenous Margins

Notes: Each series is a local projection of one failure measure on another over horizons t to $t+4$, on the specification of Table 4, column (6). The panel of interest is the projection of exogenous exposure on lagged endogenous exposure: under the exclusion restriction it should be flat. The contemporaneous horizon is partly definitional, since exogenous failures are a subset of all failures, so the informative horizons are $t+1$ onward. Both axes are in raw coefficient units, one failure measure per unit of the other, and are not scaled to the benchmark.

6. Mechanism: Where the Loss Lands

6.1. When the Spillover Appears

Table 5 interacts each margin with an indicator for the crisis years: the Savings and Loan crisis, 1987 to 1992, the Global Financial Crisis, 2007 to 2010, and the 2023 episode. Each interaction is instrumented with the corresponding interaction of exogenous exposure. The crisis main effect is not entered because county and state-year effects make it a linear combination of the fixed effects.

Both margins have to be interacted. Columns (2) and (4) interact own exposure alone and send the direct effect to zero outside crisis years. Column (5) interacts both and the picture inverts. Own exposure carries -0.028 outside crisis years and its interaction is indistinguishable from zero, while neighbour exposure carries 0.0009 outside and -0.216 inside. At each margin's benchmark the direct effect is 0.15 percentage points outside a crisis against 0.30 inside, and the indirect effect nothing outside against 0.30 inside. All four first stages clear conventional thresholds.

The heterogeneity is therefore on the neighbour margin, and the crisis-specificity of the direct effect in columns (2) and (4) is an artefact of interacting one treatment of two. A county absorbs its neighbours' failures in normal times and cannot in a crisis. The failures arriving in those episodes are also the geographically diversified ones, so calendar and composition cannot be separated here.

6.2. Where the Loss Lands

Figure 9 decomposes the aggregate response on the two BEA measures of county income. Each component enters as its contribution to the growth of its side total, so the components sum to the aggregate change at every horizon by construction. It is priced at mean non-zero exposure, 11.23 points, because it splits one impact across components rather than reporting an effect to be read against the baseline.

On the residence side the loss is net earnings and almost nothing else. Net earnings carry 90 per cent of the impact contraction, asset income 13, and transfers run the other way by 3, so the three close on the side total. Whatever cushioning the public system provides after a bank failure, it is not a transfer response in these counties.

The place-of-work side locates the loss differently. In point estimates the whole contraction sits on proprietors' income, which carries the side total on its own, while wages and wage supplements are flat and faintly positive. On the neighbour margin the ordering reverses: wages and wage supplements fall and proprietors' income does not. Inside the county the loss is not on payrolls, which is where a demand shock would land first;

Table 5: The crisis interaction, one margin and both

Outcome lags ($t-1$ to $t-3$) Margins interacted with the crisis window	Personal Income Growth _{c,t} (%)				
	In None (1)	In Own (2)	Out None (3)	Out Own (4)	Out Both (5)
Panel A: Coefficient estimates					
Failed Deposit Ratio, $x_{c,t}$	-0.034** (0.017)	0.004 (0.012)	-0.042** (0.016)	-0.006 (0.012)	-0.028*** (0.006)
Failed Deposit Ratio SLX, $Wx_{c,t}$	-0.083*** (0.031)	-0.091*** (0.032)	-0.072** (0.029)	-0.079*** (0.030)	0.0009 (0.013)
$x_{c,t} \times$ Crisis window		-0.083** (0.037)		-0.079** (0.036)	-0.029 (0.032)
$Wx_{c,t} \times$ Crisis window					-0.216*** (0.047)
Panel B: Effects at benchmark exposure, column (5) (pp)					
Direct, outside crisis windows	—	—	—	—	-0.147*** (0.034)
Direct, inside crisis windows	—	—	—	—	-0.303* (0.166)
Indirect, outside crisis windows	—	—	—	—	0.001 (0.019)
Indirect, inside crisis windows	—	—	—	—	-0.303*** (0.073)
Panel C: Model summary and identification					
Observations	112,176	112,176	122,949	122,949	122,949
RMSE	4.059	4.061	4.267	4.268	4.270
Wald, x	4,658.8	5,764.6	4,835.7	5,903.1	7,786.8
Wald, Wx	135.2	191.8	118.8	154.4	760.0
Wald, $x \times$ Crisis		1,050.8		1,041.0	1,894.8
Wald, $Wx \times$ Crisis					192.8
Wu-Hausman	4.226	6.282	5.063	6.143	9.047
Panel D: Fixed effects and errors					
County FE	✓	✓	✓	✓	✓
State $\times$ Year FE	✓	✓	✓	✓	✓
Standard errors	Conley (1999), 300-mile cutoff				

Notes: Two-stage least squares throughout; crisis windows are 1987–1992, 2007–2010 and 2023. Each interacted treatment is a further endogenous regressor instrumented by the corresponding interaction of exogenous exposure, so every column is exactly identified. The crisis main effect is not entered because county and state-by-year effects render it a linear combination of the fixed effects. Columns (1) and (2) add three lags of the outcome and run on the lag-complete sample; columns (3) to (5) carry the headline control set, and column (3) reproduces Table 4, column (6). Panel B prices column (5) at each margin's median non-zero exposure, 5.36 points for x and 1.41 for Wx , with the inside-crisis rows formed from the full coefficient covariance. Conley (1999) spatial standard errors, 300-mile cutoff, in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.




Figure 9: Decomposition of the Income Response

Notes: Panel A decomposes the response of personal income measured by place of residence into net earnings, dividends, interest and rent, and transfers. Panel B decomposes the response of earnings measured by place of work into wages, wage supplements and proprietors' income. Each series is a component's contribution to the growth of its side total, so the components sum to the total at every horizon by construction. Both panels follow the specification of Table 4, column (6) and carry the same spatial-lag controls. The decomposition is priced at mean non-zero exposure, 11.23 points, rather than at the median used elsewhere, so its axis is roughly twice the scale of the baseline exhibits.

across the boundary it is. That the neighbour margin runs through payrolls and that it appears only in crisis windows are the same fact from two directions, since a regional labour market absorbs the displaced in normal times and cannot when every employer is contracting at once. The design separates the two margins. It does not identify the channel on either.

7. Conclusion

This paper estimates the local real effects of U.S. bank failures on a design that treats a county’s own exposure and its neighbours’ separately. The spatial structure produces an error that runs in two directions at once, and its sign is known before any estimate is run. Banks expand into markets adjacent to those they already serve, so branch networks are contiguous and a county’s own exposure moves with its neighbours’. A specification without the SLX term therefore overstates the county-level effect by two-fifths, 0.059 against 0.042 per point of exposure, while dropping the 0.072 that arrives from neighbours’ exposure altogether. The point carries beyond bank failure to any shock transmitted through a firm’s geographic footprint.

