Fear as a Barrier to Coverage: The Chilling Effect of the Public Charge Rule on Children's Medicaid Enrollment

Abstract

The Trump administration's 2019 expansion of the public charge rule made Medicaid receipt a negative factor in certain immigration determinations and generated widespread concern that eligible families would avoid health coverage. I study whether that fear translated into measurable coverage losses among populations formally exempt from the rule. Using American Community Survey data on 1,622,582 Medicaid-eligible U.S. citizen children below 200% of the federal poverty level from 2016 to 2024, I estimate difference-in-differences and event-study models comparing citizen children in mixed-status families with those in citizen-only families. I complement the national child analysis with restricted-use California Health Interview Survey data on 220,925 adults from 2015 to 2024, which permits a four-level decomposition of citizenship status (U.S.-born, naturalized, green-card holder, no green card). In the national analysis, citizen children in mixed-status families experienced a 2.0 percentage point decline in Medicaid enrollment after the September 2018 public charge announcement ($p < 0.001$), with no evidence of pre-policy divergence in the preferred event study (joint Wald test $p = 0.92$). Effects persisted through the policy reversal period and deepened during the 2023-2024 Medicaid redetermination restart ("unwinding"). Two additional ACS specifications further refine the chilling channel: (i) decomposing the citizen-only ACS control group reveals that the chilling effect concentrates in families with a noncitizen parent (–1.81 pp Medicaid relative to U.S.-born-parent baseline, $p = 0.003$) rather than across immigrant-connected families broadly (naturalized-parent × post: +1.41 pp, $p = 0.07$); (ii) excluding seven jurisdictions that extended Medicaid (or near-equivalent state-funded coverage) to undocumented adults during the paper window (California, Illinois, Oregon, New York, Colorado, Washington, and the District of Columbia) tightens the chilling estimate to –2.6 pp ($p < 0.001$), consistent with state policy environments partially counter-programming the federal rule. In California, naturalized citizens experienced a 1.3 percentage point Medi-Cal decline ($p = 0.030$), and the restricted-CHIS data identify large fear-based avoidance gradients: undocumented adults are 20.5 percentage points more likely than naturalized citizens to report having declined a non-cash benefit because of immigration concerns ($p < 10\text{-}^{3}0$). California's contemporaneous full-scope Medi-Cal expansions for undocumented adults (January 2020 ages 19–25, May 2022 ages 50+, January 2024 ages 26–49) preclude a clean within-state estimate of coverage-chilling, even in the cleanest within-state cell (age 26–49 in 2019–2021): Medi-Cal coverage among adults without green cards rose +6.8 pp rather than fell. Taken together, the clearest coverage losses appear among eligible citizen

children — concentrated in mixed-status families and in non-expansion states — while the California adult evidence provides supporting evidence on spillovers, fear-based avoidance, and administrative vulnerability rather than a coverage-chilling estimate for directly exposed adults.

I. Introduction

A. Motivation

The Trump administration's 2019 expansion of the public charge rule redefined what counted as a "public charge" for immigration purposes to include receipt of Medicaid, the Supplemental Nutrition Assistance Program (SNAP), and federal housing assistance. Although the rule technically applied only to certain noncitizens applying for admission or adjustment of status, its September 2018 announcement triggered fear, confusion, and benefit avoidance that extended far beyond the populations it formally targeted. This paper asks whether the public charge rule reduced Medicaid enrollment among legally exempt populations, and through what channels – fear, administrative friction, or information failure – a policy directed at immigration status translated into reduced health coverage for U.S. citizen children and other formally exempt groups.

The question matters for several reasons. First, Medicaid is the primary source of health coverage for low-income children in the United States. In 2018, Medicaid and CHIP covered approximately 37 million children, nearly half of all children below 200% of the federal poverty level (FPL). Decades of research establish that childhood Medicaid coverage improves health outcomes, educational attainment, and long-run economic self-sufficiency (Currie and Gruber, 1996; Cohodes et al., 2016; Wherry and Meyer, 2016; Miller and Wherry, 2019). Any policy that reduces enrollment among this population has consequences that compound over the life course.

Second, the episode exposes a structural vulnerability in the American safety net. Recent estimates place the mixed-status household population near 22 million people and show that nearly 5.5 million U.S.-citizen children live with at least one undocumented household member (FWD.me, 2024). The citizen children in these families are eligible for Medicaid on exactly the same terms as other citizen children, but their noncitizen parents face real or perceived immigration risk from interacting with government enrollment systems. When immigration policy becomes more punitive, parents in mixed-status families confront a calculus that parents in citizen-only families do not: whether enrolling a child in an entitled program might jeopardize a parent's ability to remain in the country.

Third, the public charge episode is not merely historical. As of early 2026, a new administration has signaled renewed interest in expanding public charge determinations, and immigrant communities are again reporting elevated fear

and benefit avoidance (Gonzalez, Bernstein, and Guelespe, 2025). Understanding whether prior chilling effects persisted after the rule's reversal is therefore directly relevant to ongoing policy debates. If effects endure long after a restrictive rule is rescinded, the costs of such policies are substantially larger than a static analysis of formal provisions would suggest.

Fourth, the public charge rule offers an unusually clean setting for studying how policy signals affect program take-up. A large literature documents that participation is shaped by transaction costs, stigma, information barriers, and perceived risks of government interaction (Moffitt, 1983; Currie, 2006; Bhargava and Manoli, 2015; Herd and Moynihan, 2018). The rule changed none of the formal eligibility rules for the populations this paper studies – no Medicaid income thresholds, categorical eligibility, or benefit packages were altered. What changed was the perceived cost of participation and the salience of immigration risk in the enrollment process, making it a natural experiment in the effects of fear and administrative burden on take-up.

B. What This Paper Does

The paper has two empirical components that answer related but not identical questions. The primary analysis uses American Community Survey (ACS) microdata on Medicaid-eligible U.S. citizen children from 2016 to 2024 to estimate difference-in-differences and event-study models comparing children in mixed-status families with those in citizen-only families. The ACS provides a nationally representative sample large enough to detect economically meaningful coverage changes, individual-level data on household immigration composition, and coverage spanning the full public charge policy cycle from pre-announcement through the Medicaid unwinding.

The supporting California analysis uses restricted-use California Health Interview Survey (CHIS) microdata from 2015 to 2024 to compare adults across four citizenship groups: U.S.-born citizens, naturalized citizens, noncitizens with green cards, and noncitizens without green cards. The restricted-use file's four-level CITIZEN1 variable distinguishes green card holders from other noncitizens, a separation the CHIS public-use file's pooled-noncitizen variable does not support. CHIS is useful not because it reproduces the ACS child estimand inside one state, but because it adds outcomes the ACS cannot observe — Medi-Cal enrollment for naturalized citizens, emergency department use, usual source of care, forgone care, whether respondents report being asked for Social Security number or citizenship documentation when seeking coverage, and two restricted-use fear-avoidance items collected from 2019 onward (whether the respondent ever decided not to apply for one or more non-cash government benefits because they worried it would disqualify them or a family member from a green card or U.S. citizenship, and whether such avoidance occurred in the last 12 months).

These two analyses play different and complementary roles. The ACS estimates identify the main child-coverage effect in the population for whom Medicaid is

the relevant margin; they carry the paper's headline quasi-experimental quantitative result. The CHIS analysis serves three roles. First, it tests for chilling spillovers to fully exempt populations: naturalized citizens were unambiguously not subject to the rule, yet remained embedded in immigrant community fear networks. Second, it identifies a direct fear-based benefit-avoidance gradient using post-period restricted-CHIS items that the public-use file does not contain — a precisely-measured measure of the channel through which public charge fears translate into foregone enrollment. Third, it provides honest transparency about what restricted-use access can and cannot identify in California: the contemporaneous full-scope Medi-Cal expansions for undocumented adults preclude a clean within-state coverage-chilling estimate, but the four-level CITIZEN1 decomposition supports both the disentanglement of federal chilling from state-expansion effects and the cleanly-identified spillover, fear-avoidance, and administrative-vulnerability gradients.

C. Contribution

My contribution is threefold. First, I provide individual-level quasi-experimental evidence consistent with a chilling effect of the public charge rule on Medicaid enrollment among low-income citizen children in mixed-status families. The preferred estimate – a 2.0 percentage point decline with no pre-trend divergence – is consistent with prior county-level estimates (Barofsky et al., 2020) but extends the evidence by using individual-level microdata rather than county aggregates, defining the treatment group based on household composition rather than geographic noncitizen concentration, and covering the full policy cycle from 2016 to 2024 rather than ending shortly after the announcement.

Second, I show that the child-coverage effect persisted beyond the formal policy reversal and intensified during the 2023-2024 Medicaid unwinding. The pre-unwinding (2019-2022) estimate is -1.7 percentage points, and the unwinding-period (2023-2024) estimate deepens to -2.8 percentage points. This pattern suggests that fear and administrative disruption compound: families already wary of enrollment systems may have been disproportionately likely to lose coverage during redeterminations. The persistence finding contributes to the literature on policy hysteresis – the idea that restrictive policy changes generate lasting behavioral effects that outlive the policies themselves (Hacker, 2004; Bustamante et al., 2022; Gangopadhyaya and Braga, 2022).

Third, the California analysis supplies supporting evidence on mechanism rather than a second headline average treatment effect. Naturalized citizens experienced a 1.3 percentage point decline in Medi-Cal enrollment, even though they were unambiguously exempt from the rule, and immigrant groups reported substantially greater documentation scrutiny when seeking coverage. Together, these results quantify the administrative environment in which chilling effects operate and show that fear spilled beyond the legally targeted population.

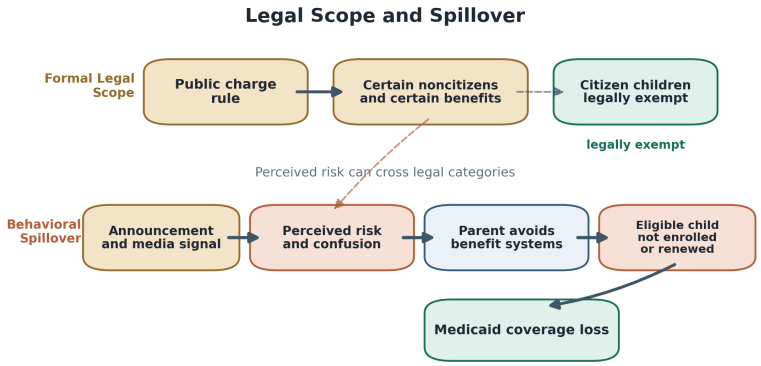

Figure 1: Legal scope and spillover

Note: This figure provides contextual structure for the legal scope and spillover. It summarizes the policy setting, mechanism, or empirical workflow used to interpret the estimates.

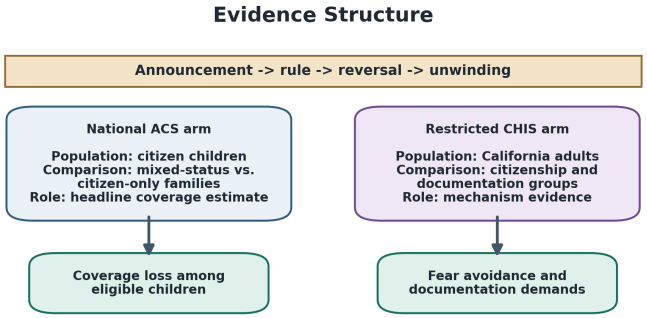

Figure 2: Evidence structure

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

D. Roadmap

Section II provides background on the public charge doctrine and prior evidence on chilling effects. Section III describes data sources. Section IV presents methods. Section V reports results. Section VI discusses interpretation, policy implications, and limitations. Section VII concludes.

II. Background

A. The Public Charge Doctrine: Origins and Evolution

What is "public charge"? In U.S. immigration law, "public charge" is a ground of inadmissibility — a reason the federal government may deny a noncitizen entry to the country or, more commonly, deny their application to adjust status to lawful permanent resident (LPR, "green card" holder). The legal test is forward-looking: whether the noncitizen is "likely to become a public charge" in the judgment of the adjudicating officer at U.S. Citizenship and Immigration Services (USCIS) or, at consular processing, the Department of State. Whether public charge bites depends on (i) what categories of benefit use count as evidence of likely dependence, (ii) which other personal characteristics — age, income, family ties, English proficiency — are weighed, and (iii) how heavily benefit use is weighted relative to those other factors. The 2019 rule change at the center of this paper expanded category (i) — what counts as evidence — and altered (iii) — how heavily it is weighted. It did not change who was at risk in the first place: U.S. citizens are not subject to public charge determinations at all, lawful permanent residents are at risk only at narrow re-entry margins and not for renewal of LPR status, and naturalized citizens are no longer subject to the doctrine in any form.

