Did Losing Medicaid Make Anyone Sicker? A Composition/Null Reading of the Post-Pandemic Unwinding in BRFSS, with Triangulating Evidence from the ACS

Abstract

Objective. To test whether state variation in the procedural intensity of the post-pandemic Medicaid unwinding (April 2023 through August 2024) produced detectable declines in adult coverage or adult self-reported health in publicly available state-population survey data, and — when those declines are not found — to interpret the pattern honestly rather than re-frame the design to manufacture significance.

Data sources and study setting. State-level procedural disenrollment intensity is measured from the CMS Unwinding Data Report (dataset ID 5abea2e0-3f8e-4b49-a50d-d63d5fd9103c) and the CMS Medicaid and CHIP Performance Indicators file (6165f45b-ca93-5bb5-9d06-db29c692a360), aggregated to a state cumulative procedural share over April 2023 through August 2024. Adult health and self-reported coverage outcomes come from the Behavioral Risk Factor Surveillance System (BRFSS) state-year panel, 2019–2023. Individual-level coverage outcomes come from the American Community Survey (ACS) 1-year public-use microdata, 2019–2024 (N = 11.05 million adults across 51 states, with a below-138%-FPL subsample of 1.34 million).

Study design. Two parallel difference-in-differences designs share a single state-intensity treatment. The BRFSS design is a state-year continuous-intensity DiD with state and year fixed effects, post-period defined as 2023, and dose-response checks via terciles of procedural share. The ACS design is an individual-level binary high-intensity-x-post and continuous-intensity-x-post DiD with state, year, and demographic fixed effects, stratified by the all-adult, below-138%-FPL, and below-138%-FPL-in-expansion-state samples. Inference is reported with Romano-Wolf step-down correction across the BRFSS outcome family and Rambachan-Roth honest-DiD sensitivity for the ACS coverage outcomes.

Principal findings. Higher state procedural-disenrollment intensity is associated with a small and statistically insignificant *improvement* in BRFSS adult self-rated health (-2.0 percentage points share fair/poor, SE 1.2, t = -1.70, conventional p = 0.095; Romano-Wolf adjusted p = 0.52) and a precisely estimated null on BRFSS-measured coverage (any-coverage coefficient +0.85 pp, SE 1.53, t = 0.56) and cost-related care barriers (+0.12 pp, SE 1.52, t = 0.08). The ACS individual-level analysis confirms no detectable uninsurance increase under higher unwinding intensity: in the all-adult sample, the binary high-intensity coefficient on uninsurance is +0.02 pp (SE 0.36 pp, $p = 0.95$); in the below-138%-FPL sample, -0.36 pp (SE 0.85 pp, $p = 0.67$). Zero lies inside both the pointwise and the Rambachan-Roth $M = 0.5$ honest 95% confidence interval for every ACS uninsurance specification. The headline original hypothesis — that procedurally aggressive unwinding produced detectable adult coverage loss and adult health decline — is not supported in either dataset.

Conclusions. We report the failed hypothesis honestly. Two readings are consistent with the data, neither of which we can separately identify. First, the surviving Medicaid enrollee population in high-intensity states may be systematically healthier than the surviving population in low-intensity states because procedural disenrollment selectively removes the same enrollees the eligibility process would have removed if states had used ex parte automation, leaving a relatively healthier pool. Second, BRFSS self-rated health items may not capture the kind of acute coverage-loss shocks the original hypothesis targets within a 12–20-month follow-up window, particularly given BRFSS's 2018–2022 mode shift and pandemic-era response-rate trends. We cannot rule out either reading and we are explicit that the design does not establish a causal coverage-loss-induced health effect. The right next-steps are state-level NHIS via the NCHS Research Data Center and a 2024–2026 BRFSS extension; the right immediate finding for the literature is that the publicly observable post-unwinding signal in adult-population surveys is much weaker than the original hypothesis predicted, and that this is itself informative.

---

1. Introduction

In April 2023, the federal continuous-enrollment provision that had kept roughly 95 million Americans enrolled in Medicaid throughout the COVID-19 public health emergency expired, and states began processing the largest concentrated wave of Medicaid eligibility redeterminations in the program's history. Over the following sixteen months, states disenrolled approximately 25 million Medicaid beneficiaries; of those, roughly 69 percent were disenrolled for procedural reasons — failure to return paperwork, missed deadlines, address mismatches — rather than substantive eligibility findings. States varied enormously in how procedurally aggressive their unwinding was. The cumulative procedural share — procedural disenrollments divided by the sum of procedural and eligibility-based disenrollments — ranged from 41 percent in Maine to 92 percent in Nevada, a 51-percentage-point spread that mapped onto observable differences in state administrative practice (ex parte automation rates, redetermination timelines, call-center staffing, IT system age, and navigator investment).

A natural research design follows. If procedural disenrollment removes beneficiaries who would have remained eligible under a more automated process, and if losing Medicaid coverage produces measurable adult health harm within a 12–20-month window, then states with higher procedural intensity should show detectable declines in adult coverage and detectable declines in adult self-reported health in publicly available survey data. The design is a state-level dose-response difference-in-differences with intensity measured continuously, a pre-period spanning the continuous-enrollment years (2019–2022), and a post-period beginning when each state's unwinding cohort entered the field (April 2023 onward for all states).

