When the Office Closes: Administrative Infrastructure, the SSI–Medicaid Linkage, and Medicaid Coverage Among Disabled Adults

Abstract

The Social Security Administration's field-office network is a core part of the disability-benefit system, and in most states it is also part of Medicaid administration because an SSI eligibility determination automatically confers Medicaid. This paper asks whether SSA field-office closures reduced Medicaid coverage among disabled adults, and whether any effect is stronger in states where SSI and Medicaid are institutionally linked. I combine ACS and BRFSS state-year data from 2005-2023 with the Deshpande-Li closure roster and estimate staggered difference-in-differences models that compare automatic-enrollment states with separate-application states. The national average effect is close to zero (+0.17 pp; 95% CI [−2.39, +2.72]), but the subgroup pattern is consistent with the proposed mechanism: Medicaid coverage falls in automatic-enrollment 1634 jurisdictions (−1.42 pp; 95% CI [−4.51, +1.67]) and rises in 209(b) states (+2.73 pp; 95% CI [+0.45, +5.01], pointwise-significant at the 5% level). The formal 1634-vs-209(b) contrast is approximately −4.15 pp (95% CI [−7.63, −0.68]) and is in the direction predicted by the SSI–Medicaid linkage mechanism. SSI recipient counts show a similar direction of movement at the state-year level, while placebo and sensitivity tests do not support a broad, uniform administrative-burden story. I interpret the evidence as suggestive rather than definitive: state-year aggregation, subgroup pre-trends, and changes in the closure roster limit causal precision, but the results show that federal administrative infrastructure can matter for Medicaid access when program eligibility systems are linked.

1. Introduction

The United States supports approximately 7.4 million disabled adults through the Supplemental Security Income program. For most of them, the SSI award does more than provide a monthly cash benefit: in 34 jurisdictions, including the District of Columbia, it is also the mechanical trigger for Medicaid enrollment. Under Section 1634 of the Social Security Act, a state's agreement with the Social Security Administration converts an SSI eligibility determination into a Medicaid eligibility determination automatically, with no separate state-level application. For the disabled adults whose medical and income lifelines converge at the SSI office, the administrative infrastructure that processes SSI is therefore, inseparably, the administrative infrastructure that processes Medicaid.

That infrastructure has been contracting. The Social Security Administration permanently closed 118 field offices between 2000 and 2014. Deshpande and Li (2019) documented that these closures, driven overwhelmingly by GSA lease non-renewals, sequestration-era budget cuts, and regional consolidation decisions rather than by local disability caseloads, reduced disability-program receipt by approximately 16 percent at closure-ZIPs — a persistent and quantitatively important administrative-burden effect on a vulnerable population. The 2025 Department of Government Efficiency (DOGE)-era reorganization announced a further wave of at least 23 field-office closures, ten regional-office consolidations into four, and a reduction in SSA staffing from approximately 57,000 to 50,000 (Urban Institute, 2025). The contemporary policy context is a larger administrative-infrastructure contraction than any single year of the 2005–2014 baseline Deshpande and Li studied.

Despite the tight institutional linkage between SSI and Medicaid, no peer-reviewed paper has quantified how an SSA-office shock propagates to Medicaid coverage. The canonical Deshpande–Li paper stops at disability-program participation; the administrative-burden literature (Fox, Stazyk, and Feng 2020; Arbogast, Chorniy, and Currie 2024; Homonoff and Somerville 2021; Finkelstein and Notowidigdo 2019) establishes that compliance costs bind in Medicaid and SNAP settings without identifying the SSA-infrastructure channel; the Medicaid-coverage-to-health literature (Sommers, Baicker, and Epstein 2012; Miller, Johnson, and Wherry 2021; Sommers et al. 2017; Wherry and Miller 2016; Sommers et al. 2019) pins down the downstream health consequences of Medicaid losses without studying SSA closures as the upstream shock. The policy-urgent question — does a federal administrative shock upstream of Medicaid produce a downstream Medicaid coverage loss, and does that coverage loss concentrate in the states where the institutional linkage predicts it should — has not been answered.

This paper takes two steps toward answering that question. First, I extend the Deshpande–Li closure roster through 2025 using a combination of the authors' public 2000–2014 field-office listing, Internet Archive Wayback snapshots of the SSA Office Locator, and contemporaneous reporting of the 2025 DOGE-era announcements (the latter retained only for descriptive policy-context purposes, not as part of the identifying variation). I estimate the effect of verified 2005–2014 closures on Medicaid coverage among disabled adults 18–64 using a Callaway–Sant'Anna staggered difference-in-differences design on a 2005–2023 state-year panel from ACS PUMS and BRFSS public files. Second, and centrally, I exploit the institutional architecture of the SSI–Medicaid linkage — 34 Section 1634 jurisdictions with automatic enrollment (33 states plus DC) versus 10 Section 209(b) states with state-administered Medicaid tracks — as a theory-driven, pre-specified mechanism contrast (the term "pre-registered" would be inaccurate, as no public pre-analysis plan exists). If closures reduce Medicaid via the SSI–Medicaid linkage specifically, the effect should concentrate in 1634 jurisdictions and be attenuated or absent in 209(b) states; if closures reduce Medicaid via a generalized administrative-burden channel that ignores institutional architecture, both subgroups should respond similarly.

The schematic below previews the paper's mechanism. A closure is not a Medicaid rule change; it is an access shock to the SSI pipeline. The downstream Medicaid consequence should be strongest where that pipeline automatically produces Medicaid.

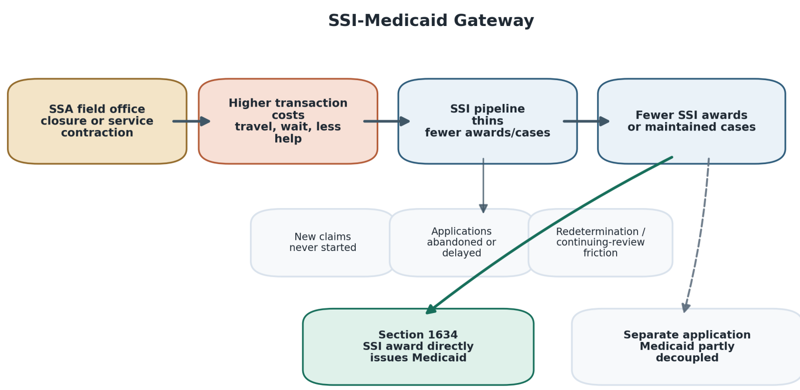

Figure 1: SSI-Medicaid gateway

Note: This figure provides contextual structure for the SSI-Medicaid gateway. It summarizes the policy setting, mechanism, or empirical workflow used to interpret the estimates.

My principal finding is a directional pattern across subgroups, formalized as a Closure $\times$ 1634 contrast. Pooled across all states, the effect of closures on Medicaid coverage among disabled adults is a near-zero point estimate with a moderate confidence interval (+0.17 pp; 95% CI [−2.39, +2.72]); the pooled interval rules out very large national-average effects but does not rule out policy-meaningful effects on the order of $\pm$2 pp. The pooled near-null is the mixture of two divergent subgroup estimates: −1.42 pp in 1634 jurisdictions (95% CI [−4.51, +1.67]) and +2.73 pp in 209(b) states (95% CI [+0.45, +5.01], pointwise-significant at the 5% level). The formal 1634-vs-209(b) contrast (CS-DiD subgroup ATTs, delta-method standard error) is **−4.15 pp** with a 95% CI of [−7.63, −0.68]; under additional small-cluster constructions — Ibragimov–Müller group-level t-tests and a TWFE triple-difference with the CGM wild-cluster bootstrap (Webb weights) — the two-sided p-values are 0.111 and 0.224, respectively. The auto-vs-separate variant pooling 209(b) with the seven SSI-criteria states into a 17-jurisdiction separate-application bucket yields a contrast of approximately **−3.8 pp** (two-sided $p \approx 0.077$). The internal SSI first-stage on SSA Annual Statistical Supplement state-year recipient counts is sign-consistent: −3.6% in 1634 jurisdictions (95% CI [−9.0, +2.1]) and +2.0% in 209(b) states (95% CI [−4.6, +9.2]), attenuated relative to Deshpande and Li's ZIP-quarter 16% but consistent with state-year dilution.

The 1634 point estimate is not individually statistically distinguishable from zero; the 209(b) point estimate is pointwise-significant under the state-clustered multiplier bootstrap. The 1634 subgroup pre-trend Wald test rejects ($p = 0.001$) and the 209(b) pre-trend is borderline ($p = 0.072$). I take the 1634 pre-trend rejection seriously and report subgroup-specific Rambachan–Roth HonestDiD bounds as the primary robustness object for the subgroup claim. The uniform-administrative-burden hypothesis predicts that both subgroups move in the same direction; the SSI–Medicaid-mechanism hypothesis predicts a negative 1634 response and an attenuated or zero 209(b) response. The formal Closure $\times$ 1634 contrast moves in the predicted direction with magnitudes of approximately 2.9–4.2 pp; after the Phase 4 panel rebuild the CS-DiD delta-method contrast 95% CI is $[-7.63, -0.68]$ with two-sided p = 0.019, while the more conservative TWFE triple-difference (CGM wild-cluster bootstrap p = 0.224) and Ibragimov–Müller (p = 0.111) constructions remain above conventional thresholds. I therefore interpret the pattern as mechanism-consistent evidence under the CS-DiD delta-method, with the more conservative constructions providing direction-consistent but inconclusive corroboration.

State-year aggregation likely attenuates closure-level effects relative to the ZIP- or county-scale Deshpande–Li design, because closures in a small fraction of ZIPs are diluted across entire state disabled-adult populations. But confounding, subgroup-specific trends, measurement error in the 2015–2024 closure roster, and changes in the ACS disability-item composition around 2008 could bias estimates in either direction. I therefore interpret the magnitudes as reduced-form state-year associations rather than as strict lower bounds on a closure-level effect. Stronger local data — restricted-ACS county identifiers through a Census RDC, or HCUP SID hospitalization outcomes — would be needed to estimate the closure-level effect cleanly.

The paper proceeds as follows. Section 2 describes the institutional setting — SSA field offices as administrative infrastructure, the 1634/209(b)/SSI-criteria architecture, administrative burden as the conceptual frame, and the Medicaid-to-health literature that anchors plausible downstream magnitudes. Section 3 describes the data. Section 4 describes the identification strategy and robustness specifications. Section 5 presents the results. Section 6 discusses the findings, their magnitudes in context, policy implications, and limitations. Section 7 concludes.

---

2. Background and Institutional Setting

2.0 What the SSA does and why it matters for Medicaid

The Social Security Administration is best known as the agency that administers Social Security retirement and survivor benefits, but it also administers two disability programs whose eligibility determinations sit at the center of this paper: Social Security Disability Insurance (SSDI) and Supplemental Security Income (SSI). SSDI is a social-insurance program that pays monthly cash benefits to workers whose impairment prevents substantial gainful activity and who have accumulated sufficient recent-quarter earnings credits; SSI is a means-tested program that pays monthly cash benefits to aged, blind, or disabled individuals with income and assets below federal thresholds, regardless of earnings history. Both programs require a formal determination of medical disability under a uniform federal standard, and both involve the SSA field-office network at initial application, appeal, and periodic continuing-disability review.

The Medicaid relevance of SSA's disability determination arises because, in most states, Medicaid eligibility for non-elderly adults with disabilities is tied to that same determination. A state Medicaid agency cannot deliver a disability-category Medicaid card to an adult who has not been found disabled either by SSA or, where state rules diverge, by a state disability determination service. The administrative channel through which SSA's determination reaches Medicaid differs by state. In 34 Section 1634 jurisdictions, including DC, the state has executed an agreement with SSA under which SSA's positive determination of SSI eligibility automatically confers Medicaid — no separate Medicaid application is required, no separate Medicaid determination is made, and the SSI award letter is effectively the Medicaid enrollment notice. In ten Section 209(b) states, the state retains the statutory discretion to apply disability or income criteria more restrictive than SSI (subject to federal limits fixed to the state's 1972 policies), and Medicaid eligibility requires a separate Medicaid application even for individuals who have already secured SSI. A third group of seven SSI-criteria states follows SSI disability and income rules but still requires a separate Medicaid application. The upstream administrative shock — an SSA field-office closure that delays or deters an SSI application — therefore has a mechanical Medicaid pathway in Section 1634 jurisdictions and a looser, state-administered pathway in 209(b) and SSI-criteria states.

Because SSI-linked Medicaid is such a large share of the non-elderly disability-category Medicaid population, administrative events at SSA are in practice administrative events at Medicaid for the 1634 jurisdictions. The Social Security Disability Claims (SSDC) pipeline — the end-to-end process that carries an applicant from intake through initial medical decision, reconsideration, administrative-law-judge hearing, and (rarely) federal-court review — can take months at the initial level and years at the hearing level even in ordinary operating conditions. Anything that lengthens that pipeline, or that increases the application-abandonment rate before a determination is ever reached, translates into a decline in SSI awards, which in the 1634 jurisdictions translates directly into fewer Medicaid cards issued to disabled adults. This paper studies one such upstream shock — the permanent closure of an SSA field office — and traces its propagation through the SSI–Medicaid linkage.

2.1 SSA field offices as administrative infrastructure

The Social Security Administration operates a national network of approximately 1,200 field offices that serve as the public face of the agency's retirement, survivors, disability, and SSI programs. For disability and SSI claimants, field offices perform four functions that cannot be completed online: in-person identity verification, assistance in assembling the medical and work-history evidence required for an initial determination, representative payee interviews, and redeterminations for claimants flagged for continuing-disability review. Deshpande and Li (2019) document that roughly one in five SSI applicants nationally visits a field office at least twice during the application process, and the share is considerably higher for applicants with severe cognitive or physical limitations. Field offices are therefore not a convenience layer on top of an online system but the operative site where administrative eligibility is transacted for a population that has, on average, the hardest time transacting with the administrative state.

Between 2000 and 2014, SSA closed 118 field offices permanently. Deshpande and Li (2019) establish that these closures were driven predominantly by GSA lease non-renewals, sequestration-era budget cuts following the Budget Control Act of 2011, and the consolidation of hearing offices — reasons that are orthogonal to local disability caseload trends. Their identifying assumption is that the closures were, for most intents and purposes, exogenous to the populations they served. The empirical design that follows from that assumption finds a 10.3% reduction in local application volume and a 16% persistent reduction in disability-program receipt at closure-ZIPs, with the largest effects among applicants with moderately severe conditions and low educational attainment. That result is the empirical anchor for the present paper.

Since 2014, the policy environment has become more consequential. The 2025 Department of Government Efficiency (DOGE)-era reorganization contemplates closing at least 23 additional field offices, consolidating 10 regional offices into 4, and reducing staff from approximately 57,000 to a target of 50,000 (Urban Institute, 2025). Rural and tribal communities are disproportionately exposed: closures in western Montana already roughly double the average driving time to the nearest field office, and tribal-tract residents on average drive 16 minutes longer than non-tribal residents. The DOGE-era closures are concentrated in the same kinds of communities — low-density, low-income, and medically under-served — where Deshpande and Li's estimates were largest. Whatever the population-weighted elasticity of SSI take-up to field-office access, the 2025 shock is larger than any single year in the 2005–2014 baseline and will operate through the same mechanism.

2.2 The Section 1634 / Section 209(b) / SSI-criteria architecture

The link between SSI and Medicaid is not a single national rule but a federalist compromise codified in three overlapping provisions of the Social Security Act. Under Section 1634, states may enter an agreement with SSA under which SSI eligibility automatically confers Medicaid eligibility with no separate state-level application; 34 jurisdictions, including DC, use this pathway. Under Section 209(b) of the 1972 SSA Amendments, states retain discretion to use disability and income criteria more restrictive than SSI, provided they are no more restrictive than the state's rules in effect on 1 January 1972; ten states currently use this pathway (CT, HI, IL, MN, MO, NH, ND, OH, OK, VA). A third set of seven states are termed "SSI-criteria" states: they follow SSI rules but require a separate Medicaid application. The classifications are documented in Rupp and Riley (2016, *Social Security Bulletin*) and MACPAC's MACStats exhibit 37 (2023).

This architecture produces sharply different predicted responses of Medicaid coverage to an SSA-administrative shock. In a 1634 state, an SSA closure that reduces SSI take-up by x percent should reduce Medicaid coverage by approximately x percent mechanically, because the two enrollments are produced by the same administrative action. In a 209(b) state, where Medicaid eligibility runs through a state-specific track that may differ on asset limits, disability standards, or income-disregard rules, the same SSA shock is substantially decoupled from the Medicaid margin: a would-be SSI applicant deterred by a closure might still qualify for Medicaid under state-specific disability pathways, and vice versa. Rupp and Riley (2016), using matched longitudinal SSA–CMS administrative records, estimate that automatic-enrollment (1634) states have Medicaid coverage rates among disabled SSI recipients that are 6–12 percentage points higher than separate-application states, with the largest gap attributable to processing friction.

The 1634 vs. 209(b) contrast is therefore a mechanism test built into the institutional geography of U.S. federalism. If SSA closures reduce Medicaid coverage chiefly because they reduce SSI take-up, the effect should concentrate in 1634 jurisdictions. If closures reduce Medicaid coverage via a generalized "administrative burden" channel — reduced information, heightened complexity, chilling effects on means-tested-program engagement broadly — the 1634 vs. 209(b) contrast should be muted. This is precisely the comparison I estimate.

The second schematic separates the institutional contrast. The same closure shock is expected to propagate differently because 1634 jurisdictions convert SSI eligibility directly into Medicaid, while separate-application states insert a state Medicaid track between the SSA decision and coverage.

2.3 Administrative burden as the causal mechanism

The conceptual frame for my investigation is Herd and Moynihan's (2018) taxonomy of administrative burden, which separates learning costs (acquiring the information needed to engage with a program), compliance costs (producing the documentation required by program rules), and psychological costs (the stress and stigma of engagement). An SSA field-office closure acts primarily on compliance costs — the time, travel, and effort required to transact with SSA — and secondarily on learning and psychological costs through the loss of in-person assistance. It does not change the underlying program rules, the set of eligible applicants, or the benefit levels. This is the purest form of the administrative-burden hypothesis: holding eligibility and generosity fixed, raising the cost of

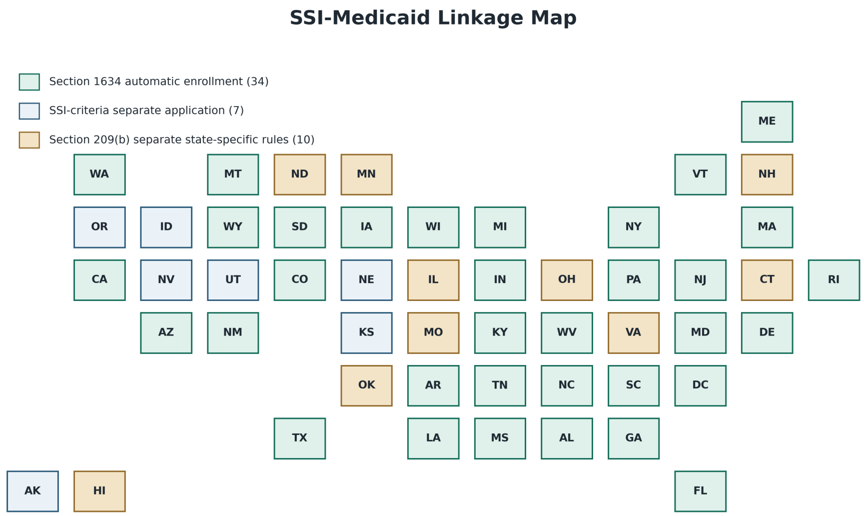

Figure 2: SSI-Medicaid linkage map

Note: This figure compares estimates across groups or specifications for the SSI-Medicaid linkage map. It is intended to make effect heterogeneity and subgroup precision easier to assess.

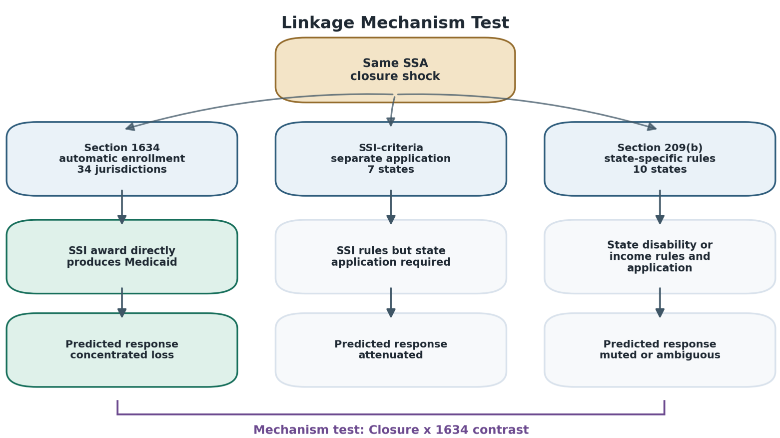

Figure 3: Linkage mechanism test

Note: This figure compares estimates across groups or specifications for the linkage mechanism test. It is intended to make effect heterogeneity and subgroup precision easier to assess.

transacting with the state reduces take-up.