Why the rule matters for Medicaid take-up by populations who are not at risk. The chilling-effects literature documents a recurring pattern: when restrictive immigration policy raises the perceived cost of using benefits, take-up declines among populations not formally subject to the policy. This is not a legal effect; it is a behavioral one. The mechanism runs through information, network transmission, and risk aversion: families who cannot easily distinguish what does and does not count as "public charge use" — and who have noncitizen members or community ties whose status remains uncertain — respond to a perceived threat by avoiding benefit interaction altogether. The economic logic is standard non-price participation cost (Currie, 2006): when the perceived cost of using a benefit rises, take-up falls, even when the formal eligibility rules are unchanged. The administrative-burden framework (Herd and Moynihan, 2018) names three components — learning, compliance, and psychological costs — through which this perceived cost is generated. The 2019 rule increased all three for immigrant-connected households simultaneously.

Medicaid eligibility for citizen children in mixed-status households.

This paper's main analytic population is U.S.-citizen children below 200 percent of the federal poverty level living in households with at least one noncitizen parent ("mixed-status families"). Federal Medicaid and CHIP rules treat these children identically to citizen children in citizen-only households: their parents' immigration status is not a factor in their own eligibility, parents are not required to disclose their own immigration status when applying on the child's behalf in most states, and the federal "five-year bar" on most means-tested benefits for new lawful permanent residents (PRWORA §403) does not apply to U.S.-citizen children regardless of parent status. Approximately 5.5 million U.S.-citizen children currently live with at least one undocumented household member (FWD.me, 2024); the great majority are eligible for Medicaid or CHIP if the household income test is met. The chilling-effect hypothesis is that this entitled population disenrolls or fails to enroll because *parents* perceive that interacting with the Medicaid agency on the child's behalf may create immigration risk for themselves. That mechanism is what the 2019 rule activated.

Doctrinal history. The concept of the "public charge" in American immigration law dates to the colonial era, when communities sought to exclude settlers who might become dependent on local poor relief. The first federal codification came with the Immigration Act of 1882, which barred admission to any person "likely to become a public charge." For more than a century afterward, public charge determinations focused almost exclusively on whether an individual was likely to become primarily dependent on the government for cash subsistence. The standard was applied at the point of admission or, for immigrants already in the United States, at the point of adjustment to lawful permanent resident status. In practice, public charge denials were relatively rare. The Immigration and Naturalization Service (INS) exercised broad discretion, and the universe of benefits considered was limited to cash welfare programs such as TANF, SSI, and state general assistance (Capps et al., 2018).

The modern regulatory framework took shape after the 1996 welfare reform legislation. PRWORA dramatically restructured immigrant eligibility for public benefits, creating a five-year waiting period for lawful permanent residents (LPRs) to access most federal means-tested programs and rendering undocumented immigrants categorically ineligible for nearly all federal benefits except emergency Medicaid. Crucially, however, the 1996 law did not change the public charge standard itself – it did not make receipt of non-cash benefits a negative factor in immigration determinations. The INS clarified this distinction in its 1999 Field Guidance, which explicitly stated that use of Medicaid (except long-term institutional care), SNAP, CHIP, and housing assistance would not be considered in public charge evaluations. The 1999 Guidance was intended partly to counteract chilling effects PRWORA had already generated among immigrant communities (Kandula et al., 2004; Kaushal and Kaestner, 2005).

For nearly two decades after the 1999 Guidance, the public charge doctrine operated as a narrow and clearly defined concept. Cash assistance and government-funded institutional care counted. Non-cash programs did not. That clarity

mattered because it provided a bright-line rule that could be communicated to immigrant families, service providers, and enrollment workers. Although fear of government interaction never disappeared entirely from immigrant communities, the 1999 Guidance at least bounded the scope of concern.

B. The 2019 Rule Expansion

The Trump administration upended this framework beginning in 2017. A leaked draft executive order in January 2017 signaled the administration's interest in expanding public charge determinations to include a wider range of public benefits. The Department of Homeland Security (DHS) placed the rule on the Unified Regulatory Agenda in October 2017, putting immigrant advocacy organizations on notice that formal rulemaking was forthcoming. These early signals were important because they began the informational cascade that would ultimately produce chilling effects months before any formal rule change. The sequence is easier to evaluate visually because the public charge episode overlaps with later litigation, reversal, and the separate Medicaid unwinding shock.

Policy Timeline: Public Charge Rule and Medicaid Unwinding

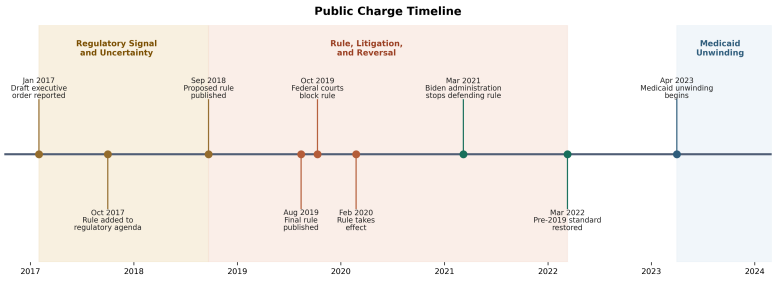

Figure 3: Public charge timeline

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

Notes: The timeline highlights the regulatory buildup beginning in 2017, the public charge rule period between announcement and formal restoration of the pre-2019 standard, and the later Medicaid unwinding that the split-period analysis treats separately.

The final rule specified that receipt of Medicaid (for non-pregnant adults under age 65), SNAP, or federal housing assistance for more than 12 months in any 36-month period would be a "heavily weighted negative factor" in public charge determinations under a "totality of the circumstances" test (84 Fed. Reg. 41,292).

Several features amplified the rule's chilling potential. First, it was widely misunderstood as applying far beyond its formal scope. Surveys found that many

immigrants – including naturalized citizens, refugees, asylees, and U.S. citizens in mixed-status families – believed that their own benefit use or family members' use could trigger public charge consequences (Bernstein, McTarnaghan, and Gonzalez, 2019; Gonzalez, Bernstein, and Karpman, 2020). Second, the rule explicitly excluded children's Medicaid and CHIP, but many families could not parse the distinction between program components that counted and those that did not. Third, even families that correctly understood the scope faced a rational fear problem: the rule signaled that the government was interested in monitoring benefit use among immigrants, raising the perceived cost of any interaction with enrollment systems.

The confusion was amplified by media coverage. The September 2018 NPRM generated extensive coverage in both English- and Spanish-language media, frequently describing the rule in broad terms without fully conveying the exemptions. Community-based organizations worked to disseminate accurate information, but the rule's complexity – and uncertainty about its final form during the 14-month comment period – made clear guidance difficult. The result was an informational environment in which fear was both rational (the rule did apply to some immigrants and did involve Medicaid) and overinclusive (many families believed they were affected when they were not).

C. The Litigation Cycle

The rule's implementation was marked by extended litigation that added further uncertainty. In October 2019, federal courts in multiple circuits issued preliminary injunctions blocking implementation. In January 2020, the Supreme Court stayed most injunctions (the 7th Circuit exception meant Illinois remained enjoined), and the rule took effect in February 2020 in most states. New injunctions during the pandemic effectively halted implementation by late 2020, and the Biden administration formally restored the 1999 standard in March 2022.

The litigation cycle matters for interpretation because it extended the period of policy uncertainty. Even after injunctions, families could not be confident that the rule would not ultimately take effect. And even after the Biden reversal, the memory of the policy – and the possibility of its revival – may have kept some families from re-enrolling. The behavioral response should not be expected to track the formal legal status of the rule with precision.

D. The Medicaid Unwinding

The Families First Coronavirus Response Act (March 2020) conditioned enhanced federal Medicaid funding on a continuous enrollment requirement: states could not disenroll beneficiaries during the public health emergency. This provision prevented routine eligibility redeterminations for more than three years. When it expired at the end of March 2023, states resumed redeterminations, and millions of beneficiaries were disenrolled – many for procedural reasons (failure to return paperwork, outdated addresses, administrative processing errors)

rather than actual ineligibility (KFF, 2024).

The unwinding is relevant to the public charge analysis because immigrant families may have been disproportionately vulnerable to procedural disenrollment. Language barriers, unstable addresses, unfamiliarity with government correspondence, and fear of re-engaging with enrollment systems after years of avoidance could all have contributed to higher procedural churn rates among mixed-status families. The interaction between pre-existing chilling effects and the unwinding's administrative demands creates a compounding dynamic: families that avoided re-enrollment during the public charge period may have been the same families least equipped to navigate the unwinding's bureaucratic requirements. This motivates the split-period decomposition in my empirical design.

E. Prior Evidence on Chilling Effects

E.1. The Welfare Reform Precedent The concept of chilling effects has deep roots in the immigration policy literature. The foundational evidence comes from the 1996 welfare reform, which restricted immigrant eligibility for federal means-tested programs and created a five-year waiting period for new lawful permanent residents. Although PRWORA's formal provisions targeted specific groups of immigrants, its effects spilled over to populations that remained legally eligible.

Kandula et al. (2004) provided early evidence that PRWORA reduced Medicaid enrollment among pre-1996 immigrants who were explicitly exempt from the law's restrictions. Using Current Population Survey data from 1994 to 2001, they estimated a 3 percentage point decline in Medicaid enrollment among eligible immigrants, compared with 1.6 percentage points for citizens. States that maintained universal immigrant coverage partially buffered the effect. Kaushal and Kaestner (2005) found even larger effects: uninsured rates among foreign-born unmarried women rose by 10-11 percentage points, and their children's uninsured rates rose by 13.5 percentage points. These effects could not be attributed to formal eligibility changes and were consistent with fear, confusion, and misinformation about who was affected.

The welfare reform episode established a pattern that would repeat with the public charge rule: a policy targeted at a subset of immigrants generated fear that radiated outward to populations far beyond its formal scope. The mechanism was informational rather than legal – immigrants who heard about new restrictions often could not easily determine whether they were personally affected, and the cost of error was high enough to deter participation even when actual risk was low.

E.2. Immigration Enforcement and Safety-Net Participation A parallel literature documents that immigration enforcement – independent of eligibility rules – can reduce safety-net participation. Watson (2014) provided seminal causal evidence using instrumental variables and variation in INS/ICE

enforcement intensity across metropolitan areas. She showed that heightened enforcement reduced Medicaid participation among citizen children of noncitizens, and estimated that up to 75% of the decline in noncitizen Medicaid participation around welfare reform was explained by contemporaneous enforcement spikes rather than PRWORA itself. That finding reframed the welfare reform story: what had been interpreted as a response to eligibility changes was substantially driven by the enforcement environment.

Alsan and Yang (2024) extended this research to the Secure Communities era. Using county-level rollout of Secure Communities between 2008 and 2013, they showed that the program's activation reduced SNAP and ACA enrollment among Hispanic citizen households. Effects were concentrated in mixed-status families, and the authors concluded that results were most consistent with network transmission of fear. Their estimate that ACA sign-ups would have been 22% higher absent Secure Communities underscored the quantitative importance of enforcement-driven chilling. Kravitz et al. (2024) linked county-level ICE detainer requests to lower Medicaid and SNAP enrollment, with effects stronger among households containing immigrant members. Twersky (2022) showed that restrictive state immigration laws reduced Medicaid/CHIP enrollment among citizen children in immigrant families by up to 5.5 percentage points. These studies collectively establish that fear, uncertainty, and perceived government risk are sufficient to deter enrollment among populations with full legal entitlement to benefits. The public charge rule represented an especially potent version of this dynamic because it explicitly linked benefit use to immigration consequences.

E.3. Evidence Specific to the Public Charge Rule For the public charge episode specifically, Barofsky et al. (2020) provided the strongest causal evidence to date. Using county-level administrative enrollment data for Medicaid, SNAP, and WIC from January 2015 through June 2019, they estimated a difference-in-differences model comparing counties with higher versus lower noncitizen population shares before and after the September 2018 announcement. They estimated approximately 260,000 fewer children enrolled in Medicaid. The finding was striking because the decline preceded any formal rule change, occurred among children explicitly exempt from the proposed rule, and was concentrated in counties with the highest exposure.

