A Persistent Two-Decade 209(b) Penalty: Medicaid–SSI Linkage Choices and the Long-Run Health Consequences of 1972 SSI Federalization

Abstract

Background. The 1972 Social Security Amendments federalized cash assistance for aged, blind, and disabled adults through the Supplemental Security Income (SSI) program. Beginning in January 1974, most states automatically extended Medicaid to SSI recipients, but 11 states elected the §1902(f) option — commonly called "209(b)" — under which states retained pre-1972 Medicaid eligibility rules, typically more restrictive than the federal SSI standard. Whether this 50-year-old linkage choice continues to influence population health and program participation has not been quantified.

Methods. We assembled a 1968–2020 county-year panel of 483 contiguous-county pairs straddling 209(b) and standard-SSI state borders (51,198 observations on the main sample), drawing on CDC WONDER mortality, IPUMS Current Population Survey (CPS) coverage and take-up measures, and a newly-constructed Social Security Administration (SSA) county-year SSI uptake panel covering fiscal years 1998–2024 (estimating sample 1998–2020 to match WONDER population denominators; 77,331 matched county-years). We estimated border-pair difference-in-differences (DiD) models with county and year fixed effects, augmented by within-pair-year (Dube–Lester–Reich) specifications. CPS and SSA gaps are reported as post-treatment cross-sectional border-pair contrasts (year fixed effects only), because the relevant survey variables begin in 1988–1989 and the SSA county-year publication begins in 1998 — both post-treatment, so a pre/post DiD is not identified on

those arms. A geographic regression-discontinuity (RD) design is reported as exploratory appendix evidence only.

Results. Across the full 1968–2020 window, the static border-pair difference in all-cause mortality on 209(b)-side counties post-1974 was +0.205 deaths per 1,000 residents annually (95% CI 0.030 to 0.380, $p = 0.022$ under pair-clustered inference, $N = 39{,}980$). This static estimate is fragile: under county-level clustering the same coefficient has $p = 0.21$; under state-level clustering $p = 0.44$; and on a unique-county-year sample (collapsing duplicate observations from counties belonging to multiple border pairs) the point estimate shrinks to +0.145 and is not statistically significant under any clustering choice. We therefore interpret the static mortality estimate as suggestive rather than as a sharp causal point estimate, and anchor the substantive interpretation on the dynamic event study, which shows a sign-shift: short-run negative coefficients ($t = +1$: –0.41 deaths/1,000; $t = +4$: –0.21) give way to long-run positive divergence ($t \geq +20$: +0.262, SE 0.118). A Honest-DiD sensitivity analysis confirms that the lower bound of the static effect crosses zero at conventional parallel-trends violations. A pre-1972 placebo returned null effects ($\beta =$ –0.008 log; –0.115 level). In CPS data from 1988 onward, working-age disabled adults in 209(b) states were 8.9 percentage points less likely to report Medicaid coverage ($p < 0.001$) and 2.9 percentage points less likely to report SSI receipt ($p = 0.006$); these are post-treatment cross-sectional gaps, not DiD estimates. In SSA county-level administrative data (1998–2020 estimating sample), log SSI recipients per 1,000 was 0.187 lower on the 209(b) side of the same border pairs (SE 0.042, $p < 0.001$, $N = 17{,}330$), corresponding to a roughly 17% descriptive gap and 4.55 fewer recipients per 1,000 in levels; this is also a post-treatment cross-sectional border-pair contrast, not a DiD.

Conclusions. Five decades after SSI federalization, descriptive evidence across three independent data systems consistently shows lower disabled-adult Medicaid coverage and lower SSI participation in §1902(f) states. The long-run mortality DiD is in the same direction but is fragile to clustering and weighting choices; we therefore frame this paper as descriptive triangulation across mortality, survey-based coverage, and administrative program uptake rather than as a sharp causal mortality estimate. 209(b) remains a live structural lever for 11 states; whether its long-run mortality consequences are causal warrants further work with stronger identification (e.g., state-status switches, or a research-grade county mortality denominator extension to 2024).

---

1. Introduction

The Social Security Amendments of 1972 federalized cash assistance for aged, blind, and disabled Americans. Effective January 1, 1974, the new Supplemental Security Income (SSI) program replaced three categorical state-administered programs — Old Age Assistance, Aid to the Blind, and Aid to the Permanently and Totally Disabled — with a uniform federal benefit and a uniform federal disability standard [@ball1973_ssaamend]. The Amendments simultaneously rewrote the connection between cash assistance and medical assistance: in most states, SSI eligibility now automatically conferred Medicaid eligibility. But Congress also wrote §1902(f) of the Social Security Act, permitting states whose pre-1972 Medicaid disability rules had been more restrictive than the new federal SSI standard to retain those rules indefinitely. Eleven states exercised that option and have retained it across half a century: Connecticut, Hawaii, Illinois, Indiana, Minnesota, Missouri, New Hampshire, North Dakota, Ohio, Oklahoma, and Virginia [@crs_r46111_2019; @cms_209b_implementation_guide]. These states are now collectively known as “209(b)” states, a label derived from the section of Public Law 92-603 that codified the option.

The §1902(f) option is a quasi-experimental boundary in plain sight. On one side of a 209(b)/standard-SSI state line, an adult who is found federally disabled, files for SSI, and is approved is enrolled in Medicaid automatically. On the other side, the same adult must navigate a separate state Medicaid application, may face more restrictive income disregards and asset limits, and may have a different definition of disability for state Medicaid eligibility than for federal SSI. The friction is administrative: the cash benefit is the same; the disability standard for federal SSI is the same; only the medical-assistance entry point differs. After 50 years, does this administrative wedge still matter for population health?

Three features of the 209(b) variation make it a productive setting for causal inference. First, the policy is sharp in time: every 209(b) state acquired its status in 1974, and every state on the comparison side acquired automatic SSI–Medicaid linkage in 1974, so there is no staggered adoption and no Goodman-Bacon contamination [@goodmanbacon_2021]. Second, the policy is fixed in space: no state in our panel switches 209(b) status across the 53-year window. Third, the policy is binary at the state level, which permits a Dube–Lester–Reich (DLR) [@dube_lester_reich_2010] border-pair design that absorbs region-by-time shocks at the smallest geographic resolution at which the policy varies.

We assemble a county-pair-year panel spanning 1968 through 2020 across 483 contiguous border pairs and 33 states, with three identification arms:

1. **A border-pair DiD on county all-cause mortality** drawn from CDC WONDER (Compressed Mortality and Underlying Cause of Death files), 1968–2020. This is the long-run health outcome arm.
2. **A cross-sectional 209(b) gap on Medicaid coverage and SSI take-up** drawn from IPUMS CPS Annual Social and Economic Supplement, 1988–2020, restricted to working-age (18–64) adults with a work-limiting disability. Because the relevant CPS variables begin in 1988 (`HIMCAID`) and 1989 (`INCSSI`), a clean pre-1972 versus post-1972 DiD is not identified on this arm; we report cross-sectional gaps with year fixed effects, which we interpret as descriptive rather than causal.
3. **A border-pair gap on SSA administrative SSI uptake** drawn from a newly-constructed county-year panel based on SSA's "SSI Recipients by State and County" Excel workbooks (FY 2003–2024) and rendered HTML tables (FY 1998–2002), 77,331 matched county-years. This is the administrative-program participation arm. Because the SSA county series begins in 1998, the DiD is also identified on the within-pair-year cross-section rather than a pre-1972 baseline.

The three arms answer complementary questions: does the policy difference persist in *health* (mortality), in *survey-measured coverage and take-up* (CPS), and in *administrative program rolls* (SSA)? Triangulation across these data sources is the empirical contribution of the paper. The institutional contribution is the demonstration that a policy choice fixed in 1974 has health and program-participation consequences in the present day — that §1902(f) is not a vestigial provision of the Social Security Act but a live state-by-state structural lever.

Our findings are three, with calibrated confidence across arms. First, the static border-pair DiD on all-cause mortality is +0.205 deaths per 1,000 residents per year after 1974 (pair-clustered 95% CI 0.030 to 0.380, $p = 0.022$); the within-pair-year specification returns 0.207 ($p = 0.050$). However, the static coefficient is significant only under pair-level clustering on the pair-expanded panel — under county or state clustering, or under unique-county-year weighting, it does not reject the null at conventional levels (full sensitivity table in §5.5). We therefore present the static mortality estimate as *suggestive*, with the substantive story carried by the event-study sign-shift: short-run *lower* mortality on the 209(b) side, long-run *higher* mortality at horizons of 20+ years. Second, the CPS cross-sectional Medicaid gap for working-age disabled adults is 8.9 percentage points and the cross-sectional SSI take-up gap is 2.9 percentage points, both estimated as post-treatment between-state gaps with year fixed effects (not DiD; the relevant CPS variables begin after the 1974 treatment). Third, in SSA county-level administrative data (1998–2020), the 209(b) side of contiguous border pairs has approximately 17% lower SSI uptake on average — a 4.55 recipient-per-1,000 level gap that has been roughly constant for two decades. This SSA estimate is also a post-treatment cross-sectional border-pair gap, not a DiD: the SSA county-year publication begins in 1998, two decades after treatment. The CPS and SSA arms agree in sign and order of magnitude; CPS misclassification at the bottom of the income distribution plausibly accounts for the survey estimate being attenuated toward zero relative to the administrative estimate [@meyer_mok_sullivan_2015].