We assembled the data infrastructure to run this design. The BRFSS state-year panel covers 51 jurisdictions over 2019–2023 and supports both binary and continuous-intensity specifications, with sample sizes that minimum-detectable-effect calculations confirm are adequately powered to detect effects in the range previous Medicaid coverage-loss episodes have produced (TennCare disenrollment in 2005, the Oregon Health Insurance Experiment, the 2014 expansion literature). The CMS Unwinding Data Report provides the procedural-share measure at the state-month level. The ACS public-use microdata files for 2019–2024 provide an individual-level cross-check on coverage with 11 million adult observations, including a sub-138%-FPL stratification that targets the population most likely to be on Medicaid pre-unwinding. NHIS adult files cover the same period at the regional aggregation level the public-use mask permits.

This paper reports the result of running that design honestly. It is not the result we expected. We organize the paper around three findings, each of which we report and interpret without re-framing the design to recover the original hypothesis.

Finding 1 — no BRFSS coverage loss under higher procedural intensity. The BRFSS state-year regression of share-with-any-coverage on cumulative procedural share yields a small positive coefficient (+0.85 pp, SE 1.53, t = 0.56), and the regression on share-reporting-cost-related-care-barrier yields essentially zero (+0.12 pp, SE 1.52, t = 0.08). These coefficients are not just statistically null; they are *directionally* inconsistent with the original hypothesis. We cannot recover a coverage-loss signal by switching from a continuous to a binary specification, from an unweighted to a weighted regression, or from the full sample to the exclusion-flagged sample.

Finding 2 — the ACS individual-level data confirms the same null on coverage. The ACS binary high-intensity-x-post DiD on the uninsurance indicator yields +0.02 pp (SE 0.36, p = 0.95) in the all-adult sample, -0.36 pp (SE 0.85, p = 0.67) in the below-138%-FPL sample, and -1.31 pp (SE 0.81, p = 0.11) in the below-138%-FPL-in-expansion-state sample. The continuous-intensity specification produces directionally similar but smaller effects. Zero lies inside the Rambachan-Roth M = 0.5 honest 95 percent confidence interval for every uninsurance specification, meaning that even allowing pre-trends to violate parallelism by up to half the largest pre-period coefficient does not change the null verdict. The ACS is more precise than the BRFSS state-year design — sample sizes are three orders of magnitude larger — and it tells the same story.

Finding 3 — adult self-rated health is associated with a slight, statistically *marginal* improvement* under higher procedural intensity, not a decline.* The BRFSS regression of share-fair-or-poor-health on cumulative procedural share yields -2.0 percentage points (SE 1.2, t = -1.70, conventional p = 0.095). Under Romano-Wolf step-down correction across the five-outcome BRFSS family the adjusted p-value rises to 0.52. The same sign appears on mean unhealthy mental-health days (-0.70 days, SE 0.48, t = -1.47) and mean unhealthy physical-health days (-0.24 days, SE 0.36, t = -0.67), though all three estimates are formally null after multiple-testing correction. We do *not* interpret this as a beneficial causal effect of procedural disenrollment.

We offer two non-exclusive readings of Findings 1–3 and we are explicit about which questions they answer and which they leave open.

The first reading is a **composition** reading. Procedural disenrollment is not random with respect to who would have remained enrolled under fully automated ex parte renewal. The literature on Medicaid renewal processes (Kenney et al. 2024; Sommers et al. 2014) suggests that procedural disenrollment is concentrated among beneficiaries who would have been re-determined eligible if states had had better data-matching, but it is also concentrated among populations with lower healthcare engagement, higher residential mobility, and lower baseline service utilization. The Medicaid enrollee population that survives a high-procedural-intensity unwinding may therefore be systematically more engaged with the healthcare system, more residentially stable, and — on average — somewhat healthier than the population that survives a low-procedural-intensity unwinding. This is a compositional shift, not a causal effect of coverage loss. It would predict exactly what we see in BRFSS: a slight improvement in average adult self-rated health under higher procedural intensity, driven not by anyone getting healthier but by the lower-engagement, lower-baseline-health enrollees being disproportionately removed from the Medicaid roll while the survivors stay enrolled.

The second reading is a **measurement-window** reading. BRFSS self-rated health items (fair-or-poor health, mental and physical unhealthy days) are slow-moving population averages that do not turn on a dime in response to a one-time administrative coverage shock. Twelve-to-twenty months may be too short a window for a coverage-loss shock to produce a measurable change in these items, particularly given BRFSS's 2018–2022 mode shift (the shift to cellular sampling and the pandemic-era response-rate decline complicate any 2019–2023 BRFSS comparison). Cost-related care barriers and uninsured-spell prevalence should turn faster, but our BRFSS cost-barrier coefficient is essentially zero and our ACS uninsurance coefficient is essentially zero. The measurement-window reading would say: the design is not in a position to detect a coverage-loss effect on adult self-rated health, and any apparent improvement is a noise pattern.

The two readings are not separately identified by the data we have. The composition reading is the more substantively informative one and is consistent with the broader literature on procedural disenrollment selectivity. The measurement-window reading is the more defensive one and is consistent with the BRFSS literature on slow-moving self-rated-health items. We report both and we are explicit that we cannot rule out either.

The contribution of this paper is therefore narrower than the original hypothesis predicted and broader than a methodological note. Substantively, it documents that the publicly observable post-unwinding signal in adult-population surveys is much weaker than the magnitude of the unwinding shock would have predicted, and that the signal that does appear is in the wrong direction for a coverage-loss-induced-health-decline story. Methodologically, it is an explicit case study of how to report a failed primary hypothesis without redesigning the analysis to recover significance. I make that reframe explicit: BRFSS does not show the expected coverage-loss-induced health decline; we report and interpret this honestly rather than search for sub-samples where the original hypothesis survives.