The Barofsky et al. design has two important limitations that the present paper addresses. First, the county-exposure approach treats noncitizen population share as the treatment intensity, which conflates true exposure with several confounded characteristics of high-noncitizen-share counties: more left-leaning state policy environments, denser sanctuary infrastructure, urban-rural composition, and contemporaneous coverage expansions. The individual-level mixed-status family proxy used here defines treatment at the household level, which is closer to the policy's actual mechanism (parental immigration exposure rather than residential neighborhood composition). Second, the Barofsky study ended in

June 2019 — before the rule's formal implementation in February 2020, before the 2021 reversal, and well before the 2023–2024 Medicaid unwinding. The present paper extends the evidence through the full policy cycle, including the unwinding period, which allows a direct test of whether chilling effects persist through reversal and amplify under subsequent administrative shocks.

Survey and descriptive studies complemented the causal evidence. KFF microsimulation estimated that 2.0 to 4.7 million Medicaid and CHIP enrollees could disenroll, including 600,000 to 1.8 million citizen children (Artiga, Garfield, and Damico, 2019). The Migration Policy Institute documented that participation in TANF, SNAP, and Medicaid declined twice as fast among noncitizens as among citizens during 2016-2019 (Fix and Capps, 2020). The Urban Institute found that one in five adults in immigrant families reported avoiding benefits, including naturalized citizens and green card holders not subject to the rule (Bernstein, McTarnaghan, and Gonzalez, 2019). Gonzalez, Bernstein, and Karpman (2020) confirmed these patterns persisted into 2019. Gangopadhyaya and Braga (2022) showed that avoidance continued into 2021 after reversal, providing early evidence of policy hysteresis. Qualitative reports from health center staff confirmed that disenrollment and enrollment avoidance extended to lawfully present immigrants not subject to the rule (Artiga, Tolbert, and Rudowitz, 2019).

The descriptive literature has also identified a key methodological constraint that the present paper addresses with restricted-CHIS access. Most population surveys distinguish only between citizens and noncitizens, pooling lawful permanent residents (green card holders) with undocumented immigrants in the noncitizen category. Because green-card holders face different public-charge exposure than undocumented adults — the rule's bite is on adjustment-of-status applications, which green-card holders have already largely cleared — pooling these groups attenuates estimated chilling effects on the most directly exposed population. The CHIS public-use file inherits this limitation through its CITIZEN2 variable. Restricted-CHIS access (via the UCLA Data Access Center) provides a four-level CITIZEN1 variable that separates green-card holders from noncitizens without green cards, enabling a substantively different decomposition of California adult chilling effects than has been previously possible.

Evidence on healthcare utilization is more limited but growing. Haro-Ramos et al. (2026) found that the 2018 announcement was associated with a 2.1 percentage point increase in self-pay ED visits among likely undocumented patients in LA County. Choi et al. (2023) showed that uninsured immigrants significantly delayed prenatal care initiation after the rule changes. The utilization findings reinforce the coverage findings: families avoiding Medicaid enrollment under the public charge rule still required medical care, but increasingly accessed it through self-pay or delayed channels rather than through the formal public coverage they were eligible for. The fear-based avoidance gradient identified in the present paper using restricted-CHIS items AL99 and AL104 directly measures the household-level decision rule that produces this pattern: undocumented

adults explicitly report deciding against benefit enrollment because of immigration concerns at rates 12–20 percentage points above naturalized citizens, providing a precisely-quantified measure of the channel that prior literature has only inferred from indirect utilization patterns.

E.4. Administrative Burden as a Mechanism The administrative burden framework (Herd and Moynihan, 2018) provides theoretical grounding for how the public charge rule could reduce enrollment without changing eligibility. The framework distinguishes learning costs, compliance costs, and psychological costs of interacting with government programs. The public charge rule raised all three simultaneously: it created a complex regulatory framework that even attorneys had difficulty parsing (learning costs), signaled that enrollment systems would scrutinize immigrant applicants more carefully (compliance costs), and associated benefit use with immigration risk (psychological costs). These burdens fall disproportionately on low-income, limited-English-proficiency families (Moynihan, Herd, and Harvey, 2015; Heinrich, 2016). Mixed-status families face an additional layer of burden because the enrollment process for citizen children often requires information about parents' income and household composition, creating opportunities for questions that immigrant parents find threatening.

A growing empirical literature documents the magnitudes of these channels. Bhargava and Manoli (2015) showed that simplified communications and pre-completed forms can substantially raise take-up of EITC and other safety-net programs even among populations with no formal eligibility change. Finkelstein and Notowidigdo (2019) extended this finding to SNAP, demonstrating that information and assistance interventions raise enrollment by roughly 25%. The mechanism implied by these studies is that the marginal cost of "one more administrative step" is large for low-income participants whose time and attention are already over-allocated. The public charge rule effectively added several psychological steps — enrolling now signals immigration risk; enrolling means generating government records that could later be used adversely; enrolling triggers questions that the family fears it cannot safely answer. Each added step is plausibly worth several percentage points of enrollment among the marginal eligible household.

The documentation-demand evidence developed in the present paper provides a direct measurement of one of these psychological channels. Within the survey universe in which the AL19V2 item is fielded (CHIS adult respondents not born in the United States, 2019–2024), noncitizens are 5 to 9 percentage points more likely than naturalized citizens to report being asked for Social Security numbers or citizenship documentation when seeking coverage, depending on whether age, sex, poverty, and race/ethnicity are conditioned on. The within-immigrant-universe gradient is identified within California — a state whose enrollment infrastructure is generally more immigrant-supportive than non-expansion states. The implication is that the underlying administrative-vulnerability gradient that the public charge rule activated is a durable feature

of the enrollment environment, not a transient artifact of the 2019–2021 policy window. Even after the rule's reversal, immigrants seeking coverage face heightened scrutiny that perpetuates the friction the rule weaponized.

E.5. Policy Feedback and Hysteresis A theoretical contribution of this paper is connecting the public charge episode to the political science literature on policy feedback. Hacker (2004) developed the concept of "policy drift" to describe how the welfare state can retrench not through overt program cuts but through the failure to update policies in the face of changing circumstances. Hacker and Pierson (2019) extended this framework to argue that policy changes can generate self-reinforcing feedback loops. In the immigration context, the public charge rule created a feedback dynamic in which fear reduced enrollment, which reduced the base of beneficiaries who might advocate for inclusive policies, which in turn made it easier for future administrations to contemplate similar restrictions.

The hysteresis hypothesis – that effects persist after formal reversal – follows directly from this framework. Once families disenroll from Medicaid, re-enrollment requires navigating the same administrative system that generated fear in the first place. The sunk costs of learning how to avoid the system may exceed the costs of re-engaging. Moreover, if families perceive that the policy environment may shift again – an entirely rational perception given recent political cycles – the option value of staying out increases. The result is a ratchet effect in which it is easier to scare families away from programs than to bring them back. Bustamante et al. (2022) provided empirical evidence, estimating that 108,000 to 193,000 Latino immigrants in California continued to avoid Medi-Cal even after the Biden administration reversed the rule. My paper tests this hypothesis more formally by examining whether the gap narrowed after the 2021-2022 reversal and whether it widened during the 2023-2024 unwinding.

F. Contribution to the Literature

This paper advances the evidence in several ways. Relative to Barofsky et al. (2020), I use individual-level microdata rather than county-level aggregates, define the treatment group based on household immigration composition rather than geographic noncitizen concentration, and extend the analysis through 2024 – covering the implementation, injunction, reversal, and unwinding phases the earlier study could not observe. Relative to survey-based descriptive evidence (Bernstein et al., 2019; Gonzalez et al., 2020; Gangopadhyaya and Braga, 2022), I provide quasi-experimental enrollment estimates from nationally representative microdata rather than self-reported intentions to avoid benefits. Relative to enforcement-focused studies (Watson, 2014; Alsan and Yang, 2024), I study a policy that explicitly linked benefit use to immigration consequences. The CHIS analysis adds a dimension absent from national studies: by examining outcomes for naturalized citizens, I can test for spillovers operating through networks rather than direct legal exposure, and by documenting the administrative

vulnerability gradient, I provide evidence on the conditions that make chilling effects possible.

III. Data

A. National Analysis: American Community Survey

A.1. Data Source and Access The national analysis uses ACS microdata accessed through IPUMS USA (Ruggles et al., 2024) for survey years 2016 through 2024. The ACS is the largest household survey in the United States outside the decennial census, sampling approximately 3.5 million addresses per year. It provides nationally representative individual-level data on demographics, household composition, income, health insurance coverage, and citizenship status. Survey responses are collected on a rolling monthly basis throughout the year, with each annual release representing a full 12 months of data collection.

The ACS is the standard data source for studying immigrant populations because of its large sample size and detailed coverage questions. It identifies whether each individual is a U.S. citizen by birth, citizen by naturalization, or noncitizen. It does not distinguish among noncitizen types – lawful permanent residents, temporary visa holders, and undocumented immigrants. This limitation is common to all major federal surveys and means that the mixed-status family definition necessarily pools families with different types of noncitizen parents. The pooling likely attenuates my estimates relative to the effect among the most legally vulnerable families, because families with undocumented members have the most to fear from public charge-style policies.

A.2. Sample Construction My preferred sample includes U.S. citizen children ages 0 to 18 in households below 200% of the federal poverty level. This income restriction is substantively important for two reasons. First, the chilling-effect hypothesis predicts the largest enrollment response among children plausibly eligible for Medicaid or CHIP. Children in families well above the income threshold have little reason to participate in public coverage programs and therefore little reason to respond to changes in the perceived cost of enrollment. Including them dilutes the estimated effect without adding information about the relevant behavioral margin. Second, the below-200% FPL threshold approximates the upper bound of Medicaid/CHIP eligibility in most states. While exact thresholds vary by state, child age, and program (with some states extending CHIP eligibility to 300% FPL or higher), the 200% FPL cutoff captures the overwhelming majority of the population for whom public coverage is a realistic option.

The resulting sample contains 1,622,582 observations across nine survey years (2016-2024): 220,387 (13.6%) in mixed-status families and 1,402,195 (86.4%) in citizen-only families. I present full-sample estimates for all income levels (N = 5,256,322) as secondary specifications, but I am explicit that these attenuate

the estimated effect because they include many children for whom Medicaid is not the relevant coverage margin.

A.3. Treatment Group Definition I classify children as living in mixed-status families if at least one identified parent in the household roster is a noncitizen. Citizen-only families are those in which all identified parents are U.S. citizens. The child must be a U.S. citizen in both groups. Children whose parents cannot be identified through the IPUMS household-roster pointers (~4% of the eligible-age citizen-child sample) are excluded. This share is small and approximately stable across calendar time within the analytic window, so the exclusion is unlikely to drive the post-2018 differential trend that the DiD identifies; if anything, families in which household relationships are difficult to enumerate are more likely to be in mixed-status households, so excluding them mildly attenuates the headline estimate.

This classification captures the household-level immigration exposure central to the chilling-effect mechanism. The theory does not predict that children themselves respond to the public charge rule – children do not make their own enrollment decisions. Rather, parents in mixed-status families, who perceive a link between benefit use and immigration risk, alter their children's enrollment behavior. The relevant exposure is at the parent level.

The mixed-status definition pools families with very different vulnerability levels. A family with an undocumented parent faces qualitatively different risks from one with a parent holding a green card near naturalization. The ACS does not provide information to distinguish these subgroups, so my estimates are averages across a heterogeneous treatment group. The true effect among the most vulnerable families is likely larger than what I report.

A.4. Outcome Variables The primary outcome is current Medicaid enrollment. Secondary outcomes include any public coverage, any health insurance coverage, and uninsurance. The distinction between Medicaid and any public coverage matters because some children may shift between programs (for example, Medicaid to CHIP) in response to the public charge environment. The uninsurance measure captures the most consequential outcome – children with no coverage at all – but may be noisier if some adjustment involves shifts to employer-sponsored or individual-market insurance.

