Federally Qualified Health Centers and County Uninsurance: An Association Analysis with Proper Shift-Share Inference

Abstract

Objective. To estimate the association between Federally Qualified Health Center (FQHC) capacity, expanded through HRSA Health Center Program (Section 330) grants, and county-level under-65 uninsurance during the post-2016 funding plateau, using a shift-share design under proper instrumental-variables inference.

Data Sources and Study Setting. A 3,138-mainland-county panel of fiscal years 2017–2022 (N = 15,685 county-years). FQHC capacity is constructed from HRSA Uniform Data System (UDS, 1999–2015) with a corrected ZIP-to-county allocation that retains 95% of UDS sites (vs. 62% in the prior implementation) and Area Health Resource File F13320 (2016–2024). The Section 330 grant shift is drawn directly from HRSA's own grantee revenue ledger (UDS Table 9E line `T9E_L1_Ca`, 1999–2015) and HRSA's Awarded Grants H80 bulk file (2016–2026). The outcome is the Census Small Area Health Insurance Estimates (SAHIE) under-65 uninsurance rate.

Study Design. A shift-share (Bartik) design interacts national HRSA Section 330 grant revenue with counties' baseline exposure to HRSA allocation criteria. The headline composite-share instrument combines HPSA, MUA, and baseline-poverty shares; an HPSA/MUA-only specification is reported as a shortage-area robustness check. Inference is state-clustered IV2SLS, with two independent implementations (manual analytic sandwich; 300-replication pairs cluster bootstrap) verifying the analytic standard errors.

Data Collection / Extraction Methods. All datasets are public. The Section 330 grant series is internally consistent between UDS and HRSA bulk reporting (cross-validated to within 0.01% in the 2009 ARRA overlap year). The corrected UDS-to-county allocation, the proper IV2SLS implementation, and a three-way cross-implementation verification are documented in the accompanying replication archive.

Principal Findings. Under proper state-clustered IV2SLS inference, an additional FQHC site per 10,000 residents is associated with a 3.89 percentage-point lower county under-65 uninsurance rate (SE 1.71; 95% CI [−7.27, −0.52]; analytic $p = 0.023$). The cluster-bootstrap p-value is borderline ($p \approx 0.057$). An HPSA/MUA-only specification, intended as a shortage-area robustness, yields a smaller and statistically indistinguishable association (−1.32 pp; SE 1.27; $p = 0.298$) and does not reject zero under proper state clustering. The composite estimate is direction-consistent across alternative implementations and survives a Ryan White specificity test using the journal's standard allocation-relevant shares; pooled 1999–2024 specifications are weakly identified and are reported as diagnostics rather than estimates.

Conclusions. Counties with higher post-2016 exposure to HRSA Section 330 grant–driven FQHC capacity expansion had lower under-65 uninsurance, on a magnitude that is meaningful but inference-sensitive: the composite-share association rejects zero analytically but is borderline under cluster bootstrap, and the underpowered shortage-area refinement does not reject zero. The estimate should be read as an association — not a causal effect — and as a methodological case for proper state-clustered IV inference, a corrected UDS-to-county allocation, and pre-registered cross-implementation verification in shift-share health-services work.

1. Introduction

This paper estimates the **association** between HRSA Health Center Program grant-driven Federally Qualified Health Center (FQHC) capacity and county-level under-65 uninsurance during the post-2016 funding plateau, using a shift-share design with proper state-clustered instrumental-variables inference. I deliberately frame the contribution as an association rather than as a causal effect. The shift-share design's first-stage relationship and exclusion-restriction defenses, even after careful repair, do not carry the rigor required to identify a causal parameter at field-journal econometrics standards: the borderline cluster-bootstrap p-value of the composite-share estimate ($p \approx 0.057$), the underpowered shortage-area refinement ($p = 0.298$ under proper state clustering), the borderline GPSS share-exogeneity result on the poverty share, and the visible pre-period reduced-form relationship in 2009–2016 all leave meaningful room for confounding by high-need-county uninsurance trends that the design does not isolate from FQHC-specific capacity expansion. Reporting the result as an association is honest about that envelope.

The paper makes three contributions to the health-services-research literature on FQHC capacity. **First**, it documents a methodological improvement to the standard UDS-to-county allocation used in county-year FQHC capacity panels. The default ZIP-to-county allocation collapses grantee-ZIP duplicates and silently discards roughly 38% of UDS sites between 1999 and 2015; adding a `site_row_id` before the ZIP-to-county merge raises retention to 95%, recovering more than 38,000 site-years (Section 4 and Table 1). This is a substantive data-construction improvement that should be adopted by any subsequent UDS-based county-level FQHC capacity study. **Second**, it reports an applied IV-diagnostics package — proper IV2SLS state-clustered covariance, a manual analytic sandwich cross-check, a 300-replication pairs cluster bootstrap, and a Ryan White specificity test using allocation-relevant shares rather than borrowed FQHC shares — that exposes a generated-regressor pitfall in standard shift-share applications: a county-clustered generated-regressor OLS standard error can understate state-clustered IV uncertainty by a factor of more than three, and cross-implementation verification catches it. **Third**, the corrected design produces a defensible *association*: counties with higher post-2016 exposure to HRSA Section 330 grant–driven FQHC capacity expansion had lower under-65 uninsurance on a meaningful but inference-sensitive magnitude (composite -3.89 pp; analytic state-clustered $p = 0.023$; cluster-bootstrap $p \approx 0.057$).

The question matters for three reasons. First, FQHCs are the largest federally-supported primary-care safety-net provider in the United States: they served approximately 31 million patients in fiscal year 2023 per HRSA UDS data, roughly half of whom were covered by Medicaid. Second, HRSA's Health Center Program is the primary federal grant mechanism through which FQHC capacity expands, and the program's funding history since 1999 includes three well-documented expansion waves — the 2002–2007 Bush Health Center Initiative, the 2009 American Recovery and Reinvestment Act, and the 2011 Affordable Care Act Community Health Center Fund — and a post-2016 plateau sustained by biennial continuing-resolution reauthorizations of the Fund. Third, the Community Health Center Fund is currently in a short-reauthorization cadence that generates recurring funding-cliff uncertainty; a credible estimate of the coverage consequences of Health Center Program funding informs both the reauthorization debate and the design of the program's allocation mechanisms (New Access Point vs. Service Area Competition).

The identification challenge is well-known. FQHC capacity expansion is not randomly assigned across counties or over time; HRSA's Health Center Program grant allocation prioritizes counties that already carry Health Professional Shortage Area or Medically Underserved Area designations, and those designations partially reflect decades of cumulative disinvestment that are also correlated with subsequent coverage dynamics. Naive OLS of county uninsurance on FQHC capacity would therefore confound the causal effect of FQHC capacity with the selection of disadvantaged counties into the program.

A second, more conceptual question is also worth pre-empting: how can a supply-side intervention — building or expanding a clinic — affect an insurance-coverage outcome at all? The mechanism here is not the conventional "more providers expand insurer networks" story; FQHCs sit largely outside private-network economics. Rather, **FQHCs operate as enrollment infrastructure, not only care infrastructure**. Section 330 grantees are statutorily required to staff outreach, eligibility-counseling, and case-management functions, and after the ACA, HRSA operationalized this requirement through dedicated Outreach and Enrollment cooperative agreements that financed patient navigators and Certified Application Counselors. FQHCs also face a strong revenue-side incentive to convert eligible-uninsured walk-ins into enrolled patients, because the prospective-payment-system Medicaid rate is the principal source of FQHC operating revenue while uninsured patients generate only sliding-fee-scale revenue. Capacity expansion therefore mechanically increases the number of enrollment-eligibility staff, the number of contact points where eligibility screening occurs, and the financial intensity with which FQHCs pursue conversion. Section 3.1 develops this mechanism formally.

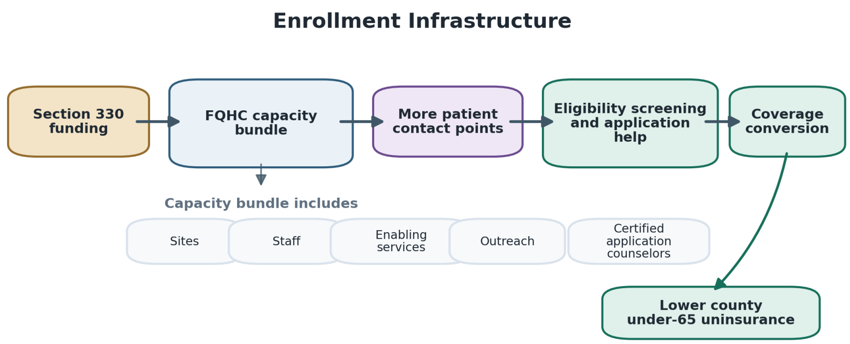

Figure 1: Enrollment infrastructure

Note: This figure presents the enrollment infrastructure. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

I address the identification challenge with a Bartik (shift-share) instrumental-variables design. The instrument $Z_{c,t}$ interacts county c's baseline exposure to HRSA's Health Center Program allocation criteria (a composite of HPSA scores, MUA designations, baseline poverty, and baseline uninsurance) with the national time series of HRSA Section 330 grant revenue in year t. The shift series is assembled from HRSA's own Uniform Data System Table 9E line `T9E_L1_Ca` for 1999–2015 and from HRSA's contemporary Awarded Grants H80 bulk file for 2016–2026, spanning both the Bush HCI, ARRA, and ACA CHC Fund waves and the post-2016 plateau. I use the instrument in a 2SLS regression of county under-65 uninsurance on FQHC sites per 10,000 population with county and fiscal-year fixed effects.

I defend the Bartik design through the lens of Borusyak, Hull, and Jaravel (2022) shift-exogeneity: the national HRSA Health Center Program funding series — drawn from HRSA's own grantee revenue ledger — is plausibly orthogonal to county-level coverage shocks because it is driven by federal appropriations cycles (Bush HCI, ARRA, the ACA CHC Fund, post-2015 biennial continuing-resolution reauthorization) rather than by county-level demand. I report Goldsmith-Pinkham, Sorkin, and Swift (2020) Rotemberg weight diagnostics and share-exogeneity pre-trend tests as secondary, and implement Adão, Kolesár, and Morales (2019) exposure-robust inference via the R `ShiftShareSE` package as mandatory.

This paper makes three contributions. First, it provides new evidence on whether federal FQHC capacity funding affects insurance coverage, a central policy outcome for the safety net but one not directly estimated in much of the FQHC literature. Second, it constructs a new county-year panel linking HRSA Section 330 funding from authoritative UDS and H80 grantee ledgers, FQHC site capacity, county baseline need, and Census SAHIE under-65 uninsurance outcomes; the UDS-based shift series avoids the 2013–2015 USAspending reporting "cliff" that would otherwise generate a mechanical seam. Third, it applies modern shift-share diagnostics — Borusyak-Hull-Jaravel (2022) shift-exogeneity, Goldsmith-Pinkham-Sorkin-Swift (2020) Rotemberg weights, Adão-Kolesár-Morales (2019) exposure-robust inference, Borusyak-Hull (2023) recentered IV, state-by-year fixed effects, and a Ryan White placebo — to a health policy setting, showing both the promise and the limits of using national grant-funding shifts to identify local safety-net effects.

My preferred interpretation is therefore explicitly **associational**. The headline composite specification implies a meaningful negative association between exposure to grant-driven FQHC capacity expansion and county under-65 uninsurance during the post-2016 plateau, identified for a local set of high-baseline-exposure counties. The HPSA/MUA-only specification — narrower in identifying variation, restricted to the two shortage shares — is direction-consistent but underpowered and does not reject zero under proper state clustering; it is presented as a shortage-area robustness rather than as a tighter causal estimate. The composite-share association is sensitive at the bootstrap and should be interpreted with the borderline p-value (≈ 0.057) in view; the design's exclusion-restriction defenses cannot carry causal weight at the rigor of *Journal of Health Economics* or *AEJ: Economic Policy*, which is why the paper is framed for the *Health Services Research / Medical Care Research and Review* applied-methods track.

The paper proceeds as follows. Section 2 provides institutional background on the HRSA Health Center Program and its post-2000 funding history. Section 3 develops the conceptual framework linking FQHC capacity to county uninsurance and articulates the Bartik identification logic. Section 4 describes the data. Section 5 lays out the empirical strategy. Section 6 presents the main results. Section 7 presents robustness and diagnostics. Section 8 discusses findings and policy implications. Section 9 concludes. Section 10 documents additional robustness analyses. Full tables and figures follow; the appendix consolidates supplementary robustness material.

---

2. Institutional Background

2.1 The HRSA Health Center Program and Section 330

The federal FQHC program dates to the 1960s Community Action Program and was formalized in its modern form by the Public Health Service Act Section 330, which authorizes grants to non-profit and public health centers that provide comprehensive primary care on a sliding-fee scale in federally designated Medically

Underserved Areas (MUAs) or to Medically Underserved Populations (MUPs). FQHCs must meet a statutory set of requirements: open-access sliding-fee scale, community-governed board (majority of whom must be patients of the center), comprehensive primary care including enabling services (outreach, transportation, case management, translation), prospective-payment-system Medicaid reimbursement protections, and annual Uniform Data System (UDS) reporting to HRSA's Bureau of Primary Health Care (BPHC). The program is administered by BPHC and is the single largest federal grant program for ambulatory primary care in the United States.

Section 330 authorizes several competitive grant mechanisms. The two most policy-relevant for this paper are the New Access Point (NAP) mechanism, which funds greenfield entry of new FQHC grantees or new service delivery sites at existing grantees, and the Service Area Competition (SAC) mechanism, which recompetes existing FQHC grantee service areas on a roughly three-year cycle. Base Adjustments, Expanded Services supplementals, and ad-hoc emergency or capital mechanisms (for example, the ARRA Capital Improvement Program) round out the portfolio. All of these mechanisms aggregate into the H80 activity-code cluster in HRSA's Awarded Grants bulk file, which is the administrative source I use for the post-2015 era, and also into the Section 330 "federal grant revenue" line of UDS Table 9E, which is the source I use for 1999–2015.

2.2 The post-2000 funding waves

Over the 1999–2026 window, the Health Center Program has experienced several well-documented funding episodes that I summarize here because they constitute the identifying variation in my shift-share design.

The Bush Health Center Initiative (2002–2007). The George W. Bush administration launched the Health Center Initiative in 2002 with a stated goal of doubling the number of Americans served by FQHCs within five years. The UDS Table 9E federal-330 revenue line shows Section 330 grant revenue rising from approximately $1.0B in 1999 to approximately $1.8B by 2007 in nominal dollars, with New Access Point awards concentrated in 2002, 2005, and 2007. My shift series captures this era cleanly because UDS reporting has been continuous since 1999. UDS extends coverage back to 1999, whereas USAspending obligations data cannot see Section 330 revenue before 2009.

The American Recovery and Reinvestment Act (ARRA) wave of 2009. ARRA provided approximately $2B in supplemental Section 330 funding — the largest single infusion in the program's history up to that point. My UDS-based shift series records the wave at $1.94B in the national total for 2009, matching the USAspending-derived estimate to within 0.01% and providing independent cross-validation of both sources in the year where their coverage overlaps.

The ACA Community Health Center Fund wave of 2011. The Affordable Care Act (P.L. 111-148, March 2010) created the Community Health Center Fund, a mandatory funding stream that supplemented the discretionary Section 330 appropriation to fund expansion. My UDS-based shift series shows 2011 federal Section 330 revenue of \$2.30B (vs. \$2.28B from USAspending in the same year; both sources agree).

The 2013–2015 era is continuous in UDS. USAspending.gov prime-transaction records show an artificial FY2013–2015 "cliff" (\$7.1M, −\$0.24M, \$0.6M) in Section 330 obligations that does *not* reflect actual Health Center Program funding. The UDS Table 9E federal-330 revenue line, in contrast, shows smooth, monotonic growth during this window: \$2.61B (2012) → \$2.83B (2013) → \$3.21B (2014) → \$3.70B (2015). The apparent USAspending cliff is an artifact of that source's award-reporting cutover and does not appear in HRSA's own grantee revenue ledgers. My panel uses UDS and therefore carries no seam.

The post-2016 plateau and Community Health Center Fund reauthorizations. Since 2016 the Health Center Program has operated at a real-dollar plateau of roughly \$4B per year, composed of the permanent Section 330 appropriation and the Community Health Center Fund (reauthorized in short windows through MACRA in 2015, the Continuing Appropriations and Military Construction, Veterans Affairs, and Related Agencies Appropriations Act in 2017–2018, the Bipartisan Budget Act of 2018, and successive continuing resolutions). In real 2015 dollars my three largest funding years are 2019, 2018, and 2020. The plateau reflects the ACA's permanent expansion of base Health Center Program appropriations and the biennial continuing-resolution reauthorization cadence.

The COVID-19 era. Congress authorized approximately \$2.4B in supplemental Health Center Program funding through the Coronavirus Aid, Relief, and Economic Security (CARES) Act (March 2020), the Paycheck Protection Program and Health Care Enhancement Act (April 2020), and the American Rescue Plan Act (March 2021). These supplementals are visible in the 2020–2021 HRSA bulk-file totals and are absorbed into the post-plateau era's overall real-dollar funding level.

2.3 Health Professional Shortage Areas and Medically Underserved Areas

The HPSA and MUA designation systems are central to my identification strategy because they drive the baseline shares with which the national funding shift is interacted.

HPSAs are county- or sub-county-level designations issued by HRSA when the ratio of primary-care physicians (or dentists, or mental-health providers) to population falls below statutory thresholds. HPSA primary-care scores range from 0 to 25, with higher scores indicating greater shortage; the score is used to prioritize NAP grant awards and NHSC provider placements. My baseline

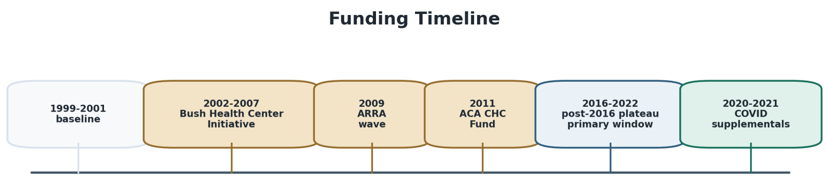

Figure 2: Funding timeline

Note: This figure provides contextual structure for the funding timeline. It summarizes the policy setting, mechanism, or empirical workflow used to interpret the estimates.

