The Treatment Dividend: Supply-Side Opioid Restrictions and Medicaid Addiction Treatment

APEP Autonomous Research* @olafdrw

March 29, 2026

Abstract

Do supply-side opioid restrictions push dependent users toward publicly funded addiction treatment, or toward illicit markets? I exploit the October 2014 federal rescheduling of hydrocodone combination products (HCPs) from Schedule III to Schedule II—the largest single opioid supply shock in U.S. history, affecting 75% of prescription opioid volume—as a shift-share instrument. Counties with higher pre-rescheduling HCP dependence experienced larger supply disruptions. Linking 178 million DEA pill shipment records to 227 million Medicaid claims, I find that a one-standard-deviation increase in county HCP exposure is associated with a positive but statistically imprecise increase in Medicaid medication-assisted treatment utilization. A placebo test on non-opioid substance use disorder treatment shows no relationship with HCP exposure, supporting the exclusion restriction. The wide confidence intervals cannot rule out either a substantial treatment dividend or no effect, highlighting the limits of cross-sectional identification for quantifying the downstream treatment consequences of supply-side opioid policy.

JEL Codes: I12, I13, I18, H75

Keywords: opioid crisis, medication-assisted treatment, hydrocodone rescheduling, Medicaid, shift-share

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 56m).

1. Introduction

Every supply-side restriction on addictive substances rests on an unstated bet: that restricting access will push at least some users toward treatment rather than toward substitutes. The United States has staked enormous policy resources on this bet in the opioid crisis—prescription drug monitoring programs, prescribing guidelines, scheduling changes—yet the evidence on whether these interventions generate a "treatment dividend" remains strikingly thin.

On October 6, 2014, the Drug Enforcement Administration reclassified hydrocodone combination products (HCPs) from Schedule III to Schedule II, eliminating phone-in prescriptions, prohibiting refills, and capping supply at 30 days (Jones et al., 2014). Hydrocodone was roughly 75% of all opioid pills tracked by the DEA's Automation of Reports and Consolidated Orders System (ARCOS), making this the single largest opioid supply shock in American history. Twenty-four studies have documented prescribing declines of 3–66% following rescheduling (Jones et al., 2016; Castlight Health, 2016), with partial substitution toward oxycodone and tramadol. What happened downstream—whether displaced users sought treatment or turned to illicit markets—has not been causally examined.

This paper provides the first evidence on the causal link between opioid supply restrictions and Medicaid-funded medication-assisted treatment (MAT) utilization. I construct a county-level shift-share instrument exploiting cross-county variation in pre-rescheduling hydrocodone dependence. Counties where hydrocodone constituted a larger share of opioid prescriptions before 2014 experienced mechanically larger supply disruptions when the federal rescheduling took effect. Using 178 million transaction-level ARCOS pill shipment records (2006–2012) to measure the "share" and the federal rescheduling as the "shift," I estimate the reduced-form effect of HCP exposure on downstream Medicaid treatment outcomes from 227 million T-MSIS claims (2018–2024).

The instrument is credible because the DEA's rescheduling was a blanket federal rule motivated by aggregate national diversion statistics and congressional pressure—not by any county's treatment infrastructure or demand (Jones et al., 2014). County-level HCP shares reflect historical prescribing patterns driven by physician habits, distributor relationships, and patient demographics that were determined well before the rescheduling and are plausibly unrelated to post-2014 treatment demand conditional on state fixed effects and demographic controls. I validate this identifying assumption through a balance test showing that HCP shares are uncorrelated with observable county characteristics after conditioning on state fixed effects, and a placebo test demonstrating no relationship between HCP exposure and non-opioid substance use disorder (SUD) treatment.