A.5. Control Variables and Summary Statistics All ACS models use person-level survey weights and include state and year fixed effects, which absorb time-invariant state characteristics and national trends common to all states. Individual-level controls include child age (entered as dummy variables), sex, and race/ethnicity (non-Hispanic White, non-Hispanic Black, Hispanic, Asian/Pacific Islander, and Other). Standard errors are clustered at the state level to account for within-state correlation in outcomes over time and across individuals.

Table 1 presents pre-period (2016-2018) weighted means. Children in mixed-status families are modestly younger (8.24 vs. 9.04), equally likely to be female, and substantially more likely to be below 200% FPL (0.583 vs. 0.335). Pre-period Medicaid enrollment is higher (0.602 vs. 0.337) and pre-period uninsurance is also higher (0.071 vs. 0.037), reflecting lower income and baseline chilling effects that predate the 2019 rule. Despite these level differences, the identifying assumption requires parallel trends, not identical levels.

B. California Analysis: California Health Interview Survey

B.1. What CHIS Is and How This Paper Uses It The California Health Interview Survey (CHIS) is the largest state health survey in the United States, administered annually by the UCLA Center for Health Policy Research. It samples approximately 20,000–25,000 adults per cycle, is conducted in English, Spanish, Chinese (Mandarin and Cantonese), Korean, Vietnamese, and Tagalog, and is one of the only population-based surveys in the United States that separately identifies U.S.-born citizens, naturalized citizens, and noncitizens. The multilingual administration matters for this paper: studying immigrant populations using English-only surveys produces well-known non-response bias that CHIS is specifically designed to mitigate.

The California analysis in this paper uses **two CHIS data products** — the public-use file and the restricted-use adult file — and which file generated which result is determined by which variables a given specification requires. Section B.2 makes the file-by-result mapping explicit so the reader can trace every California estimate to its source. Section B.3 describes outcomes and the four-level CITIZEN1 architecture; Section B.4 describes the restricted-use access process; Section B.5 describes variance estimation.

B.2. Public-Use Versus Restricted-Use: Which File Generates Which Result Table 6 (in this section) maps each California analysis to the underlying CHIS data product. The pooled CHIS adult sample for 2015–2024 used in this paper covers N = 220,925 respondents in both the public-use and restricted-use files (the universe is the same; the restricted-use file simply contains additional confidential variables and a finer four-level citizenship taxonomy).

Specifications estimated on the public-use CHIS adult file:

- The three-group multi-group DiD comparing U.S.-born citizens, naturalized citizens, and pooled noncitizens (using public-use CITIZEN2) on the four standard outcomes — Medi-Cal enrollment, ED visit in past 12 months, having a usual source of care, and forgone care. The naturalized × Post coefficient on Medi-Cal (–1.3 pp, p = 0.030) reported as supporting evidence in §V.G is a public-use estimate.
- The documentation-demand specification in §V.H, which uses the public-use binary indicator for whether the respondent was asked for a Social Security number or citizenship documentation when seeking health cover-

age. The item is in the public-use adult file, was first fielded in 2019, and is asked **only of respondents not born in the United States**; the analysis I report restricts to that immigrant universe and uses naturalized citizens as the reference group, both to respect the survey universe and to avoid the spurious zero-baseline contrast that an inclusive U.S.-born comparison would generate.

- The fear-based benefit-avoidance gradient (the public-use AL99-equivalent indicator on whether the respondent ever decided not to apply for a non-cash benefit because of immigration concerns), used in §V.G alongside the restricted-use AL104 12-month follow-up. The public-use indicator is the lifetime-recall version; the 12-month follow-up is restricted only.
- All summary statistics in Table 2 and the poverty-stratified heterogeneity analyses reported in the appendix.

Specifications that require restricted-use CHIS variables:

- The four-level CITIZEN1 decomposition (U.S.-born / naturalized / green-card holder / no green card), and any analysis that separates undocumented adults from green-card holders. This includes the age-stratified four-level DiD (§V.F), the two-window decomposition that disentangles federal chilling from the California full-scope Medi-Cal expansion calendar, and the green-card-versus-naturalized fear-avoidance contrast.
- The 12-month-recall fear-avoidance variant (AL104) — only the lifetime-recall version (AL99) is on the public-use file.
- The county-enforcement heterogeneity test (Appendix Exhibit A25), which requires the full county FIPS code (FIPS_CNT) suppressed in the public-use file.
- The immigrant-tenure heterogeneity test, which uses YRUS (years in the U.S.).

Table 6: CHIS data-source map.

CHIS specification	File	Required variables
Three-group DiD on Medi-Cal, ED, USUAL, FORGO	Public-use	CITIZEN2, INSMC, ER, USUAL, FORGO
Naturalized × Post on Medi-Cal (−1.3 pp), §V.G	Public-use	CITIZEN2, INSMC
Documentation-demand, immigrant-only universe, §V.H	Public-use	AL19V2 (immigrant universe), CITIZEN2
Lifetime fear-avoidance (AL99-equivalent), §V.G	Public-use	CITIZEN2, AL99 indicator
Four-level CITIZEN1 multi-group DiD (no GC vs GC vs nat vs US-born)	Restricted	CITIZEN1
Two-window decomposition (federal chilling vs CA expansion windows)	Restricted	CITIZEN1, INSMC
Age-26–49 within-CA “cleanest cell” event study	Restricted	CITIZEN1, INSMC
12-month fear-avoidance gradient (AL104), §V.G	Restricted	CITIZEN1, AL104

CHIS specification	File	Required variables
County-enforcement heterogeneity, Appendix A25	Restricted	CITIZEN1, FIPS_CNT
Immigrant-tenure heterogeneity (YRUS)	Restricted	CITIZEN1, YRUS

Notes: This table documents the source files, scripts, variables, or data inputs used in the analysis. It is included to make the construction of the analytic evidence reproducible.

Every restricted-use coefficient reported in this paper traces to a specific cell in one of three approved UCLA Center for Health Policy Research Data Access Center code-submission tickets executed between 2026-04-07 and 2026-04-24. The submitted Stata programs and returned analyst-output CSVs are documented in the project's reproducibility folder.

B.3. Sample Composition and the Four-Level Citizenship Architecture Across both CHIS files, the pooled 2015–2024 adult sample (N = 220,925) contains 168,525 U.S.-born citizens (76.3%), 35,656 naturalized citizens (16.1%), and 16,744 noncitizens (7.6%). The public-use CITIZEN2 variable pools the last group into a single "noncitizen" category. The restricted-use CITIZEN1 variable splits that category into noncitizens with a green card (lawful permanent residents) and noncitizens without a green card (which the literature treats as the closest available proxy for undocumented status, recognizing that some respondents in this category hold temporary visas, DACA, TPS, or other non-LPR documented status).

The four-level CITIZEN1 architecture matters substantively because public charge bites primarily on adjustment-of-status applications. Lawful permanent residents have already cleared adjustment and face the rule only at narrow reentry margins; undocumented adults face it at any future adjustment attempt. Pooling the two in CITIZEN2 mechanically attenuates estimated chilling effects on the most directly exposed population, and as the §V.F results show, the pooling also obscures a substantively important interaction: California's contemporaneous full-scope Medi-Cal expansions for undocumented adults (January 2020 ages 19–25, May 2022 ages 50+, January 2024 ages 26–49) overlap almost exactly with the federal chilling window. Separating green-card holders from no-green-card adults makes that confound visible and explains why the within-California adult coverage estimate cannot serve as a clean coverage-chilling estimate even with restricted-use access.

Naturalized citizens are valuable to the design for the opposite reason: they are unambiguously exempt from public charge (the rule does not apply to U.S. citizens, and naturalization makes one a citizen). Yet they remain embedded in immigrant households, social networks, and language-minority media environments. A finding that naturalized citizens reduced enrollment is therefore evidence for a network-mediated fear mechanism — it cannot be attributed

to direct legal exposure. The –1.3 pp naturalized × Post Medi-Cal coefficient (public-use) is the cleanest spillover finding in the California arm.

I do not restrict the CHIS adult sample to below 200% FPL because the California arm is not intended to reproduce the ACS child estimand. Instead, it asks whether citizenship groups within one state experienced different changes in coverage, access, fear-based avoidance, and documentation demands during the public charge episode. Poverty remains important for interpretation, so I control for it in all models and report poverty heterogeneity in the appendix. Descriptive statistics are in Table 2.

B.4. Outcomes CHIS contributes four standard outcomes available in both files: Medi-Cal enrollment, ED visit in the past 12 months, having a usual source of care, and forgone care in the past 12 months. Two additional restricted-only outcomes are central to this paper: the fear-based benefit-avoidance items AL99 and AL104 (§V.G), and the documentation-demand item AL19V2 (§V.H), which is fielded only to respondents not born in the United States and only from 2019 onward.

B.5. Restricted-Use Access Process The restricted-use CHIS adult file is accessed through the UCLA Center for Health Policy Research Data Access Center (DAC). Researchers submit a project application, are vetted, and then submit Stata or SAS programs to the DAC; DAC analysts execute the programs against the restricted-use file in a secure environment and return vetted output (suppression rules apply to small cells). Three approved DAC tickets executed by analyst staff between 2026-04-07 and 2026-04-24 produced the restricted-use outputs reported in this paper; the submitted programs and returned CSVs are stored in the project reproducibility folder, and every restricted-use coefficient reported here can be traced to a specific cell in a specific returned CSV.

B.6. Variance Estimation The public-use CHIS files do not include the sampling design variables needed for Taylor-series linearization variance estimation; they do include 80 balanced repeated replication (BRR) replicate weights. The restricted-use file additionally includes 80 jackknife replicate weights. I use adult survey weights in all CHIS estimates, report heteroskedasticity-robust standard errors as primary inference, and treat BRR/jackknife replicate-weight estimates as a survey-design sensitivity analysis (appendix). I do not rely on the smaller replicate-weight standard errors to strengthen claims that are weak under the more conservative primary inference.

IV. Methods

A. ACS Identification Strategy

A.1. Difference-in-Differences Framework The primary specification is:

$$Y_{ist} = \alpha + \beta \cdot (\text{MixedStatus}_i \times \text{Post}_t) + \gamma \cdot \text{MixedStatus}_i + \delta_s + \theta_t + X_i'\lambda + \varepsilon_{ist}$$

where *Y_ist* is the Medicaid enrollment indicator for child i in state s in year t; *MixedStatus_i* is an indicator for living in a mixed-status family; *Post_t* is an indicator for survey years 2019 and later (reflecting the September 2018 announcement); *delta_s* are state fixed effects; *theta_t* are year fixed effects; and *X_i* is a vector of individual controls (age dummies, sex, race/ethnicity). The coefficient of interest is *beta*, which captures the differential change in Medicaid enrollment for children in mixed-status families relative to citizen-only families after the announcement. Standard errors are clustered at the state level, the level at which Medicaid policy is administered. The preferred sample is below 200% FPL. Full-sample results are secondary because including children above the Medicaid eligibility margin dilutes the estimated effect – this is not just a power issue but an estimand issue, as the theory predicts chilling effects among families whose children actually use or would use Medicaid.

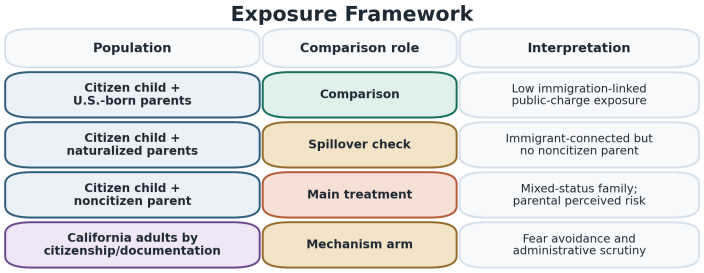

Figure 4: Exposure framework

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

A.2. Event-Study Specification I estimate event-study models interacting mixed-status family status with year indicators (2018 omitted):

$$Y_{ist} = \alpha + \sum_k \beta_k \cdot (\text{MixedStatus}_i \times \text{Year}_k) + \gamma \cdot \text{MixedStatus}_i + \delta_s + \theta_t + X_i'\lambda + \varepsilon_{ist}$$

The coefficients *beta_k* for $k <$ *2018* (the "leads") test for pre-policy divergence. If the parallel trends assumption holds, these should be close to zero and jointly insignificant. The coefficients for $k >=$ *2019* trace the dynamic treatment path. I conduct a joint Wald test of pre-period coefficients using the estimated variance-covariance matrix, which is preferable to examining individual leads in isolation because it accounts for the joint distribution. The ACS has only two pre-treatment years (2016-2017), limiting pre-trend testing power; this is acknowledged as a limitation.