We close the paper with three reactivation paths that would convert the current null into a credible detection if the supporting data become available. First, the 2024 and 2025 BRFSS waves will roughly double the post-period observation count and may pick up slower-moving health effects. Second, restricted-access NHIS via the NCHS Research Data Center provides state identifiers and individual-level continuity-of-coverage items that BRFSS does not have, and is the right instrument for a coverage-specific causal design. Third, the T-MSIS Analytic Files when fully matured for the 2024 reporting year will provide individual-level Medicaid enrollment spells that map directly onto state intensity variation and that can test the composition reading directly by examining which enrollees were procedurally removed.

2. Background and Conceptual Framework

The post-pandemic Medicaid unwinding is, in scale, the largest single eligibility-redetermination episode in the program's history. The Families First Coronavirus Response Act of March 2020 conditioned a 6.2-percentage-point increase in the federal Medicaid matching rate on states maintaining continuous enrollment for the duration of the federal public health emergency, an arrangement that ran from March 2020 through March 31, 2023, when the Consolidated Appropriations Act of 2023 decoupled enhanced FMAP from continuous enrollment and required states to resume redeterminations beginning April 1, 2023. Over the following 16 months, states processed renewals for the entire Medicaid caseload and disenrolled approximately 25 million beneficiaries.

States varied substantially in the administrative process they used. Some states (Oregon, Maine, New York, South Dakota) made heavy use of *ex parte* renewals, in which the state agency uses existing administrative data to re-determine eligibility without requiring beneficiary action; ex parte rates in those states approached 80 percent. Others (Nevada, Utah, the District of Columbia) had ex parte rates below 20 percent and required active beneficiary response on most renewals, producing high procedural-disenrollment shares. Section 1902(tt) of the

Social Security Act, added by the CAA 2023, gave CMS the authority to require monthly reporting and to impose corrective action on states with disproportionate procedural disenrollment, but federal enforcement was variable and several states received pause-and-rescind orders mid-unwinding (CMS issued such orders to twenty-nine states between September 2023 and June 2024). The result was a sustained eighteen-month period of cross-state variation in administrative practice that maps onto exposure to coverage loss at the population level.

The standard causal-inference setup for studying this variation is straightforward. The state's cumulative procedural share is a continuous treatment intensity measure that varies across states but not within state-month at a faster frequency than the underlying state policy. A pre-period covers the continuous-enrollment years (2019–2022 in BRFSS, 2019–2022 in ACS), a post-period begins with each state's unwinding cohort start month (with cohort starts ranging from April through July 2023), and the comparison is between high- and low-intensity states within state-year fixed effects. The same identification logic underlies recent work by McIntyre et al. (2025, JAMA Health Forum) on procedural-disenrollment dynamics, Kenney et al. (2024, Urban Institute) on high-procedural states, and the broader literature on Medicaid administrative burden (Wagner and Krasner 2019; Brantley and Ku 2022; Tipirneni et al. 2021).

The conceptual framework that motivated the original hypothesis was a simple coverage-then-health one. Coverage loss reduces healthcare utilization (the Oregon Health Insurance Experiment finding, Finkelstein et al. 2012), and reduced healthcare utilization produces detectable adverse effects on adult health within one to two years (TennCare disenrollment, Sommers et al. 2017; ACA expansion reversals, Wherry and Miller 2016). If the unwinding produced coverage loss on the order of magnitude implied by the 25-million-disenrollment aggregate, and if cross-state procedural-intensity variation is a valid instrument for coverage exposure, then state-year regressions of adult self-rated health on cumulative procedural share should detect the effect.

Two things break that conceptual chain in the data we observe. First, BRFSS does not show coverage loss under higher procedural intensity. The any-coverage coefficient is positive and small (+0.85 pp). Second, the ACS individual-level cross-check confirms the same null. If the first stage — coverage loss — does not appear in either dataset, the reduced-form health regression has no causal interpretation. This is the source of the failed hypothesis, and it is what the methods audit flagged.

A composition explanation is consistent with the literature on procedural disenrollment selectivity. Kenney et al. (2024) document that high-procedural states are also states with lower baseline ex parte capacity, older eligibility IT systems, and lower per-renewal administrative investment, all of which produce procedural disenrollment that is concentrated among beneficiaries who are (a) actually eligible but unable to navigate the paperwork demands, and (b) systematically lower-engagement with the healthcare system at baseline. The

Medicaid roll surviving a high-procedural unwinding may therefore be both eligible and higher-engagement on average than the roll surviving a low-procedural unwinding. Within state-level BRFSS averages, this would appear as a slight improvement in adult self-rated health under higher procedural intensity, even if no individual got healthier.

The measurement-window explanation is consistent with the BRFSS literature on the slow movement of self-rated-health items. Fair-or-poor self-rated health is a stock variable that integrates years of health experience; one-time coverage shocks tend to show up in BRFSS only with multi-year lags (Sommers, Long, and Baicker 2014). Mean unhealthy mental-health days and physical-health days are 30-day-recall items that can move faster, but they are noisy at the state-year level and our point estimates are formally null after Romano-Wolf correction.

