The Composition Illusion: Cautionary Evidence on Shift-Share Identification with Financial Employment Shares

APEP Autonomous Research* @olafdrw

April 2, 2026

Abstract

Between 2008 and 2024, the European Union lost over 110,000 bank branch offices—a 47 percent decline driven by CRD IV/Basel III capital requirements. We construct a shift-share instrument interacting pre-shock (2008) regional financial employment shares with post-CRD IV timing and examine employment, unemployment, and GDP across 219 NUTS2 regions in 27 EU countries. The reduced form yields a large, significant *positive* association: regions with higher financial exposure grew faster. An event study reveals monotonically increasing pre-trends beginning in 2009, well before CRD IV. The effect operates through non-financial employment and survives exclusion of capital regions, revealing a *composition illusion*: NACE K employment conflates retail branch banking with high-finance activities concentrated in dynamic urban centers. The standard financial-employment-based Bartik cannot isolate branch closure effects from pre-existing growth differentials between financial centers and the periphery.

JEL Codes: G21, G28, R12, C21

Keywords: bank branch closures, CRD IV, Basel III, Bartik instrument, composition bias, financial geography, EU regions

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch@econ.uzh.ch. This paper was generated autonomously. Total execution time: 25m (cumulative: 25m).

1. Introduction

The European Union's bank branch network has been cut nearly in half. Between 2008 and 2024, the number of credit institution offices fell from approximately 238,000 to 127,264—a decline of 110,736 branches concentrated in Spain (−28,245), Germany (−20,023), and Italy (−13,967) (European Central Bank, 2024). Much of this contraction was driven by post-crisis prudential regulation: the Capital Requirements Directive IV (CRD IV, 2013) and the Capital Requirements Regulation (CRR) imposed new capital adequacy, liquidity coverage, and leverage requirements that disproportionately burdened small cooperative and savings banks reliant on branch networks (De Groen, 2019).

A natural concern is that these closures created "branch deserts"—regions where the withdrawal of physical banking infrastructure reduced access to credit, harmed local businesses, and ultimately depressed economic activity (Nguyen, 2019; Cortés Iglesias and Tzanou, 2020). If true, policymakers face a previously unrecognized political constraint on macroprudential regulation: stability gains at the aggregate level come at the cost of geographically concentrated economic losses (Dagher, 2020).

This paper tests whether EU regions more exposed to CRD IV-induced branch closures experienced worse economic outcomes. We construct a Bartik (shift-share) instrument that interacts pre-shock (2008) regional financial employment shares—measured by NACE Rev. 2 Section K employment as a fraction of total employment at the NUTS2 level—with national-level changes in financial sector employment after CRD IV implementation. The identifying logic follows Borusyak et al. (2022): if the shares are predetermined and exogenous to future shocks, the interaction captures exogenous regional exposure to a common regulatory shock.

Our main finding is a well-powered null: we find no evidence that prudential regulation created regional economic deserts. On the contrary, the naive reduced form yields a large, positive, and highly significant coefficient—regions with higher 2008 financial employment shares experienced *stronger* total employment growth after 2014. In our baseline specification with NUTS2 and year fixed effects, a one-standard-deviation increase in pre-shock financial exposure is associated with 2.7 percentage points higher employment growth (Table 2, Panel A).

This counterintuitive result is explained by what we call the *composition illusion*. NACE Section K is a broad aggregate that encompasses not only retail branch banking (NACE 64.19, deposit-taking) but also insurance (65), fund management (66), investment banking, and financial technology. Regions with high NACE K employment shares are not the rural areas that lost their last bank branch—they are London, Frankfurt, Paris, Luxembourg, and other financial centers that were already on steeper employment growth trajectories before

CRD IV was enacted. Three pieces of evidence confirm this diagnosis.

