The Supply-Side Effects of Healthcare Privatization: Evidence from Puerto Rico*

Dayanara Diaz Vargas†

June 7, 2026

Abstract

We estimate the causal effects of healthcare delivery privatization on the local healthcare workforce, hospital capacity, and patient-health outcomes by exploiting Puerto Rico's 1993 reform, which transferred the island's public delivery system to private managed-care organizations, with implementation staggered across 78 counties between 1994 and 2000. Using the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), we find that the reform contracted the local healthcare sector without raising unemployment: healthcare employment per resident fell by 13.0 percent and hospital beds per resident fell by 7.8 percent, while effect on the unemployment rate was null, indicating that displaced workers exited the labor force—possibly due to out-migration to the U.S., where credentials transfer. Patient-health outcomes did not deteriorate on average; infant mortality fell by 17 to 27 percent of its pre-reform mean in counties where the reform's bite was greatest (urban, more populous, more healthcare-intensive). A back-of-envelope welfare calculation suggests that fiscal cost savings to the public health system—roughly $967 million in 2025 dollars over the four-year window—were sufficient to make the reform's net welfare positive even before valuing the patient-health gains, which add up to $1.75 billion more at the EPA Value of a Statistical Life.

JEL classification: H51, I11, I18, J21, J45.

Keywords: healthcare privatization; managed care; Medicaid; healthcare workforce.

*I would like to thank all faculty members that helped me along the way with their feedback: Dr. Shantannu Khana, Dr. Mindy Marks, Dr. James Dana.

†Northeastern University

1 Introduction

A government that pays for healthcare must decide whether to deliver it through its own facilities and salaried staff, or private operators under contract. Therefore, the question is whether contracting delivery to private operators lowers costs without lowering quality, and its answer depends on the incentives those operators face (Hart et al., 1997) and on how they are paid—in particular, on whether their prices are set by regulators or by the market (Gaynor et al., 2015). The evidence suggests that where regulators set prices above marginal cost, competition among providers raises quality; the English National Health Service is the leading example of this (Gaynor et al., 2015; Cooper et al., 2011; Gaynor et al., 2013). Where prices are instead set in the market, and private operators choose both what to charge and how much healthcare to deliver, theory makes no clear prediction, and the evidence is mixed (according to Gaynor et al. (2015). This is concerning, since public healthcare typically represents a large part of fiscal expenditures. For example, in the United States, Medicaid spent $919 billion in fiscal year 2024; about half of that—$459 billion—flowed through 291 private managed-care organizations to roughly 66 million low-income enrollees across 42 states and the District of Columbia (Kaiser Family Foundation, 2025, 2026). The Department of Veterans Affairs operates a $112.6 billion medical-care budget and now spends about twenty percent of it on care purchased from private community providers, up from twelve percent in 2014.[1] In England, the Department of Health and Social Care received £204.7 billion for 2024/25 and used £24 billion of it—about twelve percent—to buy care from providers outside the NHS.[2]

In each of these programs, the move toward private delivery changes several things at once: facility ownership, provider compensation structure, and patient access. The existing

[1]VA community-care obligations grew from $7.9 billion in 2014 to $18.5 billion in 2021 (in 2021 dollars), driven by the 2014 Veterans Choice Program and the 2018 MISSION Act, which expanded eligibility for veterans to receive care outside the VA system (Congressional Research Service, 2025; Congressional Budget Office, 2021; Farmer, 2023).

[2]This contracting arrangement traces to the 1991 NHS internal market and was broadened under the 2012 Health and Social Care Act (King's Fund, 2026; National Audit Office, 2026).

causal evidence, however, has studied these changes only one at a time. No study has estimated what happens when an entire publicly operated delivery system is handed to private operators paid by capitation, the public alternative is closed, and patient access is restricted—all together. We study one such case. Puerto Rico's 1993 healthcare reform made exactly this system-wide change, rolling it out across 78 counties between 1994 and 2000. Therefore, we exploit this setting to ask how privatizing a publicly operated delivery system affects the local healthcare labor market, while also examining its effects on hospital capacity and on patient health.

Four bodies of work have studied pieces of this change, each holding the rest of the system fixed. Research on the English NHS (Propper, 2018; Cooper et al., 2011; Gaynor et al., 2013; Bloom et al., 2015) measures how competition affected quality within a delivery system that stayed public. Chan et al. (2023) measure what happens when individual veterans are routed to private hospitals, with the VA itself left in place. The hospital-privatization literature (Duggan et al., 2023; Eliason et al., 2020; Andreyeva et al., 2024) studies the conversion of a single hospital, dialysis facility, or provider inside an otherwise mixed system. And the Medicaid managed-care literature (Duggan, 2004; Aizer et al., 2007; Wallace, 2023; Geruso et al., 2023; Marton et al., 2014) studies capitation and network restrictions inside a public insurance program whose ownership does not change. In every case, one element of the bundle moves while the others are held fixed. None of these studies estimates what happens when ownership, payment, and the public option change together.

Puerto Rico's 1993 reform differs from each of these settings. Unlike the NHS reforms, it transferred the delivery system itself rather than rearranging how a still-public system bought care. Unlike the VA expansion, it closed the public option rather than adding a private alternative alongside it. Unlike a standalone Medicaid managed-care reform, it turned public providers into private ones rather than changing the payment terms of providers who were already private. And unlike the hospital-privatization studies, it changed the entire system at once rather than a single facility within it.

Four features of the setting strengthen the analysis. First, the staggered rollout treated counties in different years, giving within-island variation in treatment timing across all 78 counties. Second, each region was served by a single contracted managed-care organization, and every insurer covered the same standardized benefit package, so residents could neither choose their insurer nor decide when their region was privatized; this leaves little room for individuals to sort into treatment on unobserved characteristics. Third, because Puerto Rico is an island, its healthcare market is self-contained and separated from the mainland, so the not-yet-treated counties that serve as controls are free of the cross-state patient and provider flows that complicate designs built on the contiguous United States. Fourth, federal Medicaid financing in Puerto Rico is capped rather than open-ended. Under Section 1108 of the Social Security Act, the Commonwealth receives a fixed annual federal allotment and a fixed matching rate; the fifty states, by contrast, draw on an open-ended match that rises with their own spending and is indexed to per-capita income (Medicaid and CHIP Payment and Access Commission, 2020; Kaiser Family Foundation, 2024), which may distort their spending decisions and incentivize federal-draw-down maximization. Because spending above the allotment draws no further federal funds, Puerto Rico's government bears the marginal cost of care. Thus, the fiscal savings we attribute to privatization therefore measure resource savings, not the matching-rate-driven responses that confound Medicaid analyses set in U.S. states.

Furthermore, this setting is informative about privatization at a scale the facility-level evidence cannot reach. When a single hospital is privatized inside a mostly public system, the patients it no longer serves can fall back on the public facilities that remain, and displaced workers can be absorbed elsewhere in that system; the measured effect is therefore a partial-equilibrium object that these outlets muffle (Duggan et al., 2023). But, Puerto Rico closed the public option across the whole system. Therefore, our estimates reflect the full equilibrium response to privatization—including the healthcare workforce's exit to the U.S. mainland—rather than the attenuated effect of converting one facility while the rest of the system absorbs the displaced. A government privatizing an entire delivery system faces exactly this closed environment, and as privatization scales from individual facilities toward whole systems, the system-wide effect, not the facility-level effect, becomes the one that matters. The treatment is, however, a bundle. The reform changed ownership, payment, and access at the same time, and the access change itself had two sides, widening coverage while introducing gatekeeping through primary-care networks. Because all of this moved together, the estimates capture the combined effect of the bundle, not the separate contribution of any one channel.

We find that the reform shrank the local healthcare sector in per-resident terms, but patient health did not fall in proportion. Healthcare employment per thousand residents dropped by 13.0 percent relative to its pre-reform mean, and Medicare-certified hospital beds per thousand fell by 7.8 percent. The unemployment rate, however, did not change at conventionally statistically significant levels. Instead, the adjustment appeared in labor-force participation, which fell by 3.5 percent. This leads us to hypothesize that healthcare workers used their medical credentials to leave for the U.S. mainland rather than remaining unemployed on the island. Additionally, patient health did not deteriorate on average. Infant mortality shows no significant change in the full county panel, but rather it falls by 17 to 27 percent of its pre-reform mean in the subsamples where the reform's effect was largest—urban, more populous, and more healthcare-intensive counties. A back-of-the-envelope calculation suggests that the fiscal savings to the public health system—about $967 million in 2025 dollars over the four years after the reform—were by themselves enough to make its net welfare positive, even before counting any value from the health improvements.

Two of these findings speak to current debates over privatization. The first concerns a prediction from work on insurance and input markets (Finkelstein, 2007; Dillender, 2022; Hackmann et al., 2025): that expanding coverage stimulates the supply of healthcare. Here that prediction reverses. When coverage expansion arrives bundled with privatization, capitation, and network restrictions, the sector contracts rather than grows. This contraction is specific to a system-wide transfer of ownership; it cannot be inferred from reforms that merely widen a private outside option, such as the MISSION Act. The second concerns the Hart et al. (1997) account of privatization, on which private operators cut quality in order to cut costs. The data here do not fit that account: the privatized system used fewer healthcare resources per resident without a matching loss in measured patient health, and quality in fact improved where the reform's effect was largest. Whether privatizing a whole system harms patients therefore appears to depend on two things—the quality of the pre-reform public system across different areas, and whether prices are set by regulators or by the market, the same distinction Gaynor et al. (2015) use to organize the wider evidence.

2 Setting

2.1 The Pre-Reform Public Delivery System

The Commonwealth of Puerto Rico operated an integrated public healthcare delivery system organized along the regionalized principles of Arbona and Ramirez de Arellano (1978): a four-tier hierarchy, universal access at the point of service, and direct public employment of clinical staff. The first tier comprised 78 *Centers of Diagnostic and Treatment* (CDTs), one per county, delivering primary medical, prenatal care, dental, and pharmaceutical services through salaried government clinicians. The second through fourth tiers comprised five Area Hospitals, six Regional Hospitals, and the Puerto Rico Medical Center—a quaternary referral complex in San Juan housing the Commonwealth's academic medical training programs. By 1990, the public delivery system operated roughly 40 percent of the Commonwealth's hospital beds, accounted for 38 percent of medical visits and 40 percent of inpatient days, and served approximately 1.8 million residents—about 52 percent of the population (Instituto de Administración y Política de Salud, 1995).

2.2 The 1993 Reform

The Puerto Rico Health Insurance Administration Act (Act 72 of September 7, 1993) created an institution in charge of being the sole purchaser of medical-indigent health coverage and restructured care delivery along three margins simultaneously (Commonwealth of Puerto Rico, 1993; World Health Organization and Pan American Health Organization, 1998).[3] *Ownership* transitioned from public to private: the 78 CDTs and 11 Area and Regional Hospitals were sold or transferred under long-term lease to private operators contracted through the new managed-care organizations (Instituto de Administración y Política de Salud, 1995); the Commonwealth retained a residual subset of supra-tertiary facilities within the Medical Center complex, but the operational delivery network had transferred to the private sector by the close of the rollout. *Payment* transitioned from fixed salaries paid by the Department of Health to capitated managed-care payments paid by ASES-contracted insurers at terms negotiated bilaterally between insurer and provider (Commonwealth of Puerto Rico, 1993; World Health Organization and Pan American Health Organization, 1998). *Access* transitioned from open-access public-facility entitlement to insurer-restricted networks with primary-care gatekeeping; cost-sharing was introduced where none had previously existed, with copayments of \$1–\$4 per office visit and \$0.50–\$3 per prescription scaled by enrollee income (Instituto de Administración y Política de Salud, 1995; World Health Organization and Pan American Health Organization, 1998). ASES partitioned the Commonwealth into eight administrative health regions and contracted each region to a single managed-care organization through competitive bidding, creating regional monopolies.

Implementation was staggered sequentially between February 1994 and July 2000 (World Health Organization and Pan American Health Organization, 1998; Brugueras Fabre, 2002).

[3] Eligibility was extended to residents with household incomes below 200 percent of the Puerto Rico poverty threshold, in addition to the existing Medicaid-eligible population. The covered-benefits package included inpatient and outpatient hospital care, physician and ambulatory services, pharmacy, dental, mental-health and substance-abuse treatment, and preventive services, issued without exclusions for pre-existing conditions or waiting periods (Commonwealth of Puerto Rico, 1993; Instituto de Administración y Política de Salud, 1995; World Health Organization and Pan American Health Organization, 1998; Centers for Medicare and Medicaid Services, 2000, 2001, 2003).

We aggregate the eight regional rollout dates to five annual treatment cohorts indexed by calendar year of transition to managed-care operation: Cohort 1 (1994; $n = 10$ counties), Cohort 2 (1995; $n = 12$), Cohort 3 (1996; $n = 30$), Cohort 4 (1998; $n = 25$), and Cohort 5 (2000; $n = 1$, San Juan).[4] Figure 1 maps the cohort assignment across the 78 counties.

3 Data

The empirical analysis combines six data sources, aligned to the 78-county cross-section and to overlapping annual time windows within 1990–2001. Labor-market outcomes are constructed from the U.S. Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW) at the county-quarter establishment level, providing average monthly employment, total quarterly wages, and active establishment counts by NAICS code; the healthcare sector is the union of NAICS 621 (Ambulatory Health Care Services), NAICS 622 (Hospitals), and NAICS 623 (Nursing and Residential Care Facilities), with NAICS 621 reported separately because it maps most directly onto the privatized municipal clinic network. Hospital-capacity outcomes are constructed from the Centers for Medicare and Medicaid Services Provider of Services (POS) public-use files, aggregated from the provider-year to the county-year level. Patient-health outcomes (the crude birth rate, infant mortality rate, low- birth-weight share, and government and private hospital birth shares) are constructed from the Puerto Rico Department of Health Vital Statistics annual records; total county population and the crude net migration rate are constructed from the same source.

[4]The peripheral-to-urban rollout pattern concentrates Cohorts 1–2 in rural and outer-island regions while Cohort 5 covers the capital metropolitan area; the empirical implications are taken up in Section 5. Two contemporaneous federal-policy contingencies bear on identification. The Section 1108(g) territorial Medicaid cap and the fixed Federal Medical Assistance Percentage applicable to Puerto Rico produce a federal contribution per enrollee substantially below what Puerto Rico would receive under the income-conditioned FMAP formula applied to the fifty states (Medicaid and CHIP Payment and Access Commission, 2020; Kaiser Family Foundation, 2024); this constraint pre-dates the reform and is absorbed by the county fixed effects in our specification. The federal Section 936 corporate-tax credit for Puerto Rico manufacturing was phased out beginning in 1996 over the same period as the reform's rollout; we test for contamination of the main estimates in Section 7. Financing detail (Commonwealth General Fund share, federal Medicaid, county contributions, SCHIP) and the post-2002 *Mi Salud/Vital* successor programs (Pan American Health Organization, 2007; Centers for Medicare and Medicaid Services, 2026) are reported in Appendix 10.

Worker-margin outcomes (the labor-force- participation rate, the employment-to-population ratio, and the unemployment rate) are constructed from the Puerto Rico Department of Labor and Human Resources monthly county-level labor-force survey. Two robustness data sources support the patient-health and access-channel analyses respectively: the National Center for Health Statistics National Vital Statistics System (NVSS) natality public-use files provide individual-level birth records used to corroborate the county-aggregate Vital Statistics outcomes, and the Centers for Disease Control and Prevention Behavioral Risk Factor Surveillance System (BRFSS) provides individual-level adult microdata on self-reported insurance coverage and access used to test the access-channel first stage of the reform. Appendix 10 reports the full data construction, the QCEW annualization procedure, the CMS ownership-classification panel, the CNMR demographic- balancing residual, the BRFSS cohort-restriction window, and the NVSS conception-year exposure indicator.

The Vital Statistics population series is the denominator for all per-thousand-resident outcomes. The series updates annually in step with the changing population during the rollout window; it is also the official population series used by the Commonwealth's statistical and fiscal apparatus during the study window. Per-capita scaling places labor-market, supply- side, and patient-health outcomes on the welfare-relevant denominator and matches the unit conventions of the literature (Hackmann et al., 2025; Wallace, 2023). One scope limitation is that the CMS POS file identifies only Medicare-certified providers, which excludes the 78 CDTs delivering the primary-care services that did not require Medicare certification under the pre- reform regulatory structure. The CMS panel therefore captures the secondary, tertiary, and supra-tertiary hospital infrastructure but not the ambulatory-clinic layer; the QCEW NAICS 621 outcome is the empirical counterpart for the ambulatory-layer privatization shock. The balanced-cohort QCEW healthcare panel covers 59 counties; the balanced CMS panel covers 40 counties; the Vital Statistics and PR Department of Labor panels cover all 78 counties.

See Table 1 for summary statistics and Table 2 to understand the pre-treatment means of relevant outcomes. Specifically, this table reports cohort-specific pre-reform (1990–1993) means of population, public-system intensity, healthcare and placebo-industry employment, labor- market activity, demographic flows, and patient-health covariates, with one-way ANOVA p-values for the joint null of cohort-mean equality. Cohorts differ systematically according to their urban-intensive (as described in Section 2.2).

4 Theoretical Framework

The three margins the reform moved jointly—ownership, payment, and access give us different hypotheses for per-capita healthcare employment.

The three margins the reform moved at once -ownership, payment, and access- do not all push the same way, and because they changed together the data recover their net, not any one alone.

Consider ownership first. Hart et al. (1997) show that when quality cannot be fully written into a contract, shifting a service from public to private hands sharpens the operator's incentive to cut costs and dulls its incentive to sustain quality: the operator keeps what it saves by spending less, yet cannot be penalized for decreasing quality where the contract does not specify. Because labor is the largest variable input in healthcare delivery, this cost-cutting falls first on staffing, so the ownership channel predicts $\Delta L^{\text{ownership}} \leq 0$. The same mechanism predicts that per-patient noncontractible quality q (i.e. the dimensions of care a payer cannot observe and write into a contract) decreases under private operation, $\Delta q \leq 0$, most visibly where the reform bit hardest. This carries one exception: if the public system was itself underinvesting in quality relative to the private sector, as competition has been found to do within continuing public systems (Propper, 2018), privatization can raise quality instead. Patient health therefore may increase or decrease, depending on the the case.

The payment margin pushes the same way. Under prospective per-member capitation, the managed-care operator becomes the residual claimant on the cost of network labor (since it keeps whatever it does not spend) which gives it a direct incentive to economize on the inputs it can adjust (Duggan, 2004; Aizer et al., 2007; Marton et al., 2014); absent an offsetting rise in the capitated rate, $\Delta L^{\text{capitation}} \leq 0$.