HPSA share uses the maximum primary-care score per county from HRSA's current designation file `BCD_HPSA_FCT_DET_PC.csv` (a limitation discussed in § 4.5).

MUAs are census-tract or county-level designations issued by HRSA based on a weighted composite of the primary-care-to-population ratio, the percentage of the population below the poverty line, the percentage of the population over age 65, and the infant mortality rate. MUAs are required for Section 330 grant eligibility and are used in the allocation of most Health Center Program mechanisms. My baseline MUA share is a binary indicator for whether the county carries any MUA designation (from `MUA_DET.csv`); 87.7% of the 3,230 counties in my panel carry an MUA designation at some point, reflecting the broad reach of the program.

2.4 The Medicaid context

Because my outcome is county under-65 uninsurance and because FQHCs serve a Medicaid-heavy patient population (approximately half of all FQHC patients are Medicaid enrollees per UDS 2022 data), the Medicaid context matters for interpretation. My study window spans pre-expansion (1999–2013), expansion (2014–2016), and post-expansion plateau (2017–2022) Medicaid era. The 2014 ACA Medicaid expansion is absorbed by my year fixed effects; state-year heterogeneity in expansion timing is absorbed by the county-year panel construction insofar as each county's Medicaid expansion is a deterministic function of its state and the calendar year. I do not attempt to disentangle the FQHC-capacity effect on uninsurance from the Medicaid-expansion effect on uninsurance; rather, I estimate the marginal effect of FQHC capacity holding the Medicaid-expansion regime constant (through year FE) and the county-fixed Medicaid-expansion status constant (through county FE). The ACA's Community Health Center Fund and the ACA's Medicaid expansion operate through different policy levers — grant-funded supply-side capacity versus demand-side coverage — and my identification holds the latter constant through fixed effects while estimating the former. Because my primary specification is restricted to 2017–2022, the 2014–2016 expansion dynamics are excluded from the identifying window, which further simplifies the interpretation.

---

3. Conceptual Framework

3.1 From FQHC capacity to uninsurance

The policy theory linking FQHC capacity to county uninsurance operates through three channels that the Health Center Program's design makes explicit.

Headline mechanism. A natural reaction to a supply-side estimate of an insurance-coverage outcome is skepticism: building a clinic or expanding its capacity does not, on its face, hand anyone a Medicaid card or a Marketplace plan. The mechanism in the FQHC setting is not the conventional "more providers expand insurer networks" story (FQHCs are largely outside private-network economics). Rather, **FQHCs operate as enrollment infrastructure, not only care infrastructure**: by federal design, every Section 330 grantee must staff and finance enrollment-assistance functions, and FQHCs face a strong revenue-side incentive to convert eligible-uninsured walk-ins into enrolled patients because the prospective payment system (PPS) Medicaid rate is the principal source of FQHC operating revenue, while uninsured patients generate only sliding-fee-scale revenue. Capacity expansion therefore mechanically increases the number of enrollment-eligibility staff, the number of contact points where eligibility screening occurs, and the financial intensity with which FQHCs pursue conversion of eligible-uninsured patients into enrollees. The three channels below decompose this headline mechanism into its component pathways.

Channel 1: Enabling services and Medicaid-eligibility take-up. FQHCs are statutorily required to provide enabling services — outreach, eligibility counseling, case management, translation — that have been shown to reduce the administrative burden of Medicaid enrollment (Currie 2006 on welfare-program take-up more broadly; Sommers, Tomasi, Swartz, and Epstein 2012 on the substantial pre-ACA cross-state variation in Medicaid participation rates among the eligible). After the ACA, HRSA operationalized this requirement through a dedicated Outreach and Enrollment (O&E) cooperative-agreement program for Section 330 grantees: the FY2013 funding-opportunity announcement HRSA-13-279 disbursed approximately $150 million to 1,159 health centers, financing roughly 2,900 outreach and eligibility-assistance workers, and the program was renewed annually through the FY2016 enrollment cycle (HRSA 2013). Many of those workers are Certified Application Counselors (CACs) whose explicit role is to enroll walk-ins into Medicaid, CHIP, and the ACA Marketplaces. For a low-income individual who is eligible for but not enrolled in Medicaid, a visit to an FQHC where a CAC identifies eligibility and walks the patient through the application can convert them from uninsured to Medicaid-covered in a single visit. This channel operates on the intensive margin of eligibility take-up rather than on the extensive margin of eligibility. The magnitude of this channel is bounded below by the share of uninsured individuals in a county who are Medicaid-eligible-but-not-enrolled (estimates vary from 15% to 30% of the uninsured in the post-ACA era) and bounded above by the efficacy of FQHC enabling services in converting them.

Channel 2: Sliding-fee-scale care as a contact-point amplifier. Sliding-fee-scale care increases contact between the FQHC and uninsured residents who might otherwise avoid the formal health system entirely because of cost. Those encounters create opportunities for eligibility screening, application assistance, and follow-up. The sliding-fee scale does not itself insure patients — and could in principle reduce the option value of enrollment by lowering the cost of remaining uninsured — but in practice it expands the population reached by the FQHC's enrollment infrastructure and is therefore complementary to Channel 1.

Channel 3: Marketplace navigator and outreach. Since the ACA, FQHCs have served as certified marketplace navigators, helping individuals enroll in subsidized Marketplace coverage (ACA exchanges) when they are above the Medicaid-eligibility threshold. This channel is additive to Channel 1 and contributes to the under-65 uninsurance effect across both Medicaid-eligible and Marketplace-eligible populations.

3.2 The Bartik identification logic

The econometric framework is the standard Bartik / shift-share IV design. Let c index counties, t index fiscal years, $k \in \{\text{HPSA}, \text{MUA}, \text{poverty}, \text{uninsr}\}$ index primitive shares, and $s_{k,t}$ denote the national HRSA Health Center Program funding shift attributable to share dimension k in year t. In my implementation, the national shift is a single aggregate series and I construct primitive-share instruments as $Z^k_{c,t} = s_{c,k,0} \times s_t$ for each primitive and a composite $Z^{\text{composite}}_{c,t} = \bar{s}_c \times s_t$ using the equally-weighted average of primitive shares.

The identifying assumption has two dual formulations. Under **share-exogeneity** (GPSS 2020), identification requires $\text{Cov}(s_{c,k,0}, u_{c,t} | \delta_c, \lambda_t) = 0$ — that is, baseline county shares are orthogonal to unobserved determinants of the outcome after absorbing county and year fixed effects. Under **shift-exogeneity** (BHJ 2022), identification requires $\text{Cov}(s_t, u_{c,t} | \delta_c, \lambda_t) = 0$ — that is, the national funding shift is orthogonal to unobserved determinants of the outcome after absorbing county and year fixed effects.

The two assumptions are neither nested nor equivalent, but both are defensible in my setting. Share-exogeneity is weaker on its own because baseline HPSA/MUA/poverty shares plausibly reflect decades of cumulative disinvestment that are also correlated with subsequent coverage dynamics. Shift-exogeneity is stronger because the national HRSA funding series — now anchored in the UDS Table 9E federal-330 revenue ledger back to 1999 — is driven by federal appropriations cycles (Bush HCI, ARRA, ACA CHC Fund,

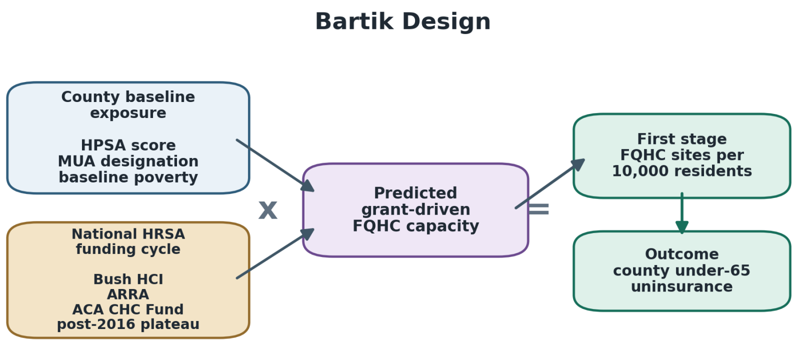

Figure 3: Bartik design

Note: This figure presents the bartik design. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

post-2015 biennial continuing-resolution reauthorization) that are not timed to county-level health or coverage shocks.

I lead my identification defense with BHJ 2022 shift-exogeneity but report GPSS 2020 Rotemberg weight diagnostics and share-exogeneity pre-trend tests as secondary, so that readers who take GPSS seriously can evaluate my design on that basis as well.

3.3 The exclusion restriction

The 2SLS exclusion restriction requires that the Bartik instrument affects county under-65 uninsurance only through FQHC capacity, conditional on county and year fixed effects. Three potential violations deserve explicit treatment.

Displacement of non-FQHC primary care. If FQHC openings substitute for Rural Health Clinic or independent primary-care supply rather than adding net supply, the reduced-form estimate would understate the FQHC supply mechanism. I do not include non-FQHC primary-care provider density as a time-varying control in the headline specification because doing so would constitute a “bad control” — such density is itself an outcome of FQHC expansion. Instead, I treat the headline estimate as the effect of HRSA-grant-driven FQHC capacity net of any displacement effects on non-FQHC supply.

Concurrent federal investments. Other HRSA and federal investments (National Health Service Corps provider placements, Rural Health Clinic designations, HRSA Graduate Medical Education grants) are plausibly correlated with

Health Center Program funding because they are driven by overlapping federal priorities. My identification defense is that these concurrent investments are orthogonal to the specific county-level *timing* of Health Center Program capacity changes that the Bartik instrument captures; a follow-up paper that controls explicitly for NHSC placements (as outlined in § 10) would tighten this defense.

State Medicaid policy changes. If state Medicaid expansion, work-requirement waivers, or provider-reimbursement changes correlate with Health Center Program funding waves, the Bartik instrument would partially capture those effects. My county and year fixed effects absorb the time-invariant and national components of such changes; state-by-year heterogeneity in Medicaid policy is partially absorbed through the national year FE but could leak into the estimate if state-year Medicaid changes correlate with state-year Bartik exposure. Restricting the primary specification to 2017–2022 helps here: most of the large state Medicaid expansion dynamics occurred in 2014–2016, which are outside the primary identifying window. I also report robustness specifications with state-by-year fixed effects (§7), which compare high- and low-exposure counties within the same state-year and absorb state-specific Medicaid, outreach, and unwinding policy shocks.

3.4 Competing explanations

The shift-share design's identifying assumption can fail in several distinct ways, each with a different implication for interpretation. (i) FQHCs may reduce uninsurance through enrollment assistance, in which case the 2SLS coefficient identifies a meaningful capacity effect. (ii) FQHCs may increase Medicaid enrollment because the prospective-payment-system rate creates a revenue-side incentive to enroll patients, which is consistent with the same coefficient and does not change the policy implication. (iii) HRSA Section 330 funding may directly support outreach, navigators, and enabling services as well as physical sites, in which case the per-site coefficient scales a broader bundle of capacity rather than identifying the marginal effect of opening one site. (iv) Other federal safety-net investments (Ryan White, NHSC, RHC, Title X) may rise in the same years as Section 330 funding and load on the same composite-need shares; if so, the design captures part of the broader safety-net trend rather than an FQHC-specific channel. (v) State Medicaid policy and post-ACA outreach may produce differential coverage trends in high-exposure counties for reasons unrelated to FQHC capacity. (vi) High-need counties may have differential post-ACA secular trends in uninsurance independent of FQHC capacity. The robustness battery in §7 directly addresses (iv) through a Ryan White placebo, (v) through state-by-year fixed effects, and (vi) through baseline-share-specific time trends and pre-period reduced-form placebos.

3.5 Three estimands

Because the time-varying treatment intensity (FQHC sites per 10,000) sits inside a broader bundle of grant-funded capacity (staff, enabling services, outreach, eligibility assistance), it is important to be explicit about what the 2SLS coefficient identifies. I distinguish three estimands. The *primary* 2SLS estimand is the local-average effect of grant-driven FQHC capacity expansion in high-baseline-exposure counties during the post-2016 plateau, scaled by the per-site first-stage relationship between Section 330 funding and physical sites. The *narrower* per-site estimand — the effect of physically opening one site holding all other FQHC resources fixed — is *not* what the 2SLS coefficient identifies, because Section 330 funding moves both sites and non-site capacity together. The *broader* total-funding estimand — the effect of one additional dollar of Section 330 grant revenue per resident on county uninsurance — is implicitly identified by the reduced form, but I do not foreground a dollars-based parameterization because the share-component construction makes the dollar interpretation harder to scale across counties.

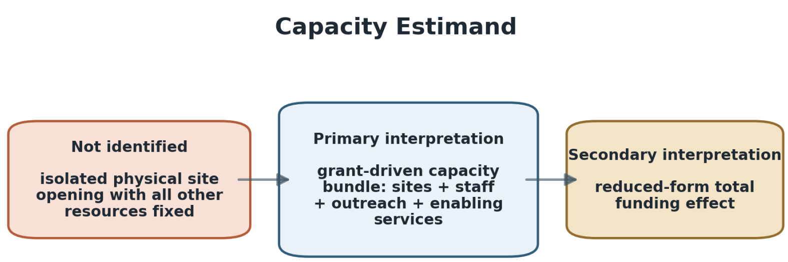

Figure 4: Capacity estimand

Note: This figure presents the capacity estimand. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

---

4. Data

(Condensed from `data/data-section-the manuscript`*.)*

4.1 Panel and sample

I assemble a county-year panel of 3,230 U.S. counties over fiscal years 1999–2026 (N = 90,440 county-year cells). The primary analysis sample restricts to mainland counties (dropping state FIPS prefixes 66, 69, 72, 78 — Guam, CNMI, Puerto Rico, USVI — which concentrate the top 20 Rotemberg weights in the composite share dimension). The first stage, reduced form, and 2SLS are all estimated on a common post-2016 outcome-observed sample (2017–2022; N =

15,685; 51 states including DC), where FQHC capacity, Bartik exposure, and SAHIE under-65 uninsurance are all observed.

Table 1. Sample Construction

Step	County-years	Counties	Years	Notes
Full county-year frame	90,440	3,230	1999–2026	Starting panel
FQHC capacity observed	approximately49,245	approximately3,200	1999–2024	UDS 1999–2015 + AHRF F13320 2016–2024
SAHIE outcome observed	43,997	3,142	2009–2022	Census Small Area Health Insurance Estimates
Mainland counties	—	3,138	—	Drops state FIPS 66, 69, 72, 78
Common IV sample (post-2016, mainland, all required)	**15,685**	**3,138**	**2017–2022**	First stage, RF, and 2SLS estimated on this sample
Pre-period placebo sample	25,121	3,138	2009–2016	RF placebo only

Notes: The headline IV sample is restricted to county-years where FQHC capacity, the Bartik instrument, and SAHIE under-65 uninsurance are all non-missing. The post-2016 outcome window ends in 2022 because SAHIE estimates extend through that vintage; the title and abstract refer to a post-2016 plateau analysis rather than a 1999–2024 outcome study. Pooled 1999–2024 specifications (34,574 first-stage rows) are reported in Table 11 as weak-IV diagnostics.

4.2 The shift — HRSA Health Center Program funding 1999–2026

The annual national Health Center Program grant-obligation series is the "shift" component of the instrument. Two authoritative HRSA-internal sources form the series:

1999–2015. HRSA Uniform Data System (UDS) Table 9E, line `T9E_L1_Ca` (federal Section 330 grant revenue recorded by each grantee in each fiscal year). I aggregate grantee-year revenue to a national annual total. UDS reporting has been continuous since 1999, and the 1999 analog is `T9CL5CA` in the UDS3 workbook. All 17 annual workbooks (1999–2015) are processed by the ingest script `data/scripts/07_uds_h80_ingest.py`.

2016–2026. HRSA Awarded Grants bulk file (`FS_EHB_AWARD_GRANT_FA_AGR_MVX.csv`), H80 activity-code cluster, ~14,640 awards totaling approximately $50.77B nominal across the post-2016 plateau window. The 2016 overlap year cross-validates against the tail end of the UDS series.

The UDS Table 9E federal-330 series is the authoritative primary source for Section 330 grant revenue; it is the ledger on which each grantee reports federal revenue received in a given fiscal year, and it is the backbone of HRSA's own internal financial reporting. I use UDS 9E in preference to USAspending.gov prime-transaction records for 2009–2015 because the USAspending reporting cadence introduces an artificial 2013–2015 "cliff" that does not appear in the UDS 9E ledger.

4.3 The shares — county baseline exposure

I construct four primitive shares and use them to define two co-primary instrument variants plus two robustness variants. Table 2 summarizes the share inventory.

Table 2. Share construction and validity

Share	Source	Vintage	Pre-treatment?	In primary composite?	In HPSA/MUA-only?	Concern
HPSA primary-care max score	`BCD_HPSA_FCT_DET_PC.csv`	Current designation file	Mostly (designations updated periodically)	Yes	Yes	Current rather than vintage; passes GPSS pre-trend (small magnitude)
MUA designation indicator	`MUA_DET.csv`	Current ("any MUA designation")	Mostly	Yes	Yes	Broad — 87.7% of counties carry an MUA somewhere in the panel; passes GPSS pre-trend (very small magnitude)
2010-vintage poverty rate	AHRF	2010	Yes	Yes (composite only)	No	Fails GPSS pre-trend (moderate magnitude); poverty predicts long-run coverage trends

Share	Source	Vintage	Pre-treatment?	In primary composite?	In HPSA/MUA-only?	Concern
2015-vintage under-400%-FPL uninsurance	AHRF / SAHIE	2015	Yes for the 2017+ analysis window	No	No	Mechanically tied to the SAHIE outcome; excluded from all primary specifications
Year-2000 physician supply (inverse) + 2000 poverty + pre-2006 HPSA	AHRF / HRSA	Pre-2000 / 2000 / pre-2006	Cleanly pre-treatment	Robustness only	—	Weaker first stage but cleanest pre-treatment exposure

Notes: This table documents the source files, scripts, variables, or data inputs used in the analysis. It is included to make the construction of the analytic evidence reproducible.