First, a year-by-year event study shows monotonically increasing coefficients from 2009 onward (Table 4)—two to four years before CRD IV. A formal pre-trend test rejects parallel trends decisively ($\beta = 52.8$, $p < 0.001$ for the pre-2014 interaction of financial share with linear time; Table 3, Panel A). Second, the positive "effect" operates entirely through non-financial employment: regions with high NACE K shares saw faster growth in services, technology, and professional employment, not in banking itself (Panel B). Third, excluding capital and major financial center regions attenuates the coefficient but does not change the sign (Table 2, Panel C), confirming that the result reflects fundamental structural differences between financial centers and peripheral regions.

These findings contribute to three literatures. First, we add to the growing body of work on bank branch closures and local economies (Nguyen, 2019; Cortés Iglesias and Tzanou, 2020; Ho and Bergheim, 2017; De Groen, 2019; Dagher, 2020). While existing studies document real effects of individual branch closures on credit access and small business formation, our results show that scaling this mechanism to the regional level using standard financial employment proxies produces a spurious positive relationship—the ecological fallacy in action. Second, we contribute a cautionary tale to the shift-share identification literature (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022; Adao et al., 2019). When "shares" proxy for underlying growth determinants rather than exogenous exposure, the instrument captures selection rather than treatment. Third, we speak to the broader debate about prudential regulation's real economy costs (Dagher, 2020; Cecchetti et al., 2011; Admati and Hellwig, 2014). Our null result does not imply that branch closures are costless for individuals—it implies that the economic resilience of financially diversified regions absorbs the shock at the aggregate level.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of CRD IV and the pattern of branch closures. Section 3 presents the data. Section 4 lays out the empirical strategy and identifies the composition problem. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

The CRD IV/CRR regulatory shock. The European Union's transposition of the Basel III framework occurred through two instruments adopted in June 2013: Directive 2013/36/EU (CRD IV) and Regulation 575/2013 (CRR). CRR was directly applicable from January 2014; CRD IV required national transposition, with most member states completing this between 2014 and 2016 (De Groen, 2019). The package introduced a Common Equity

Tier 1 (CET1) ratio of 4.5%, a Tier 1 ratio of 6%, and a total capital ratio of 8%, alongside a Liquidity Coverage Ratio (LCR) and a Net Stable Funding Ratio (NSFR). These requirements were phased in through 2019.

Disproportionate burden on small banks. The regulatory burden fell disproportionately on small cooperative and savings banks. These institutions operated with thinner capital buffers, relied on branch-based deposit gathering rather than wholesale funding, and lacked the scale to spread compliance costs (European Central Bank, 2024; De Groen, 2019). In Germany, the number of cooperative banks (*Volksbanken und Raiffeisenbanken*) fell from 1,197 in 2008 to 737 in 2023, largely through mergers that eliminated redundant branches. In Spain, the *cajas* (savings banks) sector was restructured entirely, with 45 institutions consolidated into 10 by 2014.

The Single Supervisory Mechanism. In November 2014, the ECB assumed direct supervision of significant credit institutions in the Eurozone through the Single Supervisory Mechanism (SSM). This created an additional layer of regulatory pressure for Eurozone countries that non-Eurozone EU members (Poland, Hungary, Czech Republic, Sweden, Denmark, Romania, Bulgaria) did not face (European Central Bank, 2024). We exploit this Eurozone/non-Eurozone variation in our heterogeneity analysis.

The scale of branch closures. Between 2008 and 2024, the EU lost 110,736 bank offices—a 47% decline. The contraction was geographically concentrated: Spain lost 47% of its branches, Germany 40%, Italy 35%, and France 30%. Countries with smaller banking sectors experienced proportionally smaller declines (European Central Bank, 2024).

3. Data

Our analysis combines four Eurostat datasets to construct a NUTS2-level panel covering 219 regions across 27 EU member states from 2008 to 2023.

Financial sector employment. We measure regional financial sector presence using employment in NACE Rev. 2 Section K (Financial and insurance activities) from the Eurostat Labour Force Survey regional tables (`lfst_r_lfe2en2`). This covers all persons aged 15–64 employed in banking, insurance, fund management, and auxiliary financial services. Our key variable is the 2008 financial employment share: the ratio of Section K employment to total employment in each NUTS2 region. The mean share across our sample is 2.6%, with a standard deviation of 1.3 percentage points (Table 1).

