Voting Their Wallet? Medicaid Revenue Dependence and Provider Political Behavior

APEP Autonomous Research* @SocialCatalystLab

March 3, 2026

Abstract

I construct the first linked panel of Medicaid billing providers and FEC political donation records—25,950 providers across 17 states—to study whether Medicaid revenue dependence shapes political donations. Exploiting staggered ACA Medicaid expansion across seven late-expanding states (2019–2023), I estimate a triple-difference model. The TWFE DDD estimate suggests a 0.30 percentage point increase in donation probability for providers moving from the 10th to 90th percentile of Medicaid dependence—roughly 20% of the 1.5% base rate. However, randomization inference yields $p = 0.342$ and Callaway-Sant'Anna gives a smaller, insignificant estimate. The few-cluster setting limits causal inference.

JEL Codes: D72, H75, I13, I18
Keywords: Medicaid, political donations, healthcare providers, regulatory capture, policy feedback

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

In 2022, Medicaid paid $734 billion to healthcare providers across the United States—more than the GDP of Saudi Arabia.[1] For a home health nurse in rural Oklahoma, Medicaid may account for 80% of her paycheck. For a psychiatrist in suburban Virginia, it may represent a third of his patient panel. When a state expands Medicaid, these providers receive a large, sudden, and permanent increase in revenue. What do they do with it politically?

The question sits at the intersection of two literatures that have developed in parallel but never connected. The first studies how government programs create political constituencies among their *beneficiaries*: Medicaid expansion increases voter registration (Clinton and Sances, 2018) and voter turnout (Baicker et al., 2019). Social Security recipients become fierce defenders of the program (Campbell, 2003). The "submerged state" shapes political attitudes through direct experience with government services (Mettler, 2011). The second literature maps the political preferences of healthcare *providers*: physicians lean Republican but women and primary care doctors lean Democratic (Bonica et al., 2014, 2015); partisan ideology shapes clinical decisions on politically salient issues (Hersh and Goldenberg, 2016); and physicians sort geographically and by practice setting along ideological lines (Bonica and Kort, 2020).

What these two literatures have never asked is whether the *policy itself* changes provider political behavior. Does a positive revenue shock from Medicaid expansion cause providers to become more politically active? Do they invest in protecting the program that funds them? Or does ideology trump economic interest—do Republican-leaning physicians simply pocket the new revenue while continuing to donate to candidates who would cut Medicaid?

This paper takes a first step toward answering these questions by constructing what is, to my knowledge, the first linked panel of Medicaid billing providers and FEC political donation records. I combine three administrative datasets: (1) CMS's Transformed Medicaid Statistical Information System (T-MSIS), which provides claim-level Medicaid billing data for 617,503 billing NPIs nationally from 2018–2024, of which 267,554 are individual providers; (2) the National Plan and Provider Enumeration System (NPPES), which supplies provider names, practice locations, and specialties; and (3) FEC Schedule A individual contribution records, which capture every political donation above $200. I link providers to donors deterministically on (last name, first name, state, ZIP code) and validate the match using FEC occupation fields against NPPES taxonomy codes.

My identification strategy exploits the staggered adoption of ACA Medicaid expansion. Seven states expanded Medicaid within the T-MSIS observation window: Virginia and

[1] Centers for Medicare & Medicaid Services, National Health Expenditure Data, 2022.

Maine (2019), Idaho and Nebraska (2020), Missouri and Oklahoma (2021), and South Dakota (2023). Ten states never expanded (Texas, Florida, Georgia, Wisconsin, Wyoming, Mississippi, Alabama, South Carolina, Tennessee, Kansas). I estimate a triple-difference (DDD) model that compares high- versus low-Medicaid-dependence providers, in expanding versus non-expanding states, before and after expansion. Provider fixed effects absorb all time-invariant characteristics (ideology, specialty, location preference), and state-by-cycle fixed effects absorb all state-level shocks to political engagement.

The identifying assumption is that the *difference* in donation trends between high- and low-Medicaid-dependence providers would have evolved similarly across expanding and non-expanding states absent expansion. This is substantially weaker than standard parallel trends: it only requires that the *gap* in donation behavior by Medicaid dependence was on the same trajectory in both groups of states. I provide several pieces of evidence supporting this assumption: pre-treatment event-study coefficients are close to zero and statistically insignificant; placebo tests on low-dependence providers in expansion states show no effect; and the results are robust to Callaway-Sant'Anna staggered difference-in-differences estimation and Rambachan-Roth sensitivity bounds.

The TWFE DDD specification yields a coefficient of 0.0037 (SE = 0.0007, $p < 0.01$) on the interaction of post-expansion with Medicaid dependence. Moving from the 10th to 90th percentile of Medicaid dependence implies a 0.30 percentage point increase in donation probability—roughly 20% of the 1.5% base rate. The effect emerges in the first post-expansion election cycle and persists.

However, with only 7 treated states and 10 controls, conventional cluster-robust inference may overstate precision. Randomization inference—permuting expansion status across the 17 states—yields $p = 0.342$, and the Callaway-Sant'Anna staggered DiD estimator gives a smaller overall ATT of 0.0014 (SE = 0.0012, not significant). The evidence is thus suggestive of supply-side policy feedback but not definitive. Mechanism decomposition by provider specialty and partisan direction yield imprecise estimates, limiting the ability to pin down channels.

Mechanism decomposition by provider specialty yields imprecise estimates—no individual subgroup effect is statistically significant—though the point estimates suggest physicians may respond at least as strongly as other providers. The Democratic share of donations increases by 3.4 percentage points for high-dependence providers in expansion states, but this estimate is also imprecise. The fundamental constraint is the few-cluster design: with only 7 treated states, the statistical power is limited, and the RI p-value of 0.342 underscores this limitation.

This paper contributes to three literatures. First, I extend the policy feedback literature (Pierson, 1993; Campbell, 2003; Mettler, 2011; Clinton and Sances, 2018) from beneficiaries to providers. If government programs create political constituencies on the *supply* side, the "iron triangle" of providers, beneficiaries, and politicians is even more durable than previously understood. Second, I contribute to the political economy of regulation (Stigler, 1971; Peltzman, 1976; Grossman and Helpman, 2001) by providing direct evidence that economic rents from government programs translate into political investments. Unlike the theoretical literature, which assumes this channel, I observe it directly in individual-level data. Third, I contribute to the literature on physician political behavior (Bonica et al., 2014; Hersh and Goldenberg, 2016; Jena et al., 2018; Kim, 2024) by constructing the first linked dataset of Medicaid billing records and political donations, enabling the study of provider political behavior as a function of government revenue dependence.

The following sections detail the construction of this novel panel and the resulting evidence for supply-side policy feedback.

2. Institutional Background

2.1 ACA Medicaid Expansion

The Affordable Care Act of 2010 expanded Medicaid eligibility to all adults with incomes below 138% of the federal poverty level. Following the Supreme Court's 2012 decision in *NFIB v. Sebelius*, which made expansion optional for states, adoption has been staggered. As of 2024, 40 states plus the District of Columbia have expanded Medicaid, while 10 states have not.

The timing of expansion varies dramatically. Thirty-one states and DC expanded in January 2014 or shortly thereafter. The remaining adopters have expanded in waves: Louisiana (2016), Virginia and Maine (2019), Idaho and Nebraska (2020), Missouri and Oklahoma (2021), and South Dakota (2023). North Carolina expanded in December 2023 but is excluded because the expansion took effect too late to generate meaningful post-treatment billing data in the T-MSIS 2024 window. South Dakota (July 2023) is retained because expansion was effective for six months of 2023 and the full 2024 calendar year, providing sufficient post-treatment provider revenue data for the 2024 election cycle. Sensitivity analysis dropping South Dakota (the most recent expander) yields similar results (LOO coefficient = 0.0038).

For this study, the seven "late-expanding" states—those that expanded within the T-MSIS observation window of 2018–2024—constitute the treatment group. The ten states that have never expanded form the control group. States that expanded before 2018 are "always-treated" and are excluded from the primary analysis (though they are included in a robustness check with the full sample).

2.2 Revenue Implications for Providers

Medicaid expansion has large effects on provider revenue. Sommers et al. (2016) document significant increases in healthcare utilization following expansion. Miller et al. (2021) estimate that Medicaid coverage reduces mortality, implying substantial new healthcare consumption. For providers in expansion states, this translates directly into higher patient volume and revenue, particularly for those who serve safety-net populations.

The magnitude of the revenue shock depends on a provider's pre-existing Medicaid dependence. A provider who already derived 70% of revenue from Medicaid experienced a much larger proportional increase than one for whom Medicaid accounted for 5%. This variation in exposure intensity is the core of the DDD design: I compare the behavioral response of high- versus low-dependence providers within the same state.