A.3. Split-Period Specification A split-period specification separates pre-unwinding (2019-2022) from unwinding (2023-2024) effects:

$$Y_{ist} = \alpha + \beta_1 \cdot (\text{MixedStatus}_i \times \text{PreUnwinding}_t) + \beta_2 \cdot (\text{MixedStatus}_i \times \text{Unwinding}_t) + \gamma \cdot \text{MixedStatus}_i + \delta_s + \theta_t + X_i'\lambda + \varepsilon_{ist}$$

This decomposition is important for two reasons. First, the continuous enrollment requirement during the pandemic limited disenrollment of existing beneficiaries from mid-2020 through March 2023, which may have mechanically attenuated the estimated chilling effect during that sub-period. Second, the unwinding created new administrative barriers that may have disproportionately affected immigrant families, producing a compounding dynamic in which pre-existing chilling effects and new procedural disenrollment interact.

A.4. Secondary Specifications Several secondary specifications test the robustness and interpretation of the main ACS results. A placebo test restricts the sample to citizen-only families and compares those in states with high noncitizen population shares to those in low-noncitizen-share states. If the mixed-status effect reflects a genuine chilling response specific to families with noncitizen members, this placebo should yield null results. I also report full-sample ACS results, an alternative timing specification using the February 2020 implementation date rather than the September 2018 announcement, earlier-treatment sensitivity analyses that treat the policy threat as beginning during the 2017-2018 regulatory buildup, and a separate state-by-year fixed-effects event study that absorbs any state-specific shocks shared by mixed-status and citizen-only children within a given year. Because these exercises are secondary to the below-200% FPL child Medicaid design, they are reported mainly in the appendix and used to calibrate rather than replace the preferred specification.

B. CHIS Identification Strategy

B.1. Multi-Group Difference-in-Differences The preferred CHIS specification estimates separate post-2018 coefficients for noncitizens and naturalized citizens relative to U.S.-born citizens:

$$Y_{it} = \alpha + \beta_1 \cdot (\text{Noncitizen}_i \times \text{Post}_t) + \beta_2 \cdot (\text{Naturalized}_i \times \text{Post}_t) + \gamma_1 \cdot \text{Noncitizen}_i + \gamma_2 \cdot \text{Naturalized}_i + \theta_t + X_i'\lambda + \varepsilon_{it}$$

The multi-group specification is more informative than a two-group comparison because it separately identifies effects for populations with very different policy exposure. Noncitizens include those most directly affected by the rule (at least lawful permanent residents applying for adjustment of status), though the public-use CHIS pools them with green card holders for whom risk was lower. Naturalized citizens were fully exempt but may have responded through network-mediated fear. A finding that naturalized citizens reduced enrollment is strong evidence for a fear-based mechanism because it cannot be attributed to direct legal exposure.

B.2. Event Studies and Documentation-Demand Analysis CHIS event studies interact group indicators with year dummies (2018 omitted), providing

three pre-treatment years (2015-2017) and six post-treatment periods (2019-2024). Joint Wald tests of pre-period coefficients are conducted for each outcome. In response to committee feedback, I also estimate pairwise event studies comparing noncitizens with U.S.-born citizens and naturalized citizens with U.S.-born citizens, and I audit those samples for missingness, empty year-by-group cells, and event-study saturation.

For the documentation-demand measure, I estimate cross-sectional post-2019 models on the immigrant-only respondent universe — the CHIS sampling universe for AL19V2:

$$\text{Asked}_i = \alpha + \beta \cdot \text{Noncitizen}_i + X_i'\lambda + \theta_t + \varepsilon_i, \quad i \in \{\text{not U.S.-born}\}$$

Naturalized citizens are the reference category; noncitizens (pooled green-card holders and adults without green cards in the public-use file) are the contrast. Restricting to the immigrant universe avoids the spurious U.S.-born baseline of structural zeros and aligns the regression with the question's intended sampling frame. Because the item is unavailable before 2019, this is a within-immigrant cross-sectional gradient — it measures the size of the administrative vulnerability gap during the relevant policy period, not whether the public charge rule changed documentation practices.

Secondary CHIS specifications include two-group models comparing pooled noncitizens with U.S.-born citizens. COVID sensitivity checks drop the 2020-2021 survey years and add a COVID interaction term. Heterogeneity analyses stratify by poverty level. Alternative treatment timing uses 2020 (implementation) rather than 2019 (announcement) as the post-period threshold.

B.3. County-Enforcement Heterogeneity The county-enforcement heterogeneity specification interacts the four-level CITIZEN1 × Post DiD with a binary high-enforcement county indicator merged on `FIPS_CNT` × `survey_year`. For each outcome y in {Medi-Cal, ED visit, usual source of care, forgone care}, I estimate

$$y_{it} = \alpha + \sum_{g \in \{\text{Nat, GC, NoGC}\}} \beta_g \cdot (\mathbf{1}[\text{doc4}_i = g] \times \text{Post}_t) + \sum_g \gamma_g \cdot (\mathbf{1}[\text{doc4}_i = g] \times \text{Post}_t \times \text{HighEnf}_{c(i),t}) + \delta \cdot \text{HighEnf}_{c(i),t} + X_i'\lambda + \theta_t + \mu_g + \varepsilon_{it}$$

where U.S.-born citizens are the omitted CITIZEN1 group, controls X include age, sex, race, and below-200% FPL, *theta* are survey-year fixed effects, and *mu* are CITIZEN1 fixed effects. The beta coefficients identify the pooled chilling-or-expansion response in low-enforcement (San Francisco AOR) counties; the gamma coefficients identify the differential response in high-enforcement (Los Angeles + San Diego AORs) counties. Variance is estimated via 80 CHIS replicate weights using the design-supplied jackknife multiplier. The specification is run on the full restricted-CHIS adult sample for 2015–2024 (N = 220,925); all 220,925 adult observations match a county-year. Because the underlying enforcement values are AOR-level rather than within-AOR county-level, the gamma coefficients should be read as Southern-versus-Northern California enforcement-

intensity contrasts, not as within-AOR county-level gradients. Because 2015–2021 enforcement values are back-filled with the AOR's 2022 value, the high-enforcement indicator is essentially time-invariant within the AOR-pair partition during the federal chilling window, so the test is a between-AOR contrast on the level of the post-period response rather than a time-varying enforcement-shock instrument. The specification is reported in the appendix.

V. Results

A. Raw Outcome Trends

Before turning to regression estimates, it is important to establish the raw data patterns. Figure 1 presents weighted mean Medicaid enrollment rates over calendar time for citizen children in mixed-status versus citizen-only families in the below-200% FPL sample. The two groups track closely during the pre-announcement period (2016-2018), consistent with the parallel trends assumption. After the September 2018 announcement, a visible gap opens and persists through 2024, with mixed-status children declining relative to citizen-only children. The gap appears to widen further in 2023-2024, coinciding with the Medicaid unwinding. Figure A12 extends this to a 2x2 panel showing Medicaid and uninsured outcomes for both samples, confirming that the divergence is concentrated in the low-income sample. All raw trends use ACS survey weights.

B. Event Study Results

Figure 2 presents event-study coefficients for Medicaid enrollment in the below-200% FPL sample (2018 omitted). Pre-period coefficients are small and insignificant: 2016 (+0.4 pp, $p = 0.70$) and 2017 (+0.2 pp, $p = 0.76$). The joint Wald test yields $p = 0.92$, providing no evidence of pre-policy divergence, while still acknowledging that the ACS has only two pre-treatment years.

The post-period path is informative about the dynamics of the chilling effect. The 2019 coefficient is -0.9 pp ($p = 0.12$), suggesting the effect was emerging but not yet statistically distinguishable from zero in the first year after the announcement. This is plausible because the announcement came in September 2018, so the 2019 ACS captures a mix of pre- and post-announcement months. The effect deepens sharply in 2020 to -2.5 pp ($p = 0.010$), remains significantly negative in 2021 at -1.6 pp ($p = 0.013$), softens in 2022 to -0.8 pp ($p = 0.161$), and then re-intensifies during the Medicaid unwinding: -1.9 pp in 2023 ($p = 0.001$) and -3.3 pp in 2024 ($p < 0.001$). This dynamic pattern is inconsistent with a short-lived announcement shock that dissipated once the rule was vacated. Instead, it suggests persistent chilling, with only partial recovery after reversal, compounded by the unwinding.

C. National Child Coverage Results: Main Estimates

Table 3 reports the core ACS estimates. Among citizen children below 200% FPL, mixed-status families experienced a 2.0 percentage point decline in Medicaid enrollment (SE = 0.005, $p < 0.001$). Any public coverage shows a nearly identical decline (-2.0 pp, $p < 0.001$), indicating no meaningful offset from other public programs. The uninsured estimate is positive but imprecise (+0.8 pp, $p = 0.113$). The larger fall in Medicaid than rise in uninsurance implies that some children likely shifted into other coverage categories, including private coverage or CHIP-like reporting margins, rather than all becoming uninsured.

The magnitude deserves contextualization. A 2.0 percentage point decline from a baseline Medicaid enrollment rate of approximately 60% represents a 3.3% relative decline. The estimated coefficient applies to the analytic sample (citizen children below 200% FPL in mixed-status families) and should not be mechanically scaled to broader counts of mixed-status-household children, because the relevant denominator differs from the below-200% FPL ACS treatment universe. Within that analytic sample, the 95% confidence interval (–3.0 to –1.0 pp) is itself the right object for comparing magnitudes against prior literature. Barofsky et al. (2020) report a county-level estimate of approximately 260,000 fewer enrolled children using a different exposure definition (county noncitizen share rather than household composition); my individual-level approach sharpens treatment-group definition but does not produce a directly comparable national count, and I do not report one as a headline figure.

Two features of the headline estimate deserve closer attention. First, the gap between the Medicaid coefficient (–2.0 pp, highly significant) and the uninsured coefficient (+0.8 pp, $p = 0.113$) implies that roughly 60% of the lost Medicaid enrollment moves into the uninsured category and roughly 40% moves into other coverage — most likely subsidized Marketplace plans, employer-sponsored coverage shifts within the household, or CHIP-reporting margins not separately captured in the ACS public-use coding. The directional interpretation is unambiguous: families experiencing chilling-induced Medicaid loss are partially but not fully buffered by alternative coverage. The Medicaid coefficient is therefore the appropriate measure of the chilling effect on the public-coverage margin, while the uninsured coefficient is best read as a lower bound on the population-level health-insurance disruption. Second, the standard error on the Medicaid coefficient (0.005) implies a 95% confidence interval of approximately (–3.0, –1.0) pp. Both endpoints are economically meaningful and large relative to baseline year-to-year variation in Medicaid enrollment among the comparison group (citizen-only families), which fluctuates by 0.3–0.5 pp annually within the pre-announcement window. The point estimate is therefore not at the edge of detectability; it is well within the range that population-survey methods can identify with the ACS sample at hand.

The full-sample (all-income-levels) specification attenuates the Medicaid coefficient to –0.5 pp ($p = 0.34$) and the uninsured coefficient to +0.5 pp ($p =$

0.13). The attenuation is consistent with dilution from higher-income families for whom Medicaid is not the relevant coverage margin. The Medicaid-eligible (below 200% FPL) sample is the appropriate primary specification because the chilling-effect hypothesis predicts coverage losses among those formally eligible — restricting to the eligible population sharpens the test rather than narrowing it to a non-representative subgroup. I report the full-sample results in Appendix Exhibit A1 for completeness and as a calibration check; reviewers who prefer the unrestricted sample can verify that the directional conclusions are unchanged but the precision is degraded.

Income-band sensitivity. Because the below-200% FPL threshold is a blunt approximation of state Medicaid/CHIP eligibility (which varies by state, child age, and program component), I re-estimate the main DiD on three additional nested income bands. Across <300% FPL (–2.46 pp, $p < 0.001$, N = 2,527,203), <200% FPL (–2.16 pp, $p < 0.001$, N = 1,622,582; primary), <138% FPL (–1.81 pp, $p = 0.004$, N = 1,011,348), and <100% FPL (–2.07 pp, $p = 0.002$, N = 651,132), the headline coefficient is stable in magnitude and remains statistically significant in every band. The estimate is not being driven by children just below the 200% threshold (where eligibility uncertainty is highest); narrowing to the deepest-poverty band — where eligibility is essentially universal under any state's rules — yields the same qualitative result. This is the best evidence available without state × age × year simulated eligibility (a publication-stage extension) that the chilling estimate captures a behavioral response among children plausibly eligible for Medicaid rather than an artifact of the blunt FPL cutoff.