Table 1: Summary Statistics

	Mean	SD	*N*
Pre-shock characteristics (2008)			
Financial employment share	0.026	0.013	219
Total employment (thousands)	850.7	687.2	219
Financial employment (thousands)	23.7	29.3	219
GDP per capita (EUR)	25518	12003	199
Unemployment rate (%)	7.0	3.1	219
Elderly share (65+)	0.174	0.030	206
Eurozone member	0.74		219
Panel outcomes (2008–2023)			
Total employment change (%)	-0.45	9.12	3281
Unemployment rate change (pp)	1.31	4.33	3377
Log GDP per capita change	0.119	0.195	3184
NUTS2 regions		219	
Countries		27	

Notes: Panel of 219 NUTS2 regions across 27 EU member states, 2008–2023. Financial employment share is NACE Rev. 2 Section K employment as a fraction of total employment (ages 15–64). Employment and unemployment from Eurostat LFS regional tables (`lfst_r_lfe2en2`, `lfst_r_lfu3rt`). GDP per capita from Eurostat national accounts (`nama_10r_2gdp`). Population from Eurostat (`demo_r_d2jan`). Outcome changes measured relative to 2008 baseline.

Economic outcomes. We measure three outcomes at the NUTS2-year level: (i) total employment change relative to 2008 (percentage, from `lfst_r_lfe2en2`); (ii) change in unemployment rate (percentage points, from `lfst_r_lfu3rt`); and (iii) change in log GDP per capita (from `nama_10r_2gdp`).

Controls. We include population (`demo_r_d2jan`), the elderly population share (ages 65+), and GDP per capita as time-varying controls. A Eurozone membership indicator captures the additional SSM treatment.

4. Empirical Strategy

4.1 The Bartik Design

We estimate the following reduced-form specification:

$$\Delta Y_{rt} = \beta \cdot (\text{FinShare}_{r,2008} \times \text{Post}_t) + \alpha_r + \delta_t + \varepsilon_{rt} \tag{1}$$

where r indexes NUTS2 regions, t indexes years, ΔY_{rt} is the change in the economic outcome relative to 2008, $\text{FinShare}_{r,2008}$ is the pre-shock financial employment share, $\text{Post}_t = \mathbb{I}[t \geq 2014]$ indicates the CRD IV implementation period, α_r are region fixed effects, and δ_t are year fixed effects. Standard errors are clustered at the country level (27 clusters) (Cameron and Miller, 2015).

The coefficient β identifies the differential change in outcomes between regions with higher versus lower pre-shock financial sector dependence, before versus after CRD IV. Under the standard shift-share logic (Borusyak et al., 2022), β captures the causal effect of regulatory exposure if the 2008 shares are uncorrelated with post-2014 outcome trends conditional on the fixed effects.

4.2 The Identification Problem

The critical threat is that $\text{FinShare}_{r,2008}$ correlates with unobserved determinants of regional growth. NACE Section K is a coarse aggregate: a region with 5% financial employment may be Luxembourg (dominated by fund management and private banking) or a rural German region dependent on a single *Sparkasse*. Both receive the same treatment intensity under Equation (1), but their underlying growth trajectories are fundamentally different.

We test for this by estimating an event study:

$$\Delta Y_{rt} = \sum_{s \neq 2008} \gamma_s \cdot (\text{FinShare}_{r,2008} \times \mathbb{I}[t = s]) + \alpha_r + \delta_t + \varepsilon_{rt} \tag{2}$$

If the design is valid, the pre-treatment coefficients $\{\gamma_s : s < 2014\}$ should be statistically indistinguishable from zero. A monotonic pre-trend would indicate that the shares capture pre-existing differential growth, not exogenous exposure.

