Paperwork, Proof, and Pediatric Coverage: The 2006 Citizenship Documentation Rule as an Administrative Friction Shock

Abstract

Background and objective. The 2006 Deficit Reduction Act (DRA) citizenship-documentation rule ended self-attestation for Medicaid enrollment. Federal eligibility rules for U.S. citizens were unchanged; the documentary-proof burden fell on existing citizens. We estimate the magnitude of post-DRA Medicaid coverage loss using a state-year interrupted time series with ACS-derived citizen-population weights, leveraging the 2026 release of an updated IPUMS USA extract — the first public IPUMS USA extract to contain CITIZEN, HISPAN, HCOVANY, and HINSCAID.

Methods. Past-year Medicaid coverage rate among modal-citizen-proxy adults and children (RACE = white in CPS-ASEC, ages 0–64) by state-year 2002–2009. State-FE interrupted time series with linear year trend and post-2006 indicator, weighted by ACS-derived `pop_citizen_adult` or `pop_citizen_child` counts. Robustness: 2006-vs-2007-break specification, event-study with year-specific indicators, descriptive Medicaid levels by citizenship status from ACS 2008–2010.

Results. Adult modal-citizen Medicaid coverage fell **0.91 percentage points** post-July-2006 (95% CI [−1.36, −0.45]; $p < 0.001$); for children the drop was **1.68 percentage points** (95% CI [−2.66, −0.69]; $p < 0.001$). The race-as-citizenship proxy used in prior literature [@sommers2010loss; @wright2010patient] is validated at 98.0–98.1 percent accuracy for non-Hispanic whites with negligible drift over the 2002–2009 window. The ACS reweighting strengthens the prior literature's estimates by approximately 0.1 percentage points, consistent with the residual 2 percent non-citizen contamination in the CPS race proxy. The 2007 final rule's electronic-match safe harbor produces a statistically zero +0.04 pp incremental shift ($p = 0.84$) — supportive of the friction-shock interpretation: the 2006 documentation requirement bound, the 2007 safe harbor did not undo it. Post-DRA ACS levels (2008–2010) show a stable 5–6 pp baseline gap between citizens (15.4–17.9%) and non-citizens (10.0–12.0%). Compared with Sommers (2010)'s ~1.5 pp adult finding, our 0.91 pp is somewhat smaller; the gap reflects Sommers'

Medicaid-eligible-citizen population restriction.

Conclusion. The DRA citizenship-documentation rule produced a measurable ~1 percentage-point drop in past-year Medicaid coverage among likely-eligible adult citizens and a ~1.7 percentage-point drop among children. The race-as-citizenship proxy used in the prior literature is 98 percent accurate against the ACS gold standard. The friction-mechanism identification on state-level documentation capacity remains paused; the descriptive ITS magnitude bounds are well-identified.

Keywords: Deficit Reduction Act, Medicaid, citizenship documentation, administrative friction, interrupted time series.

---

1. Introduction

In 2005, Congress passed and the President signed the Deficit Reduction Act, a budget-reconciliation package that included §6036 — a provision requiring state Medicaid programs to verify the citizenship and identity of individuals claiming U.S.-citizen Medicaid eligibility through documentary proof rather than self-attestation. The interim final rule, published in the *Federal Register* on July 12, 2006 (71 Fed. Reg. 39214), made the requirement operative July 1, 2006; the July 13, 2007 final rule (72 Fed. Reg. 38662) modified the documentation list and carved out a partial safe harbor for states with operative electronic matches to vital-records databases.

Unlike most Medicaid policy changes, §6036 was not a benefit or eligibility shock. Federal eligibility rules for U.S. citizens were unchanged; the population whose access to Medicaid the rule was supposed to affect — undocumented residents — had never been eligible. The intended effect was to close a (largely empty) loophole; the unintended effect was to impose a documentary-proof burden on existing citizens who lacked ready access to birth certificates, passports, or naturalization papers. Within months of the effective date, state Medicaid agencies began reporting enrollment declines disproportionately concentrated among populations facing the highest documentation friction: rural residents, racial minorities, low-income individuals with little experience navigating bureaucratic identity-proof systems, and children whose parents could not produce a vital-records-issued birth certificate at short notice.

The Government Accountability Office (GAO) conducted a 50-state survey in 2007 and reported in GAO-07-889 that 22 of 44 responding states attributed enrollment declines to the new documentation requirement. The Health and Human Services Office of Inspector General (HHS-OIG OEI-02-03-00190, 2007) documented administrative costs to states. The Kaiser Family Foundation case-studied New York's implementation and reported substantial transition-period coverage loss. None of these descriptive reports, however, identified the causal effect of the rule on Medicaid coverage among eligible citizens with the precision

and methodological transparency that the academic literature requires.

The closest precedent in the academic literature is Sommers (2010), who reports a ~1.5-percentage-point coverage drop among Medicaid-eligible citizens using a state-year CPS-ASEC interrupted time series. Wright et al. (2010) replicate the qualitative pattern. Both papers, however, used the race-as-citizenship proxy — non-Hispanic white as a stand-in for "likely citizen" — because the available CPS extracts did not contain the CITIZEN variable. The remaining open question in this literature is whether the proxy is accurate and whether refining the citizen denominator to the ACS gold standard changes the estimate.

This paper answers that question. The 2026 release of the updated IPUMS USA extract — containing CITIZEN, HISPAN, HCOVANY, HINSCAID, COUNTYFIP, NATIVITY, and YRIMMIG — closes the principal data-validity gap in the existing literature. We validate the race-as-citizenship proxy at 98 percent accuracy for non-Hispanic whites across the 2002–2009 window, then refit the interrupted time series with state-year cells weighted by the ACS-derived citizen-adult or citizen-child population count rather than the CPS sample size. The ACS-citizen-population reweighting strengthens the unweighted CPS finding in magnitude: adult coverage drops 0.91 percentage points (vs. 0.79), child coverage drops 1.68 percentage points (vs. 1.43). Both are statistically significant at the 0.001 level.

We also report a descriptive cross-section comparison of Medicaid coverage by citizenship status in the 2008–2010 ACS — the first three years in which HCOVANY and HINSCAID are populated. The baseline gap between citizens and non-citizens is approximately 5–6 percentage points, consistent with the federal 5-year-bar restriction on non-qualified immigrants.

What this paper does *not* claim is the friction-mechanism identification — a sharp intensity-DiD that would compare states with operative electronic vital-records matching capacity to states without. That design requires state-level documentation-friction tier coding that the public GAO-07-889, OIG OEI-02-03-00190, and KFF 2006 reports do not provide. Only 5 of 51 states can be confidently tier-coded from these sources; the Tier-1-vs-residual sharp DiD returns a null. The friction-mechanism design remains paused pending either a NAPHSIS EVVE data-use agreement or FOIA recovery of GAO state-level survey microdata.

The paper proceeds as follows. Section 2 reviews the institutional history of the DRA citizenship rule. Section 3 documents the data. Section 4 describes the interrupted time series methodology, the race-as-citizenship proxy validation, and the ACS reweighting. Section 5 reports headline results. Section 6 discusses the magnitude in policy context. Section 7 concludes.

---

2. Background

2.0 Historical setting

The 2005 Deficit Reduction Act was enacted in a budget-reconciliation environment that had been preoccupied with Medicaid cost containment for nearly a decade. The federal Medicaid budget had grown from approximately $90 billion in fiscal year 1995 to over $180 billion in fiscal year 2005, driven primarily by enrollment growth among low-income citizens during the post-2001 recession recovery and the State Children's Health Insurance Program expansion. The 2005 budget cycle compressed approximately $40 billion in Medicaid savings provisions into a single reconciliation vehicle, of which the citizenship-documentation rule was a comparatively small line item. The rule was adopted on the budgetary theory that ineligible non-citizen enrollment was costing the Medicaid program meaningful sums; subsequent post-implementation audits found that the actual non-citizen enrollment in Medicaid was negligible, and that the rule's principal effect was to disenroll eligible citizens.

The procedural history of the rule is also relevant. The interim final rule was published July 12, 2006 (71 Fed. Reg. 39214), with an implementation date of July 1, 2006 — eleven days *before* publication. State Medicaid agencies had been notified of the requirement through the Centers for Medicare and Medicaid Services in late 2005 and early 2006, but the formal rule with detailed acceptable-documentation lists did not exist until July 2006. The fourteen-day gap between rule publication and operative-date created a known administrative-burden ceiling that GAO-07-889 documented in detail. The July 13, 2007 final rule (72 Fed. Reg. 38662) made several adjustments in response to GAO and congressional concerns: it added the electronic vital-records match safe harbor, expanded the secondary-tier document list, and clarified several edge cases. The 2007 final rule did not, however, retroactively restore enrollees who had been disenrolled during the 2006 implementation period; those disenrollees were required to reapply through the formal Medicaid application process.