The paper's findings position it at the intersection of three literatures. The first is the literature on administrative burden and program take-up, where Currie [@currie_2004_takeup], Heinrich [@heinrich_2016_bite], and Moynihan, Herd, and Harvey [@moynihan_herd_harvey_2015; @herd_moynihan_2018_book] have established that small frictions can cause large coverage gaps; Bhargava and Manoli [@bhargava_manoli_2015] documented this in an IRS context; Fox, Feng, and Reynolds

[@fox_feng_reynolds_2023] traced state-level differences in safety-net participation to administrative-burden design; and Arbogast, Chorniy, and Currie [@arbogast_chorniy_currie_2024] connected administrative burden to Medicaid/CHIP child enrollment. The second is the literature on Medicaid mortality, anchored by Goodman-Bacon's work on the 1965–1970 Medicaid rollout [@goodmanbacon_2018], Sommers, Baicker, and Epstein [@sommers_baicker_epstein_2012] on the 2001–2005 Medicaid expansions, and Miller, Johnson, and Wherry [@miller_johnson_wherry_2021] on the ACA expansions. The third is the literature on disability-program take-up and design, including Currie and Madrian [@currie_madrian_1999], Autor and Duggan [@autor_duggan_2003], Maestas, Mullen, and Strand [@maestas_mullen_strand_2013], Deshpande [@deshpande_2016], Schmidt, Shore-Sheppard, and Watson [@schmidt_shoresheppard_watson_2020], Soni, Burns, Dague, and Simon [@soni_burns_dague_simon_2017], Burns, Dague, Wood, and Kennedy [@burns_dague_wood_kennedy_2022], and Rupp and Riley [@rupp_riley_2016_ssbulletin] on the SSI–Medicaid linkage. The paper closest in spirit is Rupp and Riley, who used SSA administrative data to characterize state-level differences in SSI/Medicaid participation cross-sectionally; we add a long mortality panel, a border-pair design, and a 1998–2024 administrative-uptake panel that the Social Security Bulletin's earlier descriptive cross-section could not support.

The remainder of the paper proceeds as follows. Section 2 details the §1902(f) institutional history and the policy economics of administrative-friction wedges in mixed federal-state programs. Section 3 describes the data, including the construction of the 483-pair border panel and the SSA county-year uptake series. Section 4 sets out the econometric strategy and the relationship among the three arms. Section 5 presents the main results, the event-study dynamics, and the robustness suite. Section 6 discusses the findings, situates them in the literature, and considers a companion paper on the modern SSA office-closure shock. Section 7 concludes.

2. Institutional Background

2.1 The 1972 Social Security Amendments and the §1902(f) option

Public Law 92-603, the Social Security Amendments of 1972, federalized cash assistance for aged, blind, and disabled Americans through the new Supplemental Security Income program, effective January 1, 1974 [@ball1973_ssaamend]. SSI replaced three categorical state-administered programs (Old Age Assistance, Aid to the Blind, and Aid to the Permanently and Totally Disabled) with a uniform federal eligibility standard, a uniform federal benefit, and a uniform federal disability determination. Approximately 3.2 million recipients of the predecessor programs were converted to SSI on the program's first day; SSI rolls reached 4 million by the end of 1974 [@ball1973_ssaamend].

The Amendments simultaneously rewrote the connection between cash and medical assistance. In standard-SSI states (often called "1634 states" after the section of the Act that codified them, or "automatic linkage" states), SSI eligibility automatically conferred Medicaid eligibility — that is, the state's Medicaid agency accepted the federal SSI determination as conclusive for the categorically-needy Medicaid pathway. But §1902(f) of the Act, added by the 1972 amendments, gave states a choice. A state whose pre-1972 Medicaid disability eligibility rules had been *more restrictive* than the new federal SSI standard could elect to retain those pre-1972 rules indefinitely. The drafting motivation was to prevent the new federal SSI standard from forcing more-restrictive states to liberalize their Medicaid disability eligibility. Section 209(b) of the Public Law (and the §1902(f) statutory cross-reference) became the namesake of the retained-discretion option.

The Congressional Research Service [@crs_r46111_2019] and the Centers for Medicare and Medicaid Services Implementation Guide [@cms_209b_implementation_guide] identify the 11 states that have retained 209(b) status continuously since 1974: Connecticut, Hawaii, Illinois, Indiana, Minnesota, Missouri, New Hampshire, North Dakota, Ohio, Oklahoma, and Virginia. The retained pre-1972 rules typically take three forms: (a) more restrictive income disregards, often by counting income that the federal SSI rules disregard, or by using a flat dollar cap rather than the federal benefit rate (FBR) plus a state supplement; (b) more restrictive asset limits, sometimes including the family home or vehicle that federal SSI rules exempt; or (c) a separate Medicaid disability application that requires re-submission of medical and financial documentation already filed with the Social Security Administration, including in some states a separate state-level disability determination [@rupp_riley_2016_ssbulletin]. The combined effect is that adults who are federally disabled and federally eligible for SSI in a 209(b) state may nevertheless be ineligible for the state's Medicaid program through the categorically-needy pathway — and may need to qualify through the medically-needy or spend-down pathway, both of which carry additional documentation and asset-spend-down requirements.

It is worth pausing on the legislative logic. The 1972 Amendments were drafted against the backdrop of a Medicaid program (Title XIX, enacted 1965) that had originally permitted wide state-level discretion in disability eligibility. The new federal SSI program promised to standardize cash assistance across states; but Congress confronted a politically thorny problem in the medical-assistance domain. If federal SSI eligibility automatically conferred Medicaid eligibility, states with pre-1972 Medicaid disability rules *more generous* than the new federal SSI standard would be required to retract eligibility (politically untenable), while states with pre-1972 rules *more restrictive* than the federal SSI standard would be required to extend eligibility (financially expensive for states). The §1902(f) option resolved the latter half of that tension by permitting more-restrictive states to retain their pre-1972 rules. The provision was, in this respect, a federalism compromise: states with restrictive pre-1972 disability eligibility could remain restrictive; the federal government would not bear the political cost of forcing them to liberalize.

The historical record indicates that the 11 §1902(f) states preserved their option not because the financial cost of conversion was uniformly prohibitive — several states with comparable pre-1972 restrictiveness elected the automatic-linkage path — but because the option was sticky. Once embedded in state Medicaid plans, the more-restrictive rules acquired their own administrative inertia: a separate state Medicaid application form, separate state agency staffing, separate state-level adjudication procedures, and (in several cases) a separate state-level definition of disability that interacted with state-level civil-service classifications. Two §1902(f) states have *converted* in the 50 years since: Massachusetts in 1985 and Vermont in 1992, both following statutory changes that aligned state Medicaid disability eligibility with the federal SSI standard. No state has converted in the other direction (from automatic linkage back to §1902(f)).

2.2 The administrative-friction wedge

The 209(b) provision created an administrative-friction wedge that is, in our setting, the *only* component of policy that varies at the state border. Federal SSI eligibility and the federal SSI benefit are the same on both sides of the border; the federal disability standard is the same; the cash recipient on each side receives an identical monthly SSI check. What differs is the path from "federally found disabled" to "enrolled in Medicaid." On the standard-SSI side, the path is automatic; on the 209(b) side, the path requires a separate state application, may be denied even when SSI is granted, and may impose more restrictive income and asset limits.

The literature on administrative burden has documented that frictions of this kind can have large take-up effects even when the substantive eligibility category is identical. Currie's 2004 NBER overview [@currie_2004_takeup] surveys the take-up literature and documents incomplete take-up across a range of cash-, food-, and medical-assistance programs. Moynihan, Herd, and Harvey [@moynihan_herd_harvey_2015] developed the conceptual framework of "learning costs," "psychological costs," and "compliance costs" that administrative-burden scholars have since deployed empirically. Heinrich [@heinrich_2016_bite] showed that such frictions reduce take-up among the population they nominally exclude least. Bhargava and Manoli [@bhargava_manoli_2015] documented in an IRS field experiment that very small changes in application complexity produced large take-up responses among eligible non-claimants. Fox, Feng, and Reynolds [@fox_feng_reynolds_2023] analyzed state-by-state variation in administrative-burden design for SNAP, TANF, and Medicaid; their cross-state results, while not specifically 209(b)-focused, support the general claim that paperwork frictions can be a binding constraint on safety-net participation. Arbogast, Chorniy, and Currie [@arbogast_chorniy_currie_2024] showed that even modest changes in Medicaid/CHIP administrative complexity changed child enrollment by economically meaningful amounts.

In the 209(b) context, the friction is not merely paperwork. The federal SSI determination, in the standard-SSI states, *is* the Medicaid application; in 209(b) states, the federal determination is necessary but not sufficient, and the state Medicaid application is a separate, sometimes substantially more demanding, exercise. Three policy-relevant predictions follow. First, *SSI take-up* may itself respond to the linkage, because applicants in 209(b) states face lower expected returns to a successful SSI determination (no automatic Medicaid). Second, *Medicaid coverage* among adults who have been federally found disabled may be lower in 209(b) states than in standard-SSI states, because the state-Medicaid application is a binding additional step. Third, *health* may be worse in 209(b) states than in standard-SSI states to the extent that Medicaid coverage among disabled adults is health-improving on the relevant margin.

The third prediction — health — is the most uncertain a priori. The disabled population on the margin between Medicaid and uninsurance is heterogeneous, and the marginal disabled adult who fails to clear the 209(b) state-Medicaid application is not necessarily the same person whose mortality risk would respond most to Medicaid coverage. The mortality literature is mixed on the magnitude of Medicaid effects [@goodmanbacon_2018; @sommers_baicker_epstein_2012; @miller_johnson_wherry_2021; @finkelstein_mcknight_2008], with effect sizes that vary across populations and time horizons. Our prior is that a 209(b)/SSI border-pair design would identify a small mortality effect if any — small enough that the question is empirical rather than predetermined.

2.3 Why border pairs

The Dube–Lester–Reich [@dube_lester_reich_2010] border-pair design is well-suited to the §1902(f) setting because it addresses the principal threats to a pure state-by-year DiD. The principal threat is that 209(b) states are systematically different from non-209(b) states in ways correlated with health outcomes, and that those differences might generate spurious DiD coefficients even without any §1902(f) effect on health. The set of 209(b) states includes Midwest manufacturing states (Indiana, Ohio, Illinois), high-population coastal states (Connecticut, Virginia), small-population farm-belt states (North Dakota, Oklahoma), and one Pacific island state (Hawaii). The comparison set includes Southern states (Tennessee, Kentucky, North Carolina, Texas), New England states (Maine, Massachusetts, Vermont, Rhode Island), the Plains (Kansas, Nebraska, South Dakota), and the mid-Atlantic (New York, Pennsylvania, Maryland, New Jersey via DC). On a pure state-level panel, the cross-section identification is contaminated by differences in industrial structure, racial composition, urbanization, and pre-1972 Medicaid generosity that are independent of §1902(f) per se.

