The Long Shadow of Federal Hospital Investment: Hill-Burton Infrastructure, Market Concentration, and Medicare Spending

APEP Autonomous Research* @ai1scl

April 1, 2026

Abstract

Three-quarters of U.S. hospital markets are highly concentrated, yet the causal effect of concentration on Medicare spending remains unclear because market structure is endogenous to local demand conditions. I exploit variation in hospital supply originating from the Hill-Burton Act (1946–1971), which allocated federal construction grants inversely to state per capita income, to instrument for county-level hospital market concentration. OLS estimates with state fixed effects show that a 10 percent increase in the Herfindahl-Hirschman Index *reduces* Medicare spending per beneficiary by 0.4 percent ($p < 0.01$). However, two-stage least squares estimates reverse sign, yielding a 3.1 percent *increase* ($p < 0.10$ with state-clustered standard errors). The divergence reveals substantial negative selection bias: rural, low-cost counties are mechanically concentrated. Balance tests indicate the instrument correlates with demographic covariates, so the IV should be interpreted as bounding the direction of bias rather than as a clean causal estimate.

JEL Codes: I11, I18, L11

Keywords: hospital market concentration, Hill-Burton Act, Medicare spending, instrumental variables, healthcare competition

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 41m).

1. Introduction

The United States spends over $4.1 trillion annually on health care—roughly 18 percent of GDP—and hospital services account for the largest share of that expenditure. A growing body of evidence documents that hospital market concentration is associated with higher prices for privately insured patients (Cooper et al., 2019; Gaynor et al., 2015), and hospital mergers have intensified over the past two decades (Fulton, 2017; Gaynor and Town, 2012). Yet a basic question remains surprisingly difficult to answer: does hospital market concentration increase spending on the nation's largest public insurance program, Medicare?

The difficulty is endogeneity. Markets with few hospitals tend to be rural, low-income, and medically underserved—characteristics that independently predict lower health spending. Naive regressions of spending on concentration therefore confound market power with selection into concentrated markets, producing estimates that are biased toward zero or even negative. To make progress on this question, researchers need a source of variation in hospital supply that is plausibly independent of current demand conditions.

This paper proposes an instrument rooted in one of the largest federal health infrastructure programs in American history. The Hospital Survey and Construction Act of 1946—commonly known as the Hill-Burton Act—allocated billions of dollars in federal grants for hospital construction over a quarter century (United States Congress, 1946). Critically, the formula channeled funds inversely to state per capita income: poorer states received more generous allocations per capita. Because hospitals are durable capital with high fixed costs, the spatial distribution of hospitals created by Hill-Burton persists decades later, shaping the competitive landscape of hospital markets today. I use the inverse of 1950 state per capita income as an instrument for county-level hospital market concentration, measured by an equal-share Herfindahl-Hirschman Index (HHI).

The core finding is a stark divergence between ordinary least squares (OLS) and instrumental variables (IV) estimates that illuminates the nature and magnitude of selection bias. OLS regressions with demographic controls and state fixed effects yield a negative coefficient: a 10 percent increase in HHI is associated with a 0.4 percent *decrease* in standardized Medicare spending per beneficiary ($p < 0.01$). This negative relationship—seemingly suggesting that monopoly is cheap—is consistent with the selection story: concentrated markets are disproportionately rural and low-cost. When I instrument for HHI using Hill-Burton-era income variation, the coefficient flips sign and increases in magnitude by an order of magnitude. The two-stage least squares (2SLS) estimate implies that a 10 percent increase in HHI *raises* Medicare spending by 3.1 percent ($p < 0.01$), with a first-stage F-statistic of 28.2 under heteroskedasticity-robust inference (8.4 under the more appropriate state-level clustering).

I emphasize at the outset that the IV estimate should not be taken as a precise causal effect. Balance tests reveal that the instrument—inverse 1950 state per capita income—predicts several observable covariates, including county-level Medicare risk scores ($t = 17.6$), poverty rates ($t = 23.5$), and racial composition ($t = 16.1$). This correlation is unsurprising: the same historical poverty that attracted Hill-Burton funds also shaped demographic trajectories through channels unrelated to hospital supply. The exclusion restriction is therefore threatened. I present the IV results as providing suggestive evidence about the *direction* of selection bias—negative and substantial—rather than as a clean causal estimate of market power on spending.

The contribution of this paper is primarily diagnostic. The sign reversal between OLS and IV bounds the magnitude of selection bias that plagues cross-sectional studies of hospital competition and spending. The OLS estimate provides a plausible lower bound on the effect of concentration (biased downward by negative selection), while the IV provides an upper bound (potentially contaminated by other channels through which historical poverty affects current spending). The truth plausibly lies between them, and the gap itself—spanning nearly 0.35 log points—quantifies how misleading naive correlations can be.

Several additional results sharpen the interpretation. Decomposing Medicare spending into subcategories, I find that the negative OLS association is largest for outpatient spending (-0.084, $p < 0.01$) and smallest for inpatient spending (-0.014, $p < 0.05$), consistent with the fact that inpatient prices are more tightly regulated under Medicare's prospective payment system. Durable medical equipment (DME) spending and home health spending—categories that are not primarily hospital-delivered—show near-zero associations with HHI, serving as informative placebos. Heterogeneity analysis reveals that the negative OLS coefficient is five times larger in rural counties (-0.076) than in urban ones (-0.015), further confirming that the OLS result reflects rural selection rather than a true competitive effect.