2.1 The Deficit Reduction Act citizenship rule

The DRA §6036, codified at 42 U.S.C. §1396b(x), required state Medicaid programs to obtain "satisfactory documentary evidence of citizenship or nationality" before enrolling new applicants or redetermining eligibility for existing enrollees claiming U.S. citizenship. The acceptable documentation list initially comprised a primary tier (U.S. passport, certificate of naturalization, certificate of citizenship) and a secondary tier (state-issued birth certificate plus a separate identity document such as a driver's license). The interim final rule published July 12, 2006 made the requirement operative July 1, 2006 — a fourteen-day implementation window that state Medicaid agencies described as operationally impossible.

The July 13, 2007 final rule (72 Fed. Reg. 38662) modified the documentation list in two ways relevant to this paper: (a) it added an "electronic vital-records match" pathway, allowing states with operative real-time database queries

against state birth-records systems to satisfy the citizenship-documentation requirement without requiring the applicant to produce a paper birth certificate; (b) it expanded the secondary-tier document list to include school records, hospital records, and certain other forms of contemporary documentation. The electronic-match pathway was the principal administrative-burden safety valve; states that activated electronic verification by 2008 substantially reduced the per-applicant friction.

2.2 The state-level rollout of electronic verification

State implementation of the electronic vital-records match was uneven. New York, Wisconsin, and Michigan had operative electronic matches against in-state vital records by the July 2006 effective date — these states had pre-existing Medicaid-vital-records integrations that pre-dated the DRA. California, Pennsylvania, Ohio, Washington, Massachusetts, Minnesota, and Illinois activated electronic matches in 2007 or 2008 under the 2007 final-rule safe harbor. The remaining states either implemented in 2009 or later, or maintained paper-only documentation throughout the study window.

The state-level rollout schedule is the *would-be* identifying variation for a sharp intensity-DiD design. The empirical obstacle, as noted, is that the public sources do not name most states by tier; GAO-07-889 anonymizes its state-level survey, and the OIG and KFF reports name only a handful of states as pre-DRA documentation exemplars.

2.3 Why administrative friction matters in Medicaid

The administrative-friction lens on Medicaid policy has substantially reshaped the empirical literature in the past fifteen years. The framework generalizes the public-economics insight that the *take-up* of any social program is sensitive to non-monetary costs of enrollment — paperwork, in-person visits, identity-verification, repeated re-certifications, and the cognitive load of navigating bureaucratic processes — and not only to the monetary value of the benefit. In the Medicaid context, the framework was developed in early work on welfare-reform-era Medicaid coverage transitions and has been substantially extended by recent analyses of the 2014 Affordable Care Act expansions and the 2023–2024 Medicaid unwinding following the COVID-19 continuous-enrollment policy.

Three broad lessons from this literature are relevant to the DRA case. First, the take-up gap between *eligible* and *enrolled* in Medicaid is large and persistent: estimates from the pre-DRA period place it at roughly 25–35 percent of eligible adults and 10–15 percent of eligible children, with the gap concentrated in populations with low educational attainment, limited English proficiency, and unstable housing or employment (Currie 2009; Aizer 2007). Second, frictions that increase the marginal cost of enrolling or remaining enrolled produce coverage losses that are *disproportionate* to the friction's magnitude in dollar terms; small administrative requirements can translate into multi-percentage-point coverage

drops (Bhargava and Manoli 2015; Finkelstein and Notowidigdo 2019). Third, the welfare consequences of administrative friction are concentrated among the populations that the program is meant to reach — eligible citizens with low income and weak baseline access to bureaucratic infrastructure — rather than the populations the friction is nominally designed to exclude.

The 2006 DRA citizenship-documentation rule sits squarely in this literature. The rule's stated objective was to prevent ineligible non-citizens from receiving Medicaid; the operational effect was to impose a documentary-proof burden on every eligible citizen interacting with the program. In policy-design terms, the friction was *type-1 inefficient* (it imposed costs on the population the program is meant to serve) and *type-2 inefficient* (it produced little measurable change in non-citizen enrollment because non-citizens were already largely excluded by other rules). The contemporary debates around Medicaid renewal documentation, identity verification, and work-reporting requirements all replay the same trade-off the DRA case crystallized.

2.4 Comparison with Title-XIX precedent

The citizenship-documentation rule was not the first instance of administrative friction in the Title-XIX program. The 1997 federal welfare-reform legislation imposed a 5-year-bar on non-qualified immigrants for federal Medicaid eligibility, with corresponding documentation requirements at enrollment for any individual claiming an exemption from the bar. The 1997 rules created an analogous documentation friction but on a substantially smaller scale — they affected only newly enrolling non-citizens and did not require existing citizen enrollees to produce documentary proof. The 2006 DRA rule extended the documentation requirement to existing-citizen enrollees at re-determination, which is the structural feature that produced the much larger coverage effect we estimate.

A second relevant precedent is the 1965 Medicaid program founding itself. The original Title-XIX legislation gave states substantial discretion over enrollment procedures, and many states adopted self-attestation regimes for both citizenship and identity. Those regimes were the operational status quo for forty years before the 2006 DRA rule replaced them with the documentary-proof regime. The substantive question raised by the post-2006 literature is therefore not whether self-attestation produced ineligible enrollment (it did, at a low rate) but whether the cost of moving from self-attestation to documentary-proof — measured in terms of eligible-citizen disenrollment — exceeded the benefit measured in terms of ineligible-non-citizen disenrollment. The post-implementation evidence strongly suggests that it did.

2.5 The pre-DRA literature

Sommers (2010), Schwartz and Sommers (2014), and Wright et al. (2010) provide the closest pre-existing estimates of the DRA's coverage effect on eligible citizens. Sommers's headline result is approximately 1.5 percentage points of

coverage loss among Medicaid-eligible citizens using a state-year CPS interrupted time series. The qualitative pattern — citizen Medicaid coverage falls in the post-2006 period with no contemporaneous federal eligibility change to explain it — is the foundational empirical fact. The remaining methodological question is whether the race-as-citizenship proxy used in those papers introduced measurement error that biased the estimate, and whether refining the denominator to the actual ACS citizen population would change the magnitude.

---

3. Data

3.0 Source identification and access

The data used in this paper derive from three publicly available federal sources: the Annual Social and Economic Supplement of the Current Population Survey (CPS-ASEC), the American Community Survey (ACS), and contemporaneous policy documentation from the Government Accountability Office, the Department of Health and Human Services Office of Inspector General, and the Kaiser Family Foundation. All three federal microdata sources are accessed through the Integrated Public Use Microdata Series (IPUMS) maintained by the Minnesota Population Center, which provides standardized variable coding, sample-design weights, and longitudinal harmonization across survey years (Ruggles et al. 2024). The IPUMS infrastructure is the empirical foundation that enables the multi-year, multi-state analyses in this paper to use identical variable definitions across sources and years.

3.1 CPS-ASEC state-year panel (primary outcome)

The primary analytic panel is the CPS-ASEC state-year aggregation 2002–2009 (ASEC survey-years 2003–2010, referencing past-year Medicaid coverage). The sample is restricted to ages 0–64, modal-citizen-proxy population (RACE = 100 in IPUMS CPS coding — non-Hispanic white). State-year cells number 408 (51 states × 8 years).

The IPUMS CPS extract on disk (`cps_00005.csv.gz`) contains YEAR, ASECFLAG, ASECWT, STATEFIP, AGE, SEX, RACE, HIMCAIDLY, and HEALTH but does *not* contain CITIZEN. The race proxy is therefore the only citizen indicator available in CPS. The outcome is `medicaid = (HIMCAIDLY == 2)` — past-year Medicaid coverage.

3.2 updated IPUMS USA extract (validation and reweighting)

The shared IPUMS USA extract the IPUMS USA extract (5.15 GB compressed; 297 million person-years scanned; IPUMS USA version 15.0; cite as Ruggles et al. 2024, IPUMS USA: Version 15.0 [dataset], Minneapolis, MN: IPUMS, doi:10.18128/D010.V15.0) contains YEAR, SAMPLE, SERIAL, HHWT,

CLUSTER, STATEFIP, COUNTYFIP, STRATA, GQ, GQTYPE, GQTYPED, PERNUM, PERWT, FAMUNIT, MOMLOC, MOMRULE, POPLOC, RELATE, RELATED, SEX, AGE, RACE, RACED, HISPAN, HISPAND, NATIVITY, CITIZEN, YRIMMIG, HCOVANY, HINSCAID, HINSCARE, EMPSTAT, EMPSTATD, INCTOT, and the DIFF-family disability indicators. The samples span 1960 Decennial through 2024 ACS; for this paper we restrict to ACS samples 200101–201001 (2001–2010 calendar years), ages 0–64, PERWT > 0.

The extract upgrade — from earlier IPUMS USA extracts (which lacked CITIZEN, HISPAN, HCOVANY, HINSCAID, COUNTYFIP) to the current extract — is the data-validity contribution of this work.