A second finding concerns which failures identify the effect rather than how large it is. The design turns on an asymmetry between two quantities usually treated as one: a bank’s economic weight in a county and that county’s weight in the bank need not move together, and failures where the second is small are failures a county cannot plausibly have caused but is still exposed to. Instrumenting on those failures moves the estimate away from zero rather than toward it, doubling the direct effect from 0.021 to 0.042 percentage points per point of exposure. Reverse causality predicts the opposite. Estimated by bin of the county’s share of the failing bank, the profile carries no gradient at all: the most locally concentrated failures do not carry the larger effects contamination predicts, and no bin is distinguishable from the exogenous one.

What the estimand is follows from what makes a failure exogenous. A bank a county cannot bring down is one for which that county is a small part, and those failures arrive disproportionately in systemic episodes. Most failures are not of that kind, which is why work estimating on realised failures is dominated numerically by the locally concentrated ones.

A third finding concerns when the neighbour effect exists at all. Interacting both margins with the crisis windows, the indirect effect is a precise zero outside them and 0.30 percentage points inside, while the direct effect does not detectably differ across the two. A county absorbs its neighbours’ failures in normal times and cannot in a crisis, and

the failures arriving in those episodes are also the ones for which the county was a small part, so calendar and composition cannot be separated here.

The loss is persistent in a specific sense. The annual growth response is carried by the failure year and the year after, at 0.22 and 0.23 percentage points at benchmark exposure, after which the growth rate returns to its path and briefly overshoots it. The level does not follow. The shortfall deepens to 0.68 points by the fourth year and recovers only about half of that by the eighth, where it stands at 0.29 points and is still distinguishable from zero. A county recovers its growth rate long before it recovers its level, and within the eight years this design can see it does not recover the level in full. Panel fixed-effects local projections carry Nickell bias whether or not the specification includes lagged outcomes, and the split-panel jackknife moves the estimates further from zero, so these magnitudes are a lower bound.

Where the loss lands is clearer on the residence side, where net earnings carry almost the whole of the direct contraction and transfers do not move at all, than on the production side. The design separates the two margins cleanly; it does not identify the channel on either, and why the failures the instrument uses cost more per point of exposure is left open.

A county loses income when its neighbours carry a failure, and in a systemic episode that loss is as large as the one its own exposure causes. A response scaled to the exposed county alone is therefore too large for that county and reaches too few of the others, and the geography it should cover is known from the branch network before any estimate is run. The level shortfall is still open at eight years, so the horizon over which the cost is judged is longer than the one over which it is usually measured.

Appendix A. Measurement

This appendix prices the measurement choices footnoted in Section 4. Figure A.10 re-estimates the dynamic response under each scaling of exposure, including the income scaling standard in this literature, and Figure A.11 moves the measurement vintage back one year at a time to bound pre-failure deposit erosion.

The choice of denominator changes the size of the shock being measured, not the shape of the response. On impact the instrumented direct effect is about a quarter of a percentage point under all three scalings, and in each the response recovers to zero by roughly $t+3$ and turns positive thereafter. The indirect effect is larger per point of exposure under all three scalings; at each margin's own median the indirect effect is the smaller of the two at impact. The ordinary least squares panels are uniformly smaller and recover sooner, the same gap the baseline reports, present under every scaling.

Nothing in the estimates depends on the measurement vintage. Across the four snapshots the instrumented direct effect moves only between -0.22 and -0.24 percentage points and the indirect between -0.098 and -0.104 ; under ordinary least squares the direct effect sits at -0.11 throughout and the indirect at -0.05 . The spread across rungs is an order of magnitude smaller than the confidence interval at any one of them. This is what bounds the deposit-flight concern: depositors leaving ahead of a failure understate the recorded deposit stock most severely at the snapshot closest to the event, so if flight were driving the measure the rungs would differ systematically. They do not.

Appendix A.1. Alternative Exposure Measures

Table A.6 re-estimates the headline two-endogenous specification under three denominators: the county deposit base, which is the specification of record; lagged county personal income; and the county branch count. The direct effect is all but identical under the deposit and branch scalings, at -0.042 against -0.040 , and roughly twice as large under the income scaling, while the indirect effect exceeds the direct effect under all three. Because the income and deposit denominators share no common measurement error with each other, a result that survives both is not division bias in either.




Figure A.10: Failure Exposure Under Alternative Denominators

Notes: Income growth responses over eight years when failure exposure is scaled three ways: by the county deposit base (the paper's measure), by county personal income (the literature's), and by county branches. Left column OLS, right column 2SLS. Each series is evaluated at the median non-zero value of its own denominator's own-county exposure, so impact magnitudes are comparable across rows. The income panels use a wider vertical scale than the other four.




Figure A.11: Measurement-Timing Ladder *Notes:* The direct and indirect effects when the failing bank's deposits, and the county base that scales them, are measured at successively earlier snapshots, from the failure year to three years before it. Each rung is the same failure event on an older balance sheet. Panel A is 2SLS, Panel B is OLS. The paper's estimates use the pre-failure snapshot rung, which reproduces Table 4, columns (6) and (4). Effects are in percentage points of annual income growth for a county at the median non-zero exposure. Bars are 95 per cent confidence intervals.

Table A.6: OLS and IV estimates under alternative failure-exposure measures

Exposure measure	Personal Income Growth _{c,t} (%)		
	Failed Deposit Ratio (1)	Failed Deposit-to- Income Ratio (2)	Failed Branch Ratio (3)
Panel A: Coefficient estimates			
Failed Deposit Ratio, $x_{c,t}$	-0.042** (0.016)	-0.097** (0.043)	-0.040** (0.016)
Failed Deposit Ratio SLX, $Wx_{c,t}$	-0.072** (0.029)	-0.193*** (0.044)	-0.070*** (0.023)
Panel B: Model summary and identification			
Observations	122,949	122,949	122,949
RMSE	4.2671	4.2703	4.2673
Wald (excluded instruments), x	4,835.7	271.1	2,820.0
Wald (excluded instruments), Wx	118.8	62.9	165.4
Wu-Hausman	5.063	9.208	6.816
Panel C: Fixed effects and errors			
County FE	✓	✓	✓
State × Year FE	✓	✓	✓
Standard errors	Conley (1999), 300-mile cutoff		

Notes: Each column is the headline two-endogenous 2SLS specification re-estimated with a different scaling of failure exposure; the dependent variable is the same throughout. Column (1) is the specification of record and reproduces Table 4, column (6). Both x and Wx are instrumented with the corresponding 5% marginal-market measures, z and Wz . W denotes queen-contiguity neighbour averages. Mean outcome 2.2161. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Appendix B. Spatial Clustering of Failure Exposure

This appendix maps the Failed Deposit Ratio across counties, one decade at a time, in Figure B.12, Figure B.13, Figure B.14, Figure B.15 and Figure B.16. All five share the colour scale used in Figure 1, so shading is comparable across decades. Exposure falls in contiguous blocks of counties in each decade rather than in isolated counties, and the clustering intensifies across the sample. This is the descriptive basis for treating spatial dependence and interference as first-order features of the design.




Figure B.12: Mean Failed Deposit Ratio by County, 1981 to 1989

Notes: County exposure is the deposits of banks failing in each year scaled by the county deposit base, constructed from pre-failure branch footprints and averaged over the period. The colour scale is common to all five decade maps and to Figure 1.