The access margin, by contrast, is ambiguous, because the reform changed access in two directions at once. By extending effective insurance coverage to the medical-indigent population, it raised the demand for insured care, which supply-stimulus logic (Finkelstein, 2007; Dillender, 2022; Hackmann et al., 2025) predicts will expand the local healthcare labor market. But by restricting provider networks and channeling patients through primary-care gatekeepers, it worked to suppress utilization, which pulls in the opposite direction. Therefore, the effect depends on which mechanism dominates.

Summing the three channels gives the reform's net effect,

$$\Delta L^{\text{reform}} = \underbrace{\Delta L^{\text{ownership}}} \leq 0 + \underbrace{\Delta L^{\text{capitation}}} \leq 0 + \underbrace{\Delta L^{\text{access}}}_{\gtreqless 0},$$

. Our estimate recovers the channels' combined effect, not the contribution of each, because they interact: the cost-cutting incentive under capitation, for instance, is likely stronger when the operator also owns the facility. In this regard, we rely on cross-sectional variation to understand which channel dominates, because each channel is strongest in a different type of county. The access channel's expansionary force should be strongest where the most residents stood to gain coverage: poorer, higher-unemployment counties with many uninsured residents before the reform. The ownership and capitation channels should contract employment most where the privatized public sector was largest: denser, more healthcare-intensive counties. Because these subsamples are small, we cannot test the contrasts formally, and read them only as suggestive patterns in the point estimates.[5]

[5]See Appendix 10 reports the formal derivations of the ownership and capitation predictions, the quality corollary, and the mobility margin.

5 Empirical Strategy

The empirical analysis exploits the staggered rollout of the reform across five cohorts of counties ($\mathcal{G} = \{1994, 1995, 1996, 1998, 2000\}$) to identify the average effect of treatment on the treated (ATT). Let $Y_{i,t}$ denote an outcome measured for county i in year t, let $G_i \in \mathcal{G}$ denote the treatment cohort of county i, and let $D_{i,t} = \mathbf{1}\{t \geq G_i\}$ be the post-treatment indicator. Because every county in Puerto Rico is treated by 2000, the analysis has no never-treated comparison group; identification relies instead on *not-yet-treated* counties—units in cohorts $g' > g$ that have not yet transitioned to managed-care operation by year t. The estimand of interest is the group-time average treatment effect $\mathrm{ATT}(g,t) = \mathbb{E}[\, Y_{i,t}(g) - Y_{i,t}(\infty) \mid G_i = g\,]$ for $t \geq g$ (Callaway and Sant'Anna, 2021), where $Y_{i,t}(g)$ and $Y_{i,t}(\infty)$ are the potential outcomes under treatment-in-cohort-g and no-treatment respectively. Identification rests on three assumptions: (i) conditional pre-treatment parallel trends for the not-yet-treated comparison group, $\mathbb{E}[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) \mid G_i = g] = \mathbb{E}[Y_{i,t}(\infty) - Y_{i,t-1}(\infty) \mid G_i > t]$; (ii) no anticipation, $Y_{i,t}(g) = Y_{i,t}(\infty)$ for $t < g$; and (iii) staggered, irreversible adoption. Under (i)–(iii), $\mathrm{ATT}(g,t)$ is identified by the doubly robust difference-in-differences estimator

$$\widehat{\mathrm{ATT}}(g,t) \;=\; \mathbb{E}\big[Y_{i,t} - Y_{i,g-1} \,\big|\, G_i = g\big] \;-\; \mathbb{E}\big[Y_{i,t} - Y_{i,g-1} \,\big|\, G_i > t\big]\,. \tag{1}$$

The regression tables report the simple weighted average of $\widehat{\mathrm{ATT}}(g,t)$ across cohorts and post-treatment periods, with weights proportional to the number of post-treatment county-year observations; the event-study figures report the dynamic aggregation $\widehat{\theta}_\ell$ for $\ell \in \{-4, \ldots, +4\}$ with the reference period $\ell = -1$. Inference uses the multiplier bootstrap of Callaway and Sant'Anna (2021) clustered at the county level with 1,000 replications.[6]

The main specification is the Callaway and Sant'Anna (2021) estimator with not-yet-treated comparison units. Because the estimand and the parallel-trends assumption are estimator-specific, Section 7 re-estimates every main outcome with three alternative estimators—

[6]See Appendix 10 for the dynamic-aggregation equation and the technical details of the doubly robust three-step procedure.

the canonical two-way fixed-effects event-study, the extended two-way fixed-effects estimator of Wooldridge (2025), and the imputation estimator of Borusyak et al. (2024)—as an estimator-class robustness check.

The validity of the design turns on pre-treatment trends, not on pre-treatment levels. Because earlier-treated counties are systematically more rural than later-treated ones, the rollout produces systematic cohort imbalance on observable pre-treatment levels, but under Callaway and Sant'Anna (2021) such imbalance does not threaten identification so long as the not-yet-treated comparison group satisfies assumption (i). We assess that assumption directly through the event-study pre-trend coefficients reported alongside each main result in Section 6.

Three further threats bear on the interpretation of the estimates rather than on the design itself, and Section 7 takes up each in turn: the IRS Section 936 corporate-tax phase-out beginning in 1996, which is contemporaneous with the rollout; the disproportionate leverage of San Juan as the sole 2000-cohort county; and within-island spillovers in labor and patient flows across treated and not-yet-treated counties.

6 Results

6.1 First Stage and Supply Contraction

The reform's first stage operated cleanly on the ownership margin. Table 3 reports the estimates on the share of births delivered at government and at private hospitals—the cleanest available proxy for which ownership class operated the inpatient delivery system in each county-year. The government-hospital share fell by 14.79 percentage points ($p < 0.01$), or 23.6 percent of its 62.55 percent pre-reform mean, and the private-hospital share rose by an offsetting 15.19 percentage points ($p < 0.01$). The two estimates are statistically indistinguishable in absolute magnitude and opposite in sign: the reorganization of births is essentially a one-for-one substitution from public to private facilities.[7] The first-stage event study in Figure 1 confirms the timing: pre-treatment coefficients on both series are statistically indistinguishable from zero through $\ell = -1$, and the substitution opens within the first year of treatment and widens through $\ell = +3$.

Coincident with the ownership transition, the local healthcare sector contracted relative to population. Table 4 reports the QCEW labor-market estimates for the healthcare aggregate (NAICS 621+622+623) and, separately, for ambulatory primary care (NAICS 621)—the subsector mapping most directly onto the privatized county clinic network. Healthcare employment per thousand residents fell by 0.511 employees, or 13.0 percent of its pre-treatment mean ($p < 0.01$); the contraction was bigger in ambulatory primary care, at 17.3 percent ($p < 0.01$). Establishments per thousand fell by 6.6 percent in the aggregate ($p < 0.01$) and by 4.6 percent in NAICS 621 ($p < 0.10$), so the contraction operated on both the intensive (employment) and extensive (firm-count) margins. The wage-per-employee estimate is statistically indistinguishable from zero in both panels: real compensation per healthcare worker did not respond to the reform.[8] Figure 2 reports the corresponding six-panel event study for employment, establishments, and the wage bill per employee across both industry definitions: pre-treatment coefficients are statistically flat through $\ell = -1$, and the post-reform contraction in employment and establishments opens within the first year while the wage-per-employee series remains flat throughout.

The contraction in the healthcare labor market is paralleled by a contraction in physical hospital capacity, on the narrower CMS panel of forty Medicare-certified counties. These estimates are reported alongside the patient-health outcomes in Table 5, Columns 3–4: Medicare-certified hospital beds per thousand fell by 7.8 percent ($p < 0.01$) and Medicare-

[7]The two shares are not exact complements because a small tail of births, typically under one percent, is recorded as non-hospital or with unspecified ownership.

[8]The contraction is in per-capita terms. In raw $\log(1 + y)$ levels, healthcare employment in treated counties grew by approximately 19 percent over the window, so the per-capita decline reflects population growth that outpaced the local healthcare workforce; the per-capita normalization is the welfare-relevant scaling and the convention of the closest comparison literatures (Wallace, 2023; Hackmann et al., 2025). Appendix Table A.7 reports the per-capita and log-level estimates side by side.

certified services per thousand by 4.9 percent ($p < 0.05$). The bed and service event studies appear in Figure 3, Panels (e)–(f), where the post-treatment contraction opens by $\ell = +2$.

6.2 Worker-Side Adjustment Margins

The labor-market contraction did not reflect onto wages or onto the unemployment rate. Table 6 reports three margins along which displaced healthcare workers and broader local-labor-market spillovers might adjust. The labor-force-participation rate fell by 1.11 percentage points, or 3.5 percent of its 31.895 percent pre-treatment mean ($p < 0.01$); the BLS-standard unemployment rate, by contrast, is statistically indistinguishable from zero against a pre-reform mean of 17.770 percent.[9] The crude net migration rate moves from a pre-treatment mean of +1.06 residents per thousand per year to an estimated post-treatment flow of −14.65 per thousand ($p < 0.05$), and the crude birth rate falls by 3.1 percent (−0.564 per thousand, $p < 0.05$). Total county population, however, is not lower in treated units: the point estimate on the population level is +5.1 percent ($p < 0.01$).[10] Read together, the two demographic estimates describe slower demographic growth rather than absolute population decline: treated counties continue to grow in levels, but with a population mix increasingly weighted toward non-working-age residents, consistent with the parallel declines in labor-force participation and the crude birth rate. The crude-net-migration and crude-birth-rate event studies appear in Figure 3, Panels (c)–(d); the migration signal opens within two years of treatment and the crude-birth-rate trajectory bends downward from $\ell = 0$.

The crude net migration rate is the one primary outcome whose sign is sensitive to estimator choice, and we therefore read its magnitude as suggestive; the four-estimator

[9]The unemployment rate is not reported in Table 6; Online Appendix Table OA.1 reports it alongside the labor-force-participation and employment-to-population rates, and Appendix 10 reports the labor-force-participation event study.

[10]Online Appendix Table OA.20 reports the population estimate in levels and logarithms; Online Appendix Table OA.21 reports an explicit demographic-balancing identity check. The population-level and crude-net-migration estimates do not arithmetically reconcile, because the migration rate is constructed as a residual of three separately measured flows (Equation 3) and is more sensitive to compounding measurement error than the level outcome, and because the C&S aggregation weights differ across panels of different length (the migration panel loses each county's first observation, for which Pop_{t-1} is unavailable).

comparison is reported in Section 7. The labor-force-participation result, which uses the same 78-county panel and is the behavioral counterpart of the migration story, is robust across all four estimators and is the firmer basis for the labor-supply-margin reading.

6.3 Patient Health, Heterogeneity, and Access

The short-horizon patient-health response is null on average. Table 5, Columns 1–2, reports the county-aggregate Vital Statistics outcomes whose biological response horizons fall within the four-year post-reform window: the infant mortality rate fell by 1.205 deaths per thousand live births (9.5 percent of pre-mean, statistically insignificant) and the low-birth-weight share rose by 0.06 percentage points (statistically insignificant). Both event studies, in Figure 3, Panels (a)–(b), are flat in both the pre- and post-treatment windows.[11] The patient-health null coexists with the supply-side contraction documented above: the privatized system delivered fewer per-capita inputs without a detectable loss of the patient outcomes whose response horizons fall within the panel.

Table 7 reports heterogeneity by pre-reform public-hospital exposure, partitioning the 78 counties at the median 1990–1993 government-hospital birth share. The first-stage substitution is larger in the below-median exposure group (-15.49 percentage points, $p < 0.01$) than in the above-median group ($+2.41$ percentage points, statistically insignificant), implying that counties with high pre-reform government shares had less room to fall in absolute terms; the labor-market estimate attenuates similarly in the above-median group. Patient health shows no differential response on this exposure cut. The infant-mortality response is bigger, however, along four other pre-reform classifiers: in the urban (above-median population density), larger (above-median population), higher-healthcare-intensity, and workplace-center (above-median commuting inflow) subsamples, the infant mortality rate falls by between

[11] Individual-level NVSS specifications on birth weight, prenatal-visit count, trimester of prenatal-care initiation, and Apgar scores are similarly null (Online Appendix Table OA.22). Long-horizon adult cause-specific mortality outcomes have response horizons that exceed the post-reform panel and are reported, but not interpreted as identified treatment effects, in Online Appendix Table OA.25.

17 and 27 percent of its pre-treatment mean, statistically significant across all four cuts.[12] These are the same cells in which the first-stage ownership transition is most pronounced: patient health improves where the reform had the most bite. The cross-sectional pattern is therefore difficult to reconcile with a pure quality-deterioration reading of the Hart et al. (1997) privatization mechanism (Section 4). Figures 4 and 5 report the supporting heterogeneity event studies—by baseline poverty rate and by pre-reform government-hospital exposure, respectively—each splitting the panel at the median and tracing employment, infant mortality, and net migration separately for the two halves.

The access margin moved in the expansionary direction. Across age, education, and employment subgroups, self-reported insurance coverage in the BRFSS rises by between 3 and 10 percentage points, estimated under the Wooldridge (2025) extended two-way fixed-effects estimator because the post-1996 BRFSS window admits only the later cohorts.[13] The coverage expansion is the empirical counterpart of the access channel combined with the labor-market contraction documented above, it indicates that the contractionary supply-side equilibrium obtained notwithstanding expanded coverage. This is the configuration in which the supply-stimulus prediction of Finkelstein (2007), Dillender (2022), and Hackmann et al. (2025) reverses sign once coverage expansion is bundled with delivery-system privatization, capitation, and network restriction.

[12]Appendix Table A.6 and Online Appendix Tables OA.9, OA.10, and OA.11 report the four cuts; the above-median infant-mortality ATTs are −2.03 (density, $p < 0.05$), −2.56 (population, $p < 0.01$), −2.63 (healthcare employment, $p < 0.01$), and −3.12 (commuting inflow, $p < 0.01$), in deaths per thousand live births. Heterogeneity by baseline poverty (Online Appendix Table OA.23) and pre-reform unemployment (Online Appendix Table OA.24) is underpowered relative to these four cuts.

[13]Online Appendix Tables OA.14–OA.19 report the six BRFSS subgroup estimates; the "has health plan" point estimates range from +2.8 percentage points (ages 65 and older, where near-universal Medicare leaves little room to expand) to +10.5 percentage points (ages 45–64), and are significant in five of the six subgroups. Appendix 10 discusses the subgroup pattern and the single-cohort-exposure limitation that requires the ETWFE estimator on this panel.

6.4 Welfare Bounds

6.5 Welfare Bounds

6.6 Welfare Bounds

6.7 Welfare Bounds

The reform cut per-capita inputs without a proportional loss of measured patient quality, so its welfare effect depends on whether the resources it freed up outweigh any change in patient health. We compare the two in dollar terms with a back-of-the-envelope, partial-equilibrium calculation that establishes the sign and rough magnitude of the net effect rather than a structural estimate; we report a range, not a point.

The net gain is the supply cost the smaller delivery system avoids plus the dollar value of the infant deaths it averts:

$$W = \underbrace{\frac{P\,T}{1000}\Big(|\hat{\beta}_{\text{bed}}|\,c_b + |\hat{\beta}_{\text{emp}}|\,\bar{w}\Big)}_{\text{avoided supply cost}} + \underbrace{v\,\frac{|\hat{\beta}_{\text{IMR}}|}{1000}\,B_H\,T}_{\text{value of averted infant deaths}}. \tag{2}$$

The first term applies the two estimated per-thousand reductions to the island's population P over the four post-reform years ($T = 4$): $\hat{\beta}_{\text{bed}} = -0.265$ Medicare-certified beds per thousand (Table 5) and $\hat{\beta}_{\text{emp}} = -0.511$ healthcare workers per thousand (Table 4). A lost worker is priced at the 1990–1993 mean annual healthcare wage $\bar{w}$, and a lost bed-year at $c_b = \$62{,}756$, the wage bill scaled by the U.S. HCRIS total-cost-to-wage-bill ratio of 2.1. The second term values the infant deaths averted by the heterogeneity-cell mortality reduction $\hat{\beta}_{\text{IMR}}$ (Section 6.3) at a value of a statistical life (VSL), v. Because that reduction appears only in the high-exposure half of counties, we apply it to the high-exposure births B_H alone, about half the island total. Following **?**, we set v to the U.S. EPA value of \$10 million (2024 USD), or \$3.19 million in 1982–1984 dollars.

We compute W under two choices, giving four cells. The supply estimate is either the headline C&S estimate or the more conservative §936-conditional estimate of Section 7; the health term is either zero, under the full-panel infant-mortality null, or the high-exposure reduction averaged over the four heterogeneity cuts of Section 6.3. The cells run from $262.5 to $845.0 million over the window in 1982–1984 dollars, or $844 million to $2.72 billion in 2025 dollars. Appendix Table A.4 lists each cell in both price bases and separates the two terms.[14]

Three caveats affect the result. First, we value averted infant deaths at an adult value of a statistical life; one scaled to infants' much longer life expectancy would be higher, so the cells with a health benefit are conservative. Second, the cost saving is a social gain only if the resources the reform released were absorbed elsewhere in the economy, which the decline in labor-force participation (Section 6.2) is consistent with but does not establish. Third, the four-year window omits adult cause-specific mortality, which responds over a longer horizon.

Even so, the net effect is positive in all four cells, and it does not require the health benefit: with the health term set to zero, welfare is still $300.8 million under the headline supply estimate and $262.5 million under the §936-conditional one.[15] The resource saving drives the result; improved patient health, where we value it, only adds to it.

7 Robustness

This section reports six robustness exercises. Each varies one input to the main specification—the estimator, the sample, the industry, the conditioning covariate, or the outcome normalization—while holding the others at the main-text defaults. Figure **??** reports a forest-plot summary of the six specifications applied to healthcare-employment-per-1,000; the corresponding Appendix tables report the full specifications for every primary outcome.

[14] Online Appendix Table **??** replaces the one-half weight with empirical pre-reform high-exposure birth shares and adds a wage-reabsorption discount; the cells move by under three percent.

[15] The bed and employment terms are not fully separable, since the bed price already embeds labor cost; the bed reduction alone yields a positive saving, so the result does not depend on combining them.