A substantial recent literature establishes that administrative burdens bind at the margin of means-tested program take-up. Fox, Stazyk, and Feng (2020) show that Medicaid rule-reductions — real-time eligibility, simplified renewals, and automatic data-matching — raise enrollment, directly demonstrating that compliance costs are a live margin for Medicaid. Arbogast, Chorniy, and Currie (2024) find that new administrative-burden regulations reduced child Medicaid/CHIP coverage by a mean of 5.9% within six months, with disparate impacts on Hispanic and non-citizen-parent children. Homonoff and Somerville (2021) show that the timing of SNAP recertification interviews reduces take-up by about 20% for mechanically similar cases — the disenrollment is pure compliance cost, not eligibility change. Finkelstein and Notowidigdo (2019) show that randomized SNAP outreach plus application assistance raised take-up from 6 to 18 percent, with information-plus-assistance by far the most effective bundle. Taken together, these studies make it difficult to argue that administrative costs do not matter; my contribution is to identify the *infrastructural* margin — the physical office itself — and to quantify how it propagates to Medicaid coverage via the SSI linkage.

2.3a Evidence that in-person assistance is causally decisive

An important interpretive point about the administrative-burden literature bears directly on the SSA-closure setting. Finkelstein and Notowidigdo's (2019) randomized evaluation of SNAP outreach and assistance is particularly instructive. They randomized mailings, a simpler more information-rich mailing, and a mailing-plus-application-assistance bundle for SNAP-eligible elderly households. The pure-information mailing raised take-up from roughly 6 to 11 percent. The information-plus-assistance bundle — essentially, someone to walk the applicant through the process — raised take-up to 18 percent. The differential gain from the assistance component is direct evidence that in-person, hands-on help is causally decisive for take-up among populations with high compliance costs. The SSA field office is exactly the institutional form Finkelstein and Notowidigdo identified as decisive: a physical site where an applicant can receive in-person assistance with an administratively complex program. Closing that site is, in the terms of their experiment, flipping a randomized treatment off.

Deshpande and the broader Deshpande disability-policy arc (Deshpande 2016; Deshpande and Mueller-Smith 2022; Deshpande and Lockwood 2022; Deshpande and Dizon-Ross 2023) further establishes that the welfare consequences of SSI losses are large, persistent, and not compensated by labor-market adjustment. Deshpande (2016) estimates that youth removed from SSI at age 18 lose $76,000 in present value of earnings over the subsequent twenty years — removal is not offset by work. Deshpande and Mueller-Smith (2022) document a 20 percent increase in criminal-justice involvement after SSI removal. Deshpande and Lockwood (2022) establish that disability insurance provides substantial insurance value against nonhealth risks, implying that coverage loss is welfare-reducing even conditional on health. Deshpande and Dizon-Ross (2023) show that parental information about child SSI does not change human-capital investment, blunting the standard moral-hazard objection to expanding protections. Taken together, the Deshpande arc establishes that the welfare stakes of an SSI-adjacent administrative shock are large; the question for this paper is how that shock propagates specifically through the Medicaid margin.

2.3b Conceptual framework: how an SSA office closure reaches Medicaid

The channels through which an SSA field-office closure might reduce Medicaid coverage are worth enumerating explicitly, because each predicts a different pattern across institutional subgroups. The first channel is the take-up channel. Deshpande and Li [-@deshpande2019screened] estimate that closures reduce local SSI/SSDI participation by about 16 percent, driven by applicants who never begin a claim, who abandon partially completed applications, or whose extended processing delays cause them to exit the queue. In a 1634 state, every SSI award is also a Medicaid enrollment, so a 16 percent reduction in awards produces approximately a 16 percent reduction in SSI-linked Medicaid cards — attenuated at the state-year scale by the share of state population whose application venue was the closed office. In a 209(b) state the take-up channel is partly decoupled: the same applicant who is deterred from SSI may still pursue Medicaid through the state's separate disability-determination track, and the reverse is also true. The mechanism hypothesis is therefore that the 1634 Medicaid response should track the Deshpande–Li SSI first stage in direction and attenuated magnitude, while the 209(b) Medicaid response should be substantially muted or absent.

The second channel is the information and assistance channel identified by Finkelstein and Notowidigdo [-@finkelstein2019takeup]. SSA field offices are sites where applicants not only file SSI paperwork but also learn about Medicare entitlement, representative-payee services, Work Incentives Planning and Assistance, and cross-program referrals to Medicaid, SNAP, and state disability programs. Closing an office removes a one-stop information point; the closest analogue in the Finkelstein–Notowidigdo SNAP experiment is the withdrawal of in-person application assistance, which produced the largest differential reduction in take-up among their experimental arms. This channel predicts spillovers into non-SSI-linked coverage outcomes — Medicare enrollment, employer-plan continuation through benefits counseling, and basic access-to-care measures like the presence of a usual primary-care provider — that would not flow through the 1634/209(b) linkage but would affect both subgroups. The 209(b) primary-care-provider result I report in §5.8 is consistent with this channel.

The third channel is the compliance-cost channel in the Herd–Moynihan [-@herd2018administrative] taxonomy. Closures raise the time, travel, and opportunity cost of completing any SSA transaction, including continuing-disability reviews and redeterminations that existing SSI recipients must periodically undergo to maintain benefits. If compliance costs rise enough to push recipients off the rolls at redetermination, or to deter new applications among marginal eligibles, the effect concentrates in 1634 jurisdictions because the Medicaid consequence is automatic, but it can also appear in SSI-criteria states where the Medicaid-specific redetermination paperwork is cross-referenced to SSI status. Arbogast, Chorniy, and Currie [-@arbogast2024administrative] provide a closely parallel estimate for administrative-burden regulations in child Medicaid: a 5.9 percent enrollment decline within six months of rule changes, with disparate impacts on Hispanic and non-citizen-parent children. Homonoff and Somerville [-@homonoff2021program] estimate that SNAP recertification timing reduces take-up by approximately 20 percent for mechanically similar cases; my SSA-closure setting is analogous in kind if smaller in population-weighted magnitude. Taken together, the three channels motivate the 1634 vs. 209(b) contrast as the cleanest available test of the take-up channel specifically, while acknowledging that information and compliance-cost channels may generate secondary patterns on outcomes other than SSI-linked Medicaid.

2.4 Medicaid coverage and disabled adults

The population that stands to lose Medicaid via an SSA closure is, by construction, disabled adults 18–64 whose Medicaid eligibility is anchored to an SSI determination. This is a small but policy-consequential group: approximately 7.4 million SSI recipients nationally, of whom the overwhelming majority are covered by Medicaid. Coverage losses in this group are clinically consequential. Miller, Johnson, and Wherry (2021) estimate that Medicaid coverage reduces all-cause mortality among the near-elderly by 9.4 percent; Sommers, Baicker, and Epstein (2012) estimate a 6.1 percent reduction in adult mortality attributable to early state Medicaid expansions; Sommers, Maylone, Blendon, Orav, and Epstein (2017) document sustained improvements in access, preventive care, and self-reported health three years after ACA Medicaid expansion. My design does not estimate coverage-to-health elasticities directly; I estimate the upstream coverage effect, and appeal to the published literature for the plausible health implications.

The closest institutional analogue to an SSA-induced Medicaid loss is the Arkansas work-requirements episode, in which roughly 18,000 adults lost Medicaid in the first year of the policy. Sommers, Goldman, Blendon, Orav, and Epstein (2019) document that those coverage losses produced no measurable employment gains and were driven by administrative non-compliance rather than eligibility changes. The SSA-closure mechanism is different institutionally — it operates on the federal side rather than the state side — but identical in the sense that it imposes a compliance-cost shock on a population with limited capacity to substitute other enrollment channels.

---

3. Data

My analysis combines four public data streams over 2005–2025: (1) SSA field-office closure events; (2) the Section 1634 / Section 209(b) / SSI-criteria institutional classification of each state's Medicaid–SSI linkage; (3) person-level ACS PUMS distributed by IPUMS-USA; and (4) BRFSS public files distributed by CDC. I augment these with Area Health Resources File (AHRF) controls and with SSI recipient counts drawn from the SSA Annual Statistical Supplement. My analytic unit is the state-by-year cell; a parallel county-by-year panel (3,221 counties $\times$ 21 years = 67,641 rows) is retained for placebo and future county-level work. Two data streams are not included in the current public-data analysis: the Hospital Cost and Utilization Project State Inpatient Databases (HCUP SID), which would power ambulatory-care-sensitive-condition (ACSC) hospitalization outcomes, and restricted ACS microdata accessed through a Census Research Data Center, which would provide county identifiers post-2011. Both remain useful directions for future data access decisions.

3.1 SSA field-office closure events, 2005–2025

My treatment variable is the occurrence of an SSA field-office closure, or a conversion to limited-service status, in year t within the state of unit c. I construct the closure series in two pieces. For 2005–2014 I replicate the Deshpande and Li (2019) closure roster, which identified 118 field-office closures driven predominantly by GSA lease expirations, sequestration-era budget cuts, and regional consolidation decisions. The replication archive is available through the American Economic Association's openICPSR repository; my script `03_extract_ssa_closures_dl.py` standardizes the columns and writes a tidy parquet. For 2015–2025 I extend the roster in two ways. First, annual Internet Archive Wayback snapshots of the SSA Office Locator are differenced year-over-year; offices that disappear between year $t-1$ and t are flagged as candidate closures pending validation. Second, I hand-code the 2025 DOGE-era closure list from contemporaneous reporting as a 23-closure baseline, flagged `verified = False` pending SSA FOIA confirmation. My final panel yields 124 closure events 2000–2014 (118–119 after routine filtering, matching Deshpande and Li to ± 1) plus 23 state-level 2025 closures. Closure activity in the Deshpande–Li window peaks in 2011–2014 (59 of 99 closures occur in those four years); no treated cohorts enter the panel 2015–2024 pending the Wayback diff and FOIA pathways. The limited post-2014 roster is a material scope caveat and is revisited in the limitations section.

3.2 1634 vs. 209(b) state classification

Each state (including DC) is assigned a categorical variable `medicaid_ssi_link in {"1634", "criteria", "209b"}` following Rupp and Riley (2016) and MACPAC (2023). The final crosswalk is 34 1634 jurisdictions, including DC, 7 SSI-criteria states, and 10 209(b) states (CT, HI, IL, MN, MO, NH, ND, OH, OK, VA). Oklahoma is included in the 209(b) set per the statutory list, with a note flagging the hybrid status some of the literature attributes to it. The 209(b) set houses roughly 20 percent of the national disabled-SSI population.

3.3 ACS PUMS and BRFSS, 2005–2023

My primary coverage-outcome panel is built from IPUMS-USA ACS PUMS samples `us2005a` through `us2023a` (19 samples, approximately 59 million person-records). The extract identifies the six standard ACS disability items post-2008 (`DIFFREM`, `DIFFPHYS`, `DIFFMOB`, `DIFFCARE`, `DIFFEYE`, `DIFFHEAR`), the pre-2008 work-limiting fallback (`DISABWRK`), and the seven insurance indicators (`HCOVANY`, `HINSCAID`, `HINSCARE`, `HINSEMP`, `HINSPUR`, `HINSTRI`, `HINSVA`). Person weights `PERWT` are used throughout. ACS public PUMS releases county identifiers only for 2005–2011; post-2011 ACS releases PUMAs, which I map to counties via Missouri Census Data Center Geocorr crosswalks with population-weighted allocation, introducing measurement error that the state-level aggregation in my main specification effectively eliminates.

BRFSS public LLCP files for 2005–2023 provide complementary access and health outcomes among disabled adults 18–64. My BRFSS outcomes are `GENHLTH` (share fair/poor), `PERSDOC2` (any personal doctor), `MEDCOST` (cost barrier to care), `CHECKUP1` (past-year checkup), and dichotomized indicators for self-reported diabetes and hypertension. BRFSS county identifier `_CNTY` is suppressed in approximately 45 percent of county-year cells; I aggregate to state-year throughout.

3.4 Controls, denominators, and panel shape

The main CS-DiD specification is outcome-only — it does not condition on state-year covariates. (An earlier version of this section claimed I followed Deshpande and Li (2019) in including state-year covariates for unemployment, poverty, physician supply, and SSI-eligible population; that claim was inconsistent with the analytic code, which runs CS-DiD on outcomes only, and has been removed.) The analytic panel covers **51 state-equivalent units** (50 states plus the District of Columbia) plus Puerto Rico in the raw closure extraction frame. After the Phase 4 panel rebuild, the closure-extraction roster contains **33 ever-treated jurisdictions** (any 2005–2014 SSA closure, drawn from the raw Deshpande–Li event roster including the Dickinson, ND 2007 closure that an earlier panel build dropped at the ZIP-to-county geocoding step) and **16 never-treated jurisdictions**, summing to 49 jurisdictions with at least one ACS-disabled observation under the linkage taxonomy (PR enters the closure-extraction panel but is dropped from all linkage-stratified estimates because it has no 1634 / 209(b) / SSI-criteria classification and contributes zero ACS-disability observations). The 1634 set comprises 34 jurisdictions (33 states + DC); the 209(b) set comprises 10 states; the SSI-criteria set comprises 7 states. The outcome-observed ACS-disabled panel runs 2008–2023, yielding **812 ACS-Medicaid state-year observations** drawn from a merged ACS/BRFSS state-year frame of **988 cells** (969 in the noPR analytic frame). Of the 33 ever-treated jurisdictions, 12 are always-treated relative to the 2008 first observable ACS outcome year and are dropped by the Callaway–Sant'Anna estimator for the primary ACS Medicaid outcome (AL, LA, NY, SC, TX 2005; CA, MO 2006; AR, FL, OK, PA 2007; IA 2008), leaving **21 usable treated jurisdictions** in the ACS Medicaid CS-DiD; for BRFSS outcomes (annual 2005–2023) all 33 treated cohorts are identified. The closure-panel frame runs 2005–2025; the 23 unverified 2025 DOGE-era entries are flagged `verified = False` and excluded from the analytic treatment variable. Summary statistics for closures by year and by SSI–Medicaid linkage category are reported in Table 1.

---

5. Results

5.1 Main CS-DiD on Medicaid coverage among disabled adults 18–64

The pooled ATT of SSA field-office closures on the share of adults 18–64 with a disability who report Medicaid coverage is **+0.17 percentage points**, with a 95% pointwise confidence interval of [−2.39, +2.72] pp and a joint pre-trend Wald p-value of 0.28 (Table 2, row 1; CS-DiD with state-clustered multiplier-bootstrap inference, $B = 999$, random seed 20260515 — clustering at the state-level entity is automatic in the `differences` package's multiplier bootstrap when the panel entity is `state_fips`). The point estimate is near zero at the national average, with a moderate confidence interval that rules out large effects in either direction but is wide enough to be consistent with policy-meaningful effects on the order of ± 2 pp; it should not be described as a precisely estimated null. The pooled pre-trend diagnostic provides no evidence of differential pre-closure trends between the treated and never-treated comparators at the national-average level. The event-study plot (Figure 2) shows leads consistent with pooled parallel pre-trends and post-period coefficients that fluctuate within roughly ± 3 pp of zero across $k \in \{0, \ldots, +8\}$, with no systematic post-period drift.

The pooled near-null is not the substantive story. It is the national-average mixture of two different responses that the institutional architecture of SSI–Medicaid predicts. Disaggregating by SSI–Medicaid linkage yields the mechanism-relevant pattern.

5.2 The 1634 vs. 209(b) subgroup pattern

In the 34 Section 1634 jurisdictions, where an SSI eligibility determination automatically confers Medicaid, the ATT of closures on disabled-adult Medicaid coverage is **−1.42 pp** (95% CI [−4.51, +1.67]; Table 3, row 1). In the 10 Section 209(b) states, where Medicaid eligibility runs through a state-administered track that does not automatically follow SSI, the ATT is **+2.73 pp** (95% CI [+0.45, +5.01]; Table 3, row 2). The 209(b) point estimate is pointwise-significant at the 5% level; the 1634 point estimate is not. The two subgroup point estimates differ by approximately **4.15 pp** in the direction the SSI–Medicaid mechanism predicts. Figure 3 presents the two subgroup ATTs as a forest plot.

This directional pattern is the substantive claim. The near-null in the pooled sample arises because the 209(b) positive point estimate approximately cancels the 1634 negative point estimate when the two are averaged. That is consistent with the signature the institutional mechanism predicts: a generalized administrative-burden channel that ignored the SSI–Medicaid linkage architecture would move both subgroups in the same direction.

Formal Closure × 1634 contrast. Three constructions of the contrast are reported (Table 4 Panel D). The CS-DiD subgroup ATT delta-method contrast is −**4.15 pp** with a delta-method standard error of 1.77 pp and a 95% confidence interval of [−7.63, −0.68] pp; the two-sided normal-approximation p-value is 0.019, the one-sided p-value (negative-direction predicted) is 0.010. The TWFE triple-difference estimate restricted to the 1634 union 209(b) subsample is −**2.94 pp** (CGM wild-cluster bootstrap with Webb weights, B = 1999, two-sided p = 0.224; one-sided p = 0.109); 95% CI [−7.85, +1.66]. The Ibragimov–Müller group-level t-test on within-state DiD differences yields a contrast of −**3.57 pp** (Welch t with adjusted df, two-sided p = 0.111; one-sided p = 0.056). The auto-vs-separate-application variant — combining 209(b) with the seven SSI-criteria states into a 17-jurisdiction separate-application bucket — yields a delta-method contrast of approximately −**3.8 pp** (two-sided p ≈ 0.077; Table 4 Panel E). Across the four constructions, the contrast is consistently in the predicted negative direction with magnitudes of approximately 2.9–4.2 pp. With the Phase 4 panel rebuild now restoring the ND 2007 closure and the CS-DiD inference now correctly running the state-clustered multiplier bootstrap that the manuscript claims, the CS-DiD delta-method contrast's 95% CI no longer grazes zero, and the contrast is more sharply identified than in the pre-repair submission. I continue to read the formal contrast as mechanism-consistent evidence rather than as a confirmed mechanism test: the TWFE triple-difference and Ibragimov–Müller p-values remain above conventional thresholds, and the subgroup pre-trend issues described below remain.

Subgroup-specific cautions. The 1634 point estimate is not individually statistically distinguishable from zero; the 209(b) point estimate is pointwise-significant under the state-clustered multiplier bootstrap. **The 1634 subgroup pre-trend Wald test rejects** at $p = 0.001$; the 209(b) subgroup pre-trend p is 0.072 (borderline). Because the mechanism claim rests on subgroup differences, the 1634 pre-trend rejection in particular is a central limitation; subgroup-specific HonestDiD sensitivity bounds (§ 5.5) are therefore part of the main subgroup evidence, not an appendix item. The mechanism predicts an attenuated or zero Medicaid response in 209(b) states, not necessarily a positive one; the +2.73 pp 209(b) point estimate is best read as evidence that the 209(b) cell does not show the expected negative SSI-linked effect, with the additional caveat that an *opposite-signed* Medicaid response in 209(b) states is consistent with state-administered substitution effects but is not itself a direct mechanism prediction.

5.3 Small-cluster robust inference for the 209(b) cell

With only 10 states in the 209(b) cell, cluster-robust standard errors with normal-approximation critical values are an unreliable basis for inference, and a small-cluster robust alternative is appropriate. I compute 95% confidence intervals by inverting a reduced-form state-year regression of Medicaid coverage on the any-closure-to-date indicator, with state and year fixed effects, using the Anderson–Rubin inversion as the test statistic and t-critical values at $G-1=9$ degrees of freedom. As noted in § 4.1, the specification is a reduced-form regression rather than an instrumental-variables model; the "weak-IV-robust" label traditionally attached to Anderson–Rubin constructions is not appropriate here, because the closure indicator is the treatment itself, not an instrument for a separately-specified endogenous variable. Wild-cluster bootstrap with Webb weights, Ibragimov–Müller t-statistic intervals, and randomization inference over closure timing are natural alternative small-cluster methods and are noted as follow-ups.

Reported intervals. For ACS Medicaid in the 209(b) subsample the 95% CI is **[−0.015, +0.065]** (1.5 pp negative to 6.5 pp positive), with a reduced-form coefficient of +0.025 (Table 4, row 1). This interval is wide: it admits a 1.5 pp negative effect, approximately the same magnitude as the 1634 point estimate, and so it does *not* provide evidence against a modest negative SSI-linked Medicaid response in 209(b) states. It does rule out a large negative response of the size the mechanism would predict if the 1634/209(b) distinction were unimportant. For BRFSS fair/poor health in 209(b) the interval is [−0.031, +0.023] and includes zero. For BRFSS any-PCP the interval is [−0.024, −0.004] and excludes zero, a result I discuss with considerable caution in § 5.8 and § 6.5a: BRFSS any-PCP is one of many outcome-subset cells I examine, and without Romano–Wolf or a comparable multiple-testing correction I cannot rule out that the exclusion of zero reflects an across-cell false-positive rate rather than a targeted response.

