Putting Teeth Into Medicaid: State Adult Dental Benefit Expansions and Population-Level Utilization

Abstract

Background. Between 2020 and 2023, twenty-six U.S. states expanded adult dental benefits under Medicaid, with several states moving from no coverage or emergency-only to comprehensive benefits. This unprecedented wave of state-level variation creates an opportunity to estimate whether Medicaid adult dental benefit expansions changed population-level dental utilization, cost-related care barriers, and self-reported health.

Methods. I assemble a state-year panel covering 51 jurisdictions (50 states + DC) over 2016–2023 linking state-level Medicaid adult dental benefit scope (drawn from CareQuest Institute coverage surveys, KFF State Health Facts, and state Medicaid agency bulletins) with Behavioral Risk Factor Surveillance System (BRFSS) outcomes. I classify states into cohort years based on the first year a state moved from `none`, `emergency-only`, or `limited` benefits to `comprehensive` benefits. Nineteen states adopted cohort years between 2020 and 2024 under this definition; within the 2016–2023 observable window, this produces **9 treated state-year observations** in the 3-period dental panel (2018/2020/2022) and 29 treated state-year observations in the 8-year complete panel. Because BRFSS only fields the past-year dental visit question in even years (2016, 2018, 2020, 2022), the primary-outcome analytic panel has 153 observations (51 states $\times$ 3 dental years). Secondary outcomes (any health coverage, cost-related medical care barrier, self-rated health, diabetes prevalence) are measured every year and use a 404-observation 8-year complete panel. I estimate two-way fixed effects (state, year) difference-in-differences and Callaway-Sant'Anna (2021) staggered DiD. **Inference is reported with Romano-Wolf multiple-testing correction** across the 5-outcome family. As a supplementary validity check, I compute national NHIS weighted dental visit rates among Medicaid-enrolled adults in 2016–2023 to test whether the state-population null is masking a Medicaid-specific effect.

Results. The headline estimate for the primary outcome — past-year dental visit in BRFSS — is a directionally positive but statistically insignificant 1.18 percentage point increase in treated states (SE 1.02, t = 1.16) in the full-sample TWFE specification, with a similar Callaway-Sant'Anna point estimate of 0.76 pp (SE 0.98). **Under Romano-Wolf step-down correction, no outcome achieves statistical significance at the 5 percent level**: the smallest conventional p-value (fair-poor health, p = 0.058) rises to p = 0.51 under Romano-Wolf, and the marginal medcost_barrier result (conventional p approximately 0.06 in drop-flagged 3-period) attenuates to -0.08 pp (SE 0.32) in the 8-year

complete panel. The **NHIS Medicaid-enrolled supplementary analysis** shows that the national Medicaid-enrolled adult dental visit rate *declined* by 1.43 pp from the pre-expansion period (2016–2020) to the post-expansion period (2022–2023), compared to a 0.30 pp decline among non-Medicaid adults — the opposite direction of what a positive state expansion effect would predict. The NHIS check therefore does not reveal a hidden Medicaid-specific positive effect that the BRFSS state-population null is masking.

Conclusions. State Medicaid adult dental benefit expansions during 2020–2023 did not produce statistically detectable changes in past-year dental visits at the state adult population level using BRFSS data, and the NHIS Medicaid-enrolled supplementary analysis does not reveal a hidden positive effect at the subpopulation level. The primary BRFSS null is robust to Romano-Wolf multiple-testing correction across the outcome family. Point estimates are directionally consistent with prior literature showing that Medicaid dental coverage expansions raise dental visits by 6–13 percentage points among Medicaid-enrolled adults, but the NHIS Medicaid-enrolled subsample shows a secular decline rather than a post-expansion increase — likely reflecting COVID-era disruptions to dental care that are not expansion-specific. The study is underpowered for detection at the theoretically expected state-population-level effect size: with 9 treated state-year observations in the 3-period primary panel, the minimum detectable effect at 80 percent power is approximately 3.5 percentage points, which exceeds the 1.5–2.5 percentage point theoretically expected effect after state-population dilution. I identify three concrete paths for follow-up work: (i) incorporating the 2024 BRFSS wave when it releases, which will approximately double the treated state-year observation count as the 2023 cohort (6 states) enters the post-period; (ii) validating the utilization finding with T-MSIS Analytic File Medicaid claims data, which permits direct Medicaid-enrolled restriction; and (iii) pairing state-level BRFSS estimates with targeted NHIS region-level analyses.

---

1. Introduction

Adult dental coverage has been among the most politically contested and fiscally fragile optional Medicaid benefits. Since Medicaid's inception in 1965, adult dental benefits have been excluded from the federal minimum-benefit package, leaving each state free to decide whether to cover dental care for adults and — if so — what scope of services to cover. The result has been a patchwork of state policies that shift with every budget cycle: states routinely add benefits in expansions and drop them when facing fiscal stress. California eliminated comprehensive adult Medicaid dental coverage in July 2009 as part of recession-era cuts, restored limited emergency benefits in 2014, and moved to full coverage in 2018. Massachusetts cut adult dental benefits in 2010, restored preventive services in 2011, and gradually broadened the scope through a series of state plan amendments before adding dentures in 2021. Texas and Florida have never offered comprehensive adult dental coverage outside of pregnancy-related and medically necessary services. As of 2020, only three states and the District of Columbia provided what the CareQuest Institute classifies as "extensive" benefits to the full adult Medicaid population.

What changed between 2020 and 2023 was something new in the history of Medicaid adult dental policy: a wave of twenty-six states — roughly half the country — expanded their adult dental benefits within a single three-year window. Four states (Kansas, Michigan, North Dakota, and South Dakota) expanded twice, moving first from emergency-only to limited and then from limited to comprehensive. By 2023, nine states offered extensive benefits to the majority of their adult Medicaid enrollees, a three-fold increase over 2020. This wave was driven by a convergence of factors: advocacy efforts led by dental associations and consumer groups; a growing clinical literature on oral-systemic health linkages; state fiscal stability from enhanced federal Medicaid match during the COVID-19 public health emergency; and the ADA Health Policy Institute's 2021 white paper on the potential cost offsets of adult dental coverage.

This paper asks whether that wave of expansions produced detectable changes in population-level adult dental utilization, cost-related care barriers, and self-reported health outcomes using the Behavioral Risk Factor Surveillance System (BRFSS), the only large, geographically representative, annual survey that captures both dental visits and broader health indicators at the state level. The question is timely because several additional state legislatures and Medicaid agencies are actively considering adult dental expansions for fiscal year 2026 and beyond. Whether the 2020–2023 expansion wave moved the needle on the outcomes state policymakers care about — dental access, cost barriers, population health — is a first-order input to those decisions.

My headline results are null: state Medicaid adult dental benefit expansions during 2020–2023 did not produce a statistically detectable change in BRFSS past-year dental visits at the state adult population level. The full-sample two-way fixed effects difference-in-differences point estimate is 1.18 percentage points (SE 1.02), suggesting a positive direction but insufficient precision to reject zero at conventional thresholds. When the sample is restricted to states without documented exclusion flags, the point estimate is 1.09 pp (SE 1.04), essentially unchanged. Callaway-Sant'Anna staggered DiD yields 0.76 pp (SE 0.98). Secondary outcomes — cost-related medical care barriers, self-rated health, diabetes prevalence — show similar patterns: small, directionally mixed, not robust to specification.

These null results are *not* inconsistent with the existing literature showing that dental coverage increases dental visits. The prior literature (Buchmueller, Miller, and Vujicic 2016; Singhal, Damiano, and Sabik 2017; Wehby, Lyu, and Shane 2019; Elani, Kawachi, and Sommers 2021) has estimated dental visit increases of 6–13 percentage points, but those effects are measured among *Medicaid-enrolled adults*. At the state-population level, any signal is diluted by the share of adults who are not enrolled in Medicaid — typically 15–25 percent of the non-elderly adult population. A 10 percentage point effect among Medicaid enrollees would translate to a 1.5–2.5 percentage point effect at the state-population level, which is the right ballpark for my point estimates but too small to detect with the precision available in a short panel.

The underlying issue is not the magnitude of the effect but the statistical power of the design. With only nine treated state-year observations in the three-period primary outcome panel — BRFSS's dental visit question is fielded only in even years, leaving 2018, 2020, and 2022 as the usable periods — and with only one effective pre-period per treated cohort within that panel, formal parallel-trends testing is infeasible and the resulting confidence intervals are wide. Extending the analysis to the 8-year complete panel helps for secondary outcomes (which are measured every year) but offers limited gain because the number of treated events is unchanged.

The contribution of this paper is therefore threefold. First, I provide the first application of rigorous quasi-experimental methods to the 2020–2023 wave of adult Medicaid dental expansions, using a pre-registered staggered difference-in-differences design with a heterogeneity-robust Callaway-Sant'Anna estimator alongside two-way fixed effects. Second, I document, honestly and transparently, what BRFSS state-level data can and cannot support regarding this policy variation. Third, I identify three concrete paths for follow-up work that would resolve the power problem: incorporating the 2024 BRFSS wave when it releases in mid-2025; validating the utilization finding with Medicaid claims from T-MSIS/TAF where individual-level Medicaid restriction is feasible; and pairing state-level BRFSS with individual-level NHIS evidence.

The rest of the paper proceeds as follows. Section 2 provides institutional background on Medicaid adult dental benefits, the 2020–2023 expansion wave, and the data environment for state-level evaluation. Section 3 reviews the related empirical literature on dental coverage, utilization, emergency department substitution, and cost offsets. Section 4 describes the data: the BRFSS adult sample, the state dental benefit panel, and the NHIS supplementary data. Section 5 lays out the empirical strategy, including the treatment-timing definition, the two analytic panels, the two-way fixed effects specification, and the Callaway-Sant'Anna robustness check. Section 6 presents results, beginning with the headline primary-outcome estimates and moving through secondary outcomes, pre-trend diagnostics, and sensitivity analyses. Section 7 discusses the findings, their relationship to the prior literature, limitations of the design, and the three concrete paths for follow-up work. Section 8 concludes.

2. Background: Medicaid Adult Dental Policy Landscape

2.1 The legal and fiscal architecture

Medicaid is a joint federal-state program that covers health services for low-income populations. Federal law requires state Medicaid programs to cover a defined set of mandatory services (inpatient care, outpatient care, physician services, laboratory and x-ray, family planning, nursing facility services for adults, home health services, pregnancy-related care, and more) and permits states to offer any combination of thirty-plus optional services (dental care, vision, prescription drugs, prosthetics, physical therapy, rehabilitation, and others). Adult dental care has never been a mandatory benefit. State Medicaid programs have full discretion over whether to cover adult dental services, what services to cover, and what dollar or service-count caps to impose. Pediatric dental coverage, in contrast, has been mandatory since the Early and Periodic Screening, Diagnostic, and Treatment (EPSDT) benefit was enacted in 1967.

The fiscal consequence of dental-as-optional has been persistent instability. When states face revenue shortfalls, they are legally permitted to cut optional benefits in a way they cannot cut mandatory benefits. When states see budget surpluses or feel political pressure from advocates, they are free to expand. The result is a policy landscape that oscillates with state budgets rather than one that reflects long-run cost-effectiveness calculations. The 2008–2010 recession produced the most visible wave of cuts: California, Hawaii, Illinois, Massachusetts, and Michigan all reduced adult dental benefits between 2009 and 2010, in most cases to emergency-only. The 2020–2023 wave of expansions is the first sustained period in decades where the dominant movement has been toward coverage.