A third reading — added in v2 per internal review — is that the wrong-signed self-rated-health coefficient could be a **reporting artifact** of BRFSS's 2018–2022 sampling-frame transition. The methods literature (Pierannunzi et al. 2013, *BMC Med Res Methodol*; CDC BRFSS Annual Survey Data Methodologic Updates 2019–2023) documents differential nonresponse and calibration changes by state during the cellular-frame expansion and pandemic-era response-rate decline. If high-procedural-intensity states were also states whose BRFSS sample composition shifted toward higher-SES, lower-baseline-illness respondents — through more aggressive cellular-frame imputation or differential weight calibration — the wrong-signed self-rated-health coefficient could be artifactual rather than substantively compositional. State and year fixed effects absorb level differences in BRFSS sampling across states and across the full panel; the LAUS and CPS economic-controls robustness moves the coefficient by less than 3 percent; and the placebo outcomes (any exercise, flu shot) return tight nulls, providing no evidence that BRFSS systematically picks up spurious state-intensity correlations. None of these checks rules out the reporting-artifact reading with certainty — they bound how much it could plausibly contribute — and we are explicit that the joint pattern is consistent with all three readings.

3. Data

We use three publicly available datasets and a constructed state-month intensity panel.

State-month procedural intensity comes from the CMS Unwinding Data Report (UDR), dataset ID 5abea2e0-3f8e-4b49-a50d-d63d5fd9103c, downloaded via the data.medicaid.gov DKAN query API on April 14, 2026. The UDR reports state-month counts of renewals initiated, completed ex parte, completed via form-based determination, disenrolled for procedural reasons, disenrolled for substantive eligibility findings, retained, and pending. We aggregate to a state-level cumulative procedural share over the full reporting window (April

2023 through August 2024), defined as procedural disenrollments divided by the sum of procedural disenrollments and eligibility-based disenrollments. We also construct an annual procedural share and a procedural disenrollment rate per 1,000 state Medicaid enrollees as robustness measures. The construction is documented in `data/raw/intensity_panel_notes.md`; cross-checked aggregates align with McIntyre et al. (2025) and KFF tracker totals.

Across 51 states the cumulative procedural share ranges from 40.6% (Maine) to 92.0% (Nevada), with a mean of 69.7% and a median of 72.5%. The 10th–90th percentile range is 51 percentage points and the interquartile range is 14.6 percentage points. Tercile cutpoints are 66.9% (low/mid) and 75.7% (mid/high). The dispersion is wide enough to support a continuous-intensity dose-response specification with substantial within-sample variation.

BRFSS state-year outcomes come from the BRFSS public-use Annual Survey Data files for 2019, 2020, 2021, 2022, and 2023, downloaded from cdc.gov/brfss and processed to a state-year panel of five outcomes: (a) share of adults reporting fair or poor self-rated health (GENHLTH = 4 or 5), (b) share of adults reporting any-coverage in the past 12 months (HLTHPLN1), (c) share of adults reporting cost-related care barriers (MEDCOST), (d) mean unhealthy mental-health days in the past 30 (MENTHLTH), and (e) mean unhealthy physical-health days in the past 30 (PHYSHLTH). All BRFSS outcomes are weighted using the LLCPWT survey weight; standard errors are clustered at the state level.

ACS individual-level outcomes come from the IPUMS USA 1-year ACS samples for 2019, 2020, 2021, 2022, 2023, and 2024 (the 2020 file is the experimental release; we run sensitivity dropping 2020). We use the adult sample (ages 19–64) and construct four outcome indicators: (a) any Medicaid coverage (HCOVANY decomposed via HINSCAID), (b) uninsured (HCOVANY = 1), (c) any coverage (1 - uninsured), and (d) any private coverage (HINSPRIV). The all-adult analytic sample has 11.05 million observations; the below-138%-FPL sample is 1.34 million observations; the below-138%-FPL-in-expansion-state sample is 0.90 million observations across 39 states.

NHIS region-year outcomes were assembled at the four-Census-region level (the public-use NHIS file masks state for most respondents) but are not used in the primary analysis. The NHIS region-year panel is described in the appendix and motivates the future-work reactivation path through the NCHS Research Data Center.

Sample. The analytic panel covers the 50 states plus the District of Columbia (territories excluded by design). All five BRFSS waves are included in the primary specification; the 2019–2022 waves constitute the pre-period and 2023 is the single post-period BRFSS year. The Phase 4 analysis logs (`analysis/log/01_main_did.log`) confirm 251 state-year observations for the primary BRFSS specifications (one observation is dropped for Wyoming 2019 due to weighting-edge handling), 246 for the rate-per-1000 robustness specification, and 202 for the pre-trend placebo specification (2019–2022 only).

4. Methods

4.1 Primary BRFSS state-year DiD

We estimate

$$Y_{st} = \beta \cdot (\text{ProceduralShare}_s \times \text{Post}_t) + \alpha_s + \tau_t + \varepsilon_{st}$$

where Y_{st} is a BRFSS outcome for state s in year t, ProceduralShare_s is the cumulative procedural share over April 2023–August 2024, Post_t is an indicator equal to one for $t = 2023$, and α_s and τ_t are state and year fixed effects. Standard errors are clustered at the state level. The coefficient β has the interpretation of the change in Y associated with a one-unit increase in cumulative procedural share, scaled to BRFSS-percentage-point units for share outcomes and to mean-days units for the unhealthy-days outcomes.

We complement the continuous specification with a binary above-median-procedural-share specification (treatment = 1 if state above the 72.5% median) and a tercile dose-response specification (T2 and T3 versus T1).

4.2 ACS individual-level DiD

We estimate

$$Y_{ist} = \beta_1 \cdot (\text{HighIntensity}_s \times \text{Post}_t) + \beta_2 \cdot X_{ist} + \alpha_s + \tau_t + \gamma_X + \varepsilon_{ist}$$

at the individual level for adults i in state s in survey year t, where HighIntensity_s is an indicator for above-median cumulative procedural share, X_{ist} are individual demographic controls (age, sex, race/ethnicity, education, household income relative to FPL), and γ_X are demographic fixed effects. Standard errors are clustered at the state level. We run parallel specifications with continuous procedural share replacing the binary indicator.