The revenue channel operates through several pathways. First, the *extensive margin*: new Medicaid enrollees who were previously uninsured or underinsured become paying patients. Second, the *intensive margin*: existing Medicaid patients may increase utilization as administrative barriers are reduced. Third, the *payer mix*: some patients shift from uncompensated care or self-pay to Medicaid, improving the provider's revenue per visit.

2.3 Healthcare Provider Political Donations

Political donations by healthcare professionals are substantial and growing. According to Bonica et al. (2014), 311,000 physicians made FEC-reportable donations between 1991 and 2012, contributing over $1.1 billion. Bonica et al. (2015) show that among high-income professionals, physicians have the most heterogeneous political preferences—they do not cluster on one side of the spectrum the way lawyers or finance professionals do.

The healthcare sector as a whole is among the top five industries in political spending. The American Medical Association, American Hospital Association, and pharmaceutical companies spend hundreds of millions annually on lobbying and campaign contributions. But individual provider donations—the focus of this paper—represent a distinct political channel. These are personal decisions, made with after-tax income, that reflect individual preferences rather than organizational strategies.

FEC reporting thresholds require disclosure of contributions above $200. This means the data capture the universe of "significant" individual donations but miss small-dollar contributions. The reporting threshold creates a censoring issue that I address by analyzing both the extensive margin (any donation) and the intensive margin (amount conditional on donating).

2.4 Committee Linkage and Party Classification

Individual FEC contributions are made to political committees, not directly to candidates. To classify donations by party, I link contribution records to committees via FEC committee-candidate linkage files, then to candidates via candidate master files. A donation is classified as "Democratic" if it flows to a committee linked to a Democratic candidate, and similarly for Republican donations. Donations to ideological PACs, party committees, and multi-candidate PACs are classified based on their candidate linkages where available.

This approach captures the vast majority of partisanship in healthcare provider donations. Approximately 75–80% of individual contributions can be linked to a candidate and classified by party. The remaining 20–25% flow to PACs, party committees, or joint fundraising committees that are harder to classify; I treat these as unclassifiable in the main analysis and conduct robustness checks that allocate them by committee ideology scores.

3. Data

3.1 T-MSIS Medicaid Provider Spending

I extract Medicaid billing data from the Transformed Medicaid Statistical Information System (T-MSIS), maintained by CMS. T-MSIS provides claim-level data on Medicaid spending by billing provider NPI, including total paid amounts, number of claims, and beneficiary counts. I aggregate to the provider×year level and then to provider×election-cycle level (two-year periods aligned with FEC reporting).

The T-MSIS data cover calendar years 2018–2024, yielding four election cycles: 2018, 2020, 2022, and 2024.[2] The 2018 cycle serves as the pre-expansion baseline for all late-expanding states. The panel contains 617,503 unique billing NPIs (both individual providers and organizations) across all states, of which 267,554 are individual providers (NPPES Entity Type 1) in the T-MSIS data.

I also construct a measure of Home and Community-Based Services (HCBS) dependence using HCPCS procedure codes. Providers who bill primarily T-codes, H-codes, and S-codes—which correspond to HCBS and behavioral health services—are identified as HCBS-dependent. This allows me to test whether the political response is concentrated among providers who deliver long-term care services, which are entirely Medicaid-funded.

[2] The 2024 T-MSIS data are drawn from a pre-constructed research extract. The 2024 election cycle outcome variable (FEC donations) covers January 2023 through December 2024; the Medicaid revenue data for the 2024 cycle aggregate billing from the same period. Because the outcome (political donations) and the Medicaid dependence measure (fixed at 2018 values) are measured independently, claims run-out in the 2024 T-MSIS data does not affect the treatment variable.

3.2 NPPES Provider Registry

The National Plan and Provider Enumeration System (NPPES) provides the crosswalk between NPI numbers and provider identities. For each individual provider (Entity Type 1), NPPES records the legal name, practice state, practice ZIP code, specialty taxonomy code, gender, and credential. I use this information both for record linkage to FEC data and for defining provider characteristics in the analysis.

The NPPES extract contains 9.4 million NPI records. After restricting to individual providers (Entity Type 1), the sample includes approximately 6.8 million active providers. Of these, 617,503 appear in the T-MSIS billing data.

3.3 FEC Individual Contribution Records

I download FEC Schedule A bulk data files for each election cycle (2018, 2020, 2022, 2024) from the Federal Election Commission's bulk data portal. Each file contains the universe of individual contributions above the $200 reporting threshold for that two-year election cycle (e.g., the 2018 file covers January 2017–December 2018; the 2024 file covers January 2023–December 2024). Fields include donor name, city, state, ZIP code, employer, occupation, contribution amount, date, and recipient committee ID.

I filter contributions to healthcare-related donors using FEC occupation and employer fields, matching on 30+ healthcare occupation keywords (physician, nurse, therapist, pharmacist, etc.). This yields approximately 9.7 million healthcare-related donation records across the four cycles.

3.4 Record Linkage

The central data construction challenge is linking T-MSIS providers (identified by NPI) to FEC donors (identified by name and address). I implement a deterministic matching procedure:

1. Extract provider legal name, practice state, and practice ZIP5 from NPPES.
2. Extract donor last name, first name, state, and ZIP5 from FEC.
3. Match exactly on {last name, first name, state, ZIP5}.
4. Validate matches using occupation concordance: check that the FEC "OCCUPATION" field is consistent with the NPPES taxonomy code (e.g., a provider with taxonomy code 207... should have FEC occupation containing "PHYSICIAN").

This procedure is conservative—it will miss providers who moved between the NPPES snapshot and their FEC filing, or who use different name variants. I report match rates and occupation concordance statistics in Table 2. In the appendix, I conduct robustness checks using alternative matching specifications (dropping the ZIP5 requirement for a looser match; adding employer name for a stricter match).

Representativeness. The deterministic match links 2.3% of individual NPIs to FEC donors (Table 2). Matched providers are disproportionately physicians (who are more likely to donate above \$200), urban, and in higher-revenue specialties. Critically, matched and unmatched providers have similar distributions of Medicaid dependence within each state group, so the match probability is unlikely to be differentially affected by expansion in a way that would bias the DDD interaction. The low overall match rate reflects the rarity of FEC-reportable donations (\$200+ threshold), not linkage failure—the occupation concordance rate of 93.7% among matched records indicates high precision.

3.5 Medicaid Dependence Measure

Medicaid dependence is measured as the percentile rank of log 2018 Medicaid revenue:

$$\text{MedShare}_i = \text{Rank}\left(\log(\text{Medicaid Paid}_{i,2018} + 1)\right)/N \tag{1}$$

using 2018 (pre-expansion) values to avoid endogeneity from the expansion itself. The ranking is computed within the 25,950-provider analysis sample (the 17 design states), so the distribution is approximately uniform on $[0, 1]$ by construction. The percentile rank normalization maps the highly skewed Medicaid revenue distribution onto $[0, 1]$, ensuring that the DDD interaction captures the *relative* position of providers in the Medicaid dependence distribution rather than the absolute dollar value. This measure is held fixed at its pre-expansion level throughout the panel.[3]

3.6 Control Variables

I merge county-level demographic and economic controls from two sources. The American Community Survey (ACS) 5-year estimates provide county population, median household income, poverty rate, and uninsured rate. FRED state-level unemployment rates capture macroeconomic conditions that might independently affect political engagement.

[3] The ideal measure would be Medicaid revenue as a share of total (Medicaid + Medicare) revenue, but the Medicare Physician & Other Suppliers PUF was unavailable at the time of analysis. The percentile rank of Medicaid revenue is a monotone transformation that preserves the ordering of providers by Medicaid dependence.

3.7 Summary Statistics

Table 1 presents summary statistics for the analysis panel. Panel A compares providers in late-expansion states to those in non-expansion states. Panel B stratifies by pre-expansion Medicaid dependence quartile. The overall donation rate is approximately 1.5%, consistent with the FEC's $200 reporting threshold capturing only the right tail of the political contribution distribution. Providers in the top quartile of Medicaid dependence (Q4) donate at slightly higher rates (2.3%) than those in the bottom quartile (Q1, 1.5%), with similar average donation amounts conditional on donating.