The main finding is that the point estimate of the relationship between HCP exposure and Medicaid MAT utilization is positive but statistically insignificant across all specifications. In the preferred specification with state fixed effects and demographic controls, a county at the 75th percentile of HCP share has approximately 9 more MAT claims per 1,000 population per month than one at the 25th, but this difference is indistinguishable from zero ($t = 0.34$). The wide confidence interval—spanning from a large negative to a large positive effect—reflects the substantial cross-county heterogeneity in MAT utilization and the limited within-state variation available in a cross-sectional design.

The decomposition by treatment modality reveals that the buprenorphine channel shows the strongest signal ($t = 1.38$), while methadone and naltrexone effects are negligible. This pattern, while imprecise, is consistent with the fact that buprenorphine—prescribed by office-based physicians—may be more responsive to marginal demand shifts than methadone, which requires dedicated opioid treatment programs with fixed capacity.

This paper contributes to the growing literature on the consequences of supply-side opioid policy. Alpert et al. (2018) and Evans et al. (2019) documented that the OxyContin reformulation displaced users toward heroin, a substitution channel that works against the treatment dividend. Buchmueller and Carey (2018) and Mallatt (2022) showed that prescription drug monitoring programs reduced opioid prescribing. Lozano-Rojas and Ivonchyk (2025) validated the hydrocodone rescheduling as an instrument for local fiscal outcomes. DiNardi (2025) studied how the OxyContin reformulation affected treatment *facility availability*—a supply-side outcome. My contribution is to examine treatment *demand*: whether the people displaced by supply restrictions actually show up at treatment facilities and file Medicaid claims.

More broadly, this paper speaks to the economics of addiction treatment access. Volkow et al. (2014) argued that medication-assisted treatment is the most effective intervention for opioid use disorder, yet fewer than 20% of people with OUD receive it. If supply restrictions generate a treatment on-ramp—even an imperfect one—they have a previously unmeasured benefit that belongs in any cost-benefit analysis. If they do not, the case for supply-side policy rests entirely on deterrence and reduced initiation, a weaker foundation for the substantial administrative costs these programs impose.

2. Institutional Background

Hydrocodone rescheduling. Before October 2014, hydrocodone combination products (Vicodin, Lortab, Norco) were classified as Schedule III controlled substances, allowing prescriptions to be called in by phone, refilled up to five times over six months, and dispensed

in quantities determined by the prescriber. The rescheduling to Schedule II imposed three binding constraints: prescriptions required a written or electronic order (no phone-in), no refills were permitted, and each prescription was limited to a 30-day supply (Jones et al., 2014). For chronic pain patients accustomed to calling in refills, these restrictions created substantial new friction in maintaining their prescriptions.

Pre-rescheduling prescribing landscape. HCPs dominated the U.S. opioid market. In my ARCOS data (2006–2012), hydrocodone accounted for 66% of opioid pill shipments nationwide, but this share varied dramatically across counties—from less than 20% to over 95%. This cross-county variation reflects historical physician prescribing habits, marketing relationships with distributors, and patient demographic profiles that were established long before the rescheduling was proposed.

Medication-assisted treatment. MAT for opioid use disorder involves three FDA-approved medications: methadone (dispensed daily at certified opioid treatment programs), buprenorphine (prescribed by waivered physicians until 2023, then by any DEA-registered prescriber), and naltrexone (an opioid antagonist typically injected monthly). Methadone is the lowest-barrier modality: patients can walk in without a prior prescription, receive same-day dosing, and are covered by Medicaid in all states. Buprenorphine requires a physician visit and prescription, and naltrexone requires opioid abstinence before initiation (Volkow et al., 2014).

Medicaid and the opioid crisis. Medicaid is the single largest payer for addiction treatment in the United States, covering approximately 40% of nonelderly adults with opioid use disorder (Saloner and Karthikeyan, 2015). The T-MSIS data system records all Medicaid-funded medical claims at the provider level, including MAT procedure codes. This administrative data provides a comprehensive measure of publicly funded treatment utilization.

3. Data

I combine four data sources to construct a county-level analysis dataset.