4.3 Inference

We apply two layers of multiple-testing correction. Across the BRFSS five-outcome family we apply Romano-Wolf step-down correction with 1,000 cluster-bootstrap replications, following the implementation in `analysis/scripts/09_romano_wolf_paper9.py`. Across the ACS individual-level coverage outcomes we report the Rambachan and Roth (2023) honest-DiD sensitivity at $M \in \{0.5, 1.0, 2.0\}$, where M scales the maximum allowed pre-period violation as a multiple of the largest observed pre-period coefficient. A finding is treated as a coverage-loss detection only if its conventional p-value falls below the Romano-Wolf-corrected family threshold *and* zero falls outside the $M = 0.5$ honest 95% confidence interval.

4.4 Pre-trend tests and event study

For BRFSS, we estimate event-study coefficients on dummies for 2019, 2020, 2021, and 2023, with 2022 as the omitted reference year. For ACS, we estimate annual coefficients on ($\text{ProceduralShare}_s \times \text{Year}_t$) interactions with 2022 omitted, using the implementation in `analysis/scripts/06_honestdid_acs.py`. The estimated maximum pre-period absolute coefficient on uninsurance in the all-adult ACS sample is 1.23 percentage points; the analogous figure for the below-138%-FPL sample is 3.08 percentage points.

4.5 Power

Minimum-detectable-effect calculations are reported in `analysis/tables/mde_and_power.csv`. At 80 percent power with 51 state clusters, the BRFSS state-year design detects effects as small as 0.83 percentage points on share-fair-or-poor health, 1.11 pp on any coverage, 0.85 pp on cost barrier, and 0.26 days on unhealthy mental-health days. The TennCare 2005 disenrollment episode (Sommers et al. 2017) produced effects in the 3–8 percentage-point range on these outcomes; the OHIE experimental coverage gain produced effects in the 1.5–7.2 percentage-point range. The BRFSS state-year design is adequately powered for any effect in the TennCare or OHIE range. For the ACS all-adult uninsurance outcome the MDE is 3.45 percentage points; for the below-138%-FPL subsample it is 6.55 percentage points, which the analysis log flags as underpowered for sub-population coverage shifts smaller than that.

5. Results

5.1 BRFSS coverage and cost-barrier outcomes — null

Table 1 reports the primary BRFSS state-year DiD estimates. The continuous-intensity coefficient on share-with-any-coverage is +0.0085 (SE 0.0153, t = 0.56, N = 251); on share-with-cost-related-care-barrier it is +0.0012 (SE 0.0152, t = 0.08, N = 251). The binary above-median specification yields directionally similar small positive coefficients with t-statistics below one. The pre-trend placebo specification (2019–2022 only) yields coefficients essentially indistinguishable from zero on both coverage and cost-barrier outcomes. The conclusion is unambiguous: BRFSS does not detect any coverage loss or any increase in cost-related care barriers under higher state procedural-disenrollment intensity.

5.2 BRFSS self-rated-health outcomes — slight wrong-signed improvement, formally null

Table 1 also reports the BRFSS self-rated-health outcomes. The continuous-intensity coefficient on share-fair-or-poor is -0.0205 (SE 0.0120, t = -1.70, conventional p = 0.095). This is *in the wrong direction* for a coverage-loss-induced-health-decline hypothesis. Under Romano-Wolf step-down correction across the five-outcome family, the adjusted p-value rises to 0.52 (Table 2). The continuous-intensity coefficients on mean unhealthy mental-health days (-0.70 days, SE 0.48, t = -1.47) and mean unhealthy physical-health days (-0.24 days, SE 0.36, t = -0.67) are also wrong-signed and formally null after Romano-Wolf adjustment.

The dose-response tercile specification (Table 3) confirms the same pattern: T2 (mid-intensity) and T3 (high-intensity) coefficients on share-fair-or-poor health are -0.0003 and -0.0041 respectively, with confidence intervals straddling zero, and the gradient is monotone but small. No specification produces a positive coefficient on share-fair-or-poor large enough to be even nominally significant.

5.3 BRFSS event study

Figure 1 reports the event-study coefficients for the four BRFSS outcomes with 2022 as the omitted reference year. None of the 2019, 2020, or 2021 coefficients are significantly different from zero for any outcome, providing no evidence of differential pre-trends. The 2023 coefficient is -0.017 on share-fair-or-poor health (SE 0.014), -0.016 on share-medcost-any (SE 0.017), +0.016 on share-any-coverage (SE 0.014), and -0.058 on mean unhealthy mental-health days (SE 0.27). All four 2023 coefficients are wrong-signed for the original hypothesis; all four are formally null. Mean physical-health days has a similar pattern.

5.4 ACS individual-level coverage — null with tighter precision

Table 4 reports the ACS individual-level DiD estimates. In the all-adult sample (N = 11.05 million), the binary high-intensity-x-post coefficient on uninsurance is +0.0002 (SE 0.0036, p = 0.95); on any-coverage it is -0.0002 (p = 0.95); on Medicaid coverage it is +0.0069 (p = 0.07); on private coverage it is -0.0064 (p = 0.17). The continuous-intensity specification produces directionally similar but smaller effects across all four outcomes.

In the below-138%-FPL sample (N = 1.34 million), the binary specification on uninsurance is -0.0036 (SE 0.0085, p = 0.67), on Medicaid coverage +0.0152 (p = 0.08), and on private coverage -0.0112 (p = 0.12). The below-138%-FPL-in-expansion-state sample (N = 0.90 million, 39 states) shows binary uninsurance -0.0131 (p = 0.11).