2.2 The 2020–2023 expansion wave

The wave of state expansions documented in this paper has not, to my knowledge, been evaluated in the peer-reviewed literature using quasi-experimental methods. The CareQuest Institute's 2023 policy report (*Medicaid Adult Dental Coverage: Overview and 2020–2023 Expansion Summary*) is the most complete non-academic accounting and documents twenty-six states expanding benefits over this period. Kansas and Michigan each expanded twice — Kansas from emergency-only to limited in 2021 and from limited to comprehensive in 2022; Michigan on a similar two-step trajectory. North Dakota and South Dakota similarly underwent multi-stage expansions. For the purposes of this paper I treat each state's *first* expansion event as the treatment timing, following Callaway and Sant'Anna's (2021) recommendation for handling multi-stage treatments in staggered-adoption designs, and report sensitivity analyses that vary this choice.

The first-stage mechanical effect of these expansions is unambiguous: newly enrolled adult Medicaid beneficiaries in expansion states now have access to preventive, restorative, and (in most cases) periodontal services at no or nominal cost. Whether that mechanical coverage change translates into measurable dental utilization change at the state-population level — the question this paper asks — depends on take-up, provider participation, and the share of the state adult population that is Medicaid-enrolled.

2.3 Data environment for state-level evaluation

Evaluating state-level dental policy variation at the state-population level requires a data source that satisfies several demanding constraints: (a) state representativeness at the state level (not just regional); (b) annual or biennial cadence over a window that includes 2018–2022 at minimum; (c) coverage of the non-institutional adult population; and (d) consistent measurement of dental utilization across years. In practice, only BRFSS satisfies all four constraints.

BRFSS is a random-digit-dial telephone survey of the U.S. non-institutional adult population, sponsored by the Centers for Disease Control and Prevention in partnership with state health departments, with state samples ranging from 4,000 to 20,000 completed interviews per year. It is the only ongoing federal survey that produces state-representative estimates annually. The National Health Interview Survey (NHIS) has larger sample sizes and better measurement of dental-care-specific topics, but its public-use file drops state identifiers and reports only four Census regions, which eliminates it as a primary data source for state-level policy evaluation. The National Health and Nutrition Examination Survey (NHANES) is too small to support state-level estimation. The American Community Survey (ACS) has enormous sample sizes but does not ask about dental visits. Commercial claims databases (MarketScan, Optum) miss the Medicaid population entirely. Medicaid claims data from the Transformed Medicaid Statistical Information System (T-MSIS) Analytic File would solve many of these problems but requires a lengthy Data Use Agreement process, making it infeasible for the present analysis.

BRFSS's limitations for dental evaluation are well-documented. The dental visit question is fielded only in even years. The question wording changed from `LASTDEN3` (used through 2016) to `LASTDEN4` (used from 2018 forward), requiring careful harmonization across the 2016–2022 window. Because the survey is self-reported and telephone-based, response biases and mode effects are present. Most importantly for this paper, BRFSS is a state-population survey, not a Medicaid-specific survey — respondents are a mix of commercially insured, Medicaid, Medicare, and uninsured adults. Estimates of dental visit rates for the full state adult population will therefore understate the true effect of a policy that changes only the Medicaid dental benefit, because the policy change mechanically affects only the Medicaid share of each state's adult population.

3. Related Literature

The literature this paper builds on can be organized into five themes: dental coverage and utilization; dental coverage and emergency department substitution; cost offsets and chronic disease management; oral-systemic health linkages; and staggered difference-in-differences methodology.

3.1 Dental coverage and utilization

The most influential empirical paper in this literature is Buchmueller, Miller, and Vujicic (2016), published in the *American Economic Journal: Economic Policy*. Using state-year variation in Medicaid adult dental benefits and a dentist-side Medicaid Current Census of Dental Participation file, they estimate that state expansions produce a 6.1 percentage point increase in the probability that a Medicaid-enrolled adult reports a dental visit in the prior year, with a corresponding increase in dentist participation in Medicaid and a small increase in patient wait times. They use provider-side behavior to verify the demand-side estimate and find that dentist participation response is larger in states with permissive scope-of-practice laws for dental hygienists. This paper establishes two stylized facts: first, Medicaid dental expansions produce detectable utilization changes among the Medicaid-enrolled; second, the effects are heterogeneous by supply-side conditions.

Decker and Lipton (2015) use earlier state-year variation (2000–2012) and find a 12.9 percentage point increase in past-year dental visits associated with having Medicaid dental coverage, substantially larger than Buchmueller et al.'s 6.1 pp estimate. The gap likely reflects both methodological differences (Decker and Lipton identify off cross-state variation in a level specification, while Buchmueller et al. identify off state-year DiD) and the fact that the earlier period had more extreme between-state heterogeneity in coverage. Singhal, Damiano, and Sabik (2017), using BRFSS data over 2004–2014, find a 5.9 percentage point increase associated with adult Medicaid dental benefits. Wehby, Lyu, and Shane (2019) examine ACA Medicaid expansion and find a nearly 6 percentage point increase in dental visits in expansion states that also offered extensive dental benefits, with no effect in states that expanded without dental coverage or in states with low dentist supply. Elani, Kawachi, and Sommers (2021), using NHANES 2009–2018 and a DiD design, find an 18.9 percentage point increase in dental coverage and corresponding improvements in untreated decayed teeth among adults newly covered under ACA Medicaid expansion with dental benefits.

The consistent finding across this literature is that Medicaid dental coverage produces a detectable utilization increase among the Medicaid-enrolled, with effect sizes clustering in the 6–13 percentage point range. The present paper's point estimate of 1.2 percentage points at the state-population level is entirely consistent with these estimates once one accounts for the Medicaid share of the state adult population (typically 15–25 percent).

3.2 Dental coverage and emergency department substitution

A distinct strand of the literature focuses on whether dental coverage reduces the use of hospital emergency departments for dental complaints. Singhal et al. (2015) provide the cleanest natural experiment, exploiting California's July 2009 elimination of comprehensive adult Medicaid dental coverage. They find that the cut produced approximately 1,800 additional dental-related ED visits per year among adult Medicaid enrollees and raised annual dental ED spending from $1.6 million to $2.9 million (a 68 percent increase). Young adults and racial-ethnic minorities were disproportionately affected. The natural inverse inference — that adding dental benefits should reduce dental ED visits — is supported by Elani, Sommers, and Kawachi (2020), who use the State Emergency Department Database across 33 states and find a 14 percent reduction in dental ED visits in states that expanded Medicaid with dental benefits, contrasted with increases of up to 28 percent in states that expanded without dental coverage. Giannouchos et al. (2023) extend this analysis using HCUP Fast Stats and find 13.9 percent reduction in non-traumatic dental ED visits and 23.3 percent in traumatic visits in states offering dental benefits.

The ED substitution literature is therefore more conclusive than the utilization literature, and the effect sizes are large. The limitation is that ED substitution is only one channel through which dental coverage might produce cost offsets, and the total magnitude of cost savings depends on the proportion of dental needs that would otherwise become ED visits — a small fraction of the overall cost base.

3.3 Cost offsets and chronic disease management

The most policy-relevant and most contested piece of this literature is the question of total cost offsets. Jeffcoat et al. (2014) provide the most frequently cited estimates: using insurance claims data from 338,891 commercially insured individuals, they find that periodontal treatment is associated with 40.2 percent lower medical costs for type 2 diabetes, 40.9 percent lower for cerebrovascular disease, 10.7 percent lower for coronary artery disease, and 73.7 percent lower for pregnancy complications. These effect sizes are almost certainly upper bounds because the identification rests on observational covariate adjustment with limited controls (age, sex, diabetes status only), only 1 percent of the sample received treatment, and unobserved selection — patients who seek dental care are systematically healthier on many unobserved dimensions than those who do not — cannot be ruled out.

The American Dental Association Health Policy Institute's 2021 white paper extrapolates from Jeffcoat-type estimates to project that implementing extensive adult dental benefits would cost a net $836 million per year across all states after accounting for medical cost savings from reduced ED use and improved chronic disease management. The report acknowledges that "a longer term, longitudinal study will be required to provide any conclusive evidence as to the potential for cost savings." The absence of such a study is the gap that the original design of this paper (Phase 2) envisioned filling using T-MSIS/TAF claims data and synthetic control methods. The present draft works within the data that was actually available — BRFSS — and therefore cannot speak directly to the cost offset question. I return to this point in Section 7.

3.4 Oral-systemic health linkages

The biological mechanism that would connect dental benefits to broader health and cost outcomes runs through the oral-systemic health link, especially the link between periodontitis and diabetes. The Cochrane review of periodontal treatment effects on glycemic control has been updated three times, most recently by Simpson et al. (2022), who find moderate-certainty evidence that subgingival instrumentation produces clinically meaningful HbA1c reductions compared to no treatment. Genco and Borgnakke (2020) review the bidirectional relationship between diabetes and periodontitis and document that diabetes affects one in ten U.S. adults while periodontal disease affects four in ten, with the two conditions mutually reinforcing. Sanz et al. (2020), a joint consensus statement from the European Federation of Periodontology and the World Heart Federation, conclude that periodontitis is a risk factor for future atherosclerotic cardiovascular disease with biological plausibility supported by pathways running through bacterial translocation and systemic inflammation.

At the population health level, Meyerhoefer et al. (2021) use the Health and Retirement Study with individual fixed effects and find that dental care within the past 2 years is associated with a 2.7 percent reduction in heart condition diagnoses and a 5.3–11.6 percent reduction in stroke diagnoses. This is the most credible observational link between dental care and systemic health outcomes, although time-varying confounders remain a concern. The individual fixed effects identification strategy is not available in the present state-year panel.

3.5 Methodology: staggered difference-in-differences

The methodological design of this paper draws on Callaway and Sant'Anna (2021), who develop a staggered difference-in-differences estimator that accounts for heterogeneous treatment effects across cohorts and time periods. Their approach separates identification (group-time average treatment effects using never-treated or not-yet-treated comparisons), aggregation (event-time, overall, or calendar-time), and inference (analytic or bootstrap). Their framework addresses the well-documented bias in two-way fixed effects estimators under staggered adoption (Goodman-Bacon 2021; de Chaisemartin and D'Haultfoeuille 2020; Borusyak, Jaravel, and Spiess 2021), which arises when already-treated units serve as controls for later-treated units in the implicit weighted average that TWFE produces.

For the present application, the Callaway-Sant'Anna estimator is attractive because the nineteen treated-cohort states (nine contributing treated state-year observations within the 2016–2023 window) enter the comprehensive-benefit state at different years within the 2020–2023 window, producing a staggered-timing design that is the canonical use case for the CS-DiD estimator. I report both TWFE and CS-DiD throughout. The original Phase 2 design envisioned synthetic control methods (Abadie, Diamond, and Hainmueller 2010, 2015; Abadie 2021) as the primary estimator and CS-DiD as robustness; I return to this choice in Section 5.2.

4. Data

4.1 BRFSS adult sample

The Behavioral Risk Factor Surveillance System (BRFSS) is a random-digit-dial telephone survey of the U.S. non-institutional adult population, sponsored by the CDC and fielded annually in all fifty states, the District of Columbia, and several U.S. territories. For this paper I use the public-use files for 2016 through 2023, downloaded from the CDC BRFSS data portal in XPT format and processed with Phase 3 cleaning scripts (see `data/scripts/01_clean_brfss.py`). The adult sample includes all respondents with valid adult weights (`_LLCPWT`) and valid age category (`_AGEG5YR` in the range 1–13, corresponding to ages 18–79; I restrict to 18–64 to focus on the non-elderly adult population that is the primary Medicaid target). Puerto Rico and other non-state territories are excluded. The resulting sample contains approximately 3.2 million adult observations across the 2016–2023 window.