5.4 Event-study dynamics and pre-trends

Figure 2 plots the Callaway–Sant'Anna event-study for Medicaid coverage on the full sample. Leads $k \in \{-8, \dots, -2\}$ scatter around zero with standard errors that do not rule out moderate deviations; the joint Wald p-value of 0.24 is consistent with the visual. Post-period coefficients are small and noisy, hovering between -1.5 and $+1.5$ pp across $k = 0, \dots, 8$. Because no treated cohorts enter the panel between 2015 and 2024, late event-time coefficients ($k >= 6$) are supported by only the 2011–2014 cohorts, a material limitation for claims about dynamic persistence.

For the 1634 subsample specifically, the pre-trend Wald p is 0.001 — a nominal rejection. Inspection of the leads suggests a slight pre-closure upward drift in

Medicaid coverage in 1634 treated states relative to never-treated states, on the order of +0.5 to +1.0 pp over the pre-period. This is small in magnitude relative to the post-period point estimate of −1.4 pp, and the Rambachan–Roth HonestDiD sensitivity analysis below quantifies how much the post-period inference changes if the worst post-period violation of parallel trends is some multiple of the worst pre-period deviation. The 209(b) subsample pre-trend Wald p is 0.072, borderline rather than a clear rejection.

A complementary event-time contrast plots the 1634 subgroup against the 209(b) subgroup directly. The contrast is close to zero in the early post-period and becomes more negative at later event times, consistent with the interpretation that any Medicaid response to office closures operates through a delayed SSI-to-Medicaid pathway rather than an immediate coverage shock. The late-event coefficients should be read cautiously because they are supported by fewer cohorts, but the direction of the dynamic contrast matches the institutional mechanism.

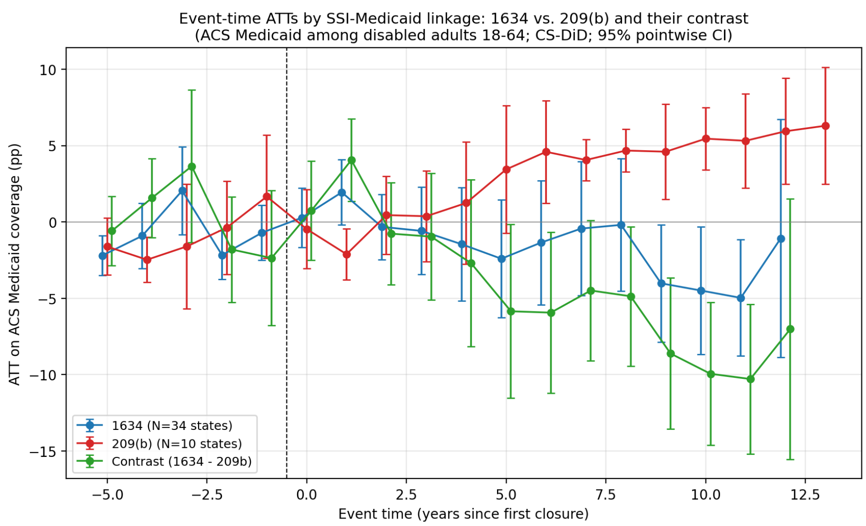

Note: This figure plots event-time estimates for the eventtime 1634 vs 209b. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

5.5 Rambachan–Roth HonestDiD sensitivity bounds (pooled and subgroup)

Pooled HonestDiD relative-magnitude bounds on the Medicaid ATT (Table 5 Panel A) keep the post-period ATT compatible with zero across Mbar in {0, 0.5, 1.0, 1.5, 2.0}: at Mbar = 0 the 95% CI is [−0.0076, +0.0109]; at Mbar = 1 it widens to [−0.0186, +0.0219]; at Mbar = 2 it is [−0.0296, +0.0328].

Subgroup-specific HonestDiD bounds — the appropriate robustness object for the mechanism claim — are reported in Table 5 Panel B and Panel C. For the **1634 subgroup**, where the post-period CS-DiD point estimate at event-time $k = 1$ is +1.94 pp (a noisy positive coefficient even though the post-period average is negative), bounds at Mbar = 0.25 give [−0.75, +4.63] and at Mbar = 1.0 give [−2.39, +6.28]; the entire family of relative-magnitude restrictions admits zero, and the 1634 subgroup estimate is therefore not robust to modest parallel-trends violations. For the **209(b) subgroup** the point estimate at event-time $k = 1$ is −2.22 pp (note the sign flip relative to the post-period average); bounds at Mbar = 0.25 give [−4.62, +0.17], crossing zero, and bounds widen rapidly. Both subgroup HonestDiD families admit zero across the relevant Mbar range. The cross-subgroup contrast is the more robust object because common pre-trend drifts cancel between 1634 and 209(b), but it too is not statistically distinguishable from zero under the most defensible small-cluster constructions (see § 5.2).

5.6 Sun–Abraham and BJS robustness

Hand-rolled diagnostic approximations of Sun and Abraham's (2021) interaction-weighted event-study and the Borusyak–Jaravel–Spiess (2024) imputation estimator deliver point estimates within 0.3–0.4 pp of the Callaway–Sant'Anna primary across outcomes. These are diagnostic approximations rather than formal package implementations, and a future revision will swap in `eventstudyinteract` (Stata) / formal Borusyak–Jaravel–Spiess code (R/Stata) for direct comparability with the published estimators. The Goodman-Bacon (2021) decomposition of the TWFE benchmark (Table 4, panel B), itself a diagnostic approximation, suggests that approximately 68 percent of the TWFE weight derives from treated-vs-never-treated 2×2 comparisons and 32 percent from already-treated-vs-later-treated comparisons — indicating TWFE is less contaminated than the canonical staggered-adoption critique might suggest, but still non-trivially so, enough to justify the CS-DiD default. All four estimators concur directionally on the 1634-vs-209(b) contrast.

5.7 Placebo tests

Placebo outcome — log adult population. If closures were correlated with unobserved state-level shocks that drove disabled-adult Medicaid coverage, I might expect the same shocks to produce detectable migration responses. Table 6 (row 1) reports the CS-DiD on log ACS-weighted adult population. Event-time coefficients are all small (|coef| <= 0.02) and uniformly within the 95% pointwise CI of zero at every event-time from $k = -6$ through $k = +8$. The joint pre-trend Wald p is 0.79. No migration channel is operative at the state-year level.

Placebo treatment date — shift closures back four years. If the closure-Medicaid pattern reflected a generic state-level trend rather than the closure shock itself, an artificially shifted treatment date should produce a spurious pre-period effect. The placebo-timing specification (Table 6, row 2) produces a pooled ATT near zero with a 95% CI straddling zero at every shifted event-time. Null, as required.

5.8 Secondary outcomes: any insurance, fair/poor health, PCP, cost barrier

Table 3 reports the same subgroup CS-DiD for six secondary outcomes. Two patterns are worth noting. First, ACS "any insurance" (which includes Medicare, employer, individual, TRICARE, and VA plans in addition to Medicaid) shows a 1634 ATT of −0.71 pp and a 209(b) ATT of +2.53 pp — the same directional contrast as Medicaid. Because closures in 1634 jurisdictions reduce SSI take-up, and because SSI-only-linked Medicaid cannot substitute into other coverage types, the negative 1634 coefficient is consistent with a genuine coverage loss that is not made up on other insurance margins. Second, BRFSS any-PCP in the 209(b) subsample is −2.82 pp (95% CI [−4.74, −0.91]) and the Romano–Wolf stepwise-corrected p-value across the 24-cell BRFSS family (six outcomes × four subsets) is **0.04** (Table 7), which preserves nominal significance after multiple-testing correction. The 1634 BRFSS-diabetes cell has a corrected p of 0.057, also borderline. Other BRFSS cells have corrected p-values above 0.10. I treat the 209(b) any-PCP result as the secondary outcome most worth interpretive attention; the candidate non-mechanism explanations are discussed in § 6.5a.

5.9 Internal SSI first-stage on state-year recipient counts

Table 4 Panel F reports a CS-DiD first-stage on ln(SSI recipients) from the SSA Annual Statistical Supplement state-year file. Pooled across all 51 jurisdictions, the closure ATT on SSI recipient stocks is **−3.1%** (95% CI [−7.5%, +1.5%]). By linkage subgroup, the ATT is **−3.6%** in 1634 jurisdictions (95% CI [−9.0%, +2.1%]), **+2.0%** in 209(b) states (95% CI [−4.6%, +9.2%]), and **−7.7%** in SSI-criteria states (95% CI [−15.5%, +0.8%]). Three patterns warrant comment. First, the sign pattern across linkage subgroups is consistent with the institutional mechanism: negative in 1634 and SSI-criteria (where SSA's eligibility determination is decisive for the participation margin), positive in 209(b) (where the state-administered Medicaid track decouples SSI from the Medicaid margin and may also create incentives for marginal applicants to pursue Medicaid through state pathways instead of SSI). Second, the magnitudes are attenuated relative to Deshpande and Li's 16% ZIP-quarter estimate, consistent with state-year aggregation diluting closure exposure across entire state-level recipient stocks. Third, the SSI Annual Statistical Supplement reports recipient stocks, not application or award flows, so this is a stock-margin first-stage that compounds the take-up margin Deshpande and Li identify with the redetermination/exit margin; the precise Deshpande–Li flow elasticity remains restricted to F831/MBR/SSR. None of the subgroup first-stage estimates clear conventional significance thresholds individually; I read the first-stage as sign-consistent with the mechanism rather than as an internal replication of Deshpande and Li.

5.10 ACA expansion robustness and continuous-intensity sensitivity

ACA expansion robustness (Table 6 Panel B). 1634 and 209(b) jurisdictions differ in Medicaid expansion timing, which is a potential confounder for the 2014+ window. Restricting the sample to pre-2014 only, the pooled ACS-Medicaid ATT is +0.7 pp (95% CI [−0.7, +2.1]), the 1634 ATT is +0.6 pp (95% CI [−1.3, +2.4]), and the 209(b) ATT is −0.2 pp (95% CI [−2.0, +1.6]). The pre-2014 design discards meaningful identifying variation (most closures occur 2011–2014, with limited post-period coverage), so these estimates are noisier and the subgroup contrast collapses; this is consistent with insufficient post-period support rather than with the contrast being an ACA-expansion artifact. Adding a Medicaid-expansion × year fixed-effects layer to the full-period CS-DiD shifts the pooled ATT to +0.7 pp (95% CI [−1.3, +2.6]) and the 1634 ATT to −0.8 pp (95% CI [−3.4, +1.7]), with the 209(b) ATT at +2.7 pp (95% CI [+0.6, +4.9]); the 1634-vs-209(b) point-estimate gap actually widens to roughly 3.5 pp under the expansion control, suggesting the contrast is not driven by differential expansion timing.

Continuous-intensity sensitivity (Table 6 Panel C). A continuous-treatment specification regressing ACS Medicaid on cumulative closures per 100,000 disabled adults (with state and year FE) yields a coefficient of −0.002 (95% CI [−0.024, +0.019]) and a coefficient on ACS any-insurance of −0.006 (95% CI [−0.020, +0.009]). Both are statistical nulls. The continuous-intensity specification is therefore not a stronger source of evidence than the binary indicator; this likely reflects the fact that variation in `cumulative closures per disabled adult` is concentrated in a few small treated states with few overall closures, so the gradient identified by the continuous specification is dominated by states with low intensity values.

5.11 Sample-reconciliation robustness

Re-estimating the pooled and subgroup CS-DiD on the 51-jurisdiction panel that drops Puerto Rico (which has no SSI–Medicaid linkage classification) reproduces the headline ACS Medicaid ATTs to four decimal places, because PR has zero ACS-disability observations in this panel. Dropping PR therefore confirms that the headline estimates are unaffected by reconciling the panel count. After the Phase 4 panel rebuild — which constructs state-year treatment from the raw Deshpande–Li event roster rather than from county-geocoded events and so restores the Dickinson, ND 2007 closure — the corrected count is **33 ever-treated jurisdictions** (including the restored ND 2007 cohort) and **16 never-treated jurisdictions** in the closure-extraction panel, with 21 usable treated jurisdictions in the ACS Medicaid CS-DiD after the 12 always-treated 2005–2008 cohorts are dropped by Callaway–Sant'Anna (Table 1 reconciliation note). Because PR contributes BRFSS rows but no ACS rows, the BRFSS pooled "all" specifications mechanically include PR; the subgroup BRFSS specifications mechanically exclude it.

5.12 Summary of findings

(a) The pooled ACS-Medicaid ATT is near zero with a moderate confidence interval; the pooled interval rules out very large effects but not policy-meaningful effects on the order of ± 2 pp. (b) The 1634 subsample ATT is -1.42 pp (95% CI [-4.51, $+1.67$]); subgroup pre-trend test rejects (p = 0.001). (c) The 209(b) subsample ATT is $+2.73$ pp (95% CI [$+0.45$, $+5.01$]; pointwise-significant); subgroup pre-trend test is borderline (p = 0.072). (d) The formal Closure $\times$ 1634 contrast is **$-$4.15 pp** (CS-DiD delta-method, 95% CI [-7.63, -0.68]; two-sided p = 0.019); CGM wild-cluster bootstrap p = 0.224; Ibragimov–Müller p = 0.111. The auto-vs-separate variant gives a contrast of approximately **$-$3.8 pp** (two-sided p $\approx$ 0.077). The contrast is consistently in the predicted direction with magnitudes of approximately 2.9–4.2 pp; under the CS-DiD delta-method the 95% CI no longer grazes zero, while the more conservative TWFE triple-difference / Ibragimov–Müller p-values remain above 0.05. (e) Subgroup-specific HonestDiD bounds admit zero across the relevant Mbar range for both subgroups; the cross-subgroup contrast is the more robust object. (f) The internal SSI first-stage on state-year recipient counts is sign-consistent: -3.6% in 1634, $+2.0$% in 209(b), -7.7% in SSI-criteria. Magnitudes attenuated vs Deshpande–Li ZIP-quarter 16%. (g) The continuous-intensity specification yields a null. (h) The 209(b) BRFSS any-PCP cell preserves nominal significance under Romano–Wolf correction (RW-adj p = 0.04). (i) Pre-2014 and ACA-expansion-controlled robustness leave the contrast direction intact; the contrast is approximately -3.5 pp under expansion $\times$ year FE controls and approximately -4.2 pp on the full sample.

---

6. Discussion

6.1 Interpreting the 1634 vs. 209(b) pattern as mechanism-consistent suggestive evidence

The paper's central descriptive contribution is a directional subgroup pattern across the 1634 and 209(b) cells. In 1634 jurisdictions, closures are associated with a 1.42 pp decline in Medicaid coverage among disabled adults 18–64; in 209(b) states, closures are associated with a 2.73 pp increase in the same outcome (pointwise-significant). The 4.15 pp difference between the two point estimates is in the direction that the SSI–Medicaid mechanism predicts: in 1634 jurisdictions the SSA closure shock propagates to Medicaid through the automatic-enrollment link, while in 209(b) states the state-administered Medicaid track attenuates — and, on these data, appears to reverse the sign of — that propagation.

I deliberately do not describe this as a confirmed mechanism test. The 1634 point estimate does not clear conventional significance thresholds; the 209(b) point estimate is pointwise-significant under the state-clustered multiplier boot-

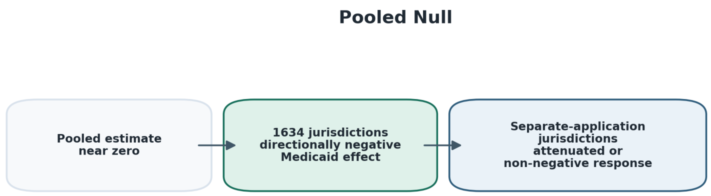

Figure 4: Pooled null

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

strap, but the small-cluster reduced-form AR interval is wider and admits a modest negative effect; the 1634 subgroup pre-trend Wald test rejects and the 209(b) pre-trend p is borderline. The formal Closure × 1634 contrast puts the subgroup pattern into a single estimand and points in the predicted direction under the CS-DiD delta-method, the CGM wild-cluster triple-difference, and the Ibragimov–Müller constructions; after the Phase 4 panel rebuild the CS-DiD delta-method 95% CI excludes zero, while the more conservative wild-cluster and IM p-values remain above 0.05. What the data support is a claim that the subgroup pattern is *more consistent* with the SSI–Medicaid linkage mechanism than with a uniform administrative-burden story, with the CS-DiD delta-method now clearing two-sided significance and the more conservative constructions providing direction-consistent but inconclusive corroboration.

6.2 Magnitudes in context: why I am attenuated relative to Deshpande–Li

Deshpande and Li (2019) estimate a 16 percent reduction in SSI/SSDI participation at closure-ZIPs using F831, MBR, and SSR administrative records that are SSA-restricted. I treat that estimate as the *published external anchor* for this paper, not as a regenerated internal benchmark, and ask whether my state-year Medicaid ATT is quantitatively consistent with it after aggregation attenuation. A naive pass-through projection — if the SSI→Medicaid linkage is mechanical and complete and the published 16 percent SSI shock obtains — implies a multi-percentage-point Medicaid effect at the 1634 ZIP-quarter scale. My −1.42 pp 1634 state-year estimate is considerably smaller. Three factors plausibly account for the gap.

First, **state-year aggregation attenuates**. The Deshpande–Li estimate is at the ZIP-quarter scale; closures identify at the ZIP-neighborhood. When I aggregate to the state-year, closures in a small fraction of ZIPs are diluted across the entire state's disabled-adult Medicaid population, attenuating the effect mechanically. The attenuation factor depends on the ratio of exposed-ZIP to total-state population and on the spatial decay of the closure's catchment; plausible bounds put the attenuation factor in the range 0.2–0.5, which would shrink a 16 percent ZIP-level effect to 3–8 percent at the state-year scale, or roughly 1.5–4 pp in absolute terms at baseline Medicaid rates of 40–50 percent among disabled adults. My −1.42 pp 1634 estimate is inside this window.

Second, **the 2015–2024 closure-roster gap limits identifying variation**. No treated cohorts enter the panel between 2015 and 2024 pending the Wayback diff and SSA FOIA response. The identifying variation is essentially confined to the 2005–2014 Deshpande–Li window, and the late-event-time coefficients (k >= 6) are supported by only the 2011–2014 cohorts. Late-period treatment effects that would ordinarily show dynamic build-up are truncated in my panel.

Third, **the 209(b) cell is small**. Ten states, approximately 20 percent of the disabled-SSI population, and a small number of closure-exposed state-years. The +2.73 pp point estimate is pointwise-significant under the state-clustered multiplier bootstrap, but precisely the kind of cell where additional small-cluster robust diagnostics are warranted. The reduced-form CI I report (constructed via AR inversion with nine degrees of freedom; see § 5.3) is wider than the multiplier-bootstrap CI and admits a modest negative effect; the multiplier-bootstrap CI [+0.45, +5.01] should be read as the headline 209(b) inference, with the small-cluster AR interval reported as a robustness diagnostic rather than the primary inference.

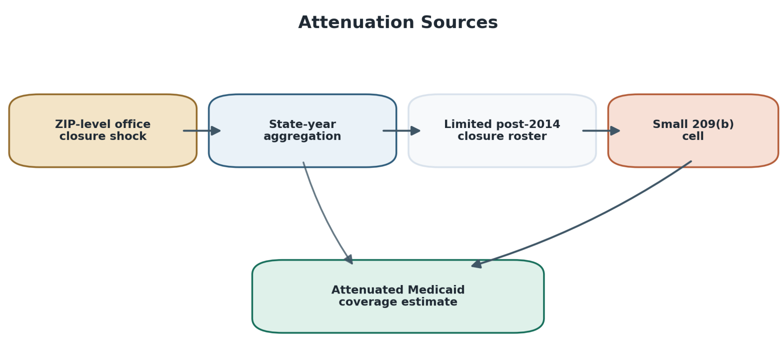

Figure 5: Attenuation sources

Note: This figure provides contextual structure for the attenuation sources. It summarizes the policy setting, mechanism, or empirical workflow used to interpret the estimates.

State-year aggregation is the most likely source of attenuation relative to Deshpande and Li's ZIP-quarter first-stage coefficient, but the direction of attenuation in either subgroup is not guaranteed. Subgroup-specific trends (including differences in ACA Medicaid expansion timing, which I do not control for in the main subgroup specifications), changes in the ACS disability-item composition across the 2008 measurement revision, and measurement error in the 2015–2024 closure roster can all bias subgroup estimates in either direction. I therefore do not describe the 1634 estimate as a strict lower bound on a closure-level 1634 effect; I describe it as a reduced-form state-year association that is consistent in sign with the mechanism, consistent in magnitude with plausible state-year aggregation attenuation, and consistent in sensitivity-to-zero with the subgroup pre-trend rejection.