Appendix B.1. Spatial Dependence

Branch networks span clusters of adjacent counties, and local labour and credit markets cross county boundaries, so county income growth plausibly inherits spatial correlation even absent failure exposure. Table B.7 reports a ring correlogram: Moran's I for county personal income growth computed separately within each 50-kilometre distance ring, so that each entry prices dependence at that range alone rather than cumulatively. Moran's I is 0.35 in the nearest ring, declines monotonically with distance, becomes statistically indistinguishable from zero between 400 and 450 kilometres, and turns significantly




Figure B.13: Mean Failed Deposit Ratio by County, 1990 to 1999
Notes: As in Figure B.12.




Figure B.14: Mean Failed Deposit Ratio by County, 2000 to 2009
Notes: As in Figure B.12.




Figure B.15: Mean Failed Deposit Ratio by County, 2010 to 2019

Notes: As in Figure B.12. This map is reproduced in the main text as Figure 1.




Figure B.16: Mean Failed Deposit Ratio by County, 2020 to 2023

Notes: As in Figure B.12. The final period covers four years rather than a full decade, so its average is taken over a shorter window.

Table B.7: Spatial autocorrelation in county income growth: ring correlogram of Moran's I by distance band

Distance ring (km)	Moran's I	Links	z	p
0–50	0.3515	11,836	22.04	< 0.001
50–100	0.2821	39,886	36.72	< 0.001
100–150	0.2001	64,066	31.67	< 0.001
150–200	0.1422	86,894	25.09	< 0.001
200–250	0.0840	107,626	18.04	< 0.001
250–300	0.0563	126,868	11.78	< 0.001
300–350	0.0239	144,970	5.89	< 0.001
350–400	0.0069	161,260	2.04	0.025
400–450	0.0047	177,058	1.58	0.111
450–500	−0.0155	189,990	−3.78	< 0.001
500–550	−0.0295	202,738	−7.59	< 0.001
550–600	−0.0395	213,600	−11.08	< 0.001

Notes: Ring correlogram for county real personal income growth. Within each ring, counties are neighbours if their great-circle centroid distance falls inside the stated band, so the rings are disjoint and each row measures dependence at that range alone rather than cumulatively out to it. Weights are row-standardised binary contiguity within the ring. *Links* is the number of county pairs entering the statistic. The statistic is computed on the 2010 cross-section of 3,012 counties, and the z statistic and p -value are set against a bootstrap null of 199 draws; p -values below 0.001 are reported as < 0.001. Source: 0_02_03_ring_correlogram.

negative beyond that. That is the signature of a regional structure with a finite range rather than a uniform national correlation. Dependence at the ranges the design cares about is therefore economically meaningful and spatially bounded, which motivates treating it as first-order: the SLX term in the specification and distance-based inference rather than an assumption of independence across neighbouring counties.

Appendix C. First Stage Results

Table C.8 reports the first stages behind columns (5) and (6) of Table 4. Exogenous exposure enters own-county exposure almost one for one, at 0.988 in the one-endogenous design and 0.936 in the two-endogenous one, which is what the construction implies. The cross terms are the informative entries: exogenous neighbour exposure predicts realised neighbour exposure at 1.09, while own exogenous exposure enters that equation at 0.058, so each margin is identified largely by its own instrument. Wald statistics on the excluded instruments are 4,836 for own exposure and 119 for the neighbour margin.

Table C.8: First-stage estimates: exogenous-failure exposure and the endogenous regressors

Second-stage column of Table 4	(5)	(6)	
Dependent variable	$x_{c,t}$ (a)	$x_{c,t}$ (b)	$Wx_{c,t}$ (c)
Panel A: First-stage coefficients			
Exog. Failed Deposit Ratio, $z_{c,t}$	0.988*** (0.017)	0.936*** (0.019)	0.058** (0.025)
Exog. Failed Deposit Ratio SLX, $Wz_{c,t}$		0.317 (0.195)	1.09*** (0.092)
Panel B: Model summary			
Observations	122,949	122,949	122,949
Dependent variable mean	0.203	0.203	0.199
RMSE	2.302	2.296	1.089
Wald (excluded instruments)	3,462.7	4,835.7	118.8
Panel C: Fixed effects			
County FE	✓	✓	✓
State-Year FE	✓	✓	✓
Standard Errors	Conley (1999), 300-mile cutoff		

Notes: Each column reports the first stage of the corresponding 2SLS column of Table 4, estimated on the same sample and with the same fixed effects; the specification carries no covariates beyond the fixed effects, and standard errors are extracted from the same fitted objects. x is the Failed Deposit Ratio; z is exposure to exogenous failures, defined as failures of banks holding at most 5% of their deposits in the county; W denotes the queen-contiguity neighbour average. Column (a) is the single first stage of the one-endogenous-regressor specification; columns (b) and (c) are the two first stages of the specification with the SLX term, each on both instruments. Wald statistics test the excluded instruments equation by equation under the same variance estimator as the reported standard errors. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix D. Inference

The standard errors in Section 4 are Conley-type spatial HAC estimates; this appendix asks whether the reported significance survives alternative inference assumptions. Figure D.17 redraws the headline responses under each estimator and Table D.9 prices them on the headline coefficients; Figure D.18 reports the weak-identification-robust confidence set and Figure D.19 the randomisation placebo.




Figure D.17: Standard-Error Ladder

Notes: The direct and indirect effects at the benchmark under seven variance estimators, with the point estimates held fixed so that only the interval moves. The rungs are Conley at 100, 300 and 600 miles and clustering on county, county and year, year, and crisis episode. The direct effect remains significant on every rung; the indirect effect loses conventional significance under the coarsest clusterings. Estimates are scaled to median non-zero exposure, so the vertical axis reads in percentage points of annual income growth for a county at the benchmark. The Conley and cluster-by-county rungs are not independent checks: the Conley estimator sums scores within each coordinate before applying its kernel, so it nests clustering on the county.

Figure D.18 inverts the Anderson–Rubin statistic over the two coefficients jointly, so the set is valid whatever the first stage delivers. At benchmark exposure the projected intervals are $[-0.42, -0.01]$ on the direct margin and $[-0.22, -0.02]$ on the indirect, against a Conley interval of $[-0.54, -0.11]$ for their sum. The set is bounded, and the robust interval is about a third wider than the Conley one and contains it.