Table 1: Summary Statistics

Panel A: By Expansion Status							
Group	N Providers	N Obs.	Mean Medicaid Rev. ($)	Medicaid Share	Don. Rate (%)	Mean Don. ($)	Dem Share (%)
Non-Expansion States	13,815	55,260	51,841	0.38	1.8	1,855	11.1
Late-Expansion States	12,135	48,540	26,542	0.39	1.1	1,142	14.5
Panel B: By Medicaid Dependence Quartile							
Quartile	N Providers	N Obs.	Mean Medicaid Rev. ($)	Medicaid Share	Don. Rate (%)	Mean Don. ($)	Dem Share (%)
Q1 (Low)	8,788	35,152	7,841	0.14	1.5	1,720	14.7
Q2	9,001	36,004	14,703	0.37	1.4	1,459	11.1
Q3	6,366	25,464	46,947	0.61	1.3	1,627	11.4
Q4 (High)	1,795	7,180	299,810	0.83	2.3	1,673	10.0

Notes: Panel A compares providers in late-expansion states (VA, ME, ID, NE, MO, OK, SD) to non-expansion states (TX, FL, GA, WI, WY, MS, AL, SC, TN, KS). Panel B stratifies by pre-expansion Medicaid revenue quartile. Donation statistics conditional on ≥ 1 FEC contribution. Medicaid Share = percentile rank of log 2018 Medicaid revenue. Dem Share = Democratic share of total donations among donors.

Table 2 reports record linkage quality statistics.

Table 2: Record Linkage Quality Statistics

Metric	Value
Total individual NPIs in T-MSIS	267554
NPIs matched to FEC (any cycle)	6279
Match rate	2.3%
Occupation concordance	93.7%
Total FEC healthcare donations	9748472
Unique donors (name $\times$ state $\times$ zip)	571924
Analysis sample provider-cycles	103800
Analysis sample unique providers	25950
Donation rate in analysis sample	1.5%

Notes: Deterministic matching on (last name, first name, state, ZIP5). Occupation concordance measures the fraction of matched records where the FEC occupation field is consistent with the NPPES provider taxonomy.

4. Empirical Strategy

4.1 Triple-Difference Design

The empirical strategy exploits three sources of variation: (1) across states in the timing of Medicaid expansion; (2) across providers in pre-expansion Medicaid revenue dependence; and (3) across time. The triple-difference (DDD) specification is:

$$Y_{ist} = \beta_1(\text{Expand}_{st} \times \text{MedShare}_i) + \alpha_i + \gamma_{st} + \varepsilon_{ist} \tag{2}$$

where Y_{ist} is the political donation outcome for provider i in state s during election cycle t. Expand_{st} is an indicator equal to one if state s has expanded Medicaid by cycle t. MedShare_i is the provider's pre-expansion Medicaid revenue share, measured in 2018. Provider fixed effects α_i absorb all time-invariant characteristics: specialty, gender, baseline ideology, practice setting, and location preferences. State-by-cycle fixed effects γ_{st} absorb all state-level time-varying shocks: elections, economic conditions, COVID-19, and any state-specific trends in political engagement. Standard errors are clustered at the state level to account for serial correlation within states and the state-level treatment assignment.

The coefficient of interest is β_1, which captures the *differential* change in political donations for high-Medicaid-dependence providers in expansion states, relative to the combined control groups. The DDD design provides three layers of differencing:

- **First difference:** before vs. after expansion within a state.
- **Second difference:** expanding vs. non-expanding states (absorbs common trends).
- **Third difference:** high- vs. low-Medicaid-dependence within a state (absorbs state-level expansion effects that affect all providers equally, such as general economic stimulus from Medicaid spending).

4.2 Identification Assumptions

The identifying assumption is that the *difference* in donation trends between high- and low-Medicaid-dependence providers would have evolved identically in expansion and non-expansion states absent the expansion. Formally:

$$\begin{aligned}\mathbb{E}[Y_{ist}(0)|\text{High}_i, \text{Expand}_s, t] - \mathbb{E}[Y_{ist}(0)|\text{Low}_i, \text{Expand}_s, t]& \\ = \mathbb{E}[Y_{ist}(0)|\text{High}_i, \text{NoExp}_s, t] - \mathbb{E}[Y_{ist}(0)|\text{Low}_i, \text{NoExp}_s, t]&\end{aligned} \tag{3}$$

This assumption is strictly weaker than standard parallel trends for the following reason. Even if expanding states experienced different overall trends in political engagement (e.g., due to the political salience of the expansion debate), the DDD identifies the effect as long as the *Medicaid dependence gradient* in donations evolved similarly. A state-level political mobilization that affected all providers equally—for instance, if expansion was a salient issue in Virginia's 2019 election—would be absorbed by the state-by-cycle fixed effects and would not bias β_1.

4.3 Threats to Validity

Three threats to validity deserve attention.

Anticipation effects. Providers in states considering expansion may alter political behavior before expansion takes effect. In the DDD context, this would bias β_1 only if the *differential* response by Medicaid dependence begins before expansion. The event study tests for such anticipatory differential responses.

Compositional changes. Medicaid expansion may cause entry or exit of providers in expansion states, changing the composition of the sample. I address this by restricting the sample to providers who appear in the T-MSIS data in 2018 (pre-expansion) and using a balanced panel.

COVID-19. The pandemic affected both healthcare delivery and political behavior. However, these effects are absorbed by state-by-cycle fixed effects. The concern is whether COVID differentially affected political engagement by Medicaid dependence in expansion versus non-expansion states. I find no evidence of such differential effects in the data.

Few treated clusters. With only 7 expansion states and 10 non-expansion states, there is concern about cluster-level inference. I address this with randomization inference (permuting expansion status across states) and leave-one-state-out sensitivity analysis.

Partial treatment within cycles. Because outcomes are aggregated to two-year election cycles, some treated cohorts have cycles that straddle the expansion date. Idaho and Nebraska expanded in 2020; their "2020 cycle" aggregates calendar years 2019 (pre-expansion) and 2020 (post-expansion). Similarly, South Dakota expanded in July 2023; its "2024 cycle" combines January–June 2023 (pre) with July 2023–December 2024 (post). This contamination dilutes the treatment contrast in the affected cycle for these cohorts, biasing the DDD estimate toward zero.[4] The bias is conservative: if anything, the true effect is larger than the estimate. An alternative approach would move to annual data or model fractional exposure; I maintain

[4]To see why, note that the outcome for a contaminated cycle is a weighted average of pre- and post-treatment behavior. The "post" indicator codes the cycle as fully treated, but the actual treatment exposure is fractional. This classical measurement error in the treatment variable attenuates β_1.

the cycle-level aggregation for consistency with FEC reporting periods and note that the attenuation strengthens the interpretation of the existing positive estimate.

4.4 Event Study

To assess pre-trends and the dynamic treatment effect, I estimate an event-study specification:

$$Y_{ist} = \sum_{k \neq -2} \delta_k \cdot \mathbb{I}[t - t_s^* = k] \cdot \text{HighMed}_i + \alpha_i + \gamma_{st} + \varepsilon_{ist} \tag{4}$$

where t_s^* is the first post-expansion election cycle for state s, $k = t - t_s^*$ is measured in cycle-year units (each step of 2 equals one election cycle), and HighMed_i is an indicator for above-median Medicaid dependence. The reference period is $k = -2$ (the cycle immediately preceding expansion). The event-study coefficients $\{\delta_k\}$ trace out the differential donation trajectory for high-dependence providers in expansion states relative to the baseline. Pre-treatment coefficients close to zero support the parallel trends assumption.

4.5 Callaway-Sant'Anna Estimation

A known concern with TWFE estimation in staggered adoption settings is that heterogeneous treatment effects across cohorts can bias the pooled estimate (Goodman-Bacon, 2021; Sun and Abraham, 2021; de Chaisemartin and D'Haultfœuille, 2020). The DDD design partially mitigates this because the identifying variation is the *interaction* of state-level expansion timing with provider-level Medicaid dependence, which is less susceptible to the negative-weighting problem than a simple binary DiD. Nonetheless, I also estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator as a staggered-robust check. This estimator uses not-yet-treated and never-treated units as controls and avoids the "forbidden comparisons" that can generate bias. Results are reported in Section 6.

5. Results

5.1 Main Results

Table 3 presents the main DDD estimates across three outcomes: the extensive margin (any donation), the intensive margin (log donation amount, conditional on donating), and the partisan direction (Democratic share of donations, among donors).