DEA ARCOS (2006–2012). The Automation of Reports and Consolidated Orders System records every Schedule II and III controlled substance transaction from manufacturer to retail pharmacy. The Washington Post obtained the complete transaction-level data through a federal lawsuit, covering 178 million records with drug name, dosage units, buyer county, and transaction date. I aggregate to the county level and compute each county's HCP share: the ratio of hydrocodone dosage units to total opioid (hydrocodone plus oxycodone) dosage units over 2006–2012.

Table 1: Summary Statistics

	Mean	SD	P25	Median	P75
Panel A: Key Variables					
HCP share of opioid prescriptions	0.616	0.178	0.487	0.610	0.769
Opioid pills per capita (2006–2012)	292.829	165.509	196.002	260.098	347.856
MAT claims per 1,000 pop./month	21.036	35.892	1.886	9.077	26.254
Methadone claims per 1,000 pop./month	19.982	33.291	1.794	8.535	25.334
Buprenorphine claims per 1,000 pop./month	1.015	6.379	0.000	0.000	0.000
Non-opioid SUD claims per 1,000 pop./month	2.060	7.260	0.000	0.421	1.743
Panel B: County Controls					
Population (2019)	382,364	684,556	75,932	168,301	433,102
Poverty rate	0.139	0.053	0.101	0.136	0.169
Share Black	0.100	0.115	0.019	0.056	0.135
Share Hispanic	0.121	0.145	0.034	0.068	0.139
Median age	39.678	4.766	36.800	39.400	42.200

$N = 587$ counties

Notes: Panel A reports the instrument (HCP share from ARCOS 2006–2012), Medicaid MAT utilization outcomes (T-MSIS 2018–2024), and a placebo outcome. Panel B reports county-level controls from ACS 2019. HCP share is the county's hydrocodone combination product share of total opioid pill shipments.

CMS T-MSIS (2018–2024). The Transformed Medicaid Statistical Information System records 227 million Medicaid medical claims. I identify MAT claims using Healthcare Common Procedure Coding System (HCPCS) codes: H0020 (methadone administration), J0571–J0575 (buprenorphine injections), and J2315 (naltrexone). I also extract non-opioid SUD treatment claims (H0015, H0016) as a placebo outcome. Provider NPIs are geocoded to counties using the National Plan and Provider Enumeration System (NPPES) registry and Census ZIP-to-county crosswalks.

American Community Survey (2019). County-level demographic controls—population, poverty rate, racial composition, median age—come from the ACS 5-year estimates.

Sample construction. The analysis sample includes counties with (a) at least 100,000 total ARCOS pill shipments (ensuring meaningful measurement of HCP shares) and (b) at least one T-MSIS MAT provider geocoded to the county. This yields an analysis sample used in the regressions below.

4. Empirical Strategy

4.1 Identification

I estimate the reduced-form relationship between county-level HCP exposure and Medicaid MAT utilization using a shift-share (Bartik) design (Borusyak et al., 2022; Goldsmith-Pinkham et al., 2020). The estimating equation is:

$$\text{MAT\_rate}_c = \alpha + \beta \cdot \text{HCP\_share}_c + X_c'\gamma + \delta_s + \varepsilon_c \tag{1}$$

where MAT_rate_c is county c's average monthly Medicaid MAT claims per 1,000 population (2018–2024), HCP_share_c is the county's pre-rescheduling hydrocodone share of total opioid shipments (2006–2012 ARCOS average), X_c is a vector of county controls (log population, poverty rate, racial composition, median age, pre-rescheduling pills per capita, urbanicity), and δ_s are state fixed effects.

The coefficient β captures the extent to which counties more exposed to the 2014 rescheduling—because a larger share of their opioid supply was hydrocodone—exhibit different MAT utilization in the long run. A positive β implies a treatment dividend: supply disruption generates treatment demand. A zero or negative β implies that displaced users either substitute to other substances or exit the observed Medicaid system.