This paper contributes to several literatures. First, it extends the extensive work on hospital competition and prices. Cooper et al. (2019) document a strong positive relationship between hospital concentration and private-insurance prices; Gowrisankaran et al. (2015) and Ho and Lee (2017) develop structural models of hospital-insurer bargaining that rationalize this pattern. Most of this literature focuses on private prices, where hospitals have bargaining leverage. Whether concentration also raises Medicare spending—where prices are administratively set—is less studied and mechanistically distinct, operating through volume, intensity, and site-of-service margins rather than price. Kessler and McClellan (2000) find that hospital competition improves quality and reduces costs for Medicare heart attack patients, but their identification relies on patient-level variation in proximity to hospitals.

Second, the paper connects to the literature on geographic variation in Medicare spending.

Wennberg et al. (2002), Skinner (2011), and Cutler et al. (2019) document large and persistent differences in per-beneficiary spending across regions, attributing them to supply-side factors including physician practice styles and hospital capacity. My results suggest that market structure—itself a consequence of historical supply-side investments—is a contributing factor.

Third, the paper builds on the small but growing literature using Hill-Burton as a source of identification. Finkelstein (2007) uses Hill-Burton funding to study the aggregate effects of Medicare's introduction on health spending. Almond et al. (2006) exploit Hill-Burton hospital desegregation requirements to study black-white convergence in infant mortality. I use the same historical episode but focus on a different mechanism: the long-run persistence of hospital supply and its effect on market structure.

The remainder of the paper proceeds as follows. Section 2 describes the Hill-Burton Act and the institutional features that make it relevant for contemporary hospital markets. Section 3 describes the data. Section 4 presents the empirical strategy. Section 5 reports results. Section 6 discusses implications and limitations.

2. Institutional Background

2.1 The Hill-Burton Act and Hospital Construction

The Hospital Survey and Construction Act of 1946, sponsored by Senators Lister Hill and Harold Burton, established the first major federal program for hospital construction in the United States. The legislation responded to a widely perceived shortage of hospital beds, particularly in rural and low-income areas, that had been documented during World War II mobilization (United States Congress, 1946). Between 1947 and 1971, the program distributed approximately \$3.7 billion in federal grants (roughly \$45 billion in 2024 dollars), funding the construction or renovation of nearly 6,800 health facilities across all 50 states.

The allocation formula is the key institutional feature for this paper's identification strategy. Federal funds were distributed to states based on a formula that weighted state population and the *inverse* of state per capita income. Poorer states received more federal dollars per capita, with the explicit goal of equalizing hospital capacity across regions. Within states, funds were further allocated to areas designated as having the greatest bed shortages. The formula was not discretionary—it was codified in statute and applied mechanically, reducing concerns about political targeting at the margin.

2.2 Persistence of Hospital Infrastructure

Hospitals are among the most durable forms of physical capital in the economy. A hospital built in the 1950s with Hill-Burton funds may still operate today, or its closure may have permanently reduced capacity in a market where replacement is prohibitively expensive. The high fixed costs of hospital construction—land acquisition, building, licensing, certificate-of-need requirements—create substantial barriers to entry and exit. As a result, the geographic distribution of hospitals established during the Hill-Burton era has proven remarkably persistent, even as demand conditions have shifted over seven decades.

This persistence is not merely anecdotal. The first-stage regression in my analysis confirms that 1950 state per capita income—a key determinant of Hill-Burton allocations—remains a strong predictor of county-level hospital market concentration in 2024. With heteroskedasticity-robust standard errors, the first-stage F-statistic is 28.2; with state-clustered standard errors (the more conservative choice, since the instrument varies at the state level), the effective F drops to 8.4—relevant but below the Stock and Yogo (2005) threshold for strong instruments. Markets in states that were relatively poor in 1950 tend to have more hospitals per capita today, and consequently lower HHI.

2.3 Medicare Spending and Hospital Market Structure

Medicare spending per beneficiary varies enormously across the United States, with per-capita spending in the highest-cost regions roughly double that of the lowest-cost regions (Wennberg and Cooper, 1999). Unlike private insurance markets, where hospital concentration can directly raise negotiated prices (Cooper et al., 2019), Medicare sets payment rates administratively through the Inpatient Prospective Payment System and the Outpatient Prospective Payment System. Hospitals are largely price-takers with respect to Medicare.

However, concentration can still affect Medicare spending through several non-price channels. Hospitals with market power may engage in more aggressive coding and upcoding of diagnosis-related groups, increasing effective reimbursement per admission. Concentrated markets may have higher utilization rates—more imaging, more procedures, longer stays—because competitive pressure to restrain costs is absent. Site-of-service shifting, where procedures migrate from lower-cost physician offices to higher-cost hospital outpatient departments, may also be more prevalent when hospitals face less competition. These mechanisms suggest that the effect of concentration on Medicare spending, if it exists, operates through volume and intensity margins rather than through prices per se.

3. Data

I construct a county-level cross-sectional dataset for 2019 (pre-COVID) by merging four administrative and survey sources. The unit of observation is a U.S. county or county-equivalent.