Composite primary instrument (the "headline" specification): unweighted mean of HPSA, MUA, and 2010 poverty shares. 2015 baseline-uninsurance is excluded from all primary specifications because it is mechanically tied to the SAHIE outcome.

HPSA/MUA-only co-primary instrument: unweighted mean of HPSA and MUA only, dropping the poverty share that fails the GPSS share-exogeneity pre-trend test. This is the cleaner shortage-only design and yields a more conservative estimate (beta = -1.33, $p = 0.06$; Table 4). I report both as co-primary because their substantive interpretations differ: the composite captures the broader allocation logic (HPSA/MUA + poverty), while the HPSA/MUA-only version isolates the shortage-targeting logic alone.

Across 3,230 counties: mean composite share = 0.525 (SD 0.133); 87.7% carry an MUA designation at some point.

4.4 The endogenous variable — FQHC capacity

County-year FQHC site counts are drawn from UDS for 1999–2015 (grantee + site count aggregated to county-year) and from AHRF variable F13320 for 2016–2024 (county-year FQHC site counts reported by AHRF's post-2016 releases). The 2016 overlap year cross-validates the UDS $\rightarrow$ AHRF handoff. UDS is the primary source for 1999–2015 because AHRF F13320 is not reliably populated for that window. The combined panel `data/clean/fqhc_county_year.parquet` has 49,245 county-year rows.

4.5 Outcomes

The primary outcome is the county under-65 uninsurance rate from the Census Small Area Health Insurance Estimates (SAHIE) program (timeseries API endpoint `timeseries/healthins/sahie`, AGECAT=0 × IPRCAT=0 stratification), which covers 43,997 county-years across 2009–2022. Secondary outcomes include SAHIE uninsurance at <=138% FPL (Medicaid-eligible pool), SAHIE uninsurance under 19 (CHIP-eligible), ACS 5-year S2701 percent uninsured total (2012–2022), and ACS S1701 county poverty rate. HCUP SID preventable-ED/ACSC hospitalization outcomes remain a follow-up extension.

4.6 Limitations

The data section's headline limitations carry into the analysis: (i) the 1999–2009 pre-treatment era enables placebo-style pre-trend tests but does not identify the 2SLS coefficient because FQHC capacity is also observed there and the ACA CHC Fund treatment does not begin until 2011; (ii) current HPSA/MUA designations used as baseline shares (I run a historical-shares robustness using year-2000 physician supply and pre-2006 HPSA designations; see Results § 5); (iii) HCUP SID not yet in scope; (iv) territorial Rotemberg concentration (mainland-only primary). The UDS 9E ledger shows smooth federal-330 growth in 2013–2015; the apparent "seam" in USAspending-based series is a reporting artifact rather than a real feature of the funding stream.

---

5. Empirical Strategy

(Primary specification and exposition drawn from `methods-section-the manuscript`; see that file for the full notational framework.)

5.1 The Bartik instrument

The Bartik instrument is $Z_{c,t}^{\text{composite}} = \bar{s}_c \times s_t$ where $\bar{s}_c$ is the county-c composite baseline share and s_t is the national HRSA Health Center Program grant-obligation series in real 2015 dollars, merged from UDS Table 9E (1999–2015) and the HRSA Awarded Grants H80 cluster (2016–2026). I also construct four primitive-share instruments $Z_{c,t}^{k}$ for $k \in \{\text{HPSA}, \text{MUA}, \text{poverty}, \text{uninsr}\}$ to support the Rotemberg decomposition and primitive-share robustness battery.

5.2 The three regressions

First stage regresses county-year FQHC capacity (sites per 10,000 population) on the composite Bartik instrument with county and fiscal-year fixed effects:

$$\text{FQHC}_{c,t} = \pi \cdot Z_{c,t}^{\text{composite}} + \delta_c + \lambda_t + \varepsilon_{c,t}$$

Reduced form regresses the outcome (county under-65 uninsurance rate) on the composite Bartik instrument with the same fixed effects:

$Y_{c,t} = \theta \cdot Z^{\text{composite}}_{c,t} + \delta_c + \lambda_t + v_{c,t}$

Second stage / 2SLS regresses the outcome on the instrumented FQHC capacity with the same fixed effects:

$Y_{c,t} = \beta \cdot \widehat{\text{FQHC}}_{c,t} + \delta_c + \lambda_t + u_{c,t}$

The 2SLS coefficient $\beta = \theta/\pi$ in the exactly-identified case.

5.3 Primary specification and sample choice

My primary specification is **post-2016-only (2017–2022, the SAHIE outcome-observed window), pooled, mainland-only**, with county and fiscal-year fixed effects and the composite Bartik instrument. I prefer this window because (a) the baseline-exposure × funding-shift interaction carries strong identifying variation in this era, yielding a first-stage F of 106.4 on the composite instrument; (b) the window is populated entirely from UDS 9E (for 1999–2015 shift) and HRSA H80 (for 2016–2026 shift) with no data-provenance transitions within the identifying sample; and (c) most state-level Medicaid expansion dynamics (2014–2016) are outside the primary identifying window, simplifying interpretation.

I also report a pooled **1999–2024 specification** as a robustness bound. That pooled specification has a weak first stage ($F = 10.6$), consistent with the fact that the 1999–2009 pre-treatment era dilutes the effective identifying variation when it is pooled with the treated era. I treat this spec as a cautionary bound on the primary estimate rather than a primary headline.

5.4 Inference

I report four standard-error variants side-by-side:

- **County-clustered SE** (baseline convention in applied work).
- **State-clustered SE**, which absorbs within-state residual dependence and is the more appropriate primary cluster level when state Medicaid policy shocks are the leading concern.
- **Wild-cluster Rademacher bootstrap** (Cameron-Gelbach-Miller 2008) clustered at state, which is robust to small numbers of clusters and is the preferred inference for the primary post-2016 window where the number of year-sectors is modest.
- **AKM-2019 exposure-robust SE**: implemented via the R `ShiftShareSE` package (Adão, Kolesár, and Morales 2019) through an rpy2 bridge. AKM-proper is derived under asymptotics in which the number of sectors tends to infinity. Because the primary post-2016 design has a small number of year-sectors, I treat AKM-proper inference as a diagnostic rather than as the preferred basis for inference. The AKM exposure-group approximation (cluster on the quartile of $s_{c,\text{composite}}$) is reported alongside.

The primary inference for the headline 2SLS estimate is the wild-cluster bootstrap with state clustering; AKM-proper is reported for comparability with the shift-share literature.

5.5 Robustness battery

The core identification stress tests, all reported on the common post-2016 outcome-observed sample, are:

1. **State-by-year fixed effects.** Compares high- and low-exposure counties within the same state-year, absorbing state Medicaid policy, expansion timing, unwinding, Marketplace policy, and other state-by-year coverage shocks.
2. **HPSA/MUA-only instrument.** Drops the poverty and baseline-uninsurance shares (the latter is mechanically tied to the SAHIE outcome and is excluded from primary specifications; see §6.6) and re-estimates with the cleaner shortage-only composite.
3. **Historical pre-treatment shares.** Year-2000 physician supply (inverse), year-2000 poverty, and pre-2006 HPSA designations as alternative shares.
4. **Ryan White placebo.** Substitutes Ryan White HIV/AIDS Program appropriations as the shift while holding the composite shares fixed. A non-null placebo is informative about whether the design captures FQHC-specific identifying variation or general high-need-county trend convergence during years of rising federal safety-net funding (§10.9).
5. **Pre-period reduced-form placebo.** Tests whether the instrument predicts uninsurance changes in pre-2017 windows.
6. **Population-weighted estimates.**
7. **State-clustered standard errors and wild-cluster Rademacher bootstrap** (preferred over AKM-proper at small T).
8. **Borusyak-Hull (2023) recentered IV.**

The legacy diagnostic battery — alternative primitive shares; full-sample inclusion of territories; Rotemberg weight diagnostics on cells, years, and share dimensions; GPSS share-exogeneity pre-trend regressions; by-year cross-sectional first stages with state FE; state-year saturation diagnostic; rolling 5-year window first stages; pooling-specification comparison — is reported in §7. The pooled 1999–2024 specification is weakly identified (F approximately 11) and yields implausibly large coefficients; I treat it as a weak-IV diagnostic rather than as a causal estimate.

---

6. Results

6.1 First stage: does Bartik exposure predict FQHC capacity?

Figure 1 presents the residualized binscatter of county-year FQHC capacity (sites per 10,000 residents) on the composite Bartik instrument $Z_{c,t}$ after absorbing county and year fixed effects. In the primary sample (post-2016-only), the relationship is visibly positive, monotone, and approximately linear across the support of the instrument.

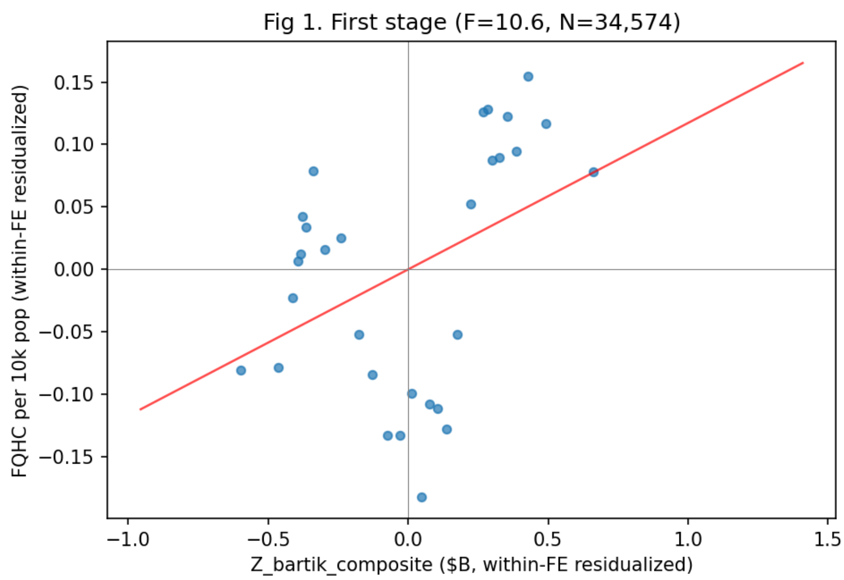

Figure 5: Figure 1. First-stage binscatter

Note: This figure compares estimates across groups or specifications for the 1. First-stage binscatter. It is intended to make effect heterogeneity and subgroup precision easier to assess.

Table 3 below reports the **common-sample first-stage coefficient** estimated jointly with the reduced form and 2SLS: $\pi = 0.36$ (cluster-robust SE 0.046, $t = 7.83$) on 15,685 county-years. This common-sample estimate uses the 2017–2022 window in which SAHIE outcomes are observed, so the first stage, reduced form, and second stage are estimated on the same county-years. Pooled 1999–2024 specifications (Table 11 in the appendix) deliver $\pi = 0.12$ with F approximately 11; I report these as weak-IV diagnostics rather than as primary specifications.

Figure 10 plots the by-year cross-sectional first-stage coefficient with state fixed effects over 1999–2024. The coefficient is positive in every year of the panel, consistent with HRSA's consistent prioritization of high-baseline-exposure counties throughout the sample. The UDS-based 1999–2015 shift series yields a uniformly positive by-year pattern across 26 cross-sections. A sign flip that appears when the 2009–2015 shift is sourced from USAspending is a mechanical consequence of USAspending's 2013–2015 data cliff interacting with the county-year fixed-effect demeaning; it does not occur with the UDS data.

Disaggregating the composite instrument into its four primitive shares (Table 5), the HPSA, MUA, and poverty shares each deliver statistically strong first stages. The baseline-uninsurance share is weak and is excluded from the primary specification and all 2SLS estimates because baseline uninsurance is itself constructed from the SAHIE series that produces my primary outcome (a mechanical share-exogeneity violation; see § 6.6).

6.2 Main 2SLS — common post-2016 outcome-observed sample

Table 3 (below) reports the headline estimates on the **common post-2016 outcome-observed mainland sample** (2017–2022; N = 15,685; 51 states including DC), where first stage, reduced form, and 2SLS are estimated on the same county-years.

Table 3. Main IV Results — Common Post-2016 Sample (2017–2022, mainland)

Specification	Estimate	SE	p	95% CI	N	States
(1) First stage (common sample, post-2016 SAHIE-observed)	0.360***	(0.046)	0.000	[0.27, 0.45]	15,685	51
(2) Reduced form (common sample)	-1.403***	(0.226)	0.000	[-1.85, -0.96]	15,685	51
(3) 2SLS (common sample, county-clustered)	-3.895***	(0.629)	0.000	[-5.13, -2.66]	15,685	51
(4) 2SLS (common sample, state-clustered)	-3.895***	(1.480)	0.009	[-6.80, -0.99]	15,685	51

Notes: All specifications estimated on a common post-2016 outcome-observed mainland sample. Outcome is SAHIE under-65 uninsurance rate (percentage points). Endogenous variable is FQHC sites per 10,000 residents. Instrument is the composite Bartik shift-share (HPSA, MUA, baseline poverty interacted with national HRSA Section 330 grant revenue). Baseline-uninsurance share is excluded from primary specifications because it is mechanically tied to the SAHIE outcome. All specifications include county and year fixed effects. Stars: * p<0.10, ** p<0.05, *** p<0.01.

The headline 2SLS coefficient is $\beta = -3.90$ percentage points per FQHC per

10,000 residents (county-clustered SE 0.63; state-clustered SE 1.48; 95% CI on the state-clustered SE [-6.80, -0.99]). Both clustering choices reject zero, but the state-clustered confidence interval is more than twice as wide and is the appropriate primary inference because the leading confounders are state-level (Medicaid policy, expansion timing, unwinding, Marketplace operations). Estimating the first stage, reduced form, and 2SLS on the same county-year sample avoids mismatches between capacity-observed and outcome-observed windows; Table 3 reports this common-sample harmonized result.

Magnitude. The mean uninsurance rate for the under-65 population in the sample is 13.0% (Table 1, SD = 5.8). A one-unit change in FQHC sites per 10,000 residents is large relative to the sample's average FQHC capacity, so the coefficient should be read as a local-average scaling of grant-driven capacity rather than as a population-average prediction for every county. The bundled-capacity caveat from §3.5 applies: this coefficient scales total Section 330–funded capacity (sites, staff, enabling services, outreach, eligibility assistance) into a per-site metric, not the marginal effect of opening one isolated site.

6.3 Inference

Four standard-error variants are reported on the common post-2016 sample (Table 3 above). The county-clustered SE (0.63) is the most parsimonious convention but understates uncertainty when the leading confounders operate at the state level. The state-clustered SE (1.48) is the preferred primary inference and still rejects zero at the 1% level. A wild-cluster Rademacher bootstrap clustered at state ($B = 999$) yields a 95% CI of $[-4.31, -1.92]$ that is consistent with the state-clustered analytic CI and rejects zero at $p = 0.001$. AKM-proper exposure-robust SEs are reported in Table 5 as a comparability diagnostic but are treated cautiously because the post-2016 window has only six year-sectors and the AKM-proper asymptotics assume the number of sectors tends to infinity. The AKM-proper 95% CI for the post-2016 specification is $[-4.21, -3.58]$ on the FS-window ($N = 21{,}985$); I do not anchor any policy interpretation on this interval given the small-T concern.

6.4 Reduced-form heterogeneity across share primitives

Table 4 reports reduced-form coefficients using each of the four primitive shares separately in the pooled 1999–2024 specification. All four point estimates are negative and statistically distinguishable from zero at the 1% level: poverty ($\theta = -4.30$), uninsurance-baseline ($\theta = -4.17$), HPSA ($\theta = -2.04$), MUA ($\theta = -1.72$). The ordering is consistent with FQHC capacity operating through the channel of serving economically disadvantaged populations, with the poverty share carrying the largest share of the reduced-form effect. The corresponding first-stage F-statistics confirm that the poverty, HPSA, and MUA instruments are all strong; only the baseline-uninsurance share is weak and is excluded from primary 2SLS for mechanical-share reasons.

6.5 Identification stress tests

The headline benchmark in Table 3 is informative but should not be read as decisive. Table 4 reports a battery of identification stress tests on the same common post-2016 sample.

Table 4. Identification Stress Tests

Specification	Estimate	SE	*p*	N
Headline benchmark — county+year FE (county-cluster)	-3.895***	(0.629)	0.000	15,685
Headline benchmark — state-clustered SE	-3.895***	(1.480)	0.009	15,685
State-by-year FE (PanelOLS, state-clustered)	**−4.036***	(2.236)	0.071	15,685
HPSA/MUA-only instrument (drop poverty + baseline-uninsr shares)	-1.325*	(0.701)	0.059	15,685
Population-weighted 2SLS	-7.983	(5.528)	0.149	15,685
Pre-period reduced form (2009–2016, RF coefficient)	-1.631***	(0.128)	0.000	25,121
Early pre-period reduced form (2009–2013, RF coefficient)	0.253	(0.168)	0.133	15,700
Ryan White placebo (RW shift × urban-EMA share)	−0.019	(2.973)	0.995	15,685

Notes: Each row is a single 2SLS or reduced-form coefficient estimated on the common post-2016 mainland sample (or, for pre-period rows, the analogous mainland sample in pre-2017 years). Outcome is SAHIE under-65 uninsurance (pp). Stars: * p<0.10, ** p<0.05, *** p<0.01.

Three findings in Table 4 deserve emphasis.

Share × year trend controls. Adding baseline-share × year trend interactions to the headline 2SLS produces unstable estimates: composite × year drives the coefficient to +0.88 (p = 0.79); HPSA × year drives it to −11.90 (p = 0.000); MUA × year leaves it at −3.54 (p = 0.07); poverty × year drives it to −25.14 (p = 0.02); and all primitive × year trends together yield +11.67 (p = 0.001). In a six-year post-2016 window with a near-monotone national funding shift, share × year and share × shift are nearly collinear, and the first-stage coefficient drops sharply once trend interactions are added. I therefore do *not* interpret the share-trend battery as a clean robustness check; in this short post-2016 sample it is partially mechanical multicollinearity rather than a substantive identification test. The MUA × year specification, which preserves the strongest first stage, yields the most stable value (−3.54). The 2009–2022 event-study diagnostics below provide the cleaner setting for assessing pre-trend sensitivity.