Figure D.19 reassigns exposure at random, so the true effect is zero by construction

Table D.9: Standard errors under alternative variance estimators

Variance estimator	Personal Income Growth _{c,t} (%)						
	Conley spatial HAC			Clustered			
	100mi	300mi	600mi	County	County × Year	Year	Crisis episode
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Coefficient estimates							
Failed Deposit Ratio, $x_{c,t}$	−0.042*** (0.014)	−0.042** (0.016)	−0.042*** (0.014)	−0.042*** (0.014)	−0.042*** (0.014)	−0.042*** (0.014)	−0.042** (0.016)
Failed Deposit Ratio SLX, $Wx_{c,t}$	−0.072** (0.032)	−0.072** (0.029)	−0.072* (0.038)	−0.072*** (0.027)	−0.072 (0.050)	−0.072 (0.050)	−0.072 (0.059)
Panel B: Model summary and identification							
Observations				122,949			
RMSE				4.2671			
Wald (excluded instruments), x				4,835.7			
Wald (excluded instruments), Wx				118.8			
Wu–Hausman				5.063			
Panel C: Fixed effects							
County FE	✓	✓	✓	✓	✓	✓	✓
State × Year FE	✓	✓	✓	✓	✓	✓	✓

Notes: One fitted model, re-summarised under seven variance estimators, so the point estimates, sample and first-stage statistics are identical across columns by construction and are reported once. Column (2) is the headline. Crisis-episode clusters are the windows defined in the text. The indirect coefficient loses conventional significance under the coarsest clusterings, which is the expected direction: coarser clusters mean fewer of them, and the indirect margin has less independent variation to spend. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.




Figure D.18: Anderson–Rubin Confidence Set *Notes:* The weak-identification-robust confidence region over the direct and indirect coefficients jointly, with projected intervals for each and for their sum. The set is constructed by inverting the Anderson–Rubin statistic on a grid and is valid whatever the strength of the first stage. Axes are in raw coefficient units. At benchmark exposure the projected intervals are $[-0.42, -0.01]$ for the direct effect, $[-0.22, -0.02]$ for the indirect effect and $[-0.61, -0.06]$ for their sum, against a Conley interval of $[-0.54, -0.11]$ for the sum. The set is bounded, since no accepted grid cell lies on an edge, and the weak-identification-robust interval is about a third wider than the Conley one and contains it.

and the share of draws rejecting at 5 per cent is the empirical size of the test. Neither arm attains its nominal level: the Gaussian-field draws reject 7.5 per cent of the time, 75 of 1,000, and the year-permutation draws 10 per cent, 20 of 200, with both intervals excluding 5 per cent. Both arms fit on ten sampled years, about 29,000 county-years against the headline's 122,949, so these are sizes at a quarter of the estimation sample rather than the size of the reported test.




Figure D.19: Permutation Placebo

Notes: The distribution of the headline estimate when exposure is permuted across counties, so the true effect is zero by construction. The vertical line marks the estimate on the true assignment. The exercise reports the empirical size of the test rather than its power. The horizontal axis is in raw coefficient units.

The large-cluster ladder in Table D.10 re-estimates the headline specification under progressively coarser clusterings. The direct effect stays significant on every rung. The indirect effect reaches significance only at the two finest partitions, where the critical value falls with the degrees of freedom rather than the standard error falling.

Table D.10: Large-cluster standard errors

	Personal Income Growth _{c,t} (%)				
	(1)	(2)	(3)	(4)	(5)
Panel A: Coefficient estimates					
Failed Deposit Ratio, $x_{c,t}$	−0.042* (0.016)	−0.042** (0.013)	−0.042** (0.012)	−0.042*** (0.012)	−0.042** (0.016)
Failed Deposit Ratio SLX, $Wx_{c,t}$	−0.072 (0.034)	−0.072 (0.045)	−0.072 (0.044)	−0.072* (0.038)	−0.072* (0.035)
Panel B: Model summary and identification					
Clusters, G	4	6	8	12	20
Observations	122,949	122,949	122,949	122,949	122,949
Wald (excluded instruments), x	4,835.7	4,835.7	4,835.7	4,835.7	4,835.7
Wald (excluded instruments), Wx	118.8	118.8	118.8	118.8	118.8
Panel C: Fixed effects					
County FE	✓	✓	✓	✓	✓
State-Year FE	✓	✓	✓	✓	✓

Notes: Bester–Conley–Hansen large-cluster standard errors from k -medoids partitions of county centroids, on the headline sample. Columns differ only in the number of clusters the partition returns, reported in the G row; the point estimates are those of Table 4, column (6), and are held fixed across columns. Inference is a t with $G - 1$ degrees of freedom, so the critical values — and hence the stars — differ by column: a coefficient can carry more stars at $G = 20$ than at $G = 4$ on an identical estimate and a similar standard error. Validity requires the partitions to be approximately independent, which contiguous blocks in a nationally synchronised crisis satisfy only approximately; column (1), at four clusters, should be read as illustrative. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix E. Weight Structure

This appendix answers the third threat in Section 4, that the spatial structure is wrong, and carries the fixed-effects saturation evidence for the second. Figure E.20, Table E.11 and Table E.12 trace the estimates across saturations; Figure E.21 and Table E.13 vary the weight matrix in kind rather than in bandwidth, and Figure E.22 traces the dynamic response under each.

Queen contiguity is selected rather than assumed. Among the operators exogenous to the failure it is the only contiguity relation defined for every county in the panel, and all seven are estimated on the same 122,949 county-years; the bank-network operator runs on 108,940, since the remainder carry no network link. It fits better still but is endogenous to the shock the design identifies from, which is why it appears in Table E.13 to be argued against rather than chosen.




Figure E.20: Effects Across Fixed-Effect Saturations

Notes: The direct and indirect effects as the time dimension is saturated from no effects through year effects to state-by-year effects, with county effects throughout and OLS and 2SLS arms shown separately. The series with and without the SLX term are plotted together, so the figure shows how much of the attenuation in Table 4 is the fixed effects and how much is the SLX term. The vertical axis reports raw coefficients, per one point of exposure, and is not scaled to the benchmark.

Table E.11: Fixed-effects ladder, OLS

Fixed effects	Personal Income Growth _{c,t} (%)				
	None (1)	County (2)	County + Year (3)	County + Div.-Year (4)	County + State-Year (5)
Panel A: With the SLX term					
Failed Deposit Ratio, $x_{c,t}$	-0.027*** (0.008)	-0.034*** (0.008)	-0.034*** (0.006)	-0.023*** (0.005)	-0.021*** (0.005)
Failed Deposit Ratio SLX, $Wx_{c,t}$	-0.087*** (0.025)	-0.113*** (0.022)	-0.103*** (0.014)	-0.048*** (0.010)	-0.037*** (0.006)
Constant	2.24*** (0.117)				
Observations	122,991	122,991	122,991	122,991	122,949
RMSE	5.534	5.420	4.885	4.660	4.267
Panel B: Without the SLX term					
Failed Deposit Ratio, $x_{c,t}$	-0.042*** (0.011)	-0.054*** (0.011)	-0.051*** (0.007)	-0.029*** (0.005)	-0.025*** (0.005)
Constant	2.22*** (0.116)				
Observations	122,991	122,991	122,991	122,991	122,949
RMSE	5.535	5.422	4.887	4.660	4.267

Notes: Ordinary least squares at five fixed-effect saturations; “Div.-Year” is Census division by year. Panel A carries the queen-contiguity SLX term, Panel B omits it, so reading down a column shows what omitting the SLX term does to the own-county coefficient at that saturation. Column (5) is the headline saturation and reproduces Table 4, columns (4) and (3). Conley (1999) spatial standard errors, 300-mile cutoff, in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table E.12: Fixed-effects ladder, 2SLS