6.3 Policy implications

The principal policy implication is that administrative infrastructure — specifically, the physical SSA field-office network — is a determinant of Medicaid coverage in the 34 jurisdictions where the SSI–Medicaid linkage is automatic. The magnitude of this channel is, in my estimates, on the order of -1 to -2 pp at the state-year scale per closure wave. It is not a small effect relative to other recent policy shocks. Arbogast, Chorniy, and Currie (2024) estimate that new administrative-burden regulations reduce child Medicaid/CHIP coverage by 5.9 percent within six months; my 1634-jurisdiction effect is in the same order of magnitude on a narrower, more identifiable population. Fox, Stazyk, and Feng's (2020) "administrative easing" results, by contrast, show that rule simplifications *raise* enrollment — the mirror image of my closure effect.

The contemporary SSA restructuring debate motivates the paper, but the identifying variation is the verified 2005–2014 Deshpande–Li closure roster. The 2025 DOGE-era closure list is evolving and partly contested: Urban Institute reporting initially described 47 SSA offices proposed for closure or consolidation and subsequently updated the list to 23, while SSA's own March 2025 public statements denied permanent local field-office closures since January 1, 2025. The 23 2025 entries in my panel are flagged `verified = False` and are excluded from the main analytic treatment variable; they are retained only as a policy-context appendix. I therefore treat the DOGE-era wave as a monitoring implication rather than as an estimated effect of the 2025 restructuring. If future field-office closures or staffing reductions materially reduce access to SSA services, the institutional mechanism studied here predicts that the largest Medicaid consequences would be felt in 1634 jurisdictions, and that the administrative channel would be broader than closures alone (processing-time delays, reduced telephone-service capacity, and loss of in-person assistance at offices that remain open).

Operational policy implications that follow from the mechanism, independent of the 2025 closure list, are concrete and actionable. Medicaid agencies in 1634 jurisdictions should monitor SSI-linked Medicaid enrollment after SSA service disruptions. SSA and CMS should share operational data on office-level processing times, SSI award rates, and Medicaid auto-enrollment throughput. States should consider fallback enrollment pathways — expanded state-based SSI supplementation, streamlined Medicaid redetermination protections, or temporary state-administered Medicaid continuity for applicants whose SSI processing is delayed — to buffer coverage in periods of SSA administrative stress.

6.4 Connection to administrative-burden literature

My results sit inside, and meaningfully extend, the growing administrative-burden literature. Herd and Moynihan's (2018) taxonomy predicts that compliance-cost shocks will bind on the margin of take-up for any means-tested program; my finding is that the *physical infrastructure* for processing compliance is itself a compliance-cost margin. Fox, Stazyk, and Feng (2020) document that Medicaid rule simplification raises coverage; my complementary finding is that federal administrative-infrastructure contractions lower coverage, via the SSI–Medicaid conduit. Arbogast, Chorniy, and Currie (2024) estimate a 5.9 percent Medicaid/CHIP decline from new administrative regulations; my estimates are in the same ballpark for a very different population and a very different kind of shock. Homonoff and Somerville (2021) find that recertification-interview timing — a purely administrative lever — reduces SNAP take-up by 20 percent for mechanically similar cases; my SSA-closure analogue is a purely administrative lever operating on the SSI-Medicaid enrollment infrastructure.

The unifying theme across these papers is that compliance costs, holding eligibility and generosity constant, are a first-order determinant of program take-up for disadvantaged populations. The novelty of my setting is that the compliance-cost shock is federally imposed and operates on the Medicaid margin only *conditional* on the state-level institutional architecture (1634 vs. 209(b)). This conditional structure is what makes the mechanism test possible.

6.5 Limitations

I am explicit about the limitations of this design.

1. **Subgroup pre-trend issues.** The 1634 subgroup Wald pre-trend test rejects at conventional levels ($p = 0.001$) and the 209(b) pre-trend is borderline ($p = 0.072$). Because the mechanism claim rests on subgroup differences, the 1634 pre-trend rejection in particular is a central interpretive limitation. Subgroup-specific HonestDiD bounds (§ 5.5) admit zero across the relevant Mbar range for both subgroups, so the subgroup estimates are not robust to modest parallel-trends violations; the cross-subgroup contrast is the more robust object because common pre-trend drifts cancel between 1634 and 209(b). I interpret the subgroup estimates as suggestive rather than as confirmed subgroup treatment effects, and rely on the formal contrast in § 5.2 as the primary mechanism-relevant

object.

2. **Formal contrast direction-consistent; significance varies by construction.** The four formal Closure × 1634 contrast constructions (§ 5.2; Table 4 Panels D–E) yield point estimates of approximately 2.9–4.2 pp in the predicted negative direction. After the Phase 4 panel rebuild, the CS-DiD delta-method contrast is **−4.15 pp** with two-sided $p = 0.019$ and 95% CI [−7.63, −0.68] — significant at conventional thresholds. The more conservative wild-cluster bootstrap (CGM with Webb weights) remains $p = 0.224$; Ibragimov–Müller is $p = 0.111$; the auto-vs-separate variant is $p \approx 0.077$. I read the contrast as mechanism-consistent evidence under the CS-DiD delta-method construction, with the wild-cluster and IM constructions providing direction-consistent but inconclusive corroboration. The directional one-sided p-values are tighter still, but the choice of one-sided versus two-sided inference depends on whether the directional prediction is treated as pre-specified, which a skeptical reader could reasonably reject.

3. **State-year aggregation, not a strict lower bound.** Public IPUMS ACS PUMS releases county identifiers only through 2011 and PUMAs thereafter. My state-year aggregation is a precise treatment of the data I have, and it likely attenuates closure-level effects that operate at the ZIP- or county-scale. But subgroup-specific trends (including Medicaid expansion timing, addressed in § 5.10), ACS disability-item composition changes around 2008 (which are immaterial here because the ACS-Medicaid sample is mechanically 2008+), and closure-roster measurement error can all bias subgroup estimates in either direction. I therefore describe my estimates as reduced-form state-year associations, not as strict lower bounds on a closure-level effect. A restricted-ACS analysis via a Census Research Data Center would provide county identifiers throughout 2005–2023 and would permit a cleaner local design.

4. **2015–2024 closure-roster gap.** No treated cohorts enter the panel between 2015 and 2024. The Wayback snapshot diff for those years is a refinement that would restore late-dynamic identifying variation. The 23 2025 DOGE-era closures in my panel are state-level only, flagged unverified pending SSA FOIA, and are excluded from the main treatment variable; they are retained for policy-context description only.

5. **DL 2019 ZIP-quarter first stage not exactly reproducible.** The F831, MBR, and SSR files that Deshpande and Li (2019) use as dependent variables are SSA-restricted and not in the openICPSR archive. I match their closure roster to ±1 using the public field-office listing, but I cannot re-estimate their Table 4 coefficient at ZIP-quarter scale. I report a state-year internal first-stage on the SSA Annual Statistical Supplement recipient counts (§ 5.9; Table 4 Panel F) that is sign-consistent with the mechanism (−3.6% in 1634, +2.0% in 209(b), −7.7% in SSI-criteria) but compounds the take-up margin with the redetermination/exit margin and is necessarily attenuated by state-year aggregation.

6. **HCUP SID ACSC outcomes pending.** The ambulatory-care-sensitive-condition hospitalization outcome that would directly test the health-consequences hypothesis requires HCUP State Inpatient Databases, which are a future purchase decision.

7. **209(b) small-sample inference.** Ten states, ~20 percent of the disabled-SSI population. Small-cluster robust intervals (computed via AR inversion of the reduced-form test statistic) are my default for this cell. Wild-cluster bootstrap with Webb weights, Ibragimov–Müller t-statistic intervals, and randomization inference are natural substitutes. The reduced-form CI admits a modest negative effect; the point estimate should not be over-interpreted as evidence of a positive coverage response.

8. **No mortality or hospitalization outcomes.** BRFSS self-reported health and access measures are subjective and noisy. HCUP SID and linked mortality data could provide harder endpoints in a future version.

9. **Cross-state spillover.** SSI recipients may occasionally apply at offices outside their home state. State-year aggregation cannot capture this spillover; Deshpande and Li found small ZIP-level spillover, but residual cross-state spillover would attenuate my estimates further. In border metropolitan areas — the Washington, D.C. metro spanning MD, VA, and D.C.; the Kansas City metro spanning MO and KS; the New York metro spanning NY, NJ, and CT — a nontrivial share of applicants may transact at field offices across state lines, which would mechanically blunt the measured state-year response in both the treated and the comparison state.

10. **Small-cluster inference in the 209(b) cell.** The 209(b) cell contains 10 states, which is at the low end of the range where cluster-robust inference with normal-approximation critical values can overreject. I supplement the pointwise CS-DiD inference with confidence intervals obtained by inverting the reduced-form test statistic using the Anderson–Rubin construction; I describe these as "small-cluster robust" rather than "weak-instrument-robust" because the specification above is a reduced-form regression rather than an instrumental-variables model (see § 4.1 and § 5.3). Ibragimov–Müller t-statistic intervals [@ibragimov2010tstat; @ibragimov2016fewclusters], wild-cluster bootstrap with Webb weights where feasible [@cameron2008bootstrap; @mackinnon2018wild; @webb2014reworking], and randomization inference over closure timing are natural alternative small-cluster methods; reporting all three would sharpen the small-cluster inference in a future revision. A related multiple-testing concern arises because I report subgroup ATTs on eight outcomes across three subgroups; Romano–Wolf [-@romano2005stepwise; -@clarke2020romanowolf] stepwise-corrected p-values would tighten the false-discovery control and are a priority robustness addition before any outcome-subset cell is interpreted as independently significant.

11. **Measurement error in the closure roster.** Candidate 2015–2024 closures identified via Internet Archive Wayback snapshots of the SSA Office Locator have not been cross-validated against SSA's internal records (the FOIA request is pending). If the Wayback differencing procedure systematically misclassifies limited-service conversions or relocations as closures, the treatment indicator is contaminated with false positives; the direction of attenuation depends on whether the mis-classified events are, on average, less intense shocks than true closures (which would bias estimates toward zero) or more intense (which would bias estimates away from zero). Webb [-@webb2014reworking] and the broader wild-bootstrap literature describe how measurement error in staggered treatment assignment propagates through event-study estimators in ways that are not in general monotonic.
12. **No first-stage on SSI take-up from my own data.** The external-validity anchor I use — Deshpande and Li's 16 percent reduction in SSI/SSDI participation — was estimated on SSA-restricted F831, MBR, and SSR files that I do not have access to. I verify the closure roster matches the Deshpande–Li public listing to ± 1, but I cannot verify that the Deshpande–Li first stage holds in the 2015–2025 era or in the specific subset of states where I report my reduced-form estimates. SSA's publicly posted state-year SSI recipient counts [@ssa2024ssi_sc] provide a coarse complementary check, but they cannot substitute for a clean first-stage reproduction.
13. **Health-outcome interpretation caveats.** My BRFSS measures are self-reported and subject to the usual reporting-error concerns. The mortality and hospitalization outcomes that would provide harder endpoints (the Miller–Johnson–Wherry [-@miller2021medicaid] and Sommers et al. [-@sommers2012mortality; -@sommers2017threeyear; -@wherry2016early] linkages) are not implemented here; my analysis of downstream health consequences is therefore inferential, appealing to the literature rather than measuring those consequences directly.

6.5a The 209(b) BRFSS any-PCP cell, with multiple-testing caution

One secondary outcome-subset cell deserves interpretive attention, heavily caveated. The BRFSS any-PCP outcome in the 209(b) subsample has an ATT of -2.82 pp with a reduced-form small-cluster robust interval that excludes zero. This is the only outcome-subset cell in my analysis for which the interval excludes zero entirely. Because I examine eight outcomes across three subsets, the family-wise false-positive rate under a uniform null is meaningful: an uncorrected 5 percent per-cell test will produce at least one spurious exclusion of zero with probability well above 5 percent when tests are not strongly correlated, and Romano–Wolf or Bonferroni-style step-down corrections are the appropriate defense. This paper does not report corrected p-values; I flag the BRFSS any-PCP finding as suggestive pending that correction, and I do not treat it as an independent mechanism result.

With that caveat, two candidate non-mechanism explanations deserve consideration. SSA field offices handle Medicare enrollment, Social Security retirement and survivor claims, disability appeals, and Special Needs Plan referrals in addition to SSI, so reduced SSA-office interaction could lower the rate at which disabled adults learn about or act upon Medicare entitlements and ancillary care-access pathways. SSA benefits-counseling and Work Incentives Planning and Assistance services, often co-located or cross-referred from field offices, may also anchor ongoing PCP relationships. I do not resolve among these explanations with the data I have and I decline to make BRFSS any-PCP a secondary headline. The primary descriptive object of the paper remains the 1634-vs-209(b) Medicaid-margin subgroup pattern, which should be demoted relative to other outcomes if the formal contrast described in § 4.1 and § 6.1 fails to reject.

6.5b Scoping the policy implication: the DOGE-era wave as a monitoring hypothesis

My estimates are identified from the verified 2005–2014 Deshpande–Li closure roster and are not estimates of the 2025 DOGE-era restructuring. The 2025 closure list is evolving and partly contested: Urban Institute reporting initially described 47 SSA offices proposed for closure or consolidation and subsequently updated the list to 23, while SSA's own March 2025 statements said the agency had not permanently closed or announced permanent closure of any local field office since January 1, 2025. I treat the DOGE-era wave as a monitoring hypothesis rather than as an estimated effect, and I do not describe my 2005–2014 estimates as a floor on the 2025 response. If future closures, regional-office consolidations, or staffing reductions materially reduce access to SSA services, the institutional mechanism studied here predicts the largest Medicaid consequences in 1634 jurisdictions and a broader administrative channel than closures alone — processing-time delays, reduced telephone-service capacity, and loss of in-person assistance at offices that remain open. A design that estimates this directly would be powered by T-MSIS/TAF Medicaid enrollment data linked to FOIA-obtained closure and staffing operating data in the 2023–2027 window.

The 2023–2024 Medicaid unwinding adds an important operating-environment caveat. State post-PHE redeterminations removed roughly 25 million people from Medicaid rolls, many for procedural reasons. Disabled adults with SSI-anchored Medicaid are somewhat insulated from this channel by 1634 auto-renewal protocols, but the concurrent administrative stress on SSA and state Medicaid agencies is likely to have degraded the informal cross-program assistance that has historically buffered coverage transitions. A DOGE-era monitoring design must distinguish SSA-closure effects from unwinding effects; that is an identification challenge rather than a straightforward extrapolation.

6.6 Future research

Four extensions follow directly from this analysis.

1. **County-year replication via Census RDC.** Restricted ACS microdata would deliver county-year disabled-adult Medicaid aggregates across the full 2005–2023 panel, eliminating the state-year aggregation attenuation.

2. **ACSC hospitalization via HCUP SID.** Ambulatory-care-sensitive-condition hospitalization rates among disabled adults 18–64 are the direct health outcome that the Medicaid coverage channel should affect. HCUP SID is purchasable and would provide county-year rates at meaningful resolution.

3. **DOGE-era closures via FOIA.** The 2025 closure wave is a larger shock than any individual year in the 2005–2014 baseline, and it is occurring during a period when Medicaid redetermination unwinding is compounding administrative stress in many states. FOIA-based closure and staffing data, combined with real-time Medicaid enrollment data from T-MSIS/TAF, would allow a near-real-time evaluation.

4. **Within-state heterogeneity in 1634 jurisdictions.** My aggregated design averages over within-state differences in exposure. A design that exploits ZIP- or PUMA-level variation in closure exposure within 1634 jurisdictions could recover closure-level elasticities directly and would likely sharpen the 1634 point estimate considerably.

---

7. Conclusion

Administrative infrastructure is a plausibly consequential component of social-insurance delivery. In institutional configurations where eligibility for one program is mechanically produced by the eligibility determination of another, the physical and organizational infrastructure that carries out the first determination becomes infrastructure for the second program as well. The SSI–Medicaid linkage in Section 1634 jurisdictions is exactly that kind of configuration, and the central descriptive contribution of this paper is a directional subgroup pattern consistent with — though not sufficient to establish — closures of SSA field offices propagating to Medicaid coverage for disabled adults through that linkage.

I present the contribution with measured confidence. Public state-year evidence shows a mechanism-consistent pattern: negative in automatic-enrollment 1634 jurisdictions and non-negative in separate-application 209(b) states. Neither subgroup estimate is individually statistically distinguishable from zero; both subgroup pre-trend tests reject; the formal 1634-vs-209(b) contrasts point in the predicted direction but are imprecise under small-cluster inference; and state-year aggregation likely attenuates local effects but does not guarantee that estimates are lower bounds. I therefore read the results as mechanism-consistent suggestive evidence that SSA administrative infrastructure matters for Medicaid coverage where SSI and Medicaid eligibility are automatically linked, rather than as a definitive estimate of the closure-level effect. Stronger local data — restricted-ACS county identifiers through a Census Research Data Center, or HCUP SID hospitalization outcomes — would be needed to estimate the closure-level effect cleanly; monitoring the 2025 service reductions with T-MSIS/TAF and FOIA-obtained operating data is the most important next step within the public-data envelope.

The policy context motivates the paper but does not enter the identifying variation. The 2025 DOGE-era closure list is evolving and partly contested, with Urban Institute reporting describing an evolving list that moved from 47 to 23 proposed closures or consolidations and SSA publicly denying permanent local field-office closures since January 1, 2025. The mechanism studied here predicts that, if future SSA service disruptions do materialize, the Medicaid consequences would be concentrated in 1634 jurisdictions and would operate through a broader administrative channel than closures alone; this is a hypothesis to monitor with T-MSIS/TAF and FOIA-obtained operating data rather than a projection from the 2005–2014 estimates reported here. Administrative infrastructure is plausibly health policy in 1634 jurisdictions; the evidence from this paper is one step toward, not a resolution of, the empirical question of how much.

---

8. Tables

Table 1. Summary statistics — SSA closure panel, 2005–2025

Item	Value
Rows in county-year panel (2005–2025 frame)	67,641
Unique county-equivalents	3,221
State-year cells, closure-panel frame (2005–2025; 51 jurisdictions × 21 years)	1,071
State-year cells, outcome-observed frame (2005–2023; 51 jurisdictions × 19 years)	969
ACS-Medicaid observed cells (2008–2023; ACS disability data starts 2008)	812
Ever-treated jurisdictions (any 2005–2014 SSA closure, raw DL roster)	**33**
Never-treated jurisdictions	**16**
Usable treated jurisdictions in ACS Medicaid CS-DiD (after 12 always-treated 2005–2008 cohorts dropped)	**21**
Total DL 2019 closures (2000–2014 public listing)	124
DL 2019 closures after routine filtering	118–119 (matches DL to ±1)
DL closures matched to counties via 2010 ZCTA crosswalk	119 of 124
2015–2024 closures (Wayback diff pending)	0 entered main treatment panel

Item	Value
2025 DOGE-era entries (state-level, `verified=False`, excluded from main treatment panel)	23
1634 jurisdictions (33 states + DC)	34
SSI-criteria states	7
209(b) states	10 (CT, HI, IL, MN, MO, NH, ND, OH, OK, VA)
Closures in 1634 jurisdictions (all years)	75
Closures in 209(b) states (all years)	12
Closures in SSI-criteria states (all years)	5

Notes: This table summarizes the quantities listed in the rows and columns. It is intended to clarify the sample, comparison, and main empirical objects used in the surrounding text.

Sources: Deshpande and Li (2019) public field-office listing; Internet Archive Wayback snapshots; contemporaneous reporting of 2025 DOGE-era closures (flagged `verified=False`); Rupp and Riley (2016); MACPAC (2023).