Estimator class. The main estimates are triangulated against three alternatives: the canonical two-way fixed-effects (TWFE) event-study, the extended two-way fixed-effects estimator of Wooldridge (2025), and the imputation-based efficient estimator of Borusyak et al. (2024). The TWFE specification is included as a contamination diagnostic: in staggered designs with treatment-effect heterogeneity, its event-time coefficients are weighted averages whose weights need not be positive, producing sign reversals where the doubly robust C&S, BJS, and Wooldridge estimators do not—a pattern each of those estimators' developers diagnoses. ETWFE, C&S, and BJS agree in sign and magnitude on the first-stage ownership transition, the healthcare-sector employment contraction, the labor-force-participation decline, and the hospital-bed contraction; TWFE produces the expected sign reversals on the outcomes where heterogeneity is most pronounced—healthcare employment, infant mortality, and the low-birth-weight share.[16] The crude net migration rate is the one outcome on which the four estimators do not agree: TWFE, ETWFE, and BJS return statistically insignificant point estimates with mixed signs (-5.69, -1.53, and $+4.51$ per thousand respectively), while the C&S point estimate of -14.65 per thousand ($p < 0.05$) is the basis for the migration discussion in Section 6.2. The migration result should therefore be read as suggestive: its sign is the C&S signature, but its magnitude is not robust across estimators. The labor-force-participation result, which uses the same 78-county panel and is the behavioral counterpart of the migration story, is robust across all four estimators and is the firmer basis for the labor-supply-margin reading.

Excluding San Juan. San Juan is the sole 2000-cohort county and accounts for approximately 11 percent of the Commonwealth's population. Excluding it attenuates the supply-side and labor-market estimates substantially while leaving the first-stage ownership-

[16] Appendix Table A.1 reports the four-estimator point estimates and bootstrap standard errors for the six primary outcomes, and the full TWFE specification (county and year fixed effects with event-time indicators, reference period omitted, standard errors clustered at the county level). Across estimators, the first-stage government-hospital-share estimate is -12.9 percent (TWFE), -26.4 percent (ETWFE), -23.6 percent (C&S), and -18.0 percent (BJS) of pre-mean, all $p < 0.01$; HC employment per thousand is -25.7 (ETWFE), -13.0 (C&S), and -16.7 (BJS), all $p < 0.01$; LFP per population is -3.6, -3.5, and -2.2 percent, all $p < 0.01$; hospital beds per thousand is -13.1, -7.8, and -5.6 percent, all $p < 0.01$.

transition estimate intact. The government-hospital share estimate moves from −14.79 percentage points in the main result to −8.22 percentage points in the excluded-San Juan sample, both $p < 0.01$.[17] The healthcare-employment point estimate attenuates from −13.0 to −5.3 percent of pre-mean and is no longer statistically distinguishable from zero; the hospital-bed estimate attenuates from −7.8 to +3.0 percent, also no longer significant. Two features of San Juan account for the attenuation. The metropolitan capital holds the largest concentration of healthcare workforce and Medicare-certified capacity in the Commonwealth; removing it removes most of the absolute mass through which a per-capita contraction can be identified. It is also the cell where the out-migration option is cheapest in the legal-portability sense (Section 4). The first-stage ownership transition depends on neither feature and survives the exclusion at conventional significance.

Within-island spillovers. The C&S framework requires the stable-unit-treatment-value assumption: outcomes for not-yet-treated comparison counties should not be affected by the treatment status of already-treated neighbors. Two sample restrictions tighten the SUTVA condition at the cost of statistical power: a non-border subset that drops counties sharing a boundary with a different-cohort treated unit, and a low-commuting-inflow subset that drops counties in the upper half of the workplace in-commuting distribution.[18] The first-stage ownership transition is robust under both restrictions (−12.14 percentage points on the non-border subset and −7.18 percentage points on the low-commuting-inflow subset, both $p < 0.05$), as are the patient-health outcomes. The labor-market estimates do not survive the non-border restriction at the resulting sample size of 27 QCEW counties: healthcare employment per thousand, the wage bill per employee, and establishments per thousand are each statistically indistinguishable from zero. The hospital-bed and services-per-thousand

[17] Online Appendix Table OA.3 reports the first-stage and patient-health side-by-side specifications; Appendix Table A.2 reports the QCEW labor-market and CMS hospital-capacity specifications; Online Appendix Table OA.4 reports the ownership-stratified CMS robustness.

[18] Appendix Table A.5 and Online Appendix Tables OA.6, OA.5, OA.7, and OA.8 report the corresponding specifications. The non-border restriction retains 27 of 78 counties on the QCEW panel, 19 of 40 on the CMS panel, and 39 of 78 on the Vital Statistics panel; the low-commuting-inflow restriction retains 64 of 78 on the Vital Statistics panel and 34 of 40 on the CMS panel.

estimates retain the direction of the main result but lose precision. The cleanest reading of the spillover-restricted estimates is power loss rather than evidence against the main estimates.

Placebo industries. The placebo specifications test whether contemporaneous shocks orthogonal to the reform—most prominently the Section 936 corporate-tax phase-out beginning in 1996—produce healthcare-sector-shaped effects on industries with no exposure to the privatized delivery system. Three industries are reported: NAICS 11 (Agriculture, Forestry, and Fishing), NAICS 48–49 (Transportation and Warehousing), and NAICS 71 (Arts, Entertainment, and Recreation). NAICS 31–33 manufacturing is excluded *a priori*, because the Section 936 phase-out is a contemporaneous federal-policy shock to Puerto Rican manufacturing during the analysis window, so the manufacturing sector is endogenous to the same policy environment as the healthcare sector and is not orthogonal to the treatment.[19] The placebo evidence is mixed. Agriculture is null on absolute-headcount margins (raw-level employment, wage bill, and establishments) and on per-capita employment and establishments; the per-capita wage-bill decline does not survive raw-level normalization, indicating population redistribution rather than an Agriculture-specific shock. Transportation contracts in both per-capita and raw-level form at a magnitude exceeding what population redistribution can produce, consistent with input-output linkage to the phasing-out manufacturing logistics chain through Section 936 sub-industries serving PR manufacturing. Arts is statistically underpowered ($n = 9$ counties). The Transportation placebo failure through a known confounder motivates the direct manufacturing-control test reported next.

Manufacturing control. The Transportation placebo failure motivates a direct test of whether the main contraction is identified in isolation from the Section 936 manufacturing phase-out. The test conditions the main healthcare-employment estimate on a pre-reform

[19]Online Appendix Tables OA.12 and OA.13 report the per-thousand-resident and raw-level placebo specifications.

manufacturing-exposure covariate (the 1990–1993 mean of NAICS 31–33 employment per thousand at the county level). Conditioning on a strictly pre-1996 baseline—rather than on the contemporaneous post-treatment manufacturing series—ensures that the covariate is not itself shocked by the policy whose contamination the test is designed to net out.[20] The C&S point estimate on healthcare employment per thousand attenuates from -14.7 percent of pre-mean (no control, $p < 0.01$) to -6.8 percent (with baseline manufacturing exposure as a conditioning covariate, statistically insignificant). The point estimate remains negative, and the attenuation is not driven by sample restriction. A reading on which the main result is purely §936 contamination is therefore not supported—the conditional estimate is negative, not zero or positive—but neither is a reading on which the result is identified in isolation from §936 at full magnitude. The supportable reading is intermediate: some of the cross-sectional heterogeneity in the main contraction loads onto pre-reform manufacturing exposure, a correlation that may reflect §936 contamination but may also reflect that high-manufacturing counties are urban and industrial and therefore overlap with the heterogeneity dimensions on which the supply-side contraction concentrates (Section 6.3). The interpretation rests on the conjunction of the policy-text-tied first stage (which this conditioning does not attenuate), the heterogeneity-as-mechanism evidence, and the negative conditional point estimate.

Per-capita normalization. The per-capita normalization is operative throughout because per-thousand-resident outcomes place labor-market and supply-side responses on the welfare-relevant denominator and match the unit conventions of the closest comparison literatures (Wallace, 2023; Hackmann et al., 2025). The contraction in healthcare employment per thousand of -13.0 percent coexists with growth in raw-level $\log(1 + y)$ employment of approximately $+19.0$ percent over the same window: the local healthcare sector grew in

[20] Appendix Table A.3 reports the conditioning specification. The sample is the 54 counties with non-missing 1990–1993 baseline manufacturing data—a tightening of the 78-county main-text panel imposed by the conditioning covariate. The C&S no-control estimate on this sample is -14.7 percent of pre-mean, economically equivalent to the 78-county main result of -13.0 percent reported in Table 4, so the sample restriction itself does not drive the conditioning result.

absolute terms but more slowly than the local population.[21] The two are not in conflict: each identifies a distinct object. The per-capita estimate identifies the change in resources available per local resident—the welfare-relevant scaling; the raw-level estimate identifies the change in absolute headcount—the mechanically larger object wherever the denominator is itself changing. The distributional finding of the paper rests on the per-capita denominator and is not a feature of the absolute headcount.

8 Conclusion

Puerto Rico's 1993 reform transferred a 78-county public delivery network of diagnostic and treatment centers, Area and Regional Hospitals, and a quaternary referral complex to private managed-care organizations paid prospective capitation by the Commonwealth's new Health Insurance Administration. The reform's first stage operated cleanly on the ownership margin, and the post-reform equilibrium delivered fewer healthcare resources per capita: healthcare employment, ambulatory primary-care employment, the establishment count, and Medicare-certified hospital beds and services all contracted relative to population. The contraction loaded onto labor-force participation rather than the unemployment rate, with parallel declines in the crude birth rate and the demographic-balancing residual—a labor-supply-margin signature consistent with the exit option that U.S. citizenship and the legal portability of medical credentials extend to the displaced provider workforce. Patient-health outcomes whose biological response horizons fall within the four-year post-reform window did not deteriorate on average and fell by between 17 and 27 percent of pre-mean in the urban, larger, higher-pre-reform- healthcare-intensity, and workplace-center counties where the first-stage transition was most pronounced.

The main empirical finding is that a system-wide delivery-system privatization can deliver less measured healthcare per capita without proportional loss of measured patient quality.

[21] Appendix Table A.7 reports per-1,000 and $\log(1+y)$ specifications for QCEW healthcare employment, establishments, and the wage bill side by side; Online Appendix Table OA.2 reports the parallel specification on $\log(1+y)$ per-1,000 normalizations and $\log(1+y)$ raw-level outcomes.

The estimates do not vindicate the Hart et al. (1997) quality-deterioration corollary in this setting: where the reform had the most bite, infant mortality fell, in some cells by more than a quarter of its pre-treatment mean. They do not vindicate the supply-stimulus prediction of the insurance-expansion literature (Finkelstein, 2007; Dillender, 2022; Hackmann et al., 2025) in its unconditional form either: the BRFSS coverage expansion is real, but the local healthcare labor market contracted rather than expanded under the bundled treatment of capitation, network restriction, and ownership transfer. The data instead support a more modest claim: an ownership-and-payment-margin reform that constrains the operator's per-enrollee budget can reduce per-capita inputs without inflicting proportional welfare losses when the pre-reform public system was operating with slack on the noncontractible quality margin (Hart et al., 1997; Propper, 2018). A back-of-envelope welfare bound places the short-horizon welfare consequence between \$844 million and \$2.72 billion in 2025 dollars over the four-year post-reform window (Section 6.7), with the cost-saving channel the binding source of welfare improvement even when the patient-health channel is valued at zero.

For the cross-system privatization debates with which the paper opened, the contribution is the counterfactual itself. The MISSION Act community-care expansion at the Department of Veterans Affairs, the independent-sector contracting at the UK Department of Health and Social Care, and the continued expansion of risk-based managed care across U.S. Medicaid programs are each a different bundle of the three institutional margins the reform moved jointly: outside-option access without ownership transfer (VA), within-system competition without capitation of the residual public delivery system (NHS), and capitation without ownership transfer (Medicaid Managed Care). The Puerto Rican evidence speaks most directly to the delivery-system-transfer counterfactual that none of the three mainland reforms cleanly identifies, and the supply-side contraction documented here cannot be inferred from the existing causal evidence on the partial bundles (Cooper et al., 2011; Gaynor et al., 2013; Chan et al., 2023; Duggan et al., 2023; Duggan, 2004; Wallace, 2023; Marton et al., 2014). The directional implication is that contracting an entire publicly operated delivery network to private capitated operators should be expected to contract per-capita inputs to the local healthcare sector; whether this delivers a welfare improvement, a welfare loss, or a wash depends on the pre-reform public-system quality, the local provider-mobility option, and the cross-sectional incidence of the contraction. The decomposition itself—what share of the contraction is attributable to ownership transfer, what share to capitation, and what share to access expansion—cannot be identified within a single-reform setting where the three margins moved jointly.

Three limitations bound the inferences this paper supports. The post-reform observation window is four years; this is sufficient for the labor-market and short-horizon patient-health outcomes that constitute the main analysis but does not support inference on long-horizon adult cause-specific mortality, on health-capital accumulation, or on the equilibrium response of the privatized delivery network once the initial four-insurer regional-monopoly arrangement matured into the multi-insurer carve-out structure of the post-2002 period (Centers for Medicare and Medicaid Services, 2002, 2004). San Juan exerts disproportionate leverage on the supply-side estimates: excluding the metropolitan capital attenuates the healthcare-employment estimate from −13.0 to −5.3 percent of pre-mean and renders the hospital-bed estimate indistinguishable from zero, though the first-stage ownership transition survives the exclusion. The contemporaneous Section 936 manufacturing phase-out is a residual contamination concern: the direct manufacturing-control test (Section 7) attenuates the main healthcare-employment estimate from −13.0 to −6.8 percent of pre-mean under pre-reform manufacturing conditioning. The contraction is not eliminated by the conditioning, but its magnitude is sensitive to it—some of the cross-sectional variation in the main contraction loads onto pre-reform manufacturing exposure, a correlation that may reflect §936 contamination or may reflect that high-manufacturing counties overlap with the urban-industrial heterogeneity dimensions on which the supply-side contraction concentrates. Identification of the healthcare-sector contraction therefore rests on the joint pattern of the policy-text-tied first stage (which is not attenuated by any robustness specification), the cross-sectional heterogeneity-as-mechanism evidence, and a manufacturing-conditional point estimate that remains negative.

9 Tables

Table 1: County Descriptive Statistics by Year

	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001
Panel A: QCEW — Healthcare Labor Market †												
Employment per 1k pop.	2.86	3.28	3.44	3.61	3.90	4.03	4.54	5.02	5.52	5.92	6.22	7.67
	(4.74)	(5.15)	(5.05)	(5.23)	(5.41)	(5.33)	(5.56)	(5.85)	(6.49)	(6.76)	(7.16)	(9.94)
Establishments per 1k pop.	0.62	0.66	0.73	0.73	0.75	0.82	0.89	0.93	0.99	1.07	1.11	1.13
	(0.54)	(0.54)	(0.57)	(0.58)	(0.57)	(0.60)	(0.62)	(0.64)	(0.65)	(0.67)	(0.68)	(0.69)
Wage bill per 1k pop. ($)	5,405	6,558	7,059	7,454	7,966	8,478	9,490	10,651	11,852	12,772	13,439	16,981
	(10,809)	(12,859)	(13,280)	(13,985)	(13,962)	(14,211)	(15,151)	(16,101)	(18,296)	(19,145)	(20,646)	(26,925)
Counties	59	59	59	59	59	59	59	59	59	59	59	59
Treated Counties	–	–	–	–	5	16	38	38	58	58	59	59
Panel B: CMS — Provider of Services ‡												
Total Providers	–	202	222	242	248	254	259	262	265	268	274	275
Government	–	42	42	48	48	48	44	44	44	45	45	43
Private	–	160	180	194	200	206	215	218	221	223	229	232
Providers per 1k pop.	–	0.08	0.09	0.09	0.09	0.10	0.10	0.10	0.10	0.10	0.10	0.10
	(–)	(0.05)	(0.06)	(0.06)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)	(0.07)
Beds per 1k pop.	–	3.17	3.15	3.21	3.18	3.18	3.16	3.10	3.08	3.18	3.20	3.21
	(–)	(3.08)	(3.06)	(3.06)	(3.04)	(3.09)	(3.14)	(3.09)	(3.07)	(3.28)	(3.26)	(3.26)
Services per 1k pop.	–	0.46	0.48	0.48	0.49	0.49	0.49	0.48	0.47	0.48	0.48	0.49
	(–)	(0.47)	(0.49)	(0.49)	(0.49)	(0.50)	(0.49)	(0.47)	(0.47)	(0.47)	(0.47)	(0.51)
FTE per 1k pop.	–	9.04	9.17	9.54	12.85	12.54	12.41	12.33	12.17	12.10	12.36	12.61
	(–)	(8.34)	(8.44)	(8.45)	(19.54)	(19.46)	(18.40)	(18.09)	(17.96)	(18.14)	(18.96)	(19.14)
Counties	–	40	40	40	40	40	40	40	40	40	40	40
Treated Counties	–	–	–	–	5	10	29	29	39	39	40	40
Panel C: VS — Demographic Flow, Hospital Ownership, and Patient Health ¶												
Population	45,228	45,502	45,884	46,435	47,253	47,684	47,863	48,792	49,147	49,571	48,935	49,228
	(59,205)	(59,564)	(60,064)	(60,778)	(61,282)	(59,393)	(58,842)	(59,597)	(59,858)	(60,151)	(59,129)	(58,987)
Crude birth rate per 1k pop.	19.38	18.45	18.54	18.66	18.08	17.46	17.42	17.42	16.19	15.69	15.64	14.64
	(2.38)	(2.01)	(2.24)	(1.98)	(1.86)	(2.11)	(2.32)	(2.19)	(1.94)	(2.21)	(1.44)	(1.48)
Crude death rate per 1k pop.	7.02	6.97	7.19	7.43	7.28	7.67	7.57	7.44	7.55	7.23	7.16	7.21
	(1.14)	(1.01)	(1.01)	(1.06)	(1.20)	(1.31)	(1.28)	(1.19)	(1.15)	(1.10)	(1.11)	(1.00)
Crude net migration rate per 1k pop.	–	-5.51	-3.04	0.66	8.56	10.90	-2.87	13.48	-1.47	1.91	-9.89	1.04
	(–)	(2.28)	(2.72)	(2.51)	(29.54)	(67.48)	(26.57)	(20.11)	(17.10)	(17.97)	(78.74)	(8.35)
Government hospital birth share (%)	65.8	65.9	65.3	64.3	63.2	59.1	51.6	45.2	36.6	33.6	30.6	32.1
	(11.3)	(11.2)	(10.9)	(10.7)	(10.7)	(12.1)	(15.2)	(15.3)	(14.6)	(18.6)	(19.8)	(22.7)
Private hospital birth share (%)	33.8	33.8	34.4	35.5	36.5	39.6	48.1	54.6	63.1	66.2	69.2	67.8
	(11.4)	(11.3)	(11.0)	(10.7)	(10.9)	(11.6)	(15.3)	(15.3)	(14.6)	(18.6)	(19.9)	(22.8)
Infant mortality rate per 1k live births	12.74	12.67	12.71	13.40	11.84	14.28	9.91	12.10	10.28	10.79	10.15	9.62
	(6.09)	(5.30)	(4.47)	(5.45)	(4.97)	(7.33)	(4.90)	(6.42)	(5.13)	(4.30)	(5.45)	(5.04)
Low birth weight rate (%)	9.28	9.25	9.48	9.64	9.85	10.12	10.49	11.06	10.98	11.36	10.96	11.48
	(1.62)	(1.79)	(2.02)	(1.66)	(1.43)	(1.90)	(1.90)	(2.32)	(2.24)	(2.46)	(2.48)	(4.07)
Counties	78	78	78	78	78	78	78	78	78	78	78	78
Treated Counties	–	–	–	–	10	22	52	52	77	77	78	78
Panel D: PR DOL — Macro Labor Market #												
Labor force / pop. (%; Figure 3)	30.5	31.4	32.0	32.2	31.4	31.8	32.4	32.5	32.3	31.3	31.5	30.9
	(4.6)	(4.7)	(5.1)	(5.1)	(4.8)	(4.9)	(5.1)	(5.0)	(5.2)	(5.2)	(4.4)	(4.3)
Unemployment rate (%; U / LF)	16.9	18.4	19.3	20.1	17.1	16.0	15.8	16.1	15.8	14.2	11.6	13.1
	(4.7)	(4.9)	(5.1)	(5.1)	(4.6)	(4.7)	(4.5)	(4.8)	(4.6)	(4.7)	(2.2)	(2.7)
Counties	78	78	78	78	78	78	78	78	78	78	78	78
Treated Counties	–	–	–	–	10	22	52	52	77	77	78	78