Fixed effects	Personal Income Growth _{c,t} (%)				
	None	County	County + Year	County + Div.-Year	County + State-Year
	(1)	(2)	(3)	(4)	(5)
Panel A: With the SLX term					
Failed Deposit Ratio, $x_{c,t}$	-0.046*** (0.013)	-0.052*** (0.018)	-0.037*** (0.012)	-0.042*** (0.015)	-0.042** (0.016)
Failed Deposit Ratio SLX, $Wx_{c,t}$	-0.103* (0.057)	-0.142** (0.060)	-0.075*** (0.026)	-0.077*** (0.027)	-0.072** (0.029)
Constant	2.25*** (0.120)				
Observations	122,991	122,991	122,991	122,991	122,949
RMSE	5.534	5.420	4.885	4.660	4.267
Wald, x	5,231	4,558	5,896	9,175	4,836
Wald, Wx	293.1	285.6	305.8	217.3	118.8
Wu-Hausman	3.042	4.854	1.355	4.970	5.063
Panel B: Without the SLX term					
Failed Deposit Ratio, $x_{c,t}$	-0.087*** (0.032)	-0.109*** (0.038)	-0.066*** (0.013)	-0.067*** (0.019)	-0.059*** (0.020)
Constant	2.23*** (0.118)				
Observations	122,991	122,991	122,991	122,991	122,949
RMSE	5.537	5.424	4.888	4.661	4.268
Wald, x	1,852	1,706	2,538	2,799	3,463
Wu-Hausman	11.72	17.15	1.414	9.034	7.446

Notes: Two-stage least squares at the same five saturations as Table E.11. Column (5) is the headline and reproduces Table 4, columns (6) and (5). Wald statistics test the excluded instruments equation by equation under the reported variance estimator. Conley (1999) spatial standard errors, 300-mile cutoff, in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.




Figure E.21: The Indirect Effect Under Alternative Weight Structures *Notes:* The indirect effect estimated under weight matrices that differ in kind rather than in bandwidth: queen contiguity, a distance band, *k*-nearest neighbours, Delaunay triangulation and a gravity operator. The vertical axis reports raw coefficients, per one point of exposure, and is not scaled to the benchmark, so magnitudes are not comparable across structures; Figure E.22 prices each at its own benchmark.

Table E.13: The indirect effect under alternative weight structures

Neighbour operator	Personal Income Growth _{c,t} (%)							
	Averaging (row-standardised)					Sum		Network
	Queen contig. (1)	Delau- nay (2)	<i>k</i> near. neigh. (3)	Income gravity (4)	Dist. band (5)	Binary sum (6)	Inv. dist. (7)	Bank network (8)
Panel A: Coefficient estimates								
Failed Deposit Ratio, $x_{c,t}$	-0.042** (0.016)	-0.041** (0.017)	-0.035** (0.015)	-0.044** (0.018)	-0.055*** (0.017)	-0.055*** (0.018)	-0.037** (0.016)	-0.055*** (0.018)
Failed Deposit Ratio SLX, $Wx_{c,t}$	-0.072** (0.029)	-0.084** (0.033)	-0.243*** (0.065)	-0.464*** (0.171)	-0.496* (0.254)	-0.003 (0.002)	-1.05* (0.562)	-1.16* (0.595)
Panel B: Model summary and identification								
Observations	122,949	122,949	122,949	122,949	122,949	122,949	122,949	108,940
RMSE	4.267	4.267	4.266	4.270	4.267	4.268	4.267	4.083
Wald, x	4,836	4,057	4,497	2,037	1,939	1,757	7,022	1,793
Wald, Wx	118.8	498.9	333.0	50.2	111.3	70.0	90.7	1,966
Wu-Hausman	5.063	4.288	5.894	11.22	6.336	15.98	10.64	9.577
Panel C: Fixed effects								
County FE	✓	✓	✓	✓	✓	✓	✓	✓
State × Year FE	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Column (1) is the headline. Averaging weights are row-standardised, so Wx is a weighted *mean* of neighbours' exposure rather than a total, and within a column θ is comparable with the own-county coefficient. That mean is taken over a different neighbourhood in each column, however, and the median non-zero Wx runs from 1.41 under queen contiguity to 0.089 under the gravity operator, so the raw coefficients in columns (1)–(5) are per unit of differently scaled quantities and are not comparable across columns; Figure E.22 prices each at its own benchmark. Sum weights in columns (6)–(7) are not standardised: Wx is a weighted *total*, so a county with more neighbours receives a larger treatment for the same neighbour exposure and θ is per unit of a quantity that scales with the neighbour count — those two magnitudes are not comparable with the rest of the row. The bank-network topology in column (8) is endogenous to the same shock the design identifies from, and is reported to be argued against rather than chosen; its smaller sample reflects counties with no network link. Conley (1999) spatial standard errors, 300-mile cutoff, in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.




Figure E.22: Dynamic Responses Under Alternative Weight Structures *Notes:* The impulse responses of the previous figure traced over horizons, one series per weight structure. Each series is scaled to the median non-zero exposure of its own weight structure, so unlike Figure E.21 the impact points are comparable across structures.

Appendix F. Dynamics

This appendix reports the cumulative counterpart of the annual responses in Section 5.2, and the sensitivity of both to the lag structure and to the Nickell correction.

Appendix F.1. Cumulative Responses

Figure F.23 converts the annual path of Figure 3 into a level shortfall, so a flat segment is a gap that is not being recovered rather than a period of no growth. On the direct margin the shortfall deepens to 0.68 percentage points by $t+4$ and stands at 0.29 at $t+8$, still distinguishable from zero.

Appendix F.2. The Lag Ladder and the Nickell Correction

This appendix reports the dynamic responses under alternative lag structures and under a split-panel jackknife correction for Nickell bias. Figure F.24 plots six panels. Panel A is ordinary least squares without lags. Panel B is the two-stage estimate without lags and is the specification of record. Panels C to E add one, two and three lags of income growth. Panel F applies the jackknife correction to the three-lag specification.

Lags do almost nothing. Panels C, D and E are close to indistinguishable from one another at every horizon on both margins, and they differ from Panel B only in attenuating the direct effect at impact slightly. The lag structure is therefore not what determines the results, which is one reason the baseline carries no lags. The other is that under a serially clustered treatment, lagged income growth is partly an outcome of earlier failure exposure, so conditioning on it absorbs part of the effect being estimated.

Excluding lags does not remove the bias that motivates including them. Mei et al. (2025) show that panel fixed-effects local projections carry Nickell bias whether or not the regressors contain lagged dependent variables, and that it distorts inference as well as the point estimates; their correction is a split-panel jackknife, which Panel F applies to the three-lag specification.