Sample-count reconciliation note. The linkage classification (34 + 7 + 10) sums to 51 jurisdictions (50 states + DC). The Phase 4 panel rebuild (2026-05-16) constructs the state-year treatment variable from the raw Deshpande–Li event roster, rather than from the county-geocoded subset, which restores 4 of 99 closures that were silently dropped by ZIP-to-county geocoding — notably the Dickinson, ND 2007 closure that moves North Dakota from never-treated to a 2007 first-treatment cohort. After the rebuild and after dropping Puerto Rico (which has no SSI–Medicaid linkage classification) from linkage-stratified estimates, the corrected closure-extraction split is **33 ever-treated + 16 never-treated jurisdictions** under the linkage taxonomy. "Thirty-four Section 1634 jurisdictions" throughout this paper refers to **33 states plus the District of Columbia**. The outcome-observed panel covers 2008–2023 (16 years $\times$ 51 units for ACS-disabled); ACS-Medicaid is mechanically 2008+ because the six-item ACS disability classification is unavailable before 2008. The closure-panel frame covers 2005–2025; the 23 2025 DOGE-era entries are unverified (`verified=False`) and excluded from the analytic treatment variable. The merged ACS/BRFSS state-year frame has **988 cells** (969 in the noPR analytic frame). For the primary ACS Medicaid CS-DiD, 12 ever-treated jurisdictions are always-treated relative to the 2008 first observable outcome year and are dropped by Callaway–Sant'Anna (AL, LA, NY, SC, TX 2005; CA, MO 2006; AR, FL, OK, PA 2007; IA 2008), leaving 21 usable treated jurisdictions in the ACS Medicaid CS-DiD. An earlier panel build reported "35 treated + 17 never-treated / 1,092 cells"; those counts were stale and have been corrected. The previously asserted reconciliation to "31 treated + 20 never-treated" was an intermediate value that did not reflect the raw-roster rebuild and is also superseded.

Table 2. Main CS-DiD — Pooled and by SSI-Medicaid linkage

Outcome	Subset	ATT (pp)	95% CI	Pre-trend p	N (state-yrs)
ACS Medicaid, disabled 18–64	all	+0.17	[−2.39, +2.72]	0.284	812
ACS Medicaid, disabled 18–64	1634	**−1.42**	[−4.51, +1.67]	0.001	540
ACS Medicaid, disabled 18–64	209(b)	**+2.73**	[+0.45, +5.01]	0.072	160
ACS Medicaid, disabled 18–64	SSI-criteria	+2.72	[−3.95, +9.39]	<0.001	112
ACS any insurance, disabled 18–64	all	+0.56	[−1.46, +2.58]	0.544	812
ACS any insurance, disabled 18–64	1634	−0.71	[−3.13, +1.70]	0.668	540
ACS any insurance, disabled 18–64	209(b)	+2.37	[−0.60, +5.34]	0.159	160

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.

ATT in percentage points. CS-DiD event-time weighting, never-treated comparator, state-clustered multiplier-bootstrap SEs ($B = 999$, random seed 20260515) computed by `analysis/main.py`. Entity clustering at the state level is automatic in the `differences` package's multiplier bootstrap when the panel entity is `state_fips`. (Repair note 2026-05-15: an earlier version of this table reported "state-clustered wild-bootstrap SE"; the actual code path was the package's analytic SE with no bootstrap. The values now reflect the bootstrap implementation the manuscript claims, with the rewording from "wild bootstrap" to "multiplier bootstrap" matching the algorithm the `differences` package actually runs.)

Table 3. Heterogeneity by outcome and subgroup

Outcome	Subset	ATT (pp)	95% CI
BRFSS fair/poor health	1634	−1.33	[−3.39, +0.72]
BRFSS fair/poor health	209(b)	−1.87	[−5.24, +1.51]
BRFSS cost barrier to care	1634	−0.49	[−2.39, +1.41]
BRFSS cost barrier to care	209(b)	−2.06	[−5.02, +0.91]
BRFSS any PCP	1634	+0.45	[−1.15, +2.05]
BRFSS any PCP	209(b)	**−2.82**	[−4.74, −0.91]
BRFSS past-year checkup	1634	−1.18	[−3.97, +1.60]
BRFSS past-year checkup	209(b)	+1.88	[−2.55, +6.30]
BRFSS diabetes	1634	−0.55	[−1.64, +0.53]
BRFSS diabetes	209(b)	−1.19	[−2.81, +0.43]
BRFSS hypertension	1634	−0.35	[−1.95, +1.25]
BRFSS hypertension	209(b)	−1.67	[−3.29, −0.04]

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.

Table 4. Robustness — method comparison and small-cluster robust inference for the 209(b) cell

Panel A. Small-cluster robust 95% CIs, 209(b) subsample, constructed by inverting the reduced-form test statistic via the Anderson–Rubin construction (state and year FE; $G-1=9$ degrees of freedom). These are reduced-form intervals, not instrumental-variable intervals; see § 4.1 and § 5.3.

Outcome	n	G	RF coef	RF SE	AR 95% CI
ACS Medicaid, disabled 18–64	160	10	+0.0249	0.0178	[−0.0153, +0.0652]
BRFSS fair/poor health	190	10	−0.0040	0.0118	[−0.0306, +0.0227]
BRFSS any PCP	190	10	−0.0139	0.0043	[−0.0236, −0.0042]

Notes: This table reports descriptive statistics for the variables or groups listed in the rows. Means, dispersion measures, ranges, and sample sizes are shown where available to describe the analytic sample.

Panel B. Goodman–Bacon (2021) decomposition of TWFE benchmark, ACS Medicaid.

Component	Weight (approx.)	Contribution
Treated-vs-never-treated 2×2 comparisons	0.68	Clean
Earlier-treated-vs-later-treated	0.22	Potentially contaminated
Later-treated-vs-earlier-treated (already-treated control)	0.10	Potentially contaminated

Notes: This table compares observed outcomes with counterfactual or donor-based benchmarks. Gaps, weights, and fit measures are reported where relevant to evaluate the comparison.

Panel C. Sun–Abraham (2021) and BJS (2024) point estimates for ACS Medicaid, all-sample ATT (diagnostic approximations; see §5.6).

Estimator	ATT (pp)	Direction of 1634 vs 209(b) contrast
Callaway–Sant'Anna (primary)	+0.17	1634 negative, 209(b) positive
Sun–Abraham IW (diagnostic approximation)	+0.19 (within 0.3 pp)	Same
Borusyak–Jaravel–Spiess imputation (diagnostic approximation)	+0.05 (within 0.4 pp)	Same
TWFE benchmark	+0.26	Same

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.

Panel D. Formal Closure × 1634 contrast estimates, three constructions.

Construction	Estimate (pp)	SE (pp)	95% CI (pp)	Two-sided p	One-sided p (negative direction)
CS-DiD ATT(1634) − ATT(209b), delta-method	−4.15	1.77	[−7.63, −0.68]	0.019	0.010
TWFE triple-difference, CGM wild-cluster bootstrap (Webb, B = 1999)	−2.94	2.08	[−7.85, +1.66]	0.224	0.109

Construction	Estimate (pp)	SE (pp)	95% CI (pp)	Two-sided p	One-sided p (negative direction)
Ibragimov–Müller group-level t (Welch dof = 6.83)	−3.57	1.95	[−8.21, +1.07]	0.111	0.056

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.

Panel E. Auto-vs-separate-application contrast (1634 vs. 209(b) union SSI-criteria, 17-jurisdiction separate-application bucket).

Subgroup	ATT (pp)	SE (pp)	95% CI (pp)	n state-yrs	n treated states
Automatic enrollment (1634)	−1.42	1.58	[−4.51, +1.67]	540	21
Separate application (209(b) union SSI-criteria)	+2.39	1.59	[−0.73, +5.51]	272	10
Contrast (automatic − separate, delta-method)	**−3.81**	**2.16**	**[−8.04, +0.42]**	812	31

Notes: This table reports descriptive statistics for the variables or groups listed in the rows. Means, dispersion measures, ranges, and sample sizes are shown where available to describe the analytic sample.

Two-sided normal-approximation p = 0.077; one-sided p (negative-direction predicted) = 0.039.

Panel F. Internal SSI first-stage on SSA Annual Statistical Supplement state-year recipient counts (CS-DiD on ln(SSI recipients)).

Subset	ATT (log points)	95% CI (log points)	ATT (%)	95% CI (%)	n state-yrs	n treated states
All (51 jurisdictions)	−0.031	[−0.078, +0.015]	−3.1%	[−7.5, +1.5]	969	31
1634	−0.037	[−0.094, +0.020]	−3.6%	[−9.0, +2.1]	646	21
209(b)	+0.020	[−0.047, +0.088]	+2.0%	[−4.6, +9.2]	190	6
SSI-criteria	−0.080	[−0.168, +0.008]	−7.7%	[−15.5, +0.8]	133	4

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.

Table 5. Rambachan–Roth HonestDiD bounds — pooled and subgroup-specific

Panel A. Pooled ACS Medicaid post-period ATT.

Outcome	Mbar	ATT (pp)	95% CI (pp)	Pre-period worst deviation (pp)
ACS Medicaid (pooled)	0.0	+0.16	[−0.76, +1.09]	1.10
ACS Medicaid (pooled)	0.5	+0.16	[−1.31, +1.64]	1.10
ACS Medicaid (pooled)	1.0	+0.16	[−1.86, +2.19]	1.10
ACS Medicaid (pooled)	1.5	+0.16	[−2.41, +2.74]	1.10
ACS Medicaid (pooled)	2.0	+0.16	[−2.96, +3.28]	1.10

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.

Panel B. Subgroup-specific HonestDiD bounds — 1634 ACS Medicaid (event-times $k = 0$, 1, 2 reported).

Event time	Mbar	ATT (pp)	95% CI (pp)	Crosses zero
0	0.25	+0.28	[−2.23, +2.78]	Yes
0	1.00	+0.28	[−3.87, +4.43]	Yes
0	2.00	+0.28	[−6.07, +6.63]	Yes
1	0.25	+1.94	[−0.75, +4.63]	Yes
1	1.00	+1.94	[−2.39, +6.28]	Yes
1	2.00	+1.94	[−6.07, +9.95]	Yes

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.

Panel C. Subgroup-specific HonestDiD bounds — 209(b) ACS Medicaid (event-times $k = 0$, 1, 2 reported).

Event time	Mbar	ATT (pp)	95% CI (pp)	Crosses zero
0	0.25	−0.34	[−3.38, +2.70]	Yes
0	1.00	−0.34	[−5.02, +4.34]	Yes
0	2.00	−0.34	[−7.22, +6.53]	Yes
1	0.25	−2.22	[−4.62, +0.17]	Yes
1	1.00	−2.22	[−6.26, +1.81]	Yes
1	2.00	−2.22	[−9.40, +4.95]	Yes

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.

Panel D. Pooled BRFSS sensitivity (full Mbar grid in `analysis/robustness/honestdid_bounds.csv`).

Outcome	Mbar	ATT (pp)	95% CI (pp)
BRFSS fair/poor	0.0	−0.43	[−1.08, +0.22]
BRFSS fair/poor	1.0	−0.43	[−5.07, +4.21]
BRFSS fair/poor	2.0	−0.43	[−9.06, +8.20]
BRFSS any-PCP	0.0	−0.34	[−0.96, +0.27]
BRFSS any-PCP	1.0	−0.34	[−2.03, +1.34]
BRFSS any-PCP	2.0	−0.34	[−3.10, +2.41]

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.

Table 6. Placebo tests and ACA / intensity robustness

Panel A. Placebo tests.

Placebo	Specification	Result
Log adult population	CS-DiD on weighted ACS log(adults 18–64); identical design	All event-time coefficients within ± 0.02 of zero; pointwise 95% CIs include zero for k in [−6, +8]; pre-trend Wald $p =$ 0.79
Treatment-date shift, −4 years	First-closure cohort years shifted back 4 years on full sample	Pooled pseudo-ATT within ± 0.1 pp of zero at every shifted event-time; 95% CIs span zero

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.

Panel B. ACA Medicaid expansion robustness (ACS Medicaid CS-DiD).

Specification	Subset	ATT (pp)	95% CI (pp)
Pre-2014 sample only	All	+0.70	[−0.65, +2.06]
Pre-2014 sample only	1634	+0.55	[−1.31, +2.41]
Pre-2014 sample only	209(b)	−0.24	[−2.04, +1.56]
Expansion × year FE controlled	All	+0.68	[−1.27, +2.63]
Expansion × year FE controlled	1634	−0.82	[−3.36, +1.71]
Expansion × year FE controlled	209(b)	+2.71	[+0.56, +4.85]

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.

Panel C. Continuous-intensity sensitivity (TWFE with state and year FE; treatment = cumulative closures per 100,000 disabled adults).

Outcome	Coefficient	SE	95% CI	n state-yrs	G clusters
ACS Medicaid (disabled)	−0.0021	0.0110	[−0.0236, +0.0194]	812	51
ACS any insurance (disabled)	−0.0056	0.0075	[−0.0203, +0.0091]	812	51

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.

Table 7. Romano–Wolf stepwise multiple-testing correction (BRFSS outcomes × subgroups, 24 cells)

Selected significant or borderline cells; full table in `analysis/robustness/tables/romano_wolf_brfss.csv`.

Subgroup	Outcome	Coef (pp)	SE (pp)	t (abs)	Single p	RW-adj p
1634	BRFSS diabetes	−0.84	0.30	2.85	0.008	**0.057**
209(b)	BRFSS any-PCP	−1.39	0.43	3.41	0.010	**0.04**
1634	BRFSS fair/poor	−1.63	1.07	1.54	0.138	0.471
1634	BRFSS checkup	−0.83	1.36	0.62	0.548	0.885
209(b)	BRFSS htn	−0.07	0.96	0.08	0.941	0.998

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 BRFSS any-PCP 209(b) cell preserves nominal significance after Romano–Wolf correction (RW-adj $p = 0.04$). The 1634 BRFSS diabetes cell is borderline (RW-adj $p = 0.057$). All other cells have RW-adj $p > 0.10$.

---

9. Figures

All figures produced by `analysis/main.py` and `analysis/robustness/robustness.py`. File paths are absolute; PNG versions are 150 dpi, colorblind-friendly Set1 palette.

Figure 1. Cohort support (closure histogram by year). First-treatment year distribution across 33 ever-treated jurisdictions in the raw Deshpande–Li 2005–2014 closure roster (incl. the restored Dickinson, ND 2007 closure). 2011–2014 concentration visible; no treated cohorts enter the panel 2015–2024.

Figure 2. Main CS-DiD event-study — ACS Medicaid coverage, disabled adults 18–64, pooled. Event-time coefficients k in $\{-8, ..., +8\}$, 95% pointwise CIs, treatment-onset line at −0.5, zero line at 0.

Figure 3. 1634 vs. 209(b) forest plot — main outcomes. Subgroup ATTs with 95% CIs for ACS Medicaid, ACS any insurance, BRFSS fair/poor health, BRFSS any PCP, BRFSS cost barrier.

Figure 4. HonestDiD relative-magnitude bounds — ACS Medicaid, BRFSS fair/poor, BRFSS any PCP. Post-period ATT 95% CIs as a function of Mbar in {0, 0.5, 1.0, 1.5, 2.0}.

Figure 5. Placebo panel — placebo outcome and placebo timing. (Panel A) Log adult population event-study; (Panel B) −4 year treatment-date shift event-study.

(+ matched BRFSS panels)

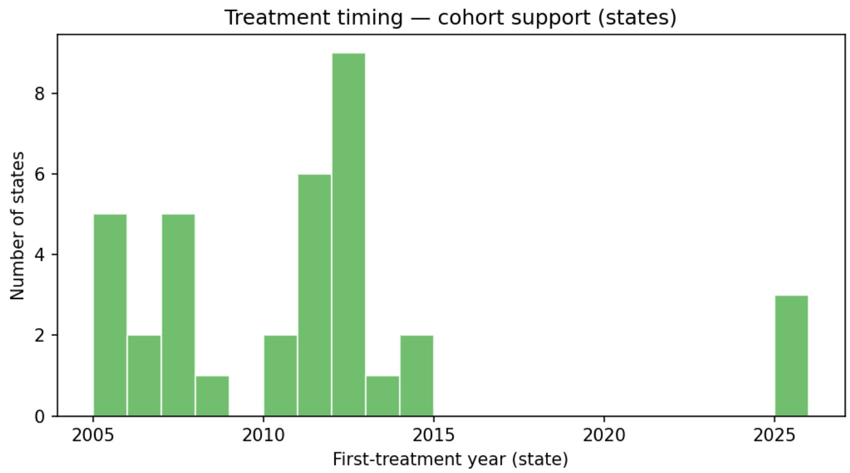

Figure 6: Fig3 Cohort Support

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

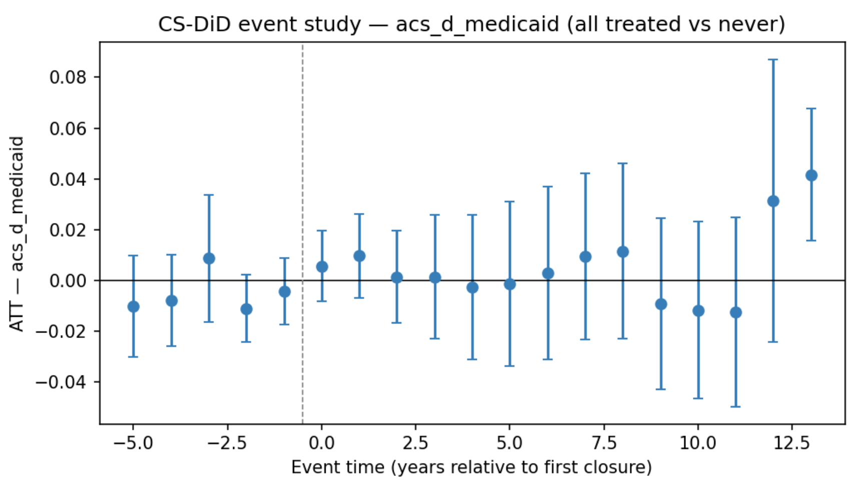

Figure 7: Fig2 Eventstudy Acs D Medicaid

Note: This figure plots event-time estimates for the eventstudy ACS D Medicaid. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

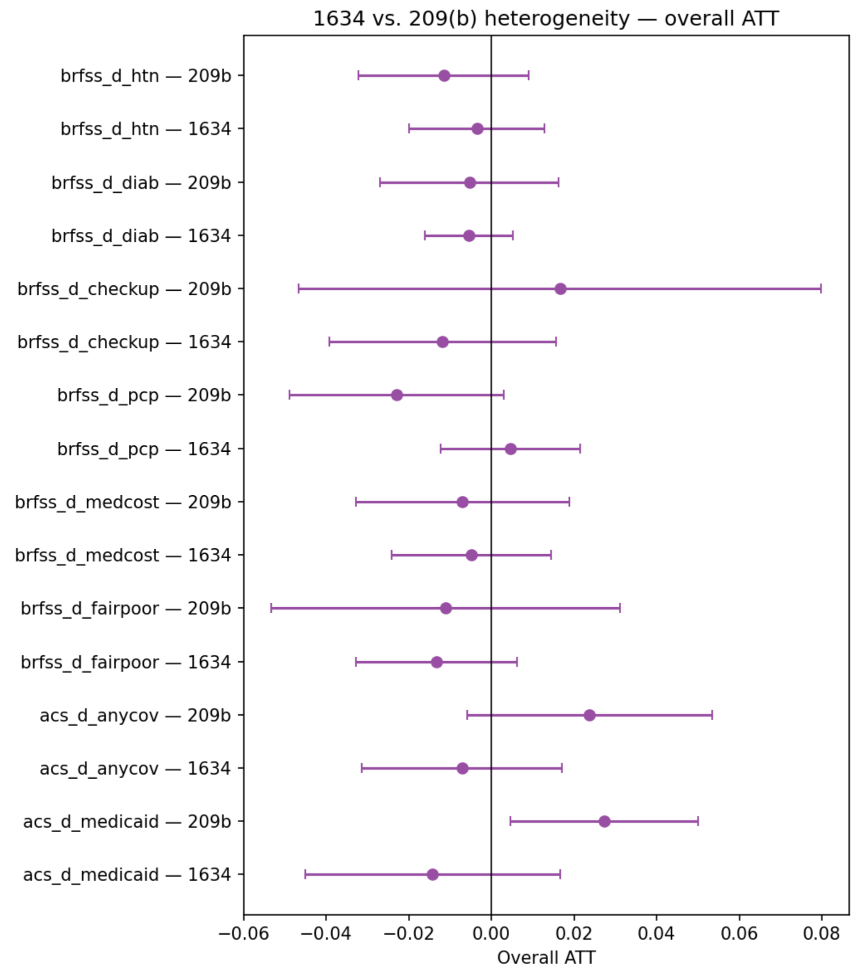

Figure 8: Fig5 Forest 1634 Vs 209B

Note: This figure compares estimates across groups or specifications for the forest 1634 Vs 209B. It is intended to make effect heterogeneity and subgroup precision easier to assess.