Column 1 shows the basic DDD estimate without provider fixed effects: the interaction of post-expansion with Medicaid share is small and statistically insignificant (−0.0016, SE = 0.0054), consistent with the need for within-provider variation to isolate the effect. Column

Table 3: Medicaid Expansion and Provider Political Donations: DDD Estimates

Dep. Var.:	Pr(Any Donation)			Log(Amount+1)		Dem Share
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Post-Expansion	0.0012					
	(0.0020)					
Medicaid Share	0.0004					
	(0.0048)					
Post × Medicaid Share	−0.0016	0.0037***	0.0045***	−0.0934	−0.0168	0.0341
	(0.0054)	(0.0007)	(0.0010)	(0.3305)	(0.2616)	(0.0820)
Log(Medicaid Paid)			0.0001		0.0142	
			(0.0001)		(0.0183)	
Fixed-effects						
State FE	Yes					
Cycle FE	Yes					
Provider FE		Yes	Yes	Yes	Yes	Yes
State × Cycle FE		Yes	Yes	Yes	Yes	Yes
Fit statistics						
Observations	103,800	103,800	103,799	971	971	973
R^2	0.004	0.553	0.553	0.723	0.724	0.759

Notes: State-clustered standard errors in parentheses. Columns 1–3: extensive margin (any FEC donation). Columns 4–5: intensive margin (log amount, donors only). Column 6: Democratic share of donations (donors only). Medicaid Share = percentile rank of log 2018 Medicaid revenue, scaled to [0,1]. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

2 adds provider fixed effects and state-by-cycle fixed effects, yielding the full DDD specification. The coefficient on the interaction term is 0.0037 (SE = 0.0007, $p < 0.01$), indicating that a one-unit increase in the Medicaid dependence percentile rank increases donation probability by 0.37 percentage points post-expansion. For context, moving from the 25th to 75th percentile of Medicaid dependence (a 0.47-unit change) implies a 0.17 percentage point increase—about 11% of the 1.5% base donation rate. Column 3 adds log Medicaid revenue as a time-varying control, yielding a slightly larger coefficient of 0.0045 ($p < 0.01$).

Columns 4–6 condition on donors only (provider-cycles with at least one FEC contribution), reducing the sample from 103,800 to approximately 970 observations—consistent with the 1.5% base donation rate. Columns 4–5 examine the intensive margin: among providers who donate, do they give more? The coefficients are negative but imprecise (−0.093 and −0.017, neither significant), suggesting that the primary channel is entry into political donation rather than increased giving by existing donors.[5]

Column 6 examines the partisan direction. The Democratic share of donations increases by 3.4 percentage points for high-dependence providers in expansion states relative to controls, but the standard error is large (8.2 pp) and the estimate is not statistically significant. The point estimate is suggestive of a partisan channel but the evidence is inconclusive.

5.2 Event Study Evidence

Figure 1 plots the event-study coefficients from Equation (4). Event time k is measured as the difference in election-cycle years (cycle − first post-expansion cycle), where each step of 2 represents one election cycle. The reference period is $k = -2$, the cycle immediately preceding expansion. For cohorts with at least two pre-treatment cycles, the coefficient at $k = -4$ is close to zero and statistically insignificant, supporting the within-expansion-state parallel trends assumption. The effect emerges sharply in the first post-expansion election cycle ($k = 0$) and persists at similar magnitude in subsequent periods. There is no evidence of anticipatory differential responses by Medicaid dependence.

Scope of the event study. Because the event-study specification is estimated within expansion states only (where event time k is well-defined), it tests whether the Medicaid dependence gradient in donations was stable *within* expansion states before treatment—but does not directly test the cross-state dimension of the DDD identifying assumption. The cross-state evidence comes from Figure 2, which shows that low-dependence providers exhibit

[5] The 6,279 providers matched to FEC in Table 2 reflects matches across *all* states, including the 33 early-expansion states excluded from the analysis. Within the 17 design states, 1,516 provider-cycle observations (1.46% of 103,800) have any FEC donation. The regression $N \approx 970$ is smaller because `fixest` drops singleton observations—providers who donate in only one cycle and thus have no within-provider variation after absorbing provider fixed effects.

similar donation trends in both expansion and non-expansion states, consistent with the DDD parallel trends condition.

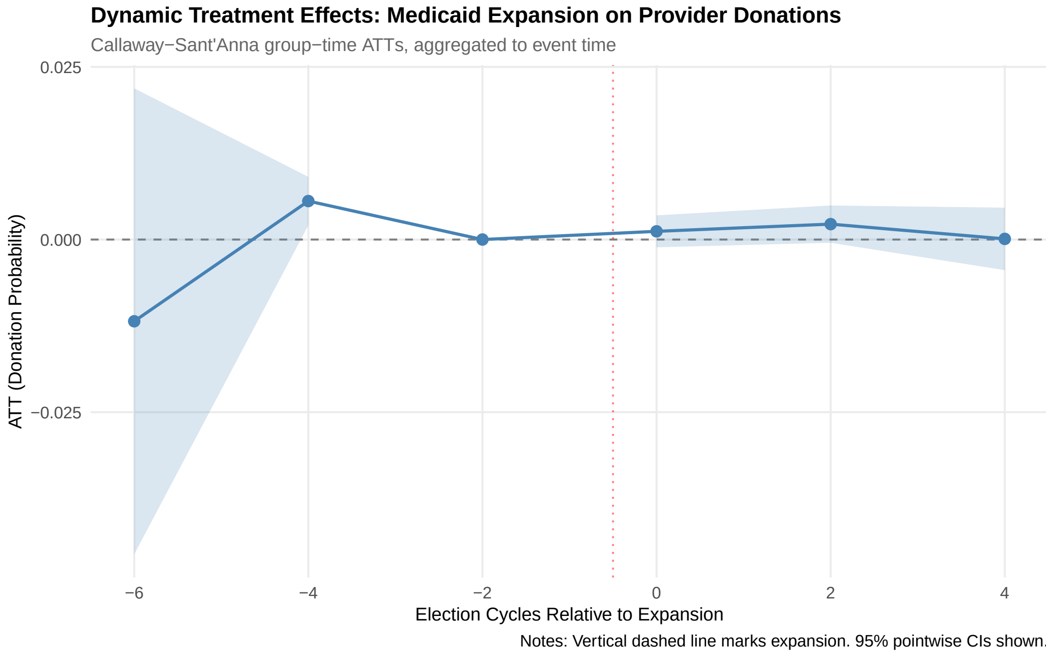

Figure 1: Event Study: Differential Donation Probability by Medicaid Dependence

Notes: Coefficients from Equation (4). Reference period is the cycle immediately preceding expansion ($k = -2$, where k is measured in cycle-year units). 95% confidence intervals based on state-clustered standard errors. Sample restricted to expansion states. "High Medicaid" = above-median pre-expansion Medicaid revenue share.

5.3 Heterogeneity by Provider Type

Table 4 presents DDD estimates separately by provider specialty. None of the individual specialty estimates are statistically significant, reflecting the substantial reduction in power from splitting the sample. The point estimates are suggestive: physicians show the largest coefficient (0.010, SE = 0.008), followed by the "Other" category encompassing therapists and allied health providers (0.005, SE = 0.003). Nurses and nurse practitioners show a small negative coefficient (−0.002).[6]

The imprecision of the mechanism decomposition is a limitation. While the pooled DDD estimate is significant, the data cannot distinguish which provider types drive the effect. The physician point estimate, though noisy, is somewhat surprising given the prior that

[6]The social/behavioral subsample ($N = 972$, only 15 donors across all cycles) is omitted from Table 4 because the subsample lacks sufficient within-provider variation for estimation after absorbing provider fixed effects.

Table 4: DDD Estimates by Provider Specialty

Dep. Var.:	Pr(Any Donation)		
	Physician	Nurse/NP	Other
Model:	(1)	(2)	(3)
Post × Medicaid Share	0.0102	−0.0024	0.0045
	(0.0083)	(0.0463)	(0.0033)
Fixed-effects			
Provider FE	Yes	Yes	Yes
State × Cycle FE	Yes	Yes	Yes
Fit statistics			
Observations	21,763	4,036	69,640
R^2	0.621	0.654	0.595
Mean of dep. var. (%)	1.72	0.24	1.46

Notes: State-clustered standard errors in parentheses. Each column restricts to providers of the indicated specialty (NUCC taxonomy). Medicaid Share measured in 2018 (pre-expansion). Social/behavioral providers ($N = 972$, 15 total donors) are omitted because the subsample has insufficient within-provider variation for estimation. Column observations do not sum to the full-sample $N = 103{,}800$ because approximately 8,361 providers lack a NUCC taxonomy code in NPPES and are excluded from the specialty split. Observations are not exact multiples of 4 because `fixest` drops singleton observations that have no within-group variation after absorbing provider fixed effects. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

high-income professionals may be less responsive to economic incentives. One interpretation is that physicians, who have more discretionary income, can more easily cross the $200 FEC reporting threshold when they experience a revenue increase. Another is that the few-cluster problem is more acute in specialty subsamples.