Means with standard deviations in parentheses. Sample is the PR balanced-panel counties (78 total); per-dataset counts below). † *Panel A (QCEW):* Overall Healthcare = NAICS 621 (Ambulatory) + 622 (Hospitals) + 623 (Nursing/Residential). All wage variables are CPI-deflated using 1982-1984 as the base year. ‡ *Panel B (CMS):* Provider counts are point-in-time totals from the CMS Provider of Services file; Government / Private split is the contemporaneous ownership classification. The trajectory rows (Always Government, Always Private, Gov. → Private switchers) classify each provider over its full panel history and are stable across years. CMS balanced panel 1991–2000: 40 counties.*CMS provider mix (provider-years, 1991–2000):* The CMS county panel summarizes 2,586 provider-year observations from 296 unique facilities. By ownership: 454 (17.6%) government, 872 (33.7%) private for-profit, 1,024 (39.6%) private nonprofit, 236 (9.1%) private-unspecified. By facility category, 36.8% are Hospitals (CMS POS code 01); the remaining provider categories are Hospital (36.8%); Home Health Agency (19.8%); Hospice (18.1%); ESRD Facility (11.9%); Ambulatory Surgical Center (7.8%); Other (5.6%). ¶ *Panel C (VS):* Crude birth rate is live births per 1,000 mid-year population. Crude death rate is deaths per 1,000 mid-year population. Crude net migration rate is the residual of the demographic-balancing equation: $[(\text{Pop}_t - \text{Pop}_{t-1}) - (\text{Births}_t - \text{Deaths}_t)]/\text{Pop}_{t-1} \times 1000$, NA in 1990 (no $t-1$ population). Government and private hospital birth shares are percentages of all live births. Low birth weight rate is the percentage of live births under 2,500g. VS balanced panel 1990–2001: 78 counties. [2pt]# *Panel D (PR DOL):* Labor force / pop. is labor force divided by total county population × 100; employment / pop. is employment divided by total county population × 100. Both use the non-standard denominator (total population) rather than the BLS-standard civilian non-institutional population aged 16+ (CNIWAP), which is not available at the county level for the 1990–2001 panel. The unemployment rate is reported as BLS-standard (unemployed / labor force × 100). PR DOL balanced panel 1990–2001: 78 counties. Across all 936 county-year observations: BLS-style U/LF mean = 16.19%; alternative U/Pop mean = 5.04%; difference = 11.16 pp; Pearson correlation between the two series = 0.878. The two series differ in level but not in trends. *Pre-reform (1990–1993) cross-county medians used as thresholds in heterogeneity figures (Main Figure 5; Appendix Figures A.5–A.7; Online Appendix Figures OA.21–OA.22):* Government hospital birth share = 63.2%; Private hospital birth share = 36.3%; Healthcare employment per 1k pop. = 1.51; Population = 30,362; Unemployment rate = 18.6%. Each median is computed across counties after first averaging the outcome over the 1990–1993 window within each county.

Table 2: Pre-Reform Balance Across Treatment Cohorts (1990–1993 Means)

Variable	Cohort 1 (1994) *Mean*	Cohort 2 (1995) *Mean*	Cohort 3 (1996) *Mean*	Cohort 4 (1998) *Mean*	Cohort 5 (2000) *Mean*	F-test *p-value*	N
Total population	43,889	34,467	37,871	45,489	443,627	<0.001	78
Log population	10.124	10.287	10.330	10.427	13.003	0.013	78
Share urban (top-20 by 1990 density)	0.600	0.417	0.433	0.560	1.000	0.638	78
Gov. hospital birth share (%)	73.57	66.70	67.17	59.30	61.64	0.003	78
HC employment per 1k pop.	4.01	2.28	2.64	2.87	29.72	<0.001	60
Labor force participation rate (%)	34.96	29.79	28.96	33.90	35.47	<0.001	78
Unemployment rate (%)	15.23	19.60	20.97	17.12	10.70	<0.001	78
Crude net migration rate (per 1k)	-3.80	-2.87	-3.25	-1.48	1.89	0.002	78
Infant mortality rate (per 1k)	13.59	11.21	13.74	12.10	13.02	0.139	78
Crude birth rate (per 1k)	19.90	19.11	19.39	17.43	17.22	<0.001	78
Agriculture, forestry & fishing emp. per 1k	1.451	16.195	8.422	15.991	–	0.393	25
Transportation & warehousing emp. per 1k	3.595	0.623	0.644	1.287	9.811	0.037	78
Arts, entertainment & recreation emp. per 1k	4.745	–	0.281	0.342	2.145	0.525	11
Counties	10	12	30	25	1		78

Cell entries are 1990–1993 means of pre-reform covariates by treatment cohort. Cohorts 1–5 correspond to reform-rollout years 1994, 1995, 1996, 1998, and 2000 respectively (no rollouts in 1997 or 1999). The F-test column reports the p-value of the one-way ANOVA test for the joint null hypothesis that all five cohort means are equal (`aov(var ~ factor(cohort))`); a small p-value indicates that the listed covariate differs systematically across cohorts and is therefore not balanced under the rollout schedule. **Sources.** Population, gov. hospital birth share, infant mortality, crude birth rate, and crude net migration rate are 1990–1993 means from the Puerto Rico Department of Health vital statistics. Healthcare employment per 1,000 residents (NAICS Overall Healthcare = 621 + 622 + 623) and the three placebo-industry rows (Agriculture/Forestry/Fishing = NAICS 11; Transportation & Warehousing = NAICS 48-49; Arts/Entertainment/Recreation = NAICS 71) are 1990–1993 means computed from the QCEW (BLS) county-quarter panel: aggregated across ownership classes, averaged to annual frequency requiring at least 3 reporting quarters per year, then divided by the population. Labor force participation and unemployment rates are 1990–1993 means from the PR Department of Labor. The urban indicator equals one for the 20 counties with the highest 1990 population density and zero otherwise; the cohort-cell entry is the share of urban counties within the cohort. The three placebo-industry rows (NAICS 11, 48-49, 71) are the same industries used in the QCEW placebo robustness analysis (Online Appendix Section 7)

Table 3: Effects of the Reform on Hospital Ownership Composition (First Stage)

	(1) Gov. Hospital Birth Share (%)	(2) Priv. Hospital Birth Share (%)
ATT	-14.79***	15.19***
	(1.94)	(1.93)
Pre-treatment mean of dep. var.	62.55	37.04
ATT as % of pre-treatment mean	-23.6%	41.0%
N counties	78	78

Estimates use the Callaway and Sant'Anna (2021) ATT estimator with not-yet-treated comparison units. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. The pre-treatment mean of the dependent variable is computed across treated counties using observations from the cohort-specific pre-treatment period (event time < 0). Government and private hospital birth shares are defined as the count of live births occurring at facilities of each ownership class divided by the total count of live births in the county-year, expressed in percent. The two shares are not exact complements because a small tail of births (typically <1%) is recorded as non-hospital or with unspecified facility ownership.

Table 4: Reform Effects on Healthcare Employment, Wages, and Establishments

	Healthcare (NAICS 621+622+623)			NAICS 621 (Ambulatory)		
	(1) Empl. per 1k	(2) Wages/Empl. per 1k	(3) Estabs. per 1k	(4) Empl. per 1k	(5) Wages/Empl. per 1k	(6) Estabs. per 1k
ATT	-0.511***	0.630	-0.050***	-0.502***	0.867	-0.034*
	(0.157)	(3.734)	(0.019)	(0.125)	(3.682)	(0.018)
Pre-treatment mean of dep. var.	3.923	38.962	0.759	2.895	38.889	0.747
ATT as % of pre-treatment mean	-13.0%	1.6%	-6.6%	-17.3%	2.2%	-4.6%
N counties	59	59	59	59	59	59

Estimates use the Callaway and Sant'Anna (2021) ATT estimator with not-yet-treated comparison units. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. The pre-treatment mean of the dependent variable is computed across treated counties using observations from the cohort-specific pre-treatment period (event time < 0). The reference period for event-time aggregation is $t = -1$ (year before treatment).

Table 5: Effects of The Reform on Patient Health (Short-Horizon) and Hospital Capacity

	Patient Health (VS)		Hospital Capacity (CMS)	
	(1) IMR (per 1k live births)	(2) Low Birth Weight (%)	(3) Beds per 1k pop.	(4) Services per 1k pop.
ATT	-1.205	0.060	-0.265***	-0.024**
	(0.928)	(0.230)	(0.065)	(0.010)
Pre-treatment mean of dep. var.	12.620	9.717	3.389	0.487
ATT as % of pre-treatment mean	-9.5%	0.6%	-7.8%	-4.9%
N counties	78	78	40	40

Estimates use the Callaway and Sant'Anna (2021) ATT estimator with not-yet-treated comparison units. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$, based on two-sided tests with cluster-robust standard errors. The pre-treatment mean of the dependent variable is computed across treated counties using observations from the cohort-specific pre-treatment period (event time < 0). The reference period for event-time aggregation is $t = -1$ (year before treatment).

Table 6: Worker-Side Adjustment Margins to The Reform

	(1) Labor force / pop. (%)	(2) Crude Net Migration (per 1k)	(3) Crude Birth Rate (per 1k)
ATT	-1.110***	-14.651**	-0.564**
	(0.241)	(7.389)	(0.225)
Pre-treatment mean of dep. var.	31.895	1.059	18.267
ATT as % of pre-treatment mean	-3.5%	-1382.9%	-3.1%
N counties	78	78	78

Estimates use the Callaway and Sant'Anna (2021) ATT estimator with not-yet-treated comparison units. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$, based on two-sided tests with cluster-robust standard errors. The pre-treatment mean of the dependent variable is computed across treated counties using observations from the cohort-specific pre-treatment period (event time < 0). The reference period for event-time aggregation is $t = -1$ (year before treatment).

Table 7: Heterogeneity in the Effects by Pre-Reform Public-Hospital Exposure

	(1) Gov. Hosp. Birth Share (%)	(2) HC Empl. per 1k	(3) Crude Net Migr. (per 1k)	(4) IMR (per 1k live births)
Below-median pre-reform government hospital birth share				
ATT	-15.49***	-0.346	-9.206	-1.454
	(2.44)	(0.233)	(10.814)	(1.101)
Pre-treatment mean	55.49	4.911	2.772	11.577
ATT as % of pre-mean	-27.9%	-7.0%	-332.1%	-12.6%
N counties	39	36	39	39
Above-median pre-reform government hospital birth share				
ATT	2.41	0.145	-23.891**	0.660
	(9.15)	(0.268)	(12.047)	(1.329)
Pre-treatment mean	71.21	2.069	-1.126	13.925
ATT as % of pre-mean	3.4%	7.0%	2121.2%	4.7%
N counties	39	23	39	39

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator with not-yet-treated comparison units. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$, based on two-sided tests with cluster-robust standard errors. counties are partitioned at the cross-county median of the 1990–1993 pre-reform government hospital birth share (median = 63.24%). *Below-median* (39 counties) = lower pre-reform reliance on the public hospital system; *Above-median* (39 counties) = higher pre-reform reliance.

10 Figures

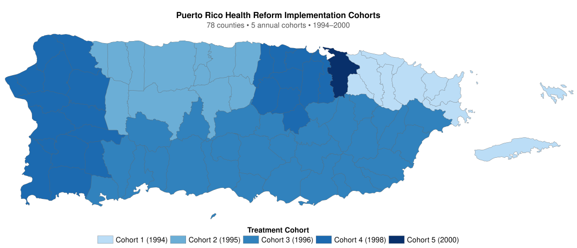

(a) Cohorts

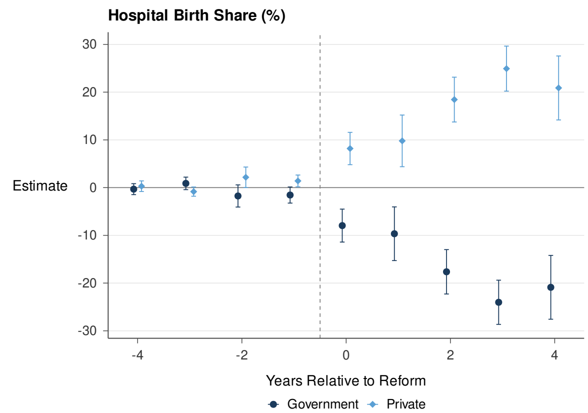

(b) Birth Share (pct)