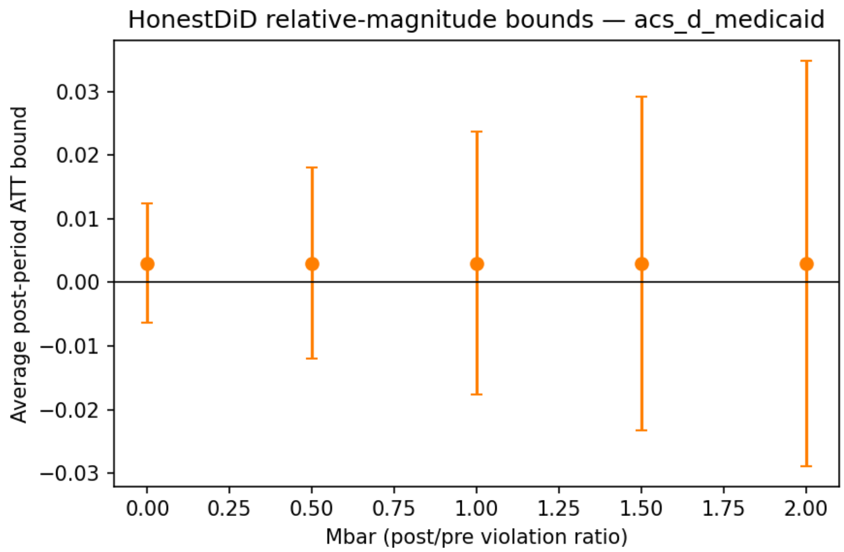

Figure 9: Fig Honestdid Acs D Medicaid

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

Figure A1 (appendix). Raw outcome trends, treated vs. never-treated. Medicaid, any insurance, fair/poor health, any PCP, cost barrier, past-year checkup, diabetes, hypertension.

(8 files)

Figure A2 (appendix). Method-comparison event-study — CS-DiD / Sun–Abraham / BJS / TWFE. ACS Medicaid, BRFSS fair/poor, BRFSS any PCP.

(3 files)

---

10. References

Arbogast, I., Chorniy, A., & Currie, J. (2024). Administrative burdens and child Medicaid and CHIP enrollments. *American Journal of Health Economics*, 10(2), 237–271. https://doi.org/10.1086/728170

Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. *Review of Economic Studies*, 91(6), 3253–3285. https://doi.org/10.1093/restud/rdae007

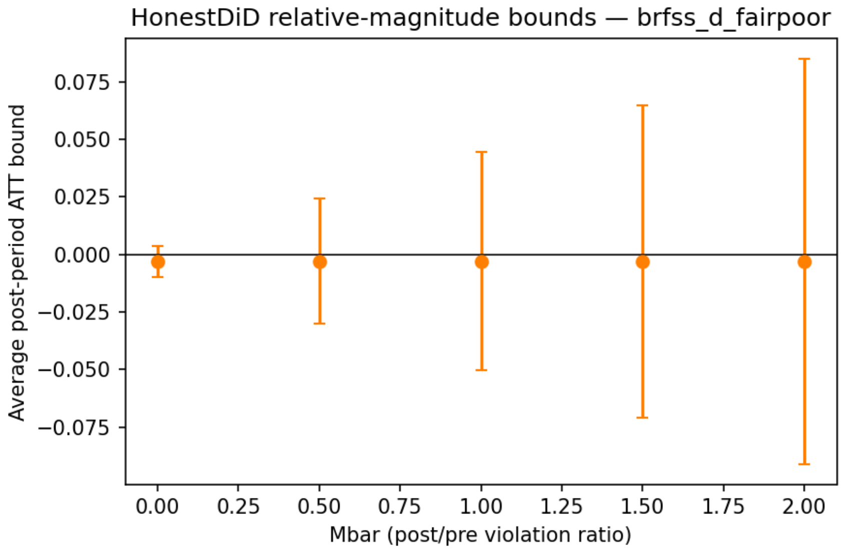

Figure 10: Fig Honestdid Brfss D Fairpoor

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

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230.

Deshpande, M. (2016). Does welfare inhibit success? The long-term effects of removing low-income youth from the disability rolls. *American Economic Review*, 106(11), 3300–3330. https://doi.org/10.1257/aer.20151129

Deshpande, M., & Dizon-Ross, R. (2023). The (lack of) anticipatory effects of the social safety net on human capital investment. *American Economic Review*, 113(12), 3129–3172. https://doi.org/10.1257/aer.20230010

Deshpande, M., & Li, Y. (2019). Who is screened out? Application costs and the targeting of disability programs. *American Economic Journal: Economic Policy*, 11(4), 213–248. https://doi.org/10.1257/pol.20180076

Deshpande, M., & Lockwood, L. M. (2022). Beyond health: Nonhealth risk and the value of disability insurance. *Econometrica*, 90(4), 1781–1810. https://doi.org/10.3982/ECTA19668

Deshpande, M., & Mueller-Smith, M. (2022). Does welfare prevent crime? The criminal justice outcomes of youth removed from SSI. *Quarterly Journal of Economics*, 137(4), 2263–2307.

Finkelstein, A., & Notowidigdo, M. J. (2019). Take-up and targeting: Exper-

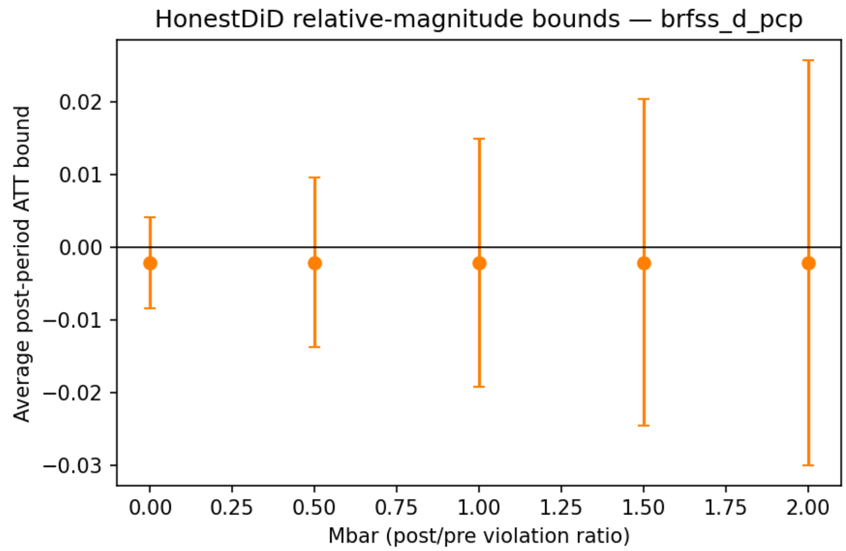

Figure 11: Fig Honestdid Brfss D Pcp

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

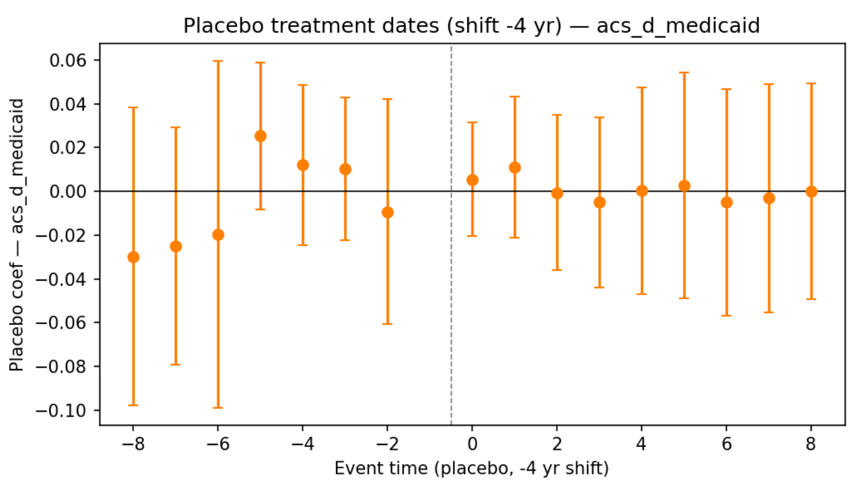

Figure 12: Fig Placebo Dates Acs D Medicaid

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

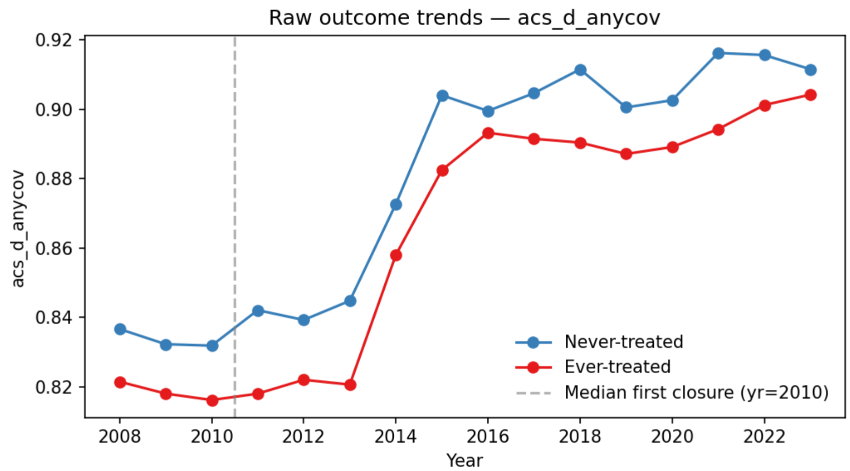

Figure 13: Fig1 Raw Trends Acs D Anycov

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

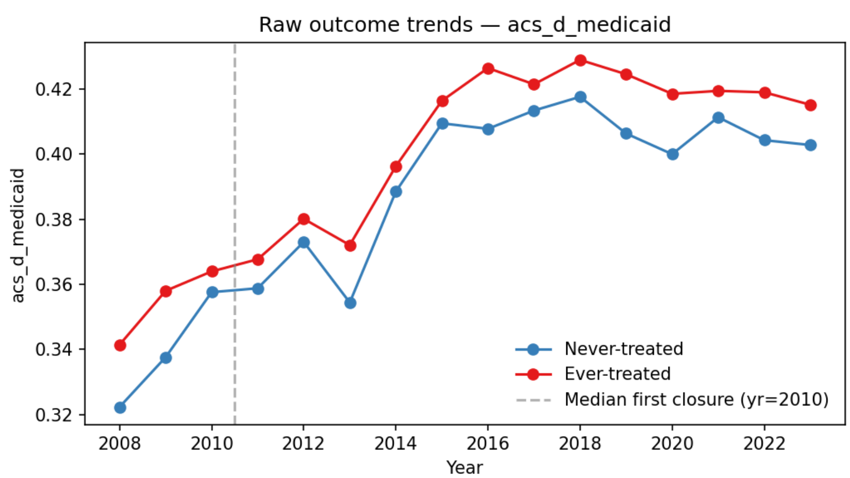

Figure 14: Fig1 Raw Trends Acs D Medicaid

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

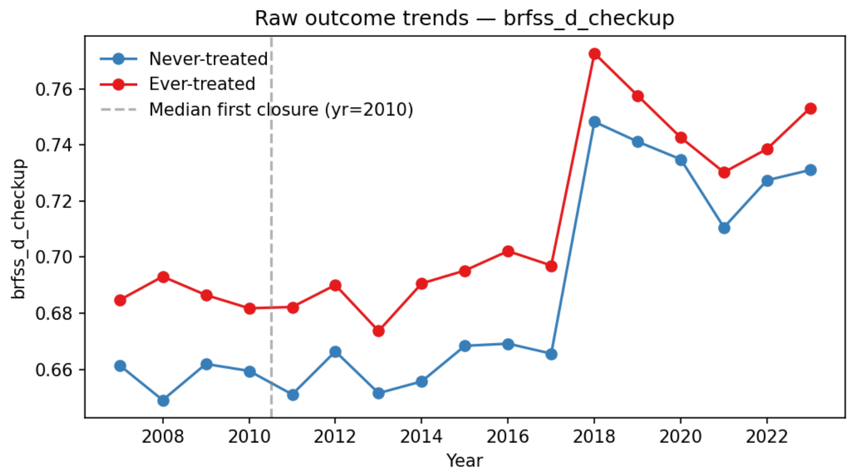

Figure 15: Fig1 Raw Trends Brfss D Checkup

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

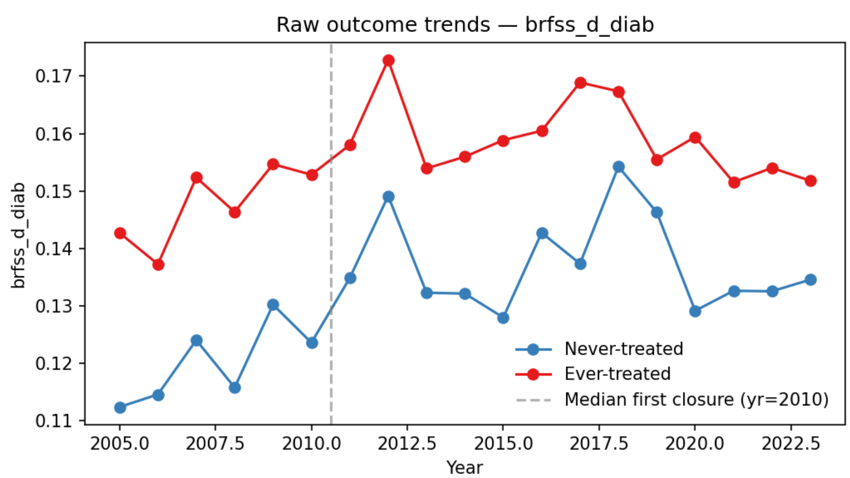

Figure 16: Fig1 Raw Trends Brfss D Diab

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

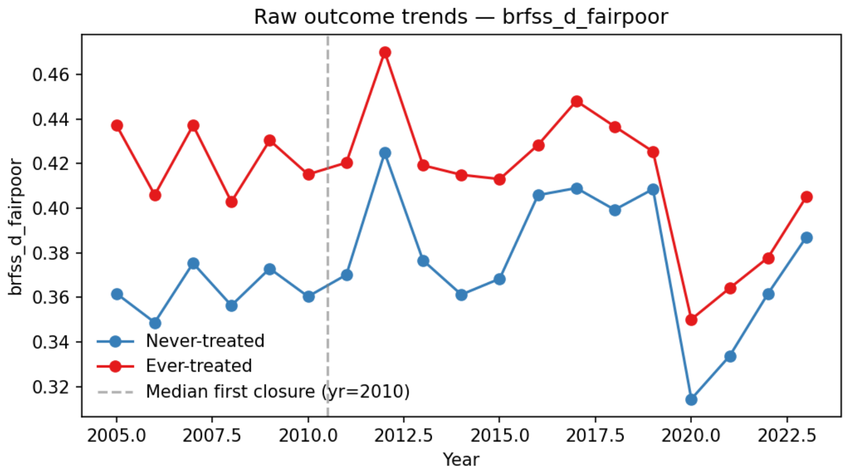

Figure 17: Fig1 Raw Trends Brfss D Fairpoor

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

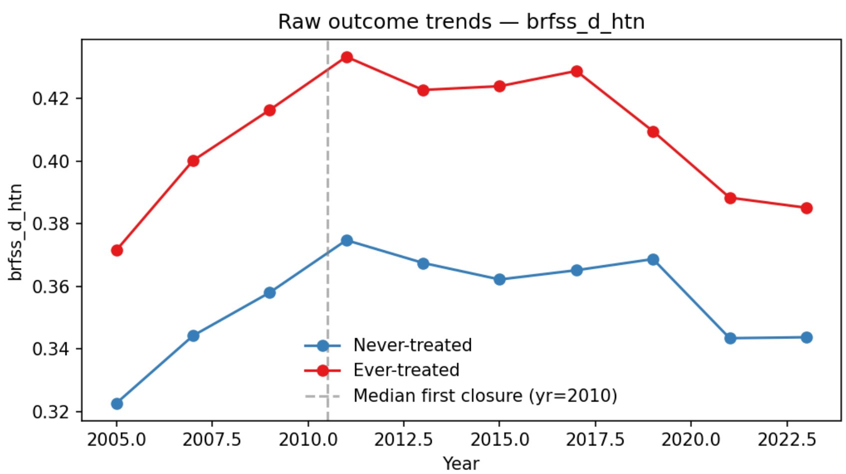

Figure 18: Fig1 Raw Trends Brfss D Htn

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

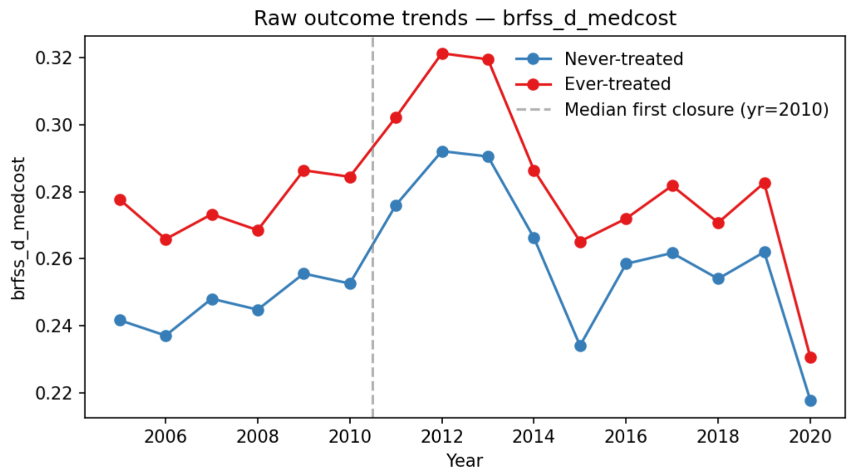

Figure 19: Fig1 Raw Trends Brfss D Medcost

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

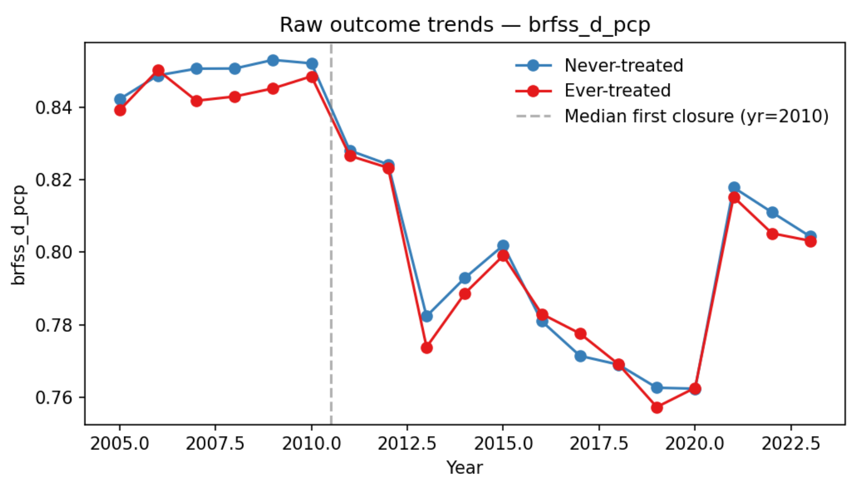

Figure 20: Fig1 Raw Trends Brfss D Pcp

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

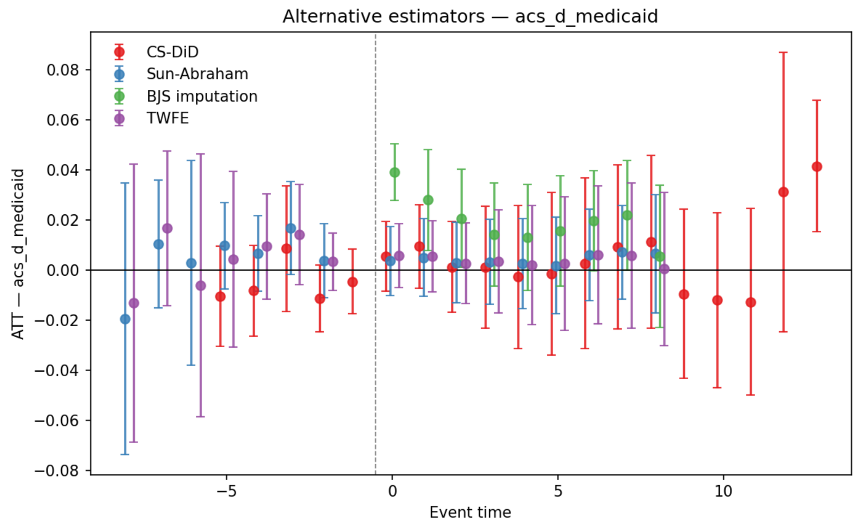

Figure 21: Fig Method Comparison Acs D Medicaid

Note: This figure reports a robustness or sensitivity check for the comparison ACS D Medicaid. It shows how the main estimate changes under alternative assumptions, samples, or specifications.