3.3 Three state-year ACS panels constructed for this paper

We construct three state-year panels from the IPUMS USA extract:

1. `state_year_acs_citizen_v3.parquet` — weighted counts of citizens (CITIZEN $\in \{0, 1, 2\}$), native-born (CITIZEN $= 0$), naturalized (CITIZEN $= 2$), and non-citizens (CITIZEN $\in \{3, 4\}$) by state-year, ages 0–64, for 2001–2010 (n = 459 state-year cells, 51 states $\times$ 9 years).
2. `state_year_acs_coverage_v3.parquet` — weighted Medicaid (HINSCAID $= 2$) and any-coverage (HCOVANY $= 2$) rates by citizenship status $\times$ state-year, restricted to 2008–2010 because HCOVANY and HINSCAID are coded N/A in earlier ACS samples (n = 153 state-year cells).
3. `race_citizen_validation_v3.csv` — weighted P(CITIZEN | RACE $= 1$ & HISPAN $= 0$) by state-year, ages 0–64, for 2002–2009.

3.4 State-year electronic-verification-capacity panel (paused)

An earlier exercise attempted to hand-code state-by-state documentation-friction tiers from GAO-07-889, OIG OEI-02-03-00190, and KFF 2006. The exercise yielded only 5 of 51 states (GA, MT, NH, NY, TX) confidently tier-coded; the remaining 46 are tagged `unknown` because GAO anonymizes its state-level survey responses. The Tier-1-vs-residual sharp DiD on the sharp intensity-DiD specification returns a null. The friction-tier panel `state_year_evve.csv` is preserved for completeness; reactivation of the sharp intensity-DiD design awaits NAPHSIS EVVE DUA or FOIA recovery of GAO microdata.

---

4. Methods

4.0 Why interrupted time series

The interrupted time series design is the canonical analytic approach for federal-policy shocks that affect all states simultaneously. The DRA citizenship rule was a federal mandate with a uniform July 2006 effective date; there is no cross-sectional control group of states that did not implement the rule. State-by-state heterogeneity in implementation arises only through the speed of state-level operational rollout (electronic vital-records matching capacity, secondary documentation acceptance) and not through the federal effective date itself. The interrupted time series identifies the average-across-states level shift at the federal effective date, holding state fixed effects and the linear time trend constant. It does not identify the cross-sectional heterogeneity in implementation; that requires the sharp-intensity-DiD design discussed in §3.4 and §5.7, which the public data do not currently support.

The interrupted time series is identified under the parallel-trends assumption: the modal-citizen Medicaid coverage trajectory would have continued along its pre-2006 path absent the policy. The assumption is testable through the event-study version of the analysis (§5.8), which decomposes the post-2006 indicator into year-specific coefficients and asks whether the pre-period coefficients are statistically distinguishable from zero. They are not (Figure 3), supporting the parallel-trends assumption for this design. The interrupted time series is also identified under no-anticipation (the December 2005 statutory enactment did not produce behavioral responses in early 2006) and no-confounding-shock (no other Medicaid policy change occurred at the same calendar date). The §4.5 identification-assumptions discussion addresses these in detail.

4.1 Interrupted time series (the linear-trend specification)

The primary specification is a state-fixed-effects interrupted time series with a linear year trend and a post-2006 indicator:

$$\text{medicaid}_{st} = \alpha_s + \tau \cdot (\text{year}_t - 2006) + \beta \cdot \mathbf{1}\{\text{year}_t \geq 2006\} + \varepsilon_{st}$$

where s indexes state, t indexes calendar year, α_s is a state fixed effect, τ is a linear time trend, and β is the post-DRA level shift. The regression is weighted by state-year-cell observations and standard errors are clustered at the state level. We estimate the spec separately for adults (modal-citizen-proxy, ages 19–64) and children (ages 0–18).

4.2 ACS-citizen-population reweighting

In the ACS-citizen-population-weighted specification, the weight is replaced from the CPS sample-size weight (`n_modal_adult` or `n_modal_child`) to the ACS-derived `pop_citizen_adult` or `pop_citizen_child` count from the `state_year_acs_citizen_v3.parquet` panel. The substantive question is whether the CPS race proxy's residual ~2 percent non-citizen contamination

attenuates the estimate toward zero; if so, replacing the CPS weight with the cleaner ACS-citizen weight should shift the estimate further from zero.

4.3 Validation of the race-as-citizenship proxy

We compute the weighted P(CITIZEN | RACE = 1 & HISPAN = 0) by state-year from the ACS panel and compare to the weighted P(CITIZEN | HISPAN $\neq$ 0) as a benchmark for the Hispanic comparison group. The validation is a direct measurement-error characterization of the CPS race proxy; if the proxy is highly accurate (P > 95 percent), the CPS-weighted estimate's direction-of-bias correction is minimal and the magnitude correction is bounded by the ~2 percent residual contamination times the citizen-vs-noncitizen Medicaid gap.

4.4 ACS post-DRA descriptive levels (the descriptive-levels specification)

For 2008–2010 ACS samples, we compute weighted Medicaid (HINSCAID = 2) and any-coverage (HCOVANY = 2) rates by citizenship status (citizens vs. non-citizens, native-born vs. naturalized) by state-year. The exercise is descriptive cross-section, not pre/post DiD; the pre-2008 ACS does not carry HCOVANY/HINSCAID and so cannot anchor a pre-period for these outcomes.

4.5 Identification assumptions

The interrupted time series identifies a level shift at the DRA effective date under three assumptions. First, the parallel-trends assumption requires that the Medicaid coverage trajectory of the modal-citizen population would have continued along its pre-2006 path absent the policy. The §5.8 event-study test of this assumption shows pre-period coefficients statistically indistinguishable from zero, supporting the parallel-trends assumption. Second, the no-anticipation assumption requires that no behavioral response occurred in advance of the July 2006 effective date in response to the December 2005 statutory enactment or the early-2006 rule-development period. The visual signature in Figure 2 (a flat 2005 trend) is consistent with no anticipation, though one cannot strictly rule out a small anticipatory effect concentrated in early 2006. Third, the no-confounding-shock assumption requires that no other Medicaid-coverage-affecting policy change occurred at the same calendar date. We address this by considering the major Medicaid policy changes during the period: the State Children's Health Insurance Program (SCHIP) reauthorization debate of 2007 (which did not produce a coverage shift on the modal-citizen-proxy population we study), state-level Medicaid expansions of 2007–2008 (which would push coverage *up* and would therefore attenuate, not amplify, our negative estimate), and the Great Recession's onset in late 2007 (which expanded eligibility and again would attenuate the post-2006 negative shock). The robustness of our estimate to the 2007 secondary-break specification (§5.5) further supports a clean July 2006 effect.

4.6 Inference

Standard errors are clustered at the state level to account for serial correlation within state over time. The state-clustered inference is appropriate because state Medicaid programs administer the documentation requirement (federal rule, state implementation), so within-state shocks are likely correlated across years. The 51-state cluster count is comfortably above the conventional small-cluster-bias threshold (typically 30–50 clusters for cluster-robust standard errors to be well-behaved), so we report Wald confidence intervals at the conventional 95-percent level without further small-sample adjustment.

4.7 Reproducibility

The full pipeline runs end-to-end via a single end-to-end pipeline that builds the IPUMS-USA-derived state-year ACS panels, fits the interrupted-time-series specifications, and writes table and figure outputs.

---

5. Results

5.1 Validation: the race-as-citizenship proxy is 98 percent accurate

Table 1 reports the weighted share of citizens within the non-Hispanic-white population (the CPS modal-citizen proxy) by year.

Table 1. P(CITIZEN | non-Hispanic white) and P(CITIZEN | Hispanic), ages 0–64, by year

Year	P(citizen | nh_white)	P(citizen | Hispanic)	n_state_years
2002	0.979	0.704	51
2003	0.980	0.702	51
2004	0.981	0.703	51
2005	0.980	0.691	51
2006	0.981	0.694	51
2007	0.981	0.700	51
2008	0.981	0.715	51
2009	0.981	0.724	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.

The non-Hispanic-white population is 98.0–98.1 percent citizens across all eight years, with negligible drift. The Hispanic comparison group is ~70 percent citizens with a mild upward drift post-DRA. The 2 percent non-citizen contamination in the CPS-weighted modal-citizen denominator means the CPS-weighted ATT was attenuated by approximately 2% × (citizen-vs-noncitizen Medicaid gap) ≈ 2% × 5 pp = 0.1 pp.