Figure 1: Policy Rollout

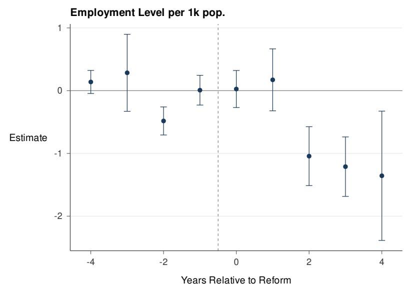

(a) Employment (per 1k) - 621+622+623

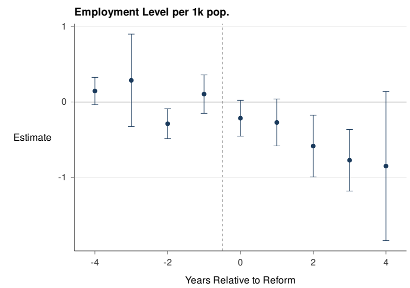

(b) Employment (per 1k) - 621

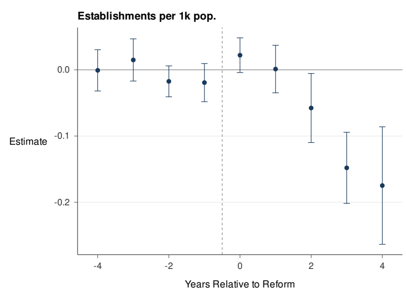

(c) Establishments (per 1k) - 621+622+623

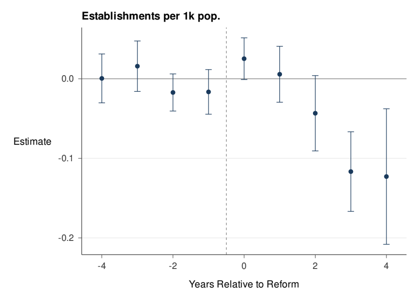

(d) Establishments (per 1k) - 621

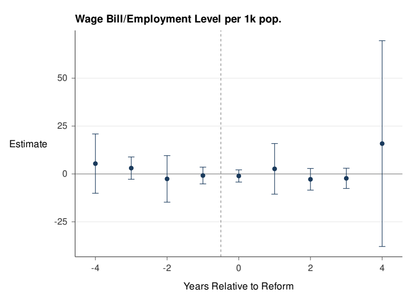

(c) Wage Bill/Employee (per 1k) - 621+622+623

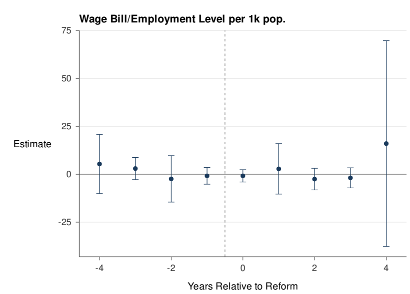

(d) Wage Bill per Employee (per 1k) - 621

Figure 2: Healthcare Labor-Market Contraction

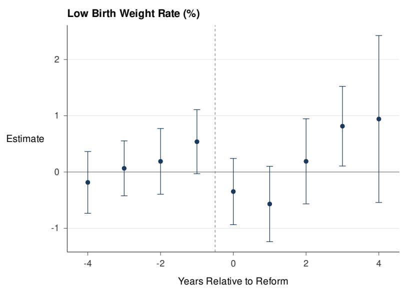

(a) Low Birth Weight

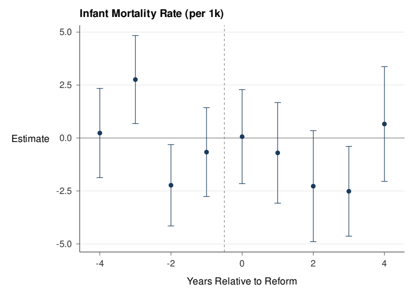

(b) Infant Mortality Rate

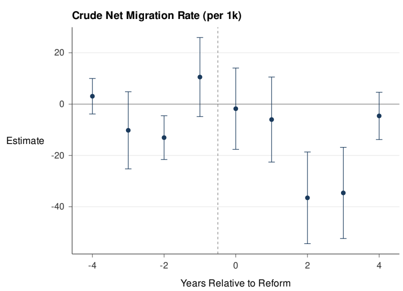

(c) Net Migration Rate

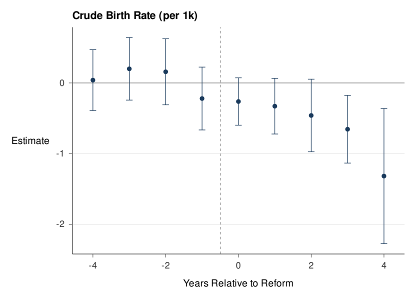

(d) Birth Rate

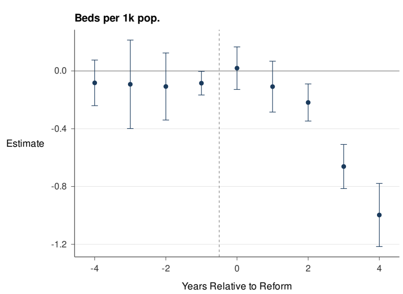

(e) Beds

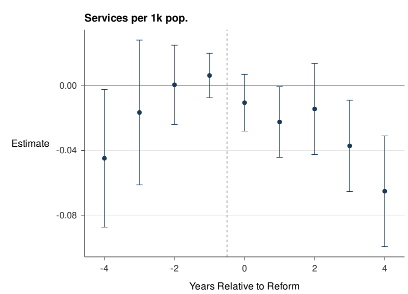

(f) Services

Figure 3: Patient Health

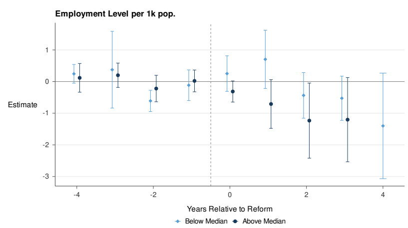

(a) Employment (per 1k) - 621+622+623

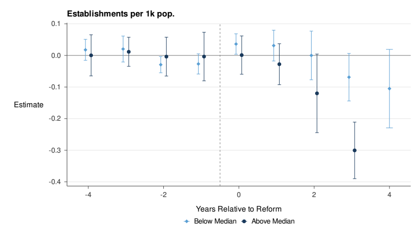

(b) Establishments (per 1k) - 621+622+623

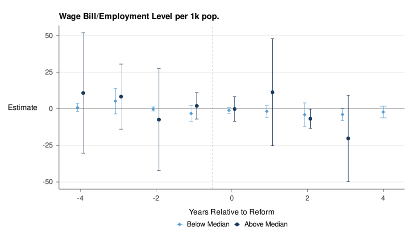

(c) Wage Bill per Employee (per 1k) - 621+622+623

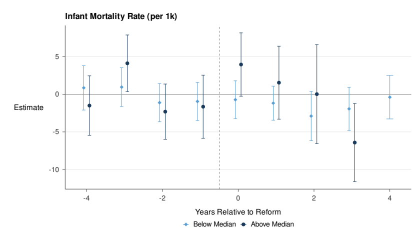

(d) Infant Mortality Rate (per 1k)

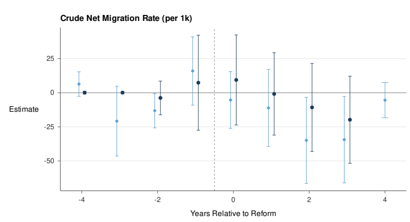

(d) Net Migration Rate

Figure 4: Heterogenous Effects by Poverty Rate

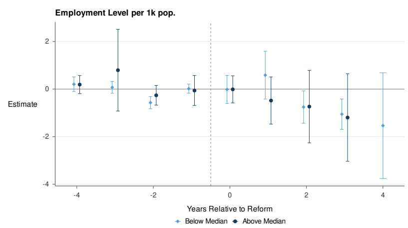

(a) Employment (per 1k) - 621+622+623

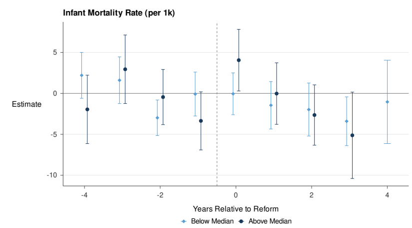

(b) Infant Mortality Rate (per 1k)

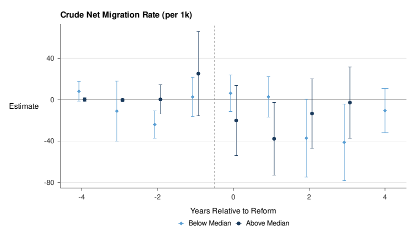

(c) Net Migration Rate (per 1k)

Figure 5: Heterogeneous Effects by Pre-Reform Gov. Hospital Exposure

References

Aizer, A., Currie, J., and Moretti, E. (2007). Does managed care hurt health? evidence from Medicaid mothers. *Review of Economics and Statistics*, 89(3):385–399.

Andreyeva, E., Gupta, A., Ishitani, C., Sylwestrzak, M., and Ukert, B. (2024). The corporatization of independent hospitals. *Journal of Political Economy Microeconomics*, 2(3):603–665.

Arbona, G. and Ramirez de Arellano, A. (1978). *Regionalization of Health Services: The Puerto Rican Experience*. Oxford University Press, New York. Published for the International Epidemiological Association and the World Health Organization.

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5):941–973.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168.

Bloom, N., Propper, C., Seiler, S., and Van Reenen, J. (2015). The impact of competition on management quality: Evidence from public hospitals. *Review of Economic Studies*, 82(2):457–489.

Borusyak, K., Jaravel, X., and Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6):3253–3285.

Brugueras Fabre, A. J. (2002). Los costos de la reforma de salud de Puerto Rico. Cited at p. 17 for the eight-region rollout date schedule.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Centers for Medicare and Medicaid Services (2000). National summary of state Medicaid managed care programs as of June 30, 2000. Technical report, U.S. Department of Health and Human Services, Baltimore, MD.

Centers for Medicare and Medicaid Services (2001). National summary of state Medicaid managed care programs as of June 30, 2001. Technical report, U.S. Department of Health and Human Services, Baltimore, MD.

Centers for Medicare and Medicaid Services (2002). National summary of state Medicaid managed care programs as of June 30, 2002. Technical report, U.S. Department of Health and Human Services, Baltimore, MD. First year of mental-health/substance-abuse carve-out to MBHOs in Puerto Rico.

Centers for Medicare and Medicaid Services (2003). National summary of state Medicaid managed care programs as of June 30, 2003. Technical report, U.S. Department of Health and Human Services, Baltimore, MD.

Centers for Medicare and Medicaid Services (2004). National summary of state Medicaid managed care programs as of June 30, 2004. Technical report, U.S. Department of Health and Human Services, Baltimore, MD. First year including Medicare dual eligibles in Puerto Rico's Reforma program.

Centers for Medicare and Medicaid Services (2026). Medicaid & CHIP in Puerto Rico: State overview. Online state-overview page.

Chan, D. C., Card, D., and Taylor, L. (2023). Is there a VA advantage? Evidence from dually eligible veterans. *American Economic Review*, 113(11):3003–3043.

Commonwealth of Puerto Rico (1993). Ley núm. 72 de 7 de septiembre de 1993: Ley de la Administración de Seguros de Salud de Puerto Rico. Laws of Puerto Rico. Puerto Rico Health Insurance Administration Act.

Congressional Budget Office (2021). The veterans community care program: Background and early effects. CBO Publication 57583, Congressional Budget Office. VHA spending on community care grew from $7.9 billion in 2014 to $17.6 billion in 2021 (in 2021 dollars); community care rose from approximately 12 to 20 percent of VHA's medical-care budget over the same period.

Congressional Research Service (2025). Department of veterans affairs FY2025 appropriations. CRS Report R48608, Congressional Research Service. VHA medical-care discretionary appropriations of $112.58 billion for FY2025; reports VHA as "the nation's largest integrated health care delivery system" with more than 1,700 sites of care.

Cooper, Z., Gibbons, S., Jones, S., and McGuire, A. (2011). Does hospital competition save lives? Evidence from the English NHS patient choice reforms. *Economic Journal*, 121(554):F228–F260.

Dillender, M. (2022). What happens to the health care labor market when demand increases? evidence from Medicaid expansions. *Journal of Human Resources*. Forthcoming.

Duggan, M. (2004). Does contracting out increase the efficiency of government programs? evidence from Medicaid HMOs. *Journal of Public Economics*, 88(12):2549–2572.

Duggan, M., Gupta, A., Jackson, E., and Templeton, Z. S. (2023). The impact of privatization: Evidence from the hospital sector. Working Paper 30824, National Bureau of Economic Research.

Eliason, P. J., Heebsh, B., McDevitt, R. C., and Roberts, J. W. (2020). How acquisitions affect firm behavior and performance: Evidence from the dialysis industry. *Quarterly Journal of Economics*, 135(1):221–267.

Farmer, C. M. (2023). The promise and challenges of VA community care: Veterans' issues in focus. *RAND Health Quarterly*, 10(3):9. VA spending on community care grew from $7.9 billion in 2014 to $18.5 billion in 2021; community care accounted for 44 percent of VA's health care services across care settings as of 2022 testimony to the House Veterans' Affairs Committee.

Finkelstein, A. (2007). The aggregate effects of health insurance: Evidence from the introduction of Medicare. *Quarterly Journal of Economics*, 122(1):1–37.

Gaynor, M., Ho, K., and Town, R. J. (2015). The industrial organization of health-care markets. *Journal of Economic Literature*, 53(2):235–284.

Gaynor, M., Moreno-Serra, R., and Propper, C. (2013). Death by market power: Reform, competition, and patient outcomes in the National Health Service. *American Economic Journal: Economic Policy*, 5(4):134–166.

Geruso, M., Layton, T. J., and Wallace, J. (2023). What difference does a health plan make? Evidence from random plan assignment in Medicaid. *American Economic Journal: Applied Economics*, 15(3):341–379.

Hackmann, M. B., Heining, J., Klimke, R., Polyakova, M., and Seibert, H. (2025). Health insurance as economic stimulus? evidence from long-term care jobs. Working Paper 33429, National Bureau of Economic Research.

Hart, O., Shleifer, A., and Vishny, R. W. (1997). The proper scope of government: Theory and an application to prisons. *Quarterly Journal of Economics*, 112(4):1127–1161.

Instituto de Administración y Política de Salud (1995). Coordinating health care reform with the U.S. territories and possessions: The case of Puerto Rico. Final Report Grant Number 18-C-90240/2-01, University of Puerto Rico, Medical Sciences Campus, San Juan, Puerto Rico. Prepared for the Health Care Financing Administration.

Kaiser Family Foundation (2024). Recent changes in Medicaid financing in Puerto Rico and other U.S. territories. KFF issue brief.

Kaiser Family Foundation (2025). 10 things to know about Medicaid Managed Care. KFF issue brief, last updated October 2025. Reports FFY 2024 figures: 291 contracting MCOs across 42 states and DC; 66 million enrollees in comprehensive MCOs; capitated MCO payments equal to 50 percent of total Medicaid spending of $919 billion.

Kaiser Family Foundation (2026). Medicaid financing: The basics. KFF issue brief, last updated March 2026. FFY 2024 totals: $919 billion total Medicaid spending, comprising $594 billion federal (65%) and $325 billion state (35%); capitated MCO payments accounted for 50% of Medicaid spending.

King's Fund (2026). The NHS budget and how it has changed. King's Fund data brief, last updated March 2026. DHSC total spending of £204.7 billion in 2024/25, of which £187 billion was day-to-day spending allocated to NHS England for health services.

Marton, J., Yelowitz, A., and Talbert, J. C. (2014). A tale of two cities? the heterogeneous impact of Medicaid managed care. *Journal of Health Economics*, 36:47–68.

Medicaid and CHIP Payment and Access Commission (2020). Medicaid and CHIP in Puerto Rico. Technical report, MACPAC.

National Audit Office (2026). Department of health and social care 2024-25: Overview. NAO Departmental Overview, January 2026. DHSC spent £219.2 billion gross (£204.7 billion after income) in 2024-25; £24.089 billion was spent on the purchase of healthcare from non-NHS bodies.

Pan American Health Organization (2007). Health systems profile: Puerto Rico—monitoring and analyzing health systems change/reform. Technical report, Pan American Health Organization/World Health Organization, Washington, D.C.

Propper, C. (2018). Competition in health care: Lessons from the English experience. *Health Economics, Policy and Law*, 13(3-4):492–508.

Wallace, J. (2023). What does a provider network do? evidence from random assignment in Medicaid managed care. *American Economic Journal: Economic Policy*, 15(1):473–509.

Wooldridge, J. M. (2025). Two-way fixed effects, the two-way Mundlak regression, and difference-in-differences estimators. *Empirical Economics*, 69:2545–2587.

World Health Organization and Pan American Health Organization (1998). La reforma del sector salud: El caso de Puerto Rico. Technical report, Pan American Health Organization, Washington, D.C.

Yagan, D. (2019). Employment hysteresis from the Great Recession. *Journal of Political Economy*, 127(5):2505–2558.

Appendix

The Appendix collects the supplementary tables referenced in the body. Each table is referenced from a footnote to the relevant subsection of the main text. The Online Appendix (Section 10) hosts the remaining robustness, sample-restriction, and subgroup specifications.

A.1 Estimator-Class Robustness

Reports four-estimator triangulation (TWFE; Wooldridge, 2025 ETWFE; Callaway and Sant'Anna, 2021 doubly robust C&S; and Borusyak et al., 2024 BJS) on the six primary outcomes. All satisfy pre-treatment parallel trends for the Labor Market outcomes. Additionally, ETWFE, C&S, and BJS agree on sign and magnitude for the first-stage hospital-ownership share, healthcare employment, labor-force participation, and hospital-bed contraction. TWFE produces sign reversals on outcomes where treatment-effect heterogeneity is most pronounced (healthcare employment, infant mortality, low-birth-weight share). The crude net migration rate is the one outcome on which the four estimators do not agree on sign.

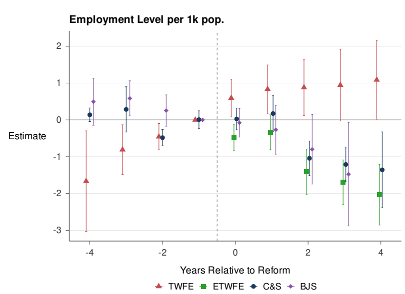

(a) Employment Level per 1k

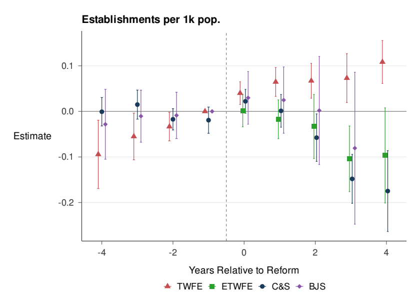

(b) Establishments per 1k

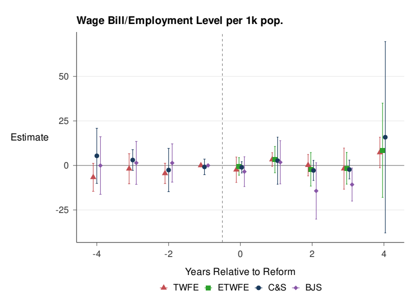

(c) Wage Bill per Employment per 1k

Figure A.1: Overall Healthcare Labor Market - Estimator Comparison

Table A.1: Four-Estimator Robustness: ATT Estimates Under TWFE, ETWFE, C&S, and BJS

	(1) TWFE	(2) ETWFE	(3) C&S	(4) BJS
Gov. Hospital Birth Share (%)				
ATT	-8.08***	-16.52***	-14.79***	-11.25***
	(1.89)	(1.60)	(1.94)	(1.69)
ATT as % of pre-mean	-12.9%	-26.4%	-23.6%	-18.0%
Pre-treatment mean	62.55			
counties	78			
HC Employment per 1k pop. (NAICS 621+622+623)				
ATT	0.868**	-1.008***	-0.511***	-0.655***
	(0.352)	(0.181)	(0.157)	(0.217)
ATT as % of pre-mean	22.1%	-25.7%	-13.0%	-16.7%
Pre-treatment mean	3.923			
counties	59			
Labor force / pop. (%)				
ATT	0.059	-1.155***	-1.110***	-0.712***
	(0.350)	(0.231)	(0.241)	(0.235)
ATT as % of pre-mean	0.2%	-3.6%	-3.5%	-2.2%
Pre-treatment mean	31.895			
counties	78			
Crude Net Migration Rate (per 1k pop.)				
ATT	-5.689	-1.525	-14.651**	4.513
	(15.281)	(2.931)	(7.389)	(3.523)
ATT as % of pre-mean	-537.1%	-144.0%	-1382.9%	426.0%
Pre-treatment mean	1.059			
counties	78			
Infant Mortality Rate (per 1k live births)				
ATT	1.973**	-1.524**	-1.205	-0.513
	(0.815)	(0.599)	(0.928)	(0.727)
ATT as % of pre-mean	15.6%	-12.1%	-9.5%	-4.1%
Pre-treatment mean	12.620			
counties	78			
Beds per 1k pop.				
ATT	-0.387	-0.443***	-0.265***	-0.189***
	(0.333)	(0.072)	(0.065)	(0.056)
ATT as % of pre-mean	-11.4%	-13.1%	-7.8%	-5.6%
Pre-treatment mean	3.389			
counties	40			

Notes: Four estimators are reported. TWFE: two-way fixed-effects regression with county and year FEs and event-time dummies. ETWFE: Wooldridge (2021) Extended Two-Way Fixed Effects. C&S: Callaway and Sant'Anna (2021). BJS: Borusyak, Jaravel & Spiess (2024). Significance levels: $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$, based on two-sided tests with cluster-robust standard errors.

A.2 Excluding San Juan

Reports the main-vs.-excluded-San-Juan specifications for the QCEW labor-market and CMS hospital-capacity outcomes. The healthcare-employment and bed-capacity estimates attenuate substantially when San Juan is removed; the first-stage ownership-transition estimate (Online Appendix Table OA.3) survives the exclusion at conventional significance.

Table A.2: Sample-Specification Robustness: Excluding San Juan and the Unbalanced-Panel Variant

	HC Empl. per 1k pop.			Beds per 1k pop.			Infant Mortality Rate		
	(1) Bal.	(2) Excl. SJ	(3) Unbal.	(4) Bal.	(5) Excl. SJ	(6) Unbal.	(7) Bal.	(8) Excl. SJ	(9) Unbal.
ATT	-0.511***	-0.168	-0.511***	-0.265***	0.090	-0.265***	-1.205	0.781	-1.205
	(0.157)	(0.191)	(0.157)	(0.065)	(0.140)	(0.065)	(0.928)	(1.237)	(0.928)
Pre-treatment mean of dep. var.	3.923	3.161	3.892	3.389	3.033	2.980	12.620	12.624	12.620
ATT as % of pre-treatment mean	-13.0%	-5.3%	-13.1%	-7.8%	3.0%	-8.9%	-9.5%	6.2%	-9.5%
counties	59	58	60	40	39	54	78	77	78

Notes: Estimates use the Callaway and Sant'Anna (2021). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$, based on two-sided tests with cluster-robust standard errors.

