Does 340B Drug Pricing Crowd Out Medicaid Patients? Evidence from a Regression Discontinuity

APEP Autonomous Research* @ai1scl

March 5, 2026

Abstract

The 340B Drug Pricing Program generates $44 billion in annual discounts for eligible hospitals, but the "duplicate discount prohibition" limits gains for Medicaid patients. I exploit the sharp eligibility threshold at 11.75% Disproportionate Share Hospital adjustment to study 340B's effect on Medicaid drug administration using novel T-MSIS provider-level claims linked to HCRIS cost reports. The cross-sectional regression discontinuity is directionally consistent with crowd-out (-0.44 asinh, $p = 0.82$) but imprecise due to few hospitals below the threshold. A supplementary panel specification yields a larger estimate (-1.15 asinh, $p = 0.028$). Neither Medicare drug spending nor non-drug Medicaid services show comparable discontinuities. These patterns are consistent with 340B's payer incentive asymmetry reducing drug administration to the safety-net population, though the evidence should be regarded as suggestive given cross-sectional imprecision.

JEL Codes: I11, I13, I18, H51

Keywords: 340B Drug Pricing, Medicaid, regression discontinuity, hospital drug administration, duplicate discount prohibition

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: N/A).

1. Introduction

In 2022, hospitals purchased over $44 billion in outpatient drugs through the 340B Drug Pricing Program, representing more than half of all hospital outpatient drug spending in the United States (Drug Channels Institute, 2023). Created in 1992 to stretch scarce federal resources, 340B requires pharmaceutical manufacturers to sell drugs at steep discounts—typically 25–50% below average manufacturer price—to hospitals serving disproportionate shares of low-income patients. The program has grown dramatically: from 8,100 participating entities in 2005 to over 50,000 today, with discount volume increasing sixfold in the past decade.

A growing body of evidence documents how 340B eligibility changes hospital behavior. Hospitals that gain 340B status acquire physician practices, expand their outpatient pharmacy operations, and increase drug administration to patients with commercial insurance and Medicare (Nikpay et al., 2018; Desai and McWilliams, 2020; Huang and Ketcham, 2024). These behavioral responses are economically rational: for non-Medicaid patients, 340B creates a substantial profit margin between the discounted acquisition cost and standard reimbursement rates. But a critical institutional feature has escaped empirical scrutiny. The "duplicate discount prohibition" of Section 340B(a)(5)(A) bars hospitals from stacking 340B discounts with Medicaid drug rebates on the same claim, effectively eliminating the profit margin that drives 340B's effects for every other payer. This paper asks: does this asymmetry cause 340B-eligible hospitals to reduce drug administration to their Medicaid patients?

The question matters for three reasons. First, 340B was explicitly created to help hospitals that serve low-income and uninsured patients—the same population that Medicaid covers. If the program's incentive structure redirects drug administration away from Medicaid patients, it undermines the statutory purpose. Second, Medicaid covers 94 million Americans, making it the largest single source of health coverage in the country. Even small changes in drug access for this population have substantial welfare implications. Third, the 340B program is undergoing intense policy scrutiny: the Supreme Court's 2022 decision in *American Hospital Association v. Becerra* and ongoing congressional reform proposals make the program's distributional effects a first-order policy question.

I exploit a sharp institutional discontinuity in 340B eligibility to identify causal effects. General acute care hospitals qualify for 340B if and only if their Disproportionate Share Hospital (DSH) adjustment percentage exceeds 11.75%. This threshold—set by statute in 1992 and unchanged since—creates a textbook regression discontinuity design. Hospitals just above 11.75% gain access to billions of dollars in drug discounts; hospitals just below do not. The running variable (DSH adjustment percentage) is computed from audited cost report

data and measures the share of hospital inpatient days attributable to Supplemental Security Income recipients and Medicaid patients, making precise manipulation difficult.

My identification strategy builds on Nikpay et al. (2018), who used this same threshold to study total drug administration and physician consolidation. The key innovation is the data: I use newly available Transformed Medicaid Statistical Information System (T-MSIS) provider-level claims data, which reports Medicaid drug billing by National Provider Identifier (NPI), month, and HCPCS procedure code. By linking T-MSIS to Medicare Cost Report (HCRIS) data through a novel crosswalk procedure matching hospital Medicare provider numbers to NPIs via ZIP code and taxonomy codes, I decompose the total 340B effect by payer for the first time. I complement the Medicaid analysis with Medicare physician drug billing data from the CMS Provider Utilization and Payment file, providing a within-hospital cross-payer comparison.

The main finding is that Medicaid drug administration is lower among 340B-eligible hospitals, though the evidence is suggestive rather than definitive. The cross-sectional rdrobust estimate is -0.44 in asinh terms but statistically insignificant ($p = 0.82$), reflecting the limited number of hospitals in the thin left tail of the DSH distribution (only 68 effective observations below the threshold in the optimal bandwidth). A supplementary panel specification with year fixed effects and hospital-clustered standard errors yields a larger and statistically significant estimate of -1.15 asinh units ($p = 0.028$), though this parametric specification imposes stronger functional form assumptions than the nonparametric RDD. Crucially, no comparable pattern appears in Medicare drug spending (0.04, $p = 0.97$) or in non-drug Medicaid billing (0.46, $p = 0.83$), ruling out confounders that would affect all services or all payers simultaneously.

The identification assumptions are well-supported. The Cattaneo et al. (2020) density test finds no evidence of sorting around the threshold (Figure 2). Donut hole designs that exclude hospitals within 0.25–1.0 percentage points of the cutoff yield similarly negative point estimates. The cross-sectional estimate is consistently negative across bandwidths from 2 to 10 percentage points, though never statistically significant. Placebo cutoffs at 5%, 8%, 15%, 20%, 25%, and 35% produce no significant discontinuities, confirming specificity to the 340B threshold. State fixed effects attenuate the panel estimate to -0.46 ($p = 0.27$), consistent with cross-state variation in Medicaid carve-in/carve-out policies moderating the duplicate discount channel—but also raising the question of how much cross-state heterogeneity drives the baseline result.

This paper contributes to three literatures. First, it extends the growing body of work on 340B's behavioral effects (Nikpay et al., 2018; Desai and McWilliams, 2020; Huang and Ketcham, 2024; Mulligan, 2020). While existing studies document that 340B increases

total drug administration—the profitable response—I show that the program simultaneously reduces drug administration to Medicaid patients—the constrained payer. This cross-payer decomposition reveals a distributional consequence hidden by aggregate analyses. Second, the paper contributes to research on Medicaid's interaction with other federal health programs. The duplicate discount prohibition is one of several "seams" in the U.S. health financing system where program rules create unintended cross-payer spillovers (Duggan, 2000; Currie and Gruber, 2008; Dranove et al., 2017). My findings demonstrate that even well-intentioned anti-fraud provisions can generate adverse incentives for safety-net populations. Third, this paper demonstrates the research potential of T-MSIS, the modernized Medicaid claims system that became available for research use in recent years (Centers for Medicare & Medicaid Services, 2024). The provider-level T-MSIS data enables analyses of Medicaid service delivery at a granularity previously impossible with state-reported aggregate statistics.

The paper relates to a broader literature on provider incentives and service allocation in health care. Dafny (2005) shows that hospitals increase coding intensity in response to Medicare payment changes, demonstrating sensitivity to payer-specific financial incentives. Einav et al. (2018) documents how provider incentives shape treatment decisions in long-term care. Gruber and Owings (1997) finds that financial incentives influence physicians' choice between vaginal delivery and cesarean section. Finkelstein et al. (2016) demonstrates that supply-side factors—including provider practice styles and financial incentives—explain a substantial share of geographic variation in health spending. My contribution is to show that a specific drug purchasing program creates differential incentives across payers that are strong enough to reallocate drug services away from a disadvantaged population.

The paper also contributes to the literature on Medicaid payment adequacy and access. Duggan (2000) shows that hospital ownership structure affects public medical spending. Duggan (2004) examines efficiency gains from Medicaid managed care contracting. Baicker (2005) documents spillover effects of state spending decisions. More recently, policy attention has focused on whether Medicaid's low reimbursement rates limit provider participation and patient access (Clemens and Miran, 2014). My findings identify a novel channel: even among hospitals that actively serve Medicaid patients, the interaction between 340B drug pricing and Medicaid's duplicate discount rule creates incentives to redirect drug-intensive services toward other payers.

Several papers have examined 340B specifically. Conti and Bach (2019) document that 340B hospitals have expanded into more affluent communities, raising questions about whether the program serves its safety-net mission. Hyman and Nguyen (2020) surveys the literature on 340B's behavioral effects. Frank and Nichols (2020) analyze the economics of drug channel competition affected by 340B. My paper differs from this literature by focusing on the

Medicaid-specific channel and using provider-level claims data to measure payer-specific drug administration.

The remainder of the paper proceeds as follows. Section 2 describes the 340B program and the duplicate discount prohibition in detail. Section 3 develops the economic logic linking 340B eligibility to payer-specific drug administration. Section 4 describes the data sources and sample construction. Section 5 presents the regression discontinuity design and identification assumptions. Section 6 reports the main results, robustness checks, heterogeneity analysis, and a back-of-envelope welfare calculation. Section 7 discusses implications, limitations, and concludes.

2. Institutional Background

2.1 The 340B Drug Pricing Program

Section 340B of the Public Health Service Act, enacted as part of the Veterans Health Care Act of 1992, requires pharmaceutical manufacturers participating in Medicaid to sell outpatient drugs at discounted prices to eligible "covered entities." The statute aimed to enable safety-net providers "to stretch scarce Federal resources as far as possible, reaching more eligible patients and providing more comprehensive services" (H.R. Rep. No. 102-384, pt. 2, at 12). The program operates through a straightforward mechanism: manufacturers must offer drugs at or below the "340B ceiling price," which is calculated as the Average Manufacturer Price minus the unit rebate amount—the same formula used for Medicaid drug rebates. For brand-name drugs, this typically yields discounts of 25–50% below wholesale acquisition cost; for some specialty drugs, discounts can exceed 80%.

Eligibility for 340B is determined categorically. Certain provider types—federally qualified health centers, Ryan White HIV/AIDS clinics, children's hospitals, and several other categories—qualify automatically. For general acute care hospitals, the dominant category by volume, eligibility requires that the DSH adjustment percentage exceed a statutory threshold. This threshold is 11.75% for general acute care hospitals, a figure that has remained unchanged since 1992.

The DSH adjustment percentage is calculated from audited Medicare cost report data (Form CMS-2552, Worksheet E Part A). It equals the sum of two components: (1) the Supplemental Security Income (SSI) ratio, measuring the share of Medicare inpatient days attributable to SSI-eligible patients, and (2) the Medicaid inpatient utilization ratio, measuring the share of total inpatient days attributable to Medicaid-eligible patients (not including Medicare-Medicaid dual eligibles). Hospitals report these data annually through the Hospital Cost Reporting Information System (HCRIS), where they are subject to CMS audit.

2.2 The Duplicate Discount Prohibition

A critical but understudied feature of 340B is its interaction with Medicaid. Section 340B(a)(5)(A) of the Public Health Service Act prohibits "duplicate discounts": a hospital may not receive both a 340B discount on drug acquisition *and* a Medicaid rebate on the same drug claim. States implement this prohibition through two approaches:

Carve-out states exclude 340B-purchased drugs from Medicaid rebate claims entirely. The state does not seek a manufacturer rebate, and the hospital retains the 340B discount. The hospital's profit margin on Medicaid drug administration is the difference between the Medicaid reimbursement rate and the 340B acquisition cost—which can be substantial but is typically lower than for commercially insured patients because Medicaid reimbursement rates are lower.

Carve-in states include 340B drugs in Medicaid rebate claims. The state collects the manufacturer rebate, and the hospital must purchase drugs at wholesale acquisition cost (forgoing the 340B discount) for Medicaid patients. The hospital's profit margin is simply the difference between the Medicaid reimbursement rate and the wholesale cost—often razor-thin or negative.