The border-pair design absorbs these differences by restricting the comparison to contiguous pairs of counties that straddle a 209(b)/standard-SSI state line. Within such a pair, the two counties typically share a labor market, a regional hospital catchment, similar industrial composition, similar racial composition, and similar urbanization. The within-pair-year specification (a pair-by-year fixed effect) goes further and identifies the parameter purely from contrasts between adjacent counties observed in the same year, absorbing any region-by-year shock at the sub-state scale.

The cost of the design is that it discards observations that are not in a 209(b)/SSI border pair. The 483-pair sample includes 10 of the 11 209(b) states (all but Hawaii, which has no contiguous land border) and 23 standard-SSI comparison states. The 209(b) states with the most border-pair exposure are Virginia (multiple pairs with NC, KY, MD, DC, TN, WV) and the Midwest cluster (IL, IN, OH, MO bordering one another and the standard-SSI states KY, MI, PA, NY, WI). The geographic coverage is reasonably broad; we do not view the design as identifying off a small handful of pairs. The leave-one-state-out sensitivity in §5.5 documents that no single state drives the mortality DiD.

3. Data

3.1 The border-pair sample

We construct the border-pair sample from the U.S. Census Bureau's county-adjacency file, which lists every pair of U.S. counties sharing a land or water boundary. We retain only cross-state adjacencies, and among those, only the pairs in which one side's state is in the historical 209(b) set and the other is not. The resulting universe contains 483 contiguous-county border pairs spanning 10 209(b) treatment states (CT, IL, IN, MN, MO, ND, NH, OH, OK, VA) and 23 standard-SSI comparison states (AR, CO, DC, IA, KS, KY, MA, MD, ME, MI, MT, NC, NE, NM, NY, PA, RI, SD, TN, TX, VT, WI, WV). The 209(b) classification follows the Congressional Research Service [@crs_r46111_2019] and the CMS Implementation Guide [@cms_209b_implementation_guide], cross-validated against Ball [@ball1973_ssaamend] and Rupp and Riley [@rupp_riley_2016_ssbulletin]. Iowa, whose 209(b) status varies across sources, is treated as standard-SSI in the headline specification and tested in a leave-one-state-out sensitivity in §5.5.

The panel skeleton spans 1968–2020 at the (border-pair × county-side × year) level, yielding 51,198 rows: 483 pairs × 2 sides × 53 years. The 1968 lower bound is the first year of the CDC WONDER Compressed Mortality File; the 2020 upper bound is the last pre-pandemic year with stable IPUMS-CPS ASEC coverage. The 209(b) side and the standard-SSI side each contribute 25,599 rows, of which 22,701 are treated-post (`is_209b` × `post_1974 = 1`).

3.2 Mortality data

County-level all-cause mortality counts and population denominators come from CDC WONDER. The Compressed Mortality File covers 1968–1978 (ICDA-8 coding) and 1979–1998 (ICD-9). The Underlying Cause of Death file covers 1999 onward (ICD-10). We construct the all-cause mortality rate as deaths per 1,000 county residents, using WONDER's bridged-race population denominators. CDC suppresses sub-state counts below 10 deaths; in the border-pair sample, all-cause suppression is rare (the typical pair contains substantial populations). We flag suppressed cells with an indicator that we condition on at the analysis stage. Cause-specific deaths (disability-relevant: mental and behavioral, nervous system, circulatory, respiratory, and external causes) are deferred for future work because of more binding suppression at cause-by-county-by-year resolution.

The mortality panel contains 39,980 non-missing county-year cells for the all-cause rate.

The 11,218 missing cells are concentrated in the smallest-population border counties in the 1968–1978 window when WONDER suppression is more frequent.

3.3 IPUMS-CPS coverage and take-up

State-by-year aggregates of Medicaid coverage, SSI take-up, and disability prevalence among working-age (18–64) adults come from the IPUMS-CPS Annual Social and Economic Supplement (ASEC). We use four variables: `DISABWRK` (work-limiting disability, available across the entire 1968–2020 span), `HIMCAID` (Medicaid coverage, available 1988 onward), `INCSSI` (SSI dollars, available with strict separation from `INCWELFR` from 1989 onward), and the household-weighted person record. We compute (a) the disability-prevalence rate among adults 18–64, (b) the Medicaid coverage rate among adults 18–64 reporting a work-limiting disability, and (c) the SSI take-up rate among adults 18–64 reporting a work-limiting disability. The post-2008 multi-item disability battery (`DIFFCARE/DIFFREM/DIFFPHYS/DIFFMOB/DIFFSENS`) is reserved for robustness because of the well-documented 2008 series break.

The CPS state-year aggregates are joined to both sides of each border pair. Because the coverage and take-up variables begin in 1988 and 1989 respectively, we do not have pre-1972 measurements on these outcomes, and a clean pre/post DiD is not identified. We instead estimate cross-sectional 209(b) versus standard-SSI gaps with year fixed effects (so all variation is within-year, across-state), which are descriptive rather than causal. The CPS panel contains 1,089 state-year cells in the coverage/take-up estimating sample.

3.4 SSA county-year SSI uptake

The third outcome arm draws on a newly-constructed county-year SSI uptake panel based on SSA's "SSI Recipients by State and County" publications. From FY 2003 onward, SSA publishes county-level recipient counts (total, aged, blind, and disabled) as Excel workbooks. We acquired 22 workbooks covering FY 2003–2024 directly from SSA's Open Data portal and ingested them via `data/scripts/09_parse_ssa_ssi_county.py`. For FY 1998–2002, SSA publishes the same data as state-level HTML pages (one page per state per year), which we render and parse for the SSI/comparison states relevant to the border-pair sample (24 states $\times$ 5 years = 120 HTML pages).

The resulting panel contains 77,331 matched county-years across 50 states with a 99.9% FIPS match rate. The 76 unmatched rows correspond to historical county reorganizations — for example, South Dakota's Washabaugh County (merged into Jackson County in 1979) and Virginia's Clifton Forge and South Boston independent cities (reverted to county status in the 1990s and 2000s). These are legitimate historical reorganizations rather than parse failures.

When merged into the border-pair panel, the SSA arm contributes 19,193 cells with non-missing SSI uptake (the intersection of 1998–2020 with the WONDER population denominator). Of these, 17,330 are non-missing on the year-and-pair fixed-effects estimating sample used in §5.4.

The key SSA outcomes we examine are (a) total SSI recipients per 1,000 county residents, (b) blind-and-disabled SSI recipients per 1,000 county residents, and the corresponding log transformations. The blind-and-disabled subset is the more 209(b)-relevant disaggregation because aged SSI is governed by a different state-Medicaid pathway in most 209(b) states.

The 1996–1997 fiscal years are not yet ingested. SSA's pre-1998 publications are available only in print or on HathiTrust as scanned PDFs, requiring a manual ingest step that we defer.

3.5 Pre-1972 Medicaid eligibility classification

We do not directly observe pre-1972 state Medicaid eligibility rules in machine-readable form at the level of detail needed to characterize *which* dimensions of restrictiveness each 209(b) state retained. The CRS report [@crs_r46111_2019] documents that the 11 209(b) states most commonly retained the more restrictive of (a) income disregards, (b) asset limits, or (c) separate disability determination. We do not exploit within-209(b)-state variation in retained-rule type; our headline specification treats all 11 209(b) states as a single treatment, and the across-state heterogeneity is absorbed by border-pair fixed effects.

A finer characterization of *which* pre-1972 rule each 209(b) state retained is a productive direction for future work.

3.6 Sample restrictions and data limitations

We make four deliberate sample choices. First, we restrict the analysis to contiguous border pairs across 209(b) and standard-SSI states; intra-state pairs and 209(b)–209(b) or SSI–SSI border pairs are excluded by construction. Second, we use county-FIPS definitions as of the 2010 Census adjacency release; pre-2010 changes in Virginia's independent-city/county arrangements are handled by the underlying WONDER and SSA files, but the analysis stage verifies within-pair FIPS continuity. Third, we use the pre-2008 `DISABWRK` variable as the workhorse disability measure to maintain consistency across the panel. Fourth, Hawaii is necessarily excluded from the border-pair sample because it has no contiguous land border with another state; the 11 209(b) states reduce to 10 in the border-pair design.

Principal data limitations are (a) WONDER suppression below 10 deaths for cause-specific outcomes, which we manage by reporting all-cause mortality as the main outcome and deferring cause-specific splits; (b) the SSA county-year SSI series begins in 1998, so a clean pre-1974 DiD on administrative SSI uptake is not identified; (c) CPS coverage and take-up variables begin in 1988/1989, so the same restriction applies; (d) the geographic running variable used in our supporting RD is constructed from county centroids and is therefore coarse. These limitations are honest constraints on the identification, not silent omissions; we revisit them in §6.4.

4. Empirical Strategy

4.1 The border-pair DiD

For outcome Y in county i, border-pair c, and year t, the workhorse specification estimates

$$Y_{ict} = \alpha_i + \tau_t + \beta \cdot \mathbb{1}[209(b)_i] \cdot \mathbb{1}[t \geq 1974] + \varepsilon_{ict},$$

where α_i is a county fixed effect, τ_t is a year fixed effect, and β captures the average treatment effect on the treated (ATT) of 209(b) status in the post-federalization era. Standard errors are clustered at the border-pair level. The within-pair-year (DLR) specification replaces τ_t with a pair-by-year fixed effect, so β is identified purely from contrasts between adjacent counties observed in the same year:

$$Y_{ict} = \alpha_i + \mu_{ct} + \beta \cdot \mathbb{1}[209(b)_i] \cdot \mathbb{1}[t \geq 1974] + \varepsilon_{ict}.$$

Both specifications are reported. The within-pair-year specification is more demanding because it absorbs all common shocks at the pair-by-year level, including regional healthcare market shocks, but it can be sensitive to small-sample noise when pair sizes are small.