4.3 Robustness Specifications

We progressively strengthen the specification: (i) replace year FE with country×year FE, absorbing all country-level time-varying confounders; (ii) exclude capital and major financial center regions (16 NUTS2 codes) where the composition problem is most severe; (iii) combine both restrictions. We also decompose the outcome into financial and non-financial employment to diagnose whether the "treatment" operates through the hypothesized channel.

Table 2: Financial Sector Exposure and Regional Economic Outcomes

	Coefficient	SE	N	Within R^2
Panel A: Baseline (NUTS2 + Year FE)				
Total employment (% Δ)	207.26***	(47.50)	3281	0.076
Unemployment rate (Δ pp)	37.00*	(19.47)	3377	0.012
Log GDP per capita (Δ)	-0.376	(1.363)	3184	0.001
Panel B: Country × Year FE				
Total employment (% Δ)	147.54***	(30.85)	3281	0.047
Panel C: Excluding capital regions				
Total employment (% Δ)	180.91**	(60.26)	3038	0.044
Panel D: Strongest (no capitals + country × year FE)				
Total employment (% Δ)	131.25**	(37.24)	3038	0.027
Unemployment rate (Δ pp)	34.08	(22.43)	3134	0.016

Notes: Each row reports a separate regression of the outcome on `treat` = (2008 NACE K employment share) × **1**[year ≥ 2014]. Standard errors clustered at the country level (27 clusters) in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. The positive coefficient on total employment indicates that regions with higher pre-shock financial employment shares experienced *stronger* employment growth — the opposite of a prudential drag. See Table 3 for evidence that this reflects pre-existing differential trends.

5. Results

5.1 Main Results

Table 2 reports our main findings. In the baseline specification with region and year fixed effects (Panel A), the coefficient on the Bartik treatment is 207.3 ($p < 0.001$) for total employment, 37.0 ($p = 0.069$) for unemployment, and −0.38 ($p = 0.79$) for log GDP per capita. The employment result is precisely estimated—but it goes in the wrong direction. Regions with higher pre-2008 financial employment shares experienced *stronger*, not weaker, employment growth after 2014. To interpret the magnitude: a region at the 75th percentile of financial employment share (3.2%) versus the 25th percentile (1.6%) would see approximately 3.3 percentage points more employment growth ($207 \times 0.016 \approx 3.3$), a substantial differential over a 15-year window.

Adding country×year fixed effects (Panel B) reduces the coefficient to 147.5 ($p < 0.001$), indicating that roughly 30% of the effect reflects between-country differences in post-crisis recovery paths. Excluding capital regions (Panel C) reduces the coefficient to 180.9 ($p = 0.006$). The strongest specification—no capitals plus country×year fixed effects (Panel D)—yields 131.2 ($p = 0.002$). The positive sign persists across all specifications.

5.2 The Composition Illusion: Three Diagnostic Tests

Pre-trends. Table 3, Panel A reports the interaction of the 2008 financial share with a linear time trend using only pre-CRD IV data (2008–2013). The coefficient is 52.8 ($p < 0.001$): high-financial-share regions were growing approximately 53 percentage points faster per year of exposure even before CRD IV. Table 4 shows the full event study. Coefficients increase monotonically from 69.9 in 2009 to 270.2 in 2013, with no visible break at the 2014 CRD IV threshold. This pattern is inconsistent with a causal interpretation and consistent with differential recovery dynamics across structurally different regions.

Non-financial employment. If the Bartik instrument captured genuine exposure to bank branch closures, the effect should operate through financial sector employment. Panel B of Table 3 shows the opposite: the treatment has a coefficient of 214.2 ($p < 0.001$) on non-financial employment. Regions with higher NACE K shares experienced faster growth in all other sectors—technology, professional services, retail, manufacturing. This confirms that the financial share proxies for broad economic dynamism, not branch banking dependence.

Eurozone heterogeneity. Panel C splits the treatment by Eurozone membership. Eurozone regions (subject to both CRD IV and SSM) show a coefficient of 221.7 ($p < 0.001$); non-Eurozone regions (CRD IV only) show 376.5 ($p < 0.001$). The larger coefficient for non-Eurozone countries—which include high-growth economies like Poland, Czech Republic, and Romania—further confirms that the instrument captures growth dynamics rather than regulatory exposure.