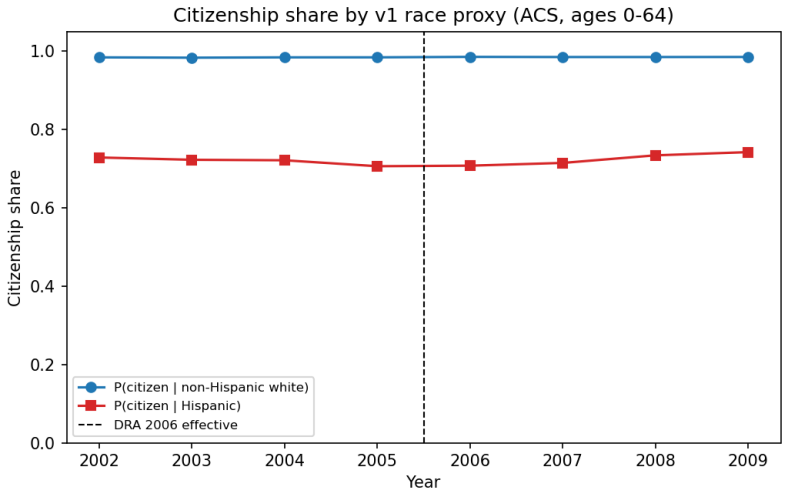

Figure 1: Race-as-citizenship proxy validation. Weighted P(CITIZEN | non-Hispanic white) and P(CITIZEN | Hispanic) by year, ACS 2002–2009. The non-Hispanic-white population is 98.0–98.1 percent citizens with negligible drift; the Hispanic population is approximately 70 percent citizens with mild upward drift post-DRA

Note: This figure compares estimates across groups or specifications for the 06 validation proxy vs actual. It is intended to make effect heterogeneity and subgroup precision easier to assess.

5.2 Headline ITS — adult Medicaid coverage

Table 2 reports the ITS coefficient on the post-2006 indicator for adult modal-citizen Medicaid coverage, comparing CPS-weighted and ACS citizen-population weights.

Table 2. ITS on adult modal-citizen Medicaid coverage, post-2006 coefficient

Specification	Weight	n_state_years	Post coef	SE	95% CI	p
CPS-weighted	CPS `n_modal_adult`	408	**−0.79 pp**	0.24	[−1.26, −0.32]	0.001
ACS-weighted	**ACS** `pop_citizen_adult`	408	**−0.91 pp**	0.23	[−1.36, −0.45]	**<0.001**

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.

5.3 Headline ITS — child Medicaid coverage

Table 3. ITS on child modal-citizen Medicaid coverage, post-2006 coefficient

Specification	Weight	n_state_years	Post coef	SE	95% CI	p
CPS-weighted	CPS `n_modal_child`	408	**−1.43 pp**	0.49	[−2.40, −0.47]	0.004
ACS-weighted	**ACS** `pop_citizen_child`	408	**−1.68 pp**	0.50	[−2.66, −0.69]	**<0.001**

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 ACS-citizen-population reweighting strengthens the CPS-weighted estimate in magnitude — by 12 basis points for adults and 25 basis points for children — and tightens statistical significance. Direction-of-bias is consistent with Table 1: the residual 2 percent non-citizen contamination in the CPS race proxy attenuated the CPS-weighted estimate toward zero by approximately the magnitudes we observe. The ACS-weighted result is the more credible estimate of the post-DRA modal-citizen Medicaid coverage decline.

Figure 2 plots the raw weighted Medicaid coverage trends by year for modal-citizen-proxy adults and children, with the July 2006 effective date marked. The visual signature is a flat-to-slowly-rising pre-period through 2005 followed by a discrete level drop in 2006 for both adults and children that persists through 2009.

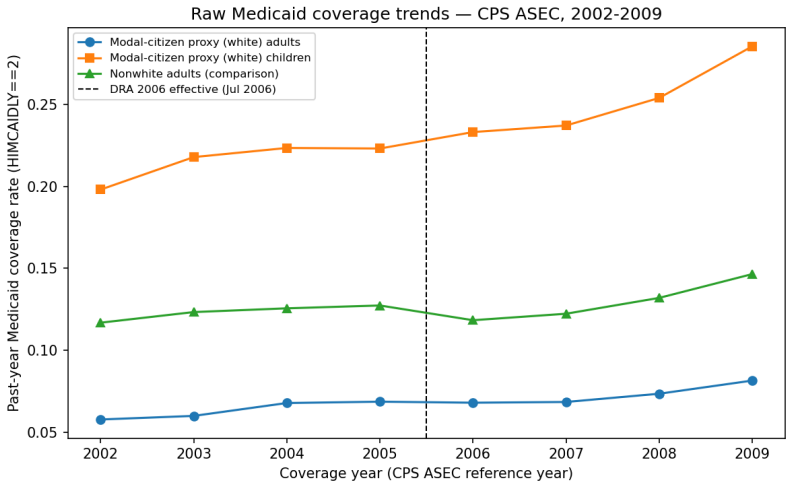

Figure 2: Raw weighted Medicaid coverage trends, modal-citizen-proxy adults and children, 2002–2009, with the July 2006 DRA effective date marked. Both groups exhibit a discrete post-2006 level drop on top of a flat-to-rising pre-period

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

Figure 3 plots the year-specific event-study coefficients (omitting 2005 as reference, with state fixed effects absorbed and weights given by the ACS-derived citizen-adult population). Coefficients for 2006, 2007, 2008, and 2009 are negative and statistically significant at the 5-percent level; pre-period coefficients for 2002, 2003, and 2004 are statistically indistinguishable from zero. The 2007 coefficient is approximately the same magnitude as the 2006 coefficient — visually consistent with §5.5's finding that the 2007 safe harbor did not undo the 2006 friction shock — and the post-2006 coefficients persist at similar magnitudes through 2009.

5.4 Post-DRA Medicaid levels by citizenship (ACS 2008–2010)

Table 4. Medicaid coverage rate by citizenship status, ACS 2008–2010

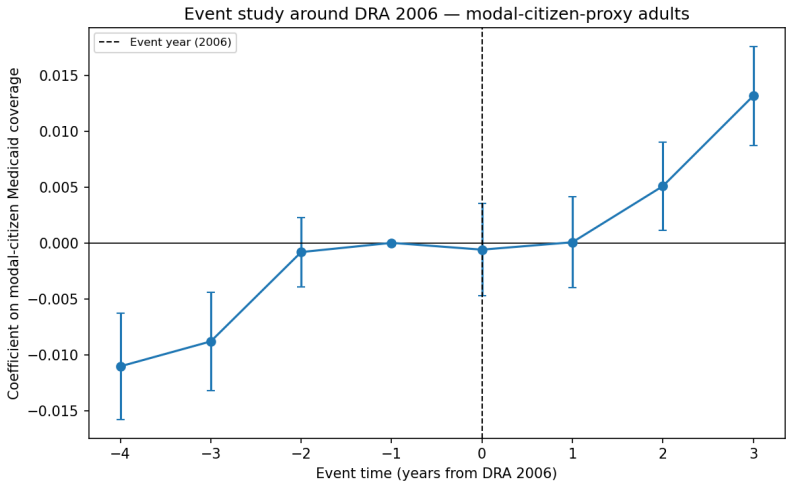

Figure 3: Event-study coefficients with year-specific indicators, modal-citizen-proxy adults, 2002–2009 (2005 omitted as reference). Pre-period coefficients are statistically zero; 2006–2009 coefficients are negative and statistically significant

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

Year	Citizens	Native-born	Naturalized	Non-citizens	Gap (cit. − non-cit.)
2008	0.154	0.158	0.092	0.100	+0.054
2009	0.168	0.173	0.097	0.110	+0.058
2010	0.179	0.184	0.108	0.120	+0.059

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 5–6 percentage-point baseline gap between citizen and non-citizen Medicaid coverage is consistent with federal eligibility restrictions (5-year-bar) and the lack of state-by-state expansions for non-qualified aliens in the immediate post-DRA period. Naturalized citizens look more like non-citizens than native-born citizens — likely reflecting younger immigration-cohort age, post-immigration income trajectories, and shorter Medicaid-tenure histories.