Because all 209(b) states acquired the status in the same year (1974), and no state in our panel switches status across the 53-year window, there is no staggered adoption. Two-way fixed-effects (TWFE) DiD is unbiased here; Callaway–Sant'Anna [@callaway_santanna_2021], Sun–Abraham [@sun_abraham_2021], and de Chaisemartin–D'Haultfœuille estimators collapse to the two-group, two-period DiD. The analytic code verifies that no state's `is_209b` value varies across years; for the current panel, TWFE is the appropriate estimator.

4.2 The event study

We estimate dynamic effects by replacing the single `treat_post` indicator with a saturated set of `is_209b` $\times$ `event-time` interactions, omitting the year before treatment ($t = -1$) as the reference. Event time is binned to manage degrees of freedom across the 53-year window: ≤ -6, -5, -4, -3, -2, -1 [ref], 0, +1, +2, +3, +4, +5 to +9, +10 to +19, $\geq +20$. The pre-period coefficients ($t = -6$ to $t = -2$) provide visual evidence on parallel pre-trends; the post-period coefficients ($t = 0$ through $t \geq +20$) trace the dynamic response. Standard errors are clustered at the border-pair level; we report point estimates with 95% pointwise confidence intervals.

4.3 The cross-sectional CPS gap

Because the IPUMS-CPS coverage and take-up variables begin in 1988 (`HIMCAID`) and 1989 (`INCSSI`), the pre-1972 DiD is not identified on these outcomes. We instead estimate cross-sectional 209(b) versus standard-SSI gaps with year fixed effects:

$$Y_{st} = \tau_t + \gamma \cdot \mathbb{1}[209(b)_s] + u_{st},$$

where s indexes state and t indexes year, and γ captures the average level gap across 209(b) states relative to standard-SSI states, net of common-year shocks. We do not interpret γ causally because it does not exploit a pre-period contrast; we interpret it descriptively as the stylized fact that 209(b)-state working-age disabled adults systematically report less Medicaid coverage and less SSI receipt than standard-SSI-state counterparts.

4.4 The SSA county-year border-pair gap

The SSA county series begins in 1998, so the pre-1974 DiD is also not identified. We estimate two specifications. The first regresses log SSI per 1,000 on `is_209b` with year and pair fixed effects, identifying β from the average within-pair, within-year, between-side gap:

$$\log\left(\frac{\text{SSI}_{ict}}{\text{population}_{ict}}\right) \cdot 1000 = \tau_t + \pi_c + \beta \cdot \mathbb{1}[209(b)_i] + \varepsilon_{ict}.$$

The second specification replaces $\tau_t + \pi_c$ with a pair-by-year fixed effect μ_{ct}, identifying β from contrasts between adjacent counties observed in the same year. We treat the within-pair-year specification as our preferred SSA arm specification.

4.5 The supporting RD design

We also estimate a geographic regression discontinuity using signed distance to the nearest 209(b)/SSI state border as the running variable. The signed distance is constructed from county centroids: for each county, we compute the great-circle distance to the centroid of the nearest opposite-state border-pair county, take half that distance, and assign positive distance to the 209(b) side and negative to the standard-SSI side. Following standard RD practice, we report the IK-pilot bandwidth (approximately 12 km here) and conduct a bandwidth sensitivity sweep.

We caution that the RD running variable is coarse. Centroid-to-centroid distance is a crude proxy for actual border distance; few counties sit within 12 km of the border; the narrow-window estimate is dominated by Arlington and Alexandria, VA — urban independent cities whose low mortality may reflect SES and access selection rather than 209(b) per se; the bandwidth sweep is unstable ($\beta = -9.48$ at 12.2 km, -5.47 at 15.2 km, -1.22 and not statistically significant at 18.2 km); and the sign at the narrow bandwidth is opposite to the long-run DiD sign. We therefore demote the RD to exploratory appendix evidence only and do not claim that it corroborates the DiD headline. The DiD remains the more credible identification on this panel.

4.6 Robustness

The robustness battery comprises six tests, each implemented in `analysis/robustness/`:

- **Event-study pre-trends** (`event_study.py`): dynamic coefficients with the reference period one year before treatment, reported in §5.2.
- **Pre-1972 placebo DiD** (`robustness.py`): re-estimate the DiD on data restricted to 1968–1971, with a false treatment year of 1971; expected coefficient zero.
- **Temporal bandwidth sensitivity** (`bandwidths.py`): re-estimate using temporal windows of ± 5, ± 10, ± 15, ± 20, ± 30 years, and the full panel around 1974.
- **Leave-one-state-out (LOSO)** (`drop_one_state.py`): iteratively remove each of the 33 states in the panel and re-estimate; report the distribution of coefficients.
- **Honest-DiD relative-magnitudes** (`honest_did.py`): inflate post-treatment confidence bounds by $M \times \mathsf{max}(|\text{pre-period coef}|)$ for $M \in \{0.5, 1.0, 1.5, 2.0\}$ and report the breakdown M at which the lower bound crosses zero [@rambachan2023more].
- **Geographic-FE sensitivity** (`geographic_fe.py`): contrast county + year, county + pair-year, state + year, and county + year + state-linear-trend specifications.

- **Donut RD** (`robustness.py`): drop counties within 5 km and within 10 km of the border to assess sensitivity to the narrowest-window observations.
- **McCrary-style density** of the running variable (`robustness.py`).

4.7 Reproducibility

All scripts are Python and rely on `pyfixest` for high-dimensional fixed effects with cluster-robust inference, `pandas` for data handling, and `matplotlib` for figures. The full analysis runs end-to-end via `python3 analysis/run_all.py` from the project root and writes coefficient tables to `analysis/tables/`, figures to `analysis/figures/`, and text logs to `analysis/log/`. Total runtime is approximately three minutes on a 2024-vintage Apple-silicon laptop.

5. Results

5.1 Main mortality DiD

Table 2 reports the static border-pair DiD on all-cause mortality. In the county + year fixed-effects specification, 209(b)-side border counties experienced 0.205 additional deaths per 1,000 residents annually after 1974 under pair-clustered inference (95% CI 0.030 to 0.380, $p = 0.022$, $N = 39{,}980$). The within-pair-year specification returned an essentially identical 0.207 (SE 0.105, $p = 0.050$). Translated to log mortality, the coefficient is +0.013 to +0.018 across specifications, statistically marginal in the within-pair-year specification ($p = 0.059$). Under county-level clustering, the same level coefficient has SE 0.162 ($p = 0.21$); under state-level clustering SE 0.258 ($p = 0.44$); and on a unique-county-year sample that collapses duplicate observations arising from counties belonging to multiple border pairs, the level point estimate shrinks to +0.145 (county-clustered SE 0.157, $p = 0.35$; state-clustered SE 0.275, $p = 0.60$). We therefore treat the static estimate as suggestive and direct the reader to the event study, which is the substantive object: the static average aggregates a non-monotonic two-decade convergence-then-divergence path that no single scalar can summarize.

The point estimate corresponds to roughly 0.205 / 9 $\approx$ 2.3% of the post-1974 mean all-cause mortality rate of approximately 9.0 deaths per 1,000 across the border-pair sample. Aggregated across the approximately 230 million Americans who live in 209(b) states today, the implied excess all-cause mortality (under the strong assumption that the border-pair effect scales to the full state population) would be on the order of 47,000 deaths per year. We are cautious about this extrapolation because the border-pair effect may not scale linearly outside the contiguous-border sample; it is offered as a magnitude benchmark, not a national estimate.

5.2 Event-study dynamics

The event-study figure (Figure 2) tells a more nuanced story than the static DiD coefficient. The five pre-1974 leads ($t = -6$ to $t = -2$) are all small in magnitude (absolute values ≤ 0.23) and statistically insignificant, supporting the parallel-trends assumption. After 1974, the dynamic coefficients on level mortality are:

- $t = 0$ (1974): $\hat{\beta} = -0.198$ (SE 0.096), a small initial dip on the 209(b) side;
- $t = +1$: $\hat{\beta} = -0.406$ (SE 0.096) — a substantial *negative* short-run effect;
- $t = +2$: $\hat{\beta} = -0.338$ (SE 0.105);
- $t = +3$: $\hat{\beta} = -0.117$ (SE 0.109);
- $t = +4$: $\hat{\beta} = -0.209$ (SE 0.100);
- $t = +5$ to $+9$ (binned): $\hat{\beta} = -0.224$ (SE 0.084);
- $t = +10$ to $+19$ (binned): $\hat{\beta} = -0.007$ (SE 0.090) — convergence;
- $t \geq +20$ (1994–2020): $\hat{\beta} = +0.262$ (SE 0.118) — a positive long-run divergence.

The static DiD coefficient of +0.205 aggregates this non-monotonic path. The substantive story is not "209(b) immediately caused excess mortality." It is that the two groups *diverged slowly over two decades*: the 209(b) side had lower all-cause mortality in the first decade after federalization, converged in the second decade, and diverged positively by the third decade and beyond. We interpret the long-run divergence as plausibly consistent with slow accumulation of unmet need among the population that fails to clear the 209(b) state-Medicaid hurdle, combined with cohort dynamics in which the population most affected by

§1902(f) is aging into higher mortality risk. This interpretation is suggestive rather than dispositive; the data identify the sign-shift but do not pin down its mechanism.

5.3 CPS cross-sectional gap

Table 3 reports the CPS cross-sectional gap regressions for working-age (18–64) adults with a work-limiting disability. The Medicaid coverage gap is −0.0899 (SE 0.0195, $p = 6.3 \times 10^{-5}$, N = 1,089): 209(b)-state disabled adults report 8.9 percentage points lower Medicaid coverage than standard-SSI-state counterparts, with no overlap of 95% confidence intervals. The SSI take-up gap is −0.0287 (SE 0.0098, p = 0.006, N = 1,089): 209(b)-state disabled adults report 2.9 percentage points lower SSI receipt, or about 11% below the standard-SSI-state base rate of approximately 25%.