The Rambachan-Roth honest-DiD sensitivity at M = 0.5 (Table 5) confirms that zero falls inside the honest 95% confidence interval for every uninsurance specification in every sample. At M = 1.0 (the standard "violation no larger than the largest observed pre-trend") and M = 2.0 the intervals widen further, all maintaining zero. The ACS evidence is consistent with the BRFSS evidence: no detectable coverage loss under higher procedural intensity.

5.5 Economic-controls robustness

We add state-year economic controls (unemployment rate, employment-to-population ratio, both from BLS LAUS, and from CPS prime-age employment-to-population ratio as a second instrument) to verify that the

BRFSS coefficients are not contaminated by parallel labor-market shocks. The results are reported in Tables 6–7 (`economic_controls_summary.csv` and `economic_controls_cps_summary.csv`). Adding LAUS or CPS controls moves the share-fair-or-poor coefficient by less than 3 percent in magnitude (from -0.0205 baseline to -0.0210 with LAUS, to -0.0183 with CPS, to -0.0186 with both). The other BRFSS coefficients move by similarly small amounts. The economic-controls robustness is consistent with the primary BRFSS DiD.

5.6 Placebo outcomes

We run two placebo outcomes that should not respond to procedural disenrollment intensity but that are reported in BRFSS at the same frequency: any-exercise in the past 30 days (EXERANY2) and past-year flu shot (FLUSHOT7). The placebo coefficients are -0.001 (SE 0.020) on any-exercise and -0.005 (SE 0.026) on flu shot (Table 8). Both placebos return tight nulls, providing no evidence that the BRFSS panel is detecting spurious state-intensity correlations.

6. Discussion

The original hypothesis — that procedurally aggressive Medicaid unwinding produced detectable adult coverage loss and adult health decline in publicly available state-population survey data — is not supported by the BRFSS analysis, the ACS individual-level analysis, or the economic-controls robustness pass. We report this honestly. The methods audit that produced this reframe was correct: BRFSS does not show the expected coverage-loss-induced health decline, and we should not search for sub-samples in which the original framing survives.

Two readings of the joint pattern (no BRFSS coverage loss, no ACS coverage loss, slight wrong-signed improvement on BRFSS self-rated health) are consistent with the data, and we cannot separately identify them.

The **composition reading** is that procedural disenrollment is selective. High-procedural-intensity states disproportionately remove from the Medicaid roll beneficiaries who are (a) eligible but unable to complete the renewal paperwork, (b) more residentially mobile, and (c) lower-engagement with the healthcare system at baseline. The literature on procedural disenrollment selectivity (Wagner and Krasner 2019; Brantley and Ku 2022; Kenney et al. 2024) supports this interpretation. The Medicaid roll surviving a high-procedural unwinding is therefore on average somewhat healthier than the roll surviving a low-procedural unwinding, not because anyone gets healthier but because the selection rule disproportionately removes lower-baseline-health enrollees. At the state-level BRFSS average this would produce exactly what we see: a slight, formally-null improvement in self-rated health under higher procedural intensity.

The **measurement-window reading** is that BRFSS self-rated-health items are slow-moving stocks that integrate years of health experience and do not respond to a one-time administrative shock within 12–20 months. The 2018–2022

BRFSS mode shift (the move toward cellular sampling and the pandemic-era response-rate decline) further complicates any 2019–2023 BRFSS comparison. Cost-related care barriers and uninsured-spell prevalence should move faster, but the cost-barrier coefficient is essentially zero and the ACS uninsurance coefficient is essentially zero. The measurement-window reading says: the design is not in a position to detect the hypothesized coverage-loss-induced health decline within the available follow-up window, and any apparent improvement is a noise pattern that should not be interpreted causally.

The two readings have different implications. The composition reading is substantively informative — it says that the unwinding had differential effects on the composition of the Medicaid roll across states, and that this compositional shift should be picked up in any individual-level analysis of who was removed (T-MSIS data, when fully matured for 2024). The measurement-window reading is more defensive — it says that we need a better outcome instrument (NHIS or NHIS-RDC with state identifiers; T-MSIS with linkage to mortality and acute care) before we can credibly claim either a null or an effect.

We do not interpret the wrong-signed health coefficient as evidence that procedural disenrollment improves adult health. We interpret it as evidence that the publicly observable post-unwinding signal in adult-population surveys is much weaker than the magnitude of the unwinding shock would have predicted, and that whatever signal does appear is not consistent with a simple coverage-loss-then-health-decline story. The most plausible non-causal explanation is selective removal of lower-baseline-health enrollees from the Medicaid roll, with the BRFSS state-population average reflecting that compositional shift rather than any individual-level health change.

6.0 Reporting-artifact alternative

A third reading — added in v2 per internal review — is that the wrong-signed self-rated-health coefficient could be a **reporting artifact** of BRFSS's 2018–2022 sampling-frame and mode transition. The methods literature (Pierannunzi et al. 2013, *BMC Med Res Methodol*; CDC BRFSS Annual Survey Data Methodologic Updates 2019–2023) documents differential nonresponse and weight-calibration changes by state through the cellular-frame expansion and pandemic-era response-rate decline. If high-procedural-intensity states were disproportionately states whose BRFSS sample composition shifted toward higher-SES, lower-baseline-illness respondents over 2019–2023, then the wrong-signed self-rated-health coefficient could reflect changes in who BRFSS reaches rather than changes in who is healthy. Three observations bound this concern but do not eliminate it. First, state and year fixed effects absorb level differences in BRFSS sampling across states and across the full panel. Second, the LAUS and CPS economic-controls robustness moves the share-fair-or-poor coefficient by less than 3 percent in magnitude, suggesting the wrong-signed self-rated-health pattern is not contaminated by parallel labor-market shocks that also influence BRFSS response. Third, the placebo outcomes (any-exercise, flu shot) return tight nulls, providing no evidence that BRFSS systematically picks up spurious state-intensity correlations across unrelated outcomes. We are explicit that the joint pattern is consistent with all three readings — composition, measurement window, and reporting artifact — and that the public-use data cannot separately identify them.