A.3 Section 936 Manufacturing Control

Reports the main healthcare-employment estimate with and without conditioning on pre-reform manufacturing exposure (1990–1993 mean of NAICS 31–33 employment per thousand at the county level). The C&S no-control estimate on the 54-county sample is -14.7 percent of pre-mean ($p < 0.01$); the C&S estimate with the strictly-pre-1996 manufacturing-exposure covariate is -6.8 percent of pre-mean (statistically insignificant). The TWFE columns are reported for direct horse-race comparison and return the positive point estimates expected under Goodman-Bacon contamination on heterogeneous-effect outcomes. The main contraction survives conditioning on §936 exposure at attenuated magnitude.

Table A.3: IRS Section 936 Manufacturing-Control Robustness: Healthcare Employment per 1,000 Residents under TWFE and Callaway-Sant'Anna, with and without Manufacturing Exposure as a Control

	TWFE		Callaway-Sant'Anna	
	(1) No control	(2) With manuf. control	(3) No control	(4) With baseline manuf. exposure
ATT on HC empl. per 1k	0.173	0.182	-0.484***	-0.224
	(0.198)	(0.215)	(0.164)	(0.233)
ATT as % of pre-mean	5.3%	5.5%	-14.7%	-6.8%
Coef. on contemp. manuf. per 1k	—	-0.003	—	—
	—	(0.013)	—	—
counties	54	54	54	54

Notes: Cluster-robust standard errors at the county level. Outcome: healthcare employment per 1,000 residents (NAICS 621+622+623). Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. **Estimators.** Columns 1–2 report TWFE event-study estimates with county and year fixed effects.

A.4 Welfare Bounds

Translates the main supply contraction and the patient-health response into comparable units (1982–1984 USD, with parenthetical translations into 2025 USD). The four cells span two parameter dimensions: the healthcare-employment-contraction estimate (main C&S versus §936-conditional C&S, the latter from Table A.3) and whether the patient-health response is valued at zero (the main IMR null) or at the high-exposure-half magnitude (the heterogeneity-cell IMR reduction averaged across density, population, healthcare-intensity, and commuting-inflow cuts of Section 6.3). The lower-bound cell is $300.8 million in 1982–1984 dollars ($967 million in 2025 dollars; $21.5 per resident per year); the upper-bound cell is $845.0 million ($2.72 billion in 2025 dollars; $60.4 per resident per year). The §936-conditional cells span $262.5 million to $806.6 million in 1982–1984 dollars ($844 million to $2.59 billion in 2025 dollars). The exercise is partial-equilibrium and back-of-envelope, not a structural welfare calculation.

Table A.4: Welfare Bounds: Cost Savings, Patient-Health Benefits, and Net Welfare

	Lower bound (ΔIMR = 0) *(1982-84 $M)*	Upper bound (ΔIMR high-exposure half) *(1982-84 $M)*
Panel A: Headline ATT (Tables 2 and 4)		
Cost saving ($M)	300.8	300.8
Health benefit ($M)	0.0	544.1
Net welfare ($M)	**300.8**	**845.0**
Per capita per year ($)	21.5	60.4
Panel B: § 936-Conditional ATT (Table A.4 col. 4)		
Cost saving ($M)	262.5	262.5
Health benefit ($M)	0.0	544.1
Net welfare ($M)	**262.5**	**806.6**
Per capita per year ($)	18.7	57.6

Notes: Partial-equilibrium back-of-envelope welfare bound; **not** a structural calculation. Body cells are in real 1982-84 USD (CPI-U all-items, 1982-84 = 100); the headline upper bound corresponds to approximately $2,716 M in 2025 USD. **Cost-saving inputs.** Δbeds/1k = -0.265 (Table 4); ΔHC empl./1k = -0.511 headline (Table 2) and -0.224 § 936-conditional (Table A.4 col. 4). Cost per bed-year ($62,756, 1982-84 USD) is the implied 1990–1993 mean hospital wage bill from CMS POS and QCEW NAICS 621+622+623, scaled by 2.1 (U.S. HCRIS hospital total-cost-to-wage-bill ratio). **Patient-health benefit.** VSL = $10 M (U.S. EPA standard, 2024 USD), rebased to 1982-84 USD as $\approx$ $3.19 M. ΔIMR = -1.47 per 1,000 live births in the high-exposure half (mean across pre-reform population density, total population, HC employment intensity, and commuting inflow; Tables A.7, OA.9–OA.11); a 0.5 multiplier rescales island-wide births to the high-exposure half.

A.5 Spillover-Excluded Sample (QCEW)

Reports the QCEW labor-market specifications on the non-border subsample ($n = 27$ counties retained out of 78). The healthcare-employment, wage, and establishment estimates are statistically indistinguishable from zero on this restricted sample; the first-stage hospital-ownership-share estimate (Online Appendix Table OA.5) survives the restriction at −12.14 percentage points ($p < 0.01$, −19.4 percent of pre-mean).

Table A.5: Spillover-Excluded Sample: QCEW Healthcare Headline Outcomes

	(1) Empl. per 1k	(2) Wages/Empl. per 1k	(3) Estabs. per 1k
ATT	-0.054	3.083	0.021
	(0.318)	(9.307)	(0.037)
Pre-treatment mean of dep. var.	3.193	42.066	0.787
ATT as % of pre-treatment mean	-1.7%	7.3%	2.7%
counties	27	27	27

Notes: Callaway and Sant'Anna (2021) doubly-robust ATT with not-yet-treated controls; reference period $t = -1$; cluster-robust standard errors at the county level via multiplier bootstrap (1,000 reps). Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. **Spillover-exclusion sample.** Restricted to the 27 counties that share no boundary with a different-cohort county.

A.6 Heterogeneity by Pre-Reform Population Density

Partitions the seventy-eight counties at the median pre-reform population density. The infant-mortality estimate in the above-median (urban) cell is −2.027 per thousand live births ($p < 0.05$, −16.9 percent of pre-mean). The first of four heterogeneity dimensions on which the IMR signal is sharper than the baseline-exposure cut reported in Table 7.

Table A.6: Heterogeneity by Pre-Reform Population Density (Urban vs. Rural Median Split)

	(1) Gov. Hosp. Birth Share (%)	(2) HC Empl. per 1k	(3) Crude Net Migr. (per 1k)	(4) IMR (per 1k live births)
Below-median pre-reform population density (urban vs. rural)				
ATT	-1.27	-0.105	-7.565	2.436
	(5.27)	(0.262)	(10.268)	(1.956)
Pre-treatment mean	67.13	2.074	0.587	13.273
ATT as % of pre-mean	-1.9%	-5.1%	-1288.3%	18.4%
counties	39	24	39	39
Above-median pre-reform population density (urban vs. rural)				
ATT	-18.21***	-0.418*	-12.513	-2.027**
	(2.36)	(0.219)	(11.042)	(0.892)
Pre-treatment mean	58.11	5.172	1.514	11.993
ATT as % of pre-mean	-31.3%	-8.1%	-826.8%	-16.9%
counties	39	35	39	39

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. The pre-treatment mean of the dependent variable is computed across treated counties using observations from the cohort-specific pre-treatment period (event time < 0). Pre-reform population density = (1990–1993 mean municipio population) / (land area in km^2). Counties partitioned at the cross-municipio median; below-median ('Rural') = 24 counties, above-median ('Urban') = 35 counties.

A.7 Per-Capita versus Raw-Level Specifications

Reports per-thousand-resident and $\log(1+y)$ raw-level specifications side-by-side for QCEW healthcare employment, establishments, and the wage bill. Healthcare employment per thousand falls by 13.0 percent ($p < 0.01$); $\log(1 + y)$ raw-level employment grows by approximately 19.0 percent ($p < 0.05$). The per-capita decline reflects population growth that outpaced the growth of the local healthcare workforce; the per-capita normalization is the welfare-relevant scaling and matches the unit conventions of the closest comparison literatures (Wallace, 2023; Hackmann et al., 2025). The distributional finding of the paper rests on the per-capita denominator and is not a feature of the absolute headcount.

Table A.7: QCEW Healthcare Headline Outcomes: Per-1,000 vs. Log-Level

	HC Employment		HC Establishments		HC Wage Bill	
	Per 1,000	Log(1+x)	Per 1,000	Log(1+x)	Per 1,000	Log(1+x)
ATT	-0.511***	0.190**	-0.050***	0.144***	0.630	0.247
	(0.157)	(0.074)	(0.019)	(0.048)	(3.734)	(0.263)
Pre-treatment mean	3.923	4.293	0.759	3.161	38.962	11.000
Implied % effect	-13.0%	19.02%	-6.6%	14.38%	1.6%	24.69%

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors.

A.8 Labor Force

Notice that the labor force contracts as the reform expands. This leads us to believe that the reform negatively affected the labor force participation rate of the population.

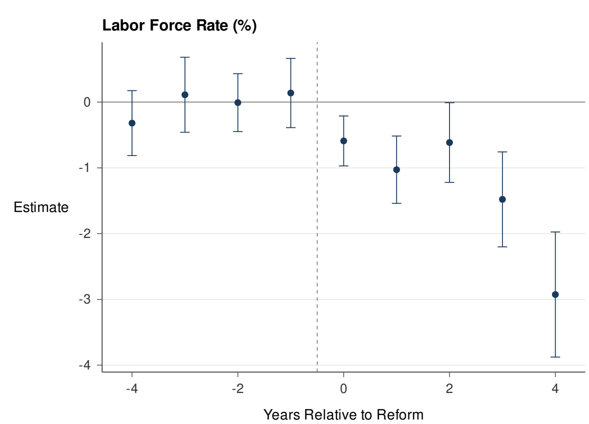

Figure A.2: Labor Force

A.9 Financing Detail

The 1993 reform replaced the Department of Health line-item appropriation that had funded the pre-reform delivery system with a multi-source aggregate funding pool flowing through ASES to the contracting insurers. Realized program financing for fiscal year 2002–2003 totaled approximately $1.282 billion, drawn in approximate proportions of seventy-three percent from the Commonwealth General Fund, fifteen percent from federal Medicaid, eleven percent from the counties, and two percent from the State Children's Health Insurance Program (SCHIP) (Pan American Health Organization, 2007). The federal Medicaid share is constrained by the Section 1108(g) territorial Medicaid cap and the fixed Federal Medical Assistance Percentage (FMAP) applicable to Puerto Rico, both of which produce a federal contribution per enrollee substantially below what Puerto Rico would receive under the income-conditioned FMAP formula applied to the fifty states (Medicaid and CHIP Payment and Access Commission, 2020; Kaiser Family Foundation, 2024). The program continues to operate under the names *Mi Salud* and *Vital*, covering approximately 1.3 million Medicaid and CHIP enrollees as of 2026 (Centers for Medicare and Medicaid Services, 2026).

A.10 Data Construction

This subsection reports the construction details for each of the six data sources summarized in Section 3.

QCEW (1990–2000). The Quarterly Census of Employment and Wages provides county-quarter establishment cells with average monthly employment, total quarterly wages, and active establishment counts by NAICS code and ownership class (federal, state, and local government, and private). The annualization procedure proceeds in two steps. Within each county-quarter cell, employment, wages, and establishments are summed across ownership classes to produce a quarter-level total. Across quarters within a county-year, employment and establishments are averaged and wages are summed; cells with fewer than three reporting quarters are dropped on the grounds that they reflect missing rather than zero healthcare activity. Per-thousand-resident outcomes are constructed by dividing each headcount by the contemporaneous Vital Statistics population in thousands. The wage-per-employee intensive-margin outcome is the annualized real wage bill divided by annualized employment, expressed in 2024 dollars using BLS CPI-U annual-average deflators. The balanced-cohort QCEW healthcare panel covers 59 counties and 649 county-year observations.

CMS Provider of Services (1991–2000). The CMS POS public-use files provide annual provider-level records of Medicare-certified facilities with full-time-equivalent counts, certified beds, services rendered, and ownership classification (government, voluntary nonprofit, proprietary). Provider-level records are aggregated to the county-year level by summing across providers operating in each county. Of 296 unique providers tracked across the 1991–2000 panel, 32 are continuously government-owned, 225 are continuously private, and 12 switched ownership within the panel window; the small switcher count underscores the partial coverage of the privatization shock at the Medicare-certified layer. The balanced CMS Track A panel covers 40 counties and 400 county-year observations.

Vital Statistics (1990–2000). The Puerto Rico Department of Health Vital Statistics annual records provide county-year totals of live births, deaths, infant deaths, and end-of-year population. These records support construction of (i) the crude birth rate, (ii) the infant mortality rate, (iii) the low-birth-weight share, (iv) the government and private hospital birth shares, (v) the crude net migration rate (CNMR), constructed as the demographic-balancing residual

$$\text{CNMR}_{i,t} \;=\; \frac{(\text{Pop}_{i,t} - \text{Pop}_{i,t-1}) - (\text{Births}_{i,t} - \text{Deaths}_{i,t})}{\text{Pop}_{i,t-1}} \times 1{,}000, \tag{3}$$

and (vi) total county population in levels and in logarithms, used as a complementary demographic outcome to (3) that is not constructed as a residual and is therefore less sensitive to compounding measurement error. CNMR estimates are reported in level units (per thousand residents per year) rather than as a percentage of the pre-treatment mean because the CNMR pre-treatment mean is close to zero (+1.06 per thousand). The CNMR is reported alongside total population because the two outcomes can in principle disagree on the sign of the local demographic response if the constituents (births, deaths, populations) are measured with component-wise error; the joint reading is informative about which margin (residual versus level) is the more reliable identification anchor in any given specification.

PR Department of Labor (1990–2000). The Puerto Rico Department of Labor and Human Resources monthly county-level labor-force survey (*Encuesta de Vivienda*) provides annualized county-year measures of the labor force, employment, and unemployment. The unemployment rate is constructed in BLS-standard form (unemployed over labor force). The labor-force-participation rate and the employment-to- population rate are constructed against total county population, since the BLS-standard 16-and-over civilian non-institutional denominator is not available at the Puerto Rican county level for the 1990–2000 panel. The maintained assumption is that the ratio of working-age civilian non-institutional population to total population is not differentially affected by treatment within the analysis window; the partial effect identified on the total-population-denominator rate is then proportional to the partial effect that would be identified on the BLS-standard rate.

NVSS Natality (1993–2000). The National Center for Health Statistics National Vital Statistics System natality public-use files provide birth-level records with infant birth weight, gestational age at delivery, the count of prenatal visits, the trimester of prenatal-care initiation, the one- and five-minute Apgar scores, and maternal demographic and educational characteristics. The empirical specification uses the conception-year treatment indicator—treated if the conception year of the recorded birth post-dates the treatment date in the county of maternal residence—rather than the delivery-year indicator, on the grounds that the relevant exposure of the prenatal-care production function is the conception-to-delivery window.

BRFSS (1996–2000). The Centers for Disease Control and Prevention Behavioral Risk Factor Surveillance System provides individual-level adult-survey microdata with self-reported insurance status, self-reported general health (a five-point ordinal scale), an indicator for whether the respondent had been unable to see a doctor in the past twelve months because of cost, and standard demographic covariates. The BRFSS Puerto Rico coverage begins in 1996, after Cohorts 1 and 2 had completed implementation; the BRFSS analysis is therefore restricted to Cohorts 3–5 (1996, 1998, and 2000), for which a pre-treatment observation window exists within the panel.

A.11 Framework Derivation

This subsection reports the formal derivations behind the three-channel decomposition (**??**) stated in Section 4.

Ownership channel. Following Hart et al. (1997), the noncontractible margin includes provider effort, time per patient, and continuity of care; the contractible margin is the per-enrollee budget that ASES pays each managed-care organization. Let $\theta \in [0,1]$ index the relative weight of quality ($\theta = 1$) versus cost ($\theta = 0$) in the operator's objective; under Hart et al. (1997), $\theta_{\text{public}} > \theta_{\text{private}}$. Letting $L^*(\theta)$ denote optimal labor input, the model implies $L^{*\prime}(\theta) > 0$, so the ownership-channel response on healthcare employment per capita satisfies

$$\Delta L^{\text{ownership}} \;=\; L^*(\theta_{\text{private}}) - L^*(\theta_{\text{public}}) \;\leq\; 0. \tag{4}$$

The same primitive admits the quality-deterioration corollary referenced in Section 4: noncontractible quality q satisfies $\Delta q \leq 0$ unless the public operator was itself underinvesting on the quality margin relative to the private optimum.

Capitation channel. The pre-reform compensation arrangement paid providers fixed government salaries; the post-reform arrangement pays managed-care organizations prospective per-member-per-month capitation, with insurers in turn contracting bilaterally with provider networks. Let $C(L)$ denote the operator's expected cost as a function of network labor input. Under salary, $C'(L)$ is borne by the principal (the Department of Health); under capitation, it is borne by the managed-care organization. As Duggan (2004) documents for U.S. Medicaid HMOs, and as Aizer et al. (2007) and Marton et al. (2014) discuss, this shift in residual cost-bearing creates an incentive to economize on input use, so the capitation margin reinforces the ownership-margin prediction in (4): $\Delta L^{\text{capitation}} \leq 0$.

Access channel. To the extent that the reform expanded effective insurance coverage among the medical-indigent population, it raised demand for insured care and, through the standard insurance-and-supply-stimulus mechanism (Arrow, 1963; Finkelstein, 2007; Dillender, 2022; Hackmann et al., 2025), predicts an expansion of the local healthcare labor market. The direction of the access-channel effect on L^* depends on whether the binding margin is demand expansion ($\Delta L^{\text{access}} > 0$) or network restriction and primary-care gatekeeping ($\Delta L^{\text{access}} \leq 0$); the sign is therefore an empirical question.

Mobility margin. The reform's labor-supply incidence is shaped by a feature distinctive to the Puerto Rican setting: the affected provider workforce holds U.S. citizenship and U.S.-credentialed professional licenses, so the legal cost of relocation to the mainland is zero. The framework therefore admits an additional adjustment margin closed off in the mainland-U.S. place-based literature (Autor et al., 2013; Yagan, 2019): out-migration of the displaced workforce. The empirical analysis traces this margin through the demographic-balancing identity (3), jointly with the level of county population. The two demographic measures are reported together because each is informative about a distinct aspect of the local labor supply, and disagreement between them carries identification content rather than being averaged away.