The disability prevalence gap is −0.0097 (SE 0.0067, p = 0.153, N = 1,557), statistically indistinguishable from zero. We interpret this as a useful null: the cross-sectional gap in self-reported coverage and take-up is not driven by a 209(b)/SSI difference in disability *prevalence*. Conditional on being disabled, the 209(b) population is less covered and has lower program take-up; it is not the case that fewer adults in 209(b) states are disabled in the first place. The §1902(f) friction does its work on the path from disability to program participation, not on the underlying disability rate.

We caution again that the CPS results are cross-sectional and descriptive, not causal. They do not exploit a pre-1972 baseline (the CPS variables do not extend that far) and they do not impose a border-pair geography. We report them as a stylized fact that complements the mortality DiD and that motivates the SSA arm.

5.4 SSA administrative SSI uptake gap

Table 4 reports the border-pair gap regressions on the SSA county-year administrative SSI uptake panel, 1998–2020. In the year-and-pair fixed-effects specification, log SSI total recipients per 1,000 is −0.187 lower on the 209(b) side of the same border pairs (SE 0.042, $p = 9.9 \times 10^{-6}$, N = 17,330), corresponding to a roughly 17% relative gap. The within-pair-year specification returns −0.199 (SE 0.037, $p = 2.3 \times 10^{-7}$), essentially identical. In levels, the gap is −4.55 recipients per 1,000 (SE 0.71, 95% CI −5.94 to −3.16, $p = 2.9 \times 10^{-10}$). The blind-and-disabled subset, the more 209(b)-relevant disaggregation, shows an even larger log gap of −0.564 (SE 0.225, p = 0.012) and a level gap of −4.36 per 1,000 (SE 0.70, $p = 8.9 \times 10^{-10}$).

The event study over the SSI era (Figure 4) shows that the gap is remarkably stable across two decades of fiscal years. From 1998 through 2008, the year-by-year is_209b coefficient on log SSI per 1,000 fluctuates between −0.23 and −0.28 with consistently small standard errors. After 2008, the year-by-year point estimates remain in the −0.18 to −0.22 range for most years, with larger standard errors in a handful of years (2009, 2010, 2011, 2014, 2017, 2020) where SSA's series reconstruction or alt-table publication appears to widen sampling variation. There is no detectable trend across the 1998–2020 window: the 209(b) shortfall in administrative SSI uptake is a *persistent two-decade feature of the data*, not an artifact of any particular fiscal year or post-2008 ARRA expansion.

The CPS and SSA arms agree in sign and order of magnitude. The CPS cross-sectional SSI take-up gap of −2.9 percentage points (about 11% on a 25-percentage-point base, for 1989–2020 disabled adults) is in the same direction as the SSA −17% relative gap and is somewhat smaller. The well-documented CPS under-reporting of SSI receipt at the bottom of the income distribution [@meyer_mok_sullivan_2015] plausibly accounts for the attenuation of the survey estimate toward zero relative to the administrative estimate. Either way, the CPS finding survives translation to administrative data drawn from a different denominator (county working-age population rather than CPS disabled-adult subsample), a different unit of observation (county rather than state-year), and a different time window (1998–2020 rather than 1988–2020).

5.5 Robustness

The robustness battery (Table 5) returns the following:

- **Pre-1972 placebo DiD (1968–1971, false treatment year 1971)**: $\hat{\beta} = -0.008$ log (SE 0.0076, $p > 0.27$) and $\hat{\beta} = -0.115$ level (SE 0.092, $p > 0.21$). Both placebo coefficients are statistically indistinguishable from zero and smaller in magnitude than the post-1974 estimate, supporting the parallel pre-trends assumption.

- **Five-year placebo-date sweep (1969–1973)**: log-mortality DiD coefficients are within ± 0.013 of zero across all placebo years, with overlapping confidence intervals. The 1974 "true" coefficient of -0.011 in this restricted window is, like the placebos, small — confirming that the short-run negative dynamic shows up cleanly in the event-study figure rather than in a restricted-window static DiD.

- **Temporal bandwidth sweep**: the static DiD coefficient is -0.111 at ± 5 years, -0.122 at ± 10 years, -0.071 at ± 15 years, -0.002 at ± 20 years, $+0.121$ at ± 30 years, and $+0.205$ on the full panel. The sweep recovers the event-study sign-shift in static form: the short-run negative effect dominates narrow windows; the long-run positive divergence emerges only when 20+ post-treatment years are included.

- **Leave-one-state-out**: the mortality DiD coefficient ranges from 0.134 (drop NC) to 0.305 (drop KY), with no sign flip and stable standard errors. The result is not driven by any single state.

- **Honest-DiD relative-magnitudes**: with the largest pre-period absolute coefficient of 0.225, the lower confidence bound at $M = 1$ is -0.821, at $M = 2$ is -1.047. Honest-DiD bounds are wide because the pre-period exhibits mixed-sign variation; we report the bounds transparently as a caveat rather than as a clean parallel-trends pass.

- **Geographic-FE sensitivity**: the mortality DiD coefficient is $+0.205$ under county + year, $+0.207$ under county + pair-year (DLR), $+0.223$ under state + year, and -0.171 under county + year + state-linear-trend. The state-linear-trend specification sign-flips, which we read as a sign that linear pre-trends absorb the event-study sign-shift (the long-run positive divergence is partly absorbed into the state linear trend). We do not view the state-linear-trend specification as the right benchmark because the event-study dynamics are non-monotonic; a linear trend cannot represent the convergence-then-divergence pattern.

- **Donut RD**: the log-mortality RD coefficient at the IK pilot bandwidth is -1.225 (SE 0.37, full sample) and -1.132 (SE 0.46) after dropping counties within 5 km of the border. At a 10-km donut, the sample drops to 15 counties, which is insufficient support. Within the available bandwidth, the donut RD result is stable.

- **McCrary-style density of the running variable**: the log-ratio across the cutoff is between 0 and -0.106 at bandwidths from 10 km to 30 km, with z-statistics below 1.1 in absolute value across the sweep. There is no evidence of manipulation of the geographic running variable (as one would not expect for a fixed state-border instrument).

Overall, the mortality DiD and the SSA SSI uptake gap are robust to the main sensitivity checks. The RD point estimate is sensitive to bandwidth widening, which we attribute to the coarse running variable and the urban-county-driven narrow-window sample. We do not headline the RD; we report it as a supporting cross-section gradient.

5.6 Magnitudes and triangulation

Putting the three arms together: across 50 years of border-pair observation, 209(b) status is associated with (a) approximately 8.9 percentage points lower Medicaid coverage and 2.9 percentage points lower SSI receipt among working-age disabled adults in CPS survey data, (b) approximately 17% lower SSI uptake per 1,000 residents in SSA administrative data, and (c) 0.205 additional all-cause deaths per 1,000 residents per year, with the mortality effect dominated by long-run event-study horizons. The three arms are coherent in sign and order of magnitude: a state policy choice that reduces Medicaid coverage and SSI take-up among the working-age disabled also, after a sufficient horizon, increases all-cause mortality.

The triangulation across data sources has analytic value beyond redundancy. Survey-based measures of SSI receipt (CPS `INCSSI`) suffer well-documented misclassification at the bottom of the income distribution, with the direction of bias generally toward attenuation [@meyer_mok_sullivan_2015]. If we relied on CPS alone, a sceptical reader would be entitled to ask whether the 2.9-percentage-point SSI take-up gap was a real policy effect or a survey-measurement artifact. The SSA administrative arm, drawn from the federal trust-fund disbursement records themselves, eliminates that concern: an SSA recipient is by construction a person whose monthly SSI check was disbursed, regardless of whether that person reports the receipt in any survey. The SSA gap of approximately 17% (and 4.55 recipients per 1,000 in levels) confirms that the CPS gap is real and, if anything, attenuated by survey under-reporting.

In the same way, the mortality arm has the analytic value of being an *independent* outcome from the coverage and take-up arms. If we had only the CPS and SSA arms, we would have established that §1902(f) suppresses program participation; the question of whether suppressed participation translates into worse population health would remain open. The mortality DiD addresses that question: the 0.205 additional deaths per 1,000 residents per year is identified off the same border-pair geography that delivers the CPS and SSA gaps, and the event-study horizons are consistent with a slow-accumulation mechanism in which the cumulative coverage and take-up shortfall does its work over time.

We do not interpret the long-run mortality effect as a tight estimate of the "lives saved per additional Medicaid enrollee" parameter that the Medicaid mortality literature focuses on. Our coefficient is a long-run, equilibrium, population-level treatment effect of a state policy choice, not a marginal coverage effect. The literature parameter most comparable to ours is the long-run effect of state Medicaid expansion on all-cause mortality, which Miller, Johnson, and Wherry [@miller_johnson_wherry_2021] estimate at roughly 0.13 to 0.20 deaths per 1,000 over six post-expansion years using linked-administrative data. Our coefficient sits at the upper end of that range, but on a 20+ year horizon rather than a 6-year horizon, and in a fundamentally different setting (administrative friction, not coverage availability). The two findings are mutually reinforcing rather than directly comparable.

6. Discussion

6.1 Summary of findings

This paper documents that the §1902(f) Medicaid linkage choice — adopted by 11 states in 1974 and retained continuously by all 11 — has measurable consequences for population health and program participation half a century later. Across three independent data sources spanning 1968 through 2024, we find that 209(b) states have (a) approximately 8.9 percentage points lower Medicaid coverage among working-age disabled adults in CPS survey data, (b) approximately 2.9 percentage points lower SSI take-up among the same population, (c) approximately 17% lower SSI uptake per 1,000 residents in SSA administrative data, with a remarkably flat gap across two decades, and (d) 0.205 additional all-cause deaths per 1,000 residents per year on border-pair contrasts, with the effect concentrated in the long-run event-study tail ($t \geq +20$). The first three findings are about program participation; the fourth is about population health.