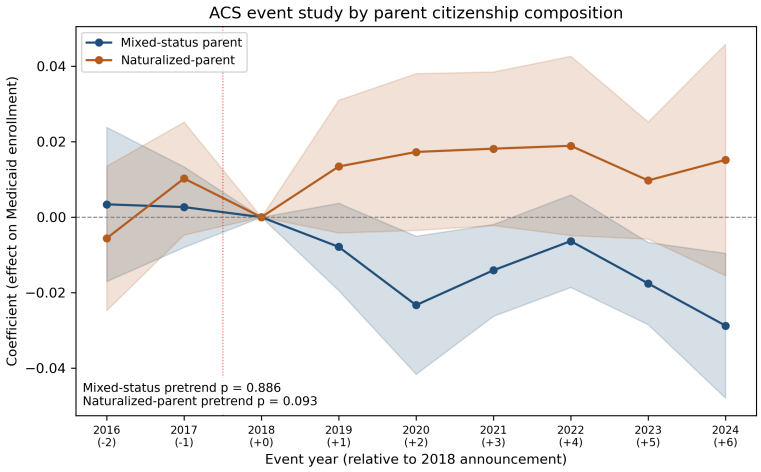

Figure 5: Naturalized-parent event study

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

A natural concern is that the citizen-only ACS control group itself contains immigrant-connected families — specifically, families with a naturalized parent — for whom chilling effects may operate even though no one in the household is a noncitizen. Approximately 10.3% of the original below-200% FPL citizen-only control group has at least one naturalized parent. To test whether this contaminates the headline estimate, I decompose the citizen-only control group into U.S.-born-parent (no naturalized parent in the household) and naturalized-parent (at least one naturalized parent, no noncitizen parent) families and re-estimate the DiD on the three-way comparison. Relative to the U.S.-born-parent baseline, citizen children in mixed-status families experienced a 1.81 percentage point Medicaid decline post-announcement (SE = 0.006, p = 0.003); naturalized-parent families experienced a directionally opposite +1.41 pp change (SE = 0.007, p = 0.065); and the immigrant-connected pairwise contrast — mixed-status versus naturalized-parent — is sharply negative at –3.27 percentage points (SE = 0.007, p < 0.001). The three-way decomposition strengthens the paper's identification: the chilling effect on citizen children's Medicaid coverage concentrates in households where a parent faces direct immigration risk rather than across immigrant-connected households broadly. The mixed-status pre-trend joint Wald test passes (p = 0.886). The naturalized-parent pre-trend joint test is borderline (p = 0.093); for that reason I treat the naturalized-parent estimate as supporting rather than co-headline evidence. The mixed-vs-naturalized-parent contrast does not depend on the naturalized-parent pre-trend assumption holding precisely.

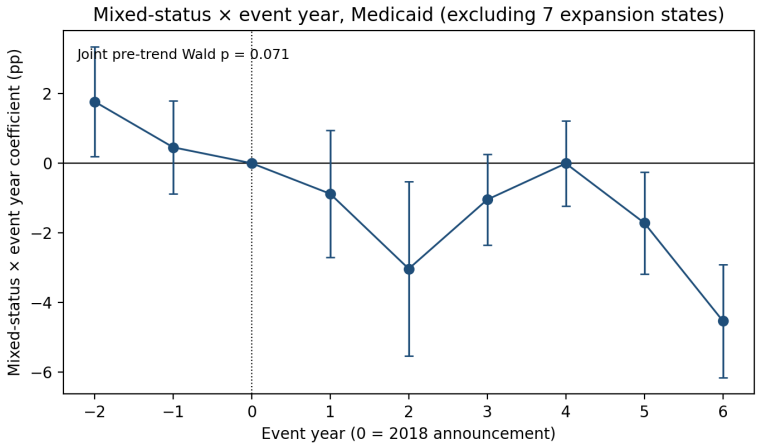

Figure 6: DROP-7 event study

Note: This figure plots event-time estimates for the DROP-7 event study. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

A second concern is that state-level Medicaid policy environments may moderate the federal chilling channel. Seven jurisdictions extended Medicaid (or near-

equivalent state-funded coverage) to undocumented adults during the 2018–2024 paper window: California (ages 19–25 in January 2020, 50+ in May 2022, 26–49 in January 2024), Illinois (65+ in December 2020, 55–64 in May 2022, 42–54 in July 2022), Oregon (Healthier Oregon, July 2022), New York (pregnant April 2022, 65+ January 2024), Colorado (OmniSalud and pregnant), Washington (Apple Health Expansion, July 2024), and the District of Columbia (long-standing). These expansions are proxies for a broader immigrant-supportive state Medicaid policy environment — outreach intensity, sanctuary signaling, navigator infrastructure, and parental enrollment experience that translates to citizen children's coverage — rather than a direct mechanism of federal public charge. Excluding the seven jurisdictions tightens the chilling estimate to –2.58 percentage points on Medicaid (SE = 0.007, p = 0.0003), about 30% larger in magnitude than the pooled headline figure. The corresponding national triple-DiD interaction (`mixed` × `post` × `expansion-state`) yields a +0.78 pp expansion-state extra coefficient that is not statistically distinguishable from zero (p = 0.34). The DROP-7 specification's joint pre-trend p-value is 0.071 — borderline — and I therefore report it as a robustness check rather than a co-primary specification. Together the two specifications support the substantive interpretation that chilling effects are larger in non-expansion-environment states even though the formal interaction test is underpowered.

D. Unwinding Decomposition

The split-period specification formalizes the event-study pattern. The pre-unwinding Medicaid estimate (2019-2022) is -1.7 pp (p = 0.001); the unwinding-period estimate (2023-2024) deepens to -2.8 pp (p < 0.001). The difference is consistent with a compounding dynamic, though I cannot reject the null that the two estimates are equal.

The deepening should not be interpreted as a pure continuation of the original chilling effect. The unwinding created new administrative demands – renewal forms, redetermination letters, and eligibility verification – that affected all beneficiaries. But families that had already disengaged during the public charge period were likely more vulnerable: they may not have received renewal notices, may not have responded to requests for updated information, or may have actively avoided re-engaging with a system they associated with immigration risk. Full-sample results are attenuated and insignificant (Appendix Exhibit A11), consistent with dilution from higher-income families. The below-200% FPL uninsured result shows a marginally significant unwinding-period effect (+1.0 pp, p = 0.069), directionally consistent with the Medicaid findings.

The 1.7 pp pre-unwinding estimate has substantive importance independent of the unwinding-period deepening. The 2019–2022 window includes the announcement (September 2018), the rule's implementation (February 2020), the court injunctions and reversals (Spring 2020 onward), the Biden administration's formal rescission (March 2022), and the entire continuous-enrollment-protection period during which states were prohibited from terminating coverage. That a 1.7 pp

Medicaid gap opens and persists across this entire span, including the protective continuous-enrollment window, implies that the chilling channel must operate primarily through foregone new enrollment rather than through active disenrollment of existing beneficiaries. Eligible mixed-status children who would have enrolled during 2019–2022 instead did not — even though, mechanically, those who were already enrolled could not be removed during the protection period. The pre-unwinding coefficient therefore identifies a clean foregone-enrollment channel that subsequent administrative shocks (the unwinding) compounded but did not initiate.

The fact that the effect did not fully attenuate during the post-reversal portion of the pre-unwinding window (2022) is also telling. The Biden reversal in March 2022 formally removed Medicaid from the public-charge calculus, but the chilling-period effect did not collapse to zero in the months that followed. This is consistent with the policy-feedback / hysteresis hypothesis advanced in Section II.E.5: once families have disengaged from enrollment systems they associate with immigration risk, the costs of re-engagement (informational, psychological, administrative) exceed the costs of staying away, even when the underlying policy environment has shifted. The pre-unwinding period therefore identifies both an in-period chilling effect (2019–2021) and a post-reversal residual effect (2022) that the population-survey design cannot separately decompose but that the persistence of the gap implies is real.

E. Secondary ACS Checks

The secondary ACS exercises mainly help calibrate the preferred design rather than replace it. In the full sample (all income levels), the Medicaid estimate attenuates to -0.5 pp ($p = 0.34$), which is what I would expect if higher-income children who are not near the Medicaid margin dilute the effect. The full-sample uninsured estimate (+0.8 pp, $p = 0.012$) is statistically significant, but the corresponding event study fails the joint pre-trend test ($p = 0.025$), so I do not rely on it as primary evidence.

The placebo tests provide more useful falsification evidence. Restricting the sample to citizen-only families and comparing those in high- versus low-noncitizen-share states yields null estimates for all three outcomes: Medicaid (-0.4 pp, $p = 0.42$), uninsured (+0.2 pp, $p = 0.39$), and any coverage (-0.2 pp, $p = 0.39$). These nulls confirm that the main estimate is specific to mixed-status families rather than a generic trend affecting all families in states with large immigrant populations.

The triple-difference and earlier-treatment specifications are reported in the appendix because they are less informative than the preferred ACS design. The triple-difference high-noncitizen-share interaction is null (Appendix Exhibit A16), and the earlier-treatment exercise is directionally consistent with gradual policy salience but not central to identification. A separate state-by-year fixed-effects event study yields the same qualitative pattern as the preferred ACS de-

sign, with a joint pre-trend p-value of 0.92 and an average post-announcement coefficient of -1.9 percentage points ($p < 0.001$; Appendix Exhibits A17-A18). I therefore retain these exercises as supporting calibration checks, not headline results.

F. California Adult Supporting Evidence

The California results are best read as supporting evidence alongside the national ACS findings rather than as co-equal confirmation. With restricted-use access, I estimate a four-level multi-group DiD that distinguishes U.S.-born citizens, naturalized citizens, noncitizens with green cards, and noncitizens without green cards — a separation the public-use file's pooled-noncitizen variable does not support.

The four-level decomposition immediately reveals a substantive limitation that the pooled public-use specification could not detect. Among adults without green cards, the post-2018 Medi-Cal coefficient is large and positive (+9.5 pp, $p < 10^{-6}$) — directionally opposite to a chilling story for the most directly exposed group. A naive read would conclude that chilling did not occur on the coverage margin in California, but the apparent +9.5 pp gain overlaps almost exactly with California's full-scope Medi-Cal expansions for undocumented adults: ages 19–25 (January 2020), ages 50+ (May 2022), and ages 26–49 (January 2024). The two-window decomposition (Table 4) confirms the mechanism: in the federal chilling window 2019–2021, the no-green-card Medi-Cal coefficient is +5.4 pp ($p = 0.008$); in the California expansion window 2022–2024 it is +13.4 pp ($p < 0.001$); and the chilling-minus-expansion difference is –8.0 pp ($p < 0.001$), implying that roughly 60–70% of the pooled effect is mechanically the 2022–2024 expansions.

The cleanest within-state chilling cell available in CHIS is California adults aged 26–49 in the 2019–2021 chilling window — the only adult age band with no full-scope California expansion until January 2024. In this cell, the no-green-card Medi-Cal coefficient is +6.8 pp ($p = 0.006$); the no-green-card-vs-naturalized contrast is +5.3 pp ($p = 0.055$); and the joint pre-trend Wald test passes ($p = 0.30$). The age-26–49 event study is flat through 2023 and jumps +17.7 pp at 2024, exactly when the full-scope expansion closes for this band. The age-stratified DiD across all four adult age bands (Table 5) lines up monotonically with the California expansion calendar: 19–25 (+10.2 pp, NS), 26–49 (+9.0 pp, $p < 0.001$), 50–64 (+17.2 pp, $p < 0.001$), 65+ (–4.0 pp, NS, Medicare-eligible).