A.12 Empirical Strategy: Technical Detail

This subsection reports the dynamic-aggregation formula and the TWFE comparator specification referenced in Section 5.

Dynamic event-study aggregation. The event-study figures aggregate the group-time treatment effects $\widehat{\mathrm{ATT}}(g,t)$ from (1) into event-time coefficients $\widehat{\theta}_\ell$ by cohort-size weights:

$$\widehat{\theta}_\ell \;=\; \sum_{g\in\mathcal{G}} \frac{N_g}{\sum_{g'} N_{g'}}\,\widehat{\mathrm{ATT}}(g,g+\ell), \qquad \ell\in\{-4,\dots,+4\}, \tag{5}$$

where N_g is the number of counties in cohort g and the reference period is $\ell=-1$. Confidence bands are constructed via the multiplier bootstrap of Callaway and Sant'Anna (2021) clustered at the county level with 1,000 replications, applied separately at each event time.

TWFE comparator specification. The canonical two-way fixed-effects (TWFE) event-study reported as estimator-class robustness in Appendix 10 regresses $Y_{i,t}$ on county fixed effects α_i, year fixed effects λ_t, and event-time indicators relative to treatment:

$$Y_{i,t} \;=\; \alpha_i + \lambda_t + \sum_{\ell\neq -1} \beta_\ell\,\mathbf{1}\{t-G_i=\ell\} \;+\; \varepsilon_{i,t}, \tag{6}$$

with standard errors clustered at the county level and the reference period $\ell=-1$ omitted.

Doubly robust three-step procedure. The $\widehat{\mathrm{ATT}}(g,t)$ estimator in (1) is the doubly robust difference-in-differences estimator of Callaway and Sant'Anna (2021), which combines an inverse-probability-weighted estimator on the not-yet-treated comparison group with an outcome-regression-based correction term. The estimator is consistent if either the propensity-score model or the outcome-regression model is correctly specified, providing a robustness margin over either component alone. Implementation uses the `did` package in R with the default propensity-score specification (logistic regression on a constant) and the default outcome-regression specification (county and year fixed effects); the group-time-by-event-time estimates are then aggregated via (5) for the event-study figures and via the post-treatment-cell-count-weighted average for the regression tables. The not-yet-treated control group is implemented as `control_group = "notyettreated"` in `did::att_gt`.

A.13 BRFSS Coverage by Subgroup

This subsection expands on the BRFSS access-channel result reported in Section 6.3. The BRFSS 1996–2000 panel restricts the analysis to Cohorts 3–5, for which a pre-treatment observation window exists within the panel (Appendix 10). The single-cohort exposure pattern within the BRFSS window means that Callaway and Sant'Anna (2021) is not implementable on this sample; the BRFSS results are reported under the Wooldridge (2025) ETWFE estimator with TWFE comparator estimates also reported. Six demographic subgroups are analyzed (Online Appendix Tables OA.14–OA.19): ages 18–29, 30–44, 45–64, and 65 and older; college graduates; and employed respondents. Coverage rises across all six subgroups, with the largest point estimate of +10.5 percentage points in the 45–64 age group ($p < 0.01$) and the smallest of +2.8 percentage points among respondents aged 65 and older ($p < 0.01$, smaller in magnitude because near-universal Medicare coverage at this age leaves limited room for the reform to expand effective insurance). The six subgroup estimates jointly rule out the no-coverage-change null even with the limited cohort exposure in the BRFSS sample, and the cross-subgroup pattern is consistent with the reform expanding effective insurance among the working-age medical-indigent population most directly targeted by the reform's eligibility criteria.

Online Appendix

The Online Appendix collects the additional sample-restriction, estimator-comparison, and subgroup specifications referenced in the main text. The exhibits are organized by topic, in the order in which they are referenced in the body.

OA.1 Macro Labor-Market Outcomes

Reports the BLS-standard unemployment rate, the labor-force-participation rate, and the employment-to-population rate side by side. The unemployment rate is statistically indistinguishable from zero against a pre-reform mean of 17.4 percent; the LFP and employment-to-population rates both fall by 3.5 percent of pre-mean. The unemployment-rate null together with the LFP decline is the labor-supply-margin signature of Section 6.2.

Table OA.1: Macro Labor-Market Effects of the reform (Online Appendix)

	(1) Unemployment Rate (%)	(2) Employment Rate (%)	(3) Labor Force Rate (%)
ATT	-0.245	-0.926***	-1.110***
	(0.285)	(0.219)	(0.241)
Pre-treatment mean of dep. var.	17.770	26.331	31.895
ATT as % of pre-treatment mean	-1.4%	-3.5%	-3.5%
counties	78	78	78

Notes: Estimates use the Callaway and Sant'Anna (2021) ATT. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, based on two-sided tests with cluster-robust standard errors.

OA.2 Log-Specification on QCEW Main

Parallel specification on $\log(1+y)$ per-1,000 normalizations and $\log(1+y)$ raw-level outcomes for the QCEW healthcare aggregate. The per-capita and raw-level estimates point in opposite directions because the local population grew faster than the local healthcare workforce during the analysis window; the main finding rests on the per-capita normalization.

Table OA.2: Log(1+x) Specification: QCEW Healthcare Headline Outcomes (Online Appendix)

	Per-1k Population (log1p)			Levels (log1p)		
	(1) Empl./1k	(2) Wages/Empl./1k	(3) Estabs./1k	(4) Empl.	(5) Wages	(6) Estabs.
ATT	0.0816***	-0.0180	0.0176*	0.1902**	0.2469	0.1438***
	(0.0274)	(0.1137)	(0.0102)	(0.0743)	(0.2635)	(0.0482)
Pre-treatment mean of dep. var.	1.2007	3.2910	0.5134	4.2928	11.0005	3.1613
ATT as % of pre-treatment mean	6.8%	-0.5%	3.4%	4.4%	2.2%	4.5%
counties	59	59	59	59	59	59

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, based on two-sided tests with cluster-robust standard errors.

OA.3 Excluding San Juan: First-Stage and Patient Health

Reports the side-by-side balanced and excluded-San-Juan estimates for the government-hospital birth share, the infant mortality rate, and the low-birth-weight share. The first-stage ownership-share estimate moves from -14.79 to -8.22 percentage points (both $p < 0.01$); the patient-health outcomes are robust to the exclusion.

Table OA.3: Excluding San Juan: Patient-Health and Birth-Share Robustness (Online Appendix)

	Gov. Hosp. Birth Share		IMR		Low BW Rate	
	(1) Bal.	(2) Excl. SJ	(3) Bal.	(4) Excl. SJ	(5) Bal.	(6) Excl. SJ
ATT	-14.79***	-8.22***	-1.205	0.781	0.060	-0.307
	(1.94)	(2.85)	(0.928)	(1.237)	(0.230)	(0.368)
Pre-treatment mean of dep. var.	62.55	62.61	12.620	12.624	9.717	9.712
ATT as % of pre-treatment mean	-23.6%	-13.1%	-9.5%	6.2%	0.6%	-3.2%
counties	78	77	78	77	78	77

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator with not-yet-treated comparison units. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, based on two-sided tests with cluster-robust standard errors.

OA.4 Excluding San Juan: CMS Provider-Stratified

Reports excluded-San-Juan robustness for CMS hospital-bed and services-per-thousand outcomes, stratified by ownership class.

Table OA.4: Excluding San Juan: CMS Provider-Structure Robustness (Online Appendix)

	Beds per 1k pop.		Services per 1k pop.		FTE per 1k pop.	
	(1) Bal.	(2) Excl. SJ	(3) Bal.	(4) Excl. SJ	(5) Bal.	(6) Excl. SJ
ATT	-0.265***	0.090	-0.024**	-0.014	-2.251**	-1.815
	(0.065)	(0.140)	(0.010)	(0.012)	(0.971)	(1.380)
Pre-treatment mean of dep. var.	3.389	3.033	0.487	0.463	10.828	10.288
ATT as % of pre-treatment mean	-7.8%	3.0%	-4.9%	-3.0%	-20.8%	-17.6%
counties	40	39	40	39	33	32

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator with not-yet-treated comparison units. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, based on two-sided tests with cluster-robust standard errors.

OA.5 Spillover-Excluded Sample: First Stage and Patient Health

The first-stage ownership-share estimate on the non-border subsample is -12.14 percentage points ($p < 0.01$, -19.4 percent of pre-mean), economically indistinguishable from the main result; the corresponding patient-health estimates are statistically indistinguishable from zero on this restricted panel.

Table OA.5: Spillover-Excluded Sample: Hospital Birth Share and Patient-Health Outcomes

	Hospital Birth Share (VS)		Patient Health (VS)	
	(1) Gov. Hosp. (%)	(2) Priv. Hosp. (%)	(3) IMR (per 1k)	(4) Low BW (%)
ATT	-12.14***	12.29***	-0.072	0.087
	(2.68)	(2.44)	(1.508)	(0.578)
Pre-treatment mean of dep. var.	62.48	37.00	12.534	9.524
ATT as % of pre-treatment mean	-19.4%	33.2%	-0.6%	0.9%
counties	39	39	39	39

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, based on two-sided tests with cluster-robust standard errors.

OA.6 Spillover-Excluded Sample: CMS Provider Capacity

CMS hospital-bed and services estimates on the spatially-restricted, SUTVA-tightened sample.

Table OA.6: Spillover-Excluded Sample: CMS Provider-Structure Outcomes

	(1) Beds per 1k pop.	(2) Services per 1k pop.	(3) FTE per 1k pop.
ATT	-0.003	-0.020	-3.278
	(0.061)	(0.023)	(2.282)
Pre-treatment mean of dep. var.	3.908	0.657	14.272
ATT as % of pre-treatment mean	-0.1%	-3.1%	-23.0%
N counties	19	19	16

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors.

OA.7 Low-Commuting-Inflow Sample: First Stage and Patient Health

Sample restricted to counties with workplace in-commuting share at or below 0.50 ($n = 64$). The first-stage ownership-share estimate is -7.18 percentage points ($p < 0.05$); the patient-health estimates are robust. Therefore, if we assume that in low-commute counties treatment-effect contamination through cross-county worker flows is minimal, then may assume that our estimates are unbiased, since the effects on the patient-health margin remain consistent with the main estimates.

Table OA.7: Low-Commuting-Inflow Sample: Hospital Birth Share and Patient Health

	(1) Gov. Hosp. Birth Share (%)	(2) IMR (per 1k live births)
ATT	-7.18**	1.327
	(3.29)	(1.392)
Pre-treatment mean of dep. var.	64.11	12.846
ATT as % of pre-treatment mean	-11.2%	10.3%
N counties	64	64

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. Sample restricted to counties with *low* workplace in-commuting share (in_share ≤ 0.50, where in_share = workers commuting in from other counties divided by total workplace count). The classifier is built from Census 2000 Journey-to-Work tabulations and read from the diagnostic CSV that all_event_studies.R writes at startup.

OA.8 Low-Commuting-Inflow Sample: CMS Provider Capacity

CMS hospital-bed and services estimates on the low-commuting-inflow subsample. CMS supply-side responses in low-commute counties are consistent with main estimates.

Table OA.8: Low-Commuting-Inflow Sample: CMS Provider-Structure Outcomes

	(1) Beds per 1k pop.	(2) Services per 1k pop.	(3) FTE per 1k pop.
ATT	0.087	-0.007	-2.113
	(0.174)	(0.011)	(1.567)
Pre-treatment mean of dep. var.	2.746	0.415	9.730
ATT as % of pre-treatment mean	3.2%	-1.7%	-21.7%
N counties	34	34	27

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$, based on two-sided tests with cluster-robust standard errors.

OA.9 Heterogeneity by Pre-Reform Population

The infant-mortality estimate in the above-median (larger) cell is -2.564 per thousand live births ($p < 0.01$, -20.8 percent of pre-mean). Second of the four heterogeneity dimensions on which the IMR signal is sharper than the baseline-exposure cut.

Table OA.9: Heterogeneity by Pre-Reform Population (1990–1993 mean) (Online Appendix)

	(1) Gov. Hosp. Birth Share (%)	(2) HC Empl. per 1k	(3) Crude Net Migr. (per 1k)	(4) IMR (per 1k live births)
Below-median pre-reform population				
ATT	-0.37	-0.154	-1.232	3.657*
	(5.78)	(0.221)	(11.833)	(2.081)
Pre-treatment mean	66.04	1.496	0.050	12.923
ATT as % of pre-mean	-0.6%	-10.3%	-2464.3%	28.3%
N counties	39	20	39	39
Above-median pre-reform population				
ATT	-17.38***	-0.300	-14.651	-2.564***
	(2.27)	(0.228)	(11.447)	(0.921)
Pre-treatment mean	59.33	5.085	1.975	12.347
ATT as % of pre-mean	-29.3%	-5.9%	-741.7%	-20.8%
N counties	39	39	39	39

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$, based on two-sided tests with cluster-robust standard errors. The classifier is computed once over the strict pre-rollout window (BASELINE_YEARS = 1990–1993) and is therefore time-invariant within county. Counties without a valid baseline value are routed to the below-median bin per the standard censoring rule. *Below-median* = 39 counties; *Above-median* = 39 counties.

OA.10 Heterogeneity by Pre-Reform Healthcare Employment

The infant-mortality estimate in the above-median (higher-pre-reform-healthcare-intensity) cell is -2.633 per thousand live births ($p < 0.01$, -20.7 percent of pre-mean). Third of the four heterogeneity dimensions on which the IMR signal is sharper than the baseline-exposure cut.

Table OA.10: Heterogeneity by Pre-Reform HC Employment (1990–1993 mean)

	(1) Gov. Hosp. Birth Share (%)	(2) HC Empl. per 1k	(3) Crude Net Migr. (per 1k)	(4) IMR (per 1k live births)
Below-median pre-reform HC employment				
ATT	-4.69	-0.084	-8.171	2.503
	(3.68)	(0.196)	(11.155)	(1.634)
Pre-treatment mean	64.79	1.136	1.440	12.572
ATT as % of pre-mean	-7.2%	-7.4%	-567.5%	19.9%
N counties	48	29	48	48
Above-median pre-reform HC employment				
ATT	-17.87***	-0.433	-7.758	-2.633***
	(2.54)	(0.268)	(13.732)	(0.830)
Pre-treatment mean	58.85	6.816	0.428	12.698
ATT as % of pre-mean	-30.4%	-6.4%	-1810.6%	-20.7%
N counties	30	30	30	30

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, based on two-sided tests with cluster-robust standard errors. Counties are partitioned at the cross-county median of the 1990–1993 mean QCEW Healthcare (NAICS 621+622+623) employment level (annualized from quarterly cells with ≥3 valid quarters) (median = 58). The classifier is computed once over the strict pre-rollout window (BASELINE_YEARS = 1990–1993) and is therefore time-invariant within county. Counties without a valid baseline value are routed to the below-median bin. *Below-median* = 48 counties; *Above-median* = 30 counties.

OA.11 Heterogeneity by Workplace Commuting Inflow

The infant-mortality estimate in the above-median (workplace-center) cell is -3.124 per thousand live births ($p < 0.01$, -26.8 percent of pre-mean). Fourth of the four heterogeneity dimensions on which the IMR signal is sharper than the baseline-exposure cut.

Table OA.11: Heterogeneity by Workplace Commuting Inflow (Low vs. High Inflow)

	(1) Gov. Hosp. Birth Share (%)	(2) HC Empl. per 1k	(3) Crude Net Migr. (per 1k)	(4) IMR (per 1k live births)
Low-inflow workplace counties				
ATT	-7.18**	-0.356	-4.777	1.327
	(3.29)	(0.218)	(9.854)	(1.392)
Pre-treatment mean	64.11	2.910	1.916	12.846
ATT as % of pre-mean	-11.2%	-12.3%	-249.4%	10.3%
N counties	64	45	64	64
High-inflow workplace counties				
ATT	-16.55***	-0.021	7.774	-3.124***
	(3.02)	(0.357)	(11.287)	(1.114)
Pre-treatment mean	55.99	6.974	-2.495	11.671
ATT as % of pre-mean	-29.6%	-0.3%	-311.6%	-26.8%
N counties	14	14	14	14

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, based on two-sided tests with cluster-robust standard errors. Workplace in-commuting share from Census 2000 Journey-to-Work tabulations: in_share = (in-commuters from other counties) / (total workplace count). Threshold cutoff at 0.50 (NOT median). *Low-inflow* = 45 counties (in_share ≤ 0.50); *High-inflow* = 14 counties (in_share > 0.50). **Note.** The 'Below-median' and 'Above-median' panels here label the low-inflow and high-inflow groups respectively. **Reading.** High-inflow workplace counties import labor-force flows from neighbors; if treatment effects differ across the inflow margin, that's evidence of cross-county worker spillovers acting on the labor-market response.

OA.12 Per-Thousand-Resident Placebo Industries

C&S estimates on three non-healthcare placebo industries: NAICS 11 (Agriculture, Forestry, and Fishing), NAICS 48–49 (Transportation and Warehousing), and NAICS 71 (Arts, Entertainment, and Recreation). NAICS 31–33 manufacturing is excluded *a priori* as endogenous to the contemporaneous Section 936 corporate-tax phase-out.

Table OA.12: QCEW Placebo Industries: Effects of the reform on Industries Outside the Healthcare Sector

Industry	(1) Empl. per 1k	(2) Wages/Empl per 1k	(3) Wages per 1k	(4) Estabs per 1k
Agriculture, Forestry & Fishing (NAICS 11)				
ATT	0.108	-6.126***	-1,476.474**	0.106
	(0.369)	(1.591)	(741.852)	(0.075)
Pre-treatment mean	12.463	30.797	–	3.463
ATT as % of pre-mean	0.9%	-19.9%	–	3.1%
counties	25	25		25
Transportation & Warehousing (NAICS 48-49)				
ATT	-0.709***	7.571	-688.266***	0.001
	(0.046)	(9.813)	(243.595)	(0.003)
Pre-treatment mean	1.475	240.942	–	0.105
ATT as % of pre-mean	-48.1%	3.1%	–	1.2%
counties	77	77		77
Arts, Entertainment & Recreation (NAICS 71)				
ATT	-0.498	1.994	-553.871	-0.032**
	(0.350)	(1.329)	(558.662)	(0.015)
Pre-treatment mean	2.107	15.207	–	0.155
ATT as % of pre-mean	-23.6%	13.1%	–	-20.6%
counties	9	9		9

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$, based on two-sided tests with cluster-robust standard errors. **Sample.** QCEW (BLS) county-quarter panel restricted to one of three placebo industries: NAICS 11 (Agriculture, Forestry, Fishing & Hunting); NAICS 48-49 (Transportation & Warehousing); NAICS 71 (Arts, Entertainment & Recreation).