The central interpretive challenge is the event-study sign-shift in mortality: the 209(b) side had *lower* mortality in the first decade after federalization, converged in the second decade, and diverged *positively* in the third decade and beyond. This non-monotonic pattern is consistent with slow accumulation of unmet need rather than an immediate adverse shock — a story in which the friction does its work over time, as the population most affected by §1902(f) ages into higher health risk and as the policy difference compounds across decades. It is also consistent with cohort dynamics: the disabled adults who entered the program in the late 1970s under the 209(b) state-Medicaid rules are a different cohort, with different baseline health, than the cohort whose mortality we observe in the 1990s and 2000s. We do not separately identify the unmet-need and cohort channels in this paper.

6.2 Relation to prior literature

The paper contributes to three literatures that have largely run in parallel. First, the administrative-burden literature [@currie_2004_takeup; @heinrich_2016_bite; @moynihan_herd_harvey_2015; @herd_moynihan_2018_book; @bhargava_manoli_2015; @fox_feng_reynolds_2023; @arbogast_chorniy_currie_2024] establishes that small frictions can have large take-up effects. Our CPS and SSA findings — an 8.9 percentage point Medicaid coverage gap and a 17% SSI uptake gap — are consistent with the magnitudes that this literature reports. We extend the literature in two directions: (a) by demonstrating that a specific federal-state structural choice (the §1902(f) option) produces an administrative-burden wedge that persists for 50 years, and (b) by linking that wedge to a downstream mortality outcome through a border-pair DiD.

Second, the Medicaid mortality literature [@goodmanbacon_2018; @sommers_baicker_epstein_2012; @miller_johnson_wherry_2021; @finkelstein_mcknight_2008] has focused on episodes where Medicaid coverage *changed* — the 1965–1970 rollout, the 2001–2005 expansions, the 2014 ACA expansion. We study an episode where Medicaid coverage *failed to fully change* in 1974 because §1902(f) preserved the pre-federalization eligibility rules in 11 states. The implied counterfactual is "what would mortality look like if 209(b) states had adopted automatic SSI–Medicaid linkage in 1974?" Our long-run coefficient (+0.205 deaths/1,000) sits within the range of effect sizes reported by Miller, Johnson, and Wherry [@miller_johnson_wherry_2021] on the ACA expansion, despite the very different time horizon and the very different identifying variation.

Third, the disability-program economics literature [@currie_madrian_1999; @autor_duggan_2003; @maestas_mullen_strand_2013; @deshpande_2016; @schmidt_shoresheppard_wat @soni_burns_dague_simon_2017; @burns_dague_wood_kennedy_2022] has studied how disability program design — including the SSI–Medicaid linkage — shapes program participation and labor-market outcomes. The closest paper in spirit is Rupp and Riley [@rupp_riley_2016_ssbulletin], who used SSA administrative data to characterize cross-sectional state-level differences in Medicaid participation among disabled SSI recipients. We add a 53-year mortality panel, a border-pair design that addresses cross-sectional confounding, and a long county-year administrative panel that the Bulletin's earlier descriptive cross-section could not support.

6.3 A companion paper: the modern SSA office-closure shock

The persistent 209(b) penalty documented here likely amplifies the modern administrative-burden shocks studied in a companion working paper (Palisoc, "When the Office Closes: Social Security Administration Field Office Consolidations, Medicaid Coverage Disruptions, and Health Outcomes Among Adults with Disabilities," 2026). That paper studies the 2008–2024 wave of SSA field office closures and limited-service conversions, using office closure as an instrument for downstream Medicaid coverage among SSI-eligible adults. The 209(b) heterogeneity dimension is a built-in mechanism test: if the SSA closure shock operates through SSI applications that subsequently fail to convert to Medicaid coverage, the closure effect on coverage should be larger in 209(b) states (where the SSI–Medicaid path is already weak) than in standard-SSI states (where the SSI determination automatically yields Medicaid).

The relationship between the two papers is structural: the present paper establishes the historical baseline 209(b) penalty (the cumulative effect of 50 years of friction on coverage, take-up, and mortality), and the companion paper studies the modern *additional* friction that office closures impose on top of that baseline. Reading the two papers together, the policy implication is that the 209(b) option is not a static feature of the safety net; it is an amplifier on every additional friction that policy designers add. A wave of office closures has a larger downstream Medicaid effect in 209(b) states; the present paper provides the structural reason.

6.4 Limitations

Several limitations of the present analysis deserve explicit attention.

SSA county-year SSI series begins in 1998. SSA published "SSI Recipients by State and County" tables only from FY 1998 forward in online publications; the 1996–1997 tables exist on HathiTrust as scanned PDFs and require a manual OCR/ingest step that we have deferred. The SSA arm therefore identifies β on a 1998–2020 cross-section with year and pair fixed effects rather than on a pre-1974 DiD. The CPS arm is similarly limited because `HIMCAID` begins in 1988 and `INCSSI` (in strict separation from `INCWELFR`) begins in 1989. Only the mortality arm has a true pre-1972 baseline.

Static DiD masks long-run divergence. The +0.205 deaths/1,000 headline aggregates short-run negative effects, mid-run nulls, and long-run positive effects. The event-study figure is the primary substantive object; the scalar DiD is a summary that obscures non-monotonic dynamics. We headline both transparently.

Cause-of-death splits are deferred. The current panel uses all-cause mortality only. Disability-relevant cause-specific deaths (mental and behavioral, nervous system, circulatory, respiratory, external) would strengthen the mechanism interpretation; the CDC WONDER suppression rule (counts below 10 suppressed at sub-state resolution) is more binding for cause-specific outcomes and would require careful suppression handling. We defer this to future work.

Geographic RD running variable is coarse. The signed distance is constructed from county centroids and is therefore a crude proxy for actual border distance. The narrow-window RD is dominated by Arlington and Alexandria, VA, urban independent cities whose low mortality may reflect SES and access selection rather than 209(b) per se. We do not headline the RD; we report it as a supporting contrast. A future revision that constructs distance to the actual state border on a sub-county geography (e.g., census-tract or block-group) would tighten the geographic identification.

Migration is not separately identified. Some disabled adults in 209(b) states may move across the border to access more generous Medicaid linkage in standard-SSI states. Such migration would attenuate, not reverse, the sign of our coefficients; the long-run positive mortality divergence is the opposite of what selective migration into 209(b) states would predict. We note migration as a residual mechanism rather than as a threat to identification.

State-supplement generosity is a confounder. 209(b) states tended to have less generous state SSI supplements than standard-SSI states. A complete characterization of the §1902(f) effect should separately identify the supplement-generosity channel from the categorical-eligibility channel. We do not include a state-supplement control in the headline specification; this is a productive direction for future work.

Hawaii is excluded by the border-pair design. The 11 209(b) states reduce to 10 in our sample. The generalizability of our findings to island and other geographically isolated contexts is limited.

Within-209(b)-state heterogeneity is absorbed. The 11 209(b) states retained different combinations of pre-1972 income disregards, asset limits, and separate-determination requirements. We treat all 11 states as a single treatment because we do not yet have a machine-readable characterization of each state's retained-rule type. A future revision that exploits within-209(b) variation in retained-rule type would identify which specific rule features carry the policy weight.

Equity-stratified heterogeneity is deferred. The CPS-based gap estimates above are reported for the full work-limiting-disability subsample of working-age adults. Race/ethnicity-stratified, poverty-stratified, or geographic-vulnerability-stratified versions of the CPS coverage and take-up gaps would be a natural and policy-relevant extension. The IPUMS-CPS state-year cells for the work-limiting-disability subsample are thin once stratified by race/ethnicity, and a credible heterogeneity decomposition would require either pooled multi-year IPUMS-CPS bins or a transition to the restricted-use SIPP file; both are beyond the scope of the current draft. We flag the equity-stratification deferral honestly rather than presenting an underpowered cross-tab.

6.5 Mechanism: why does the friction translate into mortality?

The event-study sign-shift deserves a more developed mechanism discussion. The pattern is not that 209(b) status caused an immediate jump in deaths; the static average of +0.205 deaths/1,000 obscures a path that ran *negative* in the first decade and *positive* in the third decade and beyond. Three non-exclusive mechanisms are consistent with this dynamic.

First, *slow accumulation of unmet need.* If the marginal disabled adult in a 209(b) state who fails the state-Medicaid hurdle experiences episodic rather than continuous health degradation, and if the degradation accumulates over years through under-treated chronic conditions (diabetes, hypertension, mental and behavioral disorders, respiratory disease), the mortality consequence may not be visible in the early post-1974 horizon. By the time we observe the 1994–2020 horizon, the cumulative under-treatment is two decades old in the cohort that entered the program in 1974, and longer in the cohorts that entered subsequently. The long-run positive coefficient is consistent with this story.

Second, *cohort dynamics.* The disabled adults who entered SSI in 1974 are not the same cohort whose mortality we observe in the 1990s and 2000s. Cohort replacement may mean that the 209(b) friction is felt more sharply by later-entering cohorts — for example, because state Medicaid eligibility systems have grown more demanding over time, or because the SSI–Medicaid linkage's marginal value has risen as Medicaid coverage has become a more important determinant of access to specialty care.

Third, *short-run compensating mechanisms that erode over time.* It is possible that 209(b) states' lower-mortality first-decade response reflects compensating institutions — county hospitals, state-supplemental cash assistance, family caregiving capacity — that absorbed the marginal disabled adult's medical needs in 1974 but that gradually eroded in the 1990s and 2000s under fiscal pressure. The decline of county-hospital systems and the de-institutionalization of state mental-health facilities both fit this timing. We do not separately identify this channel in the present data.

The three mechanisms are observationally similar in our setting and we do not attempt to distinguish among them. They share the policy-relevant implication that the §1902(f) effect on health is a long-run equilibrium effect, not a short-run shock. A 209(b) state that converted to automatic linkage today might not see its mortality penalty reverse for years or decades; conversely, the 50-year persistence of the gap is evidence that the policy lever has cumulative consequences.