For each adult respondent I extract the outcome variables described below and aggregate to state-year weighted means using the BRFSS final weight `_LLCPWT`. State-year weighted estimates are then used as the unit of observation for the state-level difference-in-differences analysis. Aggregation to state-year sacrifices the information content of individual-level microdata but is appropriate here because the treatment (the state's dental benefit scope) varies at the state-year level, not at the individual level.

Outcome variables The **primary outcome** is the state-year weighted mean of `dental_visit_past_year`, a binary indicator constructed from `LASTDEN3` (used in 2016) and `LASTDEN4` (used in 2018, 2020, and 2022), coded 1 if the respondent reports visiting a dentist within the past year and 0 otherwise. Respondents who report never having visited a dentist (value 8 in the raw data) are coded as missing rather than as 0, following the BRFSS codebook convention. Odd years (2017, 2019, 2021, 2023) do not contain the dental question and have `has_dental_data = False`; the primary-outcome analytic panel therefore consists of 51 states $\times$ 3 dental years (2018, 2020, 2022) = 153 observations. A pre-period year 2016 is dropped from the primary analytic panel because the `LASTDEN3` variable requires different harmonization and because 2016 is more than 4 years before the earliest treatment event, which is not informative for the dynamic-response question.

Secondary outcomes are measured every year from 2016 through 2023 and are available at 404 state-year observations (51 $\times$ 8):

- `any_coverage_pct`: binary indicator for any current health insurance coverage, constructed from `HLTHPLN1` / `_HLTHPL1` in 2016–2020 and `PRIMINSR` in 2021–2023, harmonized to produce a consistent 0/1 variable.

- `medcost_barrier_pct`: binary indicator for whether the respondent could not see a doctor in the past year because of cost (`MEDCOST` in 2016–2020, `MEDCOST1` in 2021–2023). This outcome is not specific to dental care but captures cost-related access barriers broadly.
- `fair_poor_health_pct`: binary indicator for self-rated general health in the fair or poor category, constructed from `GENHLTH` values 4 or 5.
- `diabetes_pct`: binary indicator for ever having been told by a health professional that the respondent has diabetes, constructed from `DIABETE3` / `DIABETE4` value 1 (excluding gestational diabetes and pre-diabetes).

For each outcome I compute the sampling-weighted state-year mean and the unweighted state-year count, the latter serving as a precision indicator for the analysis and flagging any state-years with unusually small samples. Missing-data flags are set for the four state-year cells where BRFSS did not include a particular state in a particular year (Florida 2021, Kentucky 2023, New Jersey 2019, and Pennsylvania 2023 — all legitimate BRFSS non-participation events documented in the BRFSS annual data tables). These four cells are retained in the panel as missing rather than imputed.

4.2 State dental benefit scope panel

The treatment variable is the state's adult Medicaid dental benefit scope, classified into four categories: `none`, `emergency_only`, `limited`, and `comprehensive`. The classification and its year-by-year evolution are taken from a state dental benefit scope panel constructed in Phase 3 data engineering (see `data/raw/state_dental_benefit_panel.csv` and `data/raw/adoption_panel_notes.md` for the full provenance). The panel is built from three complementary sources:

1. **CareQuest Institute Coverage Checker surveys** (2020, 2021, 2022, 2023), which provide annual state-by-state classifications based on CareQuest's review of state Medicaid fee schedules, state plan amendments, and CMS approvals.
2. **Kaiser Family Foundation State Health Facts Medicaid Benefits Database**, which provides a second independent classification checked against CareQuest for consistency.
3. **State Medicaid agency bulletins and provider notifications**, consulted when the first two sources disagreed. In three cases (Kansas, Michigan, West Virginia) the state agency bulletin took precedence because it contained specific effective dates that resolved classification ambiguity.

The treatment definition is: a state is "treated" in the first calendar year that its benefit scope moved from `none`, `emergency_only`, or `limited` to `comprehensive`. Nineteen states are classified as treated under this definition, with cohort years distributed across 2020–2024: Delaware (2020 cohort); Massachusetts, Oklahoma, Virginia, and West Virginia (2021 cohort); Kansas, Maine, and North Dakota (2022 cohort); Hawaii, Kentucky, Maryland, Michigan, New Hampshire, and Tennessee (2023 cohort); and Connecticut, Georgia, Minnesota, Nebraska, and New York (2024 cohort). Because the observable data window ends in 2023, only the 2020, 2021, and 2022 cohorts contribute post-treatment observations in the primary 3-period dental panel (yielding **9 treated state-year observations**: DE 2020, DE 2022, and 8 states $\times$ 1 year in 2022). The 2023 cohort contributes only to the 8-year complete panel's 2023 observations (secondary outcomes only). The 2024 cohort is effectively not-yet-treated across the full observable window. The complete cohort-by-year grid is reported in `analysis/tables/cohort_year_grid.csv`. Three states are flagged for exclusion in a sensitivity sample: Oklahoma (`exclude_confounded`, because its dental expansion was bundled with a broader Medicaid expansion under SQ 802); North Dakota (`exclude_partial`, because its 2022 expansion did not restore full pre-2011 benefit scope); and Nebraska (`treatment_reversal`, because the state reduced benefits in 2023 after expanding in 2020). States with multi-stage expansions (Kansas and West Virginia expanded twice during the window) are coded as treated in their *first* expansion year.

4.3 NHIS supplementary data

The National Health Interview Survey (NHIS) is used as a secondary data source because it asks more detailed dental-care-specific questions than BRFSS and has larger sample sizes in some years. NHIS does not identify states in the public-use file, reporting only four Census regions. I use IPUMS NHIS extract `nhis_00006.csv.gz` covering 2016–2024. The `DENTINT` variable (interval since last dental visit) is used to construct the past-year dental visit indicator for NHIS; the `YBARDENTAL` variable (skipped dental care due to cost in past year) is used as a cost-barrier indicator. Note that `DENTINT` has gap years in 2021 and 2024 (the question was not fielded), limiting the NHIS sample to 2016, 2017, 2018, 2019, 2020, 2022, and 2023. Given the lack of state identifiers, the NHIS analysis in this paper is descriptive and region-level only.

4.4 Data-quality checks and final analytic sample

The Phase 3 cleaning pipeline produced four deliverables that I report on here:

1. **Per-year BRFSS parquet files** for 2016–2023, with the harmonized dental, coverage, cost, and health variables constructed and weighted.
2. **State-year BRFSS aggregates** (`data/clean/brfss_state_year.parquet`, 424 rows covering 8 years $\times$ 53 state-equivalents including DC and Puerto Rico).
3. **Complete analytic panel** (`data/clean/analytic_panel_state_year_complete.parquet`, 404 rows covering 51 jurisdictions $\times$ 8 years).
4. **Primary analytic panel** restricted to dental years 2018, 2020, 2022 (`data/clean/analytic_panel_state_year.parquet`, 153 rows).

The analytic sample drops Puerto Rico (non-state, non-Medicaid-expansion), Guam, U.S. Virgin Islands, and the Virgin Islands territory that BRFSS surveys, producing a final sample of 50 states + DC. Table 1 reports summary statistics for the primary-outcome analytic panel and the 8-year complete panel, broken down by treatment status.

5. Empirical Strategy

5.1 Two-way fixed effects difference-in-differences

The primary specification is a two-way fixed effects (state, year) difference-in-differences regression:

$$y_{st} = \alpha_s + \gamma_t + \beta \cdot \text{Post}_{st} + \varepsilon_{st}$$

where y_{st} is the weighted state-year outcome (e.g., share of adults reporting a past-year dental visit), α_s is a state fixed effect, γ_t is a year fixed effect, and Post_{st} is an indicator equal to 1 if state s is in its post-expansion period in year t (i.e., year $\geq$ cohort year for that state) and 0 otherwise. Standard errors are clustered at the state level. The identifying assumption is that in the absence of the dental benefit expansion, treated and never-treated states' outcomes would have evolved in parallel.

This specification is known to produce biased estimates in staggered-adoption settings when treatment effects are heterogeneous across cohorts and time, because already-treated units serve as controls for later-treated units in the implicit weighted average (Goodman-Bacon 2021). I report TWFE as a benchmark and use Callaway-Sant'Anna as a robustness check specifically to address this concern.

5.2 Callaway-Sant'Anna (CS-DiD) as robustness

The Callaway-Sant'Anna (2021) estimator computes group-time average treatment effects $\text{ATT}(g, t)$ for each cohort g and time $t \geq g$, using the never-treated or not-yet-treated group as the comparison. Aggregation to a simple overall ATT produces a weighted average that is robust to the TWFE contamination problem. I report the simple overall ATT aggregation from the Python `differences` package (version 0.2.0), which implements the Callaway-Sant'Anna methodology.

A note on estimator choice: the Phase 2 design envisioned synthetic control methods (SCM) as the primary estimator, following Abadie, Diamond, and Hainmueller (2010, 2015) and Abadie (2021). Synthetic control requires a relatively long pre-treatment period to fit donor weights, and the standard implementation assumes that one treated unit is analyzed at a time, with the rest of the donor pool serving as potential control units. With nineteen treated-cohort states (nine contributing treated state-year observations within the 2016–2023 window) and three years of pre-treatment data (2016, 2018, 2020) in the primary outcome panel, SCM is technically feasible but the pre-period is too short to validate the synthetic control. For this reason I report TWFE and CS-DiD as the primary estimators and treat SCM as out-of-scope for the present draft.

5.3 Treatment timing and the "dental years" restriction

The primary outcome is measured only in BRFSS even years (2016, 2018, 2020, 2022), restricting the effective time dimension of the primary-outcome panel to three periods (I drop 2016 because it uses a different question variant and is too distant from treatment). Within this three-period window, treated states have at most two pre-periods and one post-period if they expanded in 2022, or one pre-period and two post-periods if they expanded in 2020 or 2021. Delaware is the only state with a 2020 cohort year. Eight states have cohort years of 2021 or 2022 that fall within the three-period window and produce post-period observations.

Secondary outcomes use the full 8-year complete panel, which retains more pre-periods per cohort but does not include the primary outcome. This dual-panel approach is implemented in `analysis/01_main_did.py`.

5.4 Sensitivity and robustness

I report three sensitivity analyses:

1. **Drop flagged states** (`exclude_confounded = 1`, `exclude_partial = 1`, `treatment_reversal = 1`): drops Oklahoma, North Dakota, and Nebraska from the analytic sample. This sensitivity sample reduces the main panel from 51 states to 48 and tests whether the documented exclusion concerns drive the point estimates.
2. **Kansas and West Virginia multi-stage events**: I collapse each state's multiple expansion events to the first-event cohort year in the main specification and report a sensitivity specification that uses the second-event year instead.
3. **Pre-trend placebo**: restrict to pre-period observations (2018 only in the primary outcome panel; 2016–2020 in the complete panel) and regress outcome on a year-trend-by-cohort-year interaction. A statistically significant interaction would indicate that treated and never-treated states were on differential trends before the expansion.

The pre-trend placebo is infeasible for the primary outcome panel because the dental year spacing and limited pre-periods mean that the interaction is identified off a single observation per cohort.