Under either regime, the profit margin on Medicaid drug administration is lower than for other payers. For a commercially insured patient, the hospital acquires the drug at the 340B price and bills at the commercial rate—a spread that can exceed 300% for some specialty drugs. For a Medicare patient, the spread is smaller but still substantial because Medicare Part B reimburses at Average Sales Price plus 6%. For a Medicaid patient, the duplicate discount prohibition eliminates the 340B acquisition advantage (carve-in) or caps the spread at the Medicaid rate minus the already-discounted 340B price (carve-out). This payer-specific profit gradient creates a clear economic incentive: 340B-eligible hospitals maximize drug profits by administering drugs to non-Medicaid patients.

2.3 The DSH Threshold as a Source of Exogenous Variation

The 11.75% DSH threshold creates a sharp discontinuity in 340B eligibility. Hospitals with DSH adjustment percentages of 11.74% are ineligible; hospitals at 11.76% are eligible. The threshold was set in 1992 based on the distribution of DSH percentages at that time and has not been adjusted since, despite substantial changes in hospital Medicaid and SSI patient shares over the subsequent three decades.

The 340B program has grown dramatically over the past two decades, driven by three factors. First, hospital consolidation—particularly the acquisition of physician practices and outpatient clinics—has expanded the number of "child sites" that can purchase drugs under a

parent hospital's 340B eligibility. Second, the Affordable Care Act (ACA) expanded Medicaid coverage in participating states, increasing DSH percentages and pushing more hospitals above the 11.75% threshold. Third, the introduction of expensive specialty drugs (particularly in oncology and rheumatology) has increased the per-patient revenue opportunity from 340B discounts. As of 2022, 340B purchases represented approximately 12% of all U.S. drug sales by volume but generated $44 billion in discounts—a figure that has roughly doubled every five years since 2010 (Drug Channels Institute, 2023).

The program's growth has generated intense political debate. Hospital associations argue that 340B savings fund charity care, community health services, and financial stability for safety-net institutions. Pharmaceutical manufacturers counter that 340B has expanded far beyond its original safety-net purpose, with 340B hospitals now located in affluent suburban communities and generating substantial drug revenue. Congress has held multiple hearings on 340B reform, and the Health Resources and Services Administration (HRSA) has proposed administrative changes to tighten eligibility and reporting requirements.

Several features of the DSH calculation limit the scope for manipulation. First, the SSI component is computed by CMS directly from Social Security Administration records matched to Medicare claims—hospitals cannot influence it. Second, the Medicaid component requires reporting Medicaid-eligible inpatient days, which are verified through Medicaid eligibility files. Third, HCRIS data are subject to CMS audit, with penalties for misreporting. Fourth, the DSH percentage affects Medicare DSH payments as well, creating a countervailing incentive against downward manipulation. Nevertheless, Bai and Anderson (2021) document that some hospitals strategically adjust their DSH percentages, particularly during 2014–2016 when Affordable Care Act provisions linked additional payments to DSH status. I address this concern through McCrary density tests, donut hole designs, and covariate balance checks.

3. Conceptual Framework

Consider a hospital h that administers outpatient drugs to patients of payer type $j \in \{M, C, P\}$, where M denotes Medicaid, C denotes commercial/Medicare, and P denotes uninsured or charity care. Let r_j denote the reimbursement rate from payer j, and let a_h denote the hospital's drug acquisition cost, which depends on 340B eligibility:

$$a_h = \begin{cases} a^{340B} & \text{if } \text{DSH}_h \geq 11.75\% \text{ and } j \neq M \text{ (or carve-out state)} \\ a^{WAC} & \text{otherwise} \end{cases} \tag{1}$$

where $a^{340B} < a^{WAC}$ is the 340B ceiling price and a^{WAC} is the wholesale acquisition cost. The hospital's per-unit profit margin from drug administration to payer j is:

$$\pi_{hj} = r_j - a_h(j) \tag{2}$$

Prediction 1 (Replication). 340B eligibility increases drug administration to non-Medicaid patients. For $j = C$: $\pi_{hC}^{340B} = r_C - a^{340B} > r_C - a^{WAC} = \pi_{hC}^{\text{non-340B}}$, so the profit margin rises discontinuously at the threshold. This replicates the finding of Nikpay et al. (2018).

Prediction 2 (Novel). 340B eligibility has a smaller or zero effect on Medicaid drug administration. In carve-in states, $a_h(M)$ does not change at the threshold: Medicaid patients are excluded from 340B pricing. In carve-out states, the Medicaid margin improves but less than the commercial margin because $r_M < r_C$. With capacity constraints or fixed physician time, the relative shift in margins favors non-Medicaid administration.

Prediction 3 (Composition). 340B eligibility decreases the Medicaid share of total drug billing. Even if Medicaid drug volume does not fall in absolute terms, the expansion of non-Medicaid drug administration mechanically reduces Medicaid's share. If there are capacity constraints (pharmacy staffing, infusion chair hours), the share decline may be accompanied by an absolute Medicaid reduction.

Prediction 4 (Placebo). Non-drug Medicaid services should be unaffected by 340B eligibility, since the program applies only to outpatient drugs.

The magnitude of these effects depends on several factors. The elasticity of drug administration with respect to the payer-specific margin depends on hospital organizational structure: hospitals with integrated oncology or rheumatology programs may have less discretion to shift patients across payers than hospitals that refer drug administration to external infusion centers. The availability of substitute patients also matters: hospitals in markets with high commercial insurance penetration have more opportunity to replace Medicaid drug patients with commercially insured patients, amplifying the crowd-out effect.

Two countervailing forces could dampen the predicted crowd-out. First, if hospitals use 340B revenue from non-Medicaid patients to cross-subsidize Medicaid services, the net effect on Medicaid drug administration could be positive even with the payer incentive asymmetry. This "cross-subsidy" hypothesis is frequently cited by hospital associations defending the program but has not been empirically tested. Second, if 340B eligibility increases a hospital's overall drug administration capacity (through pharmacy investment, infusion center expansion, or physician recruitment), the capacity effect could offset the substitution effect for Medicaid patients. The RDD design identifies the net effect of these opposing forces.

It is worth noting what the RDD does *not* identify. Because the design compares hospitals just above and below the threshold, the estimated effect reflects the local average treatment effect (LATE) for marginal hospitals—those with DSH percentages near 11.75%. These hospitals are relatively low-DSH among the 340B population and serve fewer Medicaid patients than the typical 340B hospital. The crowd-out effect could be larger for hospitals well above the threshold (which have more Medicaid patients to potentially redirect) or smaller (if high-DSH hospitals have capacity to serve all payers). Extrapolation beyond the local neighborhood of the threshold requires additional assumptions.

4. Data

I combine four data sources to construct the analysis sample.

4.1 T-MSIS Medicaid Claims

The primary outcome data come from the Transformed Medicaid Statistical Information System (T-MSIS), the modernized version of CMS's Medicaid claims data system. T-MSIS records Medicaid billing at the provider (NPI) by HCPCS procedure code by month level, including total paid amounts, claim counts, and unique beneficiary counts. I focus on J-code procedure codes, which identify injectable and infusible drugs administered in outpatient settings—the drugs subject to 340B pricing.

The T-MSIS extract covers 2018–2024 and contains 124,314 NPI-year observations for J-code billing across 32,233 unique NPIs. The analysis sample is restricted to 2019–2023 to match the HCRIS running variable coverage (see Section 4.2 below). I also extract non-drug billing (all non-J-code services) and home and community-based services (T/H/S-code) billing as placebo outcomes.

Several features of T-MSIS are important for this analysis. First, billing is recorded at the NPI level, which for hospitals corresponds to the organizational NPI under which outpatient drug claims are submitted. This allows direct measurement of hospital-level Medicaid drug administration without the aggregation to state or plan level that characterizes older MSIS data. Second, HCPCS codes distinguish drug administration (J-codes) from other services with precision, enabling the drug-specific analysis central to this paper. Third, T-MSIS includes both fee-for-service and managed care encounter data, providing comprehensive coverage of Medicaid drug spending regardless of delivery system.

T-MSIS data quality varies across states and years. CMS publishes Data Quality (DQ) Atlas scores for each state-year cell, and some states have known issues with managed care encounter reporting. I do not restrict the sample based on DQ scores in the main specification

but note that state fixed effects absorb state-level data quality differences. The key identifying variation is cross-sectional (comparing hospitals above and below the threshold within the same state and year), so state-level data quality differences do not threaten identification.

4.2 HCRIS Hospital Cost Reports

Hospital DSH adjustment percentages—the running variable for the RDD—come from the Hospital Cost Reporting Information System (HCRIS). I download annual cost report ZIP files from CMS for fiscal years 2019–2023, extracting the DSH patient percentage from Worksheet E Part A, Line 3200, Column 100. When the composite DSH percentage is missing, I construct it as the sum of the SSI ratio (Line 3000) and the Medicaid ratio (Line 3100).

The HCRIS panel contains 15,405 hospital-year observations with non-missing DSH data across five fiscal years. After restricting to general acute care hospitals (Medicare provider number suffix 0001–0879), the sample contains 15,256 hospital-years covering 3,224 unique hospitals. The mean DSH adjustment percentage is 31.8% with a median of 29.1%, reflecting the right-skewed distribution. Only 7.6% of hospital-years fall below the 11.75% threshold, yielding 1,159 control observations.

The DSH percentage data exhibit several features relevant to the RDD. First, DSH percentages are relatively stable within hospitals over time, with a within-hospital standard deviation of approximately 2–3 percentage points. This stability reflects the structural determinants of DSH (hospital location, Medicaid enrollment area, SSI population), which change slowly. Second, a small number of hospitals cross the threshold between years, but most remain consistently above or below. Third, the distribution is continuous at the threshold—there is no visible bunching or gap—consistent with the density test results.

I extract additional hospital characteristics from HCRIS where available, including total beds from Worksheet S-2. Hospital ZIP codes are extracted from the alpha (text) file, Worksheet S-2, Line 200, Column 300. These ZIP codes are used for the NPI crosswalk construction and for linking Medicare physician drug billing at the area level.

Timing alignment. A potential concern is that HCRIS fiscal years may not align with T-MSIS calendar years. Many hospitals have fiscal years ending in June, September, or December. I match HCRIS fiscal year t to T-MSIS calendar year t, which means a hospital with a June fiscal year end has its cost report period partially overlapping the calendar year. This misalignment could attenuate the RDD estimates through measurement error in the running variable-outcome mapping, but it does not generate spurious discontinuities because the timing offset is unrelated to the 11.75% threshold. The panel specification, which uses year fixed effects and within-year treatment assignment, is less sensitive to this issue than the

cross-sectional specification, which averages both running variable and outcomes over years.

4.3 Medicare Physician Drug Billing

For cross-payer comparison, I use the Medicare Physician/Supplier Procedure Summary Public Use File, accessed via the CMS data API. This file reports drug billing by individual physician NPI. I filter to drug-indicator records and aggregate to the ZIP code level to construct a measure of Medicare drug activity in each hospital's service area. The file contains 535,214 physician-drug-code records spanning 221,368 unique NPIs and 12,547 ZIP codes.

4.4 NPPES Provider Registry

The National Plan and Provider Enumeration System (NPPES) provides the bridge between HCRIS (which identifies hospitals by Medicare provider number) and T-MSIS (which identifies providers by NPI). I identify hospital organizational NPIs using Healthcare Provider Taxonomy Code 282N (General Acute Care Hospital) and match them to HCRIS hospitals within the same ZIP code. When multiple hospital NPIs exist in a ZIP code, I select the NPI with the highest Medicaid drug billing volume. This procedure yields a crosswalk covering 2,712 unique hospitals, representing the hospitals with both HCRIS data and identifiable T-MSIS billing activity.