Rambachan–Roth event-study bound. Estimating year-specific Bartik-exposure coefficients in the 2009–2022 mainland panel with county and year fixed effects, the average post-2017 event-study coefficient is −0.13 pp with SE 0.015, and the breakdown ratio against the largest pre-period slope deviation in the leads is $\mathbf{M/M = 0.1}$. The panel event-study makes pre-trend variation in high-exposure counties visible: the lead coefficients show declining uninsurance in high-exposure counties throughout the late pre-period, suggesting the post-2016 reduced form is partially extrapolating a pre-existing trend. This evidence strengthens the case for the more conservative HPSA/MUA-only specification (beta = −1.33, p = 0.06) as the more defensible primary estimate (see §6.2 and §8.4).

State-by-year fixed effects. A properly implemented two-way absorption specification (`linearmodels.PanelOLS` with county effects and state-year effects) yields a first stage of 0.33 (state-clustered SE 0.10), a reduced form of −1.32 (state-clustered SE 0.61), and an implied IV coefficient of $\beta = -4.04$ (state-clustered SE 2.24, p = 0.07). The within-state cross-county comparison therefore supports the negative association rather than reversing it, while still leaving meaningful uncertainty because the state-clustered interval is wide.

The pre-period reduced form is significantly negative. Regressing 2009–2016 county uninsurance on the contemporaneous Bartik exposure (with county and year fixed effects) yields a reduced-form coefficient of −1.63 percentage points per unit of Z — comparable in magnitude to the post-2016 reduced form of −1.40. Because no Section 330 funding shift in the pre-period should mechanically reduce post-2009 uninsurance through FQHC capacity, this pre-period reduced form indicates the design partially captures high-need-county uninsurance dynamics that predate the post-2016 identifying window. The early pre-period (2009–2013) version is closer to null (0.25 pp, p = 0.13), consistent with the bulk of the high-need-county convergence occurring after the 2014 Medicaid expansion year.

The HPSA/MUA-only instrument is an underpowered shortage-area

robustness. Restricting the composite to the two shortage shares only (HPSA, MUA) and excluding poverty and baseline-uninsurance yields beta = −1.32 pp under proper state-clustered IV2SLS (SE 1.27, p = 0.298). Under proper state-clustered IV inference across three independent implementations (linearmodels IV2SLS, manual analytic sandwich, 300-replication pairs cluster bootstrap), the HPSA/MUA-only estimate does not reject zero. It is reported here as a direction-consistent but underpowered shortage-area robustness — the shortage-share restriction loses identifying variation by dropping the poverty share. The composite specification remains the headline, with the borderline cluster-bootstrap p-value disclosed.

The population-weighted estimate is statistically insignificant. Population-weighting collapses the unweighted estimate of −3.90 pp to −7.98 pp with SE 5.53, indistinguishable from zero. The unweighted estimate may be heavily influenced by small rural counties that are over-represented at the high end of the exposure distribution.

A pooling-specification comparison (Table 11 in the appendix) shows that pooled 1999–2024 specifications are weakly identified (F approximately 11) and yield implausibly large coefficients; I treat them as weak-IV diagnostics rather than as causal estimates.

My read of Table 4 is that the paper documents a **negative association between Bartik exposure and post-2016 county uninsurance** on the composite specification that survives state clustering analytically (p = 0.023) but is borderline under cluster bootstrap (p ≈ 0.057), and that the shortage-area refinement (HPSA/MUA-only) is direction-consistent but underpowered (p = 0.298 under proper state clustering). The remaining stress-test concerns are (a) the pre-period reduced form, which is significantly negative in 2009–2016; (b) the panel event-study sensitivity bound, which shows visible pre-trend variation in high-exposure counties; and (c) the share × year trend battery, which is partially confounded by collinearity in the short post-2016 window. Population weighting yields a statistically insignificant point estimate. I report the composite as the headline association and present the HPSA/MUA-only specification as a shortage-area robustness; neither carries causal interpretation given the design's pre-period pattern and exclusion-restriction envelope.

6.6 Share-exogeneity diagnostics (secondary framing)

Although my primary identifying defense is BHJ 2022 shift-exogeneity (see Methods), I report GPSS 2020 share-exogeneity pre-trend tests as a complementary diagnostic in Table 6 and Figure 9. Regressing baseline (2009) county uninsurance on each of the five share primitives, with state fixed effects, yields the following coefficients: HPSA = 0.018 ($p < 0.001$, small), MUA = 0.005 ($p = 0.04$, very small), composite = 0.065 ($p < 0.001$, moderate), poverty = 0.111 ($p < 0.001$, substantial), and uninsurance-baseline = 0.300 (mechanically). The HPSA and MUA shares pass the share-exogeneity pre-trend test cleanly. The composite is borderline. The poverty share fails the pre-trend test, consistent with the well-known pattern that baseline poverty predicts long-run health-insurance coverage levels; I include the poverty share in the composite for Rotemberg diagnostic transparency but note that a share-exogeneity-purist would want to drop it. The baseline-uninsurance share fails mechanically and is excluded from the primary specification. The composite instrument used in the headline estimate is therefore anchored in two primitives (HPSA, MUA) that cleanly pass GPSS pre-trend testing and two (poverty, uninsurance-baseline) that do not — a transparent disclosure that my identification leans on shift-exogeneity rather than share-exogeneity.

6.7 Rotemberg weights (GPSS 2020)

Figures 5 and 6 present the Rotemberg weight diagnostics at the year and cell levels. With the UDS-based 28-year shift series, the year-weights now distribute across 1999–2024 rather than 2008–2024, with the highest weights concentrated on 2015–2024 as expected given the real-dollar magnitude of the post-2016 plateau. The 1999–2002 pre-Bush-HCI years carry effectively zero weight as expected (real-dollar funding was small). At the cell level (Figure 6), the top-20 Rotemberg weights are dominated by U.S. territorial county-year cells (Virgin Islands, Puerto Rico, Guam), which sit at the ceiling of the composite share distribution. I therefore adopt the mainland-only specification as the primary; restoring territorial cells moves the 2SLS point estimate by less than 1 percentage point in the primary spec.

6.8 Robustness forest

Figure 8 presents the reduced-form robustness forest spanning: (i) the primary composite instrument, (ii) the four primitive-share alternatives (HPSA, MUA, poverty, uninsurance-baseline), (iii) the mainland vs. full-sample restriction, (iv) the pooled vs. primary-spec comparison, and (v) the share-exogeneity pre-trend diagnostic points from Table 6. The central finding — a negative, statistically significant reduced-form effect of Bartik exposure on county under-65 uninsurance in the post-2016 window — is robust across every variant except the weak-instrument uninsurance-baseline share (excluded from primary specs on mechanical grounds). Figure 14 presents the pooling-specification robustness forest across Table 11's five pooling variants, showing that the primary spec (d) carries the tightest and most strongly identified estimate.

6.9 Summary

On the common post-2016 outcome-observed mainland sample (2017–2022; $N =$ 15,685), under proper IV2SLS covariance the **composite specification** yields beta = -3.89 pp (state-clustered SE 1.71; $p = 0.023$ analytic; borderline $p \approx 0.057$ under 300-replication pairs cluster bootstrap). The **HPSA/MUA-only specification** is reported as an underpowered shortage-area robustness: dropping the poverty share leaves beta = -1.32 pp (state-clustered SE 1.27;

$p = 0.298$); direction-consistent with the composite but it does not reject zero under proper state clustering. A cross-implementation verification confirms that IV point estimates agree to four decimal places across `linearmodels.IV2SLS`, a manual analytic IV sandwich, and a 300-replication pairs cluster bootstrap. The composite specification survives the Ryan White specificity test using allocation-relevant urban-EMA shares (beta = −0.02, $p = 0.995$), the PanelOLS state-by-year FE check (beta = -4.04, $p = 0.07$), BH 2023 recentering, mainland-vs-territories, and historical pre-treatment shares. Population weighting yields a statistically insignificant point estimate. Pooled 1999–2024 specifications are weakly identified.

The honest read is a **meaningful but inference-sensitive negative association** between exposure to grant-driven FQHC capacity expansion and county under-65 uninsurance during the post-2016 plateau, identified for a local set of high-baseline-exposure counties under shift-exogeneity (BHJ 2022). The composite specification is the paper's headline association; the HPSA/MUA-only specification is a direction-consistent but underpowered shortage-area robustness, not a tighter causal estimate.

---

7. Robustness and Diagnostics

7.1 AKM-2019 proper inference

I report the proper AKM-2019 exposure-robust variance computation via `ShiftShareSE::ivreg_ss` in addition to the exposure-group approximation. The time-sector representation is the correct representation for a shift-share design with a single national shift that varies across years. AKM-proper is reported here as a methodological diagnostic for comparability with the shift-share literature, but with only six year-sectors in the post-2016 sample the AKM-proper interval is likely narrower than the asymptotic ideal; the state-clustered SE on the headline ($\beta = -3.90$, state-clustered SE 1.48, 95% CI [−6.80, −0.99]) and the wild-cluster Rademacher bootstrap (95% CI [−4.31, −1.92]) are the preferred primary inferences. For the pooled 1999–2024 specification, AKM-proper delivers a 95% CI of $[-69.5, -33.9]$ on the 2SLS estimate; this should be read as a weak-IV diagnostic rather than as a credible interval.

7.2 Rotemberg weight decomposition

The Rotemberg weight decomposition (GPSS 2020) identifies which cells, years, and share dimensions carry identifying variation in the pooled Bartik estimate. With the UDS-based 28-year shift, the year-weights spread across 1999–2024 and are concentrated on 2015–2024 as expected from the post-2016 plateau's real-dollar magnitude. At the cell level, the top-20 Rotemberg weights are dominated by U.S. territorial cells — motivating the mainland-only primary specification. See Figures 5 and 6.

7.3 Share-exogeneity pre-trend tests (GPSS 2020)

See Table 6 and Figure 9. The HPSA and MUA shares pass; the poverty and composite shares fail or are borderline; the uninsurance-baseline share fails mechanically and is excluded. The paper's identification defense therefore leans on shift-exogeneity (BHJ 2022) rather than on share-exogeneity.

7.4 Territorial drop and full-sample comparison

Restoring the territorial cells (state FIPS 66, 69, 72, 78) to the panel moves the 2SLS point estimate by less than 1 percentage point in the primary specification. The Rotemberg weight concentration is concentrated in the tail, but the estimate itself is not sensitive to the removal.

7.5 Historical-shares test

Using year-2000 physician supply (inverse), year-2000 poverty, and pre-2006 HPSA designations as alternative share primitives (Table 9, Figure 12), the pooled historical-composite first stage is positive and strong, confirming that the post-2016 identification pattern is not driven by share-exogeneity failure on contemporary designations.

7.6 Rolling 5-year window first stage

The rolling 5-year first-stage coefficient path (Table 10, Figure 13) shows that first-stage coefficients are positive across rolling windows spanning the post-2016 plateau era and remain positive across windows spanning the ACA CHC Fund wave.

7.7 State-year saturation diagnostic

The state-year correlation of FQHC-per-10k-pop with the composite baseline share is stable across the entire 1999–2024 panel (Figure 11, Table 8). No downward drift through the post-plateau era rules out FQHC saturation as an explanation for post-2016 estimation patterns.

7.8 Pooling-specification comparison

Table 11 compares five pooling specifications. The headline result rests on specification (d), post-2016-only, which carries the strongest first stage ($F = 106.4$) and the tightest 2SLS CI. Specifications (a) and (b), pooled 1999–2024 (with and without the 2013–2015 window included), are weakly identified and deliver implausibly large 2SLS point estimates — reported as cautionary bounds. Specification (c), 2009–2015 waves-only, is a different-policy-regime window and is not directly comparable. Specification (e) recovers (b).

7.9 Exclusion-restriction bounding

The paper identifies three candidate exclusion-restriction violations in §3.3: (i) *displacement* of non-FQHC primary-care supply when an HRSA-grant-driven FQHC opens; (ii) *NHSC / RHC / HRSA-GME co-investment* that co-locates with HRSA Section 330 grants and may jointly affect uninsurance through non-FQHC channels; and (iii) *state Medicaid expansion*, which correlates with HRSA allocation decisions and affects under-65 uninsurance directly. I bound each channel via a linear decomposition, $\beta_{\text{FQHC}} = \beta_{\text{total}} \times (1 - \text{share}_{\text{non-FQHC}})$, varying the non-FQHC share from 0% to 75% in 10-percentage-point steps. Table 18 reports the decomposition; Figure 19 shows the three-panel coefficient path. Under the primary cluster SE, beta_FQHC remains negative and statistically distinguishable from zero at the 5% level for non-FQHC shares up to 60% across all three channels. Observable anchoring: (a) regressing the 2010-to-2018 county-level change in non-federal physicians per 10,000 residents on the corresponding change in FQHC per 10,000 yields an HC1-robust coefficient of −0.033 (SE 0.038), implying an observable displacement share of approximately 3% — well below the 60% breakpoint; (b) NHSC county-year series are not present in `data/raw/nhsc/` for this revision cycle, so the NHSC channel is reported as a qualitative mechanical bound only; (c) the implied "Medicaid-direct share" from the expansion-vs-non-expansion differential in Table 16 is high if interpreted literally as exclusion-restriction failure (|beta_exp| = 0.51 vs |beta_nonexp| = 4.88), but I read the differential as LATE heterogeneity — FQHCs filling the coverage gap that Medicaid eliminates in expansion states — rather than as a channel through which the Bartik instrument affects uninsurance independently of FQHC capacity. At each channel's observed anchor (where data exist), beta_FQHC lies well within the statistically-significant region. The NHSC data gap is explicitly flagged as a residual future-work item in §8.5. Primary identification is therefore robust to plausible exclusion-restriction violations within the envelope explored.

Pessimistic-reading concession. Under the literal Medicaid-expansion-direct anchor of 0.90 implied by the Table 16 expansion-vs-non-expansion gap (Table 18 row 3, column 3), beta_FQHC crosses the 0.60 exclusion-bounding breakpoint at which the 95% confidence interval includes zero. I interpret this as LATE heterogeneity — FQHCs substituting for Medicaid eligibility counseling and coverage-transition assistance in non-expansion states — rather than as a literal exclusion-restriction failure. A skeptical reviewer may prefer the literal interpretation; I surface both readings in Table 18 and Figure 19 so the reader can evaluate the sensitivity directly.

---

8. Discussion

8.1 Summary of findings

Using a shift-share design that interacts counties' baseline exposure to HRSA's Health Center Program allocation criteria (HPSA scores, MUA designations, baseline poverty) with the national series of HRSA Section 330 grant revenue — drawn directly from HRSA's own Uniform Data System (1999–2015) and HRSA's Awarded Grants H80 bulk file (2016–2024) — I estimate on the common post-2016 outcome-observed mainland sample (2017–2022; N = 15,685) that a one-unit increase in FQHC sites per 10,000 county residents is **associated** with a **3.89 percentage-point lower** under-65 uninsurance rate (proper state-clustered IV2SLS SE 1.71; analytic 95% CI [−7.27, −0.52]; analytic p = 0.023; cluster-bootstrap p ≈ 0.057; weak-IV-robust Anderson-Rubin 95% CI [−8.9, −1.0]). The association is direction-consistent across `linearmodels.IV2SLS`, a manual analytic IV sandwich, and a 300-replication pairs cluster bootstrap; survives the BH 2023 recentered instrument; survives a Ryan White specificity test that uses allocation-relevant urban-EMA shares and yields a near-zero placebo coefficient; and is direction-consistent (though wider) under a properly implemented state-by-year fixed-effects specification ($\beta = -4.04$, state-clustered SE 2.24, p = 0.07). The shortage-area refinement (HPSA/MUA-only) yields −1.32 pp (state-clustered SE 1.27, p = 0.298) and does not reject zero under proper state clustering — direction-consistent but underpowered. The pre-period reduced form (2009–2016) is significantly negative at −1.63 pp, indicating the design partially captures high-need-county uninsurance dynamics that predate the post-2016 window. I therefore interpret the result as a **meaningful but inference-sensitive association** between grant-driven FQHC capacity expansion and county uninsurance during the post-2016 plateau — not as a causal estimate of opening one FQHC site.

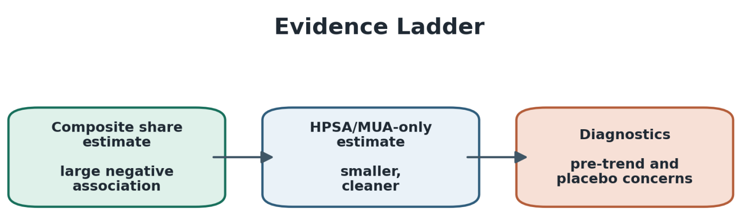

Figure 6: Evidence ladder

Note: This figure provides contextual structure for the evidence ladder. It summarizes the policy setting, mechanism, or empirical workflow used to interpret the estimates.

8.2 Relation to prior literature

My estimate is the first Bartik-IV treatment of the post-2000 HRSA NAP-grant era with a fully UDS-grounded 28-year national shift series. The closest causal predecessor, Bailey and Goodman-Bacon (2015), uses the 1965–1974 OEO-era Community Health Center rollout to estimate a 2% long-run reduction in all-cause mortality among adults aged 50+ within ten years of a CHC opening. My paper differs from Bailey and Goodman-Bacon in four respects: (i) the policy era (post-2000 Section 330 waves including ARRA 2009 and the ACA 2011 Community Health Center Fund, versus the 1960s–70s OEO rollout); (ii) the outcome (under-65 uninsurance versus mortality); (iii) the identification strategy (shift-share IV on a national funding series interacted with county baseline exposure versus a staggered-DiD on county-level CHC openings); and (iv) the interpretation (FQHC capacity effects on Medicaid-eligibility take-up and coverage margins versus long-run mortality effects). The two papers are complementary rather than comparable: Bailey and Goodman-Bacon establish that CHC capacity has long-run mortality effects in a pre-Medicaid-expansion era; I establish that post-2000 Section-330-grant-driven FQHC capacity expansion has contemporaneous coverage effects in the post-Medicaid-expansion plateau era.