The correction moves the two margins differently. The direct effect at impact rises in magnitude by roughly half but stays inside both the uncorrected interval and the interval of Panel B, so the estimate the paper stands on survives it; the indirect effect close to doubles and does not, which joins the reasons given in Section 5.1 for reporting θ alongside β rather than in place of it. Both remain significant at impact, and in both cases the uncorrected estimates understate the loss, the direction Mei et al. (2025) report across the four crisis studies they revisit. Because the ladder is flat, a correction applied to the specification of record would land close to Panel F and, since the lags attenuate the direct effect first, slightly larger in magnitude, so Panel F is a lower bound. Figure F.25




Figure F.23: Dynamic Effects of Bank Failure Exposure on County Personal Income Level

Notes: The figure reports the cumulative counterpart of Figure 3. The outcome at horizon h is the sum of annual log growth from the failure year through $t+h$, so each point is a log-level gap rather than a growth rate and a flat path indicates a level shortfall that is not recovered. The three panels report the cumulative direct, indirect and total responses. Each horizon is a separate local projection on the specification of Table 4, column (6); post-hoc summation of the annual path is not used, because it carries no valid standard error. Only the 2SLS arm is shown. Estimates are scaled to benchmark exposure. Shaded bands denote 95 per cent confidence intervals.

compares the headline, the three-lag variant and the corrected estimate at impact on a common sample.




Figure F.24: Dynamic Responses Under Alternative Lag Structures and the Nickell Correction
Notes: Each panel plots the direct and indirect responses of county income growth to failure exposure at horizons $h = 0$ to 8 , scaled to median non-zero exposure. Panel A is ordinary least squares without lags. Panel B is two-stage least squares without lags and is the specification of record. Panels C to E add one, two and three lags of income growth. Panel F applies a split-panel jackknife correction for Nickell bias (Mei et al., 2025) to the three-lag specification. Shaded bands are 95 per cent confidence intervals.




Figure F.25: Outcome Lags and the Split-Panel Jackknife

Notes: The direct and indirect effects under the headline specification, under the variant that adds three lags of the outcome, and under the split-panel jackknife correction for Nickell bias. Adding outcome lags barely moves the point estimates on a common sample. The jackknife moves them further, and in the conservative direction, so the headline is a lower bound in magnitude under that correction. The axis reports raw coefficients and is not scaled to the benchmark.

Appendix G. Exclusion Restriction: Leads, Balance, and Cascade

Section 5.4 states the lead and cascade results the exclusion restriction turns on. This appendix carries the supporting material: the full four-arm lead figure (Figure G.26) and the complete lead table (Table G.14), and the balance of the instrument on pre-period covariates (Figure G.27 and Table G.15).

The cascade tests are reported in the body at Figure 8, since the direction that matters for the exclusion restriction, whether locally concentrated failures predict later exogenous ones, is a first-order defence of the design rather than a supplementary check.




Figure G.26: Lead Tests, All Four Arms

Notes: The full version of Figure 7. The four panels cross the direct and indirect margins with the headline specification and its outcome-lag variant, and each plots the treatment leads against the instrument leads. The joint test on the treatment's direct leads cannot be computed under the headline variance estimator, which is not positive semi-definite in that cell, and the joint test on the indirect instrument leads is reported in Section 5.4 with the conditioning caveat that attends it. The vertical axis reports raw coefficients, per one point of exposure.

Table G.14: Lead tests: coefficients on leads of the treatment and of the instrument

Leads of Margin	Personal Income Growth _{c,t} (%)							
	Treatment, x				Instrument, z			
	Direct		Indirect		Direct		Indirect	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Contemporaneous exposure								
Own, t	-0.033*** (0.007)	-0.028*** (0.006)	-0.027*** (0.006)	-0.023*** (0.006)	-0.057*** (0.022)	-0.062*** (0.022)	-0.043** (0.018)	-0.049*** (0.018)
Neighbour average, t			-0.049*** (0.004)	-0.043*** (0.005)			-0.101*** (0.030)	-0.090*** (0.029)
Panel B: Leads								
Lead $t+1$	-0.018** (0.008)	-0.020*** (0.007)	-0.055*** (0.016)	-0.043*** (0.015)	0.011 (0.020)	0.011 (0.018)	-0.052 (0.035)	-0.048 (0.033)
Lead $t+2$	-0.021** (0.009)	-0.018** (0.007)	-0.045*** (0.012)	-0.043*** (0.009)	-0.020 (0.017)	-0.028* (0.016)	-0.046 (0.038)	-0.056 (0.039)
Lead $t+3$	-0.003 (0.013)	-0.004 (0.016)	0.011 (0.017)	0.012 (0.017)	0.035** (0.017)	0.031 (0.019)	0.055*** (0.011)	0.050*** (0.009)
Lead $t+4$	0.008 (0.007)	0.004 (0.007)	-0.003 (0.022)	-0.011 (0.022)	0.019 (0.012)	0.020* (0.012)	0.024 (0.033)	0.030 (0.033)
Panel C: Leads jointly zero								
χ^2	8.23	—	20.90	73.50	26.78	23.09	66.60	33.28
df	4	—	4	4	4	4	4	4
p -value	0.0834	—	0.0003	0.0000	0.0000	0.0001	0.0000	0.0000
Panel D: Model summary and controls								
Observations	98,628	108,487	98,628	108,487	98,628	108,487	98,628	108,487
RMSE	4.011	4.182	4.010	4.181	4.012	4.182	4.011	4.182
Outcome lags ($t-1$ to $t-3$)	✓		✓		✓		✓	
County FE	✓	✓	✓	✓	✓	✓	✓	✓
State $\times$ Year FE	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Reduced-form regressions of income growth on contemporaneous exposure and four leads of it. Odd-numbered columns add three lags of the outcome; even-numbered columns are the headline control set, which carries none. The instrument columns (5)–(8) are the test that matters for identification: the exclusion restriction requires no anticipatory relationship between *exogenous* exposure and prior growth. Every testable arm rejects the joint-zero null: on the direct margin the headline arm (column 6) returns $\chi^2 = 23.09$ ($p = 0.0001$) and the outcome-lag arm (column 5) $\chi^2 = 26.78$ ($p < 0.0001$); the indirect arms reject more decisively still. Because z is a subset of x , its leads are mechanically noisier, so the power of these tests differs across panels. Entries marked “—” are untestable: in column (2) the variance matrix is not positive semi-definite, so no joint statistic can be formed. Conditioning numbers vary sharply across arms — the reciprocal condition number is 0.004 in column (5) and 0.008 in column (6) against 0.043 in column (4) — so the joint statistics are not equally well determined. Coefficients on the outcome lags are reported in Table C.16. Conley (1999) spatial standard errors, 300-mile cutoff, in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.




Figure G.27: Instrument Balance on Pre-Period Characteristics

Notes: Each point is the coefficient from a regression of the instrument on one pre-period county characteristic, standardised so the coefficients are comparable across characteristics. Under the identifying restriction the threshold should not select on the local economy's prior condition. Coefficients are standardised, not scaled to the benchmark.