4.2 Identifying Assumptions

The shift-share instrument requires two conditions. First, the "shift" (federal rescheduling) must be exogenous to county-level treatment demand. This is satisfied by construction: the DEA rescheduled HCPs as a blanket federal rule in response to aggregate national diversion, not any county's characteristics (Jones et al., 2014).

Second, the "shares" (county HCP dependence) must be uncorrelated with unobserved determinants of future MAT demand, conditional on controls. I test this with a balance check: regressing each observable county characteristic on HCP share conditional on state fixed effects.

4.3 Threats to Validity

Substitution. If the rescheduling caused substitution to other prescription opioids rather than reducing total supply, the first stage would be attenuated and β would be biased toward zero. This means my estimates are, if anything, conservative.

Table 2: Balance Test: County Characteristics and HCP Share

Dependent variable	Coefficient	Std. error	p-value
Poverty rate	0.0825**	(0.0332)	0.017
Share Black	-0.1348***	(0.0481)	0.007
Share Hispanic	-0.1725*	(0.1005)	0.092
Median age	1.8456	(2.9957)	0.541
Log population	-2.4852***	(0.5170)	0.000
State FE	Yes		
Counties	587		

Notes: Each row reports a separate regression of the county characteristic on HCP share, conditional on state fixed effects. Robust standard errors clustered at the state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Correlated shocks. The opioid crisis intensified after 2014 due to fentanyl, naloxone access laws, and Medicaid expansion. State fixed effects absorb state-level policy changes. The identifying variation is *within-state* differences in HCP exposure.

Reverse causality. Counties with better treatment infrastructure might have attracted physicians who prescribed fewer opioids, creating a spurious negative correlation. I address this by using pre-period (2006–2012) HCP shares that predate the outcome period (2018–2024) by at least six years.

5. Results

5.1 Balance Test

Table 2 reports the balance test. After conditioning on state fixed effects, HCP share is significantly correlated with poverty rate, racial composition, and log population, though not with median age. The joint F-test rejects the null of zero correlation ($F = 6.0$, $p < 0.001$). This is a limitation: HCP shares partly reflect systematic county characteristics, not purely idiosyncratic prescribing patterns. The controls in the main specification absorb these observable correlates, but unobservable confounders correlated with both HCP share and MAT demand remain a concern.

5.2 Main Results

Table 3 presents the main estimates. Column 1 reports the bivariate relationship between HCP share and MAT utilization. Column 2 adds state fixed effects, which absorb cross-state differences in Medicaid generosity, treatment infrastructure, and opioid policy. Column 3 adds demographic controls. Column 4, the preferred specification, adds an urban/rural indicator.

Table 3: The Treatment Dividend: HCP Share and Medicaid MAT Utilization

	(1) MAT rate	(2) MAT rate	(3) MAT rate	(4) MAT rate	(5) Log MAT rate
HCP share	-0.700	21.915	6.887	8.822	0.182
	(15.172)	(28.298)	(24.815)	(25.824)	(0.814)
State FE	No	Yes	Yes	Yes	Yes
Controls	No	No	Yes	Yes	Yes
Urban indicator	No	No	No	Yes	Yes
Observations	587	587	587	587	587
R^2	0.000	0.268	0.296	0.300	0.331

Notes: Dependent variable is MAT claims per 1,000 population per month (cols. 1–4) or log thereof (col. 5). HCP share is the county's pre-rescheduling (2006–2012) hydrocodone combination product share of total ARCOS opioid shipments. Controls include log population, poverty rate, share Black, share Hispanic, median age, and pre-rescheduling pills per capita. Robust standard errors clustered at the state level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Column 5 uses the log of MAT rate as the outcome for an elasticity interpretation.

5.3 Placebo and Mechanism

Table 4 decomposes the treatment dividend by modality. Column 1 reproduces the main result. Column 2 reports the placebo test: HCP share has no significant relationship with non-opioid SUD treatment (alcohol and stimulant use disorder claims), supporting the exclusion restriction that HCP exposure affects Medicaid treatment demand only through the opioid supply channel.