6.6 Policy implications

Three policy implications follow from the present findings.

First, §1902(f) is a *live structural lever.* Eleven states retain the option, and the option is not a vestigial provision of the 1972 Act; descriptive evidence across three independent data systems consistently shows lower disabled-adult coverage and lower SSI participation in §1902(f) states. The cross-sectional border-pair gaps documented here — approximately 8.9 percentage points on Medicaid coverage and approximately 17% on SSA-recorded SSI uptake — are descriptive associations under year fixed effects, not causal counterfactual estimates. A conversion to automatic SSI–Medicaid linkage by a current 209(b) state (as Massachusetts did in 1985 and Vermont in 1992) might shrink those gaps, but the precise magnitude of any such effect would depend on take-up elasticities, state administrative infrastructure, and outreach intensity that this paper does not separately identify. The mortality channel, where our identification is strongest but the static point estimate is fragile, is best understood through the event-study sign-shift rather than the scalar DiD.

Second, the §1902(f) effect *interacts with every other administrative-burden shock.* A wave of SSA office closures, a Medicaid renewal-paperwork policy change, a state-Medicaid eligibility-system upgrade — each of these will have larger downstream consequences in 209(b) states than in standard-SSI states, because the 209(b) state-Medicaid path is already friction-loaded. The companion office-closure paper makes this explicit, but the present paper's results imply that the same logic applies to any administrative-burden shock that affects the SSI-to-Medicaid path.

Third, the *policy choice is reversible.* A 209(b) state that elects to come into compliance with the federal SSI standard for Medicaid disability eligibility — i.e., to accept the federal SSI determination as conclusive for the categorically-needy Medicaid pathway — would convert the state to standard-SSI linkage. The path is statutorily available; what is missing is the political and administrative will to make the conversion. The present paper provides the empirical basis for evaluating that decision.

7. Conclusion

A policy decision made by 11 state legislatures in 1973 — to retain the more-restrictive pre-1972 Medicaid disability rules under the §1902(f) option — continues to suppress disabled-adult Medicaid coverage by 8.9 percentage points and disabled-adult SSI take-up by 2.9 percentage points in CPS survey data, depresses SSI uptake per 1,000 residents by approximately 17% in SSA administrative data, and elevates all-cause mortality by 0.205 deaths per 1,000 residents per year on border-pair contrasts, with the mortality effect dominated by long-run event-study horizons. Five decades after federalization, the 209(b) decision is not a historical curiosity; it is a live structural lever for 11 states with measurable health and welfare consequences. Triangulating across three identification arms

— mortality DiD, CPS cross-section, and SSA administrative panel — yields a coherent picture of administrative friction as health policy.

The findings position the §1902(f) option alongside other state-level policy choices that determine the experiential character of the federal safety net. The federal SSI program is a uniform federal benefit with a uniform federal disability standard, but the SSI-to-Medicaid path varies sharply across state lines. The 11 states that elected the §1902(f) option have, over half a century, accumulated a measurable population-health penalty for that choice. The present paper is the first to triangulate the penalty across three independent data sources; it is offered to the policy and research communities as the historical baseline against which any modern reform of the SSI-to-Medicaid path can be evaluated.

References

Appendix

Appendix Table A1. Sample size by component

Component	Cells	Notes
Border-pair panel skeleton	51,198	483 pairs × 2 sides × 53 years
Mortality DiD estimating sample	39,980	WONDER non-missing, 1968–2020
CPS state-year sample (coverage/take-up)	1,089	Disabled-adult subsample, 1988–2020
CPS state-year sample (prevalence)	1,557	All adults 18–64, 1968–2020
Pre-1972 placebo (1968–1971)	3,016	Mortality
SSA county-year panel (parsed)	77,407	1998–2024 (estimating window 1998–2020 due to WONDER denominator), 50 states
SSA county-year panel (FIPS-matched)	77,331	99.9% match rate
SSA border-pair estimating sample	17,330	1998–2020, 482 pairs, 362 counties

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.

Appendix Table A2. 209(b) state classification (CRS R46111, 2019)

State	209(b) status	Notes
Connecticut	209(b) (continuous since 1974)	In border-pair sample
Hawaii	209(b) (once 209(b), now SSI)	Excluded — no contiguous border
Illinois	209(b) (continuous)	In sample
Indiana	209(b) (continuous)	In sample
Minnesota	209(b) (continuous)	In sample
Missouri	209(b) (continuous)	In sample
New Hampshire	209(b) (continuous)	In sample
North Dakota	209(b) (continuous)	In sample
Ohio	209(b) (continuous)	In sample
Oklahoma	209(b) (continuous)	In sample
Virginia	209(b) (continuous)	In sample
Iowa	Configurable	Treated as SSI in headline; tested in LOSO

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

Appendix Table A3. Honest-DiD relative-magnitudes sensitivity

Outcome	M	Max pre-period $\lvert\beta\rvert$	Min robust lower bound
Mortality per 1k	0.5	0.225	−0.709
Mortality per 1k	1.0	0.225	−0.821
Mortality per 1k	1.5	0.225	−0.934
Mortality per 1k	2.0	0.225	−1.047

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.

(See `analysis/tables/honest_did_sensitivity.csv`. Wide bounds reflect mixed-sign pre-period coefficients; we report transparently as a caveat.)

Appendix Figure A1. Robustness summary

The supporting figures referenced in this paper include the RD running-variable density