Table G.15: Instrument balance on pre-period county characteristics

	Dependent variable: $z_{c,t}$		
Predetermined characteristic	Coefficient	Standard error	Observations
	(1)	(2)	(3)
Panel A: Local economic conditions			
Personal Income Growth $_{c,t-1}$	0.0008	(0.0009)	119,361
Neighbour Avg. Income Growth $_{c,t-1}$	-0.002	(0.001)	119,930
Population Growth $_{c,t-1}$	-0.004	(0.003)	119,361
Panel B: Banking structure			
Bank Concentration (HHI) $_{c,t-1}$	$-9.48 \times 10^{-6***}$	(2.91×10^{-6})	122,949
County Deposit Growth $_{c,t-1}$	-0.0001***	(4.82×10^{-5})	119,347
Branch Count Growth $_{c,t-1}$	3.32×10^{-5}	(0.0001)	119,362
Panel C: All six in one regression			
Joint Wald	34.36 on 6 df, $p < 0.0001$		
Panel D: Fixed effects			
County FE		✓	
State $\times$ Year FE		✓	

Notes: Each row of Panels A and B is a separate regression of the instrument on one predetermined county characteristic, with county and state-by-year fixed effects. Panel C reports the joint Wald from a single regression carrying all six at once, which is the test the balance claim rests on. The rejections fall entirely on banking structure, deposit concentration and deposit growth, both of which move mechanically with failure waves. The three characteristics describing the local economy's prior condition — own and neighbour income growth and population growth — are all indistinguishable from zero, and those are the ones the exclusion restriction turns on. A rejection on the first group is the fingerprint of a bad control rather than of a bad instrument. Adding the rejecting characteristics to the second stage leaves the estimate where it was: at benchmark exposure the direct effect is -0.223 in the baseline, -0.222 adding lagged concentration, -0.237 adding neighbour income growth, -0.236 adding both and -0.233 adding the lagged instrument, against a baseline standard error of 0.087. Conley (1999) spatial standard errors, 300-mile cutoff. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table G.16: The headline specification with and without outcome lags

Estimator	Personal Income Growth _{c,t} (%)				
	OLS		2SLS		
	Full	Lag-complete	Full	Lag-complete	
Sample	(1)	(2)	(3)	(4)	(5)
Panel A: Coefficient estimates					
Failed Deposit Ratio, $x_{c,t}$	−0.021*** (0.005)	−0.025*** (0.005)	−0.042** (0.016)	−0.034** (0.017)	−0.040** (0.016)
Failed Deposit Ratio SLX, $Wx_{c,t}$	−0.037*** (0.006)	−0.042*** (0.003)	−0.072** (0.029)	−0.083*** (0.031)	−0.072*** (0.028)
Personal Income Growth _{c,t−1}		−0.181*** (0.043)		−0.181*** (0.043)	
Personal Income Growth _{c,t−2}		−0.050*** (0.017)		−0.050*** (0.017)	
Personal Income Growth _{c,t−3}		−0.056*** (0.016)		−0.056*** (0.016)	
Panel B: Model summary and identification					
Observations	122,949	112,176	122,949	112,176	112,176
RMSE	4.267	4.059	4.267	4.059	4.132
Wald, x			4,835.7	4,658.8	4,684.2
Wald, Wx			118.8	135.2	135.6
Wu–Hausman			5.063	4.226	5.371
Panel C: Fixed effects					
Outcome lags ($t-1$ to $t-3$)		✓		✓	
County FE	✓	✓	✓	✓	✓
State × Year FE	✓	✓	✓	✓	✓

Notes: The headline specification with and without three lags of the outcome. Adding lags costs the first three years of every county's panel, so columns (2), (4) and (5) run on the lag-complete sample; column (5) drops the lags again on that sample, isolating the effect of the specification change from the effect of the sample change. Lagged outcomes combined with county effects bias the within estimator, and in a serially clustered treatment lagged growth is itself partly an outcome of earlier exposure, which is why the headline carries no outcome lags. W denotes queen-contiguity neighbour averages. Conley (1999) spatial standard errors, 300-mile cutoff, in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

References

- 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.
- Amberg, N., Becker, B., 2024. Banking Without Branches. CEPR Discussion Paper 19450. Centre for Economic Policy Research. doi:10.2139/ssrn.4930927.
- 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.
- Ashcraft, A.B., 2006. New evidence on the lending channel. *Journal of Money, Credit and Banking* 38, 751–775. doi:10.2139/ssrn.281629. dOI: 10.2139/ssrn.281629 MAG ID: 2043431305.
- Balla, E., Mazur, L.C., Prescott, E.S., Walter, J.R., 2019. A comparison of community bank failures and fdic losses in the 1986–92 and 2007–13 banking crises. *Journal of Banking and Finance* 106, 1–15. URL: <https://www.sciencedirect.com/science/article/pii/S0378426619300950>, doi:10.1016/j.jbankfin.2019.04.005.
- Berger, A.N., Bouwman, C.H.S., 2009. Bank liquidity creation. *The Review of Financial Studies* 22, 3779–3837. URL: <https://doi.org/10.1093/rfs/hhn104>, doi:10.1093/rfs/hhn104.
- Berger, A.N., Miller, N.H., Petersen, M.A., Rajan, R.G., Stein, J.C., 2005. Does function follow organizational form? evidence from the lending practices of large and small banks. *Journal of Financial Economics* 76, 237–269. URL: <https://www.sciencedirect.com/science/article/pii/S0304405X05000139>, doi:10.1016/j.jfineco.2004.06.003.
- Berger, A.N., Udell, G.F., 2002. Small business credit availability and relationship lending: The importance of bank organisational structure. *The Economic Journal* 112, F32–F53. URL: <https://doi.org/10.1111/1468-0297.00682>, doi:10.1111/1468-0297.00682.
- Bernanke, B.S., 1983. Nonmonetary effects of the financial crisis in propagation of the great depression. *The American Economic Review* 73, 257–276. MAG ID: 3123285923.
- Bernanke, B.S., Blinder, A.S., 1992. The federal funds rate and the channels of monetary transmission. *The American Economic Review* 82, 901–921. URL: <https://www.jstor.org/stable/2117350>. publisher: American Economic Association.
- Bonfim, D., Nogueira, G., Ongena, S., 2021. “sorry, we’re closed” bank branch closures, loan pricing, and information asymmetries*. *Review of Finance* 25, 1211–1259. URL: <https://doi.org/10.1093/rof/rfaa036>, doi:10.1093/rof/rfaa036.
- Calomiris, C.W., Hubbard, R.G., Stock, J.H., 1986. The farm debt crisis and public policy. *Brookings Papers on Economic Activity* 17, 441–486. URL: <https://ideas.repec.org/a/bin/bpeajo/v17y1986i1986-2p441-485.html>. publisher: Economic Studies Program, The Brookings Institution.
- Célerier, C., Matray, A., 2019. Bank-branch supply, financial inclusion, and wealth accumulation. *The Review of Financial Studies* 32, 4767–4809. URL: <https://academic.oup.com/rfs/article/32/12/4767/5477425>.
- Chodorow-Reich, G., 2014. The employment effects of credit market disruptions: Firm-level evidence from the 2008–9 financial crisis. *Quarterly Journal of Economics* 129, 1–59. doi:10.1093/qje/qjt031. dOI: 10.1093/qje/qjt031 MAG ID: 2168967186 S2ID: ea2259346ce836d70b18f084e52ebd9d366a0554.
- Clair, R.T., O’Driscoll, Gerald P., J., Yeats, K.J., 1994. Is banking different? a reexamination of the case for regulation. *Cato Journal* 13, 345–365. URL: <https://ideas.repec.org/a/cto/journl/v13y1994i3p345-365.html>. publisher: Cato Journal, Cato Institute.
- Conley, T.G., 1999. Gmm estimation with cross sectional dependence. *Journal of Econometrics* 92, 1–45. URL: <https://www.sciencedirect.com/science/article/pii/S0304407698000840>, doi:10.1016/S0304-4076(98)00084-0.