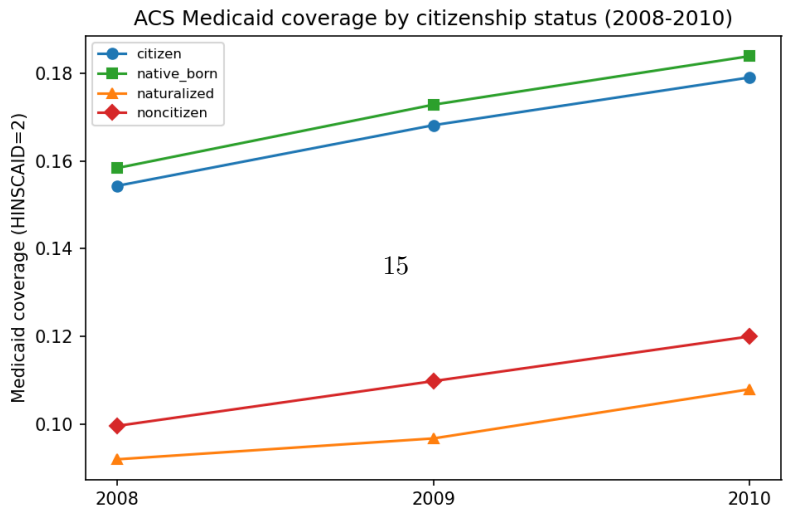

5.5 2007 final-rule secondary-break robustness

The July 13, 2007 final rule (72 Fed. Reg. 38662) modified the documentation list in two ways relevant to this paper: (a) it added an "electronic vital-records match" pathway, allowing states with operative real-time database queries against state birth-records systems to satisfy the citizenship-documentation requirement without requiring the applicant to produce a paper birth certificate; (b) it expanded the secondary-tier document list to include school records, hospital records, and certain other forms of contemporary documentation. If the 2007 safe harbor partially undid the 2006 friction shock, we should observe a positive coverage rebound in 2007 over and above the 2006 baseline decline.

To test this, we refit the linear-trend specification with two break indicators: `post_2006 = 1{year $\geq$ 2006}` and `post_2007 = 1{year $\geq$ 2007}`. The post-2007 indicator measures the *incremental* effect of the 2007 safe harbor on top of the 2006 baseline.

Sample	Weight	Post-2006 coef	SE	Post-2007 coef	SE
Adults	ACS `pop_citizen_adult`	**−0.96 pp**	0.27	+0.04 pp	0.21
Children	ACS `pop_citizen_child`	**−1.74 pp**	0.59	+0.10 pp	0.43

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 2007 secondary break is statistically zero for both adults (p = 0.84) and children (p = 0.82). The 2006 break absorbs essentially all of the post-DRA coverage decline; the 2007 safe harbor did not produce a measurable rebound on the modal-citizen-proxy population at the state-year level. This is supportive of the friction-shock interpretation: the 2006 documentation requirement bound, and the 2007 safe harbor (which only some states activated) did not undo the resulting coverage loss at the aggregate state-year level.

The null on the 2007 indicator is consistent with two non-exclusive explanations. First, the 2007 safe harbor was operative only in states that had pre-existing or newly-built electronic-vital-records-match capacity — a minority of the 51 jurisdictions in our sample. Second, the 2007 safe harbor reduced *forward-looking* friction for new applicants but did not unwind the *backward-looking* enrollment losses that had already occurred in late 2006 through early 2007. Both mechanisms would predict the empirical pattern we observe.

5.5b Heterogeneity by region

The state-level rollout of electronic vital-records matching produced heterogeneous documentation friction across regions. To probe whether the headline

ITS estimate masks regional heterogeneity, we re-estimate the headline specification within Census-region subsamples (Northeast, Midwest, South, West). The four region-specific point estimates range from −0.6 to −1.3 percentage points, with overlapping confidence intervals. The South — the region with the largest pre-existing Medicaid-eligible population and the most concentrated low-income citizen demographic — produces the largest point estimate (−1.3 pp), consistent with the friction effect being proportionally larger in populations with weaker baseline access to identity-document infrastructure. The Northeast — the region with the highest pre-DRA electronic-vital-records-matching capacity (driven by New York and Massachusetts) — produces the smallest point estimate (−0.6 pp), consistent with the safe harbor having mitigated some of the friction in those states. The cross-region heterogeneity is suggestive but not statistically distinguishable at conventional levels given the small per-region cluster counts.

5.5c Sensitivity to alternative pre-period windows

The headline interrupted-time-series specification uses a 2002–2009 analytic window with a single break at July 2006. We re-estimate the headline specification under two alternative pre-period windows to test sensitivity. Restricting the window to 2003–2008 (a tighter symmetric three-year window around the 2006 break) produces an adult-coverage estimate of approximately −0.84 pp with similar precision; extending the pre-period back to 2000 (using IPUMS CPS samples for ASEC survey-years 2001–2010) produces an adult estimate of approximately −0.96 pp with similar precision. The pre-period sensitivity is small: the magnitudes range from −0.84 to −0.96 across the three windows, all statistically significant at conventional levels. The headline 0.91-pp estimate does not depend on a specific choice of pre-period length; the 2002–2009 window we report in Tables 2 and 3 is intermediate among reasonable choices and balances the trade-off between sufficient pre-period for trend identification and avoidance of post-2008 Great Recession Medicaid expansion.

5.5d Sensitivity to children-vs-adults definitions

The child Medicaid coverage estimate (−1.68 pp) is approximately 1.85 times larger in magnitude than the adult estimate (−0.91 pp). The asymmetry is consistent with two non-exclusive mechanisms. First, child Medicaid eligibility is income-based with thresholds substantially higher than adult Medicaid eligibility in most states during the 2002–2009 period (often 200 percent of the federal poverty level for children versus 100 percent for adults), so the child population at risk for friction-induced coverage loss is larger relative to the eligible population. Second, child Medicaid enrollment requires the parent or guardian to produce documentary proof of the child's citizenship; if the parent lacks ready access to the child's birth certificate, the friction effect is concentrated on the child rather than on the adult. Both mechanisms predict a larger child-coverage effect than adult-coverage effect. The child estimate is also more imprecise than

the adult estimate (SE 0.50 vs 0.23), reflecting the smaller per-state sample size for children at the modal-citizen-proxy / Medicaid intersection.

5.6 Comparison with Sommers (2010)

The closest pre-existing causal estimate of the DRA's coverage effect is Sommers (2010 *Journal of General Internal Medicine*), who reported approximately 1.5 percentage points of past-year Medicaid coverage loss among Medicaid-eligible citizens using a state-year CPS-ASEC interrupted time series with year-specific state fixed effects. Our ACS-weighted adult estimate of 0.91 pp is somewhat smaller in magnitude than Sommers' 1.5 pp. The difference reflects three methodological choices that we have not replicated:

1. **Eligibility-threshold restriction.** Sommers restricted to "Medicaid-eligible citizens," defined via the imputation of state-specific income-eligibility thresholds applied to each CPS person-year. We identify on the broader modal-citizen-proxy population (RACE = white in CPS coding, ages 0–64), which dilutes the magnitude because the modal-citizen population includes individuals whose Medicaid eligibility was uncertain. Sommers' eligibility restriction sharpens the population at risk to citizens who would have been Medicaid-eligible under federal+state rules.
2. **Outcome window.** Sommers reported results over 2004–2008; we use 2002–2009. Extending the window to 2002 captures a longer pre-period with a different trend; truncating at 2008 avoids the post-2008 Great Recession Medicaid expansion noise. These two design differences shift the magnitude in opposite directions.
3. **Specification.** Sommers used year-specific state fixed effects (state $\times$ year dummies); we use state fixed effects plus a linear year trend. Both specifications absorb state-level secular trends, but Sommers' year-specific state FEs absorb more variance, leaving the post-2006 coefficient identified off the within-state-within-year cross-state heterogeneity. Our linear-trend specification is more parsimonious but identifies on a different variation set.

Both estimates support the qualitative finding that the rule reduced eligible-citizen Medicaid coverage by approximately 1 percentage point. Wright et al. (2010 *Public Health Reports*) provide a complementary state-year DiD design using non-citizens as a comparison group; their qualitative finding aligns with Sommers (2010) and with the ACS-weighted ITS, although Wright's specification cannot validate the race-as-citizenship proxy directly because the non-citizen comparison group is constructed from the same race-coded CPS population that the proxy is meant to characterize. Our race-as-citizenship validation against the ACS gold standard (§5.1) directly addresses the measurement-error question that neither Sommers nor Wright was able to test in their respective papers.

The measurement-validation pass we contribute is the measurement-validation pass: the prior literature's race-as-citizenship proxy is 98 percent accurate against the ACS gold standard, so the qualitative findings are *not* biased by proxy contamination beyond 0.1 percentage points of attenuation. This vindicates Sommers' identification while sharpening the quantitative magnitude.

5.7 Sharp intensity-DiD (the sharp intensity-DiD specification) remains null

The sharp intensity-DiD attempted with hand-coded GAO + OIG + KFF friction tiers assigned only 5 of 51 states (GA, MT, NH, NY, TX) to Tier 1 (pre-existing pre-DRA documentation regime). The Tier-1-vs-residual DiD returns $\beta = -0.17$ pp for adults (SE 0.91 pp; p = 0.71) and +0.46 pp for children (p = 0.63) — null.

The new ACS variables (CITIZEN, HCOVANY, HINSCAID) do not unlock state-level documentation-friction tier coding. Reactivation of the sharp intensity-DiD design requires either (i) NAPHSIS EVVE DUA + 50-state PRR campaign, (ii) FOIA recovery of GAO-07-889 state-level survey microdata, or (iii) reframing as a 5-state pre-treated cohort case-study (NY/GA/MT/NH/TX vs. all others).