Medicare spending. I measure standardized Medicare fee-for-service spending per beneficiary at the county level from the CMS Geographic Variation Public Use File (Centers for Medicare & Medicaid Services, 2019). Standardization removes geographic differences in Medicare payment rates (wage indices, practice cost indices, disproportionate share adjustments), isolating variation in utilization and intensity. The PUF also provides disaggregated spending by category: inpatient, outpatient, evaluation and management (E&M), durable medical equipment (DME), and home health.

Hospital market structure. I obtain the universe of Medicare-certified hospitals from CMS Hospital Compare (Centers for Medicare & Medicaid Services, 2024), which provides facility-level geocoded data including hospital type. I restrict to Acute Care Hospitals, Critical Access Hospitals, and VA Medical Centers—the three types that provide general inpatient and emergency services—and count the number of such facilities per county. I construct an equal-share HHI as HHI $= 10{,}000/N$, where N is the number of hospitals. This measure assumes equal market shares within a county, which is a simplification but avoids the need for discharge-level data that are unavailable at the county level. Counties with zero hospitals are excluded from the regression sample, as HHI is undefined. The median county has exactly one hospital, yielding HHI $= 10{,}000$ (a monopoly). Only 24.3 percent of counties have two or more hospitals.

Instrument. I use the inverse of 1950 state per capita personal income, obtained from the Bureau of Economic Analysis historical series. This variable directly enters the Hill-Burton allocation formula: states with lower 1950 income received more construction funds per capita. Because the instrument varies at the state level, it captures broad geographic patterns in hospital supply rather than county-specific idiosyncrasies.

Demographic controls. County-level demographic and socioeconomic variables come from the American Community Survey 2022 five-year estimates: total population, poverty rate, share aged 65 and over, median age, and Black population share. Medicare-specific controls include average Hierarchical Condition Category (HCC) risk scores and dual-eligible beneficiary shares from the CMS PUF.

Table 1: Summary Statistics

	N	Mean	SD
Panel A: Medicare Spending (2019)			
Medicare std. spending/capita ($)	3,134	10,193.37	1,658.41
Inpatient std. spending/capita ($)	3,134	2,702.02	488.65
Panel B: Hospital Market Structure			
Hospitals (all types)	3,134	1.42	2.62
HHI (equal-share)	3,134	8,492.02	2,753.8
Competitive ($\geq$ 2 hospitals)	3,134	0.24	0.43
Panel C: County Characteristics			
FFS beneficiaries	3,134	10,549.46	24,124.18
Average HCC risk score	3,134	0.96	0.11
Population	3,123	104,858.4	334,157.5
Poverty rate	3,123	0.14	0.06
Share 65+	3,123	0.2	0.05
Median age	3,123	41.6	5.39
Dual-eligible share (%)	3,124	0.19	0.08
Black share (%)	1,942	0.08	0.11

Notes: Unit of observation is county. Medicare spending from CMS Geographic Variation PUF (2019, pre-COVID). Hospital counts from CMS Hospital Compare (2024), including Acute Care, Critical Access, and VA hospitals. HHI computed as equal-share Herfindahl–Hirschman Index (10,000/N). County demographics from ACS 2022 5-year estimates.

3.1 Summary Statistics

Table 1 reports summary statistics for the 3,134 U.S. counties in the sample. Mean standardized Medicare spending is $10,193 per beneficiary, with a standard deviation of $1,658. The average county has 1.4 hospitals and a mean HHI of 8,492, reflecting the extreme concentration of most hospital markets. Three-quarters of counties have zero or one hospital; only 763 counties have two or more. The regression sample is smaller (1,620 counties) due to the requirement that counties have at least one hospital and non-missing values for all controls.

4. Empirical Strategy

4.1 OLS Specification

The baseline OLS specification estimates the cross-sectional relationship between hospital market concentration and Medicare spending:

$$\ln(Y_c) = \alpha + \beta \ln(\text{HHI}_c) + X_c'\gamma + \delta_s + \varepsilon_c \tag{1}$$

where Y_c is standardized Medicare spending per beneficiary in county c, HHI_c is the equal-share Herfindahl-Hirschman Index, X_c is a vector of demographic controls (HCC risk score, log population, poverty rate, share 65+, dual-eligible share, Black share), and δ_s are state fixed effects. The coefficient β is the elasticity of Medicare spending with respect to market concentration. Standard errors are heteroskedasticity-robust throughout.

The log-log specification has two advantages: it is scale-invariant, and β has a direct interpretation as an elasticity. I estimate the model sequentially—without controls, with controls, and with state fixed effects—to trace how the coefficient evolves as confounders are absorbed.

4.2 Instrumental Variables Strategy

The central concern with OLS is that HHI_c is endogenous. Counties with few hospitals may be concentrated for reasons (low population density, low demand, limited economic activity) that independently predict low spending. To address this, I estimate a two-stage least squares (2SLS) model using inverse 1950 state per capita income as an instrument for $\ln(\text{HHI}_c)$.