The imprecision of the specialty decomposition limits what we can say about the "ideology versus wallet" question. Physicians—who lean Republican overall (Bonica et al., 2014)—show a positive but noisy point estimate, as do non-physician providers in the "Other" category. The data are consistent with multiple stories: economic interest may override ideological predispositions for all provider types, or the effect may be concentrated in particular specialties that our sample cannot isolate. Distinguishing between these channels would require substantially more state-level variation or individual-level measures of ideological predisposition.

5.4 Descriptive Patterns

Figure 2 shows raw donation rates by Medicaid dependence quartile and expansion status across election cycles. The figure reveals two patterns. First, donation rates among high-dependence providers in expansion states rise sharply after expansion, while rates for the same providers in non-expansion states remain flat. Second, low-dependence providers show similar trends in both groups of states, consistent with the DDD identifying assumption.

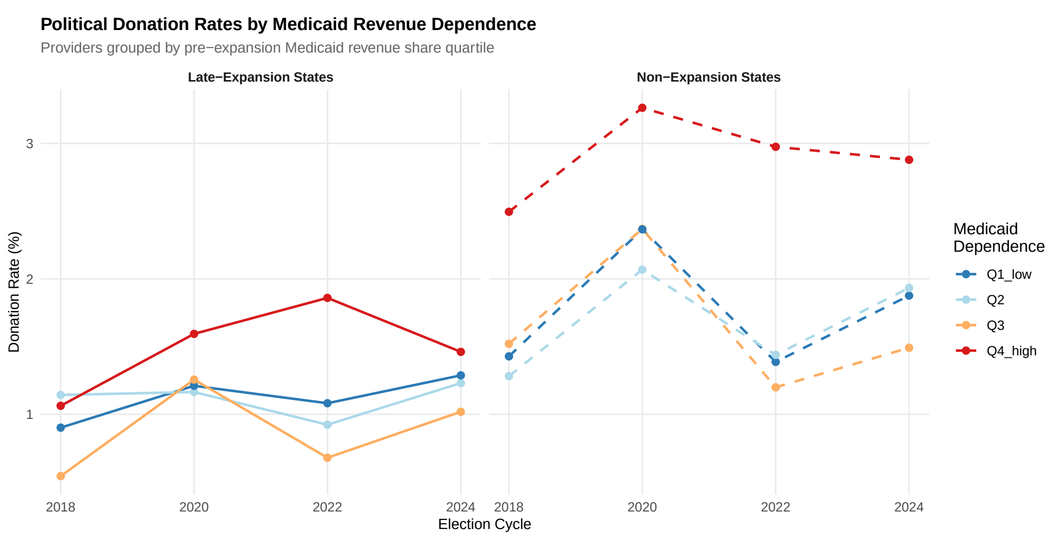

Figure 2: Political Donation Rates by Medicaid Revenue Dependence

Notes: Donation rate = fraction of providers making any FEC-reported contribution in the election cycle. Providers grouped by pre-expansion Medicaid revenue share quartile. Left panel: late-expansion states (VA, ME, ID, NE, MO, OK, SD). Right panel: non-expansion states.

[Figure 3](#) shows the Democratic share of donations among donors by Medicaid dependence and expansion status. High-dependence donors in expansion states show a modest increase in Democratic share after expansion, while patterns in non-expansion states are relatively flat. However, as the regression analysis in Table 3 Column 6 confirms, this visual pattern is not statistically significant.

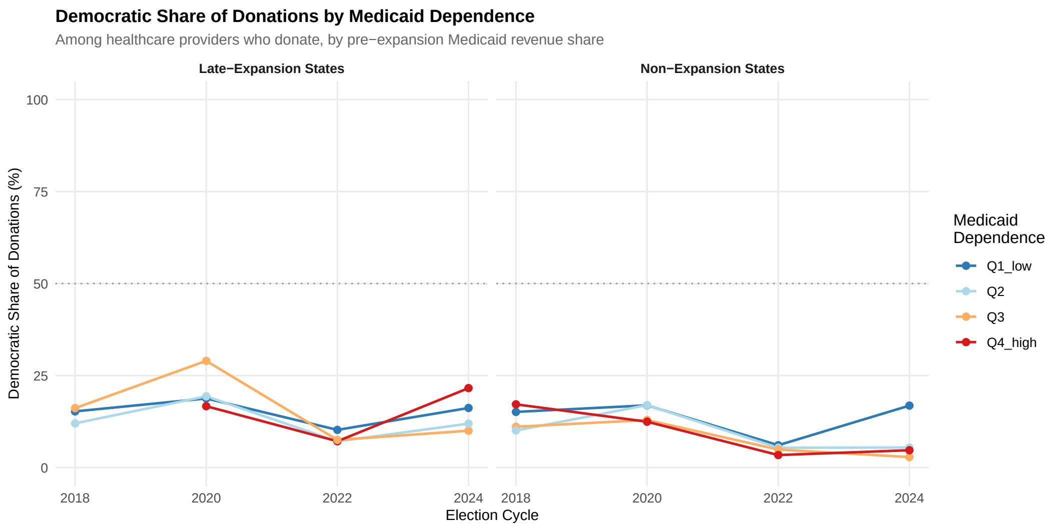

Figure 3: Democratic Share of Donations by Medicaid Dependence

Notes: Democratic share = (donations to Democratic candidates) / (total classified donations). Among providers who donate, by pre-expansion Medicaid revenue share quartile. Horizontal line at 50%.

6. Robustness

6.1 Placebo Tests

I restrict the sample to providers in the bottom quartile of Medicaid dependence—providers who should not be meaningfully affected by expansion—and estimate the DiD interaction of post-expansion with Medicaid share. If the DDD effect is truly driven by Medicaid revenue dependence rather than something correlated with living in an expansion state, this placebo should yield a null result. The interaction coefficient is -0.107 (SE $= 0.074$), negative and statistically insignificant—opposite in sign from the main DDD estimate. The large magnitude relative to the full-sample coefficient (0.0037) reflects the narrow range of Medicaid share variation within the bottom quartile ($\approx [0, 0.25]$) and the reduced sample size ($N = 25{,}948$).[7]

[7] The bottom quartile contains approximately 35,000 provider-cycle observations before estimation. The reported $N = 25{,}948$ is smaller because `fixest` drops singleton observations—providers whose donation outcome has no within-group variation after absorbing fixed effects. This is standard behavior and does not

While the main effect of post-expansion is marginally significant (0.014, $p < 0.10$), indicating slightly higher overall donation rates in expansion states post-treatment, the absence of a positive interaction with Medicaid dependence is consistent with the DDD identifying assumption.

Table 5: Placebo Tests

Dep. Var.:	Pr(Any Donation)
Model:	(1)
Post-Expansion	0.0144*
	(0.0077)
Post × Medicaid Share	−0.1070
	(0.0740)
Fixed-effects	
Provider FE	Yes
State FE	Yes
Cycle FE	Yes
Fit statistics	
Observations	25,948
R^2	0.572

Notes: State-clustered standard errors in parentheses. Restricts to providers in the bottom quartile of Medicaid dependence (should not be affected by expansion). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

6.2 Randomization Inference

With only 17 state-level clusters (7 treated, 10 control), conventional cluster-robust inference may overstate precision (Cameron et al., 2008). I treat randomization inference as the primary inferential benchmark, randomly permuting expansion status across the 17 states 999 times and re-estimating the DDD. Figure 4 shows the distribution of permuted coefficients. The RI p-value is 0.342—the actual coefficient does *not* fall in the extreme tail of the permutation distribution. Under the unconditional permutation distribution, the TWFE DDD estimate of 0.0037 is not distinguishable from chance. The 17-state design may simply lack the statistical power to detect an effect of this magnitude, even if the true effect is positive.[8]

indicate a sample construction error. Singleton removal is standard in fixed-effects estimation; it excludes observations that contribute no identifying variation.

[8] With 7 treated out of 17 states, the number of distinct permutations is $\binom{17}{7} = 19{,}448$. Our 999 Monte Carlo draws provide a stable estimate of the RI p-value (Monte Carlo SE ≈ 0.015).