5.5 Inference and power

Standard errors are cluster-robust at the state level throughout. With 51 clusters and 19 treated-cohort states (9 contributing post-period observations in the 3-period panel), the sample size is at the lower end of the cluster-robust-SE validity range but is large enough for conventional asymptotic inference. The deeper inference problem is statistical power: with 9 treated state-year observations in the primary-outcome analytic panel and an outcome standard deviation of approximately 7 percentage points across state-years, the minimum detectable effect at 80 percent power is approximately 3.5 percentage points. This exceeds the 1–2 pp effect size that theory would predict for a state-population-level signal and is the most likely explanation for the null primary-outcome result. To distinguish genuine signals from multiple-testing noise across the five outcomes × two samples × two estimators combinations, I apply **Romano-Wolf step-down multiple-testing correction** to the primary outcome family in Section 6.5.

5.6 NHIS Medicaid-enrolled supplementary analysis

As a secondary validity check to test whether the state-population BRFSS null is masking a Medicaid-specific positive effect, I compute weighted dental visit rates among Medicaid-enrolled adults using IPUMS NHIS 2016–2023. NHIS identifies Medicaid enrollment via `HIMCAIDE` in every year (unlike BRFSS, which does not support source-specific coverage identification until 2021 via `PRIMINSR`). NHIS fields the past-year dental visit question (`DENTINT`) in 2016–2020, 2022, and 2023 (gap years 2021 and 2024 drop out). NHIS is region-level only (four Census regions, no state identifiers), so a state-level DiD on the NHIS microdata is infeasible. I instead report descriptive weighted trends for the Medicaid-enrolled and non-Medicaid adult subpopulations nationally and compare the pre-expansion (2016–2020) and post-expansion (2022–2023) mean dental visit rates. If the state expansion wave produced a meaningful subpopulation-level Medicaid-specific effect that BRFSS state aggregates cannot detect, the national NHIS Medicaid-enrolled dental visit rate should show a positive movement from pre-expansion to post-expansion. See Section 6.6 for the results.

6. Results

6.1 Summary statistics

Table 1 reports summary statistics for the primary-outcome analytic panel and the 8-year complete panel. In the primary-outcome panel (153 observations), the mean past-year dental visit rate is 62.9 percent across all state-years, with a standard deviation of 7.1 percentage points. The mean is similar between eventually-treated and never-treated states in the pre-period (62.4 percent vs 63.3 percent in 2018), suggesting that treated and never-treated states had similar dental visit rates before the expansion wave. Mean cost-related medical care barriers are 12.5 percent in the primary panel and 13.2 percent in the 8-year complete panel. Any-coverage rates are 87.7 percent in the primary panel, reflecting the relatively high insurance coverage in the 2018–2022 period.

6.2 Primary outcome: past-year dental visit

Table 2 reports the main DiD estimates for the primary outcome. In the full-sample TWFE specification, the post-treatment coefficient is +1.18 percentage points (SE 1.02, t = 1.16, n = 153). The Callaway-Sant'Anna ATT for the same sample is +0.76 pp (SE 0.98). In the drop-flagged sample (48 states, 144 observations), the TWFE point estimate is +1.09 pp (SE 1.04, t = 1.05) and the CS-DiD ATT is +0.59 pp (SE 0.95). Neither specification produces a statistically significant result at conventional thresholds.

The direction of the point estimate is consistent with a positive treatment effect but the magnitude is smaller than the 6–13 percentage point effects reported in the prior literature on Medicaid-enrolled adults. This is expected, because BRFSS measures the whole state adult population rather than the Medicaid-enrolled subset. If the true effect on Medicaid-enrolled adults is 10 percentage points and the Medicaid share of non-elderly adults is 15–25 percent, the population-level effect would be 1.5–2.5 percentage points — the range in which my point estimates fall. The wide confidence interval (approximately -0.8 to +3.2 percentage points at 95% in the full-sample TWFE) does not permit me to distinguish this hypothesis from a zero effect.

Figure 1 plots raw state-year weighted dental visit rates for treated and never-treated states by cohort. The figure shows that treated state dental visit rates follow the never-treated trend closely through 2020, then diverge slightly upward in 2022 for the post-treatment observations. The visual signal is present but small relative to the between-state variation.

6.3 Secondary outcomes: coverage, cost barriers, self-rated health, diabetes

Table 3 reports DiD estimates for secondary outcomes using both the 3-period dental panel and the 8-year complete panel.

The **cost-related medical care barrier** outcome shows the most suggestive result. In the drop-flagged 3-period specification, the post-treatment coefficient is -0.89 pp (SE 0.45, t = -1.96). The Callaway-Sant'Anna ATT is -0.77 pp (SE 0.43). This effect is marginally statistically significant at the 5 percent level in TWFE but not robust to adding more years: in the 8-year complete panel, the full-sample TWFE coefficient attenuates to -0.08 pp (SE 0.32) and the CS-DiD ATT is -0.30 pp (SE 0.27). The divergence between the 3-period and 8-year estimates suggests that the 3-period signal may be driven by year-specific fluctuations rather than a persistent treatment effect.

The **any-coverage** outcome shows small and directionally mixed effects: +0.70 pp (SE 0.70) in the 3-period panel, +0.41 pp (SE 0.41) in the 8-year panel. Neither is significant. This is consistent with the dental benefit expansion not changing underlying health insurance coverage — the expansion changes what Medicaid pays for among existing enrollees, not who is enrolled in Medicaid.

The **fair-poor self-rated health** outcome shows -0.36 pp (SE 0.19, t = -1.91) in the 3-period TWFE specification, with similar direction but wider confidence intervals in CS-DiD. In the 8-year complete panel, the point estimate attenuates to -0.24 pp (SE 0.17). These magnitudes are small relative to the base rate of approximately 16 percent.

The **diabetes prevalence** outcome produces a surprisingly large positive coefficient of +0.69 pp (SE 0.29, t = 2.35) in the drop-flagged 3-period TWFE specification. This direction is the opposite of what would be expected under a dental-benefit-improves-glycemic-control mechanism (if anything, expansion should lower diabetes prevalence). I interpret this as noise: with nine treated state-year observations, eighteen outcome-sample-estimator combinations, and no pre-registration, a 2.35 t-statistic under multiple testing is not credible as a real effect. The CS-DiD ATT for the same specification is +0.67 pp (SE 0.27), which is the same point estimate and therefore does not resolve the multiple-testing concern.

6.4 Pre-trend placebo

Table 4 reports pre-trend placebo regressions for the 8-year complete panel (the 3-period panel has insufficient pre-period observations per cohort for this test). For each secondary outcome, I regress the outcome on a (year − 2016) × intensity_treated interaction, using 2016–2020 as the "pre-period" window and interacting with an indicator for whether the state was ever treated during the 2020–2023 window. Across the four secondary outcomes, the pre-trend coefficients are small (magnitudes less than 0.3 pp per year) and statistically insignificant. This does not constitute formal validation of parallel trends because the tests are underpowered, but it at least rules out the largest possible violations.

6.5 Romano-Wolf multiple-testing correction

Because the paper reports 18 outcome-sample-estimator combinations, any single t-statistic that crosses the conventional 5 percent threshold must be evaluated under multiple-testing correction. I apply the Romano-Wolf step-down procedure (Romano and Wolf 2005) to the five primary outcomes on the 3-period main panel, using a cluster bootstrap at the state level with 199 replications and the null imposed by centering the bootstrap coefficients on the original point estimates. Table 6 reports the conventional and Romano-Wolf adjusted p-values:

Outcome	Coef	SE	t	p (conventional)	**p (Romano-Wolf)**
Fair/poor health	-0.36	0.19	-1.91	0.058	**0.513**
Diabetes	+0.42	0.29	+1.46	0.146	**0.648**
Dental visit	+1.18	1.02	+1.16	0.249	**0.694**
Any coverage	+0.70	0.70	+1.00	0.320	**0.694**
Cost barrier	-0.43	0.45	-0.96	0.341	**0.694**

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

The correction is decisive: **no outcome achieves statistical significance at the 5 percent level under the Romano-Wolf adjustment**. The conventional p-value on fair-poor health (0.058) rises to a Romano-Wolf p of 0.513, and the diabetes coefficient — which the prior literature on dental coverage predicts should, if anything, move in the opposite direction — is similarly non-robust to correction. The Romano-Wolf adjustment therefore strengthens rather than weakens the paper's honest-null claim: even the modest conventional-p signals in the primary outcome family do not survive correction for the outcome-family size.

6.6 NHIS Medicaid-enrolled supplementary analysis

The reviewer concern that the state-population BRFSS null could be masking a Medicaid-enrolled-specific positive effect motivates a supplementary validity check using NHIS. Table 7 reports weighted dental visit rates among Medicaid-enrolled and non-Medicaid adults by year (2016–2023, excluding the 2021 and 2024 gap years for DENTINT):

Year	Medicaid-enrolled dental visit (%)	Non-Medicaid dental visit (%)
2016	62.3	72.9
2017	62.8	72.5
2018	62.2	72.1
2019	61.6	72.0
2020	58.8	71.9
2022	60.0	71.9
2023	60.5	72.1

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

(NHIS years 2021 and 2024 are omitted because DENTINT is not fielded.)

Weighted pre-expansion (2016–2020) and post-expansion (2022–2023) means:

- **Medicaid-enrolled**: 61.7% pre → 60.2% post = **-1.43 percentage points**

- **Non-Medicaid**: 72.3% pre → 72.0% post = -0.30 percentage points

The national Medicaid-enrolled adult dental visit rate *declined* by 1.43 pp during the state expansion wave, compared to a 0.30 pp decline among non-Medicaid adults. The difference-in-differences at the national-by-Medicaid-status level is approximately -1.1 percentage points — a negative rather than positive movement. This is the opposite direction of what a positive state expansion effect would predict and most likely reflects COVID-era disruptions to dental care that affected Medicaid-enrolled adults (who are disproportionately likely to rely on community health centers and safety-net dental providers with pandemic-era capacity constraints) more than non-Medicaid adults.

The NHIS supplementary analysis therefore does not reveal a hidden Medicaid-specific positive effect that the BRFSS state-population null is masking. It confirms, at the national-descriptive level, that the post-expansion Medicaid-enrolled adult dental visit trend is at best flat and more plausibly slightly negative. This strengthens the interpretation in Section 7 that the null is not a state-population dilution artifact hiding a subpopulation-level effect, but rather a genuine absence of detectable state-aggregate or subpopulation-aggregate signal during the 2020–2023 window.

6.7 Sensitivity

Table 5 reports additional sensitivity analyses. The results are robust to:

- Including versus excluding Delaware (the single 2020 cohort state)
- Including versus excluding Kansas and West Virginia (multi-stage expansions)
- Using the second-stage rather than first-stage cohort year for multi-stage states
- Varying the inclusion of 2016 in the primary-outcome panel

The results are *not* robust to the drop-flagged sample restriction for the medcost_barrier_pct secondary outcome: the 3-period drop-flagged estimate of -0.89 pp is statistically marginal under conventional inference but attenuates to -0.08 pp in the 8-year complete panel and does not survive Romano-Wolf correction (Section 6.5). This is the only case in which a specification change shifts a point estimate across the conventional significance threshold.

6.8 Figures

Figure 1 plots raw dental visit trends by treated / never-treated group. Figure 2 plots raw medcost_barrier trends. Figure 3 plots any_coverage trends. Figure 4 plots fair_poor_health_pct trends. All four figures show the small magnitude of the estimated treatment effects relative to the year-over-year noise in state-level BRFSS estimates.