First stage. The first-stage equation is:

$$\ln(\text{HHI}_c) = \pi_0 + \pi_1 Z_s + X_c'\pi_2 + \nu_c \tag{2}$$

where $Z_s = 1/\text{PCI}_{s,1950}$ is the inverse of state per capita income in 1950. The instrument enters at the state level, reflecting the Hill-Burton allocation formula. The prediction is $\pi_1 > 0$: states that were poorer in 1950 received more hospital construction funds, resulting in more hospitals and lower HHI—but since Z_s is the *inverse* of income, higher Z_s (lower income) should predict lower HHI (more hospitals), so $\pi_1 < 0$. The first-stage F-statistic is 28.2, well above conventional weak-instrument thresholds (Stock and Yogo, 2005; Staiger and Stock, 1997).

Second stage. The second-stage equation replaces $\ln(\text{HHI}_c)$ with the predicted values from the first stage:

$$\ln(Y_c) = \alpha + \beta^{IV} \widehat{\ln(\text{HHI}_c)} + X_c'\gamma + \varepsilon_c \tag{3}$$

Under the assumption that inverse 1950 state income affects current Medicare spending *only* through its effect on hospital market structure (the exclusion restriction), β^{IV} identifies the causal effect of concentration on spending.

4.3 Threats to Validity

Exclusion restriction. The identifying assumption requires that 1950 state income affects 2019 county Medicare spending only through hospital market concentration. This assumption is strong and almost certainly violated to some degree. States that were poor in 1950 differ from wealthy states along many dimensions—racial composition, educational attainment, industrial structure, public health infrastructure—that may independently affect health spending patterns 70 years later. The balance tests in Section 5 confirm these concerns: the instrument predicts Medicare risk scores, poverty rates, and racial composition of counties, all of which plausibly affect spending through channels other than hospital concentration.

Interpretation of IV. Given the likely violation of the exclusion restriction, I interpret the IV estimates cautiously. The key informational content is the *sign reversal* relative to OLS. If OLS is biased downward (by negative selection of low-cost counties into concentrated markets) and IV is biased upward (by residual correlation of historical poverty with current spending), then the true causal effect lies between the two estimates. The OLS coefficient of −0.043 and the IV coefficient of +0.311 bracket a wide range, but the reversal itself establishes that the naive negative association is driven by selection and that the causal effect of concentration on spending is likely positive or at worst zero.

Measurement of HHI. The equal-share HHI is an imperfect measure of market concentration because it assumes symmetric hospital size. In reality, markets with one large teaching hospital and several small community hospitals are less competitive than the equal-share measure suggests. This measurement error likely attenuates the OLS estimates toward zero, reinforcing the conclusion that the true relationship is more positive than OLS suggests. The IV estimates are consistent under classical measurement error in the endogenous variable, provided the instrument is valid (Angrist and Pischke, 2009).

5. Results

5.1 Main Results

Table 2 reports the main estimates of the relationship between hospital market concentration and Medicare spending. The progression from column (1) through (4) traces how the estimated coefficient changes as the specification is enriched, and the reversal between OLS and IV constitutes the paper's central finding.

Column (1) reports the simple bivariate regression of log Medicare spending on log HHI without controls. The coefficient is -0.007 and statistically insignificant ($p > 0.10$). In the raw data, there is essentially no relationship between concentration and spending. This null result masks two offsetting forces: concentrated markets tend to be rural and cheap (selection), while concentration may raise spending through market power (the causal effect of interest).

Adding demographic controls in column (2) reveals the selection channel. The coefficient jumps to -0.053 ($p < 0.01$). Once we compare counties with similar risk scores, populations, and poverty rates, more concentrated markets have *lower* spending. The HCC risk score alone explains much of the variation—its coefficient of 1.40 implies that a one-unit increase in average risk (roughly one standard deviation) raises spending by 140 log points. Column (3) adds state fixed effects, reducing the HHI coefficient modestly to -0.043 ($p < 0.01$). Within states, the negative relationship between concentration and spending persists, suggesting that the selection story operates at the sub-state level: within a given state, the most concentrated (typically most rural) counties have lower spending.

Column (4) presents the IV estimate, which reverses the sign. When HHI is instrumented with inverse 1950 state per capita income, the coefficient is $+0.311$ ($p < 0.10$) with standard errors clustered at the state level. This estimate implies that a 10 percent increase in concentration raises Medicare spending by about 3 percent. However, inference must be cautious: because the instrument varies only at the state level, proper clustering on 49 states yields a first-stage effective F-statistic of 8.4—below the Stock and Yogo (2005) threshold of 16.38 for 10 percent maximal IV size. The instrument is relevant but not strong under conservative inference. Column (5) adds a second instrument (inverse 1960 state per capita income) for over-identification; the coefficient is virtually unchanged at 0.311 but is no longer significant at conventional levels ($p > 0.10$).

The sign reversal between columns (3) and (4), while imprecisely estimated, is the main result. The OLS coefficient is -0.043 and the IV coefficient is $+0.311$, a gap of 0.354 log points that is suggestive of substantial negative selection bias in OLS. Even accounting for the possibility that the IV is biased upward (by residual channels through which historical poverty affects spending), the reversal establishes that the naive negative association between concentration and spending is likely driven by selection rather than a true competitive benefit.