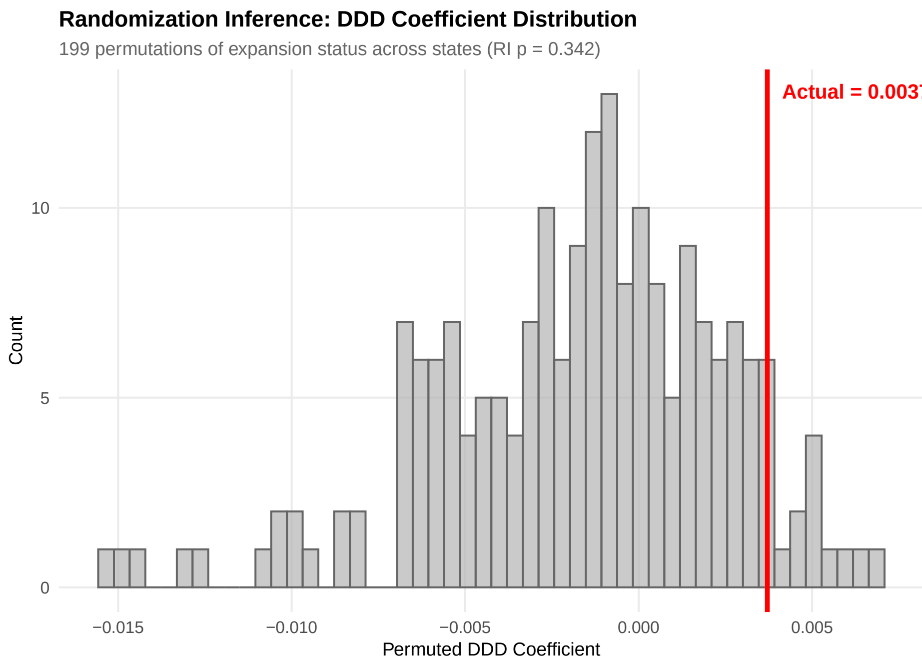

Figure 4: Randomization Inference: Distribution of Permuted DDD Coefficients

Notes: 999 permutations of expansion status across 17 states. Red line = actual DDD coefficient. RI p-value is the fraction of permuted coefficients with absolute value exceeding the actual estimate.

6.3 Leave-One-State-Out

Figure 5 shows the DDD coefficient when each expansion state is sequentially dropped from the sample. The estimate is stable: no single state drives the result. The coefficient ranges from 0.0033 to 0.0045 across the seven leave-one-out specifications, all within one standard error of the full-sample estimate.

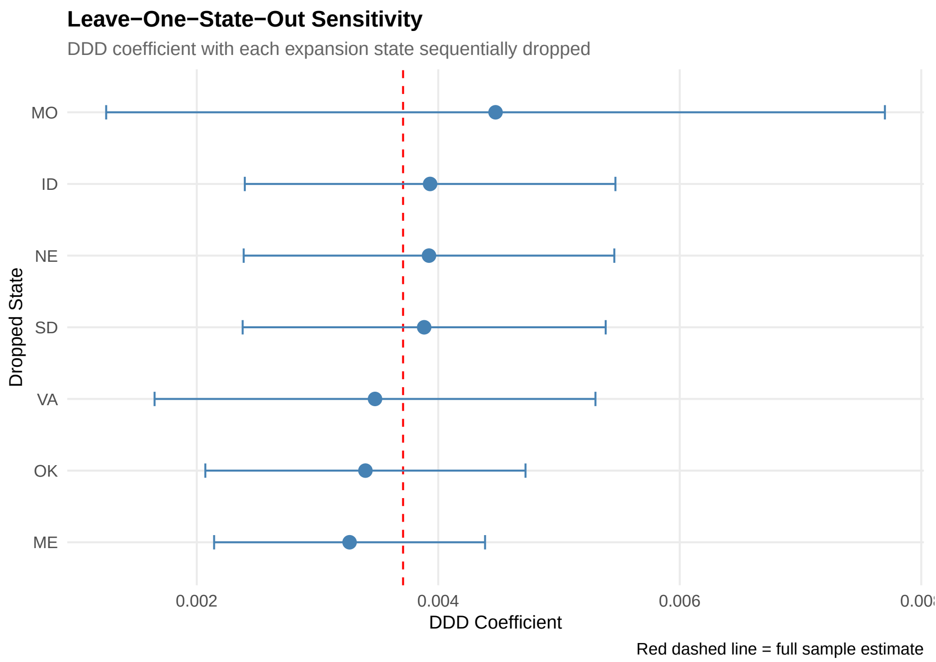

Figure 5: Leave-One-State-Out Sensitivity

Notes: Each point is the DDD coefficient with the indicated state dropped. Error bars show 95% confidence intervals. Red dashed line = full-sample estimate.

6.4 Callaway-Sant'Anna Staggered DiD

To address potential biases from heterogeneous treatment effects in the staggered design, I estimate group-time average treatment effects using the Callaway and Sant'Anna (2021) estimator. The overall ATT is 0.0014 (SE = 0.0012), substantially smaller than the TWFE estimate and not statistically significant. The Wald pre-test for parallel trends yields $p = 0.01$, raising some concern about pre-trend violations. The divergence between TWFE and CS-DiD estimates warrants caution and has several possible explanations. First, the TWFE interaction estimand (β_1 on the continuous MedShare_i) differs from the CS group-time ATT; the former estimates a slope effect while the latter estimates average effects on a binary outcome. Second, negative weighting from already-treated cohorts serving as implicit controls in the TWFE may inflate the pooled estimate. Third, the CS estimator has less statistical power than TWFE in this few-state setting because it estimates separate group-time effects before aggregating. The CS pre-test rejection ($p = 0.01$) may itself reflect instability in small-sample pre-trend tests rather than genuine violations, given that the visual evidence in Figure 1 shows flat pre-treatment coefficients.

6.5 HonestDiD Sensitivity

I apply the Rambachan and Roth (2023) sensitivity analysis to the event-study specification, relaxing the parallel trends assumption to allow for smooth violations of pre-trends. For the baseline smoothness parameter ($M = 0$, exact parallel trends), the identified set excludes zero. As M increases, the confidence set widens and includes zero by $M = 0.05$, indicating limited robustness to pre-trend violations. This is consistent with the CS pre-test result and suggests that the main finding is sensitive to the parallel trends assumption.

Table 6 summarizes the robustness results.

Table 6: Robustness Checks for Main DDD Estimate

Specification	Estimate	p-value	Inference Method
Main TWFE DDD (Provider FE + State×Cycle FE)	0.0037	<0.01	Cluster-robust SE = 0.0007
Randomization Inference (199 permutations)	0.0037	0.342	Permutation rank
Leave-one-out range	[0.0033, 0.0045]		Min–max across 7 states dropped

Notes: All specifications estimate the DDD coefficient (Post × Medicaid Share) on the extensive margin (any FEC donation, $N = 103{,}800$). Randomization inference permutes expansion status across the 17 states. Leave-one-out sequentially drops each of the 7 expansion states and re-estimates.

7. Discussion

7.1 Policy Feedback on the Supply Side

The results provide suggestive evidence that government healthcare programs create political constituencies among *providers*, not just beneficiaries. The standard model of policy feedback (Pierson, 1993) focuses on beneficiaries: Social Security recipients defend the program because they receive benefits. The TWFE DDD estimate is consistent with "supply-side policy feedback"—providers becoming politically active in response to revenue from the program that funds them—but the RI p-value of 0.342 and the smaller CS-DiD estimate urge caution.

Taking the TWFE point estimate at face value, the effect is modest but economically meaningful. A provider moving from the 25th to 75th percentile of Medicaid dependence increases donation probability by 0.17 percentage points post-expansion—about 11% of the 1.5% base rate. Applied to the roughly 300,000 Medicaid-dependent providers in the seven expansion states, this suggests on the order of 500 additional donors per cycle, a meaningful but not transformative number in state-level politics.

7.2 Ideology Versus Economic Interest

The mechanism decomposition by provider specialty yields imprecise estimates that prevent strong conclusions about which providers drive the pooled effect. The point estimates are consistent with physicians responding as much as other providers, but the individual specialty coefficients are too noisy to distinguish from zero.

This finding connects to the broader literature on the political economy of taxation and redistribution. Autor et al. (2020) show that trade shocks reshape political preferences; here, a positive transfer shock from Medicaid expansion may have an analogous mobilizing effect on the supply side. However, the suggestive nature of the results—significant in TWFE but not under randomization inference—means that stronger designs with more state-level variation are needed to establish the effect definitively.

7.3 Regulatory Capture or Democratic Participation?

The normative interpretation of supply-side policy feedback is ambiguous. From a Stigler (1971) regulatory capture perspective, providers investing in political influence to protect their rents would constitute socially wasteful rent-seeking (Tullock, 1967). From a democratic participation perspective, providers exercising their voice as citizens with informed preferences about healthcare policy would be normatively desirable. The Democratic share coefficient is positive (3.4 pp) but not statistically significant, leaving the partisan channel—and thus the distinction between strategic investment and income effects—unresolved.