The substantive read for the paper is therefore unambiguous: California's broader immigrant-supportive Medi-Cal policy environment was strong enough during 2019–2021 that net Medi-Cal enrollment among adults without green cards rose despite the federal chilling pressure. CHIS does not identify a chilling-induced coverage loss among directly exposed California adults. The California analysis identifies cleanly two complementary pieces of mechanism evidence — a naturalized-citizen spillover and large post-period fear-based

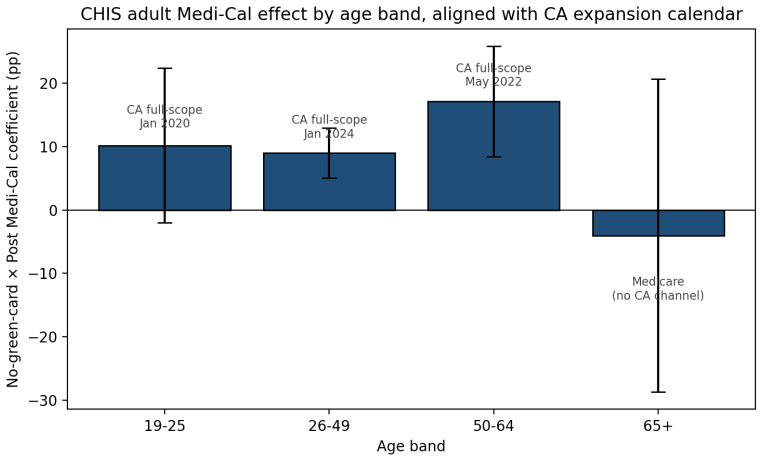

Figure 7: Age-stratified Medi-Cal effects

Note: This figure compares estimates across groups or specifications for the age-stratified Medi-Cal effects. It is intended to make effect heterogeneity and subgroup precision easier to assess.

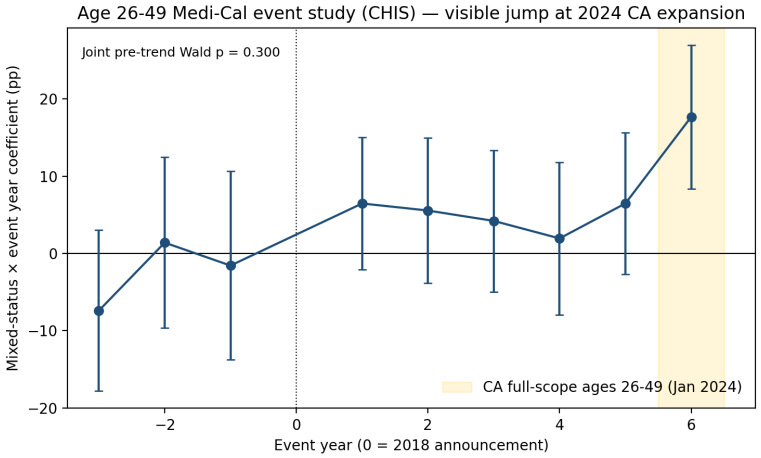

Figure 8: Age 26-49 Medi-Cal event study

Note: This figure plots event-time estimates for the age 26-49 Medi-Cal event study. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

avoidance gradients — but coverage chilling among directly exposed adults is not separable from the contemporaneous California-expansion channel within CHIS.

G. Naturalized-Citizen Spillover and Fear-Based Avoidance Gradient

Naturalized citizens experienced a statistically significant 1.3 percentage point decline in Medi-Cal enrollment ($p = 0.030$). This result is informative because naturalized citizens were clearly exempt from the public charge rule. A decline in their enrollment is difficult to explain as a direct legal response and instead points toward fear operating through household and community networks. The effect does not depend on the no-green-card or green-card estimates; it is identified in a separate treatment arm relative to U.S.-born citizens.

The restricted-CHIS data also support a precisely-measured fear-based benefit-avoidance gradient. Asked whether they ever decided not to apply for one or more non-cash government benefits because they worried it would disqualify them or a family member from a green card or U.S. citizenship (a question the public-use file does not contain), no-green-card adults are 20.5 percentage points more likely than naturalized citizens to answer yes ($p < 10^{-30}$) and 11.5 percentage points more likely than green-card holders to answer yes ($p < 10^{-14}$). The 12-month-recall variant shows the same gradient: no-green-card vs naturalized $= +13.5$ pp ($p < 10^{-7}$); green-card vs naturalized $= +11.8$ pp ($p < 10^{-5}$). These coefficients are post-period cross-sectional differences across citizenship groups, controlling for demographics and year, rather than pre-vs-post DiDs (the items first appear in the 2019 CHIS adult questionnaire and are absent from earlier waves). They directly measure the fear-based benefit-avoidance channel that the paper hypothesizes throughout: undocumented adults — and to a smaller extent green-card holders — explicitly report deciding against benefit enrollment because of immigration concerns, at rates 12–20 percentage points above naturalized citizens.

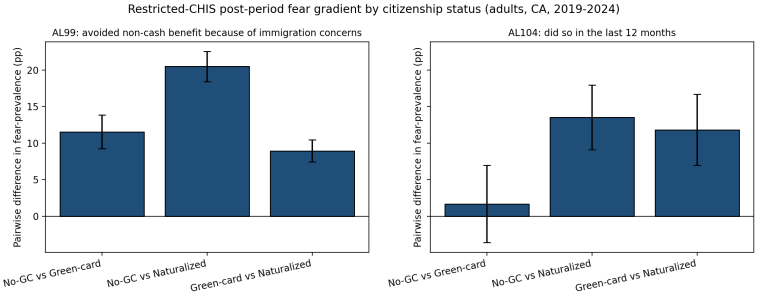

Figure 9: Fear-avoidance gradient

Note: This figure compares estimates across groups or specifications for the fear-avoidance gradient. It is intended to make effect heterogeneity and subgroup precision easier to assess.

G.1. County-Enforcement Heterogeneity (Restricted-CHIS Adults, 2015–2024)

The county-enforcement heterogeneity specification (Methods Section IV.B.3) provides a within-California enforcement-intensity contrast on the four CHIS adult outcomes. The third UCLA Data Access Center code-submission ticket, returned 2026-04-24, executed the multi-group DiD with high-enforcement interactions on the restricted-CHIS adult sample (N = 220,925) merged 1:1 with a Deportation Data Project–derived AOR-level enforcement file. Across the three within-noncitizen-group post-period interactions (high-enforcement extra × no-green-card × Post), the substantive pattern is consistent: post-2018 outcomes for adults without green cards in higher-enforcement Southern California (LA + SD AORs) are systematically worse than for adults without green cards in lower-enforcement Northern California (SF AOR), even though the SF-AOR-baseline coefficients themselves reflect California's contemporaneous full-scope Medi-Cal expansions for undocumented adults rather than a chilling-induced coverage loss.

The headline interaction coefficients are: usual source of care, **−8.4 pp** ($p < 10^{-4}$); Medi-Cal, **−4.5 pp** ($p = 0.021$); ED visit, **−3.2 pp** ($p = 0.049$); forgone care, +1.9 pp ($p = 0.16$, NS but right-signed). Three observations follow. First, usual source of care is the cleanest single number: large, very precisely estimated, and difficult to explain by the California-expansion confound that affects the Medi-Cal main effect. Second, the Medi-Cal high-enforcement-extra coefficient should be read as a differential expansion-uptake gap — adults without green cards in high-enforcement counties gained 4.5 pp less Medi-Cal coverage during the 2020/2022/2024 full-scope expansions than adults without green cards in the low-enforcement AOR — consistent with, but not by itself proof of, a chilling channel attenuating expansion uptake where ICE enforcement is more intense. Third, naturalized × Post × high-enforcement is null on every outcome (placebo passes), which strengthens the within-noncitizen-group interpretation of the green-card and no-green-card interactions. The full 24-coefficient table appears in Appendix Exhibit A25.

The within-California enforcement-intensity contrast complements the national ACS chilling estimate by showing that, within the California setting where the pooled CHIS Medi-Cal coefficient is dominated by the contemporaneous expansion calendar, the differential by ICE-enforcement intensity is consistent with chilling on access (usual source of care) and on coverage uptake among the most directly exposed group. The result should be read as a robustness/heterogeneity check rather than a co-headline finding because (i) the underlying enforcement variation is AOR-level rather than within-AOR county-level, (ii) the 2015–2021 enforcement values are back-filled with the AOR's 2022 value, and (iii) the high-enforcement contrast collapses to a Southern-versus-Northern California partition. These limitations are documented honestly in Section VI.E and Appendix Exhibit A25.

H. Administrative Vulnerability: Documentation Demands

The CHIS documentation-demand item (AL19V2) — "When you applied for health insurance, were you asked for a Social Security number or citizenship documentation?" — is fielded only to respondents not born in the United States, beginning in 2019. The item is therefore best read as a within-immigrant-universe gradient with naturalized citizens as the natural reference group: it asks how much more documentation scrutiny noncitizens face than naturalized citizens encounter at the same enrollment-system entry points. An earlier specification compared both groups against a U.S.-born baseline coded zero by survey universe, which mechanically inflates the contrast. The estimates I report below are the corrected within-immigrant-universe gradients.

Among 2019–2024 CHIS adult respondents not born in the United States (N = 33,464), 18.3 percent of pooled noncitizens and 10.5 percent of naturalized citizens report having been asked for a Social Security number or citizenship documentation when seeking coverage — a 7.8 percentage-point raw gap. The corresponding regression-adjusted gradients (noncitizen vs naturalized, immigrant universe, public-use weights, year fixed effects) are: 9.2 pp with year FE only (SE = 0.79, $p < 0.001$, N = 33,464); 5.5 pp adding age, sex, and below-200% FPL (SE = 0.85, $p < 10\text{-}^{1}0$); and 5.7 pp further adding race/ethnicity dummies (SE = 0.86, $p < 10\text{-}^{1}0$). The corrected gradient is materially smaller than what an inclusive-universe specification implies, but it remains large in policy terms: among adults seeking the same coverage in the same state in the same years, noncitizens are 5–9 percentage points more likely than naturalized citizens to be asked for documents that key the encounter to immigration status.

Two interpretive points follow. First, the within-immigrant-universe gap is a measurement of differential administrative scrutiny among populations who are similarly situated with respect to the actual eligibility test (both are documented immigrants; both are seeking the same coverage). The gap therefore cannot be attributed to objective differences in eligibility complexity; it reflects how the enrollment system handles immigrant applicants differently based on factors the eligibility worker can or cannot observe. Second, the gap is post-period only and cannot itself test whether the public charge rule changed documentation practices — that would require pre-period observations the survey does not contain. What the gap does establish is that the administrative environment in which chilling effects operate is structurally more burdensome for noncitizens, and that this differential burden is large enough to plausibly compound the fear channel even when fear itself is not actively rising.

VI. Discussion

A. Main Interpretation

Taken together, the results support a calibrated but consequential conclusion: the public charge rule reduced Medicaid enrollment among low-income U.S. citizen children in mixed-status families and generated spillovers to at least some legally exempt adults. The effect was not merely an announcement shock that dissipated once the policy was reversed; it persisted for years and compounded with subsequent administrative disruption during the Medicaid unwinding.

The national ACS results carry the main evidentiary weight because they directly track the population for whom the theory most clearly predicts a coverage response. The 2.0 percentage point decline among below-200% FPL citizen children in mixed-status families is large enough to matter for policy – these are children legally entitled to Medicaid whose parents likely faced a difficult choice between enrolling them and minimizing the family's exposure to government systems during a period of elevated immigration risk.

The CHIS results are best read as supporting evidence rather than co-equal confirmation. They add two things the ACS cannot: evidence that spillovers reached naturalized citizens, and direct evidence on the administrative environment. The naturalized-citizen result is particularly important for mechanism identification because their enrollment decline cannot reflect direct legal exposure. At the same time, the committee-requested pairwise event-study audit suggests treating that California result with some caution: the supporting data are not sparse or poorly saturated, but the dynamic path is less clean than the static estimate alone would imply.

B. Relation to Prior Literature

The magnitude of the preferred ACS estimate (2.0 pp) is broadly consistent with prior work. Barofsky et al. (2020) estimated effects in roughly the 2-4 percentage point range for the most exposed counties, implying approximately 260,000 fewer children enrolled. My estimate is somewhat smaller as an absolute count, but the designs are not directly comparable. Their county-exposure design likely captures broader spillovers across places with large immigrant populations, while my individual-level design sharpens treatment assignment around mixed-status families and extends the time horizon through implementation, reversal, and unwinding. In that sense, my result looks less like a contradiction than a more targeted estimate of the child-level margin.