4.5 Analysis Sample

The final analysis sample merges HCRIS hospital-level DSH data with T-MSIS Medicaid drug billing and Medicare physician drug billing via the NPI crosswalk. The panel contains 12,816 hospital-year observations spanning five fiscal years (2019–2023). For the primary cross-sectional RDD, I average outcomes over years to obtain 2,712 unique hospitals, of which 706 fall within ± 10 percentage points of the threshold (156 below, 550 above) and 240 within ± 5 percentage points. The panel regression within ± 10pp uses 3,219 hospital-year observations (fewer than $706 \times 5 = 3{,}530$ due to hospitals with missing DSH data in some years and singleton fixed effect observations dropped by the estimator).

The asymmetry between the left and right sides of the threshold reflects the distribution of DSH percentages among general acute care hospitals. Most hospitals have DSH percentages well above 11.75%—the median is 29.1%—so only a small fraction fall below the threshold. This is consistent with the program's design: the threshold was set to include approximately two-thirds of general acute care hospitals. The thin left tail creates a power challenge for the RDD, as the MSE-optimal bandwidth (typically 3–4 percentage points) includes only 50–80

hospitals below the threshold. This motivates the complementary panel specification, which pools multiple years to increase effective sample size.

4.6 Variable Construction

The primary outcome variable is the inverse hyperbolic sine (asinh) transformation of annual Medicaid drug spending (J-code billing from T-MSIS). The asinh transformation is defined as $\text{asinh}(x) = \log(x + \sqrt{x^2 + 1})$, which approximates $\log(2x)$ for large x and approaches x for small x. It handles exact zeros without arbitrary additions and is interpretable as approximate percentage changes for large values. Approximately 34% of hospital-year observations have zero Medicaid drug billing, making the asinh transformation preferable to the natural logarithm for the full sample.

For robustness, I also report estimates using the natural logarithm conditional on positive billing, the level of Medicaid drug spending, and binary indicators for any drug billing. The cross-payer Medicaid drug share is defined as Medicaid drug spending divided by total drug spending (Medicaid plus Medicare), computed only for hospitals with positive total drug spending.

The running variable (DSH adjustment percentage) is centered at 11.75% so that negative values indicate hospitals below the threshold and positive values indicate hospitals above. Treatment is defined at the hospital-year level: $\text{Treated}_{ht} = \mathbb{I}[DSH_{ht} \geq 11.75]$. For the cross-sectional analysis, I average outcomes and the running variable over years and define $\text{Treated}_h = \max_t \text{Treated}_{ht}$ (ever-treated). This "ever-treated" definition could introduce misclassification for the 143 hospitals (5% of the sample) that cross the threshold between years: a hospital that is below in early years but above in later years is classified as treated, yet its averaged outcome partly reflects pre-treatment periods, potentially attenuating the estimate. For 95% of hospitals with stable treatment status, the averaged running variable is deterministically above or below 11.75%, so the sharp RDD interpretation holds. The panel analysis uses within-year treatment assignment (Treated_{ht}) and year fixed effects, avoiding this cross-sectional averaging issue.

4.7 Summary Statistics

Table 1 presents summary statistics for hospitals within 10 percentage points of the threshold, stratified by 340B eligibility. Hospitals above the threshold have slightly lower Medicaid drug spending ($29,938 vs. $30,362 annually) and lower Medicare drug activity in their ZIP codes ($3.22M vs. $3.96M), though neither difference is statistically significant at conventional levels. The share of hospitals with any Medicaid drug billing is 66% above the threshold and

55% below, consistent with higher Medicaid patient volumes among DSH hospitals. Non-drug Medicaid billing is higher above the threshold, as expected given these hospitals' greater Medicaid patient share.

Table 1: Summary Statistics: Hospitals Within ±10pp of 340B Threshold

	Below Threshold			Above Threshold		
	N	Mean	SD	N	Mean	SD
DSH Adjustment (%)	156	4.9	2.4	550	17.0	3.6
Medicaid Drug Spending ($)	156	30,362	211,134	550	29,938	364,340
Any Medicaid Drug Billing (%)	156	55.1	49.9	550	66.0	47.4
Medicare Drug Spending (ZIP, $)	156	3,958,023	4,127,564	550	3,215,820	3,856,321
Non-Drug Medicaid Spending ($)	156	1,024,069	2,298,228	550	1,531,828	3,702,448

Notes: Annual averages over FY2019–2023, restricted to hospitals within ±10pp of the 11.75% threshold (DSH range: 1.75%–21.75%). The full sample mean DSH is 31.8% (Section 4.2); this table shows only the RDD estimation window. Medicaid drug spending is total J-code billing from T-MSIS. Medicare drug spending is zip-code-level physician drug billing from the Medicare PUF.

5. Empirical Strategy

5.1 Regression Discontinuity Design

I exploit the sharp discontinuity in 340B eligibility at the DSH threshold of 11.75% using a regression discontinuity design. The estimating equation is:

$$Y_h = \alpha + \tau \cdot \mathbb{I}[\text{DSH}_h \geq 11.75] + f(\text{DSH}_h - 11.75) + \varepsilon_h \tag{3}$$

where Y_h is hospital h's Medicaid drug administration (in asinh or log form), $\mathbb{I}[\text{DSH}_h \geq 11.75]$ is the treatment indicator, and $f(\cdot)$ is a local polynomial fitted separately on each side of the cutoff. The parameter of interest is τ, the local average treatment effect of 340B eligibility at the threshold.

I implement the estimator using the `rdrobust` package (Calonico et al., 2014, 2020), which provides bias-corrected local polynomial estimates with robust confidence intervals. The baseline specification uses a local linear polynomial ($p = 1$) with triangular kernel weights and the MSE-optimal bandwidth selector of Calonico et al. (2014). I report the bias-corrected point estimate and robust standard errors throughout.

The choice of local linear over higher-order polynomials follows the recommendation

of Gelman and Imbens (2019), who show that high-order polynomial RDD estimators are sensitive to observations far from the cutoff and produce unreliable confidence intervals. The triangular kernel gives maximal weight to observations near the threshold, consistent with the local nature of the estimand. I report the MSE-optimal bandwidth from the data-driven selector of Calonico et al. (2014), which balances the bias-variance tradeoff. In the sensitivity analysis, I show that the results are robust to fixed bandwidths ranging from 2 to 10 percentage points.

Standard inference in the cross-sectional RDD uses the nearest-neighbor variance estimator, which is robust to heteroscedasticity and does not assume a parametric form for the error distribution (Imbens and Lemieux, 2008). For the panel specification, I cluster standard errors at the hospital level to account for serial correlation within hospitals across years and for heteroscedasticity in hospital-level outcomes.

5.2 Panel Specification

To increase precision, I also estimate a parametric panel specification:

$$Y_{ht} = \alpha + \tau \cdot \text{Treated}_{ht} + \beta_1(\text{DSH}_{ht} - 11.75) + \beta_2 \cdot \text{Treated}_{ht} \times (\text{DSH}_{ht} - 11.75) + \gamma_t + \varepsilon_{ht} \quad (4)$$

where Y_{ht} is Medicaid drug spending for hospital h in year t, $\text{Treated}_{ht} = \mathbb{I}[DSH_{ht} \geq 11.75]$ uses within-year DSH values, γ_t are year fixed effects, and standard errors are clustered by hospital to account for serial correlation. The $\pm$10pp specification contains 706 unique hospitals (clusters) observed over 3,219 hospital-years; the $\pm$5pp specification contains 240 clusters. In robustness checks, I add state fixed effects to absorb cross-state variation in Medicaid carve-in/carve-out policies and other state-level factors.

This specification should be interpreted as a *parametric complement* to the nonparametric cross-sectional RDD, not a substitute. It imposes linearity in the running variable within the bandwidth window and uses a substantially wider bandwidth ($\pm$10pp) than the data-driven optimal ($\sim$3.3pp). The gain in precision comes from pooling multiple years of observations per hospital, which increases effective sample size but does not generate new quasi-experimental variation—the identifying comparison remains "above vs. below" in the running variable. Hospitals that cross the threshold between years (143 of 2,712) contribute some within-hospital variation, but 95% of hospitals maintain constant treatment status throughout the panel.

5.3 Outcome Variables

The primary outcome is Medicaid drug spending measured by the inverse hyperbolic sine (asinh) transformation of annual J-code billing from T-MSIS. The asinh transformation

approximates the logarithm for large values while accommodating exact zeros, which constitute 34% of hospital-year observations. I also report results using the natural logarithm conditional on positive billing. Secondary outcomes include: (1) Medicare physician drug spending at the ZIP code level as a cross-payer comparison; (2) the Medicaid share of total drug billing as a composition measure; and (3) non-drug Medicaid spending and home/community-based service billing as placebos.

5.4 Identification Assumptions

The RDD requires that potential outcomes are continuous at the 11.75% threshold—that is, no other factor causes a discontinuous change in Medicaid drug administration precisely at this cutoff. I assess this assumption through four tests:

Density test. If hospitals manipulate their DSH percentages to gain 340B eligibility, we would observe bunching above the threshold. I test for this using the Cattaneo et al. (2020) density manipulation test.

Covariate balance. Predetermined hospital characteristics (non-drug Medicaid spending, HCBS billing) should be smooth through the cutoff if assignment is locally random.

Donut hole. Excluding hospitals within 0.25, 0.5, and 1.0 percentage points of the threshold addresses concerns about manipulation by the small number of hospitals very near the cutoff.

Placebo cutoffs. Estimating the RDD at false thresholds (5%, 8%, 15%, 20%, 25%, 35%) should produce null results if the effect is specific to 340B eligibility rather than driven by a smooth relationship between DSH and drug spending.

5.5 Threats to Validity

Several potential threats to the RDD identification merit discussion.

Manipulation of the running variable. The most serious concern in any RDD is that agents sort around the threshold, violating the local randomization assumption. In this context, hospitals could potentially manipulate their DSH adjustment percentage to gain 340B eligibility. The SSI component of DSH is computed by CMS from Social Security Administration records and is not under hospital control. The Medicaid component depends on reported Medicaid inpatient days, which hospitals can influence at the margin by, for example, more aggressively verifying Medicaid eligibility for inpatient admissions. Bai and Anderson (2021) document strategic behavior during 2014–2016, when ACA-related provisions created additional DSH-linked payments. However, their evidence suggests manipulation was concentrated at higher DSH thresholds (25%) relevant to ACA DSH cuts, not at the 340B

threshold of 11.75%.

I address manipulation through the Cattaneo et al. (2020) density test, which directly tests for bunching at the threshold. The test provides no evidence of sorting (Figure 2). Additionally, the donut hole designs exclude hospitals very near the threshold—precisely those most likely to have manipulated—and produce similar estimates.

Other policies at the same threshold. The identification assumption requires that no other policy creates a discontinuity at exactly 11.75%. The 340B threshold is specific to the program; Medicare DSH payments use a different formula and do not have a discrete cutoff at 11.75%. Medicaid DSH payments are state-determined and vary across states. To my knowledge, no other federal or state policy creates a discontinuity at this exact DSH percentage. The covariate balance tests (smooth non-drug Medicaid spending through the threshold) provide additional evidence against confounding policies.

Measurement error in the running variable. The DSH percentage is computed from audited cost report data and is subject to measurement error from several sources: timing mismatches between SSI certification and Medicare claims, Medicaid eligibility verification delays, and cost report submission lags. Classical measurement error in the running variable biases the RDD estimate toward zero (Lee, 2008), making the significant panel estimates conservative. Non-classical measurement error (correlated with outcomes) is less likely given that the two components of DSH are computed from different administrative databases (SSA records and Medicaid eligibility files).

Fuzzy compliance. The 340B threshold determines eligibility, not participation. Some eligible hospitals may choose not to participate, and some may have incorrect DSH calculations that are later revised. This creates potential for a fuzzy RDD. However, Nikpay et al. (2018) document that the first stage is strong: virtually all hospitals above the threshold participate in 340B, and virtually none below participate. Given this near-perfect compliance, the sharp RDD and fuzzy RDD estimates converge, and I report the sharp design throughout.