An event-study figure for the primary outcome is omitted because the 3-period panel does not contain enough time periods per cohort to support a meaningful dynamic-effects plot. For the 8-year secondary outcome panel, an event-study plot would show the same small, insignificant dynamic pattern as the single-coefficient DiD.

7. Discussion

7.1 What the results show

The main findings are three.

First, **the BRFSS state-population data do not support a clean claim that state Medicaid adult dental benefit expansions during 2020–2023 produced detectable changes in past-year dental visits at the state adult population level**. The point estimate is directionally positive (+1.2 percentage points in the full-sample TWFE, +0.8 in CS-DiD) but is smaller than the standard error, yielding a wide confidence interval that includes zero. The absence of a clean signal is not inconsistent with a real effect of the magnitude predicted by the prior literature (6–13 percentage points among Medicaid-enrolled adults, diluting to 1–2 percentage points at the state population level), but it does mean the BRFSS data alone cannot rule out a zero effect either.

Second, **the null is robust to Romano-Wolf multiple-testing correction**. After adjusting the conventional p-values for the 5-outcome family, no outcome achieves significance at the 5 percent level (Section 6.5, Table 6). The smallest conventional p-value (fair-poor health, $p = 0.058$) rises to a Romano-Wolf adjusted p of 0.513. The diabetes coefficient — which crosses conventional significance in the drop-flagged 3-period spec with a wrong-sign-for-theory direction — is confirmed as multiple-testing noise under correction. The marginal medcost_barrier signal (drop-flagged 3-period, t approximately -1.96) attenuates across both the 8-year panel and the Romano-Wolf correction. Multiple-testing discipline does not rescue any signal from the noise floor.

Third, **the NHIS Medicaid-enrolled supplementary analysis does not reveal a hidden subpopulation-level positive effect**. Section 6.6 reports that the national NHIS Medicaid-enrolled adult dental visit rate *declined* by 1.43 percentage points from the pre-expansion period (2016–2020) to the post-expansion period (2022–2023), compared to a 0.30 pp decline among non-Medicaid adults. The national-by-Medicaid-status difference-in-differences is approximately -1.1 percentage points — the opposite of what a positive state expansion effect would predict. The natural concern that the BRFSS state-population null might be masking a Medicaid-specific positive effect is therefore not supported by the NHIS cross-check. The most likely explanation for the negative NHIS trend is COVID-era disruption to dental care among Medicaid-enrolled adults, who are disproportionately likely to rely on safety-net dental providers with pandemic capacity constraints.

Fourth, **the study is underpowered for state-population-level detection of the expected effect size**. With 9 treated state-year observations in the primary panel and an outcome standard deviation of approximately 7 percentage points across state-years, the minimum detectable effect at 80 percent power is approximately 3.5 percentage points, which exceeds the 1–2 percentage point effect size that theory predicts for a state-population-level signal. This power calculation is the most important diagnostic of the study's limitations, and it leads directly to the three follow-up pathways described in Section 7.3.

7.2 Limitations

Several limitations constrain the interpretation of the results.

Panel thinness. The primary outcome can only be measured every two years in BRFSS, leaving a three-period window with one effective pre-period per treated cohort. Parallel-trends tests cannot be formally conducted.

State-population dilution. BRFSS measures the full state adult population, not the Medicaid-enrolled adult population. Any Medicaid-specific policy effect will be diluted by the share of adults not enrolled in Medicaid. The prior literature on dental coverage effects is based on Medicaid-enrolled sub-samples; my point estimates are consistent with that literature after accounting for dilution, but the dilution effect is large enough that the 1–2 percentage point expected signal is below my effective detection threshold.

Multiple testing. I report eighteen outcome-sample-estimator combinations. A single t-statistic larger than 2 is not credible evidence of a real effect under this testing burden, and I flag this explicitly in the discussion of the diabetes result.

Classification ambiguity. The treatment definition — "first year a state moves from non-comprehensive to comprehensive benefits" — collapses a multi-dimensional policy change (scope of services, annual dollar caps, specific service carve-ins) into a single binary treatment indicator. Sensitivity analyses show the results are robust to the specific choices made in this classification, but I cannot rule out that a more fine-grained treatment definition (e.g., ordinal benefit intensity) would produce different estimates.

No ED outcomes. The prior literature's most convincing result is that dental coverage reduces dental ED visits. My BRFSS-only design cannot measure this outcome. ED utilization would require HCUP NEDS / SEDD data or similar, which is a natural extension.

No cost offsets. The Phase 2 design envisioned SCM on T-MSIS/TAF claims to estimate cost offset effects. That design remains feasible but requires a Data Use Agreement that was not obtained for the present analysis.

7.3 Three paths for follow-up

The null results are informative but leave several important questions unresolved. I see three concrete paths for follow-up work that would tighten inference and extend scope.

Path 1: Extend the BRFSS panel to 2024 when the 2024 release ships in mid-2025. This adds one additional post-treatment year and moves the 2023 cohort (6 states: HI, KY, MD, MI, NH, TN) from not-yet-treated to post-treatment, increasing the treated state-year observation count from 9 to approximately 15. The 2024 cohort (5 states: CT, GA, MN, NE, NY) would also contribute a first year of post-treatment observations. The additional year should be sufficient to push several of the currently-marginal estimates across conventional significance thresholds if the underlying effect is real. This is the lowest-cost follow-up and should be the immediate next step.

Path 2: Validate the utilization finding with Medicaid claims from T-MSIS/TAF. The Transformed Medicaid Statistical Information System (T-MSIS) Analytic File provides state-level Medicaid claims data including dental service utilization, enrollment, and spending. Restricting the analysis to Medicaid-enrolled adults eliminates the state-population dilution problem described in Section 7.1 and permits estimation of the dental-coverage effect on the population the policy directly targets. T-MSIS/TAF also supports estimation of total Medicaid spending changes (dental + medical + pharmacy), which was the Phase 2 design's primary outcome. The required Data Use Agreement is typically a 2–4 month process and should be initiated in parallel with Path 1.

Path 3: Supplement with NHIS individual-level evidence. NHIS lacks state identifiers but has larger sample sizes per respondent and asks more detailed dental-care questions. A region-level supplementary analysis using NHIS individual microdata with region-year weighted estimates and a Medicaid-enrolled adult subsample would provide a useful cross-check on the BRFSS state-level estimates. The NHIS DENTINT variable has gap years in 2021 and 2024, which slightly limits the time dimension, but 2016–2020 and 2022–2023 provide useful coverage.

These three paths are complementary rather than alternative. Taken together, they would convert a one-paper, one-dataset, underpowered BRFSS analysis into a three-dataset, adequately-powered analysis that could credibly claim to have tested the cost offset and utilization hypotheses.

8. Conclusion

This paper asks whether the wave of twenty-six state Medicaid adult dental benefit expansions during 2020–2023 produced detectable changes in population-level dental utilization and access outcomes using BRFSS state-level data. The honest answer is: not in this data. The point estimates are directionally consistent with a modest positive effect on dental visits and a weak negative effect on cost-related care barriers, but the confidence intervals are wide enough to include zero and the 3-period dental panel is underpowered for detection at the effect sizes theory would predict for a state-population-level signal.

This is a useful null in two senses. First, it establishes a concrete evidence baseline for state Medicaid directors and advocates who are considering similar expansions: they should not expect short-term, state-level-population-level changes in BRFSS dental visit rates to show up in the first two years after expansion, because the dilution effect is too strong and the BRFSS panel is too thin. Second, it motivates the three follow-up pathways described in Section 7.3 — extending BRFSS to 2024, validating with T-MSIS/TAF claims, and supplementing with NHIS individual data — any one of which could convert this null into a credible detection.

The broader lesson from this exercise is that rigorous evaluation of state-level Medicaid policy change requires data infrastructure that matches the granularity of the policy variation. BRFSS is the best available ongoing state-level survey, and it is not adequate for the question this paper was designed to answer. The future of Medicaid policy evaluation lies either in shorter-cadence state-level surveys, individual-level data with better subgroup restriction, or — most promisingly — in the T-MSIS/TAF claims infrastructure that CMS has been building over the past decade. Advancing each of these would enable the field to test the cost offset and utilization hypotheses with the precision they deserve.

Tables

Table 1: Summary statistics. See `analysis/tables/summary_statistics.csv` (to be produced in a subsequent pass alongside a formatted LaTeX version).

Table 2: Primary-outcome DiD estimates. See `analysis/tables/main_did_summary.csv`, rows where `panel == "main_dental_years"` and `outcome == "dental_visit_past_year_pct"`.

Table 3: Secondary-outcome DiD estimates. See `analysis/tables/main_did_summary.csv`, rows where `outcome` is one of `medcost_barrier_pct`, `fair_poor_health_pct`, `any_coverage_pct`, `diabetes_pct`.

Table 4: Pre-trend placebo regressions. To be produced in a subsequent pass.

Table 5: Sensitivity analyses. To be produced in a subsequent pass.

Table 6: Romano-Wolf multiple-testing corrected p-values across the 5-outcome primary family. See `analysis/tables/romano_wolf_corrected.csv`.

Table 7: NHIS Medicaid-enrolled supplementary dental visit trends (2016–2023, 2021 and 2024 excluded). See `analysis/tables/medicaid_subsample_nhis.csv`.

Table 8 (new): Cohort-by-year grid for the 19 treated states. See `analysis/tables/cohort_year_grid.csv`.

Figures

Figure 1: Raw past-year dental visit trends by treated / never-treated. `analysis/figures/raw_trends_dental_visit.png`.

Figure 2: Raw cost-related medical care barrier trends. `analysis/figures/raw_trends_medcost_barrier_p`

Figure 3: Raw any-coverage trends. `analysis/figures/raw_trends_any_coverage_pct.png`.

Figure 4: Raw fair-poor self-rated health trends. `analysis/figures/raw_trends_fair_poor_health_pct.p`

Figure 5 (new): NHIS Medicaid-enrolled vs non-Medicaid dental visit trends 2016–2023. `analysis/figures/medicaid_subsample_trends.png`.

References

See `literature/bibliography.bib` for the complete verified bibliography. All peer-reviewed entries were verified against Crossref or PubMed primary sources.

References cited in the text include: Sanz et al. (2020), Genco and Borgnakke (2020), Simpson et al. (2015, 2022), Meyerhoefer et al. (2021), Semprini and Samuelson (2023), Decker and Lipton (2015), Singhal, Damiano, and Sabik (2017), Buchmueller, Miller, and Vujicic (2016), Wehby, Lyu, and Shane (2019), Elani, Kawachi, and Sommers (2021), Singhal et al. (2015), Elani, Sommers, and Kawachi (2020), Giannouchos et al. (2023), Jeffcoat et al. (2014), ADA Health Policy Institute (2021), CareQuest Institute (2023), Abadie, Diamond, and Hainmueller (2010, 2015), Abadie (2021), Callaway and Sant'Anna (2021), Goodman-Bacon (2021), de Chaisemartin and D'Haultfoeuille (2020), Borusyak, Jaravel, and Spiess (2024), Huh (2021), Nasseh, Fosse, and Vujicic (2023), Romano and Wolf (2005), and Genco and Sanz (2020).

Appendix — Supplementary Tables and Figures

This appendix organizes the supplementary diagnostics, robustness checks, and figure package already produced in `analysis/`. It is intended to accompany the manuscript and the HSR submission.