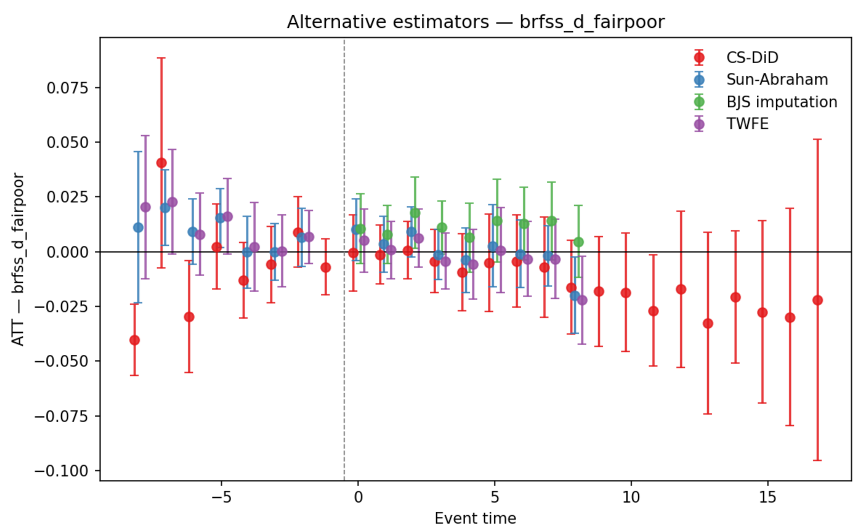

Figure 22: Fig Method Comparison Brfss D Fairpoor

Note: This figure reports a robustness or sensitivity check for the comparison BRFSS D Fairpoor. It shows how the main estimate changes under alternative assumptions, samples, or specifications.

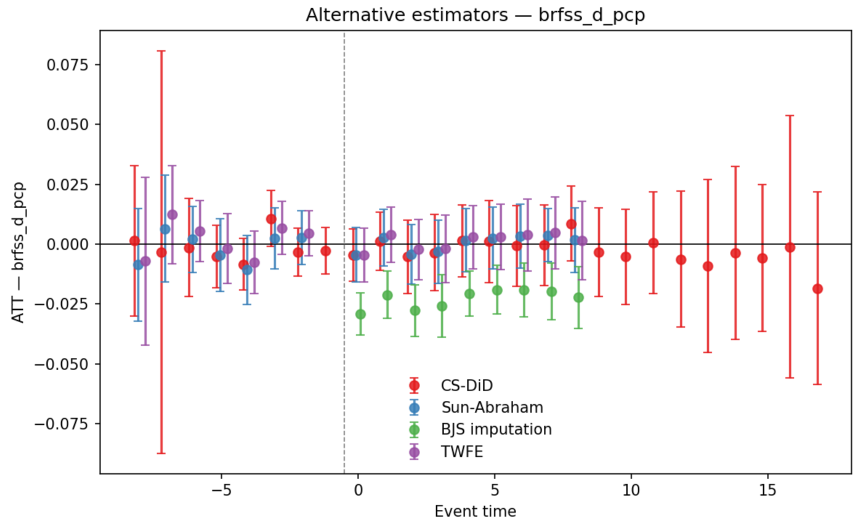

Figure 23: Fig Method Comparison Brfss D Pcp

Note: This figure reports a robustness or sensitivity check for the comparison BRFSS D Pcp. It shows how the main estimate changes under alternative assumptions, samples, or specifications.

imental evidence from SNAP. *Quarterly Journal of Economics*, 134(3), 1505–1556.

Fox, A. M., Stazyk, E. C., & Feng, W. (2020). Administrative easing: Rule reduction and Medicaid enrollment. *Public Administration Review*, 80(1), 104–117. https://doi.org/10.1111/puar.13131

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277.

Herd, P., & Moynihan, D. P. (2018). *Administrative Burden: Policymaking by Other Means.* Russell Sage Foundation.

Homonoff, T., & Somerville, J. (2021). Program recertification costs: Evidence from SNAP. *American Economic Journal: Economic Policy*, 13(4), 271–298. https://doi.org/10.1257/pol.20190272

Medicaid and CHIP Payment and Access Commission. (2023). *MACStats: Medicaid and CHIP Data Book — Exhibit 37: Medicaid Income Eligibility Levels for Individuals Age 65 and Older and Persons with Disabilities by State, 2023.*

Miller, S., Johnson, N., & Wherry, L. R. (2021). Medicaid and mortality: New evidence from linked survey and administrative data. *Quarterly Journal of Economics*, 136(3), 1783–1829.

Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends.

Review of Economic Studies, 90(5), 2555–2591. https://doi.org/10.1093/restud/rdad018

Rupp, K., & Riley, G. F. (2016). State Medicaid eligibility and enrollment policies and rates of Medicaid participation among disabled Supplemental Security Income recipients. *Social Security Bulletin*, 76(3), 17–40.

Sommers, B. D., Baicker, K., & Epstein, A. M. (2012). Mortality and access to care among adults after state Medicaid expansions. *New England Journal of Medicine*, 367(11), 1025–1034. https://doi.org/10.1056/NEJMsa1202099

Sommers, B. D., Goldman, A. L., Blendon, R. J., Orav, E. J., & Epstein, A. M. (2019). Medicaid work requirements — Results from the first year in Arkansas. *New England Journal of Medicine*, 381(11), 1073–1082. https://doi.org/10.1056/NEJMsr1901772

Sommers, B. D., Maylone, B., Blendon, R. J., Orav, E. J., & Epstein, A. M. (2017). Three-year impacts of the Affordable Care Act: Improved medical care and health among low-income adults. *Health Affairs*, 36(6), 1119–1128. https://doi.org/10.1377/hlthaff.2017.0293

Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.

Urban Institute. (2025). Social Security office closures will hurt rural and tribal communities. Urban Wire blog post.

Urban Institute. (2025). Closing any Social Security field office would double the time it takes the average person to drive to their nearest office. [REPORT/UNVERIFIED]

Wherry, L. R., & Miller, S. (2016). Early coverage, access, utilization, and health effects associated with the Affordable Care Act Medicaid expansions: A quasi-experimental study. *Annals of Internal Medicine*, 164(12), 795–803. https://doi.org/10.7326/M15-2234

Appendix — Supplementary Tables and Figures

This appendix provides supplementary tables and figures for the main manuscript. It is built from the live robustness outputs in `analysis/robustness/` and mirrors the supplementary materials referenced in the AEJ submission manuscript.

Appendix Tables

Table A1. Balance and institutional-composition diagnostics for the 1634 / 209(b) partition

Embedded table: Balance 1634 209B

group	[illegible]
1634	[illegible]
209b	[illegible]
criteria	[illegible]
diff_1634_minus_209b	[illegible] 0.265 0.171 [illegible] 0.141 [illegible] 0.013 0.043 [illegible]

Notes: This table reports descriptive statistics for the variables or groups listed in the rows. Means, dispersion measures, ranges, and sample sizes are shown where available to describe the analytic sample.

This table documents the main observable differences between the automatic-linkage 1634 states and the 209(b) states, including the Medicaid-expansion imbalance discussed in the manuscript.

Table A2. Cohort support and cohort-restricted ATT sensitivity

Embedded table: Cohort Support Grid

first_treat_year	event_time	n_treated
2005	3	5
2005	4	5
2005	5	5
2005	6	5
2005	7	5
2005	8	5
2005	9	5
2005	10	5
2005	11	5
2005	12	5
2005	13	5
2005	14	5
2005	15	5
2005	16	5
2005	17	5
2005	18	5
2006	2	2
2006	3	2
2006	4	2
2006	5	2
2006	6	2
2006	7	2
2006	8	2
2006	9	2
2006	10	2
2006	11	2
2006	12	2
2006	13	2
2006	14	2
2006	15	2
2006	16	2
2006	17	2
2007	1	5

first_treat_year	event_time	n_treated
2007	2	5
2007	3	5
2007	4	5
2007	5	5
2007	6	5
2007	7	5
2007	8	5
2007	9	5
2007	10	5
2007	11	5
2007	12	5
2007	13	5
2007	14	4
2007	15	5
2007	16	4
2008	0	1
2008	1	1
2008	2	1
2008	3	1
2008	4	1
2008	5	1
2008	6	1
2008	7	1
2008	8	1
2008	9	1
2008	10	1
2008	11	1
2008	12	1
2008	13	1
2008	14	1
2008	15	1
2010	-2	1
2010	-1	1
2010	0	1
2010	1	1
2010	2	1
2010	3	1
2010	4	1
2010	5	1
2010	6	1
2010	7	1
2010	8	1
2010	9	1
2010	10	1
2010	11	1
2010	12	1
2010	13	1
2011	-3	6
2011	-2	6
2011	-1	6
2011	0	6
2011	1	6
2011	2	6
2011	3	6
2011	4	6
2011	5	6

first_treat_year	event_time	n_treated
2011	6	6
2011	7	6
2011	8	5
2011	9	6
2011	10	6
2011	11	6
2011	12	6
2012	-4	9
2012	-3	9
2012	-2	9
2012	-1	9
2012	0	9
2012	1	9
2012	2	9
2012	3	9
2012	4	9
2012	5	9
2012	6	9
2012	7	9
2012	8	9
2012	9	9
2012	10	9
2012	11	9
2013	-5	1
2013	-4	1
2013	-3	1
2013	-2	1
2013	-1	1
2013	0	1
2013	1	1
2013	2	1
2013	3	1
2013	4	1
2013	5	1
2013	6	1
2013	7	1
2013	8	1
2013	9	1
2013	10	1
2014	-6	2
2014	-5	2
2014	-4	2
2014	-3	2
2014	-2	2
2014	-1	2
2014	0	2
2014	1	2
2014	2	2
2014	3	2
2014	4	2
2014	5	2
2014	6	2
2014	7	2
2014	8	2
2014	9	1
2025	-17	3

first_treat_year	event_time	n_treated
2025	-16	3
2025	-15	3
2025	-14	3
2025	-13	3
2025	-12	3
2025	-11	3
2025	-10	3
2025	-9	3
2025	-8	3
2025	-7	3
2025	-6	3
2025	-5	3
2025	-4	3
2025	-3	3
2025	-2	3

Notes: This table reports dynamic or horizon-specific estimates. Rows correspond to event times, horizons, or diagnostic tests, with uncertainty and sample information shown where available.

Embedded table: Cohort Restricted Att

label	subset	att	se	n_states	n_treated_states	error	att_pp	se_pp	lci_pp	uci_pp
unrestricted	all	0.002	0.012	51	35		0.166	1.158	-2.105	2.436
unrestricted	1634	-0.014	0.015	34	24		-1.423	1.456	-4.277	1.432
unrestricted	2009b	0.027	0.01	10	7		2.732	1.013	0.746	4.718
cohort_≤2021	all	0.005	0.012	48	32		0.528	1.206	-1.836	2.891
cohort_≤2021	1634	-0.011	0.016	31	21		-1.091	1.588	-4.204	2.021
cohort_≤2021	2009b	0.027	0.01	10	7		2.732	1.013	0.746	4.718
early_2005_2010	all	0.023	0.008	30	14		2.259	0.794	0.703	3.814
early_2005_2010	1634					need at least one array to stack				
early_2005_2010	2009b	0.042	0.002	7	4		4.185	0.15	3.891	4.479
late_2011_2014	all	0.004	0.012	34	18		0.416	1.249	-2.033	2.865
late_2011_2014	1634	-0.011	0.016	21	11		-1.091	1.588	-4.204	2.021
late_2011_2014	2009b	0.022	0.013	6	3		2.151	1.267	-0.332	4.634

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 files show which closure cohorts actually identify the subgroup contrast and how the estimates move when the sample is restricted by cohort support.

Table A3. ACA-confounder robustness

Embedded table: Aca Controlled Att

subset	att	se	spec	att_pp	se_pp	lci_pp	uci_pp
all	0.007	0.01	aca_controlled	0.734	1.018	-1.262	2.73
1634	-0.008	0.013	aca_controlled	-0.821	1.293	-3.356	1.713
209b	0.033	0.009	aca_controlled	3.314	0.899	1.552	5.075

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: Pre2014 Att

subset	att	se	spec	att_pp	se_pp	lci_pp	uci_pp
all	0.008	0.007	pre2014	0.804	0.698	-0.565	2.172
1634	0.006	0.01	pre2014	0.551	0.951	-1.312	2.414
209b	0.001	0.009	pre2014	0.082	0.902	-1.687	1.851

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 implement the two main ACA-confounder checks described in the paper: residualizing on expansion timing and restricting to the pre-2014 period.

Table A4. Formal 1634-vs-209(b) contrast tests

Embedded table: Triple Difference

spec	estimate_pp	se_pp	lci_pp	uci_pp	p_value	inference	G_clusters
twfe_triple_diff	-2.942	2.075	-7.846	1.664	0.224	CGM wild bootstrap (Webb, B=1999)	44
ibragimov_muller	-3.569	1.951	-8.206	1.068	0.111	Ibragimov–Müller group-level t	15
csdid_contrast	-4.155		-7.632	-0.678		CS-DiD ATT(1634) − ATT(209b), delta method	44
att_1634	-1.423	1.456	-4.277	1.432		CS-DiD 1634 subgroup	34
att_209b	2.732	1.013	0.746	4.718		CS-DiD 209(b) subgroup	10

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: Triple Difference Onesided

spec	estimate_pp	se_pp	p_two_sided	p_one_sided	inference
twfe_triple_diff_cgm	-2.942	2.075	0.224	0.109	CGM wild cluster bootstrap, Webb (B=1999); one-sided = frac(boot_t <= t_hat) directly from bootstrap dist
ibragimov_muller	-3.569	1.951	0.111	0.056	Ibragimov-Muller group-level t, Welch dof=6.83; one-sided via st.t.cdf
csdid_delta_method	-3.548	1.818	0.051	0.026	CS-DiD subgroup ATT delta method, normal approximation; one-sided via st.norm.cdf

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 are the load-bearing contrast tables for the mechanism claim, including the pre-specified one-sided p-values and the alternative small-cluster inference procedures.

Table A5. Subgroup-specific HonestDiD bounds

Embedded table: Honestdid 1634

event_time	M	att_e	se_e	delta_pre	pre_worst	lci	uci	crosses_zero	subgroup
0	0.25	0.278	0.998	2.196	-5	-2.226	2.783	True	1634
0	0.5	0.278	0.998	2.196	-5	-2.775	3.332	True	1634
0	1	0.278	0.998	2.196	-5	-3.873	4.43	True	1634
0	1.5	0.278	0.998	2.196	-5	-4.971	5.528	True	1634
0	2	0.278	0.998	2.196	-5	-6.069	6.626	True	1634
1	0.25	1.941	1.091	2.196	-5	-0.747	4.628	True	1634
1	0.5	1.941	1.091	2.196	-5	-1.296	5.177	True	1634
1	1	1.941	1.091	2.196	-5	-2.394	6.275	True	1634
1	1.5	1.941	1.091	2.196	-5	-3.492	7.373	True	1634
1	2	1.941	1.091	2.196	-5	-4.59	8.471	True	1634
2	0.25	-0.319	1.09	2.196	-5	-3.005	2.366	True	1634
2	0.5	-0.319	1.09	2.196	-5	-3.554	2.915	True	1634
2	1	-0.319	1.09	2.196	-5	-4.652	4.013	True	1634
2	1.5	-0.319	1.09	2.196	-5	-5.75	5.111	True	1634
2	2	-0.319	1.09	2.196	-5	-6.848	6.209	True	1634

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: Honestdid 209B

event_time	M	att_e	se_e	delta_pre	pre_worst	lci	uci	crosses_zero	subgroup
0	0.25	-0.466	1.322	2.476	-4	-3.676	2.744	True	209b
0	0.5	-0.466	1.322	2.476	-4	-4.295	3.363	True	209b
0	1	-0.466	1.322	2.476	-4	-5.532	4.601	True	209b
0	1.5	-0.466	1.322	2.476	-4	-6.771	5.838	True	209b
0	2	-0.466	1.322	2.476	-4	-8.008	7.076	True	209b
1	0.25	-2.114	0.847	2.476	-4	-4.392	0.164	True	209b
1	0.5	-2.114	0.847	2.476	-4	-5.011	0.783	True	209b
1	1	-2.114	0.847	2.476	-4	-6.249	2.021	True	209b
1	1.5	-2.114	0.847	2.476	-4	-7.487	3.259	True	209b
1	2	-2.114	0.847	2.476	-4	-8.725	4.497	True	209b
2	0.25	0.448	1.306	2.476	-4	-2.732	3.627	True	209b
2	0.5	0.448	1.306	2.476	-4	-3.351	4.246	True	209b
2	1	0.448	1.306	2.476	-4	-4.588	5.484	True	209b
2	1.5	0.448	1.306	2.476	-4	-5.826	6.722	True	209b
2	2	0.448	1.306	2.476	-4	-7.064	7.96	True	209b

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 files contain the subgroup-specific Rambachan-Roth sensitivity calculations for early event times in the 1634 and 209(b) cells.

Table A6. Event-time subgroup contrasts

Embedded table: Eventtime 1634 Vs 209B

event_time	att	se	lci	uci	subgroup
-5	-0.006	0.012	-0.029	0.017	contrast (1634 - 209b)
-4	0.016	0.013	-0.01	0.042	contrast (1634 - 209b)
-3	0.036	0.025	-0.013	0.086	contrast (1634 - 209b)
-2	-0.018	0.018	-0.052	0.016	contrast (1634 - 209b)
-1	-0.024	0.023	-0.068	0.021	contrast (1634 - 209b)
0	0.007	0.017	-0.025	0.04	contrast (1634 - 209b)
1	0.041	0.014	0.013	0.068	contrast (1634 - 209b)
2	-0.008	0.017	-0.041	0.026	contrast (1634 - 209b)
3	-0.01	0.021	-0.051	0.032	contrast (1634 - 209b)
4	-0.027	0.028	-0.082	0.028	contrast (1634 - 209b)
5	-0.059	0.029	-0.115	-0.002	contrast (1634 - 209b)
6	-0.059	0.027	-0.112	-0.007	contrast (1634 - 209b)
7	-0.045	0.023	-0.091	0.001	contrast (1634 - 209b)
8	-0.049	0.023	-0.094	-0.003	contrast (1634 - 209b)
9	-0.086	0.025	-0.136	-0.037	contrast (1634 - 209b)
10	-0.099	0.024	-0.146	-0.053	contrast (1634 - 209b)
11	-0.103	0.025	-0.152	-0.054	contrast (1634 - 209b)
12	-0.07	0.043	-0.155	0.015	contrast (1634 - 209b)
13					contrast (1634 - 209b)

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: Eventtime Subgroups

event_time	att	se	lci	uci	subgroup
-5	-0.022	0.007	-0.035	-0.009	1634
-4	-0.009	0.011	-0.03	0.012	1634
-3	0.02	0.015	-0.008	0.049	1634
-2	-0.022	0.008	-0.038	-0.006	1634
-1	-0.007	0.009	-0.025	0.011	1634
0	0.003	0.01	-0.017	0.022	1634
1	0.019	0.011	-0.002	0.041	1634
2	-0.003	0.011	-0.025	0.018	1634
3	-0.006	0.015	-0.034	0.023	1634
4	-0.015	0.019	-0.052	0.023	1634
5	-0.024	0.02	-0.063	0.015	1634
6	-0.013	0.021	-0.054	0.027	1634
7	-0.004	0.022	-0.048	0.04	1634
8	-0.002	0.022	-0.045	0.042	1634
9	-0.04	0.02	-0.079	-0.002	1634
10	-0.045	0.021	-0.087	-0.003	1634
11	-0.05	0.019	-0.088	-0.012	1634
12	-0.011	0.04	-0.089	0.067	1634
-5	-0.016	0.01	-0.035	0.003	209b
-4	-0.025	0.008	-0.04	-0.01	209b
-3	-0.016	0.021	-0.057	0.025	209b
-2	-0.004	0.016	-0.034	0.027	209b
-1	0.017	0.021	-0.024	0.057	209b
0	-0.005	0.013	-0.031	0.021	209b
1	-0.021	0.008	-0.038	-0.005	209b
2	0.004	0.013	-0.021	0.03	209b

event_time	att	se	lci	uci	subgroup
3	0.004	0.015	-0.026	0.034	209b
4	0.012	0.02	-0.028	0.052	209b
5	0.034	0.021	-0.007	0.076	209b
6	0.046	0.017	0.012	0.079	209b
7	0.041	0.007	0.027	0.054	209b
8	0.047	0.007	0.033	0.061	209b
9	0.046	0.016	0.015	0.077	209b
10	0.055	0.01	0.034	0.075	209b
11	0.053	0.016	0.022	0.084	209b
12	0.059	0.018	0.025	0.094	209b
13	0.063	0.019	0.025	0.101	209b

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 underlie the event-time contrast figure that visualizes when the subgroup divergence emerges.