Column (6) reports a specification using a binary indicator for "competitive" markets (two or more hospitals) as the endogenous variable. The IV coefficient of -0.459 ($p < 0.10$) on the competitive indicator implies an implausibly large spending reduction for monopoly-to-competitive transitions, further suggesting that the instrument captures channels beyond market structure alone. I take this as evidence that the IV magnitudes should be interpreted with caution, while the directional finding—that OLS is biased downward—is robust.

Balance tests. To assess the plausibility of the exclusion restriction, I regress each covariate on the instrument. The results are concerning. Inverse 1950 state income is a strong predictor of HCC risk scores ($t = 17.6$), poverty rates ($t = 23.5$), and Black population share ($t = 16.1$). States that were poor in 1950 have counties with sicker, poorer, and more racially diverse populations today. While these covariates are included as controls in the 2SLS specification, the correlations raise the possibility that the instrument affects spending through channels other than hospital supply. The IV estimate should therefore be interpreted as an upper bound on the causal effect, not as a precise point estimate.

5.2 Spending Subcategories and Placebos

Table 3 decomposes the OLS relationship between concentration and spending into five Medicare spending subcategories, using the state-fixed-effects specification from column (3) of Table 2. This decomposition serves two purposes: it illuminates which margins of spending drive the aggregate result, and it provides informal placebo tests using spending categories that should be unrelated to hospital market power.

The coefficient on inpatient spending (column 1) is -0.014 ($p < 0.05$), roughly one-third the magnitude of the total spending coefficient. This is consistent with the fact that Medicare inpatient prices are tightly regulated under the prospective payment system, limiting the scope for hospitals to increase revenue per admission. The coefficient on outpatient spending (column 2) is much larger: -0.084 ($p < 0.01$). Outpatient services—including hospital outpatient department visits, ambulatory surgery, and diagnostic imaging—are reimbursed under fee schedules that leave more room for volume and intensity adjustments. The stronger relationship between concentration and outpatient spending is consistent with concentrated hospitals engaging in more aggressive utilization of outpatient services.

Durable medical equipment (DME) spending (column 3) shows a near-zero and insignificant coefficient of $+0.015$. DME is supplied by vendors, not hospitals, so hospital concentration should have no direct effect on DME spending. This null result serves as a clean placebo:

Table 2: Hospital Market Concentration and Medicare Spending

Dependent Variable:			log_tot_spend			
		OLS		IV	IV (Over-ID)	IV
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Constant	9.290***	9.110***		4.093	4.104	7.247***
	(0.0895)	(0.1989)		(2.580)	(2.654)	(0.6264)
log(HHI)	-0.0072	-0.0531***	-0.0430***	0.3113*	0.3105	
	(0.0102)	(0.0093)	(0.0071)	(0.1816)	(0.1873)	
HCC risk score		1.401***	0.9974***	1.178***	1.179***	1.112***
		(0.0853)	(0.0643)	(0.1707)	(0.1731)	(0.1994)
log(Population)		-0.0555***	-0.0340***	0.1092	0.1088	0.0933
		(0.0091)	(0.0069)	(0.0863)	(0.0886)	(0.0690)
Poverty rate		0.0255	-0.2145**	0.5706	0.5694	0.7436*
		(0.1211)	(0.0985)	(0.3496)	(0.3539)	(0.3897)
Share 65+		-0.2981**	-0.1256*	0.0848	0.0839	0.0404
		(0.1272)	(0.0666)	(0.3349)	(0.3355)	(0.2366)
Dual-eligible		-0.6350***	-0.2388***	-0.2844	-0.2852	-0.3708
		(0.1159)	(0.0635)	(0.2958)	(0.3019)	(0.2251)
Black share		0.0806	0.0118	-0.0339	-0.0336	-0.2055
		(0.0897)	(0.0391)	(0.1355)	(0.1399)	(0.1414)
Competitive						-0.4594*
						(0.2630)
Fixed-effects						
state_fips_fac			Yes			
Fit statistics						
Observations	2,357	1,620	1,620	1,620	1,620	1,930
R^2	0.00077	0.63671	0.79735	0.62602	0.62607	0.61446
F-test (1st stage), log(HHI)				28.209	14.217	
F-test (1st stage), Competitive						23.985

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Standard errors clustered at state level in parentheses

the fact that HHI does not predict DME spending increases confidence that the relationships with inpatient and outpatient spending reflect hospital-specific mechanisms rather than spurious county-level correlations. Home health spending (column 5) is similarly null at 0.0002, providing a second placebo. Evaluation and management spending (column 4) shows a positive coefficient of +0.062 ($p < 0.01$), which may reflect the fact that physician practice patterns are shaped by the hospital environment in which they operate.

Table 3: Hospital Concentration and Medicare Spending by Category

Model:	Inpatient (1)	Outpatient (2)	DME (3)	E&M (4)	Home Health (5)
Variables					
log(HHI)	-0.0142**	-0.0842***	0.0146	0.0618***	0.0002
	(0.0057)	(0.0116)	(0.0092)	(0.0075)	(0.0194)
Fixed-effects					
state_fips_fac	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	1,620	1,620	1,620	1,620	1,615
R^2	0.69461	0.47329	0.56893	0.80930	0.72485

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

5.3 Heterogeneity

Table 4 reports the OLS coefficient from the state-fixed-effects specification estimated separately for urban and rural counties. The coefficient is -0.015 ($p < 0.01$) for urban counties and -0.076 ($p < 0.01$) for rural counties—a fivefold difference. This heterogeneity is consistent with the selection story. Rural counties exhibit a steeper negative gradient because the variation in HHI across rural areas is more strongly correlated with population density and economic activity. A rural county with three hospitals is likely a regional hub (larger, wealthier, higher-spending), while a rural county with one hospital is likely a sparsely populated area with low utilization. In urban areas, this selection gradient is muted because even the most concentrated urban markets have substantial population and economic activity.