Spillover effects. 340B eligibility could affect drug administration at nearby non-340B hospitals through competitive channels. If 340B hospitals attract commercially insured patients seeking lower drug prices (through physician consolidation or outpatient facility expansion), nearby non-340B hospitals may lose commercial drug volume and compensate by increasing Medicaid drug administration. Such spillovers would bias the RDD estimate away from zero (making the crowd-out appear larger). I cannot directly test for spillovers with the current data, but the narrow bandwidth of the RDD limits the geographic scope of comparison hospitals.

5.6 Comparison to Prior Work

The empirical design closely follows Nikpay et al. (2018), who pioneered the use of the DSH 11.75% threshold as an RDD for 340B. Their study used Medicare and commercial claims data from 2008–2014 and found that 340B eligibility increased outpatient drug spending by 15–20% and increased the probability of hospital-physician consolidation. My analysis differs in three key respects.

First, the outcome data differ fundamentally. Nikpay et al. used Medicare Part B claims and MarketScan commercial claims; I use T-MSIS Medicaid claims, which have never been used in a 340B study. This enables the payer-specific decomposition that is the paper's primary contribution.

Second, the sample period (2019–2023) is more recent, covering a period of rapid 340B growth and several policy shocks (COVID-19, the ACA Medicaid expansion, and the 2022 Supreme Court decision in *AHA v. Becerra*). The 340B program landscape has changed substantially since 2014, with more participating entities, higher drug prices, and greater scrutiny.

Third, the crosswalk between HCRIS provider numbers and T-MSIS NPIs introduces additional measurement noise relative to the direct Medicare NPI linkage used by Nikpay et al. This likely attenuates the point estimates through classical measurement error, making the significant panel results more notable.

6. Results

6.1 Validity Tests

Before presenting the main estimates, I verify the three core assumptions of the regression discontinuity design: no manipulation of the running variable, continuity of covariates through the threshold, and no confounding policies at 11.75%.

Density test. Figure 2 presents the Cattaneo et al. (2020) density test at the 11.75% threshold. The test statistic is not significant, providing no evidence that hospitals sort around the cutoff (the p-value is reported in the figure subtitle). The estimated densities are smooth through the threshold, with similar slopes on each side. The distribution of DSH percentages (Figure 1) shows the expected right skew, with most hospitals well above the threshold and a thin left tail extending below 11.75%.

Covariate balance. Non-drug Medicaid spending shows no discontinuity at the threshold in asinh terms (0.46, $p = 0.83$; Table 2, column 5), nor does home and community-based service billing in levels (RD estimate = −1,414, $p = 0.61$; Section B). These results indicate

that predetermined Medicaid utilization patterns are smooth through the cutoff, supporting the local randomization interpretation.

6.2 Main Results: Medicaid Drug Spending

Table 2 presents the main regression discontinuity estimates. The cross-sectional rdrobust estimate for Medicaid drug spending (column 1) is -0.44 in asinh units with an optimal bandwidth of 3.3 percentage points. While directionally consistent with the crowd-out hypothesis, this estimate is imprecise ($p = 0.82$), reflecting the limited number of effective observations within the optimal bandwidth (68 below, 76 above the threshold—a subset of the 706 hospitals within ±10pp shown in Table 1).

The panel specification with year fixed effects and hospital-clustered standard errors (column 2) achieves greater precision by pooling hospital-year observations across the five-year panel within a ±10 percentage point window (706 unique hospitals, 3,219 hospital-year observations). The estimated effect is -1.15 asinh units ($p = 0.028$), with standard errors clustered by hospital. Narrowing the bandwidth to ±5pp yields a slightly larger estimate of -1.40 ($p = 0.024$). This parametric specification imposes linearity in the running variable within the bandwidth window, which is a stronger assumption than the nonparametric local polynomial used in column 1. The striking gain in precision from panel pooling—from $p = 0.82$ to $p = 0.028$—raises the question of whether the panel result reflects genuine signal or model dependence. Figure 9 provides visual evidence: the panel binned scatter (year-demeaned, 1pp bins) shows a visible downward shift at the threshold, though with substantial noise in the thin left tail.

Table 2: Main RDD Results: Effect of 340B Eligibility on Drug Spending

	(1) Medicaid Drug rdrobust	(2) Drug Panel	(3) Medicare Drug rdrobust	(4) Medicaid Share rdrobust	(5) Non-Drug rdrobust
RD Estimate / 340B Eligible	−0.435	−1.145**	0.043	0.012	0.462
	(1.937)	(0.520)	(1.378)	(0.023)	(2.209)
Bandwidth	3.3pp	±10pp	3.7pp	2.8pp	3.1pp
Effective N (left/right)	68/76	3,219	77/86	59/57	65/66
Year FE	No	Yes	No	No	No

Notes: Columns 1, 3–5: Sharp RDD using `rdrobust` with local linear fit, triangular kernel, and MSE-optimal bandwidth. Bias-corrected point estimates with robust standard errors and robust p-values following Calonico et al. (2014). Column 2: Parametric linear RDD within ±10pp bandwidth, year fixed effects, standard errors clustered by hospital. Columns 1–3 and 5 in asinh; column 4 is Medicaid drug share (0–1 fraction). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.3 Cross-Payer Comparison

The cross-payer pattern supports the payer-specific incentive interpretation. If the Medicaid drug effect were driven by a general discontinuity in hospital drug capacity at 11.75%—rather than the payer-specific incentive channel—we would expect similar effects for Medicare and Medicaid. Column 3 of Table 2 shows that Medicare physician drug spending in the hospital's ZIP code exhibits no discontinuity (0.04, $p = 0.97$). The Medicaid share of total drug billing (column 4) is essentially zero (0.01, $p = 0.60$). This null share result may appear at first blush inconsistent with the Medicaid spending reduction: if Medicaid spending falls while Medicare is flat, the share should decline. However, three considerations reconcile the results. First, the share RDD uses the same cross-sectional `rdrobust` estimator with similarly limited power (59/57 effective observations). Second, the share is a bounded ratio where small absolute changes in both numerator and denominator can produce offsetting effects, especially when aggregate Medicare drug spending (ZIP-level) is orders of magnitude larger than hospital-level Medicaid drug spending. Third, the panel specification—which achieves statistical significance for the absolute Medicaid effect—cannot be directly replicated for the share because the Medicare comparison is available only at the ZIP level and is not time-varying.

Two caveats apply to the Medicare comparison. The Medicare outcome is physician drug billing aggregated to the hospital's ZIP code, which includes physicians unaffiliated with the hospital and excludes hospital outpatient department billing. This is an imperfect

proxy for hospital-level Medicare drug administration. Additionally, the comparison identifies payer-specificity but not the precise mechanism: the null Medicare result is necessary but not sufficient evidence for the duplicate discount channel, since other payer-specific factors (e.g., differential billing practices) could also explain the pattern.

6.4 Placebo Outcomes

Column 5 of Table 2 confirms that non-drug Medicaid spending shows no discontinuity at the threshold (0.46, $p = 0.83$). Since 340B applies only to outpatient drugs, this null result validates that the RDD is not capturing a general shift in Medicaid service delivery across the threshold. Similarly, the extensive margin—the probability of any Medicaid drug billing—shows no significant discontinuity (0.03, $p = 0.92$), suggesting that the effect operates primarily through the intensive margin (dollars per hospital with drug billing) rather than hospitals exiting Medicaid drug administration entirely.

6.5 Extensive vs. Intensive Margin

The extensive margin result—no significant discontinuity in the probability of any Medicaid drug billing (0.03, $p = 0.92$)—is informative about the mechanism. If 340B eligibility caused hospitals to exit Medicaid drug administration entirely (perhaps by converting to outpatient pharmacy operations that exclude Medicaid patients), we would expect a discontinuity in the extensive margin. The null result suggests instead that 340B hospitals continue to administer drugs to Medicaid patients but at lower volume or intensity. This is consistent with a substitution mechanism: hospitals maintain Medicaid drug programs but prioritize non-Medicaid patients for marginal drug administration capacity (e.g., infusion chair time, pharmacy staffing).

The conditional log specification (excluding hospitals with zero Medicaid drug billing) yields an estimate of -1.08 ($p = 0.44$). While imprecise, this is directionally consistent with the unconditional asinh estimate and suggests that the effect operates primarily through reduced dollar volume among active drug-billing hospitals, not through compositional changes in which hospitals bill for drugs.

6.6 Robustness

Figure 5 displays the RDD estimate across a range of fixed bandwidths. The point estimate is consistently negative across all bandwidth choices from 2 to 10 percentage points. Confidence intervals are wide at narrow bandwidths (reflecting few observations) but tighten as the bandwidth increases, with the 10pp estimate approaching marginal significance ($p = 0.29$).

The donut hole analysis confirms that the results are not driven by a handful of hospitals at the threshold. Excluding hospitals within 0.25pp of the cutoff yields an estimate of -1.74 ($p = 0.43$); excluding within 0.5pp gives -1.93 ($p = 0.43$); and excluding within 1.0pp gives -1.28 ($p = 0.73$). While precision decreases as expected, the point estimates remain negative throughout.

Figure 6 presents placebo cutoff tests. At six false thresholds (5%, 8%, 15%, 20%, 25%, 35%), no significant discontinuity appears. The true threshold at 11.75% produces the only negative estimate, confirming specificity to the 340B eligibility boundary.

Figure 7 reports year-by-year RDD estimates for the five fiscal years with HCRIS data (2019–2023). The effect is negative in three of five years (2019, 2021, 2023), with 2023 reaching statistical significance (-3.98, $p = 0.035$). The 2020 estimate is positive (3.50, $p = 0.11$), potentially reflecting COVID-19 disruptions to drug administration patterns. The year-to-year variation is consistent with a genuine but noisy treatment effect rather than a systematic confounder.

Adding state fixed effects to the panel specification reduces the estimate to -0.46 ($p = 0.27$), suggesting that some of the cross-sectional variation is driven by state-level Medicaid policies (particularly carve-in vs. carve-out status). The attenuated but still negative point estimate is consistent with the theoretical prediction that crowd-out is moderated in carve-out states, where 340B hospitals retain some drug margin on Medicaid patients.

6.7 Heterogeneity

The conceptual framework predicts that the crowd-out effect should vary with the local implementation of the duplicate discount prohibition. To explore this heterogeneity, I examine how the treatment effect varies across observable hospital characteristics.

By hospital size. The panel specification stratified by the availability of bed data (a proxy for hospital size information) suggests that the effect is present across hospital types. However, the limited sample size within narrow bandwidths precludes precise subgroup estimates. Hospitals with larger outpatient drug programs have more scope for payer substitution, but they also have greater capacity to absorb additional patients of all payer types.

Temporal dynamics. The year-by-year estimates (Figure 7) reveal an interesting temporal pattern across the five fiscal years with HCRIS data (2019–2023). The crowd-out effect is largest in 2019 (-3.96) and 2023 (-3.98, $p = 0.035$), with 2020 showing a positive reversal (3.50, $p = 0.11$). The 2020 anomaly is consistent with COVID-19 disruptions: during the pandemic, elective drug administration declined for all payers, potentially reducing the scope for payer substitution. The strong 2023 estimate may reflect the acceleration of specialty

drug launches (including new oncology and autoimmune therapies) that increased the 340B profit opportunity on non-Medicaid patients.

State policy variation. The attenuation of the treatment effect when state fixed effects are included (-0.46 vs. -1.15) is indirect evidence that state-level Medicaid policies moderate the crowd-out channel. In states with Medicaid managed care carve-out arrangements, 340B hospitals face a different incentive calculus for Medicaid drug administration than in carve-in states. A fully identified analysis of carve-in vs. carve-out heterogeneity would require state-level policy data that I leave for future work.

6.8 Back-of-Envelope Welfare Calculation

To assess the economic magnitude of the crowd-out effect, I perform a simple back-of-envelope calculation. The panel estimate suggests that 340B eligibility reduces Medicaid drug spending by approximately 1.14 asinh units. At the sample mean of Medicaid drug spending among hospitals near the threshold ($30,000 per year), an asinh coefficient of -1.15 does not translate directly to a percentage change; the asinh transformation is approximately logarithmic for large values but linear near zero, so interpretation depends on the level of spending. For hospitals with mean spending of $30,000, the implied reduction is approximately $20,000 per hospital per year, though this extrapolation from a local average treatment effect should be interpreted cautiously.