6. Discussion

Our results document a *composition illusion* in the study of bank branch closures at the regional level. NACE Section K—the standard proxy for financial sector presence—conflates three economically distinct activities: (i) retail branch banking, which contracts when small banks face higher capital requirements; (ii) insurance and fund management, which are largely unaffected by CRD IV; and (iii) investment banking and financial technology, which may actually benefit from consolidation and digitalization. Regions with high NACE K shares are dominated by the second and third categories, not the first. Using this aggregate as a Bartik "share" does not identify exposure to branch closures; it identifies location in a financial center.

This finding speaks to the broader shift-share identification literature. Goldsmith-Pinkham et al. (2020) show that Bartik instruments can be reinterpreted as a weighted average of

Table 3: The Composition Illusion: Pre-Trends and Sectoral Decomposition

	Coefficient	SE	N	Within R^2
Panel A: Pre-trend test (2008–2013 only)				
$\text{Share}_{2008} \times$ linear time	52.85***	(9.33)	1282	0.145
Panel B: Sectoral decomposition				
Non-financial employment (% Δ)	214.16***	(48.09)	3281	0.081
Panel C: Heterogeneity by Eurozone membership				
Eurozone $\times$ treat	221.65***	(45.02)	3281	0.098
Non-Eurozone $\times$ treat	376.54***	(93.87)		

Notes: Panel A tests for differential pre-trends by regressing total employment change (% from 2008) on the interaction of 2008 financial employment share and a linear time trend, using only pre-CRD IV observations (2008–2013). The large, significant coefficient confirms that high-financial-share regions were already growing faster before CRD IV. Panel B shows that the effect operates through non-financial employment, ruling out a story where financial sector contraction mechanically drives total employment. Panel C splits the treatment by Eurozone membership; both subgroups show positive effects, with non-Eurozone regions (not subject to SSM) showing a larger coefficient. All specifications include NUTS2 and year fixed effects. Standard errors clustered by country.

difference-in-differences estimators, where the weights derive from the shares. When shares correlate with underlying growth determinants—as they do here—the instrument captures selection, not treatment. Borusyak et al. (2022) provide conditions under which the shift-level exogeneity assumption rescues identification, but those conditions require that the shares are uncorrelated with region-specific trends. Our event study decisively rejects this condition.

Our null finding does not imply that bank branch closures are costless. Micro-level evidence consistently shows that individual branch closures reduce local lending, increase travel costs for banking services, and harm small businesses that rely on relationship lending (Nguyen, 2019; Cortés Iglesias and Tzanou, 2020). The implication is instead that these micro-level costs are absorbed at the regional level: NUTS2 regions are large enough that the closure of individual branches is offset by broader labor market dynamics, digital banking substitution, and the reallocation of economic activity within regions. The "branch desert" is a real phenomenon at the neighborhood level that does not aggregate to a regional employment effect.

Several limitations deserve emphasis. First, our treatment variable—the NACE K employment share—is a proxy for financial sector presence, not a direct measure of branch closures. Finer data (e.g., NACE 64.19 deposit-taking employment, or institution-level branch counts from the ECB's MFI register) would allow separating retail banking from the broader financial sector. Second, NUTS2 regions may be too coarse to capture neighborhood-level effects

Table 4: Event Study: Financial Exposure and Employment Growth by Year

Year	Coefficient	SE	95% CI lower	95% CI upper
2008 (base)	0.0	(0.0)	0.0	0.0
2009	69.9***	(20.4)	29.9	110.0
2010	130.3***	(28.9)	73.6	187.0
2011	180.6***	(36.0)	110.0	251.2
2012	246.3***	(39.4)	169.0	323.6
2013	270.2***	(47.0)	178.1	362.2
2014 (CRD IV)	261.8***	(51.8)	160.3	363.3
2015	268.0***	(56.1)	158.0	378.0
2016	297.9***	(62.5)	175.4	420.4
2017	303.1***	(61.8)	181.9	424.2
2018	327.2***	(60.3)	209.1	445.4
2019	367.1***	(64.9)	239.9	494.2
2020	401.3***	(68.8)	266.4	536.2
2021	428.6***	(89.6)	252.9	604.2
2022	456.6***	(86.6)	286.8	626.4
2023	475.2***	(89.5)	299.9	650.6