OA.13 Raw-Level Placebo Industries

Raw-level (rather than per-thousand-resident) specifications on the same three placebo industries. Agriculture is null in raw levels, supporting the per-capita-redistribution interpretation of the per-capita Agriculture-wage-bill decline reported in Section 7. Transportation contracts in raw levels by 94.0 percent, consistent with input-output linkage to the phasing-out manufacturing logistics chain. Arts is statistically underpowered. For context, a per-1k placebo null could in principle hide an offsetting population-denominator shock that the per-1k normalization absorbs. Raw-level placebos provide a cleaner test for absolute employment-flow effects.

Table OA.13: QCEW Placebo Industries (Raw Levels): Employment, Wage Bill, and Establishment Counts

Industry	(1) Employment	(2) Wage Bill	(3) Establishments
Agriculture, Forestry & Fishing (NAICS 11)			
ATT	14.5	12,717	0.3
	(13.8)	(20,913)	(2.0)
Pre-treatment mean	293.0	193,903	85.5
ATT as % of pre-mean	4.9%	6.6%	0.3%
N counties	25	25	25
Transportation & Warehousing (NAICS 48-49)			
ATT	-191.9***	-164,320**	1.2***
	(11.8)	(83,737)	(0.3)
Pre-treatment mean	204.1	791,798	7.4
ATT as % of pre-mean	-94.0%	-20.8%	15.9%
N counties	77	77	77
Arts, Entertainment & Recreation (NAICS 71)			
ATT	-63.7	-192,039*	-14.4***
	(51.5)	(106,043)	(3.8)
Pre-treatment mean	309.2	568,279	30.1
ATT as % of pre-mean	-20.6%	-33.8%	-47.7%
N counties	9	9	9

Notes: Three estimators are reported. TWFE: two-way fixed-effects regression with county and year FEs and event-time dummies; the overall ATT is the simple mean of the post-treatment dummy coefficients with delta-method standard error on the post-period. ETWFE: Wooldridge (2021) Extended Two-Way Fixed Effects, with overall ATT as a cohort-size-weighted average. C&S: Callaway and Sant'Anna (2021) doubly-robust ATT, aggregated to a single overall ATT. All three estimators use cluster-robust standard errors at the county level; C&S uses the multiplier bootstrap (1000 replications). Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. Raw-level outcomes: Employment is the average quarterly headcount; Wage Bill is the average quarterly real wage bill (deflated; see Section 12b for the deflator); Establishments is the average quarterly establishment count.

OA.14 BRFSS Insurance Coverage by Subgroup

BRFSS subgroup estimates of the ETWFE point estimate on the "has health plan" indicator. The coverage-expansion first stage is positive and significant in five of the six reported subgroups; the 30–44 age group is the one in which the estimate is positive but does not reach conventional significance. The pattern is consistent with the BRFSS analysis covering Cohorts 3–5 (the cohorts for which a pre-treatment BRFSS observation window is available), documented in Appendix 10.

Table OA.14: BRFSS Healthcare Access — Age 18–29

	(1) Has Health Plan	(2) Preventive Services	(3) Could Not Afford Doctor	(4) General Health (1-5)
TWFE (weighted)	-0.0339 (0.0253)	-0.215*** (0.033)	-0.0109 (0.0152)	0.256*** (0.049)
ETWFE	0.0783*** (0.0232)	-0.109* (0.056)	-0.0118 (0.0245)	-0.001 (0.052)

Notes: TWFE estimated with BRFSS survey weights (`weights_var = 'survey_weight'`); ETWFE is unweighted. Cells report ATT (SE in parentheses). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Callaway-Sant'Anna is not run for BRFSS because the dataset's 1996–2000 window contains only one fully-treated reform cohort (1998); the 2000 cohort has no valid not-yet-treated controls in the post-period. Sample: BRFSS respondents aged 18–29. Aggregated from `all_brfss_coefs.csv` produced by all_event_studies.R; same numerical conventions as Online Appendix Figure OA.15.

Table OA.15: BRFSS Healthcare Access — Age 30–44

	(1) Has Health Plan	(2) Preventive Services	(3) Could Not Afford Doctor	(4) General Health (1-5)
TWFE (weighted)	-0.0053 (0.0145)	0.094*** (0.026)	0.0190** (0.0082)	-0.040 (0.035)
ETWFE	0.0301 (0.0186)	0.025 (0.029)	0.0147 (0.0174)	-0.107** (0.044)

Notes: TWFE estimated with BRFSS survey weights (`weights_var = 'survey_weight'`); ETWFE is unweighted. Cells report ATT (SE in parentheses). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Callaway-Sant'Anna is not run for BRFSS because the dataset's 1996–2000 window contains only one fully-treated reform cohort (1998); the 2000 cohort has no valid not-yet-treated controls in the post-period. Sample: BRFSS respondents aged 30–44. Aggregated from `all_brfss_coefs.csv` produced by all_event_studies.R; same numerical conventions as Online Appendix Figure OA.16.

Table OA.16: BRFSS Healthcare Access — Age 45–64

	(1) Has Health Plan	(2) Preventive Services	(3) Could Not Afford Doctor	(4) General Health (1-5)
TWFE (weighted)	0.0891*** (0.0084)	0.133*** (0.033)	0.0117 (0.0101)	0.017 (0.040)
ETWFE	0.1052*** (0.0138)	0.100*** (0.032)	-0.0102 (0.0120)	0.102 (0.073)

Notes: TWFE estimated with BRFSS survey weights (`weights_var = 'survey_weight'`); ETWFE is unweighted. Cells report ATT (SE in parentheses). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. Callaway-Sant'Anna is not run for BRFSS because the dataset's 1996–2000 window contains only one fully-treated reform cohort (1998); the 2000 cohort has no valid not-yet-treated controls in the post-period. Sample: BRFSS respondents aged 45–64. Aggregated from `all_brfss_coefs.csv` produced by all_event_studies.R; same numerical conventions as Online Appendix Figure OA.17.

OA.16 Population in Levels and Logarithms

Reports the C&S point estimate on total county population in level and log specifications. The level estimate is +5.1 percent of pre-mean ($p < 0.01$); the log specification is consistent in sign and magnitude. The population growth estimate is the level outcome that complements the demographic-balancing residual (CNMR) reported in Table 6.

Table OA.17: BRFSS Healthcare Access — Age 65 and Older

	(1) Has Health Plan	(2) Preventive Services	(3) Could Not Afford Doctor	(4) General Health (1-5)
TWFE (weighted)	0.0393*** (0.0049)	0.037*** (0.012)	0.0001 (0.0127)	0.138*** (0.033)
ETWFE	0.0280*** (0.0087)	-0.023 (0.014)	-0.0101 (0.0144)	0.211*** (0.052)

Notes: TWFE estimated with BRFSS survey weights (`weights_var = 'survey_weight'`); ETWFE is unweighted. Cells report ATT (SE in parentheses). Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. Callaway-Sant'Anna is not run for BRFSS because the dataset's 1996–2000 window contains only one fully-treated reform cohort (1998); the 2000 cohort has no valid not-yet-treated controls in the post-period. Sample: BRFSS respondents aged 65+ (Medicare-eligible). La Reforma did not directly target this group; near-null on health-plan coverage is the expected no-effect benchmark. Aggregated from `all_brfss_coefs.csv` produced by all_event_studies.R; same numerical conventions as Online Appendix Figure OA.18.

Table OA.18: BRFSS Healthcare Access — College Graduates

	(1) Has Health Plan	(2) Preventive Services	(3) Could Not Afford Doctor	(4) General Health (1-5)
TWFE (weighted)	-0.0375*** (0.0104)	0.085*** (0.026)	0.0524*** (0.0103)	0.111*** (0.038)
ETWFE	0.0810*** (0.0212)	0.059* (0.032)	0.0305 (0.0209)	0.201*** (0.059)

Notes: TWFE estimated with BRFSS survey weights (`weights_var = 'survey_weight'`); ETWFE is unweighted. Cells report ATT (SE in parentheses). Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. Callaway-Sant'Anna is not run for BRFSS because the dataset's 1996–2000 window contains only one fully-treated reform cohort (1998); the 2000 cohort has no valid not-yet-treated controls in the post-period. Sample: BRFSS respondents with college degree. Aggregated from `all_brfss_coefs.csv` produced by all_event_studies.R; same numerical conventions as Online Appendix Figure OA.20.

Table OA.19: BRFSS Healthcare Access — Employed Respondents

	(1) Has Health Plan	(2) Preventive Services	(3) Could Not Afford Doctor	(4) General Health (1-5)
TWFE (weighted)	-0.0044 (0.0111)	0.029 (15,216.475)	0.0120 (0.0077)	-0.019 (0.034)
ETWFE	0.0460*** (0.0102)	0.016 (0.029)	-0.0233** (0.0104)	-0.144*** (0.037)

Notes: TWFE estimated with BRFSS survey weights (`weights_var = 'survey_weight'`); ETWFE is unweighted. Cells report ATT (SE in parentheses). Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. Callaway-Sant'Anna is not run for BRFSS because the dataset's 1996–2000 window contains only one fully-treated reform cohort (1998); the 2000 cohort has no valid not-yet-treated controls in the post-period. Sample: BRFSS respondents reporting employment. Aggregated from `all_brfss_coefs.csv` produced by all_event_studies.R; same numerical conventions as Online Appendix Figure OA.19.

Table OA.20: Effects of The Reform on Log County Population

	(1) Population (levels)	(2) Log Population
ATT	2,528***	0.0357***
	(690)	(0.0086)
Pre-treatment mean of dep. var.	49,898	10.4172
ATT as % of pre-treatment mean	5.1%	0.3%
N counties	78	78

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. Total county population (Col. 1) and its natural logarithm (Col. 2).

OA.17 Demographic Reconciliation

Reports the explicit demographic-balancing identity check between the population-level C&S point estimate and the sum of the implied contributions from natural increase (births minus deaths) and net migration. The residual of approximately 66 per thousand reflects (i) measurement error in the CNMR, which is constructed as a residual of three measured flows, and (ii) C&S aggregation weights that differ across panels with different observation counts (the CNMR loses each county's first observation because Pop_{t-1} is unavailable). The two demographic margins are reported jointly because the disagreement carries identification content.

Table OA.21: Demographic Margins of Adjustment: Reconciling Levels and Flows

	Population (Levels)		Net Migration	Vital Rates (Flows)	
	(1) Population	(2) Log Population	(3) CNMR (per 1k)	(4) Crude Birth Rate (per 1k)	(5) Crude Death Rate (per 1k)
ATT	2,528***	0.0357***	-14.651**	-0.564**	0.047
	(690)	(0.0086)	(7.389)	(0.225)	(0.114)
Pre-treatment mean of dep. var.	49,898	10.4172	1.059	18.267	7.293
ATT as % of pre-treatment mean	5.1%	0.3%	-1382.9%	-3.1%	0.6%
N counties	78	78	78	78	78

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. The demographic identity for any closed area between two periods is $\Delta\text{Population}_{t,t-1} = (\text{Births}_t - \text{Deaths}_t) + \text{Net Migration}_t$, and dividing by average population yields the per-1k version: $\Delta \log P \approx \text{CBR} - \text{CDR} + \text{CNMR}$. Columns (3)–(5) report the C&S ATT for each of the three flow components on the right-hand side; columns (1)–(2) report the left-hand side in absolute and proportional units.

OA.18 NVSS Individual-Level Birth-Outcome Robustness

Reports C&S estimates on individual-level NVSS natality records using the conception-year-treatment indicator. Outcomes: birth weight, the count of prenatal visits, the trimester of prenatal-care initiation, and the five-minute Apgar score. The point estimates are statistically indistinguishable from zero, consistent with the county-aggregate Vital Statistics null reported in Table 5.

Table OA.22: NVSS Individual-Level Birth Outcomes — Effects of The Reform

	(1) Birth Weight (g)	(2) Prenatal Visits	(3) Apgar (5-min)	(4) Pr(LBW)	(5) Pr(VLBW)
ATT (C&S)	-0.75	-0.18	0.002	0.0015	-0.0001
	(9.99)	(0.17)	(0.020)	(0.0036)	(0.0013)

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator. Cluster-robust standard errors at the county (municipio) level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. The variables are continuous birth-weight, prenatal-visits, and Apgar measures plus binary low/very-low-birth-weight indicators.

OA.19 Heterogeneity by Pre-Reform Poverty Rate

Partitions the seventy-eight counties at the median Census 2000 poverty rate. The healthcare-employment estimate in the above-median-poverty cell is -31.2 percent of pre-mean ($p < 0.10$); the result is directionally consistent with concentration of the contraction in slack-labor-market cells. Sample-split power limits formal inference on the access-channel-versus-ownership-channel decomposition contemplated by the framework of Section 4.

Table OA.23: Heterogeneity by Census 2000 Poverty Rate

	(1) Gov. Hosp. Birth Share (%)	(2) HC Empl. per 1k	(3) Crude Net Migr. (per 1k)	(4) IMR (per 1k live births)
Below-median poverty rate				
ATT	-17.81***	-0.295	-15.773	-1.622*
	(1.93)	(0.214)	(11.426)	(0.984)
Pre-treatment mean	57.92	5.481	2.166	11.970
ATT as % of pre-mean	-30.7%	-5.4%	-728.1%	-13.5%
N counties	39	34	39	39
Above-median poverty rate				
ATT	-2.49	-0.557*	1.249	2.066
	(4.76)	(0.290)	(11.907)	(1.972)
Pre-treatment mean	67.14	1.787	-0.037	13.265
ATT as % of pre-mean	-3.7%	-31.2%	-3378.7%	15.6%
N counties	39	25	39	39

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. Counties are partitioned at the cross-county median of the Census 2000 % of residents below the federal poverty line (median = 54.55%). The classifier is computed once over the strict pre-rollout window (baseline years = 1990–1993) and is therefore time-invariant within county. Counties without a valid baseline value are routed to the below-median bin per the standard censoring rule. *Below-median* = 39 counties; *Above-median* = 39 counties. Census 2000 is partially post-rollout for cohorts 1–4 (1994–1998) and concurrent with cohort 5 (2000), so the classifier is partially endogenous to the reform itself; yet we retain it because no comparable pre-reform county-level poverty series exists.

OA.20 Heterogeneity by Pre-Reform Unemployment Rate

Partitions the seventy-eight counties at the median pre-reform unemployment rate. The healthcare-employment estimate in the above-median-unemployment cell is -17.8 percent of pre-mean (statistically insignificant); the directional pattern is consistent with the poverty-rate cut.

Table OA.24: Heterogeneity by Pre-Reform Unemployment Rate (BLS, 1990–1993 mean)

	(1) Gov. Hosp. Birth Share (%)	(2) HC Empl. per 1k	(3) Crude Net Migr. (per 1k)	(4) IMR (per 1k live births)
Below-median pre-reform unemployment rate				
ATT	-14.23***	-0.331	-5.217	-2.617**
	(2.24)	(0.203)	(10.675)	(1.146)
Pre-treatment mean	58.72	5.027	0.059	12.332
ATT as % of pre-mean	-24.2%	-6.6%	-8888.1%	-21.2%
N counties	39	32	39	39
Above-median pre-reform unemployment rate				
ATT	-7.37***	-0.454	-11.417	3.213
	(2.13)	(0.358)	(11.903)	(2.294)
Pre-treatment mean	66.50	2.550	2.100	12.912
ATT as % of pre-mean	-11.1%	-17.8%	-543.7%	24.9%
N counties	39	27	39	39

Notes: Estimates use the Callaway and Sant'Anna (2021) doubly-robust ATT estimator. Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, based on two-sided tests with cluster-robust standard errors. Counties are partitioned at the cross-county median of the 1990–1993 mean PR DOL BLS-standard unemployment rate (unemployed / labor force $\times$ 100) (median = 18.62%). The classifier is computed once over the strict pre-rollout window (baseline years = 1990–1993) and is therefore time-invariant within county. counties without a valid baseline value are routed to the below-median bin. *Below-median* = 39 counties; *Above-median* = 39 counties.

OA.21 Adult Cause-Specific Mortality

Reports C&S estimates on adult cause-specific mortality outcomes (heart disease, cerebrovascular disease, diabetes, accidents, homicide). Several causes of death - such as cardiovascular disease, cerebrovascular disease, diabetes, and pneumonia or influenza- have biologically plausible response horizons that extend well beyond our four-year post-reform observation window. Conversely, acute causes like accidents and homicides likely reflect the true, unbiased mortality rate; because these events demand immediate attention from local providers, they are less susceptible to the confounding spillover effects of healthcare access in neighboring counties. Similarly, maternal mortality operates on a longer time horizon. Preventing maternal death depends heavily on long-term, systemic healthcare managed over decades, and is increasingly confounded by demographic shifts, such as women entering pregnancy later in life. Consequently, we select the Infant Mortality Rate (IMR) as our primary healthcare indicator. Unlike maternal health, saving a newborn relies predominantly on immediate, high-tech medical interventions, making IMR highly sensitive to short-term, localized healthcare quality. The remaining mortality estimates are reported solely for completeness and should not be interpreted as identified treatment effects.

Table OA.25: Cause-Specific Mortality Effects of The Reform

	(1) Heart	(2) Cerebro.	(3) Diabetes	(4) Pneum./Flu	(5) Perinatal	(6) Maternal	(7) Accidents	(8) Homicide
ATT	-0.096**	-0.041**	-0.010	0.009	–	0.002*	0.138	0.107***
	(0.046)	(0.020)	(0.034)	(0.021)		(0.001)	(0.101)	(0.015)
Pre-treatment mean of dep. var.	1.572	0.370	0.524	0.298	–	0.003	0.344	0.143
ATT as % of pre-treatment mean	-6.1%	-11.0%	-2.0%	3.1%	–	73.1%	40.1%	75.0%
N counties	78	78	78	78	–	78	78	78

Notes: Estimates use the Callaway and Sant'Anna (2021). Cluster-robust standard errors at the county level are obtained via the multiplier bootstrap with 1000 replications. Significance levels: $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$, based on two-sided tests with cluster-robust standard errors. **Outcomes.** Cause-specific death rates per 1,000 residents. **Caveat.**