The urban-rural decomposition also speaks to the external validity of the estimates. The aggregate negative OLS coefficient is disproportionately driven by the rural subsample, which constitutes the majority of counties but a minority of the population. For the urban counties where most Medicare beneficiaries reside, the negative selection bias is smaller, suggesting that the OLS estimate may be closer to the true causal effect in urban markets.

5.4 Robustness

Table 4 presents four additional robustness checks, all using state fixed effects. Column (3) replaces log HHI with log number of hospitals as the concentration measure; the results are mechanically equivalent given the equal-share construction. Column (4) trims outlier counties (top and bottom 1 percent of log spending), yielding a coefficient of −0.039, close to the baseline of −0.043. Column (5) clusters standard errors at the state level rather than using heteroskedasticity-robust errors; the point estimate is unchanged but the standard error increases from 0.005 to 0.007, reflecting within-state correlation. All robustness checks confirm the baseline OLS result.

Table 4: Robustness Checks

Dependent Variable:	log_tot_spend				
	Urban	Rural	log(N hosp.)	Trimmed	Clustered SE
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
log(HHI)	-0.0154***	-0.0762***		-0.0387***	-0.0430***
	(0.0042)	(0.0148)		(0.0045)	(0.0071)
Fixed-effects					
state_fips_fac	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Observations	1,142	476	1,620	1,599	1,620
R^2	0.84789	0.78985	0.79735	0.79840	0.79735

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

6. Discussion

The central finding of this paper is a diagnostic one: the naive cross-sectional relationship between hospital market concentration and Medicare spending is negative, but this negative correlation is driven by selection rather than causation. Rural, low-cost, medically simple counties are mechanically concentrated because they have few hospitals. Once this selection is addressed—even imperfectly—the estimated effect reverses sign, suggesting that concentration raises spending.

The magnitude of the selection bias, spanning 0.35 log points between the OLS and IV estimates, has important implications for how policymakers and researchers interpret the hospital competition literature. Studies that rely on cross-sectional correlations between market structure and spending—whether for Medicare or private insurance—may substantially understate the effect of concentration if they do not adequately control for the characteristics that sort low-cost populations into concentrated markets. The finding echoes a broader lesson in empirical economics: the sign of the bias matters as much as its existence, and in this case, the bias is large enough to flip the qualitative conclusion from "concentration reduces costs" to "concentration raises costs."

The placebo results provide supporting evidence that the OLS relationship is not merely a county-level artifact. The near-zero coefficients on DME and home health spending—categories unrelated to hospital market power—suggest that the negative association between HHI and hospital spending (inpatient and outpatient) reflects hospital-specific mechanisms. The larger outpatient coefficient relative to inpatient is consistent with the institutional environment: Medicare inpatient prices are regulated through DRG-based prospective payment, while outpatient services offer more scope for volume and intensity adjustments.

These results connect to the broader literature on the causes of geographic variation in Medicare spending. Skinner (2011) argues that supply-side factors—physician practice styles, hospital capacity, and local medical culture—are the primary drivers of spending variation. My results suggest that market structure, itself a product of historical supply-side investments like Hill-Burton, is one channel through which supply shapes spending. The infrastructure persistence channel—from federal grants in the 1950s to market concentration in the 2020s—illustrates how policy interventions can have effects that outlast the programs themselves by decades.

The paper has clear limitations that should temper the interpretation. The instrument is not clean: inverse 1950 state income predicts demographic and health characteristics of counties that plausibly affect spending through non-hospital channels. The equal-share HHI is a coarse measure of concentration that ignores variation in hospital size and specialty. The cross-sectional design cannot address time-varying confounders. And the state-level instrument cannot identify within-state causal effects, which may differ from the between-state effects captured by the IV. Future work could exploit within-state variation in Hill-Burton allocations, which were determined at the sub-state level based on bed shortage surveys, to construct a more granular instrument.

Despite these limitations, the paper contributes a clear piece of evidence to the hospital competition debate: the negative correlation between concentration and Medicare spending is a selection artifact. Any policy analysis that treats this correlation as causal will substantially understate the costs of hospital consolidation for the Medicare program.

7. Conclusion

Hospital markets across the United States are remarkably concentrated, and the degree of concentration reflects infrastructure investments made 70 years ago under the Hill-Burton Act. The naive negative correlation between concentration and Medicare spending is driven by selection: rural, low-cost counties are concentrated because they lack the demand to support multiple hospitals, not because their hospitals compete less aggressively. Once this selection is addressed, the evidence points toward concentration raising spending, consistent with the theoretical prediction that market power allows hospitals to increase utilization and intensity even when prices are regulated.