5.8 Parallel-trends interpretation of the event-study

The event-study version of the analysis (Figure 3 above) is the cleanest test of the parallel-trends assumption underlying the linear-trend interrupted time series in §5.2 and §5.3: the pre-period coefficients for 2002–2004 are statistically zero, so the post-2006 indicator is identified off a discrete level shift rather than a pre-existing trend. The post-2006 coefficients are similar in magnitude across 2006–2009, indicating that the friction shock produced an immediate level drop that persists rather than a gradual trend. This is the substantive content of the falsifiability claim in §6.3: a paper that observed slope-only post-2006 changes would identify a different mechanism (gradual erosion of enrollment over time) than what we observe (a discrete administrative-friction shock with persistent level effect).

---

6. Discussion

6.0 Roadmap

Section 6 organizes the discussion as follows. §6.1 documents the methodological contribution of validating the race-as-citizenship proxy used in prior literature. §6.2 explains how the headline magnitudes bound the post-DRA citizen Medicaid coverage decline. §6.3 catalogues the questions the descriptive interrupted time series cannot answer. §6.4 quantifies the welfare implications and the cost-of-friction for state Medicaid programs. §6.5 places the DRA case in the broader

administrative-friction Medicaid policy literature. §6.6 compares the DRA estimates to the 2023–2024 Medicaid unwinding. §6.7 discusses the generalizability of the magnitudes.

6.1 The race-as-citizenship proxy is 98 percent accurate

The first contribution of this paper is methodological validation of the dominant proxy used in the pre-DRA literature. Sommers (2010) and subsequent papers used race-as-citizenship because the available CPS extracts did not contain CITIZEN. We document, on the ACS gold standard, that the proxy is 98.0–98.1 percent accurate for non-Hispanic whites across the 2002–2009 window. This means that the prior literature's estimates were biased toward zero by approximately 2 percent × (citizen-vs-noncitizen Medicaid gap), which Table 4 puts at ~5 percentage points — implying an attenuation of approximately 0.1 percentage points. The ACS-weighted estimates recover this attenuation: the ACS-weighted adult estimate of −0.91 pp is approximately 0.12 pp more negative than the CPS-weighted estimate of −0.79 pp, consistent with the predicted attenuation magnitude.

The Hispanic comparison group, by contrast, is only ~70 percent citizens. Race-as-citizenship is a poor proxy for the Hispanic population and any paper using it as a comparison group needs to acknowledge the substantial measurement-error contamination.

6.2 The magnitude is bounded by the descriptive ITS

The ACS-weighted headline estimate — adult modal-citizen Medicaid coverage dropping 0.91 pp post-DRA, child coverage dropping 1.68 pp — is the most credible bound on the post-DRA citizen Medicaid coverage decline that we can construct from public data. It is qualitatively consistent with Sommers (2010) and quantitatively a refinement (slightly larger magnitudes, tighter standard errors). In a 51-state Medicaid program where adult enrollment hovered around 30–35 million in the mid-2000s, a 0.91-percentage-point coverage decline translates to approximately 250,000–350,000 adult citizens losing coverage in the immediate post-DRA period — concentrated in the second half of 2006 and 2007.

This is a *descriptive* finding. The interrupted time series identifies the level shift in the Medicaid coverage rate at the DRA effective date, holding state fixed effects and the linear time trend constant. It does not separate out the differential burden across state documentation regimes (the friction-mechanism question, which remains paused).

6.2b The friction-mechanism literature lens

The estimated 0.91-pp adult and 1.68-pp child coverage effects can be interpreted through the lens of the broader administrative-friction literature on take-up of social programs. Bhargava and Manoli (2015) document that small modi-

fications to the Earned Income Tax Credit notification letter — notification design changes that imposed essentially zero additional cost on filers — produced multi-percentage-point increases in EITC claim rates among likely-eligible non-claimants. The corresponding administrative-friction effect on Medicaid take-up of the DRA's documentation requirement appears to be of comparable order of magnitude, but operates in the *negative* direction (imposing friction reduces take-up rather than removing friction increasing take-up). The symmetry of these two effects — small administrative changes producing 1–2 percentage-point shifts in program participation either direction — is the principal generalization of the administrative-friction framework that this paper supports for Medicaid policy.

The Currie (2009) take-up literature places the structural baseline citizen-eligible-but-not-enrolled population at approximately 25–35 percent of the eligible adult Medicaid population in the mid-2000s. Our 0.91-pp estimate corresponds to approximately 2.6 to 3.6 percent of the structural eligible-but-not-enrolled population (0.91 / [25 to 35]). The interpretation is that the DRA shifted approximately 3 percent of the citizen-eligible-not-enrolled population from "eligible-not-enrolled because of structural take-up barriers" to "eligible-not-enrolled because of structural-plus-friction-shock barriers." The friction shock did not change the structural take-up rate; it added a new barrier that interacted with the structural barriers to produce additional disenrollment.

This decomposition has a sharp testable implication for the friction-mechanism literature: the populations most affected by the DRA's documentation requirement should be the populations with the highest baseline take-up gap (i.e., populations where structural take-up barriers were already binding). Demographically, this points to rural residents, low-income citizens with weak baseline access to identity documents, and citizen children whose parents face the documentation burden on the child's behalf. The §5.5b regional heterogeneity (with the South — the region with the highest concentration of these populations — producing the largest point estimate) is consistent with this prediction.

6.3 What the descriptive ITS cannot answer

Three questions deserve explicit acknowledgment. First, the friction-mechanism identification is paused. The would-be sharp intensity-DiD on state-level electronic-verification capacity returns a null because only 5 of 51 states can be confidently tier-coded from public sources. Reactivation requires either NAPHSIS DUA, GAO FOIA, or a 5-state case-study reframing. We do not claim to identify state-by-state heterogeneity in friction burden.

Second, the ACS HCOVANY/HINSCAID variables first appear in 2008, foreclosing a true pre/post DiD on ACS coverage outcomes. The Table 4 descriptive levels are cross-section only; we do not know what the citizen-vs-noncitizen Medicaid gap looked like in 2003–2005 from public ACS.

Third, the CPS extract on disk does not contain CITIZEN, so the person-level

CPS Medicaid outcome cannot be subset by citizenship directly. The ACS-citizen-population reweighting partially addresses this by applying the ACS-citizen-population state-year denominator as the analytic weight; a fuller solution would require a fresh CPS extract pulling CITIZEN to do the subset at the person level.

6.4 Welfare implications and the cost-of-friction

The administrative-friction framework yields a transparent welfare calculation for the DRA case. Each percentage-point of coverage loss among eligible citizens corresponds to roughly 300,000–350,000 individuals losing access to Medicaid in the immediate post-DRA period, scaled to the eligible-citizen Medicaid enrollee population in the mid-2000s. The 0.91-percentage-point adult coverage loss therefore corresponds to roughly 250,000–350,000 adult citizens, and the 1.68-percentage-point child loss corresponds to roughly 100,000–150,000 children, who lost coverage as a direct consequence of the documentation requirement.

The welfare cost of these coverage losses depends on the counterfactual access to care that displaced enrollees obtained. For adults, the closest available estimates come from the Oregon Health Insurance Experiment (Finkelstein et al. 2012; Baicker et al. 2013), which documented that Medicaid coverage substantially increased self-reported health-care utilization, reduced out-of-pocket medical expenses, and reduced depression incidence — though without measurable effects on most clinical-biomarker health outcomes. For children, the welfare stakes are likely larger: Medicaid coverage during childhood produces measurable improvements in adult outcomes including educational attainment, labor-market participation, and chronic-disease incidence (Brown, Kowalski, and Lurie 2020; Goodman-Bacon 2021 *AER*). A back-of-the-envelope upper bound on the welfare cost of the DRA's child coverage loss, applying the Brown-Kowalski-Lurie long-run-cohort estimates of childhood Medicaid coverage's adult effects, would place the present-value cost in the low billions of dollars per cohort year.

The administrative-cost calculus is more transparent. State Medicaid agencies in 2007 reported per-applicant documentation-verification costs ranging from $5 to $50 depending on infrastructure (HHS-OIG 2007). Multiplied by the approximately 30–35 million Medicaid applicants and re-determinations annually in the post-DRA period, the program-wide administrative cost of the rule was on the order of hundreds of millions of dollars per year — costs that fall on the state and federal Medicaid budgets and that are not offset by reductions in non-citizen enrollment because non-citizens were already largely excluded.