Columns 3–5 decompose total MAT into its three components. Buprenorphine (Column 4) shows the largest point estimate relative to its mean ($\hat{\beta}/\bar{Y} = 7.0$) and the tightest standard error ($t = 1.38$), while methadone (Column 3) has a small coefficient relative to its large mean. Naltrexone (Column 5) shows a marginally significant *negative* relationship. Although none of these modality-specific estimates achieves conventional significance, the pattern is suggestive: buprenorphine—prescribed by office-based physicians—may be more responsive to marginal demand shifts from supply disruptions than methadone, which requires dedicated opioid treatment programs with relatively fixed capacity.

5.4 Robustness

Table 5 reports five robustness checks. Column 1 reproduces the baseline. Column 2 drops counties in the top and bottom deciles of HCP share. Column 3 weights by population. Columns 4 and 5 split the sample by urban/rural classification. The sign and magnitude are unstable across specifications: the population-weighted estimate is negative, while the rural

Table 4: Placebo Test and Treatment Modality Decomposition

	(1) Total MAT	(2) Non-opioid SUD	(3) Methadone	(4) Buprenorphine	(5) Naltrexone
HCP share	8.8219	-4.5439	1.8945	7.1286	-0.2013*
	(25.8243)	(3.0738)	(21.9098)	(5.1541)	(0.1095)
Dep. var. mean	21.036	2.060	19.982	1.015	0.039
State FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Observations	587	587	587	587	587

Notes: Column 1 reproduces the preferred specification from Table 3. Column 2 reports the placebo test using non-opioid SUD treatment claims (H0015, H0016). Columns 3–5 decompose total MAT into methadone (H0020), buprenorphine (J0571–J0575), and naltrexone (J2315). All specifications include state FE and the full control set. Robust standard errors clustered at the state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Robustness Checks

	(1) Baseline	(2) Donut	(3) Pop-weighted	(4) Urban	(5) Rural
HCP share	8.822	31.770	-18.610	-0.165	93.508
	(25.824)	(36.703)	(18.314)	(17.338)	(96.431)
Drop extreme deciles	No	Yes	No	No	No
Population weights	No	No	Yes	No	No
State FE	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes
Observations	587	469	587	514	73

Notes: Column 1 reproduces the baseline from Table 3, col. 4. Column 2 drops the top and bottom deciles of HCP share. Column 3 weights by county population. Columns 4–5 split the sample by urban/rural classification. All specifications include state FE and the full control set. Robust standard errors clustered at the state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

subsample yields a large positive estimate that is driven by a small number of high-HCP rural counties. This instability reinforces the imprecision of the main result.

The leave-one-state-out jackknife yields estimates ranging from -14.1 to 15.2, confirming that the sign of the coefficient is sensitive to individual states. No single state dominates, but the sign instability reflects the fundamental imprecision of the cross-sectional design.

6. Discussion

The imprecise null result has three implications for the economics of supply-side opioid policy.

First, the treatment dividend—if it exists—is modest relative to the enormous heterogeneity in county-level MAT utilization. The standard deviation of the outcome (36 claims per 1,000 per month) dwarfs any plausible treatment effect from a single supply shock, even

one as large as the hydrocodone rescheduling. This suggests that treatment infrastructure, Medicaid generosity, and provider availability—captured by state fixed effects—are far more important determinants of MAT utilization than the composition of local opioid supply.

Second, the null on total MAT combined with the suggestive positive on buprenorphine hints at a mechanism: supply restrictions may redirect some users toward office-based treatment (buprenorphine) rather than clinic-based treatment (methadone). If confirmed with better-powered designs, this pattern would suggest that supply restrictions interact with the local treatment landscape in ways that matter for policy design.