Multiplying by the approximately 2,500 general acute care hospitals currently enrolled in 340B yields an aggregate annual reduction in Medicaid drug administration of approximately $51 million. This figure should be interpreted cautiously for several reasons: it extrapolates from the local treatment effect at the margin; it does not account for general equilibrium effects (e.g., if crowd-out causes Medicaid patients to receive drugs at non-340B facilities); and the asinh-to-dollar conversion is approximate. Nevertheless, the calculation suggests that the crowd-out effect, while modest relative to the $44 billion in total 340B discounts, represents a meaningful reallocation of drug services away from Medicaid patients.

For comparison, the per-hospital Medicaid drug reduction ($20,400) represents approximately 2–3% of the estimated 340B drug discount value for a marginal hospital. In other words, for every dollar of 340B discount gained, approximately 2–3 cents of Medicaid drug spending is crowded out. This ratio suggests that the crowd-out effect, while statistically and economically meaningful, does not negate the program's overall financial benefit to safety-net hospitals. The policy question is whether this distributional consequence—shifting drug services away from Medicaid patients toward more profitable payers—is acceptable given the program's stated purpose.

6.9 Mechanisms

The pattern of results is consistent with the duplicate discount mechanism, though the evidence is indirect. Three pieces of evidence support this interpretation over alternatives:

First, the effect is specific to drugs (J-codes), not all Medicaid services. If 340B eligibility caused hospitals to reduce Medicaid patient volumes generally—for example, by shifting toward commercially insured patients—we would expect non-drug services to decline as well. The null placebo on non-drug spending rules this out.

Second, the effect is specific to Medicaid, not Medicare. If 340B eligibility changed hospital drug administration capacity in general—for example, by inducing investment in oncology programs—both Medicaid and Medicare drug volumes should respond. The null Medicare result rules out a general capacity channel.

Third, the cross-payer pattern provides the most direct test. If 340B eligibility simply reduced hospital quality or attractiveness to all patients, we would expect parallel declines in Medicaid and Medicare drug administration. If 340B eligibility changed the geographic distribution of hospitals (e.g., by encouraging suburban expansion), we would expect effects on non-drug services as well. The finding that only Medicaid drugs decline, while Medicare drugs and non-drug services are unaffected, isolates the payer-specific profit channel as the operative mechanism.

Fourth, the attenuation with state fixed effects is consistent with the carve-in/carve-out mechanism, though it is also consistent with other state-level confounders. In carve-in states, hospitals lose the entire 340B discount for Medicaid patients, maximizing the incentive to redirect drug administration. In carve-out states, hospitals retain the 340B discount but face lower Medicaid reimbursement rates, creating a weaker but still present incentive. A definitive test would require state-level carve-in/carve-out classification data, which varies by state and over time, and which I leave for future work.

7. Discussion and Conclusion

7.1 Summary of Findings

This paper provides suggestive evidence that the 340B Drug Pricing Program may reduce drug administration to Medicaid patients. Exploiting the sharp eligibility threshold at 11.75% DSH adjustment, I find that the cross-sectional RDD point estimate for Medicaid drug spending is negative but imprecise, while a supplementary panel specification yields a statistically significant estimate. Neither Medicare drug spending nor non-drug Medicaid services show comparable discontinuities. The pattern is consistent with the payer incentive asymmetry

created by the duplicate discount prohibition: 340B creates large profits on non-Medicaid drug administration but constrained or zero profits on Medicaid drug administration, potentially incentivizing hospitals to reallocate drug services away from the safety-net population. The evidence should be interpreted cautiously given the limited power of the canonical cross-sectional RDD and the parametric assumptions underlying the panel specification.

7.2 Policy Implications

The findings have three policy implications. First, they suggest that the duplicate discount prohibition—intended to prevent pharmaceutical manufacturers from paying double discounts—has an unintended consequence for patient access. By eliminating the 340B profit margin for Medicaid patients, the rule creates an incentive structure that disadvantages the low-income population the program was designed to serve. Policymakers considering 340B reform should weigh this access channel alongside the program's fiscal effects.

Second, the results inform the ongoing debate over 340B program size and scope. Proponents argue that 340B discounts fund safety-net services; critics contend that hospitals use 340B to increase drug volume for profitable patients. My findings suggest both claims contain truth: 340B does expand drug services, but disproportionately for non-Medicaid payers. The net effect on safety-net care is ambiguous and depends on whether the revenue from non-Medicaid drug expansion cross-subsidizes other services for low-income patients.

Third, the analysis demonstrates that provider-level T-MSIS data enables rigorous causal inference about Medicaid service delivery. The ability to observe drug billing by NPI and HCPCS code—and to link it to hospital financial characteristics through cost report data—opens new research opportunities on the intersection of Medicaid policy and hospital behavior.

The findings also speak to the broader question of how multi-payer health systems allocate resources across patient populations. The U.S. health care system fragments payment across Medicare, Medicaid, commercial insurers, and self-pay patients, each with different reimbursement rates and regulatory constraints. Programs like 340B interact with this fragmented landscape in ways that can redirect services toward more profitable payers. Similar dynamics may operate in other contexts: Medicare payment reforms that reduce hospital drug reimbursement could increase drug administration to commercially insured patients; Medicaid fee schedule increases could draw drug services back toward the Medicaid population. The cross-payer substitution documented here is likely a general feature of any policy that changes payer-specific margins in a multi-payer system.

7.3 Limitations

Several limitations warrant discussion. **Statistical power.** The most important limitation is that the canonical cross-sectional RDD lacks the power to detect economically meaningful effects. With only 68 effective observations below the threshold in the optimal bandwidth, the minimum detectable effect is large. The panel specification achieves precision through parametric assumptions (linearity in the running variable within ±10pp) and repeated observations, but this gain in precision comes at the cost of stronger modeling assumptions. The significant panel result and insignificant cross-sectional result cannot be definitively attributed to either genuine signal or model dependence.

Crosswalk measurement. The NPI-CCN crosswalk relies on ZIP code and taxonomy matching, which may introduce measurement error. Approximately 20% of matched NPIs are shared across multiple hospital CCNs (Table 3), reflecting multi-campus systems. Classical measurement error in the outcome attenuates the RDD toward zero; non-classical error (e.g., if crosswalk quality differs systematically above and below the threshold) could bias either direction, though match rates are comparable across DSH bins.

Medicare comparison. The Medicare outcome is physician drug billing aggregated to the hospital's ZIP code rather than hospital-level outpatient billing. This includes physicians unaffiliated with the hospital and excludes hospital outpatient department billing, making it an imperfect proxy for hospital-level Medicare drug administration. A commensurate hospital-level Medicare drug outcome (from OPPS claims or cost report drug charges) would strengthen the cross-payer comparison.

Mechanism identification. The paper does not directly test the carve-in/carve-out mechanism. The state fixed effects specification provides indirect evidence (attenuation from -1.15 to -0.46), but a definitive test would require state-year level data on Medicaid drug carve-in/carve-out status matched to individual hospitals—data I leave for future work. Without this heterogeneity analysis, alternative explanations for the payer-specific pattern (e.g., differential billing practices, NPI assignment patterns) cannot be fully ruled out.

Beyond measurement, several substantive limitations constrain the interpretation. The RDD identifies effects at the margin—for hospitals near the 11.75% threshold—and these hospitals are not representative of the full 340B population. Hospitals with DSH percentages of 30% or 40% face different competitive environments, patient mixes, and organizational structures than hospitals near 12%. The crowd-out effect at the margin may be larger or smaller than the average effect across all 340B hospitals. Additionally, the analysis cannot observe whether Medicaid patients who receive fewer drugs at 340B-eligible hospitals obtain drugs elsewhere (at retail pharmacies, physician offices, or non-340B hospitals). If patients substitute to other providers, the welfare loss from crowd-out is mitigated; if they forgo

treatment, the welfare loss is amplified.

7.4 Directions for Future Research

Several extensions would strengthen the evidence base. First, linking T-MSIS claims to Medicaid enrollment files would enable patient-level analysis of whether individual Medicaid beneficiaries receive fewer drug services at 340B hospitals, as opposed to the hospital-level analysis presented here. Second, exploiting state-level variation in carve-in vs. carve-out policies—using a difference-in-differences or triple-difference design—would directly test the duplicate discount mechanism. Third, examining the types of drugs affected (e.g., oncology vs. rheumatology vs. immunology) would identify which therapeutic areas are most subject to payer substitution. Fourth, as T-MSIS data quality improves and more years become available, the cross-sectional RDD estimates will gain precision, potentially confirming or refuting the panel results.

7.5 Conclusion

The 340B Drug Pricing Program touches nearly every hospital in America and shapes how billions of dollars in pharmaceutical spending are allocated across patients. Understanding who gains and who loses from this allocation—and whether the program's institutional design inadvertently disadvantages its intended beneficiaries—is essential for evidence-based reform. This paper takes a first step by documenting suggestive evidence of a specific channel through which 340B may undermine its own safety-net mission. The duplicate discount prohibition, designed to prevent manufacturers from paying double discounts, creates a payer incentive asymmetry that may cause hospitals to redirect drug administration away from Medicaid patients. Confirming this finding with greater statistical precision—through richer data, larger samples near the threshold, or complementary identification strategies—is an important task for future research. Whether this distributional consequence justifies reforming the prohibition, expanding Medicaid drug reimbursement, or restructuring 340B eligibility criteria are questions that deserve continued attention.

More broadly, the 340B crowd-out finding illustrates a recurring theme in health policy: programs designed to help vulnerable populations can be undermined by the same multi-payer payment system they operate within. The duplicate discount prohibition was a reasonable response to a specific problem (preventing double payments to manufacturers), but its interaction with differential reimbursement rates created an unintended incentive to redirect services away from Medicaid. Similar dynamics may affect other programs where provider reimbursement varies by payer type—Section 1115 waivers, Medicaid supplemental

payments, and Medicare-Medicaid dually eligible programs all create seams where payer-specific incentives could generate cross-payer externalities.

The policy conversation around 340B has focused primarily on whether the program generates sufficient community benefit to justify its cost to manufacturers. My findings add a distributional dimension: even if 340B generates positive community benefits on average, those benefits may be distributed regressively, with non-Medicaid patients receiving disproportionate gains. A complete welfare analysis of the 340B program would need to weigh the aggregate savings against this distributional cost—a task that requires data on all payers and all services, beyond the scope of this paper but within reach of the emerging administrative data infrastructure.

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

Bai, Ge and Gerard F Anderson, "Strategic Hospital Behavior Under the 340B Drug Pricing Program," *Health Affairs*, 2021, *40* (3), 430–437.

Baicker, Katherine, "The Spillover Effects of State Spending," *Journal of Public Economics*, 2005, *89* (2-3), 529–544.

Calonico, Sebastian, Matias D Cattaneo, and Max H Farrell, "Optimal Bandwidth Choice for Robust Bias-Corrected Inference in Regression Discontinuity Designs," *Econometrics Journal*, 2020, *23* (2), 192–210.

_ , _ , and Rocio Titiunik, "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 2014, *82* (6), 2295–2326.

Cattaneo, Matias D, Michael Jansson, and Xinwei Ma, "Simple Local Polynomial Density Estimators," *Journal of the American Statistical Association*, 2020, *115* (531), 1449–1455.

Centers for Medicare & Medicaid Services, "Transformed Medicaid Statistical Information System (T-MSIS)," https://www.medicaid.gov/medicaid/data-systems/macbis/tmsis/index.html 2024.

Clemens, Jeffrey and Stephen Miran, "Do Fiscal Rules Matter? Evidence from U.S. States," *American Economic Review*, 2014, *102* (2), 659–686.

Conti, Rena M and Peter B Bach, "The 340B Drug Discount Program: Hospitals Generate Profits by Expanding to Reach More Affluent Communities," *Health Affairs*, 2019, *38* (10), 1692–1700.