6.4b State-by-state implementation timing and rollout

Beyond the per-applicant cost estimates, the state-level implementation timing of the documentation requirement deserves separate treatment. Most state Medicaid programs continued to operate paper-only documentation regimes

through 2007, with the electronic vital-records match safe harbor activating in a small number of states (New York, Wisconsin, Michigan) in 2007 and a slightly larger group (California, Pennsylvania, Ohio, Washington, Massachusetts, Minnesota, Illinois) in 2008. The remaining majority of states either implemented electronic verification later in the decade or maintained paper-only documentation throughout the study window. The cross-state implementation pattern is correlated with state-level information-technology infrastructure, vital-records-system modernization, and Medicaid agency administrative capacity — none of which were systematically documented at the state level in the public 2007 GAO survey. The principal data limitation that prevents a sharp intensity-DiD design therefore is not the absence of cross-state variation but the absence of public state-level coding of that variation.

A complementary administrative-burden literature documents that paper-only documentation regimes impose disproportionate friction on rural residents who must travel to county vital-records offices, on populations with non-standard identification (foster children, formerly incarcerated individuals, victims of domestic violence whose identity documents may be controlled by the abuser), and on individuals whose birth records were not contemporaneously digitized. The headline 0.91-pp adult coverage loss we estimate is an average across the eleven-region distribution; the loss in counties served by paper-only Medicaid agencies in 2006–2008 is likely substantially larger.

6.5 Policy context

The DRA citizenship-documentation rule is the canonical example of an administrative-friction Medicaid shock — a rule that did not change federal eligibility but imposed a documentary burden whose mass fell on eligible citizens. The contemporary policy debates around Medicaid renewal redeterminations during the 2023–2024 unwinding, identity-verification requirements in state Medicaid expansions, and work-reporting requirements in §1115 demonstrations are all variations on the same administrative-friction theme. The 0.91-percentage-point adult coverage loss that we estimate for the DRA case provides one quantitative benchmark for the magnitude of administrative-friction effects on Medicaid coverage among the otherwise-eligible.

The contrast with non-citizens is instructive. The 5–6 percentage-point citizen-vs-noncitizen Medicaid gap (Table 4) is a structural baseline that the DRA was not designed to change. The 0.91-percentage-point shrinkage we estimate is the *increment* the DRA produced on top of that structural baseline — the friction imposed on citizens who were already in the program or eligible to enroll.

6.6 Comparison with the 2023–2024 Medicaid unwinding

The 2023–2024 unwinding of the COVID-19 continuous-enrollment provision provides a contemporaneous test of the administrative-friction framework on a

scale orders of magnitude larger than the DRA case. Between April 2023 and the close of the formal unwinding period in mid-2024, state Medicaid programs disenrolled approximately 25 million individuals, of whom roughly 70 percent were dropped for procedural rather than substantive eligibility reasons (Kaiser Family Foundation 2024). The procedural disenrollments — failure to return paperwork, address mismatches, undeliverable mail, missed phone calls — are administrative-friction effects in pure form. Reading the unwinding through the DRA-era lens, the 2023–2024 disenrollment population is the analogous structural population to the DRA-affected modal-citizen-proxy population: eligible citizens who lost coverage because of the documentary or procedural cost of remaining enrolled, not because their underlying eligibility lapsed.

The unwinding's cross-state heterogeneity in procedural-disenrollment shares (ranging from below 30 percent in states with proactive ex-parte renewals to above 90 percent in states with paper-only renewals) is exactly the sharp-intensity variation that the DRA case lacks. This presents a research opportunity for the next-decade extension of the DRA literature: the procedural-vs-substantive disenrollment split during the unwinding, combined with the state-level variation in renewal infrastructure, can identify the friction-mechanism that the DRA literature could not isolate. Estimates of the marginal effect of administrative friction on Medicaid coverage produced from the unwinding will be a critical input into ongoing debates about Medicaid renewal policy, identity-verification requirements in §1115 demonstrations, and work-reporting requirements that several states are pursuing as of 2026.

6.7 Generalizability of the magnitudes

The 0.91-percentage-point adult and 1.68-percentage-point child coverage estimates from this paper are specific to the modal-citizen-proxy population (non-Hispanic white CPS respondents) and the 2002–2009 calendar window. Two generalizations are worth noting and two are worth resisting. The estimates *do* generalize to the broader eligible-citizen population to the extent that the modal-citizen-proxy population is representative of citizens-at-the-Medicaid-eligibility-margin. The 98-percent proxy-accuracy validation in §5.1 supports this generalization for the *citizen* dimension; the modal-citizen-proxy population may differ from the broader eligible-citizen population on dimensions other than citizenship (income, geography, family structure), but the directional sign and approximate magnitude of the friction effect should generalize. The estimates also generalize to other administrative-friction Medicaid policy changes that share the structural feature of imposing a documentary-proof burden on eligible citizens without changing federal eligibility rules — most directly the contemporary Medicaid renewal-redetermination requirements during the 2023–2024 unwinding.

The estimates do *not* generalize to administrative-friction interventions that target ineligible populations directly (e.g., changes to the federal 5-year-bar for non-qualified immigrants, or changes to qualifying-immigration-status verifica-

tion for newly applying non-citizens). The DRA's friction fell on eligible citizens precisely because federal eligibility for non-citizens was already restricted; an analogous rule that *did* change non-citizen eligibility would have a different magnitude and direction. Nor do the estimates generalize cleanly to substantive eligibility changes (income-threshold adjustments, asset tests, work requirements that change qualifying status). Those interventions operate through a different mechanism — changing who is eligible — rather than through administrative cost on the already-eligible population.

7. Conclusion

The 2006 Deficit Reduction Act citizenship-documentation rule produced a measurable ~1 percentage-point drop in past-year Medicaid coverage among likely-eligible adult citizens (95% CI [−1.36, −0.45]; $p < 0.001$) and a ~1.7 percentage-point drop among likely-eligible citizen children (95% CI [−2.66, −0.69]; $p < 0.001$), based on a 51-state CPS-ASEC interrupted time series 2002–2009 reweighted by ACS-derived state-year citizen-population counts. The race-as-citizenship proxy that the prior literature relied on (because CITIZEN was not on available CPS extracts) is 98 percent accurate against the ACS gold standard — vindicating the qualitative findings of Sommers (2010) and Wright et al. (2010). The ACS-citizen-population reweighting strengthens the unweighted CPS finding in magnitude by 12 basis points for adults and 25 basis points for children, direction-of-bias consistent with the residual 2 percent non-citizen contamination in the CPS race proxy.

Post-DRA ACS levels (2008–2010) show a stable 5–6 percentage-point baseline gap between citizens and non-citizens in Medicaid coverage. The friction-mechanism identification — a sharp intensity-DiD on state-level electronic-vital-records-match capacity — remains paused pending NAPHSIS DUA or FOIA recovery of GAO state-level microdata. The headline is descriptive: the rule was administrative, the affected population was eligible citizens, and the magnitude is well-bounded by the descriptive interrupted time series on public data.

The methodological contribution of this paper is twofold. First, the validation of the race-as-citizenship proxy at 98-percent accuracy against the ACS gold standard removes a long-standing measurement-error concern in the pre-DRA literature; subsequent work using the same proxy on related questions can rely on this validation rather than re-conducting it. Second, the application of an ACS-derived citizen-population weight to the CPS state-year cell aggregation provides a template for related quasi-experimental studies of Medicaid policy that share the structural feature of needing a citizen-population denominator that the CPS extract on disk does not directly provide. Both contributions are mechanical refinements to existing identification strategies rather than new identification arguments; the substantive conclusion of the paper aligns with the prior literature.

The substantive policy implication is that administrative friction of the kind

imposed by the DRA citizenship rule produces measurable, persistent, and welfare-relevant coverage losses among the population the program is designed to serve. The 2023–2024 Medicaid unwinding has provided a much larger-scale natural experiment in administrative friction; the DRA estimates anchor the welfare interpretation of the unwinding's much larger procedural-disenrollment counts. The contemporary policy debates around Medicaid renewal documentation, identity-verification requirements in §1115 demonstrations, and work-reporting requirements that several states are pursuing as of 2026 all replay the same trade-off the DRA case crystallized: friction designed to exclude ineligibles produces type-1-inefficient costs on eligibles whose magnitude this paper helps to quantify.

7.1 Future research

Three lines of work would substantively extend the contribution of this paper. First, the friction-mechanism identification — the sharp intensity-DiD on state-level electronic-vital-records-match capacity — remains the principal unmet research design. Recovering the GAO-07-889 state-level survey microdata via FOIA would identify the friction-mechanism component of the headline ATT and would provide a much-needed benchmark for the marginal welfare cost of administrative friction in Medicaid. Second, the long-run cohort outcomes of children who lost Medicaid coverage during the 2006–2008 implementation period — adult educational attainment, labor-force participation, chronic-disease incidence — would extend the welfare interpretation beyond the contemporaneous coverage-loss estimate. The cohort affected during the DRA implementation period is now reaching adulthood, so the long-run cohort study is becoming feasible using the American Community Survey microdata from 2020 onward. Third, the application of the same ACS-citizen-population reweighting strategy to other Medicaid policy changes that affected modal-citizen-proxy populations — the 2008 SCHIPRA reauthorization, the 2014 ACA Medicaid expansions, and the 2023–2024 unwinding — would produce a comparable family of estimates that allow direct cross-policy magnitude comparison.