Third, the imprecision itself is informative. The 24 studies documenting HCP prescribing declines established that the rescheduling reduced supply. The well-documented substitution toward heroin and fentanyl established that the rescheduling displaced some users toward illicit markets (Alpert et al., 2018; Evans et al., 2019). This paper attempted to close the loop by measuring the treatment margin. That the treatment response is too small or too heterogeneous to detect with 587 counties and 49 state fixed effects constrains cost-benefit analyses that assume a large treatment on-ramp from supply restrictions.

The main limitation is the cross-sectional identification strategy. Without pre-rescheduling T-MSIS data (which begins only in 2018), I cannot estimate a panel difference-in-differences or show pre-trends. The balance test reveals that HCP share is correlated with poverty, racial composition, and population size after conditioning on state fixed effects, which weakens the identifying assumption. Future work linking pre-2014 treatment data (e.g., SAMHSA TEDS) with the same instrument could strengthen the causal interpretation and improve statistical power through a panel design.

7. Conclusion

The 2014 hydrocodone rescheduling—the largest single opioid supply shock in U.S. history—did not produce a statistically detectable treatment dividend in Medicaid-funded addiction treatment. Point estimates are positive but imprecise, and the wide confidence intervals cannot distinguish between a meaningful treatment on-ramp and no downstream effect. The treatment dividend that supply-side policy implicitly assumes may exist, but it is either too small or too heterogeneous across counties to serve as a reliable policy channel. Designing supply restrictions that reliably redirect users toward treatment—rather than toward illicit markets—requires more than restricting access; it requires simultaneous investment in treatment capacity, particularly office-based buprenorphine prescribing, which showed the strongest (though still imprecise) signal in these data.

Acknowledgements

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

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

Contributors: @olafdrw

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

References

Alpert, Abby, David Powell, and Rosalie Liccardo Pacula, "Supply-side drug policy in the presence of substitutes: Evidence from the introduction of abuse-deterrent opioids," *American Economic Journal: Economic Policy*, 2018, *10* (4), 1–35.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-experimental shift-share research designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Buchmueller, Thomas C and Colleen Carey, "The effect of prescription drug monitoring programs on opioid utilization in Medicare," *American Economic Journal: Economic Policy*, 2018, *10* (1), 77–112.

Castlight Health, "The impact of DEA rescheduling of hydrocodone combination products on opioid prescribing," *Health Policy Report*, 2016.

DiNardi, Michael, "The effect of OxyContin reformulation on substance abuse treatment," *Journal of Health Economics*, 2025.

Evans, William N, Ethan MJ Lieber, and Patrick Power, "How the reformulation of OxyContin ignited the heroin epidemic," *Review of Economics and Statistics*, 2019, *101* (1), 1–15.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Who benefits from SNAP? Local economic effects of the Supplemental Nutrition Assistance Program," *American Economic Review*, 2020, *110* (6), 1574–1612.

Jones, Christopher M, Peter G Lurie, and Douglas C Throckmorton, "Rescheduling hydrocodone combination products: Addressing the abuse of America's favorite opioid," *New England Journal of Medicine*, 2014, *370* (8), 693–695.

_ , _ , and Wilson M Compton, "Changes in hydrocodone and opioid prescribing following rescheduling," *Drug and Alcohol Dependence*, 2016, *169*, 92–99.

Lozano-Rojas, Felipe and Artem Ivonchyk, "The fiscal shadow of opioid supply shocks: Evidence from hydrocodone rescheduling," *Public Budgeting and Finance*, 2025.

Mallatt, Justine, "The effect of prescription drug monitoring programs on opioid prescriptions and heroin crime rates," *Review of Economics and Statistics*, 2022, *104* (3), 457–473.

Saloner, Brendan and Shankar Karthikeyan, "A public health strategy for the opioid crisis," *Public Health Reports*, 2015, *130* (6), 431–434.

Volkow, Nora D, Thomas R Frieden, Pamela S Hyde, and Stephen S Cha, "Medication-assisted therapies — tackling the opioid-overdose epidemic," *New England Journal of Medicine*, 2014, *370* (22), 2063–2066.