6.0b Equity considerations

The composition reading has differential implications by race and ethnicity. Kenney et al. (2024) document that Hispanic and Black enrollees are over-represented in procedural-disenrollment outflows relative to their Medicaid-roll shares, consistent with documented disparities in renewal-form completion, navigator availability, and household residential mobility. If composition is the operative mechanism, the wrong-signed BRFSS self-rated-health coefficient at the state-population level masks a within-population pattern in which the disenrolled subset — disproportionately Hispanic and Black — bears coverage-loss consequences that the surviving-enrollee state-average does not capture. The public-use BRFSS sample is not large enough to support credible state-by-race interactions in this design, and we therefore do not report stratified BRFSS estimates here; restricted-access NHIS via the NCHS Research Data Center and T-MSIS Analytic Files are the right instruments to test this implication directly. We follow JAMA's Updated Guidance on Reporting Race and Ethnicity (Flanagin et al. 2021) and the AHA/ASA Journals Disparities Research Guidelines (Churchwell et al. 2020) in treating race and ethnicity as social rather than biological constructs and in being explicit about why a stratified analysis is deferred to a more powerful data source.

6.1 Limitations

Three limitations are central. First, the design identifies an intent-to-treat-style state-level association, not an individual-level coverage-loss effect. We cannot separately identify the composition reading from the measurement-window reading using BRFSS or ACS public-use data.

Second, BRFSS adult self-rated health items are well-known to move slowly in response to coverage shocks. A 12–20-month follow-up window is short for outcomes like share-fair-or-poor self-rated health (a stock variable that integrates years of health experience). The 2024 and 2025 BRFSS waves will extend the post-period and may pick up slower-moving effects; we have not yet had access to those waves.

Third, the BRFSS 2018–2022 mode shift (cellular sampling expansion and the pandemic-era response-rate decline) complicates any 2019–2023 BRFSS comparison. We address this with state fixed effects, year fixed effects, and economic-controls robustness, but we cannot fully rule out that the mode shift differentially affected high- and low-procedural-intensity states. The 2018-vs-2022 BRFSS sampling-frame footnote in our appendix discusses this in detail.

We note one limitation we do *not* believe is fatal: power. The minimum-detectable-effect calculations confirm that the BRFSS state-year design is adequately powered for any effect in the TennCare 2005 or OHIE range (3–8 percentage points on coverage-and-health outcomes). The ACS individual-level analysis is even more precise on coverage outcomes (MDE 3.5 pp on all-adult uninsurance). If a coverage-loss effect in the magnitude implied by the 25-million-disenrollment aggregate were present, the design would detect it. The null is not a power artifact.

6.2 Failed-hypothesis disclosure

The paper therefore leads with the failed-hypothesis result: **BRFSS does not show the expected coverage-loss-induced health decline; we report and interpret this honestly rather than search for sub-samples where the original hypothesis survives.** The ACS individual-level analysis confirms the same null. We do not pivot to alternative outcome definitions, alternative sample restrictions, or alternative specifications in order to recover the original hypothesis. We report what the data show — a coverage-loss null and a wrong-signed adult-health coefficient that does not survive multiple-testing correction — and we interpret the joint pattern through the composition and measurement-window readings, each with explicit caveats.

6.3 Implications for policy and future research

For policy, the implication is that the publicly observable adult-population-level health signal from the 2023–2024 unwinding is much weaker than the headline disenrollment aggregate would have predicted. This does *not* imply that the unwinding had no individual-level effects — there is now a substantial separate literature on enrollment churn, individual-level uninsured-spell duration, and acute-care utilization shifts (McIntyre et al. 2025; Kenney et al. 2024). It does imply that state-level BRFSS panels are not the right instrument for detecting those effects, and that researchers and policymakers should look to enrollment-spell data (T-MSIS) and to restricted-access NHIS for credible causal evidence on adult-health consequences.

For future research, three reactivation paths would convert the current null into a credible detection if the supporting data become available.

1. **2024–2026 BRFSS extension.** The 2024 and 2025 BRFSS waves will roughly double the post-period observation count (from one post-period year to three) and may pick up slower-moving health effects. The 2024 BRFSS public-use file releases in mid-2025; opportunistic extension when available.
2. **NHIS state-level via the NCHS Research Data Center.** Restricted-access NHIS provides state identifiers and individual-level continuity-of-coverage items that BRFSS does not have. The NCHS RDC application process is multi-month and requires an approved analytic protocol; it is

the right next instrument for a coverage-specific causal design.

3. **T-MSIS Analytic Files for the 2024 reporting year, when matured.** T-MSIS provides individual-level Medicaid enrollment spells that map directly onto state intensity variation and that can test the composition reading directly by examining which enrollees were procedurally removed. T-MSIS access requires a CMS Data Use Agreement, which is a separate gating step but is feasible within an 18-month application window.