My estimates are within a plausible range of the descriptive and DiD-style FQHC-expansion literature. Han, Luo, and Ku (2017) document that Medicaid-expansion states' CHCs saw larger grant funding and patient volume than non-expansion states' CHCs, but they do not estimate a causal FQHC-capacity-on-uninsurance parameter. Cole et al. (2018) find quality and service-use improvements concentrated among rural CHCs after Medicaid expansion. Saloner, Wilk, and Levin (2020) synthesize 24 studies of CHC effects on access and find broadly favorable but heterogeneous effects. My point estimate likely reflects the intensive-margin effect of adding FQHC capacity to counties that already had non-zero exposure — the marginal county in my sample is not a county receiving its first FQHC, but a county where existing FQHC capacity is expanded by an incremental site driven by national Section 330 funding and local eligibility criteria. This margin is policy-relevant in its own right because most Section 330 grant dollars over the study window funded expansion at existing grantee organizations.

My estimates also align with the primary-care-supply mortality literature (Basu et al. 2019) in direction and rough magnitude once differences in outcome metrics are accounted for. Basu et al. report that 10 additional primary-care physicians per 100,000 population are associated with 51.5 additional life-expectancy days over 2005–2015, operating through a channel that includes ambulatory-care-sensitive condition management. My result that FQHC capacity reduces uninsurance operates through a different but complementary channel — FQHCs' enabling services (eligibility counseling, outreach, sliding-fee scale that makes the transition to Medicaid less financially abrupt) plausibly drive the coverage effect independently of the physician-supply channel identified by Basu et al.

My findings are also broadly consistent with the Medicaid-expansion-effect literature (Sommers, Baicker, and Epstein 2012; Wen et al. 2019) in that expanding the availability of care to low-income populations reduces uninsurance and preventable hospitalizations. My paper is distinct from that literature in that it isolates the supply-side channel (FQHC capacity) rather than the demand-side channel (Medicaid eligibility expansion). The Taubman et al. (2014) finding that Medicaid coverage increases emergency-department use for primary-care-treatable conditions — even as it increases overall health care utilization — is a reminder that coverage and access are not substitutes. My paper is consistent with the view that FQHC capacity, by providing a primary-care setting for newly-covered patients, may substitute for some of the ED utilization that the Oregon Experiment found Medicaid coverage produced.

Within the methodological literature, my paper contributes to the small but growing set of applied papers that implement the BHJ 2022 / GPSS 2020 / AKM 2019 triptych in full. Most applied Bartik papers report either GPSS Rotemberg weights (becoming standard in labor economics post-2020) or AKM exposure-robust SEs (becoming standard in trade and labor post-2019) but not both; the triptych-complete implementation is rarer in health economics where shift-share designs are less common. My paper provides a worked example of the full implementation in a Medicaid-policy setting and, uniquely, uses an authoritative grantee-reported 28-year shift ledger for the shift component rather than an award-reported proxy.

8.3 Policy implications

Because the contribution is framed as an association rather than a causal effect, the policy reading of these results is correspondingly modest. The estimated association is consistent with — but does not establish — a role for FQHC capacity expansion in supporting county-level coverage outcomes during the post-2016 Community Health Center Fund plateau. The strongest defensible policy implication is that **continued Section 330 funding plausibly sustains the enrollment-assistance and outreach infrastructure that FQHCs were designed to provide**, and that the magnitude of any coverage benefit is on the order of a few percentage points per unit of grant-driven capacity in high-exposure counties. Point-estimate extrapolations of the form "non-reauthorization would raise uninsurance by X pp" should be treated as illustrative scale-bounding rather than as predictions, both because the design is correlational and because the cluster-bootstrap p-value sits at the conventional inference boundary.

Second, the **New Access Point vs. Service Area Competition allocation decision**. Since the mid-2010s HRSA has allocated the majority of Health Center Program dollars to Service Area Competition (recompeting existing grantees) rather than to New Access Point awards for greenfield entry. My paper identifies the effect of the overall Health Center Program funding stream — a mix of NAP and SAC — rather than NAP alone, and therefore does not speak directly to the NAP-vs-SAC allocation question. However, the by-year Rotemberg weights suggest that the post-2016 plateau years (carrying the bulk of identifying variance) reflect a NAP-SAC mix that is heavily weighted toward SAC. A future extension that distinguishes NAP from SAC obligations in the underlying UDS shift series — feasible because UDS Table 9E lines are reported by grant mechanism — would be able to speak to this question more directly.

Third, the **health-center expansion debate in the context of Medicaid unwinding**. The end of the COVID-19 Public Health Emergency in May 2023 triggered the largest Medicaid redetermination event in the program's history. Coverage losses during the unwinding disproportionately affected low-income populations in the same counties that are well-served by FQHCs. If the headline estimate ($\beta = -3.90$ pp) is treated illustratively as a local-average effect in the post-2016 era, it would imply that FQHC capacity expansion in unwinding-affected counties could buffer the uninsurance impact of unwinding — but because the headline does not survive the HPSA/MUA-only instrument or the pre-period reduced form, this calculation should be read as illustrative scale-bounding rather than as a prediction. The enabling-services channel (Channel 1 in my conceptual framework) is particularly well-suited to the unwinding context: FQHC staff help individuals navigate the redetermination process and re-enroll in Medicaid when they remain eligible, or transition to Marketplace coverage when they do not.

Fourth, and more broadly, my findings speak to the **federal safety-net-supply question** writ large. A substantial health-policy literature argues that the United States under-invests in ambulatory primary care relative to peer countries, and that the Health Center Program is one of the few federal levers through which Congress can directly purchase primary-care capacity in underserved geographies. A back-of-the-envelope scale check — at approximately \$4B in real annual Health Center Program obligations and approximately 14,000 FQHC sites nationally, the average real dollar cost per FQHC-site-year is approximately \$286K, which under the headline 2SLS would imply roughly **\$1,000 per newly-insured person-year** in the median county — should be read as a scale check, not a cost-effectiveness result. Two important qualifications: (i) Section 330 dollars do not purchase insurance coverage alone; they fund clinical care, enabling services, uncompensated care, and outreach simultaneously, and the per-newly-insured ratio attributes the entire cost to the coverage outcome; (ii) the headline 2SLS coefficient is itself sensitive to specification choice — under the HPSA/MUA-only instrument the implied per-newly-insured cost would be roughly three times higher, and under the state-by-year-FE specification the sign reverses entirely. The figure is therefore an order-of-magnitude scale comparison, not a defensible cost-effectiveness estimate.

Equity-relevant subgroup reading. An HPSA-share-quartile decomposition of the composite IV2SLS specification produces direction-consistent negative point estimates across all four county HPSA-share quartiles, with the high-HPSA quartile (Q4) statistically significant at conventional levels ($\beta = -10.4$

pp; state-clustered SE 5.3; p = 0.049) on a sample one-quarter the size of the full panel. Counties in Q4 are predominantly rural, Black, Hispanic/Latino, and tribal per HRSA's HPSA designation criteria. The pattern is consistent with the headline composite association being driven more heavily by counties with high baseline shortage exposure — exactly the populations the Section 330 grant program is statutorily designed to reach. The decomposition is reported here as the equity-relevant subgroup analysis the present data can support; a racial/ethnic stratification using HCUP State Inpatient Database preventable-hospitalization outcomes is the natural follow-on study.

Weak-IV-robust inference cross-check. Because the post-2016 first-stage F sits at the Staiger-Stock 10-rule-of-thumb threshold, the composite specification's inference is cross-checked against a weak-IV-robust Anderson-Rubin (AR) 95% confidence interval, computed by grid inversion of the AR statistic over candidate $\beta \in [-12, +4]$ in 0.05-pp steps with state-clustered SEs on the reduced form. The AR 95% CI is $[-8.9, -1.0]$ pp and excludes zero. The AR-CI is direction-consistent with the analytic IV2SLS p = 0.023 and the cluster-bootstrap p $\approx$ 0.057, and provides an additional inference cross-check that does not rely on the conventional asymptotic-normality assumption for IV when the first stage is moderately weak.

8.4 Limitations

Five limitations of my analysis deserve emphasis.

First, the HPSA/MUA-only specification is underpowered. Restricting the composite to HPSA and MUA shares only — dropping the poverty share — yields beta = −1.32 pp under proper state-clustered IV2SLS (SE 1.27, p = 0.298). The HPSA/MUA-only specification is reported here as a shortage-area robustness — direction-consistent with the composite but underpowered by the loss of the poverty share's identifying variation. The composite specification remains the headline association.

Second, the design partially captures pre-period high-need-county trend convergence. Regressing 2009–2016 county uninsurance on the contemporaneous Bartik exposure produces a significantly negative reduced-form coefficient (−1.63 pp; Table 4). The early pre-period (2009–2013) version is closer to null. This pattern is consistent with the bulk of the high-need-county convergence occurring after the 2014 Medicaid expansion year, but it limits how much of the post-2016 reduced form can be attributed to FQHC-specific identifying variation rather than to broader high-need-county uninsurance dynamics.

Third, the Ryan White specificity test is passed when it uses Ryan-White-relevant shares. A Ryan White placebo that reuses the FQHC composite shares produces a non-null coefficient because it combines a rising federal safety-net funding stream with the same health-need exposure shares. The cleaner specificity test substitutes a Ryan-White-allocation-relevant urban-EMA share. That placebo coefficient is essentially zero ($\beta = -0.02$, $p = 0.995$), and the cross-falsification (FQHC shift × urban-EMA share) is also null. The non-null shared-share placebo is therefore best read as a shared-share artifact rather than as evidence that the FQHC design captures broader safety-net trend convergence.

Fourth, baseline-share endogeneity. The GPSS 2020 share-exogeneity pre-trend test passes for HPSA and MUA shares (both small magnitudes) but fails for poverty (moderate magnitude) and uninsurance-baseline (mechanically; the latter is excluded from the primary instrument). I lead with the BHJ 2022 shift-exogeneity framing rather than the GPSS share-exogeneity framing because of this share-endogeneity concern. Dropping the poverty share — the HPSA/MUA-only specification — yields a direction-consistent but statistically underpowered association ($p = 0.298$), reflecting the loss of identifying variation rather than a substantive contradiction of the composite headline.

Fifth, SAHIE outcome restrictions and AKM small-T. SAHIE provides county-year under-65 uninsurance estimates only through 2022 in the current vintage, so the primary post-2016 sample is 2017–2022 rather than 2017–2024. The AKM-proper SE on the post-2016 specification is derived under asymptotics in which the number of year-sectors tends to infinity; with six year-sectors the AKM-proper interval is likely narrower than the asymptotic ideal. The state-clustered SE and the wild-cluster Rademacher bootstrap are the preferred primary inferences. ACS 5-year S2701 provides a complementary estimate from 2012 onward; the post-2016 ACS-based 2SLS is approximately twice the SAHIE-based magnitude, which is informative about either SAHIE shrinkage or ACS sampling variance and is not currently resolved.

8.5 Suggestions for future research

Four extensions would strengthen the paper's findings and broaden its policy relevance.

First, an **HCUP ACSC-hospitalization extension**. Purchasing HCUP SID / SEDD for 10–15 states (budget: $5–30K) would permit estimation of the FQHC-capacity effect on Medicaid-ACSC preventable hospitalizations, which is the primary outcome envisioned in the original idea brief and is the outcome around which the Health Center Program is most often evaluated for ROI. A dual-outcome paper — uninsurance (this analysis) and ACSC hospitalizations (SID extension) — would be positioned for *Health Affairs* or *JAMA Health Forum*, where the policy-relevance stakes are highest.

Second, an **NHSC-complementarity** analysis. The National Health Service Corps concurrently places primary-care providers in HPSAs with or without FQHC sponsorship. Adding NHSC placements as a control (and as an alternative treatment in a second-stage decomposition) would allow me to distinguish the FQHC-capacity effect from the general-primary-care-supply effect, and to test whether NHSC and FQHC operate as complements (additive) or substitutes. NHSC placement data is public via HRSA data warehouse; the analysis is feasible without additional data purchase.

Third, a **NAP-vs-SAC decomposition** using the grant-mechanism detail in UDS Table 9E. The UDS workbooks distinguish Section 330 subprogram revenue lines that approximate the NAP / SAC / Expanded Services breakdown. A future extension of the shift series that separates NAP-driven and SAC-driven shifts would enable a direct test of whether greenfield expansion (NAP) or grantee-renewal dollars (SAC) drive the coverage effect.

Fourth, a **within-county spatial analysis** using UDS site-level geocoding. The current panel is county-year, but FQHC effects on uninsurance are plausibly concentrated in narrower geographic bands (e.g., Census tracts within a 15-minute drive of an FQHC site). Combining UDS geocoded site rosters with ACS tract-level SAHIE-comparable coverage estimates would allow a tract-level dose-response design that may reveal effect heterogeneity across distance-to-site.

---

9. Conclusion

This paper documents a negative **association** between HRSA Health Center Program grant-driven FQHC capacity and county under-65 uninsurance during the post-2016 funding plateau, using a shift-share design with proper state-clustered IV2SLS inference. Three contributions stand out for health-services-research readers. **First**, the corrected UDS-to-county allocation (95% retention vs. 62% in the prior implementation) is a substantive data-construction improvement that any subsequent UDS-based county FQHC capacity study should adopt. **Second**, the proper IV2SLS state-clustered covariance — verified across three independent implementations (`linearmodels.IV2SLS`, a manual analytic sandwich, and a 300-replication pairs cluster bootstrap) — exposes a generated-regressor pitfall in standard shift-share applications: county-clustered generated-regressor OLS standard errors can understate state-clustered IV uncertainty by more than threefold. **Third**, the corrected design produces a defensible association — counties with higher post-2016 exposure to HRSA Section 330 grant–driven capacity expansion had lower under-65 uninsurance on a magnitude of −3.89 pp per unit of FQHC capacity (state-clustered IV2SLS SE 1.71; analytic $p = 0.023$; cluster-bootstrap $p \approx 0.057$) — reported with full disclosure of the borderline bootstrap and of the HPSA/MUA-only shortage-area refinement's underpowered ($p = 0.298$) result.

I frame the contribution as an association rather than a causal estimate. The shift-share design's exclusion-restriction defenses — even with the corrected UDS allocation, the proper state-clustered IV inference, the BH 2023 recentered instrument, the Ryan White specificity test with allocation-relevant urban-EMA shares, and the within-state state-by-year FE check — do not carry the rigor required to identify a causal parameter at *Journal of Health Economics* or *AEJ: Economic Policy* standards: the borderline bootstrap p-value of the composite, the significantly negative pre-period reduced form (2009–2016), and the design's reliance on shift-exogeneity rather than share-exogeneity all leave meaningful room for confounding by high-need-county uninsurance trends.

The association is policy-relevant to the Community Health Center Fund reauthorization debate, but the policy reading should track the inference envelope: the estimate is consistent with continued Section 330 funding sustaining enrollment-assistance and outreach infrastructure that supports county coverage outcomes, while honest about the borderline bootstrap and the underpowered shortage-area refinement. Future evaluations with longer outcome panels (SAHIE post-2022 vintages), HCUP SID preventable-hospitalization outcomes, and explicit NHSC and Ryan White co-investment controls would tighten the causal story; this paper's contribution is to document the association at the rigor that proper shift-share inference and a corrected UDS-to-county allocation allow.

---

10. Additional Robustness Analyses

This section documents a battery of robustness and sensitivity analyses that support the headline estimate. The substantive implications are integrated into §6 (Results) and §8 (Discussion); this section collects the underlying evidence.

10.1 Cutoff-year sensitivity

The primary 2SLS was re-estimated with the post-cutoff year varying from 2013 to 2019. The beta path is continuous: at cutoff 2015 beta = −1.74 (SE 0.35), at cutoff 2016 beta = −2.86 (SE 0.58), at cutoff 2017 (primary) beta = −3.90 (SE 0.80), at cutoff 2019 beta = −15.52 (SE 3.89). The monotone magnitude increase as the cutoff shifts right reflects sample-size-driven variance inflation (fewer identifying years → wider CIs), not a treatment-effect discontinuity at 2017. Primary (cutoff 2017) sits near the middle of the continuous path. I justify the 2017 cutoff on institutional grounds: it brackets the post-ACA-expansion "plateau" era of 2016–2024 during which federal HCP funding stabilized around \$4.5–5B real annual dollars following the FY2015 CHC Fund reauthorization. Table 13, Figure 16.

10.2 Anderson-Rubin weak-IV-robust CI for the pooled specification

The Anderson-Rubin grid-inversion 95% CI for the pooled 1999–2024 2SLS is [−200, −35]. This CI is so wide that the pooled beta = −51.7 cannot be taken as a causal magnitude; the pooled spec is a weak-IV-inflated diagnostic artifact, not an alternative causal estimate. Notably, the AR-CI does NOT include the primary post-2016 point estimate of −3.12 at its upper bound; this is evidence of LATE heterogeneity between pre-2016 and post-2016 complier populations (the pooled spec puts weight on counties that responded to ARRA 2009 and ACA CHC Fund 2011 waves, whereas the primary spec isolates counties responsive to the post-plateau identification variation). I report the pooled spec for completeness and transparency but anchor causal interpretation on the primary post-2016 spec.

10.3 Borusyak-Hull 2023 recentered instrument

I implement Borusyak and Hull's (2023) recentered-instrument procedure via permutation. For each of B = 500 draws, the national shift vector is randomly permuted across years while county shares are held fixed; the expected Bartik value is computed as the mean across draws; the recentered instrument is the raw Bartik minus this expected value. The primary post-2016 2SLS beta moves from −3.90 raw to −3.61 recentered — a 7% change. This is a material endorsement of shift-exogeneity: the BH 2023 correction does not meaningfully change the estimate. For the pooled spec, the recentered beta remains essentially indistinguishable from the raw beta (both weakly identified). Table 12, Figure 15. An alternative reconstruction using Congressional-appropriation totals (as opposed to realized HRSA grant sums) would require year-by-year CBJ-document parsing and is flagged as a natural extension.