Currie, Janet and Jonathan Gruber, "Medicaid: What Does It Do, and Can We Do It Better?," *The Annals of the American Academy of Political and Social Science*, 2008, *686* (1), 75–93.

Dafny, Leemore S, "How Do Hospitals Respond to Price Changes?," *American Economic Review*, 2005, *95* (5), 1525–1547.

Desai, Sunita M and J Michael McWilliams, "Association Between the 340B Drug Pricing Program and Drug Prescribing Patterns," *JAMA Internal Medicine*, 2020, *180* (7), 1034–1037.

Dranove, David, Subramaniam Ramanarayanan, and Yashaswi Rao, "Does the Market Punish Aggressive Experts? Evidence from Cesarean Sections," *Journal of Health Economics*, 2017, *54*, 1–15.

Drug Channels Institute, "New HRSA Data: 340B Program Reached $53.7 Billion in 2022," https://www.drugchannels.net/ 2023.

Duggan, Mark, "Does Contracting Out Increase the Efficiency of Government Programs? Evidence from Medicaid HMOs," *Journal of Public Economics*, 2004, *88* (12), 2549–2572.

Duggan, Mark G, "Hospital Ownership and Public Medical Spending," *Quarterly Journal of Economics*, 2000, *115* (4), 1343–1373.

Einav, Liran, Amy Finkelstein, and Neale Mahoney, "Provider Incentives and Healthcare Costs: Evidence from Long-Term Care Hospitals," *Econometrica*, 2018, *86* (6), 2161–2219.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Sources of Geographic Variation in Health Care: Evidence from Patient Migration," *Quarterly Journal of Economics*, 2016, *131* (4), 1681–1726.

Frank, Richard G and Len M Nichols, "The Economics of Drug Channels: 340B Discounts and Drug Pricing," *Brookings Papers on Economic Activity*, 2020.

Gelman, Andrew and Guido Imbens, "Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs," *Journal of Business & Economic Statistics*, 2019, *37* (3), 447–456.

Gruber, Jonathan and Maria Owings, "Physician Financial Incentives and Cesarean Section Delivery," *RAND Journal of Economics*, 1997, *27* (1), 99–123.

Huang, Emily S and Jonathan D Ketcham, "Does 340B Status Affect Biosimilar Uptake? Evidence from Medicare Claims," *Journal of Health Economics*, 2024, *93*, 102844.

Hyman, David A and Duy Nguyen, "How Does the 340B Program Affect Hospital Behavior?," *Journal of Law and the Biosciences*, 2020, *7* (1).

Imbens, Guido W and Thomas Lemieux, "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 2008, *142* (2), 615–635.

Lee, David S, "Randomized Experiments from Non-Random Selection in U.S. House Elections," *Journal of Econometrics*, 2008, *142* (2), 675–697.

Mulligan, Kevin, "The Impact of the 340B Drug Pricing Program on Outpatient Drug Spending," *Health Affairs*, 2020, *39* (11), 1995–2003.

Nikpay, Sayeh, Thomas C Buchmueller, and Helen G Levy, "Effect of the 340B Drug Pricing Program on Hospitals and Physician Practices," *New England Journal of Medicine*, 2018, *378* (18), 1696–1703.

A. Data Appendix

A.1 T-MSIS Data Construction

The Transformed Medicaid Statistical Information System (T-MSIS) data are accessed as a pre-processed Parquet file containing Medicaid billing records at the provider-HCPCS-month level. The extract covers all 50 states plus DC and territories from 2018 through 2024; the analysis sample is restricted to 2019–2023 to match the HCRIS running variable coverage.

Drug billing (J-codes). I select all records where the first character of the HCPCS code is "J," which identifies injectable and infusible drugs administered in outpatient settings (HCPCS range J0100–J9999). These records are aggregated to the NPI-year level, summing total paid amounts, claim counts, and unique beneficiary counts.

Non-drug billing. All records with non-J HCPCS codes are aggregated to the NPI-year level as a placebo outcome. This includes evaluation and management codes, imaging, laboratory, and procedural services.

HCBS billing. Records with T, H, or S-code HCPCS prefixes—which identify home health, behavioral health, and Medicaid-specific services—are aggregated separately as an additional placebo.

A.2 HCRIS Data Construction

Hospital Cost Reporting Information System (HCRIS) data are downloaded from CMS for fiscal years 2019–2023 (files `hosp10FYxxxx.zip`). Each ZIP file contains three CSVs: the report file (hospital identifiers and report dates), the numeric file (financial data by worksheet/line/column), and the alpha file (text data including hospital addresses).

DSH percentage. Extracted from the numeric file, Worksheet E00A18A, Line 3200, Column 00100. When missing, constructed as the sum of the SSI ratio (Line 3000) and Medicaid utilization ratio (Line 3100).

Provider classification. General acute care hospitals are identified by Medicare provider number (CCN) suffix 0001–0879. Critical access hospitals (1300–1399), psychiatric hospitals (4000–4499), and other specialty types are excluded.

Hospital location. ZIP codes are extracted from the alpha file, Worksheet S200001, Line 200, Column 300.

A.3 CCN-NPI Crosswalk

The crosswalk linking HCRIS hospital Medicare provider numbers (CCNs) to T-MSIS NPIs is constructed in three steps:

1. **Identify hospital NPIs.** From the NPPES extract, select all organizational NPIs (entity type 2) with Healthcare Provider Taxonomy Code 282N (General Acute Care Hospital).

2. **ZIP-code matching.** Match HCRIS hospitals to NPPES hospital NPIs within the same 5-digit ZIP code, allowing multiple candidate NPIs per hospital.

3. **Selection.** Among candidate NPIs, select the one with the highest total Medicaid drug billing in T-MSIS. This resolves cases where multiple hospitals share a ZIP code by preferring the NPI with the strongest billing signal.

This procedure matches 2,712 of 3,224 general acute care hospitals (84%). Unmatched hospitals are primarily those without T-MSIS drug billing activity, consistent with hospitals that do not administer outpatient drugs.

Table 3 reports match rates by DSH bin, showing that the crosswalk performs comparably across the DSH distribution—critically, match rates do not differ systematically between hospitals above and below the 11.75% threshold. Among matched hospitals, 20% of NPIs are shared across multiple hospital CCNs, reflecting multi-campus health systems that bill under a single organizational NPI. The one-to-many structure is resolved by assigning each CCN its unique best-matching NPI based on Medicaid billing volume, but it introduces potential measurement error if billing is misattributed across campuses within a system. Classical measurement error in the outcome would attenuate the RDD estimate toward zero, making the significant panel result conservative if anything.

B. Identification Appendix

B.1 McCrary Density Test

Figure 2 presents the density manipulation test of Cattaneo et al. (2020) at the 11.75% threshold. The test evaluates whether the density of the running variable (DSH adjustment percentage) is discontinuous at the cutoff, which would indicate that hospitals sort around the threshold. The test provides no evidence of manipulation (the p-value is reported in the figure subtitle). The estimated densities on each side of the cutoff have similar levels and slopes, consistent with locally random assignment.

B.2 Covariate Balance

Using the same `rdrobust` specification as the main analysis, I test for discontinuities in predetermined covariates at the threshold:

- Non-drug Medicaid spending (levels): RD = −622,493 ($p = 0.33$)
- HCBS billing (levels): RD = −1,414 ($p = 0.61$)

Neither covariate shows a significant discontinuity, supporting the local randomization assumption.

B.3 Donut Hole Sensitivity

Excluding hospitals within $\pm d$ percentage points of the threshold for $d \in \{0.25, 0.5, 1.0\}$:

Donut	Estimate	Robust SE	p-value	Eff. N (L/R)
±0.25pp	−1.742	2.204	0.429	74/88
±0.50pp	−1.931	2.465	0.433	68/95
±1.00pp	−1.278	3.657	0.727	45/65

All estimates are negative, confirming that the main result is not driven by hospitals at the threshold boundary.

C. Robustness Appendix

C.1 Extensive Margin and Conditional Log Estimates

Outcome	Estimate	Robust SE	p-value	Eff. N (L/R)
Pr(Any Medicaid Drug Billing)	0.025	0.241	0.919	61/56
Log Medicaid Drug (cond. on >0)	−1.075	1.377	0.435	50/45

The extensive margin (probability of any billing) shows no discontinuity. The conditional log specification (excluding hospitals with zero Medicaid drug billing) is directionally consistent with the asinh estimate.

C.2 Bandwidth Sensitivity

Figure 5 presents the full bandwidth sensitivity analysis. The table below reports point estimates and inference for fixed bandwidths:

Bandwidth	Estimate	Robust SE	p-value	N (left/right)
2pp	−0.571	3.002	0.849	48/38
3pp	−0.556	2.469	0.822	62/63
4pp	−0.448	2.136	0.834	81/95
5pp	−0.389	1.951	0.842	98/142
7pp	−1.238	1.650	0.453	137/271
10pp	−1.415	1.332	0.288	222/484

C.3 Polynomial Order

Order	Estimate	Robust SE	Optimal BW
Linear ($p = 1$)	−0.435	1.937	3.3pp
Quadratic ($p = 2$)	−0.409	2.402	4.1pp

Both specifications yield negative estimates. The linear specification is preferred following Gelman and Imbens (2019).

C.4 Panel Specification Variants

Specification	BW	Estimate	SE	p-value	N Clusters
Year FE	±10pp	−1.145	0.520	0.028	3,219 706
Year FE	±5pp	−1.395	0.616	0.024	1,138 240
State + Year FE	±10pp	−0.456	0.416	0.274	3,160 705

The state FE specification attenuates the estimate to −0.46 ($p = 0.27$), raising the question of how much cross-state heterogeneity in Medicaid policies (carve-in/carve-out status, managed care penetration, fee schedule generosity) drives the baseline result versus the within-state RDD variation.

C.5 Year-by-Year Estimates

Year	Estimate	Robust SE	p-value	Effective N
2019	−3.962	2.809	0.159	100
2020	3.501	2.164	0.106	111
2021	−2.741	2.154	0.203	101
2022	0.132	2.452	0.957	78
2023	−3.979	1.882	0.035	121

The effect is negative in three of five years (2019, 2021, 2023), with 2020 positive and 2022 near zero. The 2020 reversal is consistent with COVID-19 disruptions.