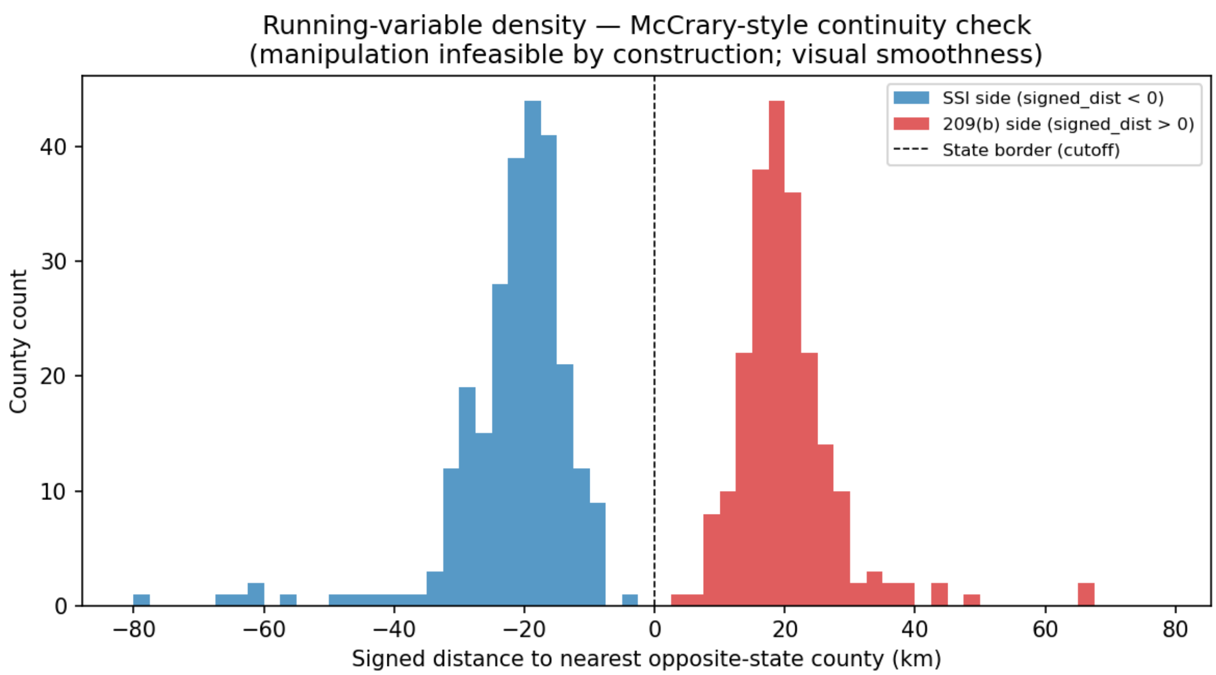

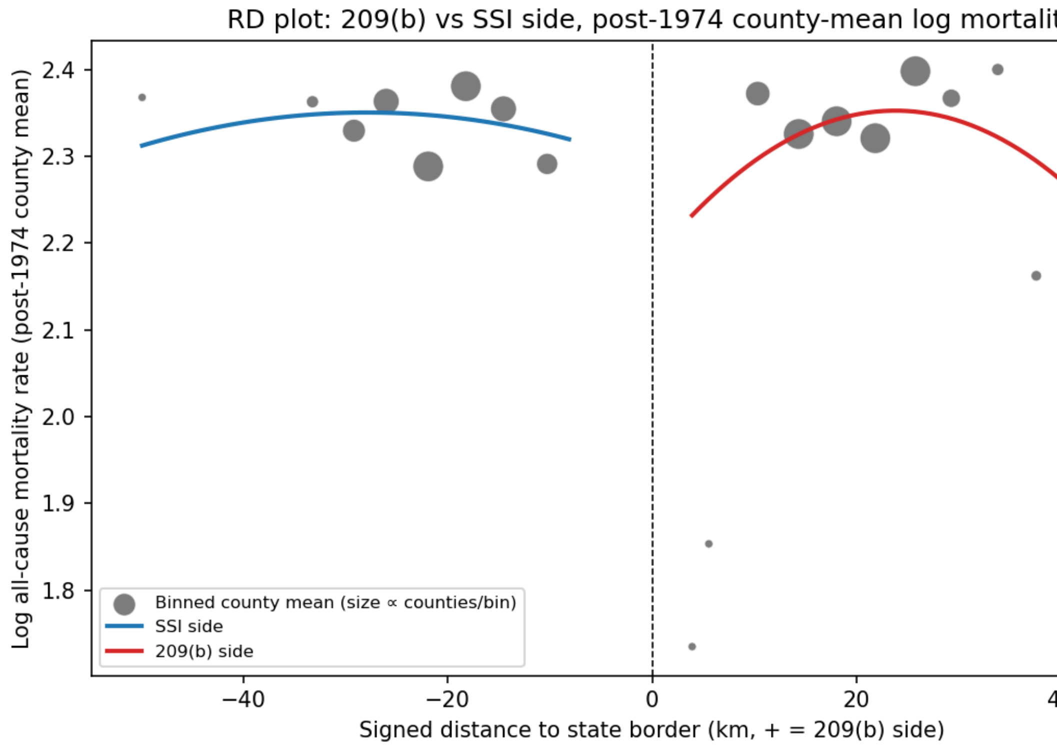

, the RD binned-mean plot

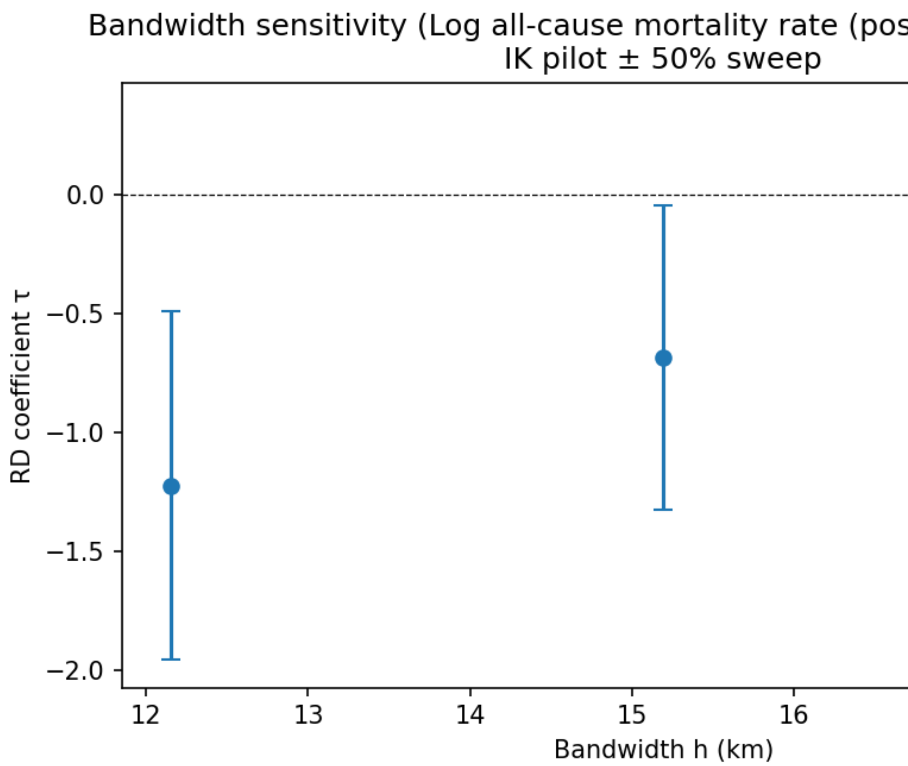

, the RD bandwidth sensitivity sweep

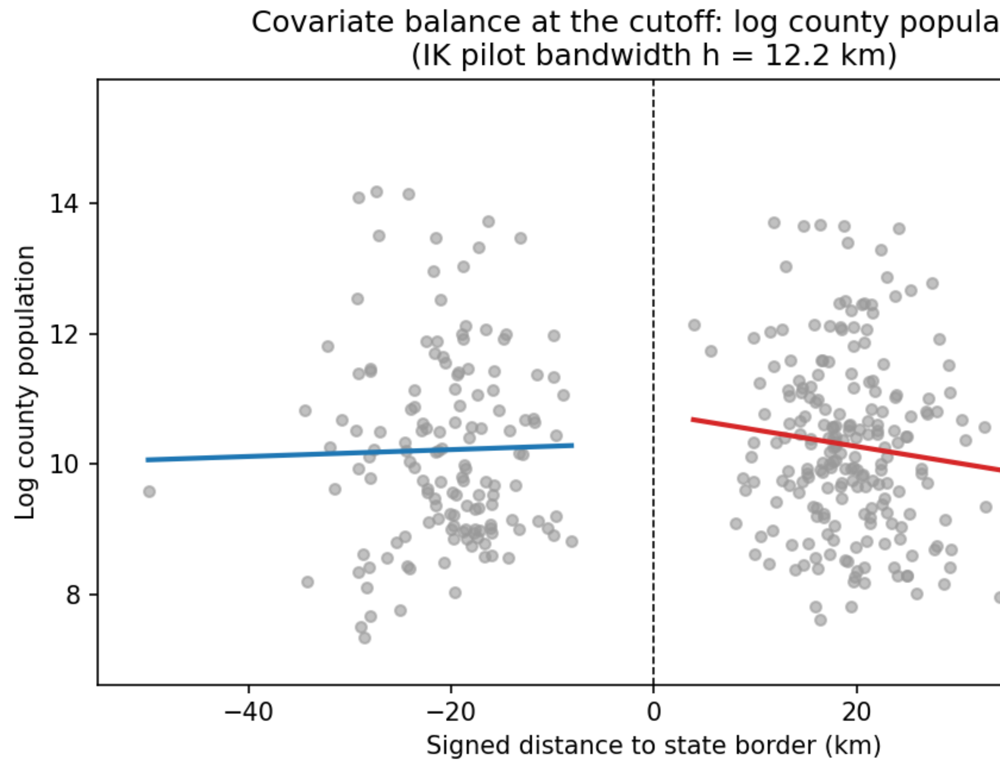

, the RD covariate balance plot

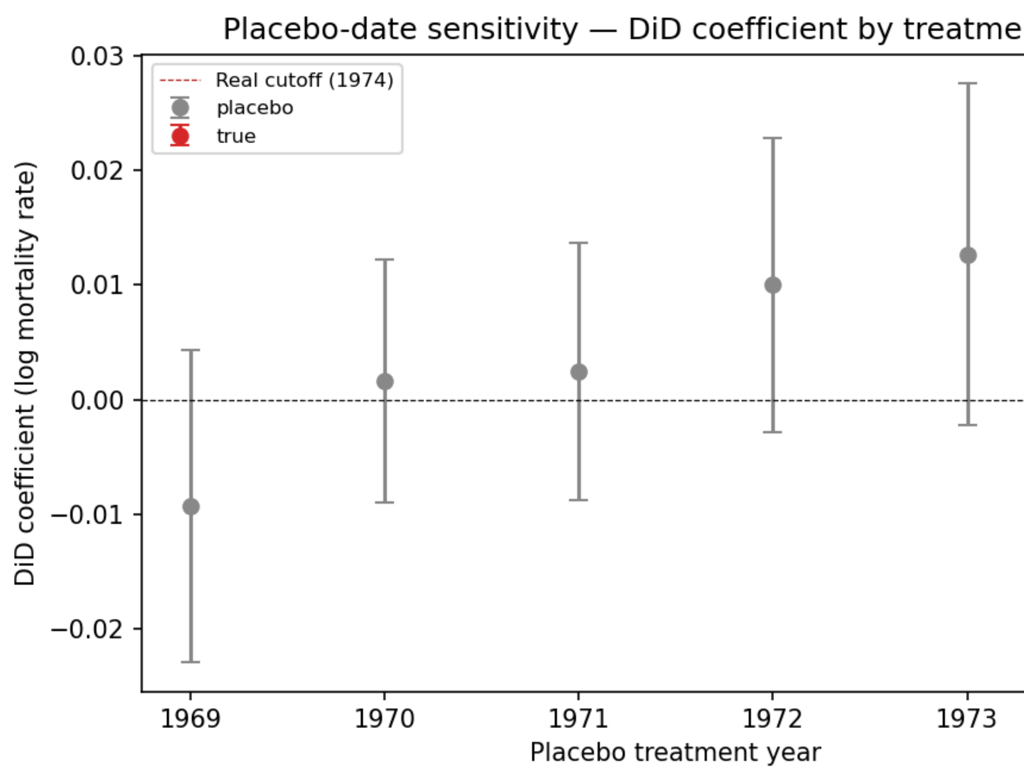

, the RD placebo-dates figure

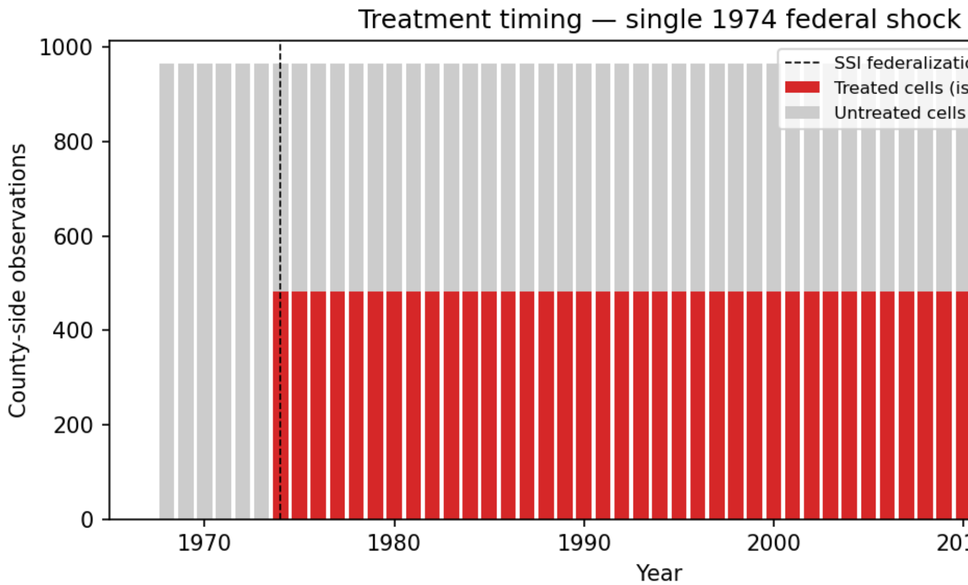

, and the treatment-timing figure

.