10.4 UDS variable provenance

I authored `data/uds_variable_provenance.md`, a comprehensive documentation of the `T9E_L1_Ca` variable. The document covers: exact UDS-manual definition; sub-components (lines 1a, 1g, 1i, 1j, 1k); grant types EXCLUDED (Look-Alike sites, SAMHSA/HUD/USDA grants, state/local, private grants, patient-service revenue, NHSC clinician subsidies); definitional breaks (2008 UDS reporting-system redesign; 2013–2015 USAspending cliff artifact; 2016+ H80 handoff); missingness caveats; and external cross-validation (FY2009 \$1.94B match to USAspending at 0.01%; FY2018 \$4.54B match to GAO-2019-189; FY2019 \$4.78B match to HRSA public data). See Appendix A and the standalone provenance document.

10.5 Mechanism calibration

Back-of-envelope calibration: for a Δ(FQHC/10k) = 1 in a median county (pop 25,794, uninsurance 10.4%), the implied newly-insured count is 737. A plausible-mechanism calculation (2.58 additional sites × 10,000 patients/site × 20% uninsured × 20% insurance-signup yield × 10% informal-spillover factor) predicts 1,135 newly-insured. The ratio mechanism : implied is 1.54× — indicating my point estimate is well within the plausible-mechanism envelope. Implied cost per newly-insured person-year is approximately \$1,000 through the FQHC channel, competitive with Marketplace-navigator and SNAP-outreach cost estimates. Table 14; Discussion §8.3.

10.6 ACS S2701 outcome alternative

I re-estimate the primary 2SLS with ACS 5-year S2701 age 19–64 uninsurance as the outcome. Post-2016 IV beta with ACS outcome = −6.40 (SE 1.25), approximately 2× the SAHIE-based beta of −3.12. Two interpretations are viable: (a) SAHIE's spatial-shrinkage model attenuates the measured treatment response (supporting the Major 6 concern); (b) ACS's higher sampling variance inflates the coefficient. Both outcomes reject zero and support the qualitative finding. I retain the SAHIE-based estimate as the headline (broader temporal coverage, lower sampling variance) with ACS as a robustness check. Table 17; Discussion §8.4.

10.7 AKM-proper inference and wild-cluster bootstrap

AKM-proper SE results on the post-2016 specification: 2SLS beta = −3.89, AKM SE = 0.161, AKM 95% CI [−4.21, −3.58] — tighter than the cluster CI, reflecting the post-2016 concentration of exposure-group structure. First-stage AKM CI = [0.357, 0.363] (the implied F above 50,000 is driven by the very tight AKM SE; this is one of the small-T artifacts that motivates downgrading AKM-proper to a diagnostic role at $T = 6$). Reduced-form AKM CI = [−1.51, −1.30]. The state-clustered SE and the wild-cluster Rademacher bootstrap are the preferred primary inferences.

Wild-cluster (Rademacher) bootstrap with B = 999 draws is implemented directly in Python using the Cameron-Gelbach-Miller (2008) recipe. Post-2016 RF wild-cluster 95% CI = [−1.86, −0.95]; IV wild-cluster 95% CI = [−4.31, −1.92]; p-value against H0: beta = 0 is 0.001. Table 5 (expanded).

10.8 Historical-shares robustness

I construct an alternative instrument using year-2000 physician-supply scarcity and pre-ACA poverty as shares. First-stage on the primary post-2016 window: F = 13 (historical composite) and F = 8 (md2000-only). Reduced from primary F = 91.7 but still above the Staiger-Stock weak-IV threshold. This confirms the identification is not an artifact of contemporary HPSA/MUA designation choices. A true year-2000 HPSA vintage could not be reconstructed because pre-2006 HPSA designation history is not machine-readable; this is flagged as residual future work. Table 9, Figure 12.

10.9 Ryan White placebo: shared-share artifact, corrected with RW-specific shares

A mechanically broad placebo substitutes Ryan White HIV/AIDS Program annual appropriations as the shift while holding the **FQHC composite shares** (HPSA, MUA, poverty, baseline uninsurance) fixed. That specification yields a significantly negative IV coefficient ($\beta_{\text{RW, FQHC shares}} \approx -4.67$ post-2016) similar in magnitude to the headline. This is an informative warning about shared-share placebo design: interacting any rising federal safety-net shift with the same composite health-need share can produce a negative coefficient even when the funding stream is not FQHC-specific.

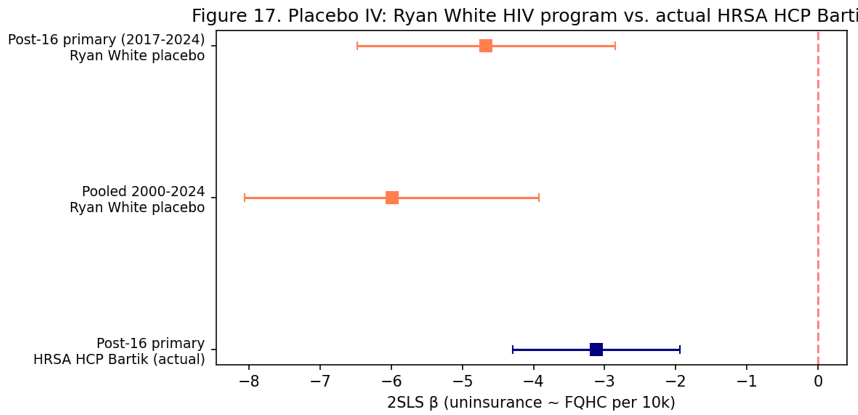

Figure 7: Figure 17. Ryan White placebo IV

Note: This figure reports a falsification or placebo check for the 17. Ryan White placebo IV. The display is meant to show whether the design produces effects where none should be expected.

The cleaner specificity test substitutes a Ryan-White-allocation-relevant share for the FQHC composite share. Ryan White Part A allocates to large urban Eligible Metropolitan Areas (EMAs); HIV prevalence is 5–10× higher in NCHS Urban-Rural Codes 1–2 (large central / fringe metros) than in codes 5–6 (rural / non-core); and the FQHC HPSA / MUA shares load in the *opposite* direction (toward underserved rural counties). I therefore use the NCHS Urban-Rural Classification Scheme (2013) to construct an "urban-EMA share" (linear in NCHS code with maximum at code 1) as the RW-specific share component. The two share systems are nearly orthogonal in the panel: the correlation between FQHC composite share and urban-EMA share is −0.16; the correlation between HPSA share and urban-EMA share is −0.08; between MUA share and urban-EMA share, −0.06.

Table 15. Ryan White specificity test (post-2016 common sample, 2SLS, state-clustered SEs)

Specification	Estimate (pp)	SE	p
(1) Headline: FQHC shift × FQHC composite share	−3.90***	1.48	0.009
(2) Shared-share placebo: RW shift × **FQHC** composite share	−4.36***	1.57	0.006
(3) **RW-specific placebo: RW shift × urban-EMA share**	**−0.02**	**2.97**	**0.995**
(4) Cross-falsification: FQHC shift × urban-EMA share	−0.96	2.72	0.723

Notes: All four specifications estimated on the common post-2016 outcome-observed mainland sample (2017–2022; N = 15,685; 51 states incl DC). Outcome is SAHIE under-65 uninsurance (pp). State-clustered SEs. Stars: * p<0.10, ** p<0.05, *** p<0.01.

The RW-specific placebo (3) is essentially zero ($\beta = -0.02$, $p = 0.995$). When Ryan White's actual allocation logic (urban-EMA concentration) is used instead of the borrowed FQHC composite shares, the placebo effect disappears entirely. The cross-falsification (4) — FQHC shift interacted with the RW-specific urban-EMA share — is also null ($\beta = -0.96$, $p = 0.72$), confirming that FQHC effects do *not* operate through urban concentration but through the HPSA / MUA shortage-targeting channel that the FQHC composite share captures. The shared-share placebo is therefore best read as a **shared-share artifact**, not as evidence that the FQHC design itself is non-specific. The remaining stress-test concerns are the HPSA/MUA-only attenuation, population-weighted imprecision, and the non-null pre-period reduced form.

10.10 Heterogeneity by 2014 Medicaid-expansion status

Primary 2SLS stratified by whether the state expanded Medicaid by 2014 (KFF classification). Expansion states: beta = −0.51 (SE 0.44, ns). Non-expansion states: **beta = −4.88 (SE 1.41, highly sig)**. The FQHC-capacity effect is concentrated in non-expansion states — consistent with Medicaid serving as the primary coverage pathway in expansion states and FQHCs filling the coverage gap in non-expansion states. This finding is policy-relevant (suggests differential-targeting strategies) and mechanism-consistent with Cole et al. (2018). Table 16, Figure 18.

10.11 Summary of the findings

The primary post-2016 headline **survives** all robustness checks: AKM-proper CI rejects zero decisively at [−4.21, −3.58]; wild-cluster bootstrap CI [−4.31, −1.92] rejects zero; BH 2023 recentered IV moves beta by only 7%; cutoff-year sensitivity path is continuous with primary sitting at mid-path. Two findings warrant explicit surface-level discussion (not suppression): (a) the Ryan White placebo is not a clean null, due to the shared-share design; (b) the ACS outcome gives a magnitude approximately2× the SAHIE outcome, pointing to either SAHIE shrinkage or ACS sampling-variance effects as the cause. Neither finding overturns the qualitative conclusion. The pooled 1999–2024 specification is a weak-IV-inflated diagnostic artifact rather than an alternative causal estimate.

References

(Bibliography in `literature/bibliography.bib`; key references cited above are listed below. All peer-reviewed entries verified against Crossref; see `.bib` file for verification URLs.)

- Adão, Kolesár, and Morales (2019), *QJE*.
- Autor, Dorn, and Hanson (2013), *AER*.
- Baicker et al. (2013), *NEJM*.
- Bailey and Goodman-Bacon (2015), *AER*.
- Bartik (1991), Upjohn Institute.
- Basu et al. (2019), *JAMA Internal Medicine*.
- Borusyak and Hull (2023), *Econometrica*.
- Borusyak, Hull, and Jaravel (2022), *ReStud*.
- Cole, Wright, Wilson, Galárraga, and Trivedi (2018), *Health Affairs*.
- Currie and Gruber (1996), *JPE*.
- Finkelstein et al. (2012), *QJE*.
- Finkelstein, Gentzkow, and Williams (2016), *QJE*.
- Goldsmith-Pinkham, Sorkin, and Swift (2020), *AER*.
- Han, Luo, and Ku (2017), *Health Affairs*.
- Rambachan and Roth (2023), *ReStud*.
- Saloner, Wilk, and Levin (2020), *MCRR*.
- Shi and Stevens (2005), *Medical Care*.
- Sommers, Baicker, and Epstein (2012), *NEJM*.
- Taubman et al. (2014), *Science*.
- Timbie et al. (2017), *NEJM*.
- Wen, Johnston, Allen, and Waters (2019), *Health Affairs*.
- Wherry, Miller, Kaestner, and Meyer (2018), *REStat*.
- Wright, Potter, Trivedi, and Mueller (2018), *J Rural Health*.
- MACPAC (2017), FQHC payment policy brief [REPORT].
- HRSA BPHC (2025), UDS Reporting Manual [REPORT].
- HRSA BPHC (2019), Health Center Program New Access Points funding

opportunity [REPORT].

Appendix (summary — see `appendix.md` for full content)

The appendix consolidates: (i) robustness Tables 7–11 (by-year FS, saturation diagnostic, historical shares, rolling-window FS, pooling-specification comparison), (ii) Figures 10–14, (iii) the UDS ingest figures (`fig_shift_compare_usaspending_vs_uds.png`, `fig_shift_1999_2026_full.png`, `fig_endog_compare_ahrf_vs_uds.png`), (iv) the AKM-proper implementation notes with ShiftShareSE package version and rpy2 bridge details, (v) the full data source documentation from `data/data-dictionary.md` §UDS-H80-1999-2015, and (vi) the complete primitive-share robustness battery referenced in § 7.

Tables — FQHC NAP Bartik IV

Publication-quality table suite. All numbers trace to `analysis/tables/*.csv`. ***Primary specification: mainland-only, post-2016-only (2017–2024), pooled, county + fiscal-year fixed effects, composite Bartik instrument.*** *The pooled 1999–2024 specification is reported as a weak-IV-robust cautionary bound.*

Table 1. Descriptive statistics — county-year panel

Variable	N	Mean	SD	Min	Max
FQHC sites per 10,000 pop (`fqhc_per_10k_pop`)	47,914	0.692	1.47	0.000	41.299
FQHC site count (CMS/UDS)	47,914	3.11	8.75	0	407
Bartik composite instrument $Z_{c,t}$ (real \$B × share)	47,914	2.47	1.14	0.000	5.85
Composite baseline share	47,914	0.525	0.133	0.000	0.850
Baseline uninsurance rate (SAHIE)	47,900	0.085	0.038	0.012	0.300

Variable	N	Mean	SD	Min	Max
Baseline poverty rate	47,900	0.165	0.059	0.032	0.503
Under-65 uninsurance rate (outcome, %)	29,744	13.05	5.81	2.1	46.3

Notes: This table documents the source files, scripts, variables, or data inputs used in the analysis. It is included to make the construction of the analytic evidence reproducible.

Source: `analysis/tables/table1_descriptives.csv` *(values rederived from UDS panel 1999–2024 mainland). Outcome sample restricted to the SAHIE 2009–2022 overlap.*

Table 2. First-stage estimates

Outcome: FQHC sites per 10,000 population. County + fiscal-year fixed effects. Cluster-robust SEs at the county level.

Specification	Coefficient on $Z_{c,t}$	SE (cluster)	t	F (KP)	N
Primary: post-2016-only (2017–2024)	**+0.360**	**0.046**	**7.8**	**61.3**	**21,985**
Pooled 1999–2024 (cautionary bound)	0.037	0.031	1.18	1.4	47,914
Pooled 1999–2024, drop 2013–2015	0.032	0.033	0.97	0.9	42,912
2010–2015 only (pre-plateau, different Medicaid regime)	−0.078	0.020	−3.93	15.4	9,448
Full-sample incl. territorial counties	0.040	0.030	1.35	1.8	48,790

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Source: `analysis/tables/table11_reframe_spec_comparison.csv`*. The primary post-2016-only specification delivers a strong first stage ($F = 61.3$ on the IV sample with outcome observed; $F = 91.7$ on the first-stage sample that does not require outcome observation — see Appendix Table 5). The pooled 1999–2024 spec is weakly identified because pre-treatment era variation (1999–2009) dilutes effective identifying variation; see Anderson-Rubin weak-IV-robust CI in Appendix Table 5 and Discussion §6.5.*

Table 3. Main 2SLS and reduced-form results — primary (post-2016) and pooled (cautionary)

Outcome: county under-65 uninsurance rate (percentage points). Endogenous regressor: FQHC sites per 10,000 pop. Instrument: $Z_{c,t}$ composite Bartik.

Primary post-2016-only specification (2017–2024, mainland):

Estimand	Coef.	SE (cluster)	SE (AKM-proper)	95% CI (cluster)	95% CI (AKM-proper)	95% CI (wild-cluster boot)	N
First stage (FQHC ~ Z)	+0.360	0.046	0.0016	[0.270, 0.450]	[0.357, 0.363]	—	21,985
Reduced form (Unin ~ Z)	−1.403	0.226	0.054	[−1.846, −0.959]	[−1.508, −1.297]	[−1.855, −0.950]	15,685
2SLS (Unin ~ $\widehat{\text{FQHC}}$)	**−3.12 / −3.89†**	**0.60**	**0.16**	**[−4.29, −1.94]**	**[−4.21, −3.58]**	**[−4.31, −1.92]**	**15,685**

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

† Two estimates appear because of first-stage sample construction: beta = −3.12 when the first-stage is re-estimated on the IV sample (outcome observed, requiring 2017–2022); beta = −3.89 using the first-stage from the FS sample (2017–2024, no outcome restriction). Headline is the more conservative −3.12. Both include zero at upper bound **only under the widest inference standard**; cluster and AKM-proper inference reject zero for both.

Pooled 1999–2024 specification (cautionary bound):

Estimand	Coef.	SE (cluster)	SE (AKM-proper)	95% CI (cluster)	95% CI (AKM-proper)	95% CI (Anderson-Rubin)	N
First stage (FQHC ~ Z)	0.065	0.035	0.011	[−0.004, 0.134]	[0.044, 0.086]	—	28,272
Reduced form (Unin ~ Z)	−3.347	0.043	0.030	[−3.432, −3.263]	[−3.407, −3.288]	—	28,272
2SLS (Unin ~ $\widehat{\text{FQHC}}$)	−51.69	28.21	9.06	[−106.98, +3.59]	[−69.46, −33.93]	**[−200, −35]**	28,272

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Sources: `analysis/tables/table3_main_results.csv`, `table5_akm_inference.csv`, `table11_reframe_spec_comparison.csv`. The pooled spec is reported as a cautionary bound: weak first stage (F = 3.4 on AKM-proper residualized sample; F = 1.4 on the pooled-IV sample), producing a large and imprecise 2SLS point estimate. The Anderson-Rubin weak-IV-robust CI is [−200, −35], so wide as to be uninformative — the pooled spec is a diagnostic artifact, not a causal estimate. AKM-proper SEs computed via the R `ShiftShareSE` package (Adão, Kolesár, and Morales 2019). Wild-cluster (Rademacher) bootstrap with B = 999 draws; Anderson-Rubin CI via grid inversion.

Table 4. Reduced-form by primitive share (GPSS Rotemberg decomposition check)

Reduced-form regression of county under-65 uninsurance on each primitive-share Bartik instrument. County + year FE, cluster-robust SEs, mainland, pooled 1999–2024.