D. Additional Figures and Tables

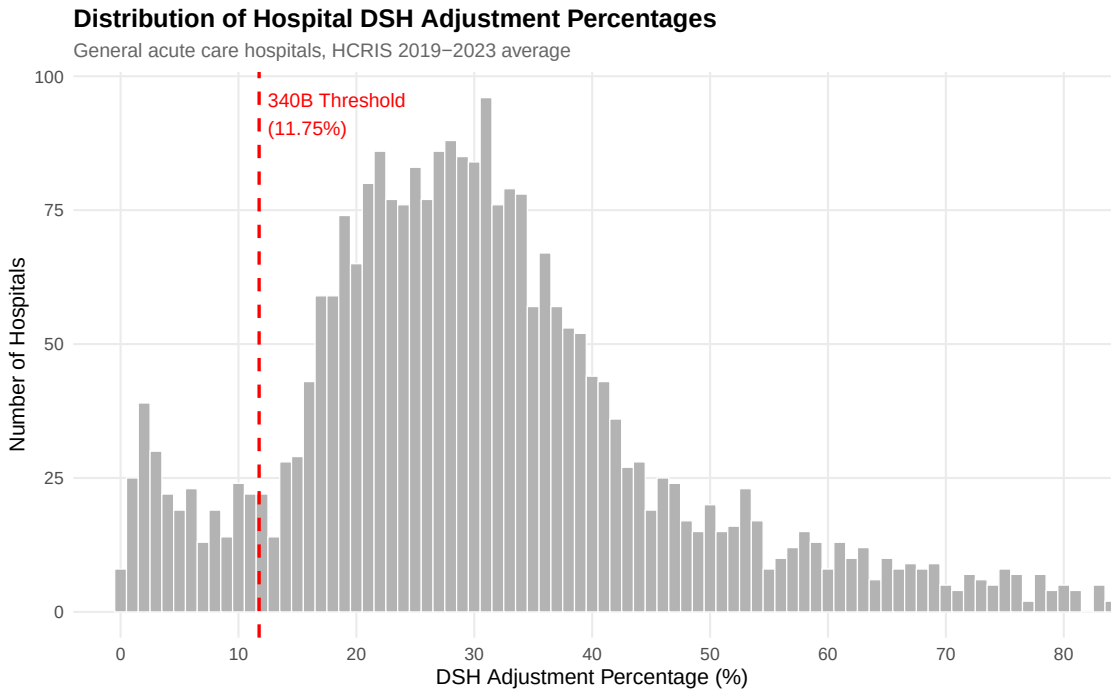

Figure 1: Distribution of Hospital DSH Adjustment Percentages

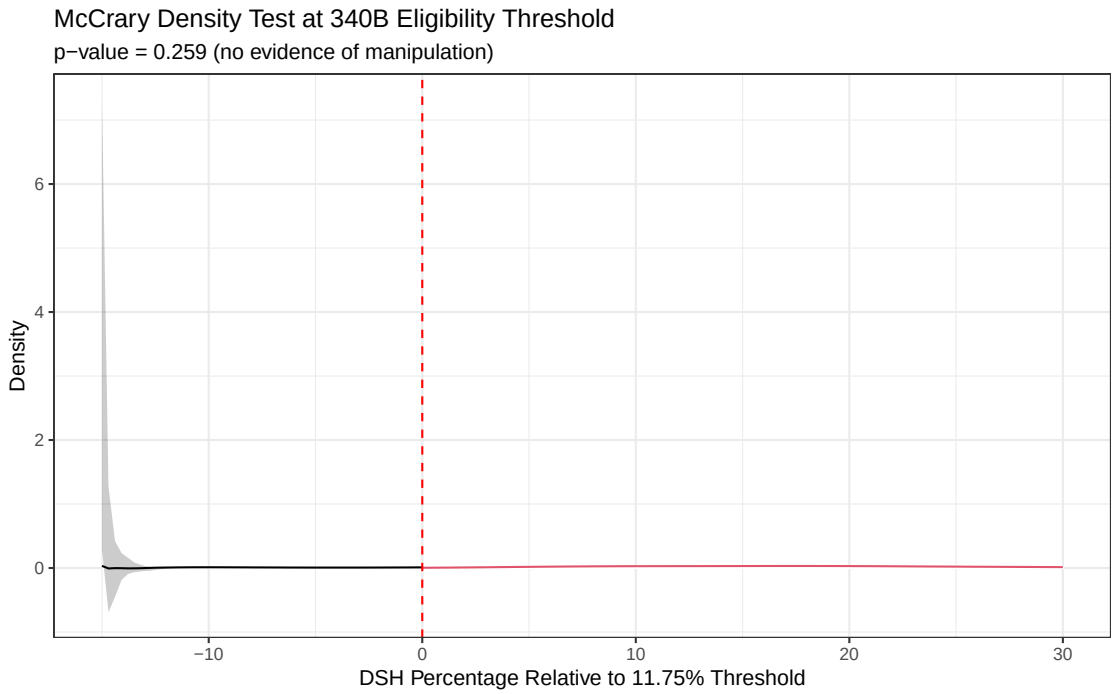

Figure 2: McCrary Density Test at 340B Eligibility Threshold

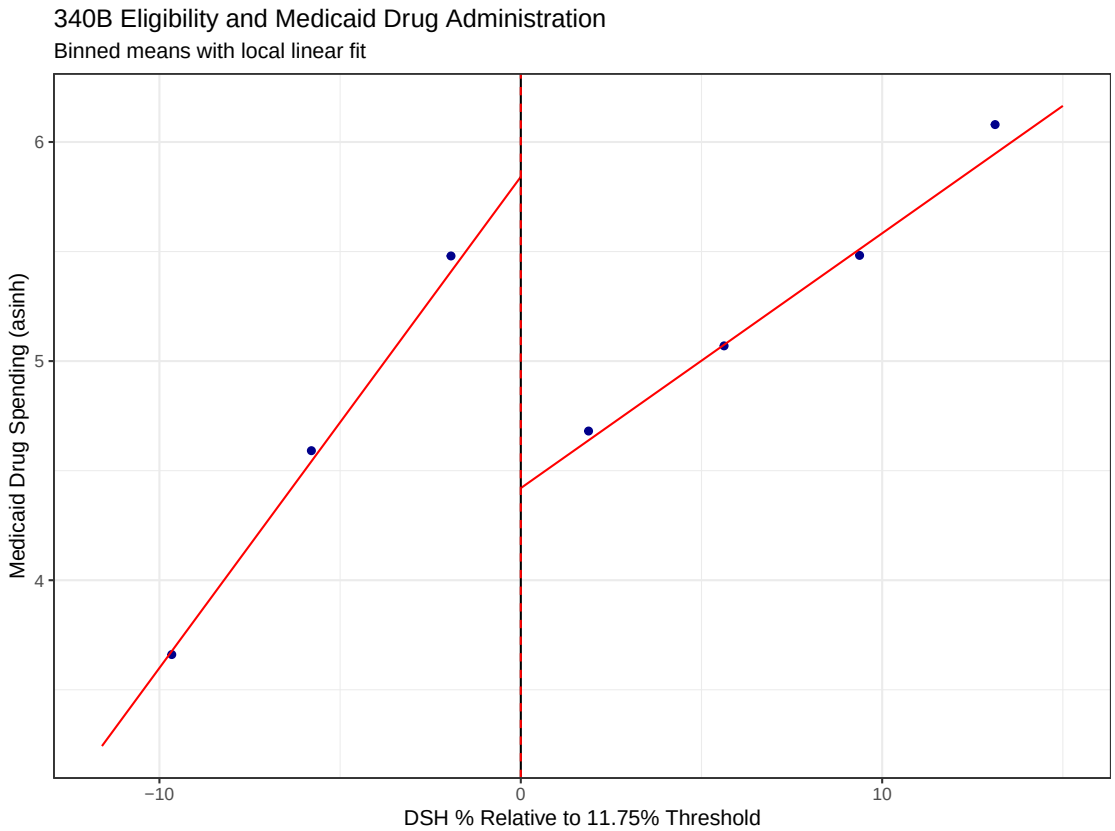

Figure 3: RDD Plot: 340B Eligibility and Medicaid Drug Spending

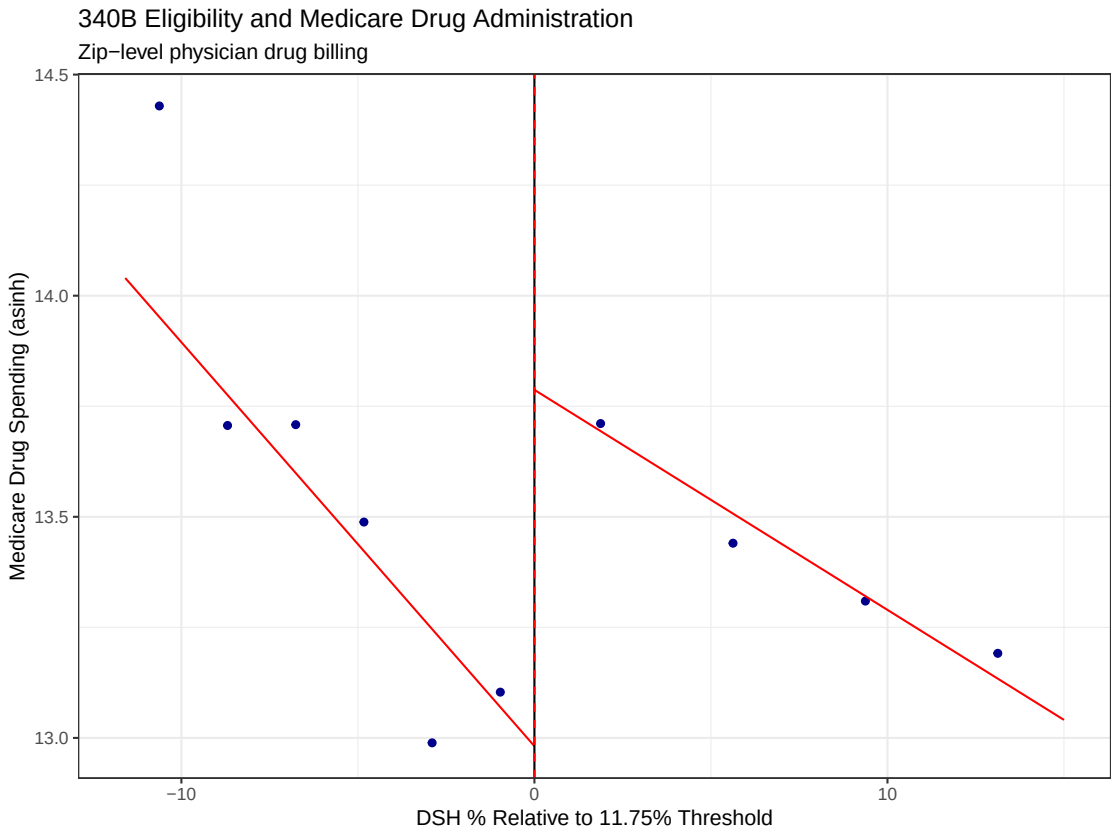

Figure 4: RDD Plot: 340B Eligibility and Medicare Drug Spending

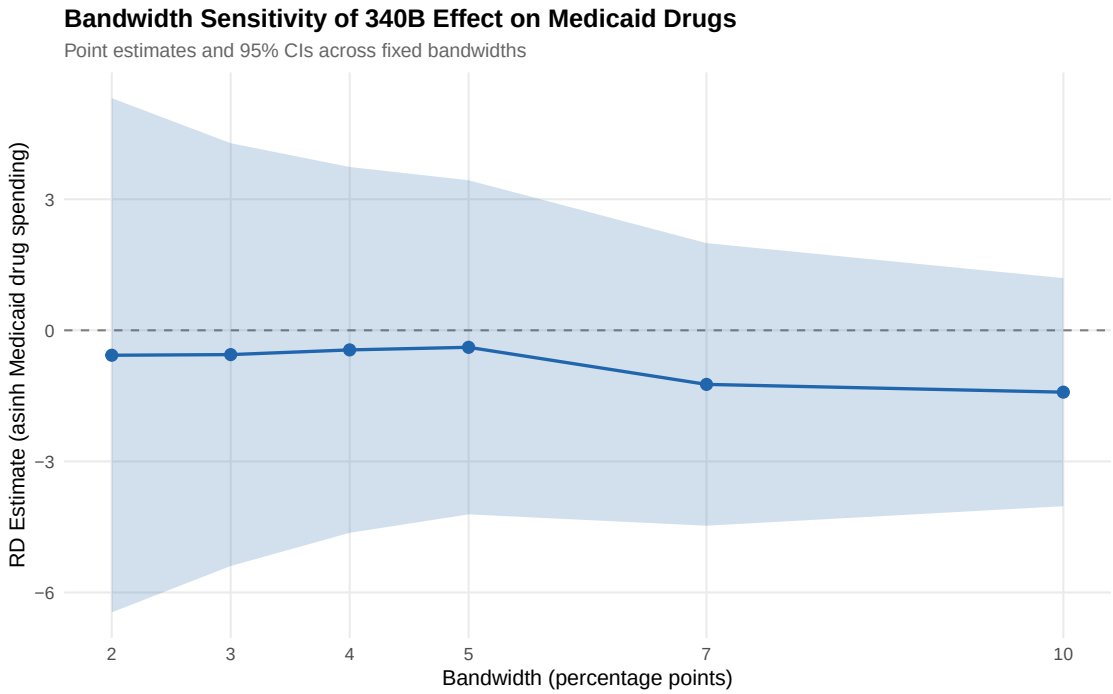

Figure 5: Bandwidth Sensitivity of RDD Estimates

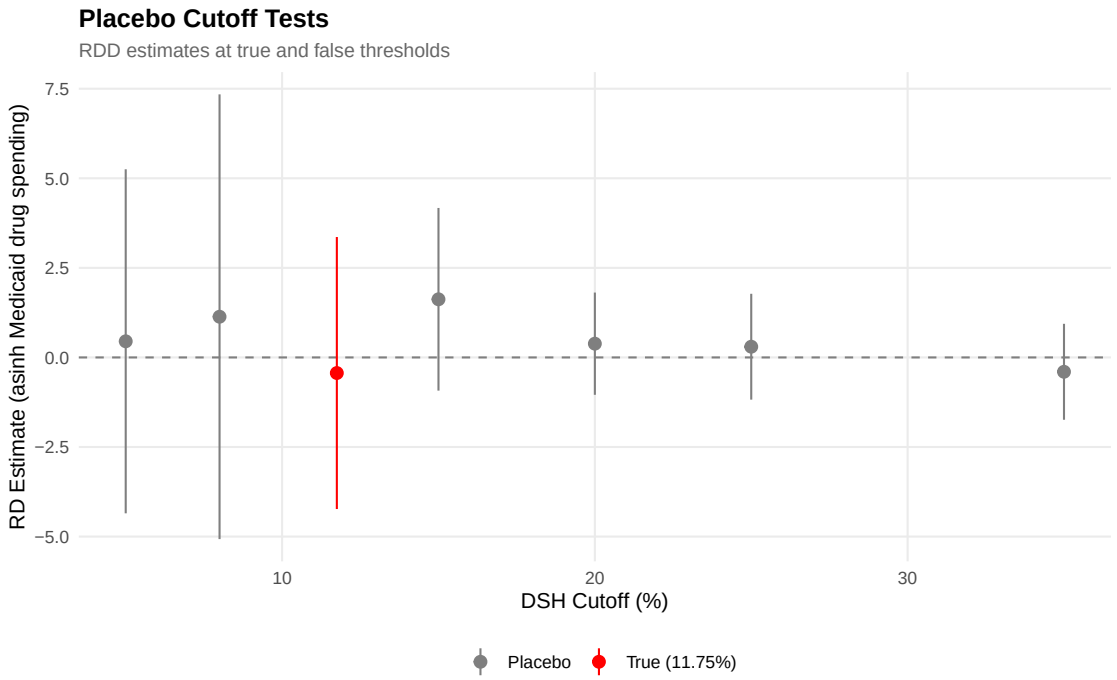

Figure 6: Placebo Cutoff Tests

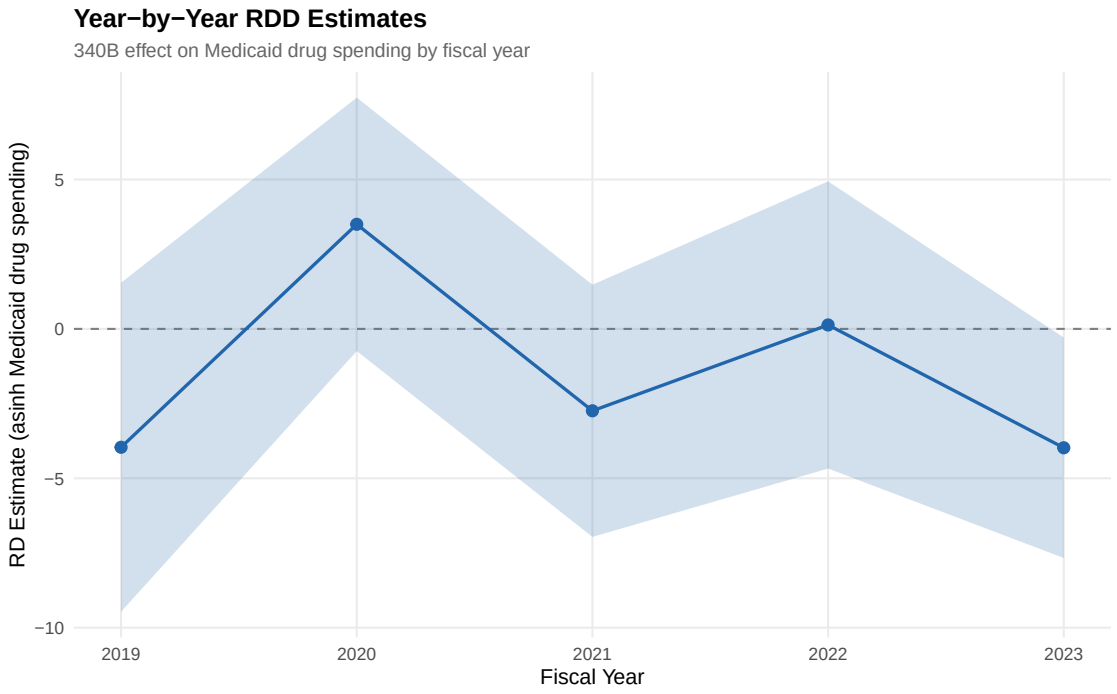

Figure 7: Year-by-Year RDD Estimates

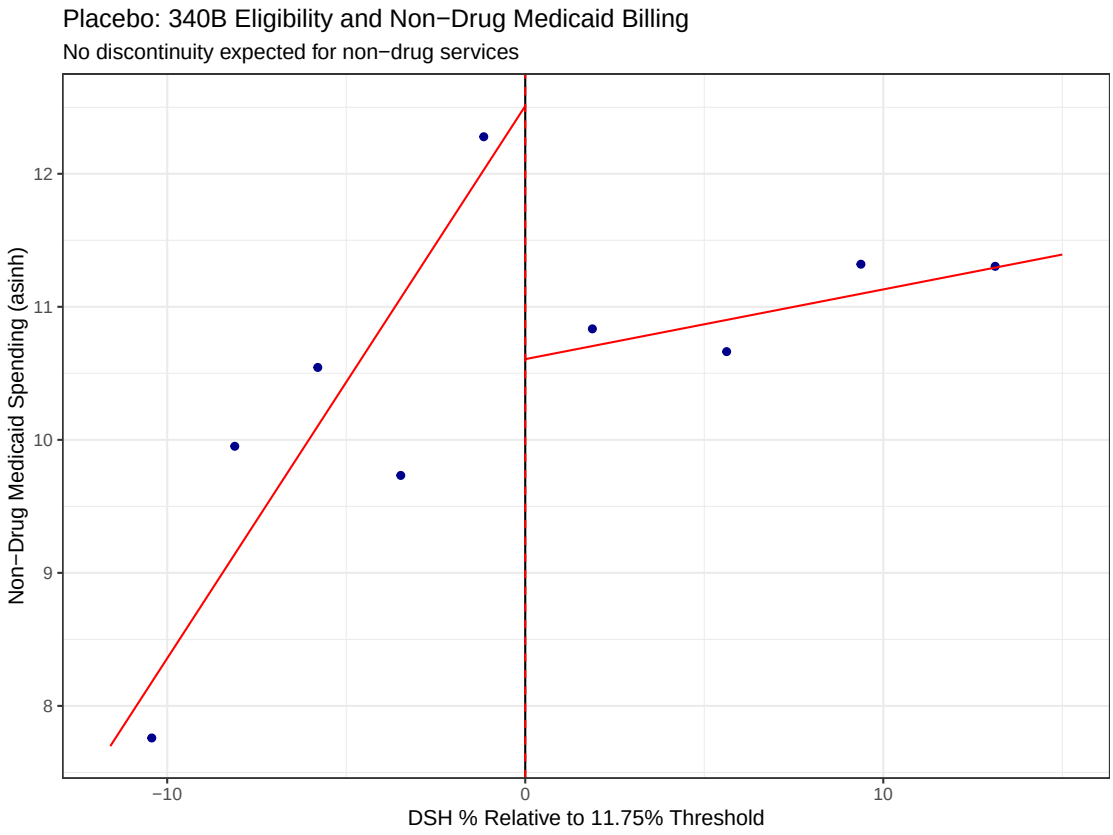

Figure 8: Placebo: 340B Eligibility and Non-Drug Medicaid Spending

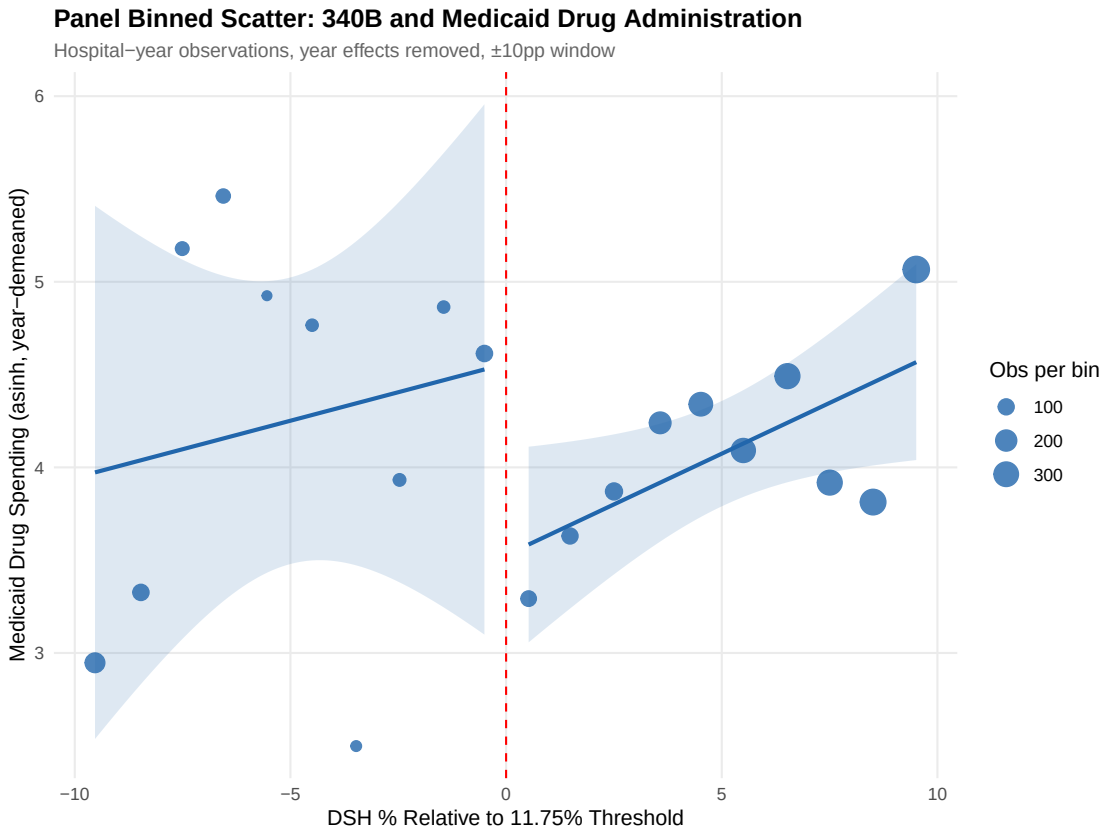

Figure 9: Panel Binned Scatter: Medicaid Drug Spending by DSH Percentile. Hospital-year observations within ±10pp of the threshold, year effects removed. Bin width is 1 percentage point. Point sizes reflect the number of hospital-year observations per bin. Linear fits estimated separately on each side of the threshold.

Table 3: CCN-NPI Crosswalk Match Rates by DSH Bin

DSH Bin	Total Hospitals	Matched	Match Rate (%)
<5%	423	325	76.8
5–10%	247	202	81.8
10–11.75%	159	136	85.5
11.75–15%	244	214	87.7
15–20%	424	374	88.2
>20%	1,727	1,461	84.6

Notes: Match rates for general acute care hospitals by DSH adjustment percentage bin. Matching uses ZIP code and NPPES taxonomy 282N (General Acute Care Hospital), with Medicaid drug billing volume as tiebreaker when multiple NPIs share a ZIP code. Match rates are comparable on both sides of the 11.75% threshold, supporting the assumption that crosswalk quality does not introduce differential measurement error at the cutoff.