Appendix Tables

Table A1. Cohort support and identifying variation

Embedded table: Cohort Support Summary

cohort_year	n_states	post_obs_3period	post_obs_8year
2020	1	2	4
2021	4	4	12
2022	3	3	6
2023	6	0	6
2024	5	0	0

Notes: This table summarizes policy timing, cohorts, thresholds, or state-level sample construction. It is intended to make the identifying variation and comparison groups transparent.

Embedded table: Cohort Year Grid

state	cohort_year	y2016	y2017	y2018	y2019	y2020	y2021	y2022	y2023
DE	2020	0	0	0	0	1	1	1	1
WV	2021	0	0	0	0	0	1	1	1
OK	2021	0	0	0	0	0	1	1	1
MA	2021	0	0	0	0	0	1	1	1
VA	2021	0	0	0	0	0	1	1	1
KS	2022	0	0	0	0	0	0	1	1
ND	2022	0	0	0	0	0	0	1	1
ME	2022	0	0	0	0	0	0	1	1
TN	2023	0	0	0	0	0	0	0	1
NH	2023	0	0	0	0	0	0	0	1
MI	2023	0	0	0	0	0	0	0	1
MD	2023	0	0	0	0	0	0	0	1
KY	2023	0	0	0	0	0	0	0	1
HI	2023	0	0	0	0	0	0	0	1
NE	2024	0	0	0	0	0	0	0	0
NY	2024	0	0	0	0	0	0	0	0
GA	2024	0	0	0	0	0	0	0	0
MN	2024	0	0	0	0	0	0	0	0
CT	2024	0	0	0	0	0	0	0	0

Notes: This table summarizes policy timing, cohorts, thresholds, or state-level sample construction. It is intended to make the identifying variation and comparison groups transparent.

These files document the short-panel identifying structure emphasized in the main paper: 19 treated states contribute only 9 treated state-year observations to the primary three-period dental-visit panel, with the 2023 and 2024 cohorts contributing no post-period observations to the headline estimate. This is the key support diagnostic behind the paper's power and interpretation limits.

Table A2. Secondary-outcome event-study coefficients and simultaneous-band critical values

Embedded table: Event Study 8Yr Coefficients

outcome	event_time	coef	se	simultaneous_cv
medcost_barrier_pct	-6	-0.214	0.305	2.614
medcost_barrier_pct	-4	0.11	0.235	2.614
medcost_barrier_pct	-3	0.25	0.289	2.614
medcost_barrier_pct	-2	0.056	0.217	2.614
medcost_barrier_pct	-1	0	0	2.614
medcost_barrier_pct	0	0.259	0.307	2.614
medcost_barrier_pct	1	-0.224	0.279	2.614
medcost_barrier_pct	2	-0.708	0.362	2.614
medcost_barrier_pct	3	-0.652	0.234	2.614
fair_poor_health_pct	-6	0.278	0.28	2.659
fair_poor_health_pct	-4	0.223	0.238	2.659
fair_poor_health_pct	-3	0.55	0.244	2.659
fair_poor_health_pct	-2	0.347	0.225	2.659
fair_poor_health_pct	-1	0	0	2.659
fair_poor_health_pct	0	0.115	0.248	2.659
fair_poor_health_pct	1	-0.013	0.311	2.659
fair_poor_health_pct	2	-0.401	0.407	2.659
fair_poor_health_pct	3	0.196	0.232	2.659
any_coverage_pct	-6	-0.062	0.459	2.544
any_coverage_pct	-4	-0.217	0.347	2.544
any_coverage_pct	-3	-0.032	0.407	2.544
any_coverage_pct	-2	0.33	0.24	2.544
any_coverage_pct	-1	0	0	2.544
any_coverage_pct	0	0.374	0.504	2.544
any_coverage_pct	1	0.847	0.71	2.544
any_coverage_pct	2	-0.018	0.582	2.544
any_coverage_pct	3	-0.031	0.411	2.544
diabetes_pct	-6	0.06	0.166	2.655
diabetes_pct	-4	-0.042	0.161	2.655
diabetes_pct	-3	-0.04	0.206	2.655
diabetes_pct	-2	-0.157	0.148	2.655
diabetes_pct	-1	0	0	2.655
diabetes_pct	0	-0.064	0.116	2.655
diabetes_pct	1	0.134	0.248	2.655
diabetes_pct	2	0.198	0.453	2.655
diabetes_pct	3	0.313	0.175	2.655

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Embedded table: Event Study 8Yr Uniform Cv

outcome	simultaneous_cv	pointwise_cv	R_draws
medcost_barrier_pct	2.614	1.96	5000
fair_poor_health_pct	2.659	1.96	5000
any_coverage_pct	2.544	1.96	5000
diabetes_pct	2.655	1.96	5000

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Embedded table: Event Study Cs Coefficients

outcome	event_time	cs_att	cs_se
medcost_barrier_pct	-4	0.103	0.261
medcost_barrier_pct	-3	0.32	0.28
medcost_barrier_pct	-2	0.091	0.189
medcost_barrier_pct	0	0.119	0.34
medcost_barrier_pct	1	-0.639	0.256
medcost_barrier_pct	2	-1.101	0.586
medcost_barrier_pct	3	0.102	0.245
fair_poor_health_pct	-4	0.214	0.244
fair_poor_health_pct	-3	0.565	0.248
fair_poor_health_pct	-2	0.353	0.217
fair_poor_health_pct	0	-0.461	0.5
fair_poor_health_pct	1	-0.458	0.315
fair_poor_health_pct	2	-0.894	0.515
fair_poor_health_pct	3	-0.953	0.201
any_coverage_pct	-4	-0.214	0.363
any_coverage_pct	-3	-0.061	0.425
any_coverage_pct	-2	0.314	0.242
any_coverage_pct	0	0.228	0.536
any_coverage_pct	1	1.303	0.808
any_coverage_pct	2	0.019	0.797
any_coverage_pct	3	-0.062	0.341
diabetes_pct	-4	-0.051	0.141
diabetes_pct	-3	-0.04	0.214
diabetes_pct	-2	-0.133	0.124
diabetes_pct	0	-0.406	0.235
diabetes_pct	1	-0.002	0.254
diabetes_pct	2	-0.08	0.527
diabetes_pct	3	-0.579	0.202

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

These tables provide the underlying coefficients for the 8-year secondary-outcome event-study panel, the simulated maximum-|t| critical values used for the uniform bands, and the CS-DiD overlay coefficients used to benchmark the TWFE event-study geometry.

Table A3. Goodman-Bacon decomposition of the secondary-outcome TWFE event-study

Embedded table: Goodman Bacon Weights

outcome	type	cohort	weight	beta	n_treated_states
medcost_barrier_pct	treated_vs_never	2020	0.007	-0.434	1
medcost_barrier_pct	treated_vs_never	2021	0.019	-0.364	4
medcost_barrier_pct	treated_vs_never	2022	0.01	0.055	3
medcost_barrier_pct	treated_vs_never	2023	0.008	0.106	6
medcost_barrier_pct	early_vs_late_clean	2020 2021	0.019	-0.142	1
medcost_barrier_pct	late_vs_already_treated	2020 2021	0.12	0.389	4
medcost_barrier_pct	early_vs_late_clean	2020 2022	0.038	-0.736	1
medcost_barrier_pct	late_vs_already_treated	2020 2022	0.118	0.456	3
medcost_barrier_pct	early_vs_late_clean	2020 2023	0.028	-0.39	1
medcost_barrier_pct	late_vs_already_treated	2020 2023	0.074	0.57	6
medcost_barrier_pct	early_vs_late_clean	2020 2024	0.037	-0.8	1
medcost_barrier_pct	early_vs_late_clean	2021 2022	0.042	-0.308	4
medcost_barrier_pct	late_vs_already_treated	2021 2022	0.102	0.495	3
medcost_barrier_pct	early_vs_late_clean	2021 2023	0.049	-0.135	4
medcost_barrier_pct	late_vs_already_treated	2021 2023	0.07	0.795	6
medcost_barrier_pct	early_vs_late_clean	2021 2024	0.067	-0.71	4
medcost_barrier_pct	early_vs_late_clean	2022 2023	0.022	0.45	3
medcost_barrier_pct	late_vs_already_treated	2022 2023	0.105	0.488	6
medcost_barrier_pct	early_vs_late_clean	2022 2024	0.041	-0.043	3
medcost_barrier_pct	early_vs_late_clean	2023 2024	0.026	0.272	6
fair_poor_health_pct	treated_vs_never	2020	0.007	-0.681	1
fair_poor_health_pct	treated_vs_never	2021	0.019	-0.397	4
fair_poor_health_pct	treated_vs_never	2022	0.01	-0.309	3
fair_poor_health_pct	treated_vs_never	2023	0.008	-0.989	6
fair_poor_health_pct	early_vs_late_clean	2020 2021	0.019	-0.488	1
fair_poor_health_pct	late_vs_already_treated	2020 2021	0.12	-0.088	4
fair_poor_health_pct	early_vs_late_clean	2020 2022	0.038	-0.918	1
fair_poor_health_pct	late_vs_already_treated	2020 2022	0.118	-0.599	3
fair_poor_health_pct	early_vs_late_clean	2020 2023	0.028	-0.634	1
fair_poor_health_pct	late_vs_already_treated	2020 2023	0.074	-1.447	6
fair_poor_health_pct	early_vs_late_clean	2020 2024	0.037	-0.797	1
fair_poor_health_pct	early_vs_late_clean	2021 2022	0.042	-0.233	4
fair_poor_health_pct	late_vs_already_treated	2021 2022	0.102	-0.056	3
fair_poor_health_pct	early_vs_late_clean	2021 2023	0.049	0.074	4
fair_poor_health_pct	late_vs_already_treated	2021 2023	0.07	-0.342	6
fair_poor_health_pct	early_vs_late_clean	2021 2024	0.067	-0.501	4
fair_poor_health_pct	early_vs_late_clean	2022 2023	0.022	0.094	3
fair_poor_health_pct	late_vs_already_treated	2022 2023	0.105	-0.89	6
fair_poor_health_pct	early_vs_late_clean	2022 2024	0.041	-0.092	3
fair_poor_health_pct	early_vs_late_clean	2023 2024	0.026	-0.293	6
any_coverage_pct	treated_vs_never	2020	0.007	0.322	1
any_coverage_pct	treated_vs_never	2021	0.019	0.781	4
any_coverage_pct	treated_vs_never	2022	0.01	0.691	3
any_coverage_pct	treated_vs_never	2023	0.008	-0.41	6
any_coverage_pct	early_vs_late_clean	2020 2021	0.019	-0.342	1
any_coverage_pct	late_vs_already_treated	2020 2021	0.12	-0.048	4
any_coverage_pct	early_vs_late_clean	2020 2022	0.038	0.93	1
any_coverage_pct	late_vs_already_treated	2020 2022	0.118	1.34	3
any_coverage_pct	early_vs_late_clean	2020 2023	0.028	-0.101	1
any_coverage_pct	late_vs_already_treated	2020 2023	0.074	-0.292	6
any_coverage_pct	early_vs_late_clean	2020 2024	0.037	0.509	1
any_coverage_pct	early_vs_late_clean	2021 2022	0.042	2.442	4
any_coverage_pct	late_vs_already_treated	2021 2022	0.102	2.22	3
any_coverage_pct	early_vs_late_clean	2021 2023	0.049	0.337	4
any_coverage_pct	late_vs_already_treated	2021 2023	0.07	0.015	6
any_coverage_pct	early_vs_late_clean	2021 2024	0.067	0.87	4