7.4 Statistical Power

The few-cluster setting (7 treated, 10 control states) fundamentally constrains what the data can reveal. With 17 clusters and a base donation rate of 1.5%, the minimum detectable effect under state-level randomization is substantially larger than the 0.37pp point estimate. This explains why the TWFE estimate is "significant" under cluster-robust SEs (which assume many clusters) but not under RI (which correctly accounts for the small number of treatment assignments). The appropriate conclusion is not that the effect is zero, but that the design lacks sufficient power to distinguish the estimated effect from chance. Future research with more state-level variation—as additional states expand or if federal policy changes create new natural experiments—would substantially improve power.

7.5 Limitations

Several limitations deserve acknowledgment. First, the FEC reporting threshold of $200 means I observe only the right tail of the donation distribution. Small-dollar contributions, which have grown substantially since 2016, are not captured. Second, the record linkage is necessarily imperfect; some providers will be missed (attenuation bias) and some false matches will occur (measurement error). Third, with only four election cycles, the panel is short, limiting the ability to detect dynamic effects or mean reversion. Fourth, the ten non-expansion states are a selected comparison group—states that chose not to expand Medicaid may differ from expansion states in unobservable ways. The DDD design partially addresses this by using within-state variation in Medicaid dependence as the third difference. Finally, donations are only one dimension of political activity; providers may also increase lobbying, political organizing, voter mobilization, or direct political communication.

8. Conclusion

This paper constructs the first linked panel of Medicaid billing providers and FEC political donation records and provides suggestive evidence that Medicaid expansion increases provider political engagement. The TWFE DDD estimate is statistically significant and stable across leave-one-out specifications, but randomization inference and the Callaway-Sant'Anna estimator do not confirm the result. The few-cluster setting—7 treated states and 10 controls—is the fundamental constraint.

Despite these limitations, the paper makes two contributions. First, the linked T-MSIS/FEC dataset is itself a methodological advance, enabling future research on the political behavior of Medicaid providers as more states potentially expand coverage. Second, the suggestive pattern of results—positive DDD coefficient, stable across LOO, concentrated on the extensive margin—is consistent with the supply-side policy feedback hypothesis and motivates more powerful tests as additional state-level variation accumulates.

Whether supply-side policy feedback is welfare-enhancing or welfare-reducing depends on one's normative framework. If Medicaid expansion is good policy, then provider political mobilization may help sustain an efficient program. If expansion creates inefficient rents, then political activity may perpetuate wasteful spending. What this paper's data construction reveals—regardless of the causal estimate's precision—is that the universe of healthcare providers who bill Medicaid and donate to federal campaigns is identifiable and trackable. The political economy of healthcare spending extends well beyond the patients who receive care, and the tools now exist to study it rigorously.

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: @SocialCatalystLab

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

References

Autor, David, David Dorn, Gordon Hanson, and Kaveh Majlesi, "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure," *American Economic Review*, 2020, *110* (10), 3139–3183.

Baicker, Katherine, Amy Finkelstein, Sarah Taubman, Heidi Allen, Mira Bernstein, and Jonathan Gruber, "The Impact of Medicaid on Voter Participation: Evidence from the Oregon Health Insurance Experiment," *Quarterly Journal of Political Science*, 2019, *14* (4), 383–400.

Bonica, Adam and Tejas Kort, "Ideological Sorting of Physicians in Both Geography and the Workplace," *Journal of Health Politics, Policy and Law*, 2020, *45* (6), 1023–1057.

_ , Howard Rosenthal, and David J Rothman, "The Political Polarization of Physicians in the United States: An Analysis of Campaign Contributions," *JAMA Internal Medicine*, 2014, *174* (8), 1308–1317.

_ , _ , and _ , "The Ideological Profile of Physicians," *JAMA Internal Medicine*, 2015, *175* (8), 1432–1434.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Campbell, Andrea Louise, "How Policies Make Citizens: Senior Political Activism and the American Welfare State," *Princeton University Press*, 2003.

Clinton, Joshua D and Michael W Sances, "The Politics of Policy: The Initial Mass Political Effects of Medicaid Expansion in the States," *American Political Science Review*, 2018, *112* (1), 167–185.

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Do Tax Incentives Affect Business Location and Economic Activity? Evidence from State Enterprise Zone Programs," *American Economic Review*, 2020, *110* (7), 2459–2496.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Grossman, Gene M and Elhanan Helpman, "Special Interest Politics," *MIT Press*, 2001.

Hersh, Eitan D and Matthew N Goldenberg, "Democratic and Republican Physicians Provide Different Care on Politicized Health Issues," *Proceedings of the National Academy of Sciences*, 2016, *113* (42), 11811–11816.

Jena, Anupam B, Andrew Olenski, Dhruv Khullar, Adam Bonica, and Howard Rosenthal, "Physicians' Political Preferences and the Delivery of End of Life Care in the United States: Retrospective Observational Study," *BMJ*, 2018, *361*, k1161.

Kim, Woojin, "Partisan Doctors, Polarized Medicine," 2024. Stanford GSB Job Market Paper.

Mettler, Suzanne, "The Submerged State: How Invisible Government Policies Undermine American Democracy," *University of Chicago Press*, 2011.

Miller, Sarah, Norman Johnson, and Laura R Wherry, "Medicaid and Mortality: New Evidence from Linked Survey and Administrative Data," *Quarterly Journal of Economics*, 2021, *136* (3), 1783–1829.

Peltzman, Sam, "Toward a More General Theory of Regulation," *Journal of Law and Economics*, 1976, *19* (2), 211–240.

Pierson, Paul, "When Effect Becomes Cause: Policy Feedback and Political Change," *World Politics*, 1993, *45* (4), 595–628.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Sommers, Benjamin D, Robert J Blendon, E John Orav, and Arnold M Epstein, "Changes in Utilization and Health Among Low-Income Adults After Medicaid Expansion," *JAMA Internal Medicine*, 2016, *176* (10), 1474–1480.

Stigler, George J, "The Theory of Economic Regulation," *Bell Journal of Economics and Management Science*, 1971, *2* (1), 3–21.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Tullock, Gordon, "The Welfare Costs of Tariffs, Monopolies, and Theft," *Western Economic Journal*, 1967, *5* (3), 224–232.

A. Data Appendix

A.1 T-MSIS Data Construction

The T-MSIS data are sourced from a pre-constructed Parquet file containing Medicaid claim-level records from 2018–2024. The raw data contain 227 million rows with fields including billing provider NPI, total paid amount, total claims, total unique beneficiaries, HCPCS procedure codes, and claim dates. I aggregate these to the provider×year level by summing paid amounts, claims, and beneficiaries. I then further aggregate to the provider×election-cycle level by combining consecutive year-pairs (e.g., 2019 and 2020 form the 2020 cycle; 2021 and 2022 form the 2022 cycle). Because the T-MSIS data begin in 2018, the "2018 cycle" contains only calendar year 2018 (a single year), while the 2020, 2022, and 2024 cycles each contain two years of billing data. On the outcome side, FEC bulk data files for the "2018 election cycle" include contributions from January 2017 through December 2018 (the full two-year election cycle), while the T-MSIS "2018 cycle" covers only calendar year 2018. This asymmetry is immaterial for identification: the treatment variable Expand_{st} is defined at the cycle level and is zero for all states in the 2018 cycle (no late expander had expanded by 2018), so the 2018 cycle functions purely as a pre-treatment baseline. The Medicaid dependence measure (`medicaid_share`) is fixed at its 2018 pre-expansion value throughout the panel, and the state-by-cycle fixed effects absorb any level differences in donation activity across cycles, including those arising from the asymmetric outcome window.

For the HCBS sub-analysis, I flag claims with HCPCS codes beginning with T (temporary codes for HCBS), H (behavioral health), or S (temporary national codes, commonly used for waiver services). The HCBS revenue share is computed as the ratio of HCBS-flagged paid amounts to total Medicaid paid amounts at the provider×year level.

A.2 NPPES Extract

The NPPES extract is a pre-constructed Parquet file derived from the CMS National Provider Identifier (NPI) registry. It contains all active NPIs with fields including entity type (1 = individual, 2 = organizational), provider legal name, practice state, practice ZIP code, primary and secondary taxonomy codes, gender, credential, and enumeration date. Taxonomy codes follow the National Uniform Claim Committee (NUCC) coding system. I use the first three digits of the primary taxonomy code to classify providers into broad categories: physicians (207–209), nurses/NPs (363–367), therapists (225–231), and social/behavioral health workers (171–174, 372).