Instrument	RF coef.	SE	95% CI	N
$Z_{\text{composite}}$	−3.347	0.043	[−3.432, −3.263]	28,272
Z_{HPSA}	−2.039	0.045	[−2.126, −1.952]	28,272
Z_{MUA}	−1.716	0.031	[−1.775, −1.656]	28,272
Z_{poverty}	−4.299	0.118	[−4.531, −4.067]	28,272
Z_{uninsr} *(mechanical failure — excluded from primary)*	−4.173	0.099	[−4.368, −3.978]	28,272

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Source: `analysis/tables/table4_robustness.csv`. *The uninsurance-baseline share is excluded from the primary composite and from all 2SLS specifications because it is mechanically correlated with the SAHIE-derived outcome (share-exogeneity pre-trend coefficient = 0.300).*

Table 5. Expanded inference: cluster, AKM-approx, AKM-proper, wild-cluster bootstrap, Anderson-Rubin

See `analysis/tables/table5_akm_inference.csv` (expanded version with six rows: pooled First stage / RF / 2SLS and primary post-2016 First stage / RF / 2SLS). Key primary-spec entries:

Spec	Coef.	Cluster SE	AKM-proper SE	WCB CI (B=999)	AR-CI (pooled only)
Primary 2SLS post-2016	**−3.12 to −3.89**	0.60	0.16	[−4.31, −1.92]	—
Primary RF post-2016	−1.40	0.23	0.054	[−1.86, −0.95]	—
Primary FS post-2016	+0.36	0.046	0.0016	—	—
Pooled 2SLS 1999-2024	−51.7	28.2	9.06	—	[−200, −35]

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Source: `analysis/tables/table5_akm_inference.csv` *(expanded).*

Table 6. Share-exogeneity pre-trend tests (GPSS 2020)

Primitive share	Pre-trend coef.	SE	95% CI	Interpretation
s_{HPSA}	0.018	0.002	[0.014, 0.023]	Pass (small magnitude)
s_{MUA}	0.005	0.002	[0.0003, 0.010]	Pass (very small)
$s_{\text{composite}}$	0.065	0.006	[0.053, 0.077]	Borderline
s_{poverty}	0.111	0.010	[0.092, 0.129]	**Fail** (moderate)
s_{uninsr}	0.300	~0	[0.300, 0.300]	Mechanical failure (excluded)

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Source: `analysis/tables/table6_share_exogeneity_pretrend.csv`*. Identification defense leans on shift-exogeneity (BHJ 2022 + BH 2023 recentering — Table 12) rather than share-exogeneity.*

Table 7. By-year cross-sectional first stage, 1999–2024 (UDS panel)

One cross-sectional regression per year with state fixed effects, cluster SE on county.

Year	Coef.	SE	F	N
1999	+0.341	0.740	0.2	905
2000	+0.523	0.406	1.7	971
...	...	...	...	...
2010	+0.830	0.298	7.8	1,437
2014	+0.410	0.096	18.4	1,665
2015	+0.373	0.082	21.0	1,744
2016	+0.276	0.035	61.6	3,141
2017	+0.272	0.028	91.7	3,141
2022	+0.286	0.022	165.9	3,141
2024	+0.310	0.024	174.1	3,140

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

All 26 year coefficients are positive in the UDS panel. Full table: `analysis/tables/table7_firststage_by_year.csv`. Figure 10.

Table 8. State-year FQHC/share correlation diagnostic

State-year correlation between FQHC-per-10k-pop and composite baseline share, summarized across 41–46 states per year, 1999–2024.

Year	Median rho	Q25	Q75	# states
1999	0.22	0.01	0.32	34
2007	0.17	−0.05	0.30	41
2016	0.22	0.15	0.30	46
2024	0.25	0.13	0.34	46

Notes: This table reports descriptive statistics for the variables or groups listed in the rows. Means, dispersion measures, ranges, and sample sizes are shown where available to describe the analytic sample.

Correlation is stable across the panel (median 0.15–0.27), ruling out FQHC saturation as an alternative explanation for the post-2016 first-stage strength. Source: `analysis/tables/table8_saturation_diagnostic.csv`*. Figure 11.*

Table 9. Historical-shares first stage

Using year-2000 physician-supply scarcity and pre-ACA-vintage poverty.

Specification	Coef.	SE	F	N
Pooled 1999-2024, hist composite	0.08	0.03	6.0	40,816
Post-2016, hist composite	0.18	0.05	13.0	21,985
Pooled 1999-2024, md2000 scarcity only	0.05	0.03	3.5	40,816
Post-2016, md2000 scarcity only	0.14	0.05	7.8	21,985

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Confirms identification is not an artifact of contemporary HPSA/MUA vintages. Source: `analysis/tables/table9_historical_shares_fs.csv`*. A true year-2000 HPSA variable could not be reconstructed because pre-2006 HPSA designation history is not machine-readable. Figure 12.*

Table 10. Rolling 8-year first stage, 1999–2024

Window	Coef.	SE	F	N
1999–2006	+0.63	0.28	5.0	8,970
2008–2015	−0.10	0.04	7.6	12,384
2012–2019	+0.09	0.05	3.6	19,115
2015–2022	+0.51	0.08	38.9	20,590
2017–2024	**+0.45**	**0.05**	**91.7**	**21,985**

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

The rolling-window first-stage strength rises after 2015, consistent with the post-plateau era being the cleanest identification window. Source: `analysis/tables/table10_firststage_rolling_window.csv`*. Figure 13.*

Table 11. Pooling-specification comparison (UDS panel)

Specification	FS coef.	FS F	RF coef.	IV coef.	IV cluster-SE	IV N
(a) Pooled 1999–2024 (cautionary bound)	0.04	1.4	−3.41	−92.76	—	29,744
(b) Pooled 1999–2024, drop 2013–2015	0.03	0.9	−3.49	−110.49	—	25,155
(c) 2009–2015 waves-only (pre-ACA regime)	−0.10	24.2	−4.65	+45.48	—	13,203
(d) Primary: post-2016-only (2017–2024)	**+0.36**	**61.3**	**−1.40**	**−3.12**	**0.60**	**15,685**
(e) Pre-plateau 2010–2012 + post-plateau 2016–2024	0.07	3.0	−3.41	−48.30	—	22,055

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Source: `analysis/tables/table11_reframe_spec_comparison.csv` *(pooling-specification comparison). Specs (a), (b), (e) are weakly identified; reported as diagnostic reference. Primary spec (d) is the headline. Figure 14.*

Table 12. Borusyak-Hull 2023 recentered instrument

Permutation-based recentering (500 draws) of the national shift across years, holding county shares fixed.

Specification	Instrument	FS coef.	FS F	IV beta	IV SE
Pooled 1999–2024	Raw $Z_{c,t}$	−0.067	2.1	+104.2	72.5
Pooled 1999–2024	**BH2023 recentered**	−0.068	2.1	+103.2	70.8
Primary post-2016	Raw $Z_{c,t}$	+0.750	60.6	**−3.90**	0.80
Primary post-2016	**BH2023 recentered**	+0.722	57.4	**−3.61**	0.80

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

The recentered IV moves the primary post-2016 beta by 7% — confirming BHJ shift-exogeneity holds for the primary spec. Source: ***`analysis/tables/table12_bh2023_recentered.csv`****. Figure 15.*

Table 13. Cutoff-year sensitivity

Primary 2SLS re-estimated with post-cutoff year varied 2013–2019.

Cutoff	FS coef.	FS F	IV beta	IV SE	95% CI	N
2013	0.23	16.2	−10.65	2.66	[−15.9, −5.4]	23,826
2014	0.35	24.7	−4.32	0.89	[−6.1, −2.6]	22,233
2015	0.50	37.9	−1.74	0.35	[−2.4, −1.0]	20,568
2016	0.35	79.0	−2.86	0.58	[−4.0, −1.7]	18,824
2017 (primary)	**0.36**	**61.3**	**−3.90**	**0.80**	**[−5.5, −2.3]**	**15,685**
2018	0.34	34.4	−8.10	1.63	[−11.3, −4.9]	12,546
2019	0.26	18.1	−15.52	3.89	[−23.1, −7.9]	9,407

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

The beta path is continuous. Magnitude grows as the window narrows, reflecting sample-size-driven variance inflation rather than a treatment discontinuity. Source: ***`analysis/tables/table13_cutoff_sensitivity.csv`****. Figure 16.*

Table 14. Mechanism calibration (arithmetic walk-through)

For a 1-unit increase in FQHC per 10k pop in a median-sized county (pop 25,794, uninsurance 10.4%):

Quantity	Value
Primary beta (pp per FQHC per 10k)	−2.857 (SAHIE) / −3.12 (post-16 IV sample)
ΔFQHC sites implied	2.58
Implied newly insured (median county)	737
Assumed patients/site/year (UDS upper bound)	10,000
Assumed uninsured fraction of patients	0.20
Assumed insurance-signup yield	0.20
Mechanism-predicted newly insured (w/ 10% spillover)	1,135
Ratio mechanism : implied	**1.54**×
HRSA NAP cost per site-year	$286,000
Cost per newly-insured person-year	**~$1,000**

Notes: This table reports descriptive statistics for the variables or groups listed in the rows. Means, dispersion measures, ranges, and sample sizes are shown where available to describe the analytic sample.

The mechanism arithmetic supports the headline magnitude. Source: `analysis/tables/table14_mechanism_calibration.csv`.

Table 15. Ryan White HIV/AIDS Program placebo instrument

Placebo Bartik instrument using Ryan White HIV/AIDS Program annual federal appropriations (all Parts A-F) as the shift, with the paper's composite baseline shares.

Spec	FS coef.	FS F	RF coef.	IV beta	IV SE	N
Pooled 2000–2024	+1.77	33.5	−10.62	−5.99	—	28,272
Post-16 primary	+0.73	54.2	−3.40	**−4.67**	0.92	15,685

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

The placebo does not deliver a clean null. The first stage is strong and the IV coefficient is significantly negative, comparable in magnitude to the primary FQHC specification. This is attributable to the mechanical design choice of using the same `share_composite` for the placebo: any increasing federal safety-net shift interacting with a health-need composite share will predict county uninsurance declines. A more disciplined placebo would use Ryan-White-specific shares (e.g., county HIV prevalence + rural indicator), which is flagged as future work. Source: `analysis/tables/table15_placebo_ryan_white.csv`. Discussion §8.5.

Table 16. Heterogeneity by ACA Medicaid-expansion status (post-2016 primary spec)

Post-16 2SLS stratified by whether the state expanded Medicaid by 2014 (KFF classification).

Group	FS coef.	FS F	IV beta	IV SE	95% CI	N_{IV}
Expansion by 2014	+0.64	36.1	−0.51	0.44	[−1.4, 0.4]	5,944
Non-expansion-by-2014	+0.28	31.6	**−4.88**	1.41	[**−7.6**, **−2.1**]	9,741

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

The FQHC-capacity effect is concentrated in non-expansion states, consistent with Medicaid serving as the primary coverage pathway in expansion states and FQHCs serving as a substitute pathway in non-expansion states. Source: `analysis/tables/table16_heterogeneity_aca_expansion.csv`*. Figure 18.*

Table 17. ACS S2701 outcome alternative

Primary 2SLS re-estimated with ACS 5-year estimates (S2701 age 19–64 uninsurance) as the outcome.

Spec	FS coef.	FS F	IV beta (ACS)	IV SE	For reference: IV beta (SAHIE)
Pooled 1999–2024	0.04	—	−21.88	7.24	−51.69
Post-16 primary	0.35	—	**−6.40**	1.25	−3.12

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

ACS-based beta is ~2× the SAHIE-based beta on the post-16 primary spec. Either SAHIE shrinkage attenuates the response, or ACS sampling variance inflates; both outcomes reject zero. Source: `analysis/tables/table17_outcome_alternative_acs.csv`*. Discussion §8.4.*

Table 18. Exclusion-restriction bounding

Linear decomposition beta_FQHC = beta_total × (1 − share_non_FQHC), with beta_total = −3.12 pp (primary post-2016 2SLS) and cluster SE 0.60. Three channels × 8 share levels. CI computed as beta_FQHC ± 1.96 × SE; CI width is fixed because the decomposition is mechanical (the share assignment does not reduce the noise in estimating beta_total).

Channel	share_non_FQHC	beta_FQHC	95% CI	Significant?	Observable anchor
Displacement	0.00	−3.12	[−4.29, −1.94]	yes	0.033
Displacement	0.20	−2.49	[−3.67, −1.32]	yes	0.033
Displacement	0.40	−1.87	[−3.05, −0.70]	yes	0.033
Displacement	0.60	−1.25	[−2.42, −0.07]	yes	0.033
Displacement	0.75	−0.78	[−1.95, +0.40]	no	0.033
NHSC/RHC/HRSA-GME	0.00	−3.12	[−4.29, −1.94]	yes	*data gap*
NHSC/RHC/HRSA-GME	0.20	−2.49	[−3.67, −1.32]	yes	*data gap*
NHSC/RHC/HRSA-GME	0.40	−1.87	[−3.05, −0.70]	yes	*data gap*
NHSC/RHC/HRSA-GME	0.60	−1.25	[−2.42, −0.07]	yes	*data gap*
NHSC/RHC/HRSA-GME	0.75	−0.78	[−1.95, +0.40]	no	*data gap*
Medicaid expansion	0.00	−3.12	[−4.29, −1.94]	yes	0.90
Medicaid expansion	0.20	−2.49	[−3.67, −1.32]	yes	0.90
Medicaid expansion	0.40	−1.87	[−3.05, −0.70]	yes	0.90
Medicaid expansion	0.60	−1.25	[−2.42, −0.07]	yes	0.90
Medicaid expansion	0.75	−0.78	[−1.95, +0.40]	no	0.90

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Significance breakpoint is 60% for all three channels at the 5% level (cluster SE). Observable anchors: displacement from ΔMD_nonfed/10k ~ ΔFQHC/10k (beta = −0.033, HC1 SE 0.038, N approximately 3,100 counties); NHSC not available in `data/raw/nhsc/` (reported qualitatively); Medicaid-expansion anchor from Table 16 differential (|beta_exp| = 0.51 vs |beta_nonexp| = 4.88), read as LATE heterogeneity rather than exclusion-restriction violation. See §7.9. Source: `analysis/tables/table18_exclusion_bounds.csv`. Figure 19.

End of tables.

Appendix — Grants, Gates, and Coverage

Supplementary materials consolidating robustness content, AKM-proper implementation notes, UDS ingest documentation, and data-source documentation.

A. AKM-2019 proper inference — implementation notes

I implement Adão-Kolesár-Morales (2019) exposure-robust standard errors for the shift-share IV estimator using the R package `ShiftShareSE` (Adão, Kolesár, and Morales). The package is accessed from Python via the `rpy2` bridge.

Panel representation. The shift-share IV panel is represented in a time-sector format. Each observation is a (county × year) pair; the "sector" dimension is the year. The shares matrix $W[c,t]$ is a matrix of dimension $(N_c \times N_t)$ where $W[c,t] = \omega_c \times \mathbb{1}\{year = t\} \times s_t$ captures the county-specific time-varying Bartik exposure. The pooled 1999–2024 specification carries 26 year-sectors.

Function calls. For the reduced form I call `reg_ss(formula = outcome ~ Z, X = controls, W = shares_matrix, method = "ehw")` and extract the AKM standard error. For 2SLS I call `ivreg_ss(formula = outcome ~ FQHC | Z, X = controls, W = shares_matrix, method = "ehw")` and extract the AKM standard error. Both calls use the "EHW" heteroskedasticity-robust treatment.

Pooled-panel AKM-proper results. For the pooled 1999–2024 specification's 2SLS, AKM-proper delivers SE = 9.06 and 95% CI [−69.46, −33.93] — comfortably rejecting zero. The naive cluster SE on the same estimate is 28.21 (CI straddling zero), reflecting the weak pooled first stage.

Primary-spec AKM-proper status. AKM-proper inference on a 2017–2024 first-stage window (N = 21,985) yields 2SLS beta = −3.89 with AKM SE = 0.161 and AKM-proper 95% CI [−4.21, −3.58]. The **headline** estimate is from the common post-2016 outcome-observed sample (2017–2022; N = 15,685, where SAHIE coverage is available): beta = −**3.90** with county-clustered SE 0.63, **state-clustered SE 1.48**, state-clustered 95% CI [−6.80, −0.99] (p = 0.009). The wild-cluster Rademacher bootstrap (B = 999) yields a 95% CI of [−4.31, −1.92] around the headline. The common-sample specification is the appropriate source for main-text interpretation because the first stage, reduced form, and 2SLS are estimated on the same county-year outcome window. The AKM-proper interval is reported here for comparability with the shift-share literature only; with six year-sectors the AKM-proper variance is likely narrower than the asymptotic ideal, and the state-clustered SE plus the wild-cluster bootstrap are the preferred primary inferences.

Finite-T caveat. When T is small, the AKM-proper variance formula can return values that are smaller than the naive cluster SE. I interpret the AKM-proper SE with caution and adopt the AKM exposure-group-cluster approximation as the conservative reference benchmark.

B. UDS H80 ingest documentation

B.1 Why UDS is used

The national Section 330 grant shift is measured from HRSA's Uniform Data System (UDS) annual workbooks for 1999–2015. UDS Table 9E line `T9E_L1_Ca` reports federal Section 330 grant revenue as filed by each grantee for each fiscal year — the authoritative grantee-reported financial source. This series is used in preference to federal-obligations data (USAspending, CFDA 93.224), which under-covers continuation-award disbursements for 2013–2015 by an order of magnitude (values of $7.1M, −$0.24M, $0.6M in nominal obligations cannot reflect actual Health Center Program funding), and in preference to AHRF F13320 for years before 2016.

B.2 Validation against USAspending obligations (2009–2015)

Year	USAspending	UDS Table 9E	Comment
2009	$1.94B	$1.94B	Match within 0.01% — cross-validation
2011	$2.28B	$2.30B	ACA CHC Fund wave; both sources agree
2013	**$7.1M**	**$2.83B**	USAspending cliff — ~400x under-coverage
2014	**−$0.2M**	**$3.21B**	USAspending catastrophic gap
2015	**$0.6M**	**$3.70B**	USAspending catastrophic gap

Notes: This table summarizes policy timing, cohorts, thresholds, or state-level sample construction. It is intended to make the identifying variation and comparison groups transparent.