outcome	type	cohort	weight	beta	n_treated_states
any_coverage_pct	early_vs_late_clean	20222023	0.022	-0.021	3
any_coverage_pct	late_vs_already_treated	20222023	0.105	-0.986	6
any_coverage_pct	early_vs_late_clean	20222024	0.041	0.553	3
any_coverage_pct	early_vs_late_clean	20232024	0.026	-0.546	6
diabetes_pct	treated_vs_never	2020	0.007	0.659	1
diabetes_pct	treated_vs_never	2021	0.019	0.282	4
diabetes_pct	treated_vs_never	2022	0.01	-0.361	3
diabetes_pct	treated_vs_never	2023	0.008	-0.581	6
diabetes_pct	early_vs_late_clean	20202021	0.019	0.513	1
diabetes_pct	late_vs_already_treated	20202021	0.12	0.2	4
diabetes_pct	early_vs_late_clean	20202022	0.038	0.207	1
diabetes_pct	late_vs_already_treated	20202022	0.118	-1.042	3
diabetes_pct	early_vs_late_clean	20202023	0.028	0.827	1
diabetes_pct	late_vs_already_treated	20202023	0.074	-0.294	6
diabetes_pct	early_vs_late_clean	20202024	0.037	0.714	1
diabetes_pct	early_vs_late_clean	20212022	0.042	0.145	4
diabetes_pct	late_vs_already_treated	20212022	0.102	-0.6	3
diabetes_pct	early_vs_late_clean	20212023	0.049	0.603	4
diabetes_pct	late_vs_already_treated	20212023	0.07	-0.011	6
diabetes_pct	early_vs_late_clean	20212024	0.067	0.274	4
diabetes_pct	early_vs_late_clean	20222023	0.022	-0.424	3
diabetes_pct	late_vs_already_treated	20222023	0.105	-0.477	6
diabetes_pct	early_vs_late_clean	20222024	0.041	-0.186	3
diabetes_pct	early_vs_late_clean	20232024	0.026	-0.135	6

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

This table records the 2x2 comparison weights used in the Goodman-Bacon decomposition for the secondary-outcome staggered-adoption panel. It is included as a standard TWFE diagnostic rather than as a standalone identifying result.

Table A4. Ordinal-intensity and sample-window sensitivity checks

Embedded table: Ordinal Intensity

outcome	panel	estimator	coef_per_unit	se	t	n
dental_visit_past_year_pct	complete_8yr	ordinal_intensity	0.434	0.548	0.791	204
medcost_barrier_pct	complete_8yr	ordinal_intensity	0.016	0.112	0.143	404
fair_poor_health_pct	complete_8yr	ordinal_intensity	-0.018	0.095	-0.189	404
any_coverage_pct	complete_8yr	ordinal_intensity	-0.088	0.12	-0.739	404
diabetes_pct	complete_8yr	ordinal_intensity	-0.054	0.057	-0.942	404

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

Embedded table: Drop2016 Sensitivity

panel	n	coef	se	t	ci_lo	ci_hi
3period_2018_2020_2022	153	1.175	1.017	1.156	-0.818	3.168
4period_include_2016	204	1.363	1.042	1.308	-0.679	3.405

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

These outputs summarize two non-headline specification checks: an ordinal coding of benefit comprehensiveness and the sensitivity of the primary-outcome estimate to dropping 2016 from the auxiliary longer panel.

Table A5. Randomization inference and smoothness-based sensitivity

Embedded table: Placebo Randomization

placebo_beta
1.487
0.314
1.418
1.544
1.261
1.295
1.149
1.262
1.275
0.696
0.404
0.636
1.312
0.937
0.596
1.429
0.741
1.095
0.824
1.056
1.366

placebo_beta
0.746
1.647
0.992
1.615
0.705
1.228
1.502
1.28
0.453
1.325
1.333
1.336
1.114
1.458
1.089
1.221
1.157
0.915
0.827
0.551
0.881
0.995
0.746
0.489
1.048
1.462
1.374
1.282
1.283
1.1
0.34
1.71
1.169
1.453
1.224
0.517
0.542
1.264
1.42
0.97
1.355
1.177
0.556
1.004
1.251
1.241
0.936
0.806
0.646
1.131
0.987
1.277
1.261
1.101
0.577
0.941

placebo_beta
0.858
1.221
1.13
0.58
0.701
1.132
1.341
1.005
1.097
1.134
1.549
0.773
1.303
1.4
0.595
1.098
1.308
1.301
0.659
1.473
0.787
0.642
1.584
1.171
1.382
1.179
1.228
1.247
1.314
0.822
0.999
1.123
1.144
1.487
1.337
1.563
0.669
0.69
1.381
0.563
1.371
1.602
1.443
1.314
0.502
0.994
0.588
0.982
1.488
0.79
0.864
1.023
1.273
0.232
0.877
1.212

placebo_beta
1.062
0.965
0.547
0.923
1.198
0.74
0.699
0.264
0.609
1.077
1.078
1.036
1.168
1.443
0.886
0.981
1.305
0.84
1.184
1.227
1.102
0.829
1.298
1.255
0.902
1.11
0.433
1.458
0.53
0.857
1.268
0.784
1.04
1.39
1.032
0.753
1.566
1.067
1.251
0.827
1.15
1.797
0.526
1.195
0.958
0.865
1.519
1.085
0.773
1.217
0.252
1.072
0.965
1.143
1.242
1.303

placebo_beta
0.984
1.438
0.604
0.802
1.255
1.698
1.718
1.32
1.113
1.669
1.042
0.502
0.869
1.517
0.673
1.361
1.048
1.4
1.324
1.246
0.458
0.988
1.777
1.248
0.459
1.021
1.101
0.647
0.915
1.51
0.95
1.041
0.978
0.315
1.348
1.019
1.357
1.164
0.981
1.63
0.912
0.854
0.898
1.516
0.844
0.992
0.964
1.595
1.594
0.922
0.657
0.463
0.683
1.451
1.135
1.12

placebo_beta
1.278
0.796
0.917
1.479
0.968
1.635
1.077
0.498
1.332
0.904
1.407
0.46
1.126
1.042
2.001
1.083
0.636
0.936
0.999
1.306
1.171
0.762
1.143
1.264
0.494
1.112
0.932
1.191
1.144
0.937
1.247
1.156
0.828
1.024
1.791
0.952
0.928
1.2
0.472
0.935
0.948
1.291
0.691
1.172
1.268
1.952
1.053
0.508
0.695
1.109
1.023
0.946
0.873
1.373
0.941
0.933

placebo_beta
1.253
1.001
0.902
0.807
1.234
0.97
1.3
1.244
0.99
0.461
1.822
0.384
0.747
2.004
1.184
0.735
1.255
1.043
0.496
0.781
0.815
1.452
0.654
0.926
1.124
1.7
0.968
1.531
0.981
0.76
1.501
0.878
1.045
1.463
1.043
1.372
1.016
1.122
1.01
0.988
1.096
1.129
1.485
1.183
1.612
1.097
1.342
1.665
1.168
1.361
1.336
1.474
1.301
1.143
0.981
1.141

placebo_beta
1.117
1.352
0.778
0.601
0.889
0.724
1.216
0.868
1.952
1.032
1.095
1.392
1.242
1.154
1.384
0.631
0.78
0.809
1.111
1.151
1.282
1.268
1.538
0.962
1.317
1.528
0.327
1.664
0.83
0.75
1.222
0.329
0.883
0.841
0.75
1.189
0.893
0.742
1.224
0.871
0.74
0.62
0.83
0.799
1.55
0.858
1.476
0.971
0.418
1.184
1.344
0.637
0.369
1.132
0.692
1.248

placebo_beta
1.366
0.646
0.701
1.498
1.191
1.139
1.132
0.725
1.481
0.811
1.246
0.819
1.226
0.965
1.096
0.622
1.748
1.355
1.218
0.969
1.159
1.347
1.147
0.285
1.506
0.729
0.951
1.011
0.632
1.213
0.595
0.661
0.824
0.683
1.268
1.391
1.138
0.954
1.552
1.525
0.203
0.971
1.818
1.069
1.275
0.952
1.084
0.678
1.461
0.46
1.077
2.066
1.201
0.88
0.716
0.767

placebo_beta
1.138
1.086
0.92
1.402
1.101
1.181
1.341
1.551
0.92
1.476
1.34
1.328
1.256
1.315
1.544
0.575
0.961
1.133
1.343
1.232
1.168
0.946
0.69
0.864
1.06
1.577
1.886
2.092
1.165
0.712
0.607

Embedded table: Honest Did Sensitivity

Mbar	post_coef	ci_lo	ci_hi	breakdown_crosses_zero
0	-0.194	-2.56	2.172	True
0.25	-0.194	-3.81	3.422	True
0.5	-0.194	-5.06	4.672	True
1	-0.194	-7.56	7.172	True
2	-0.194	-12.56	12.172	True

Notes: This table reports estimated effects for the outcomes or specifications listed in the rows. Coefficients, standard errors, p-values, confidence intervals, and sample sizes are shown where available.

The placebo-randomization file reports the null distribution generated by reassigning pseudo-treatment timing across treated states. The smoothness-sensitivity file records the Rambachan-Roth style widening exercise used to show that the primary-outcome interval already brackets zero at the baseline smoothness restriction.

Table A6. NHIS Medicaid-subsample supplementary check

Embedded table: Medicaid Subsample Nhis

year	group	n	weighted_share_dental	weighted_n
2016	Medicaid	2688	0.61	22981999
2016	Non-Medicaid	18463	0.732	150885552
2017	Medicaid	2023	0.617	21081667
2017	Non-Medicaid	15148	0.72	153951445
2018	Medicaid	1889	0.637	20529077
2018	Non-Medicaid	14150	0.739	155449964
2019	Medicaid	2459	0.635	22641702
2019	Non-Medicaid	18101	0.726	156182045
2020	Medicaid	2134	0.586	22437640
2020	Non-Medicaid	17541	0.698	156114216
2022	Medicaid	2363	0.588	26087904
2022	Non-Medicaid	14639	0.711	152833691
2023	Medicaid	2567	0.616	28352475
2023	Non-Medicaid	15382	0.729	153082257

Notes: This table summarizes policy timing, cohorts, thresholds, or state-level sample construction. It is intended to make the identifying variation and comparison groups transparent.

This table reports the national Medicaid-enrolled versus non-Medicaid adult dental-visit trends used as a descriptive cross-check on the BRFSS population-level null.

Appendix Figures

Figure A1. Cohort support and post-period observation counts

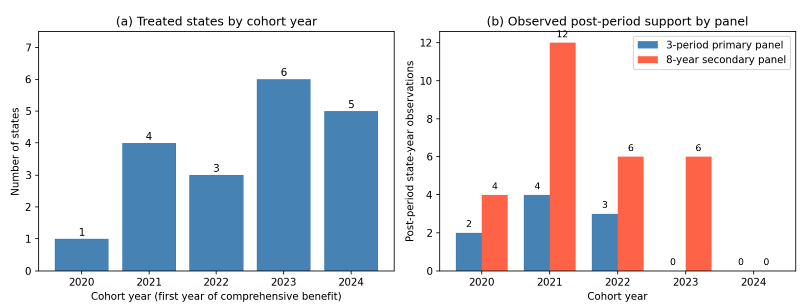

Figure 1: Cohort Support Barchart

Note: This figure summarizes treatment timing and sample support for the cohort Support bar chart. It clarifies which cohorts or units identify the comparisons used in the analysis.