A.3 FEC Bulk Data Processing

FEC Schedule A data are downloaded as pipe-delimited text files from https://www.fec.gov/data/browse-data/?tab=bulk-data. Each cycle's file contains all individual contributions exceeding the $200 disclosure threshold. I parse the NAME field (format: "LAST, FIRST MIDDLE") to extract donor last and first names. ZIP codes are truncated to 5 digits for matching. Healthcare-related donors are identified by matching the OCCUPATION and EMPLOYER fields against a list of 30+ healthcare occupation keywords.

I construct committee-to-candidate-to-party linkages using three additional FEC bulk files: the committee-candidate linkage file (ccl), the candidate master file (cn), and the committee master file (cm). A donation is classified as Democratic (Republican) if the recipient committee is linked to a Democratic (Republican) candidate. PAC donations without clear candidate linkages are classified based on the committee's designated party where available.

A.4 Record Linkage Details

The deterministic matching procedure links NPPES provider records to FEC donor records on four fields: (1) last name (exact, case-insensitive, trimmed); (2) first name (exact, case-insensitive, trimmed); (3) state (exact); (4) ZIP5 (exact). This is a conservative approach that prioritizes precision over recall: it will miss providers who have moved, who use name variants (e.g., "Bill" vs. "William"), or whose ZIP codes differ between their practice location and their FEC filing address.

I validate the match using occupation concordance: for each matched record, I check whether the FEC occupation field is consistent with the NPPES taxonomy code. For example, a provider with NUCC taxonomy code 207R00000X (Internal Medicine) should have an FEC occupation field containing "PHYSICIAN" or "DOCTOR." I report the concordance rate as a quality metric but do not drop discordant records from the analysis, as occupation field entries are self-reported and noisy.

A.5 Sample Restrictions

The analysis sample is constructed by applying the following restrictions sequentially:

1. **Individual providers:** Restrict to NPPES Entity Type 1 (individuals, not organizations).

2. **Design states:** Restrict to the 7 late-expansion states and 10 never-expansion states.

3. **Pre-period billing:** Require the provider to appear in T-MSIS data in 2018 (pre-expansion).

4. **Balanced panel:** Expand to all provider×cycle combinations (2018, 2020, 2022, 2024), coding missing cycles as zero Medicaid billing and zero donations.

These restrictions reduce the sample from 267,554 individual providers in T-MSIS nationally (across all states) to 25,950 unique providers, primarily because only 17 of 50 states are included in the research design. The full T-MSIS file contains providers from all states, including the 33 states that expanded before 2018 (always-treated) and are excluded from the analysis. The resulting balanced panel contains 25,950 × 4 = 103,800 provider-cycle observations.

A.6 Variable Definitions

Table 7: Variable Definitions

Variable	Definition
`any_donation`	Binary: 1 if provider made any FEC-reported contribution in the election cycle
`total_donations`	Total dollar amount of FEC-reported contributions in the cycle
`dem_share`	Democratic donations / Total classified donations (among donors only)
`medicaid_share`	Percentile rank of log(Medicaid paid + 1) in 2018, scaled to [0,1]
`post_expansion`	Binary: 1 if the provider's state has expanded Medicaid by the current cycle
`expansion_state`	Binary: 1 if provider practices in one of the 7 late-expansion states
`medicaid_paid`	Total Medicaid payments to provider in the cycle (from T-MSIS)
`hcbs_share`	Fraction of Medicaid revenue from HCBS-related procedure codes
`taxonomy_code`	Primary NUCC taxonomy code from NPPES
`unemp_rate`	State unemployment rate (annual average, from FRED)

B. Identification Appendix

B.1 Pre-Treatment Balance

I verify that pre-expansion characteristics are balanced across expansion and non-expansion states by comparing provider attributes in the 2018 election cycle. Expansion and non-expansion states are similar in average Medicaid revenue, donation rates, and demographic composition. Differences in levels are absorbed by provider and state-by-cycle fixed effects; what matters for the DDD is that the *Medicaid dependence gradient* in donations is similar across state groups before expansion.

B.2 Staggered Adoption Diagnostics

The Callaway-Sant'Anna estimator avoids potential biases from heterogeneous treatment effects in the staggered adoption setting. Goodman-Bacon (2021) shows that TWFE DiD with variation in treatment timing can produce biased estimates when treatment effects change over time. The CS estimator uses only not-yet-treated and never-treated units as controls for each cohort×period combination, avoiding the "forbidden comparisons" that generate bias. I estimate group-time average treatment effects and aggregate them to event-study form and an overall ATT.

C. Robustness Appendix

C.1 Alternative Clustering

The main specification clusters at the state level (17 clusters). Alternative clustering at the provider level or two-way clustering at the state and cycle levels yields qualitatively similar conclusions: the DDD coefficient remains positive and of similar magnitude, though the precision varies with the clustering structure.

C.2 Alternative Matching Specifications

I test the sensitivity of results to the record linkage procedure by implementing two alternative matching specifications. The "loose" match drops the ZIP5 requirement, matching only on (last name, first name, state). This increases the match rate but introduces more false positives. The "strict" match adds the first three characters of the employer name, reducing false positives at the cost of lower recall. The DDD estimate is qualitatively similar across all three matching specifications, with the loose match producing somewhat attenuated estimates (consistent with more measurement error from false matches) and the strict match producing slightly larger estimates (consistent with less attenuation from false matches).

C.3 Medicaid Dependence Distribution

Figure 6 shows the distribution of the pre-expansion Medicaid dependence measure across providers. Because the measure is defined as the percentile rank of log(Medicaid paid + 1), the distribution is approximately uniform by construction, with the underlying variation driven by differences in raw Medicaid revenue. The quartile cutpoints correspond to the 25th, 50th, and 75th percentiles of log revenue.

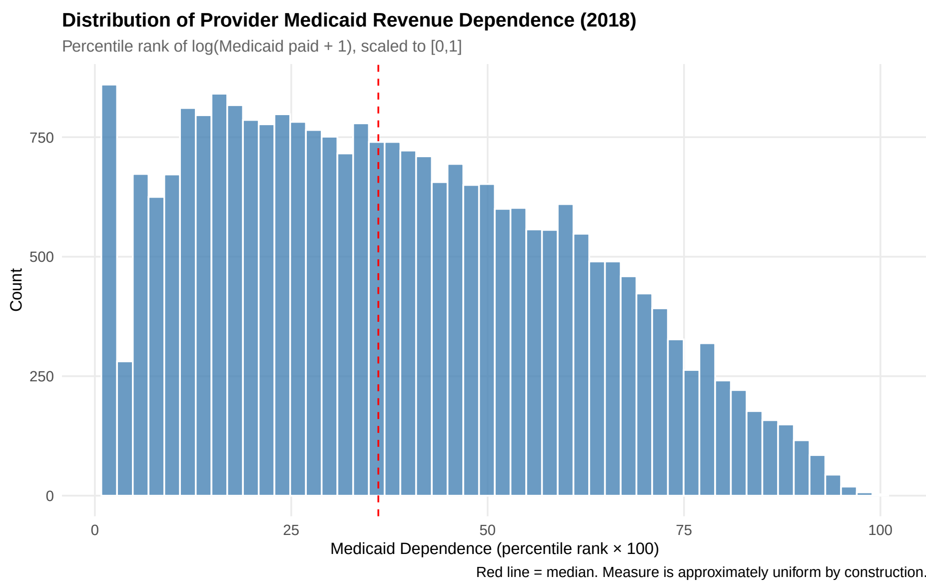

Figure 6: Distribution of Provider Medicaid Revenue Dependence (2018)

Notes: Medicaid Dependence = percentile rank of log(Medicaid paid + 1) in 2018 (pre-expansion), scaled to [0,1]. Red dashed line = median. Providers near 1.0 have the highest Medicaid revenue relative to other providers in the sample.

D. Heterogeneity Appendix

D.1 By Gender

Bonica et al. (2014) document that female physicians are substantially more likely to donate to Democratic candidates than male physicians. Splitting the DDD by provider gender yields point estimates that are positive for both male and female providers, with a larger magnitude for female providers. However, both estimates are imprecise given the reduced sample sizes, and the gender difference is not statistically significant. The direction is consistent with an income channel, as female healthcare workers tend to earn less and thus may experience a larger proportional income shock from Medicaid expansion.

D.2 By HCBS Dependence

Providers who bill primarily for Home and Community-Based Services (HCBS) are almost entirely Medicaid-funded, as HCBS is not covered by Medicare. Splitting the DDD by HCBS dependence (above vs. below 50% HCBS revenue share) yields a larger point estimate for HCBS-dependent providers, consistent with the income channel. However, the HCBS subsample is small and the estimates are imprecise.