The infrastructure persistence channel documented here suggests that the long-run consequences of large-scale public capital investments may extend far beyond their original policy objectives. Hill-Burton aimed to equalize hospital access; seven decades later, the spatial distribution of hospitals it created continues to shape the competitive landscape of health care markets and, through that channel, the cost of caring for 65 million Medicare beneficiaries.

Acknowledgements

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

Project Repository: https://github.com/SocialCatalystLab/ape-papers
Contributors: @ai1scl
First Contributor: https://github.com/ai1scl

References

Almond, Douglas, Kenneth Y Chay, and Michael Greenstone, "Civil Rights, the War on Poverty, and Black-White Convergence in Infant Mortality in the Rural South and Mississippi," *MIT Department of Economics Working Paper*, 2006.

Angrist, Joshua D and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, 2009.

Centers for Medicare & Medicaid Services, "Geographic Variation Public Use File," Technical Report, CMS 2019. Data vintage: calendar year 2019.

_ , "Hospital Compare General Information," Technical Report, CMS 2024.

Cooper, Zack, Stuart V Craig, Martin Gaynor, and John Van Reenen, "The Price Ain't Right? Hospital Prices and Health Spending on the Privately Insured," *Quarterly Journal of Economics*, 2019, *134* (1), 51–107.

Cutler, David, Jonathan Skinner, Ariel D Stern, and David Wennberg, "The Geography of Medicare," *American Economic Review: Papers & Proceedings*, 2019, *109* (5), 72–76.

Finkelstein, Amy, "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare," *Quarterly Journal of Economics*, 2007, *122* (1), 1–37.

Fulton, Brent D, "Health Care Market Concentration Trends in the United States: Evidence and Policy Responses," *Health Affairs*, 2017, *36* (9), 1530–1538.

Gaynor, Martin and Robert Town, "The Impact of Hospital Consolidation—Update," *Robert Wood Johnson Foundation Synthesis Project*, 2012.

_ **, Kate Ho, and Robert J Town**, "The Industrial Organization of Health-Care Markets," *Journal of Economic Literature*, 2015, *53* (2), 235–284.

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town, "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, 2015, *105* (1), 172–203.

Ho, Kate and Robin S Lee, "Insurer Competition in Health Care Markets," *Econometrica*, 2017, *85* (2), 379–417.

Kessler, Daniel P and Mark B McClellan, “Is Hospital Competition Socially Wasteful?,” *Quarterly Journal of Economics*, 2000, *115* (2), 577–615.

Skinner, Jonathan, “The Causes and Consequences of Regional Variations in Health Care,” *Handbook of Health Economics*, 2011, *2*, 45–93.

Staiger, Douglas and James H Stock, “Instrumental Variables Regression with Weak Instruments,” *Econometrica*, 1997, *65* (3), 557–586.

Stock, James H and Motohiro Yogo, “Testing for Weak Instruments in Linear IV Regression,” in Donald W K Andrews and James H Stock, eds., *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge University Press, 2005, pp. 80–108.

United States Congress, “Hospital Survey and Construction Act,” *Public Law 79-725*, 1946.

Wennberg, John E and Megan M Cooper, “The Dartmouth Atlas of Health Care,” *American Hospital Publishing*, 1999.

_ , Elliott S Fisher, and Jonathan S Skinner, “Geography and the Debate over Medicare Reform,” *Health Affairs*, 2002, *21* (Suppl1), W96–W114.

A. Data Appendix

A.1 Data Sources and Access

CMS Geographic Variation Public Use File (2019). The primary outcome data come from the CMS Geographic Variation PUF, calendar year 2019, downloaded from https://data.cms.gov/. This file provides county-level aggregates of standardized Medicare fee-for-service spending per beneficiary, disaggregated by service type (inpatient, outpatient, E&M, DME, home health, and others). Standardization removes geographic payment adjustments, isolating differences in utilization and intensity. The file covers all 50 states plus the District of Columbia and U.S. territories. I restrict to the 50 states and DC.

CMS Hospital Compare (2024). Hospital locations and types come from the CMS Hospital Compare General Information file. I retain three facility types: Acute Care Hospitals, Critical Access Hospitals, and VA Medical Centers. Hospitals are geocoded to counties using FIPS codes. The 2024 vintage reflects current market structure; I use it as a proxy for the 2019 market structure under the assumption that hospital entry and exit between 2019 and 2024 was minimal. A total of 4,453 facilities are matched to counties.

Bureau of Economic Analysis (1950). State per capita personal income in 1950, measured in current dollars, is obtained from the BEA Regional Economic Accounts historical tables. These data are available for all 48 contiguous states plus DC (Alaska and Hawaii were not yet states in 1950 and are excluded from the IV sample).

American Community Survey (2022). County-level demographic variables come from the ACS 5-year estimates (2018–2022), accessed via the Census Bureau API. Variables include total population, poverty rate, share of population aged 65 and over, median age, and Black or African American population share. Some counties suppress demographic variables due to small population sizes, resulting in missing values that reduce the regression sample.

A.2 Sample Construction

The initial sample consists of 3,134 counties with non-missing Medicare spending data. The regression sample is restricted by the following filters: (1) counties must have at least one hospital (Acute Care, Critical Access, or VA), reducing the sample to approximately 2,400 counties; (2) counties must have non-missing values for all demographic controls, further reducing to 1,620 counties; (3) the IV sample additionally requires the county to be in a state with 1950 income data (excluding Alaska and Hawaii). The final regression sample for the main specification contains 1,620 counties.