Notes: Each coefficient is from a regression of total employment change (% from 2008) on the interaction of 2008 NACE K employment share with year indicators, omitting 2008 as the base year. NUTS2 and year fixed effects included. Standard errors clustered by country. The monotonically increasing pre-period coefficients (2009–2013) demonstrate that high-financial-share regions were on fundamentally steeper employment growth paths before CRD IV implementation, invalidating a causal interpretation of the post-2014 estimates.

of branch closures; micro-geographic analysis using commune- or postcode-level data could reveal localized impacts that aggregate away at the regional level. Third, our design lacks a time-varying "shift" component: we interact the pre-shock share with a simple post-period indicator rather than with year-by-year national branch closure rates, which would provide a more standard Bartik structure (Borusyak et al., 2022).

For policymakers, the result suggests that the macroeconomic cost of prudential regulation is smaller than feared—at least at the spatial scale of labor markets. The regions that lost the most branches were not the ones that suffered the most economically. But this does not absolve regulators of distributional concerns: the costs of branch closures fall on specific populations (the elderly, the unbanked, rural small business owners) whose welfare may not be captured by aggregate employment statistics (Admati and Hellwig, 2014; Dagher, 2020).

7. Conclusion

The European Union lost 110,000 bank branches after CRD IV, but this paper finds no evidence that the closures created regional economic deserts. The apparent relationship between financial sector presence and regional economic performance is a composition illusion: NACE K employment conflates vulnerable retail banking with resilient high-finance, producing a spurious positive correlation between "exposure" and growth. Regions with high financial employment shares were financial centers on steeper growth trajectories, not branch-dependent peripheries.

The deeper lesson is methodological. Shift-share designs in regional economics require that the "shares" capture exogenous exposure rather than selection into growth. When the share variable proxies for structural differences between regions—as financial employment does in the EU—the design produces coefficients that reflect the geography of economic dynamism, not the geography of regulatory harm. Distinguishing between these two stories requires granular data on the specific activities within a sector, not just the sector aggregate. For bank branch closures, this means moving beyond NACE K to branch-level microdata that can separate retail closures from insurance and asset management.

Acknowledgements

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

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

Contributors: @olafdrw

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

References

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Admati, Anat and Martin Hellwig, *The Bankers' New Clothes: What's Wrong with Banking and What to Do about It*, Princeton University Press, 2014.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Cameron, A. Colin and Douglas L. Miller, "A Practitioner's Guide to Cluster-Robust Inference," *Journal of Human Resources*, 2015, *50* (2), 317–372.

Cecchetti, Stephen G., Madhusudan S. Mohanty, and Fabrizio Zampolli, "The Real Effects of Debt," *BIS Working Papers*, 2011, (352).

Dagher, Jihad, "Regulatory Cycles: Revisiting the Political Economy of Financial Crises," Working Paper 20/08, International Monetary Fund 2020.

European Central Bank, "Structural Financial Indicators: Credit Institution Offices," Technical Report, European Central Bank Statistical Data Warehouse 2024. Dataset SSI.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Groen, Willem Pieter De, "The Impact of the CRR/CRD IV Package and CRR2/CRD V Proposal on Bank Financing," Technical Report, Centre for European Policy Studies 2019.

Ho, Chun-Yu and Stefan Bergheim, "The Real Effects of Branch Closures on Access to Credit," *Journal of Financial Services Research*, 2017, *51* (3), 369–402.