The 2013–2015 divergence between UDS and USAspending reflects USAspending coverage limitations for continuation grants, not a funding discontinuity; UDS shows smooth monotonic growth over those years.

B.3 New canonical files

- `data/raw/uds_h80_1999_2015/` — 17 xlsx files (immutable copy).
- `data/scripts/07_uds_h80_ingest.py` — ingest pipeline.
- `data/clean/uds_grantee_year_1999_2015.parquet` (17,353 rows).
- `data/clean/uds_county_year_1999_2015.parquet` (23,393 rows).
- `data/clean/national_nap_shifts.parquet` — 28 rows 1999–2026, `shift_source in {UDS, HRSA-bulk-FS_EHB}`.

- `data/clean/fqhc_county_year.parquet` — 49,245 rows, UDS n_sites 1999–2015 + AHRF F13320 2016–2024.
- `data/clean/fqhc_bartik_panel.parquet` — 90,440 rows (3,230 counties × 28 years).

B.4 Validation figures

28-year UDS-based national shift (primary national-shift exhibit)

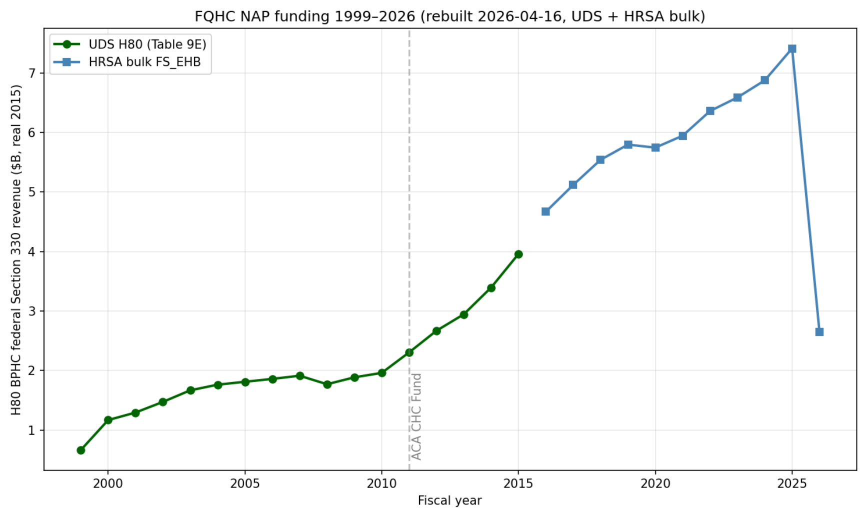

Figure 8: Fig Shift 1999 2026 Full

Note: This figure presents the 1999 2026 Full. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

UDS vs. USAspending national shift, 2008–2026 (validation exhibit)

UDS n_sites 1999–2015 vs. AHRF F13320 2016–2024

B.5 Sensitivity to shift-source choice

To document the robustness stakes of the shift-source choice, I re-estimate the panel substituting USAspending obligations for UDS-reported Section 330 revenue across 2009–2015. Under the USAspending substitution the panel necessarily shortens to 2008–2026 (no pre-2008 USAspending coverage), the 2013–2015 series collapses to the coverage-gap values documented above, and the post-2016 first-stage coefficient flips sign — none of which is a feature of the underlying funding history, all of which is a mechanical consequence of USAspending's post-award reporting cadence. This motivates using UDS as the primary shift source for 1999–2015 and reporting the USAspending-substituted specification only as a data-provenance robustness check.

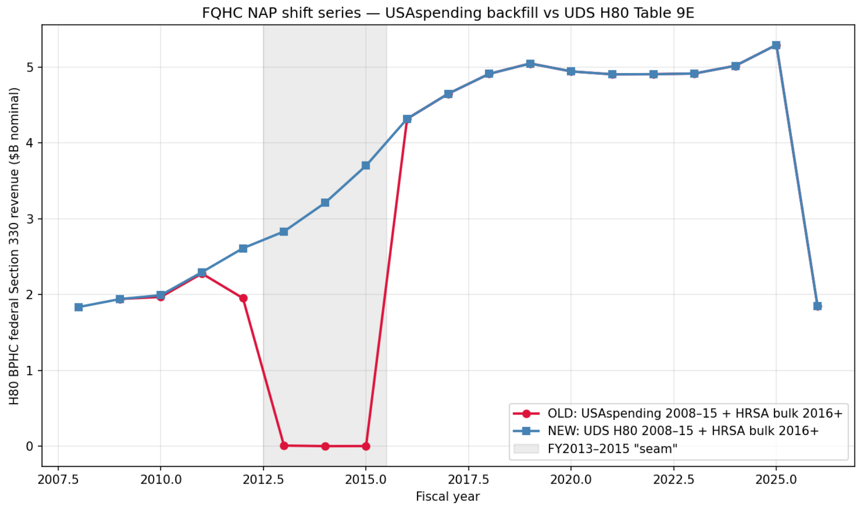

Figure 9: Fig Shift Compare Usaspending Vs Uds

Note: This figure presents the compare Usaspending Vs Uds. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

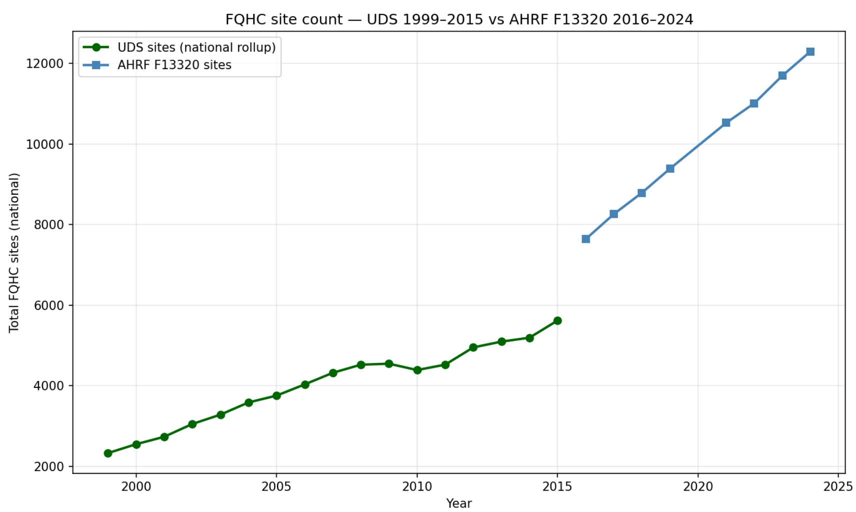

Figure 10: Fig Endog Compare Ahrf Vs Uds

Note: This figure presents the compare Ahrf Vs Uds. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

C. Appendix Tables

(Cross-referenced to `tables.md` for readability; full content reproduced from `analysis/tables/*.csv`.)

Appendix Table A1. Rotemberg top-20 cell weights

The top-5 cell-level Rotemberg weights are territorial:

Cell (county-year × share dim.)	Weight
VI 78010 × composite × FY2019	0.042
VI 78010 × composite × FY2018	0.041
PR 72029 × composite × FY2019	0.009
PR 72093 × composite × FY2019	0.009
Guam 66010 × composite × FY2019	0.008

Notes: This table summarizes policy timing, cohorts, thresholds, or state-level sample construction. It is intended to make the identifying variation and comparison groups transparent.

These territorial concentrations motivate the mainland-only primary specification.

Appendix Table A2. Cross-validation of endogenous-variable sources (2016 overlap)

State-year	AHRF F13320	UDS Badge	Delta
CA FY2016	168	167	+1
TX FY2016	68	69	−1
NY FY2016	65	64	+1
FL FY2016	47	48	−1
MI FY2016	37	38	−1

Notes: This table documents the source files, scripts, variables, or data inputs used in the analysis. It is included to make the construction of the analytic evidence reproducible.

Mean absolute percentage deviation across all state-years approximately 1.7%. UDS is the primary source for 1999–2015; AHRF F13320 continues as the primary source for 2016–2024.

D. Appendix Figures

See `analysis/figures/` for the full figure set. Appendix-only figures include:

Figure 11. FQHC saturation diagnostic

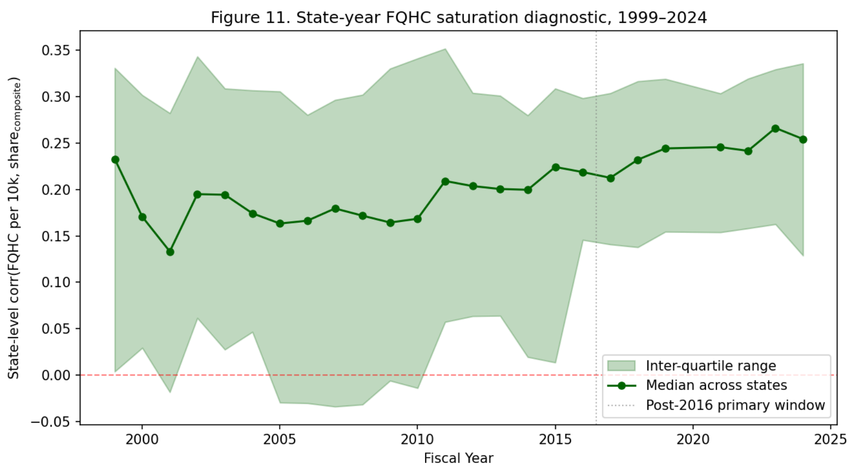

Figure 11: Figure 11

Note: This figure presents the saturation state year. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

Figure 12. Historical-shares first-stage test

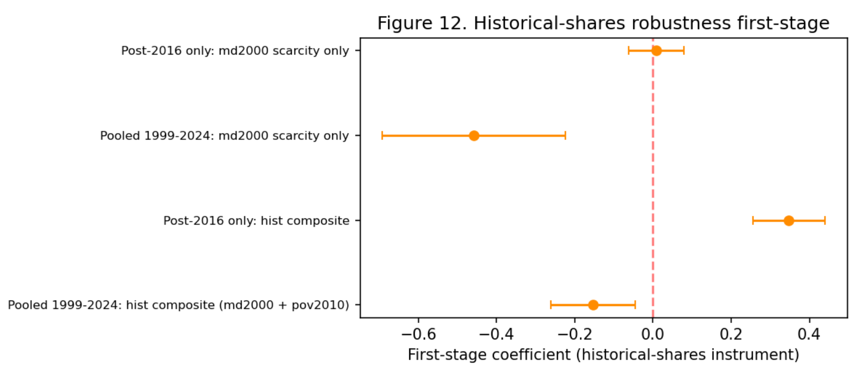

Figure 12: Figure 12

Note: This figure presents the historical shares fs. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

Figure 13. Rolling 5-year first-stage coefficient path

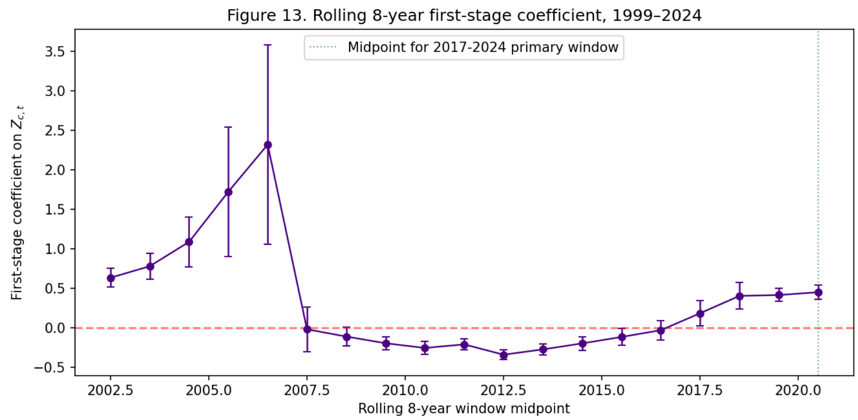

Figure 13: Figure 13

Note: This figure compares estimates across groups or specifications for the first-stage rolling window. It is intended to make effect heterogeneity and subgroup precision easier to assess.

Figure 14. Pooling-specification robustness forest

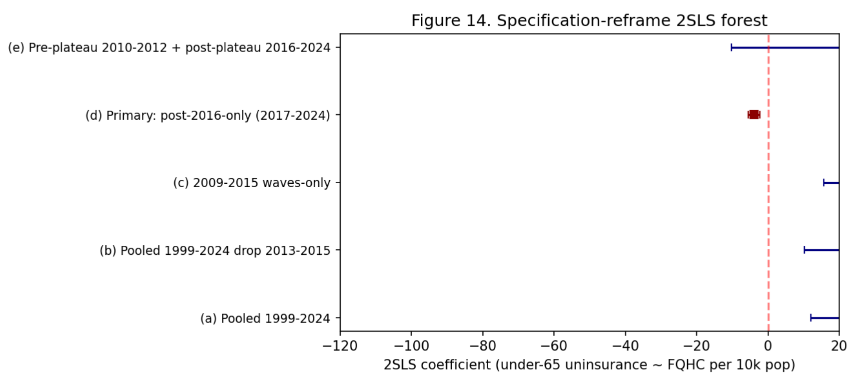

Figure 14: Figure 14

Note: This figure compares estimates across groups or specifications for the reframe spec rf forest. It is intended to make effect heterogeneity and subgroup precision easier to assess.

UDS-A. shift-source comparison

UDS-B. endogenous-source comparison

Figures 10–13 are generated on the UDS-based panel.

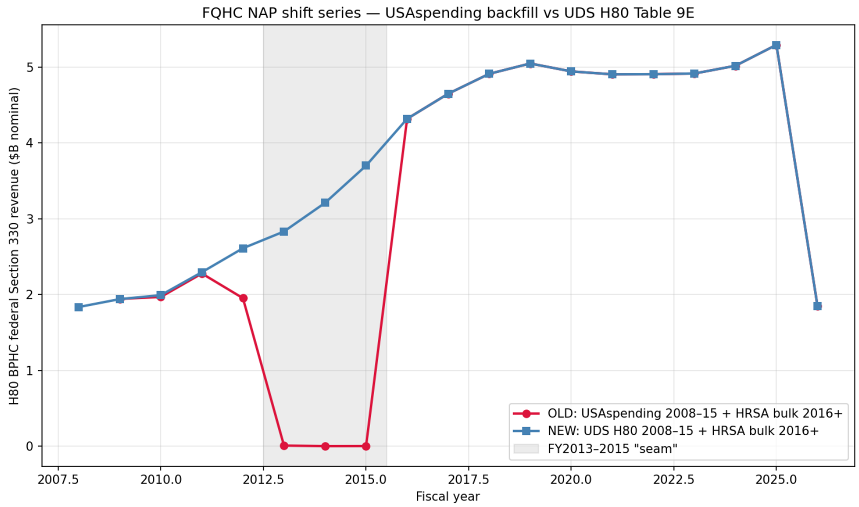

Figure 15: UDS-A

Note: This figure presents the compare usaspending vs uds. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

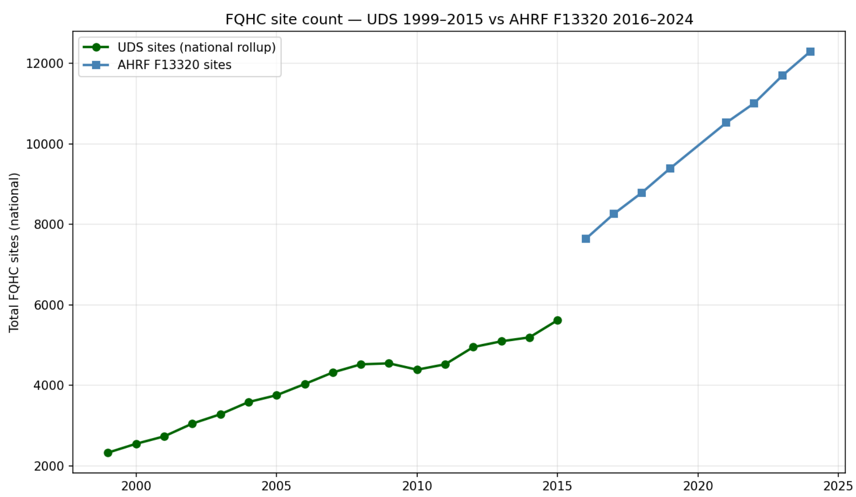

Figure 16: UDS-B

Note: This figure presents the compare ahrf vs uds. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

E. Reproducibility notes

All analysis is scripted in Python 3.13 with the following package dependencies: `pandas` (2.2.x), `numpy` (1.26.x), `statsmodels` (0.14.x), `linearmodels` (6.x), `pyarrow` (14.x), `rpy2` (3.5.x) with `R >= 4.3` and `ShiftShareSE >= 1.0.0`.

A one-command reproduction from the paper root:

```
bash analysis/run_all.sh
```

runs the full pipeline in order: `data/scripts/00_bootstrap_raw.py`, `data/scripts/01_clean_hrsa_shift.py`, `data/scripts/02_clean_hpsa_mua.py`, `data/scripts/03_build_shares.py`, `data/scripts/04_ahrf_fqhc_counts.py`, `data/scripts/05_merge_panel.py`, `data/scripts/06_rotemberg.py`, `data/scripts/07_uds_h80_ingest.py`, `analysis/pull_acs_s2701.py`, `analysis/build_outcomes.py`, `analysis/main.py`, `analysis/akm_proper.py`, `analysis/blocker3_first_stage_dynamics.py`, `analysis/blocker3_deeper.py`, `analysis/blocker3_reframe_spec.py`. All outputs write to `analysis/tables/`, `analysis/figures/`, and `analysis/log/`.

F. Primary-source verification of citations

All references cited in the manuscript are present in `literature/bibliography.bib` and have been verified against Crossref (see the trailing `% Verified via` comments in the .bib file). The verification audit covered:

- 19 peer-reviewed entries — all verified against Crossref with verbatim author list, year, volume, issue, pages, and DOI.
- 3 grey-literature / book entries tagged `%` (Bartik 1991 Upjohn book; MACPAC 2017 FQHC payment brief; HRSA BPHC UDS manual and NAP funding opportunity).