A.3 Variable Definitions

- **Standardized Medicare spending per capita**: Total standardized FFS spending divided by the number of FFS beneficiaries, from CMS Geographic Variation PUF. Standardization removes geographic payment adjustments.
- **HHI (equal-share)**: 10,000/N where N is the number of hospitals (Acute Care + Critical Access + VA) in the county. Range: 10,000 (monopoly) to 10,000/N for counties with N hospitals.
- **Inverse 1950 state PCI**: $1/\text{PCI}_{s,1950}$ where PCI is state per capita personal income in 1950 dollars.
- **HCC risk score**: Average Hierarchical Condition Category risk score for FFS beneficiaries in the county, from CMS PUF. Higher values indicate sicker populations.
- **Dual-eligible share**: Share of Medicare beneficiaries also enrolled in Medicaid, a proxy for low income.

B. Identification Appendix

B.1 First-Stage Diagnostics

The first-stage regression of ln(HHI) on the inverse of 1950 state per capita income, controlling for demographic covariates, yields an F-statistic of 28.2 with heteroskedasticity-robust standard errors. Because the instrument varies at the state level, the appropriate inference clusters standard errors on states. Under state-level clustering (49 clusters), the effective first-stage F-statistic drops to 8.4, below the Stock and Yogo (2005) critical value of 16.38 for 10 percent maximal IV size. The instrument is relevant but not strong under conservative inference, and IV estimates should be interpreted with appropriate caution regarding potential weak-instrument bias.

B.2 Balance Tests

I regress each covariate on the instrument (inverse 1950 state PCI) to assess balance. The following covariates are significantly predicted by the instrument: HCC risk score ($t = 17.6$), poverty rate ($t = 23.5$), Black population share ($t = 16.1$), share 65+ ($t = 8.3$), and dual-eligible share ($t = 12.1$). These imbalances indicate that the instrument is correlated with county characteristics that plausibly affect Medicare spending through non-hospital channels. The inclusion of these variables as controls in the 2SLS specification mitigates but does not eliminate this concern, because the controls may not fully capture the relevant confounding pathways.

B.3 Interpretation Under Imperfect Exclusion

Given the balance test results, the IV estimate should be interpreted as an upper bound under the following reasoning. If the instrument is positively correlated with unobserved determinants of spending (e.g., medical culture, provider practice styles shaped by historical poverty), then the 2SLS estimate is biased upward. Combined with the OLS estimate, which is biased downward by negative selection, the two estimates bracket the true causal effect: $-0.043 \leq \beta^{\text{true}} \leq 0.311$.

C. Robustness Appendix

The robustness checks in Table 4 address several concerns. Trimming the top and bottom 1 percent of spending eliminates the possibility that outlier counties drive the results; the coefficient moves from -0.043 to -0.039. Clustering standard errors at the state level accounts for within-state correlation in both the outcome and the instrument; the standard error increases from 0.005 to 0.007 but remains highly significant. The urban and rural subsamples show that the negative OLS relationship is present in both subgroups but is substantially larger in rural areas, consistent with the selection interpretation.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Total Medicare spending (OLS)	-0.043	0.0052	0.1556	-0.1665	0.0202	Large negative
Total Medicare spending (IV)	0.3113	0.074	0.1556	1.2055	0.2865	Large positive
Inpatient spending (OLS)	-0.0142	0.0057	0.1768	-0.0485	0.0195	Small negative
Panel B: Heterogeneous (sample splits)						
Total spending, urban (OLS)	-0.0154	0.0042	0.1377	-0.0673	0.0183	Moderate negative
Total spending, rural (OLS)	-0.0762	0.0148	0.1708	-0.2688	0.0521	Large negative

Notes: **Country:** United States. **Research question:** Does hospital market concentration, shaped by historical federal hospital construction policy, increase county-level Medicare spending per beneficiary? **Policy mechanism:** The Hill-Burton Act (1946–1971) allocated federal construction grants inversely to state per capita income, creating persistent cross-county variation in hospital supply that determines current market structure. **Outcome definition:** Log standardized Medicare fee-for-service per capita spending (CMS Geographic Variation PUF), price-adjusted to remove geographic payment differences. **Treatment:** Continuous; log Herfindahl–Hirschman Index based on equal-share hospital counts per county. **Data:** CMS Geographic Variation PUF (2019), CMS Hospital Compare (2024), BEA state personal income (1950), ACS 5-year (2022); 1,620 counties with ≥ 1 hospital and non-missing controls. **Method:** IV/2SLS with inverse 1950 state per capita income as instrument, heteroskedasticity-robust standard errors; OLS specifications include state fixed effects. **Sample:** Counties with at least one Medicare-certified hospital (Acute Care, Critical Access, or VA); excludes counties with fewer than 100 FFS beneficiaries. SDE $= \hat{\beta} \times \mathrm{SD}(X)/\mathrm{SD}(Y)$ for continuous treatment, where SD(X) and SD(Y) are cross-sectional standard deviations. Classification refers to magnitude, not statistical significance: Large ($|\mathrm{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).