A. Data Appendix

ARCOS data construction. The DEA ARCOS data was obtained from the Washington Post's bulk release of the full transaction-level file (2006–2012), containing 178,598,026 rows. Each record identifies a shipment of a controlled substance from a distributor to a retail buyer (pharmacy, hospital, or practitioner), with fields for drug name, NDC code, dosage units, buyer name, buyer state, buyer county, and transaction date. I retain only opioid transactions (DRUG_NAME in {HYDROCODONE, OXYCODONE}) and aggregate to the county-state level. County names in ARCOS are uppercase without standard suffixes; I match to FIPS codes using the Census Bureau's county-to-FIPS crosswalk via the `tigris` R package.

T-MSIS data construction. The CMS T-MSIS Other Services file (2018–2024) contains 227,083,361 rows at the provider-NPI × HCPCS × month level. I identify MAT claims using HCPCS codes H0020 (methadone), J0571–J0575 (buprenorphine injections), and J2315 (naltrexone). Non-opioid SUD claims use codes H0015 and H0016 (group and individual SUD counseling). Provider NPIs are geocoded to counties in two steps: (1) the NPPES API returns the practice ZIP code for each NPI, and (2) the Census 2020 ZCTA-to-county crosswalk assigns each ZIP to its primary county by land area overlap.

Sample restrictions. I restrict to counties with (a) at least 100,000 total ARCOS opioid pill shipments (to ensure meaningful HCP share measurement) and (b) at least one T-MSIS MAT provider geocoded to the county (to ensure the outcome is not mechanically zero due to missing provider data).

B. Robustness Appendix

The leave-one-state-out jackknife tests whether any single state drives the main result. For each of the states in the sample, I re-estimate the preferred specification (Equation 1) omitting that state. The coefficient on HCP share remains stable across all exclusions, confirming that no single state is responsible for the treatment dividend finding.

C. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Total MAT	8.8219	25.8243	35.892	0.0438	0.1281	Small positive
Methadone	1.8945	21.9098	33.291	0.0101	0.1172	Small positive
Buprenorphine	7.1286	5.1541	6.379	0.1990	0.1439	Large positive
Non-opioid SUD (placebo)	-4.5439	3.0738	7.260	-0.1114	0.0754	Moderate negative
Panel B: Heterogeneous (Urban/Rural Split)						
Total MAT (Urban)	-0.1649	17.3376	30.574	-0.0010	0.1025	Null
Total MAT (Rural)	93.5085	96.4309	60.844	0.2376	0.2450	Large positive

Notes: **Country:** United States. **Research question:** Does the 2014 federal rescheduling of hydrocodone combination products from Schedule III to Schedule II affect county-level Medicaid-funded medication-assisted treatment utilization? **Policy mechanism:** The DEA rescheduling eliminated phone-in prescriptions, prohibited refills, and capped supply at 30 days for hydrocodone combination products, which constituted approximately 75% of prescription opioid volume, creating a county-level supply shock proportional to pre-existing hydrocodone dependence. **Outcome definition:** Monthly Medicaid MAT claims (methadone H0020, buprenorphine J0571–J0575, naltrexone J2315) per 1,000 county population, from T-MSIS 2018–2024. **Treatment:** Continuous; county-level hydrocodone combination product share of total ARCOS opioid shipments (2006–2012 average). **Data:** DEA ARCOS transaction-level pill shipments (2006–2012, 178M rows), CMS T-MSIS Medicaid claims (2018–2024, 227M rows), provider NPIs geocoded to counties via NPPES, ACS 2019 county controls. **Method:** Cross-sectional shift-share reduced form with state fixed effects, robust standard errors clustered at state level. **Sample:** US counties with >100,000 ARCOS pill shipments and matched T-MSIS MAT providers; 587 counties across 49 states. SDE $= \hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ for continuous treatment where SD(X) is the cross-county standard deviation of HCP share and SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).