7. Conclusion

The post-pandemic Medicaid unwinding produced 25 million disenrollments over 18 months and substantial cross-state variation in procedural intensity. Public-data state-level adult-population analyses (BRFSS and ACS) do not detect coverage loss or adult-health decline under higher state procedural intensity. The signal that does appear — a slight, statistically marginal *improvement* in adult self-rated health under higher procedural intensity — is wrong-signed for the original hypothesis and is most plausibly attributed to selective removal of lower-baseline-health enrollees from the Medicaid roll, with explicit caveats about BRFSS measurement-window limitations and the pandemic-era mode shift. We report the failed primary hypothesis honestly. The right next-step data infrastructure for a credible causal coverage-loss-and-health design is the NCHS Research Data Center NHIS, the T-MSIS Analytic Files when matured for the 2024 reporting year, and the 2024–2025 BRFSS waves as they release.

Tables

Table 1. Primary BRFSS state-year DiD — cumulative procedural share, 2019–2023 panel. See `analysis/tables/main_did_summary.csv` for the full table. Continuous-intensity coefficients on the five primary outcomes are: share-fair-or-poor health -0.0205 (SE 0.0120, t = -1.70, N = 251); any coverage +0.0085 (SE 0.0153, t = 0.56, N = 251); cost-related care barrier +0.0012 (SE 0.0152, t = 0.08, N = 251); mean unhealthy mental-health days -0.6975 (SE 0.4757, t = -1.47, N = 251); mean unhealthy physical-health days -0.2403 (SE 0.3612, t = -0.67, N = 251). Binary above-median and pre-trend placebo specifications report directionally consistent nulls.

Table 2. Romano-Wolf step-down multiple-testing correction across the BRFSS five-outcome family. See `analysis/tables/romano_wolf_paper9.csv`. Conventional p-values range from 0.095 (share-fair-or-poor health) to 0.94 (cost barrier); Romano-Wolf step-down adjusted p-values range from 0.52 to 0.92. No BRFSS outcome achieves family-wise significance at the 5 percent level.

Table 3. Dose-response tercile DiD on BRFSS outcomes. See `analysis/tables/dose_response_terciles.csv`. Mid-intensity (T2) and

high-intensity (T3) coefficients versus low-intensity (T1) are reported with 95% confidence intervals. No tercile contrast achieves nominal significance.

Table 4. ACS individual-level DiD — binary high-intensity-x-post and continuous-intensity-x-post specifications. See `analysis/tables/acs_individual_results.csv`. The all-adult sample (N = 11.05 million) yields binary uninsurance +0.0002 (SE 0.0036, p = 0.95); the below-138%-FPL sample (N = 1.34 million) yields -0.0036 (p = 0.67); the below-138%-FPL-in-expansion-state sample (N = 0.90 million) yields -0.0131 (p = 0.11).

Table 5. Rambachan-Roth honest-DiD sensitivity for the ACS coverage outcomes at M $\in$ {0.5, 1.0, 2.0}. See `analysis/tables/honestdid_acs_summary.csv`. Zero falls inside the honest 95% confidence interval for every uninsurance specification at every M value across every sample.

Table 6. State-economic-controls robustness — LAUS, CPS, and joint specifications. See `analysis/tables/economic_controls_summary.csv` and `analysis/tables/economic_controls_cps_summary.csv`. Adding state unemployment and employment-to-population controls moves the BRFSS coefficients by less than 3 percent in magnitude.

Table 7. Minimum-detectable-effect calculations. See `analysis/tables/mde_and_power.csv`. BRFSS state-year MDE at 80 percent power: 0.83 pp on share-fair-or-poor health; 1.11 pp on any coverage; 0.85 pp on cost barrier; 0.26 days on unhealthy mental-health days. ACS individual-level MDE at 80 percent power: 3.45 pp on all-adult uninsurance; 6.55 pp on below-138%-FPL uninsurance (flagged as underpowered for sub-population shifts).

Table 8. Placebo outcomes — any exercise and past-year flu shot. See `analysis/tables/placebo_outcome_summary.csv`. Both placebos return tight nulls (-0.001 on any exercise, SE 0.020; -0.005 on flu shot, SE 0.026).

Figures

Figure 1. BRFSS event-study coefficients by event year, 2019–2023. Path: `analysis/figures/brfss_event_study.png` (rebuild needed — figure paths reference the prior `papers/unwinding-adult-health` location and must be regenerated against `future/unwinding-adult-health` before submission). Coefficients on (procedural-share x year) interactions with 2022 omitted, 95% confidence bands shown. None of the 2019–2021 coefficients are significantly different from zero for any outcome.

Figure 2. Intensity distribution across states, ranked low to high. Path: `analysis/figures/intensity_distribution.png` (rebuild needed). Cumulative procedural share from Maine (40.6%) to Nevada (92.0%), with median (72.5%) and tercile cutpoints marked.

Figure 3. Raw outcome trends, treated vs. comparison by tercile. Paths: `analysis/figures/raw_trends_share_fair_poor_health.png`, `raw_trends_share_any_coverage.png`, `raw_trends_share_medcost_any.png`, `raw_trends_mean_menthlth_days.png`, `raw_trends_mean_physhlth_days.png` (all rebuild needed). Mean outcomes by procedural-share tercile across 2019–2023, with 2022/2023 transition marked.

Figure 4. Honest-DiD bounds on the ACS uninsurance coefficient at M $\in$ {0.5, 1.0, 2.0}. Path: `analysis/figures/honestdid_acs.png` (rebuild needed). Honest 95% confidence intervals at each M, with the conventional pointwise interval shown for reference.