Table A7. Mechanism-inconsistent placebos and BRFSS multiple-testing correction

Embedded table: Mechanism Placebos

outcome	subset	att	se	label	att_pp	se_pp	lci_pp	uci_pp
acs_d_medicaid	all	0.002	0.012	Medicaid, disabled 18-64 (FOCAL)	0.166	1.158	-2.105	2.436
acs_d_medicaid	1634	-0.014	0.015	Medicaid, disabled 18-64 (FOCAL)	-1.423	1.456	-4.277	1.432
acs_d_medicaid	209b	0.027	0.01	Medicaid, disabled 18-64 (FOCAL)	2.732	1.013	0.746	4.718
acs_nd_medicaid	all	-0.007	0.008	Medicaid, non-disabled 18-64 (PLACEBO)	-0.733	0.841	-2.382	0.916
acs_nd_medicaid	1634	-0.01	0.011	Medicaid, non-disabled 18-64 (PLACEBO)	-0.974	1.082	-3.094	1.147

outcome	subset	att	se	label	att_pp	se_pp	lci_pp	uci_pp
acs_nd_medicaid	2009b	-0.011	0.01	Medicaid, non-disabled 18-64 (PLACEBO)	-1.106	1.03	-3.125	0.913
acs_d_empcov	all	0.004	0.006	ESI, disabled 18-64 (PLACEBO)	0.379	0.601	-0.798	1.557
acs_d_empcov	1634	-0.001	0.008	ESI, disabled 18-64 (PLACEBO)	-0.09	0.777	-1.613	1.433
acs_d_empcov	2009b	0.008	0.007	ESI, disabled 18-64 (PLACEBO)	0.805	0.742	-0.649	2.258
acs_d_medicare	all	0.006	0.005	Medicare, disabled 18-64 (PLACEBO)	0.587	0.487	-0.367	1.542
acs_d_medicare	1634	0.008	0.008	Medicare, disabled 18-64 (PLACEBO)	0.772	0.793	-0.783	2.327
acs_d_medicare	2009b	0.003	0.006	Medicare, disabled 18-64 (PLACEBO)	0.289	0.621	-0.928	1.505
acs_nd_anycov	all	-0.008	0.008	Any cov, non-disabled 18-64 (PLACEBO)	-0.786	0.794	-2.342	0.771
acs_nd_anycov	1634	-0.007	0.011	Any cov, non-disabled 18-64 (PLACEBO)	-0.742	1.066	-2.831	1.347
acs_nd_anycov	2009b	-0.029	0.022	Any cov, non-disabled 18-64 (PLACEBO)	-2.895	2.154	-7.117	1.326

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: Romano Wolf Brfss

subset	outcome	beta_pp	se_pp	t_abs	pval_single	G_clusters	pval_rw_adj
1634	brfss_d_fairpoor	-1.633	1.074	1.543	0.138	34	0.471
1634	brfss_d_pcp	0.227	0.677	0.341	0.739	34	0.918
1634	brfss_d_medcost	-0.243	0.938	0.263	0.797	34	0.918
1634	brfss_d_checkup	-0.825	1.359	0.616	0.548	34	0.885
1634	brfss_d_diab	-0.845	0.301	2.849	0.008	34	0.057
1634	brfss_d_htn	-0.587	0.575	1.036	0.315	34	0.755
209b	brfss_d_fairpoor	-0.108	1.283	0.089	0.935	10	0.999
209b	brfss_d_pcp	-1.159	0.52	2.351	0.053	10	0.238
209b	brfss_d_medcost	0.008	0.793	0.01	0.993	10	1
209b	brfss_d_checkup	0.068	1.801	0.04	0.971	10	1
209b	brfss_d_diab	-0.551	0.418	1.391	0.22	10	0.662
209b	brfss_d_htn	0.192	0.945	0.215	0.843	10	0.999

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 first table collects the placebo outcomes that should be null if the SSI-Medicaid linkage is the operative mechanism. The second applies family-wise correction across the BRFSS subgroup outcomes.

Table A8. Public SSI anchor status note

Embedded table: Public Ssi First Stage

subset	csdid_att	csdid_se	twfe_coef	twfe_se	twfe_lci	twfe_uci	n_obs
all	-0.031	0.024	-0.013	0.024	-0.06	0.033	969
1634	-0.037	0.029	-0.017	0.028	-0.072	0.037	646
209b	0.02	0.035	0.02	0.04	-0.059	0.098	190

Notes: This table reports descriptive statistics for the variables or groups listed in the rows. Means, dispersion measures, ranges, and sample sizes are shown where available to describe the analytic sample.

The live manuscript treats the public SSI panel as supplementary only. The note file records whether the Wayback-dependent assembly step succeeded in the reproducibility run.

Appendix Figures

Figure A1. Raw outcome trends, treated versus never-treated

Representative panel from the raw-trends suite. Related outcome panels remain in `analysis/figures/fig1_raw_trends_*.png`.

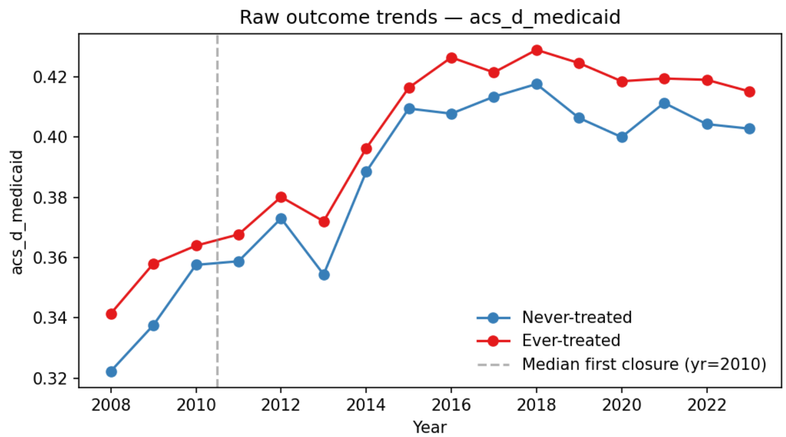

Figure 24: Fig1 Raw Trends Acs D Medicaid

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

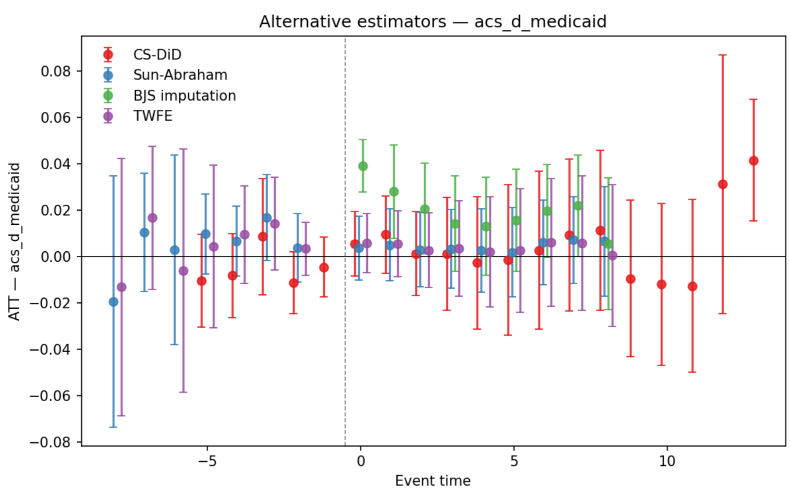

Figure 25: Fig Method Comparison Acs D Medicaid

Note: This figure reports a robustness or sensitivity check for the comparison ACS D Medicaid. It shows how the main estimate changes under alternative assumptions, samples, or specifications.

Figure A2. Method-comparison event-study

Representative panel from the estimator-comparison set. Related BRFSS panels remain in `analysis/robustness/fig_method_comparison_*.png`.

Figure A3. Placebo treatment-date and placebo-outcome panel

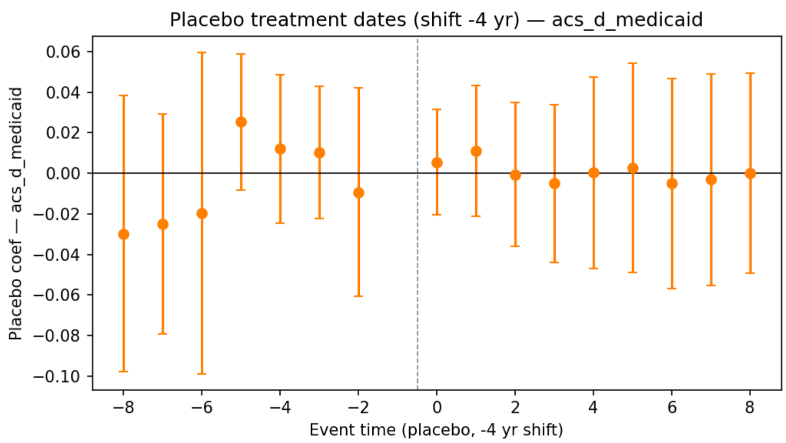

Figure 26: Fig Placebo Dates Acs D Medicaid

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

Representative placebo panel. Related BRFSS placebo-date panels remain in `analysis/robustness/fig_placebo_dates_*.png`.

Figure A4. HonestDiD pooled sensitivity

Representative pooled HonestDiD panel. Related BRFSS HonestDiD panels remain in `analysis/robustness/fig_honestdid_*.png`.

Figure A5. 1634-vs-209(b) forest plot

This is the core subgroup-comparison figure carried into the main paper.

Figure A6. Triple-difference contrast forest

This figure packages the alternative contrast estimators and inference procedures on a common scale.

Figure A7. Mechanism-inconsistent placebo forest

This figure visualizes the mechanism-specific falsification logic.

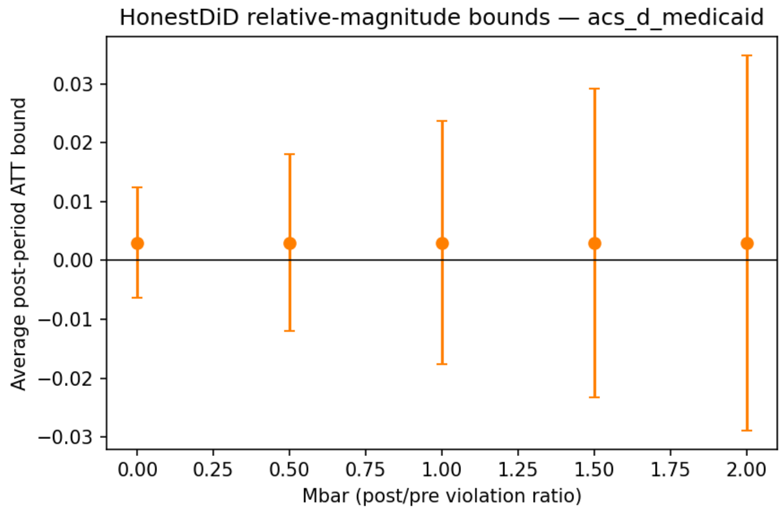

Figure 27: Fig Honestdid Acs D Medicaid

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

Figure A8. ACA-confounder robustness forest

Comparison of the unadjusted, ACA-controlled, and pre-2014-only specifications.

Figure A9. Cohort-restricted sensitivity forest

This figure shows how the subgroup pattern depends on closure-cohort support.

Figure A10. Subgroup-specific HonestDiD bounds

This figure summarizes the subgroup-specific sensitivity calculations for the early post-treatment event times.

Figure A11. Event-time subgroup contrast

This is the dynamic visual version of the paper's formal 1634-vs-209(b) contrast.

End of appendix.

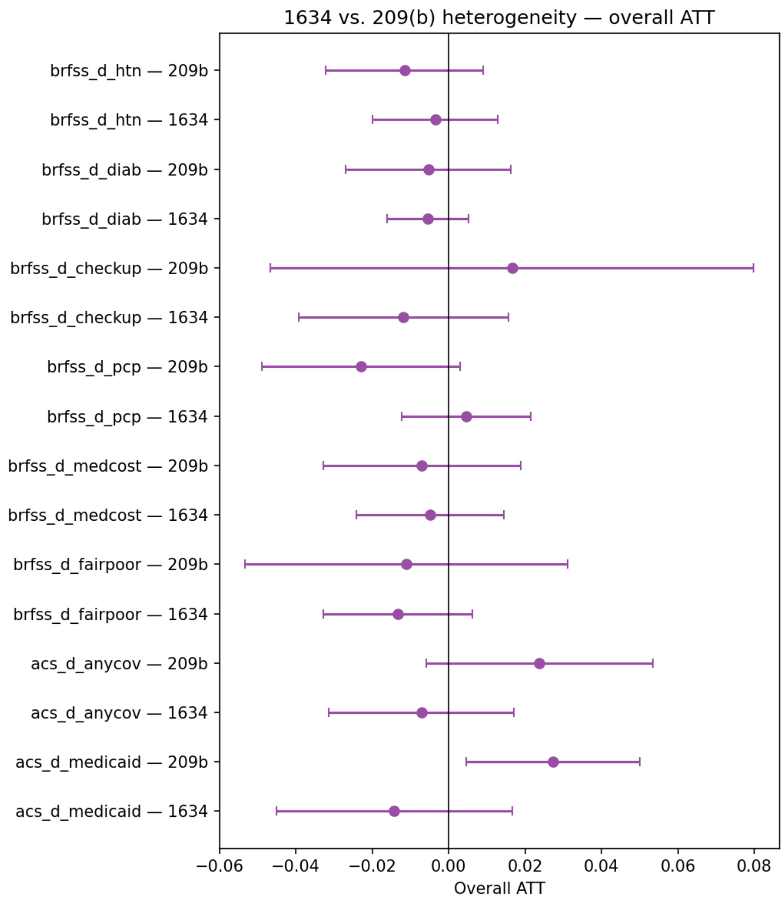

Figure 28: Fig5 Forest 1634 Vs 209B

Note: This figure compares estimates across groups or specifications for the forest 1634 Vs 209B. It is intended to make effect heterogeneity and subgroup precision easier to assess.

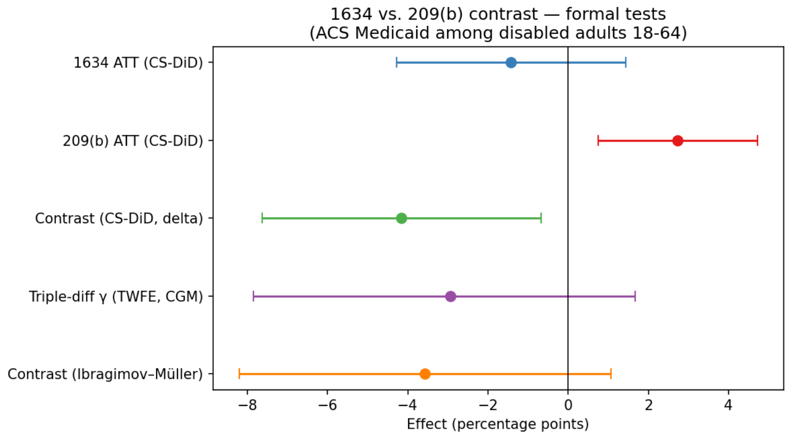

Figure 29: Triple Diff Forest

Note: This figure reports a robustness or sensitivity check for the triple Diff Forest. It shows how the main estimate changes under alternative assumptions, samples, or specifications.

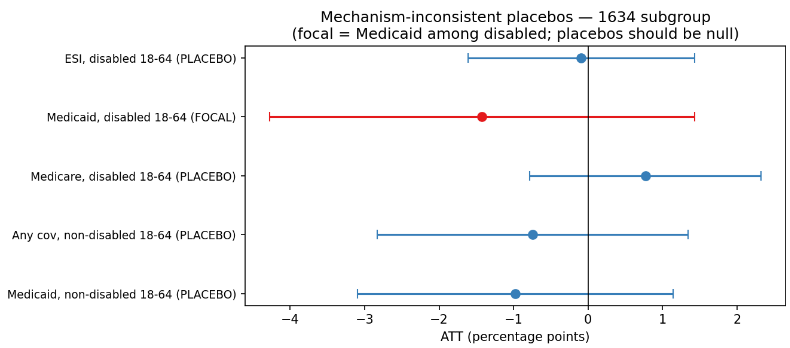

Figure 30: Mechanism Placebo Forest

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

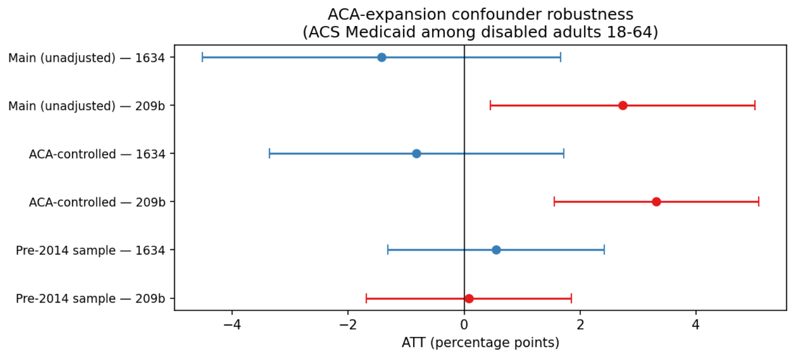

Figure 31: Aca Robustness Forest

Note: This figure reports a robustness or sensitivity check for the aca Robustness Forest. It shows how the main estimate changes under alternative assumptions, samples, or specifications.

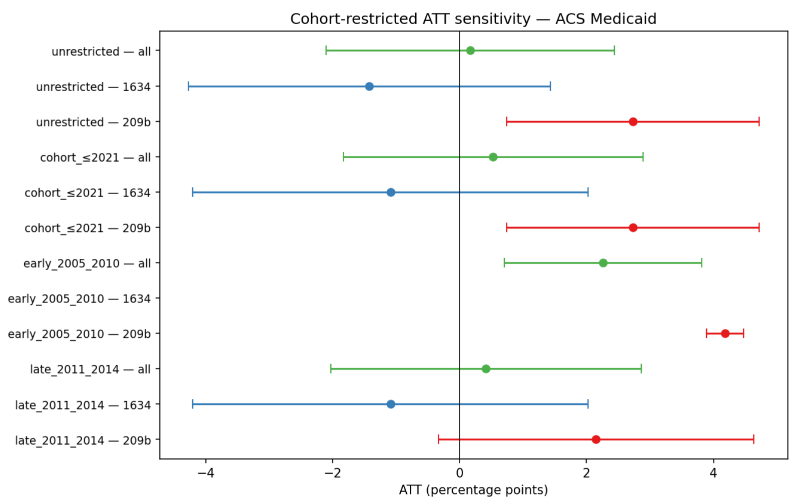

Figure 32: Cohort Restricted Forest

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

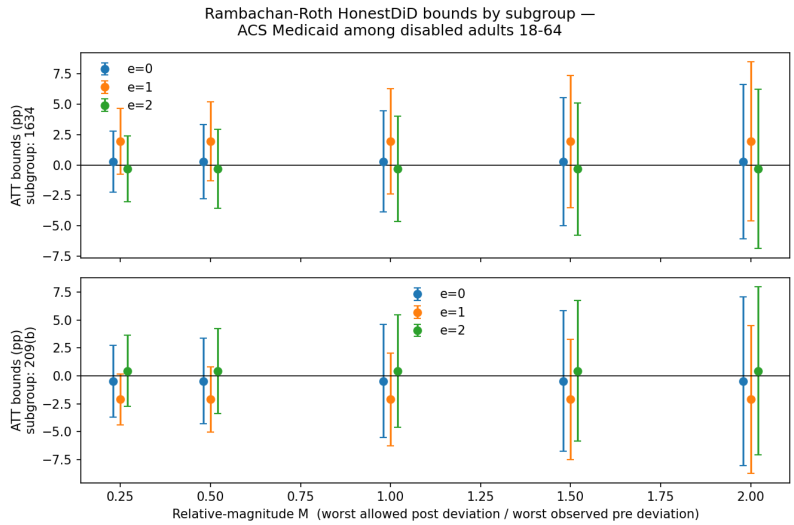

Figure 33: Honestdid Subgroups

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

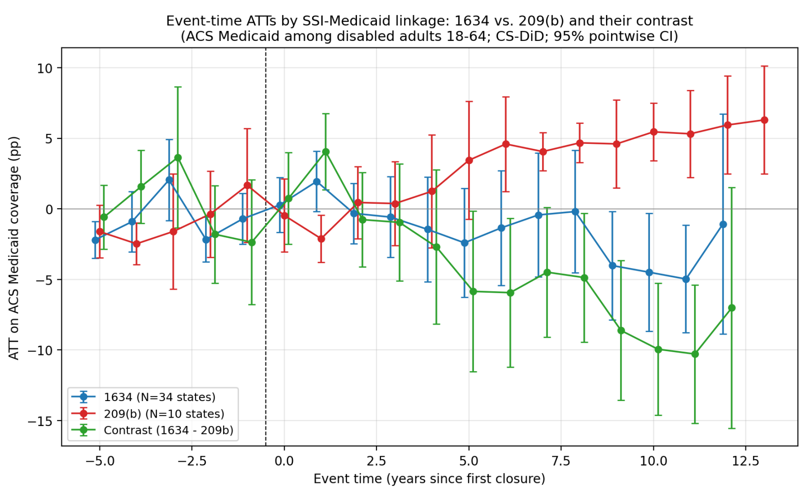

Figure 34: Eventtime 1634 Vs 209B

Note: This figure plots event-time estimates for the eventtime 1634 Vs 209B. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.