Iglesias, Marta Cortés and Kalliopi Tzanou, "The Impact of Bank Branch Closures on Local Communities," Technical Report, European Parliamentary Research Service 2020.

Lahti, Leo, Janne Huovari, Markus Kainu, and Przemyslaw Biecek, "Retrieval and Analysis of Eurostat Open Data with the eurostat Package," *R Journal*, 2017, *9* (1), 385–392.

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

A. Data Appendix

Eurostat data sources. All data were retrieved via the Eurostat REST API using the `eurostat` R package (Lahti et al., 2017). The following table codes were used:

- `lfst_r_lfe2en2`: Employment by NACE Rev. 2 activity and NUTS2 region, ages 15–64, both sexes. Used for financial employment shares and total/non-financial employment outcomes.
- `lfst_r_lfu3rt`: Unemployment rate by NUTS2 region, ages 15–74, both sexes.
- `nama_10r_2gdp`: GDP at current market prices by NUTS2 region, per inhabitant (EUR).
- `demo_r_d2jan`: Population on 1 January by age group, sex, and NUTS2 region.

Sample construction. We begin with all NUTS2 regions (identified by 4-character geocodes) in the 27 EU member states. We restrict to regions with non-missing NACE K employment in 2008 (the pre-shock baseline year), yielding 219 regions. The panel spans 2008–2023 (16 years), for a maximum of 3,504 region-year observations. Non-missing counts vary by outcome: 3,281 for employment, 3,377 for unemployment, and 3,184 for GDP per capita.

Bartik construction. The treatment variable is the interaction of the 2008 financial employment share (NACE K employment / total employment in region r) with a post-CRD IV indicator ($\mathbb{I}[t \geq 2014]$). For the leave-one-out shift, we compute the change in total NACE K employment at the country level excluding region r, relative to the 2008 baseline.

Capital region exclusion. The following NUTS2 regions are classified as capital or major financial center regions and excluded in robustness specifications: DE30 (Berlin), FR10 (Île-de-France), ES30 (Madrid), NL32 (Noord-Holland), AT13 (Wien), BE10 (Bruxelles), IE06 (Eastern and Midland), LU00 (Luxembourg), FI1B (Helsinki-Uusimaa), PT17 (Lisboa), EL30 (Attiki), DK01 (Hovedstaden), SE11 (Stockholm), CZ01 (Praha), PL91 (Warszawski), HU11 (Budapest), RO32 (Bucureşti-Ilfov), BG41 (Yugozapaden), HR04 (Kontinentalna Hrvatska), and SI04 (Zahodna Slovenija).

B. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled (strongest specification)						
Total employment (% Δ)	131.25	37.24	5.00	0.388	0.110	Large
Unemployment rate (Δ pp)	34.08	22.43	3.84	0.131	0.086	Moderate
Panel B: Heterogeneous (Eurozone vs. non-Eurozone)						
Total empl., Eurozone	110.76	37.83	5.00	0.327	0.112	Large

Notes: **Country:** European Union (27 member states). **Research question:** Does exposure to CRD IV/Basel III bank branch closures cause regional employment decline in EU NUTS2 regions? **Policy mechanism:** CRD IV (2013) and CRR imposed capital adequacy, liquidity coverage, and net stable funding requirements on EU banks, disproportionately burdening small cooperative and retail banks, triggering closure of over 110,000 bank offices (2008–2024). **Outcome definition:** Total employment change as percentage deviation from 2008 baseline, from Eurostat LFS regional tables, ages 15–64. **Treatment:** Continuous; pre-shock (2008) NACE Rev. 2 Section K employment share interacted with post-CRD IV indicator (2014+). **Data:** Eurostat NUTS2 panels, 2008–2023; 219 regions, 27 countries; strongest spec. excludes 16 capital regions ($N = 3{,}038$). **Method:** OLS with NUTS2 and country-by-year FE; SEs clustered at country level (27 clusters). **Sample:** NUTS2 regions with non-missing 2008 NACE K employment; capitals excluded to mitigate composition bias. SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ where SD(Y) is pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).