- Contreras, S., Delis, M.D., Ghosh, A., Hasan, I., 2022. Bank failures, local business dynamics, and government policy. *Small Business Economics* 58, 1823–1851. URL: <https://doi.org/10.1007/s11187-021-00478-5>, doi:10.1007/s11187-021-00478-5.
- Cowan, A.R., Salotti, V., 2015. The resolution of failed banks during the crisis: Acquirer performance and fdic guarantees, 2008–2013. *Journal of Banking and Finance* 54, 222–238. URL: <https://www.sciencedirect.com/science/article/pii/S0378426614003975>, doi:10.1016/j.jbankfin.2014.12.016.
- 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>.
- Demirgüç-Kunt, A., Kane, E., Laeven, L., 2015. Deposit insurance around the world: A comprehensive analysis and database. *Journal of Financial Stability* 20, 155–183. URL: <https://www.sciencedirect.com/science/article/pii/S1572308915000893>, doi:10.1016/j.jfs.2015.08.005.
- Diamond, D.W., Dybvig, P.H., 1983. Bank runs deposit insurance and liquidity. *Journal of Political Economy* 91, 401–419. doi:10.21034/qr.2412. dOI: 10.21034/qr.2412 MAG ID: 2993370650 S2ID: e67db834fe37e3384a801af20c11417be9562917.
- Driscoll, J.C., 2004. Does bank lending affect output? evidence from the u.s. states. *Journal of Monetary Economics* 51, 451–471. URL: <https://www.sciencedirect.com/science/article/pii/S0304393204000145>, doi:10.1016/j.jmoneco.2004.01.001.
- Duflo, E., Saez, E., 2003. The role of information and social interactions in retirement plan decisions: Evidence from a randomized experiment*. *The Quarterly Journal of Economics* 118, 815–842. URL: <https://doi.org/10.1162/00335530360698432>, doi:10.1162/00335530360698432.
- Friedman, M., Schwartz, A.J., 1963. A monetary history of the united states, 1867–1960. *The American Historical Review* 51, 100. doi:10.2307/1842177. dOI: 10.2307/1842177 MAG ID: 2571179306 S2ID: d322ea40724b23aa12d0dbddabd95ffb4ea3ddd9.
- Ge, Y., Qiu, J., 2007. Financial development, bank discrimination and trade credit. *Journal of Banking & Finance* 31, 513–530. URL: <https://www.sciencedirect.com/science/article/pii/S0378426606002202>, doi:10.1016/j.jbankfin.2006.07.009.
- Ghosh, A., 2017. Do bank failures still matter in affecting regional economic activity. *Journal of Economics and Business* 90, 1–16. doi:10.1016/j.jeconbus.2016.12.001. dOI: 10.1016/j.jeconbus.2016.12.001 MAG ID: 2561248150 S2ID: 1b8a7c893e8a39a90669c0a93f0e7e0cf09912af.
- Gilbert, R.A., Kochin, L.A., 1989. Local economic effects of bank failures. *Journal of Financial Services Research* 3, 333–345. doi:10.1007/bf00114049. dOI: 10.1007/bf00114049 MAG ID: 2046125621.
- Granja, J., Leuz, C., Rajan, R.G., 2022. Going the extra mile: Distant lending and credit cycles. *The Journal of Finance* 77, 1259–1324. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.13114>, doi:10.1111/jofi.13114. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.13114>.
- Greenstone, M., Mas, A., Nguyen, H.L., 2020. Do credit market shocks affect the real economy? quasi-experimental evidence from the great recession and "normal" economic times. *American Economic Journal: Economic Policy* 12, 200–225. URL: <https://www.aeaweb.org/articles?id=10.1257/pol.20160005>, doi:10.1257/pol.20160005.
- Huber, K., 2018. Disentangling the effects of a banking crisis: Evidence from german firms and counties. *The American Economic Review* 108, 868–898. doi:10.1257/aer.20161534. dOI: 10.1257/aer.20161534 MAG ID: 2789783506 S2ID: 952a6a00450fa1974a219dc14cd24faa5c56cac6.
- Ivashina, V., 2009. Asymmetric information effects on loan spreads. *Journal of Financial Economics* 92, 300–319. URL: <https://www.sciencedirect.com/science/article/pii/S0304405X09000208>, doi:10.2139/ssrn.849726.
- Jayaratne, J., Strahan, P.E., 1996. The finance-growth nexus: Evidence from bank branch deregulation.

- Quarterly Journal of Economics 111, 639–670. doi:[10.2307/2946668](#). dOI: 10.2307/2946668 MAG ID: 2014648357.
- Jordà, Ò., 2005. Estimation and inference of impulse responses by local projections. *American Economic Review* 95, 161–182. URL: <https://www.aeaweb.org/articles?id=10.1257/0002828053828518>, doi:[10.1257/0002828053828518](#).
- Kandrac, J., 2014. Bank Failure, Relationship Lending, and Local Economic Performance. Finance and Economics Discussion Series 2014-41. Board of Governors of the Federal Reserve System. doi:[10.2139/ssrn.2353687](#).
- Khwaja, A.I., Mian, A., 2008. Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *The American Economic Review* 98, 1413–1442. doi:[10.1257/aer.98.4.1413](#). dOI: 10.1257/aer.98.4.1413 MAG ID: 2057009456 S2ID: 1535749b3d71a1ceaf391db26c96e09decea536d.
- LeSage, J.P., Pace, R.K., 2009. *Introduction to Spatial Econometrics*. Chapman and Hall/CRC, Boca Raton, FL.
- Mei, Z., Sheng, L., Shi, Z., 2025. Nickell bias in panel local projection: Financial crises are worse than you think. *Journal of International Economics* , 104210doi:[10.1016/j.jinteco.2025.104210](#).
- Miguel, E., Kremer, M., 2004. Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica* 72, 159–217. URL: <https://www.jstor.org/stable/3598853>. publisher: [Wiley, Econometric Society].
- 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: 3a3cc06a0b40c7a667a05547feafdc63426b3c3.
- 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].
- Ranish, B., Stella, A., Zhang, J., 2024. Out of sight, out of mind: Nearby branch closures and small business growth URL: <https://www.federalreserve.gov/econres/feds/out-of-sight-out-of-mind-nearby-branch-closures-and-small-business-growth.htm>.
- Schularick, M., Taylor, A.M., 2012. Credit booms gone bust: Monetary policy, leverage cycles and financial crises, 1870–2008. *The American Economic Review* 102, 1029–1061. doi:[10.1257/aer.102.2.1029](#). dOI: 10.1257/aer.102.2.1029 MAG ID: 3124777039.