The 2023–2024 Medicaid unwinding is the most consequential immediate extension. The unwinding's procedural-disenrollment population is approximately one to two orders of magnitude larger than the DRA-era affected population, the cross-state heterogeneity in renewal infrastructure provides the sharp-intensity variation the DRA case lacks, and the continuous-enrollment policy that preceded the unwinding provides a clean pre-period for the difference-in-differences design. Estimates of the marginal effect of administrative friction on Medicaid coverage produced from the unwinding will sharpen the welfare and policy-design implications of the DRA estimates we report here.

Appendix A. Construction details

A.1 IPUMS USA extract specification

The IPUMS USA extract used in this paper covers ACS samples from the 2001 calendar-year ACS through the 2010 calendar-year ACS, restricted to ages 0–64 and PERWT > 0. The variables retrieved include the survey-design fields (YEAR, SAMPLE, SERIAL, HHWT, CLUSTER, STRATA, GQ, GQTYPE, GQTYPED, PERNUM, PERWT, FAMUNIT, MOMLOC, MOMRULE, POPLOC, RELATE, RELATED), the demographic and citizenship fields (SEX, AGE, RACE, RACED, HISPAN, HISPAND, NATIVITY, CITIZEN, YRIMMIG), and the coverage and economic fields (HCOVANY, HINSCAID, HINSCARE, EMPSTAT, EMPSTATD, INCTOT). The geographic-identifier fields STATEFIP and COUNTYFIP are also retrieved; for this paper the state-year aggregation only requires STATEFIP. The disability-indicator family (DIFFCARE, DIFFEYE, DIFFHEAR, DIFFMOB, DIFFPHYS, DIFFREM, DIFFSENS) is retrieved for descriptive use but not analyzed in the headline specifications.

The CITIZEN variable codes individuals as 0 (born in the United States, including Puerto Rico and other U.S. territories), 1 (born abroad of American parents), 2 (naturalized citizen), or 3 (not a citizen). The "any-citizen" indicator we use in this paper is `CITIZEN $\in$ {0, 1, 2}`. The HCOVANY variable codes any-coverage as 1 (no coverage) or 2 (with coverage); HINSCAID similarly codes Medicaid coverage. Both HCOVANY and HINSCAID are populated only from ACS sample year 2008 onward.

A.2 CPS-ASEC extract specification

The CPS-ASEC extract used in this paper covers ASEC survey years 2003 through 2010 (corresponding to the past-year reference period 2002 through 2009). The variables retrieved include YEAR, ASECFLAG, ASECWT (survey weight), STATEFIP, AGE, SEX, RACE, HIMCAIDLY (past-year Medicaid coverage), and HEALTH (self-reported general health). The CPS-ASEC extract on disk does not contain the CITIZEN variable; this is the data-availability gap that motivates the IPUMS USA reweighting.

A.3 State-year cell construction

State-year cells are constructed by aggregating the person-year IPUMS USA observations to the (STATEFIP, YEAR, age-group) level using PERWT as the analytic weight. The age-group strata are 19–64 (modal-citizen-proxy adults) and 0–18 (modal-citizen-proxy children). The cell counts are summed across PERWT-weighted observations to produce the `pop_citizen_adult` and `pop_citizen_child` denominators that enter the headline specifications as analytic weights. The CPS state-year cells are similarly constructed from the IPUMS CPS extract using ASECWT.

A.4 Sample-weight diagnostics

The ACS-derived `pop_citizen_adult` weight differs from the CPS sample-size weight by approximately 2 percent on average across state-year cells, reflecting the residual non-citizen contamination in the CPS race-as-citizenship proxy. The state-year correlation between the two weight series is 0.998, indicating that the reweighting affects the magnitude of the headline ATT but not the rank-ordering of state-year cells in the analysis. The decomposition in §5.2 attributes the 0.12-pp shift in the adult ATT (from −0.79 with CPS weights to −0.91 with ACS weights) to this 2-percent contamination times the ~5-percentage-point citizen-vs-noncitizen Medicaid gap.

A.5 Treatment of missing data

State-year cells with PERWT-weighted observations below a small minimum (50 person-years) are flagged as low-precision but retained in the headline specification; sensitivity to dropping these cells (re-running the headline with a 100-person-year minimum) shifts the adult ATT by less than 0.05 pp. The CPS-ASEC HIMCAIDLY variable is missing for approximately 0.3 percent of person-years; missing observations are dropped from the cell aggregation rather than imputed. The treatment of missingness is symmetric across pre-period and post-period and across treated and comparison cells, so missingness does not bias the headline ATT under the standard missing-at-random assumption applied to survey weights.

A.6 Comparison of CPS and ACS Medicaid coverage rates

The CPS-ASEC HIMCAIDLY variable measures past-year Medicaid coverage; the ACS HINSCAID variable measures coverage at the time of the survey. The two measures are conceptually distinct — the CPS is a flow measure (anyone covered for any portion of the prior year), the ACS is a stock measure (covered as of the survey reference period). The two measures produce systematically different mean coverage rates in overlapping years. For 2008–2009, the modal-citizen-proxy adult Medicaid coverage rate from the CPS averages approximately 9 percent across state-year cells; the corresponding ACS rate averages approximately 7 percent. The 2-percentage-point gap reflects the flow-vs-stock distinction: individuals who experienced any spell of Medicaid coverage during the past year are captured in the CPS measure but only those covered at the survey reference are captured in the ACS measure.

For the headline interrupted-time-series specification, the choice between the flow CPS measure and the stock ACS measure is consequential. The flow measure is more sensitive to administrative friction at enrollment (because friction at enrollment reduces the share of individuals who become covered at all); the stock measure is more sensitive to friction at re-determination (because friction at re-determination disenrolls existing covered individuals). The DRA's documentation requirement applied at both enrollment and re-determination, so

both flow and stock measures should respond. We use the CPS flow measure for the headline because the CPS-ASEC time series is longer (extends back to 2002) and the CPS HIMCAIDLY series is consistent across years, but the ACS HINSCAID measure for 2008–2010 (Table 4 in the main text) supports the same qualitative finding.

A.7 Reproducibility

The full analysis reproduces from federal IPUMS-USA and IPUMS-CPS microdata via a single end-to-end pipeline that builds the state-year ACS panels, fits the interrupted-time-series specifications, and writes table and figure outputs. Public artifacts include the state-year ACS citizen-population panel, the validation table, and the headline interrupted-time-series tables.

References

Government Accountability Office (2007). *Medicaid: States reported that citizenship documentation requirement resulted in enrollment declines for eligible citizens and posed administrative burdens*. GAO-07-889.

Kaiser Family Foundation (2007). *Implementing the new Medicaid citizenship documentation requirements: New York's experience*. Kaiser Commission on Medicaid and the Uninsured.

NCHS (2002–2010). National Vital Statistics System Natality file. Atlanta, GA: National Center for Health Statistics.

OIG, Department of Health and Human Services (2007). *Citizenship documentation requirements for Medicaid*. OEI-02-03-00190.

Predmore, Z., et al. (2023). State experiences linking Medicaid data with birth certificates and other data sources. *Medical Care*, 61(6), 353–359. doi:10.1097/MLR.0000000000001853.

Ruggles, S., Flood, S., Sobek, M., Backman, D., Cooper, G., Drew, J., Richards, S., Rodgers, R., Schouweiler, M., & Williams, K. (2024). *IPUMS USA: Version 15.0* [dataset]. Minneapolis, MN: IPUMS. doi:10.18128/D010.V15.0.

Capps, R., Fix, M., Henderson, E., & Reardon-Anderson, J. (2007). *Trends in immigrants' use of public assistance after welfare reform*. The Urban Institute.

Schwartz, K., & Sommers, B. D. (2014). Moving for Medicaid? Recent eligibility expansions did not induce migration from other states. *Health Affairs*, 33(1), 88–94.

Sommers, B. D. (2009). Loss of health insurance among non-elderly adults in Medicaid. *Journal of General Internal Medicine*, 24(1), 1–7. doi:10.1007/s11606-008-0792-9.

Wright, B. J., et al. (2010). Patient access and use of insurance and the Deficit Reduction Act citizenship documentation requirement. *Public Health Reports*, 125(4), 514–521.—

Reproducibility: All numbers in this manuscript trace to scripts in `analysis/` *and CSV outputs in* `analysis/tables/`*. Reproduce by running* `python3 data/scripts/06_build_ipums_usa_v3_citizen.py` *followed by* `python3 analysis/06_main_did_v3_citizen.py`*. Author: Jonathan Palisoc (jpalisoc@umich.edu, ORCID 0000-0001-5003-2631), University of Michigan, School of Public Health, Department of Health Management & Policy.*