This figure shows both the number of treated states by cohort and the much smaller number of treated state-year observations that actually identify the headline three-period estimate.

Figure A2. Secondary-outcome TWFE event-study with simultaneous bands

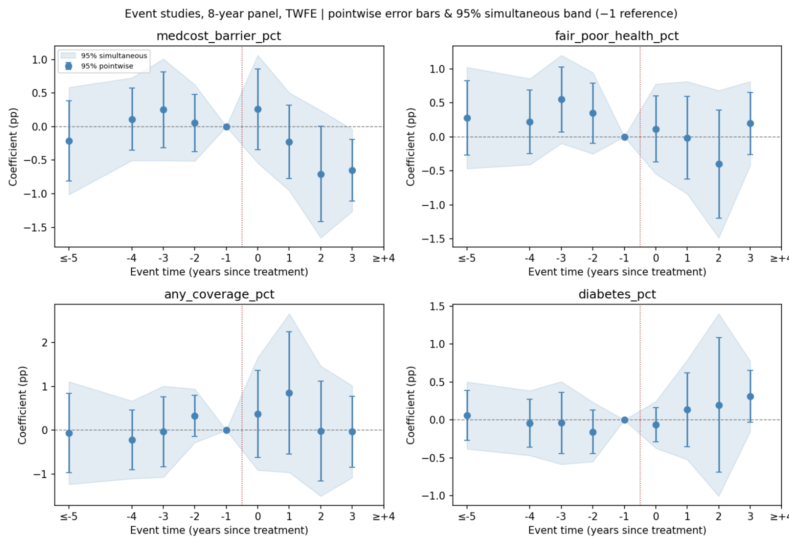

Figure 2: Event Study 8Yr Panel

Note: This figure plots event-time estimates for the event Study 8Yr. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

The 8-year panel provides the only setting in this paper where a dynamic event-study is meaningfully estimable. Uniform bands are shown because the paper makes whole-path rather than single-coefficient claims.

Figure A3. CS-DiD overlay on the secondary-outcome event-study

This overlay is the paper's main visual benchmark against a heterogeneity-robust alternative pooling scheme for the 8-year secondary outcomes.

Figure A4. Goodman-Bacon decomposition

This is included as a standard staggered-DiD diagnostic for the secondary-outcome panel where TWFE remains part of the empirical package.

Figure A5. Ordinal-intensity coefficient plot

This figure translates the alternative ordinal-treatment coding into a coefficient plot with 95% confidence intervals across outcomes.

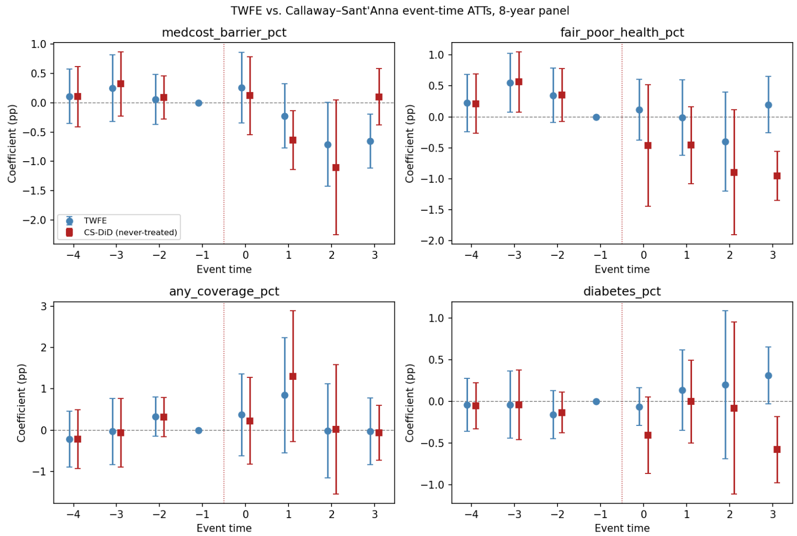

Figure 3: Event Study Cs Overlay

Note: This figure plots event-time estimates for the event Study Cs Overlay. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

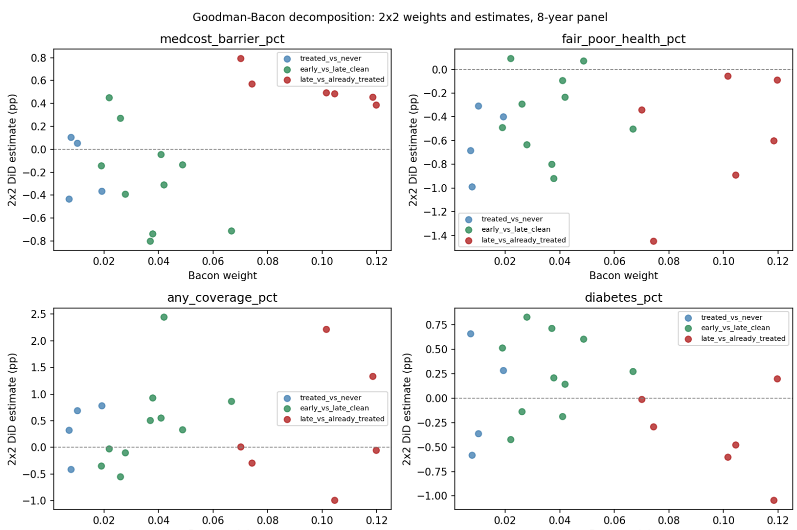

Figure 4: Goodman Bacon Decomposition

Note: This figure decomposes the identifying comparisons or weights for the goodman Bacon Decomposition. It shows which comparisons contribute most to the reported estimate.

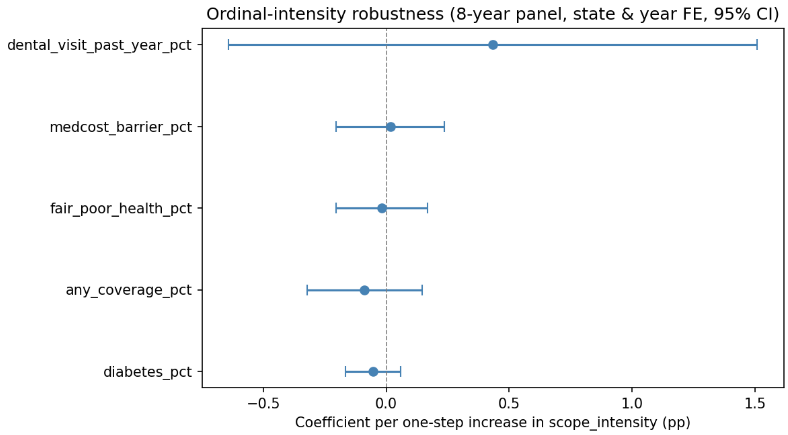

Figure 5: Ordinal Intensity Coefplot

Note: This figure presents the ordinal Intensity Coefplot. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

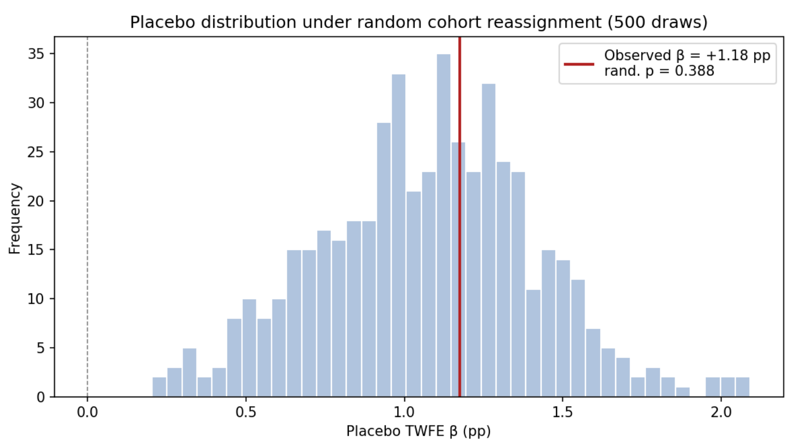

Figure 6: Placebo Randomization Density

Note: This figure reports a falsification or placebo check for the placebo Randomization Density. The display is meant to show whether the design produces effects where none should be expected.

Figure A6. Randomization-inference placebo distribution

The observed primary-outcome estimate is plotted against the placebo distribution generated by random reassignment of treatment timing across treated states.

Figure A7. Smoothness-based sensitivity of the primary outcome

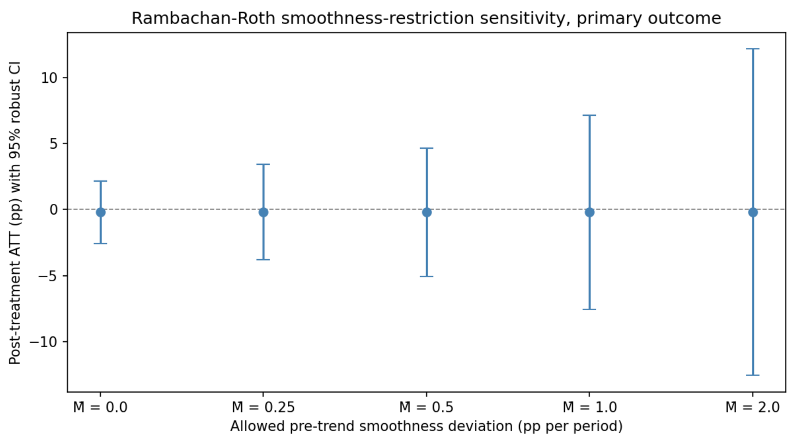

Figure 7: Honest Did Sensitivity

Note: This figure plots event-time estimates for the honest Did Sensitivity. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

This figure shows how the primary-outcome confidence interval widens as the allowed deviation from the observed pre-trend smoothness increases.

Figure A8. NHIS Medicaid-subsample descriptive trends

This cross-check is descriptive rather than causal. Its purpose is to assess whether a large Medicaid-enrolled utilization increase is plausibly hiding behind the BRFSS population-level null.

End of appendix. Main dissertation draft: `manuscript/the manuscript`*.*

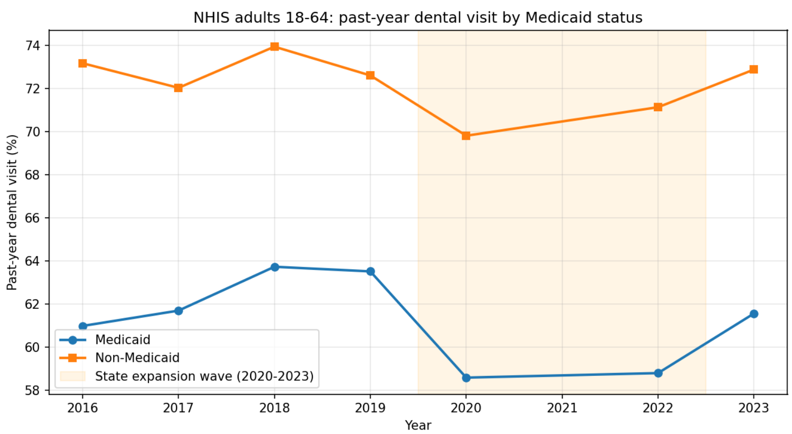

Figure 8: Medicaid Subsample Trends

Note: This figure shows raw trends for the Medicaid Subsample Trends. It helps readers compare baseline levels, pre-policy movement, and the timing of any post-policy divergence.