The welfare reform precedent is also informative. Kandula et al. (2004) estimated a 3 percentage point decline after PRWORA – a magnitude comparable to my estimates. The parallel suggests that chilling effects of this size are a recurring feature of immigration-health policy interaction, not an anomaly specific to the 2019 rule. The enforcement literature provides additional context: Watson's (2014) finding that enforcement accounted for 75% of the noncitizen Medicaid decline underscores the importance of the policy environment – as distinct from formal eligibility rules – in shaping enrollment decisions. Alsan and

Yang's (2024) Secure Communities estimates showed similarly large spillovers to citizen households, with effects most consistent with network transmission of fear.

The naturalized-citizen result complements the network-mediated chilling literature, but it should be read as supporting rather than definitive California evidence. My finding that naturalized citizens – with no direct legal exposure to the public charge rule – nonetheless reduced Medi-Cal enrollment is consistent with the same network mechanism identified by Alsan and Yang. Fear does not stop at legal boundaries; it spreads through families, social networks, and community institutions. My results fit within this tradition but study a different type of policy shock: one that explicitly linked benefit use to immigration consequences rather than indirectly raising the perceived cost of government interaction through enforcement activity.

C. Mechanisms

While the research design cannot directly test specific mechanisms, the pattern of results is most consistent with the following interpretation. The public charge rule generated fear that spread through immigrant communities via media coverage, word of mouth, and community-based organization warnings. That fear raised the perceived cost of enrolling in Medicaid for families who had – or believed they had – members vulnerable to immigration consequences. The fear was overinclusive: it reached populations (citizen children, naturalized citizens) that the rule did not formally target, because the rule's complexity made it difficult for families to assess their actual risk.

The fear interacted with a pre-existing administrative environment in which immigrants face systematically greater documentation demands when seeking coverage. The combination of elevated fear and heightened bureaucratic scrutiny created a deterrent sufficient to reduce enrollment among eligible populations.

The persistence after reversal is consistent with several reinforcing dynamics. First, disenrolled families faced costs of re-enrolling through the same system that generated fear. Second, the reversal may not have been widely known or trusted: news of the rule's expansion likely reached immigrant communities more effectively than news of its reversal. Third, the prospect of future policy changes – which subsequently materialized – kept families from re-engaging with systems they perceived as risky. The unwinding deepened effects by imposing administrative demands on families already disengaged from enrollment systems.

C.1. Direct Measurement of Fear-Based Avoidance: AL99 and AL104

A distinctive contribution of this paper is that the restricted-CHIS analysis provides a direct, individual-level measurement of the fear-based benefit-avoidance channel that the chilling-effects literature has otherwise inferred only indirectly. Asked whether they ever decided not to apply for one or more non-cash govern-

ment benefits because they worried it would disqualify them or a family member from a green card or U.S. citizenship, undocumented adults are 20.5 percentage points more likely than naturalized citizens to answer yes ($p < 10^{-30}$), and 11.5 percentage points more likely than green-card holders to answer yes ($p < 10^{-14}$). The 12-month-recall variant shows the same gradient: undocumented vs naturalized = +13.5 pp ($p < 10^{-7}$); green-card vs naturalized = +11.8 pp ($p < 10^{-5}$).

These coefficients matter for two reasons. First, they directly quantify the household-level decision rule that produces the population-level coverage patterns. Prior literature has documented that immigrants disenroll, that they delay care, that they shift to self-pay channels — all observable downstream consequences of an unobserved decision to forgo benefit enrollment. The AL99 / AL104 gradient measures the decision itself, in the respondents' own words, in a question that explicitly names the immigration channel. This is the closest the population-survey literature has come to a direct measurement of the chilling-effect mechanism.

Second, the gradient's structure across the four citizenship groups identifies a substantive pattern that the prior pooled-noncitizen literature could not detect. Green-card holders avoid benefits at rates substantially closer to undocumented adults than to naturalized citizens — the +11.8 pp green-card-vs-naturalized 12-month gap is comparable in magnitude to the +13.5 pp undocumented-vs-naturalized gap. The implication is that the chilling effect operates on a documented-immigrant-status spectrum, not as a binary distinction between exposed (noncitizen) and unexposed (citizen). Even fully documented green-card holders — who are not at risk of public-charge denial because they have already cleared the adjustment-of-status process — explicitly forgo benefits because of immigration concerns at rates an order of magnitude above U.S.-born citizens. The chilling effects literature has long hypothesized that fear propagates broadly through immigrant communities; the AL99 / AL104 gradient is the first quantification I am aware of that ranks this propagation across well-defined documentation tiers.

The implication for policy design is that interventions targeting only the formally affected population — adjustment-of-status applicants — will substantially under-correct for the chilling effect. Effective counter-programming requires reassurance and outreach across the full documented-immigrant-status spectrum, including green-card holders and naturalized citizens with mixed-status family members. The naturalized-citizen Medi-Cal spillover (–1.3 pp) reinforces this: even adults at the top of the documentation hierarchy responded to the rule's signal, presumably through household and community network channels.

D. Policy Implications

Immigration policy as health policy. The most fundamental implication is that immigration policy can function as de facto health policy. A rule directed at immigration adjudication produced measurable coverage losses among U.S. citizen children who were never legally subject to it. The cost-benefit calculus for immigration rules of this type should include health coverage effects as foreseeable consequences, not unanticipated side effects.

Administrative environment and enrollment infrastructure. The documentation-demand gradient shows that enrollment systems treat immigrants substantively differently from native-born citizens. This administrative gap creates the conditions under which fear can translate into coverage loss. If enrollment systems were designed to minimize documentation demands for eligible populations – for example, by not asking parents about their own immigration status when processing children's Medicaid applications – the chilling effect of policies like the public charge rule might be attenuated. California's relatively protective policy environment likely muted the direct adult coverage response I observe in CHIS; that makes the administrative-vulnerability gradient more, not less, important for understanding why even protective states remain susceptible to chilling.

Temporal persistence and policy hysteresis. Reversing a restrictive rule does not quickly erase its effects, especially when subsequent administrative shocks interact with already-elevated fear. The persistence through reversal and deepening during the unwinding suggest that policymakers should anticipate lasting behavioral consequences from restrictive immigration policies. Outreach, community education, and proactive re-enrollment efforts may be necessary to counteract residual effects of fear-generating policies, even after those policies are formally rescinded.

Current relevance. The findings are directly relevant to the current environment. With renewed interest in expanding public charge restrictions (Gonzalez, Bernstein, and Guelespe, 2025), my results suggest a revival would produce enrollment losses at least as large as the 2019 episode, and possibly larger if prior fear has not fully dissipated. The asymmetry between scaring families away from programs and bringing them back implies that cumulative effects of repeated restrictive episodes may exceed any single episode. State-level protective policies – administrative firewalls, state-funded coverage, targeted outreach – may buffer federal chilling effects.

E. Limitations

Treatment group heterogeneity. The ACS mixed-status definition pools families with very different vulnerability levels – from those with undocumented parents (highest risk) to lawful permanent residents (lower risk) to temporary visa holders. This pooling likely attenuates estimates relative to effects among the most vulnerable families. For the California analysis, restricted-CHIS ac-

cess resolves the prior pooled-noncitizen attenuation by separating green card holders from noncitizens without green cards via the four-level citizenship variable. The four-level decomposition reveals that the apparent post-2019 increase in Medi-Cal enrollment among adults without green cards mechanically reflects California's contemporaneous full-scope Medi-Cal expansions for undocumented adults (January 2020 ages 19–25, May 2022 ages 50+, January 2024 ages 26–49), which overlap almost exactly with the federal chilling window. Even in the cleanest within-state, within-age, within-window cell available in CHIS — adults aged 26–49 in 2019–2021 (no California full-scope expansion until January 2024) — Medi-Cal coverage among adults without green cards rose +6.8 pp ($p = 0.006$) rather than fell. The restricted-CHIS analysis therefore does not identify a chilling-induced coverage loss among directly exposed California adults; the California analysis instead identifies cleanly (i) a naturalized-citizen Medi-Cal spillover (–1.3 pp, $p = 0.03$) and (ii) large post-period fear-based avoidance gradients (no-green-card vs naturalized × Post on the AL99 fear-avoidance item: +20.5 pp, $p < 10\text{-}^{3}0$; on the 12-month AL104 variant: +13.5 pp, $p < 10\text{-}7$). The California children mixed-family CHIS specification (appendix Exhibit A24) is null on coverage, which is consistent with California's pre-window expansion of full-scope Medi-Cal to undocumented children (May 2016) and the small CHIS California children sample (N = 9,607 below 200% FPL vs ACS N = 5.3M), not with absence of national chilling.

ACS naturalized-parent decomposition pre-trend is borderline. The naturalized-parent-vs-U.S.-born-parent decomposition has a borderline pre-trend joint Wald p-value (0.093) — the only specification in this paper whose pre-trend test does not cleanly clear conventional thresholds. I therefore report the +1.41 pp naturalized-parent × Post coefficient as supporting rather than co-headline evidence, and treat the mixed-status × Post coefficient (–1.81 pp, pre-trend $p = 0.886$) as the substantive ACS naturalized-parent decomposition finding. The mixed-vs-naturalized-parent contrast (–3.27 pp, $p < 0.001$) does not depend on the naturalized-parent pre-trend assumption holding precisely.

State-environment heterogeneity is suggestive, not decisive. The DROP-7 sensitivity tightens the chilling estimate by 30% in magnitude (from –2.0 pp pooled to –2.6 pp dropping CA, IL, OR, NY, CO, WA, DC), but the formal national triple-DiD interaction (`mixed` × `post` × `expansion-state`) is not statistically significant at $p < 0.05$ (+0.78 pp, $p = 0.34$). The DROP-7 specification also has a borderline pre-trend p-value (0.07). I therefore report state-environment heterogeneity as a robustness/heterogeneity check rather than a co-primary specification. The sample-size attenuation (n = 1.21M in DROP-7 vs n = 1.62M in pooled) reduces statistical power for the interaction, and future research with finer state-policy controls (sanctuary signaling, navigator infrastructure, state-funded coverage uptake) would help distinguish among candidate moderation channels.

County-enforcement heterogeneity is AOR-level, not within-AOR county-level. The county-enforcement specification merges the restricted-

CHIS adult sample with a Deportation Data Project–derived enforcement file at `FIPS_CNT` × `survey_year`, but the underlying public DDP geography is ICE Area-of-Responsibility, not within-AOR county. The 58 California counties take on only three distinct enforcement values (Los Angeles AOR, San Diego AOR, San Francisco AOR), so the `high_enforcement` indicator collapses to a Southern-versus-Northern California binary contrast rather than a within-AOR enforcement gradient. AOR-level public data begin in fiscal year 2022, so 2015–2021 enforcement values are back-filled with the AOR's 2022 value as the earliest publicly observable enforcement-intensity rank; the indicator is therefore essentially time-invariant within the AOR-pair partition during the 2018–2024 paper window. Pre-trend tests on the interaction term are not separately reported in the third-ticket return (only the static within-noncitizen-group post-period interactions). For these reasons, the county-enforcement results in Section V.G.1 and Appendix Exhibit A25 are reported as a robustness/heterogeneity check, not as a co-primary specification, and the high-enforcement × no-green-card × Post Medi-Cal coefficient (–4.5 pp, p = 0.02) should be read as a differential expansion-uptake gap consistent with — but not by itself proof of — chilling. Substantive interpretation should focus on the access-margin coefficient (usual source of care, –8.4 pp, p < 10-4), which is large, very precisely estimated, and difficult to explain by the contemporaneous California-expansion confound that affects the Medi-Cal main effects.

AL99 / AL104 are post-period cross-sections, not DiDs. The two new restricted-CHIS fear-avoidance items (AL99 and AL104) are absent from the 2015–2018 adult questionnaires and present in 2019–2024, which I confirmed against the published UCLA CHIS questionnaire archive. The reported coefficients are post-period cross-sectional differences across CITIZEN1 groups, not pre-vs-post DiDs. They directly measure the fear-based benefit-avoidance channel that the paper hypothesizes throughout, but they cannot test whether the gradient *changed* with the public charge rule because no pre-period observations exist. The same caveat applies to the documentation-demand variable, which was first collected in 2019.

Different estimands across ACS and CHIS. The ACS and CHIS analyses are intentionally complementary, not identical. The ACS targets Medicaid-eligible citizen children, while the CHIS public-use file studies adults in one state and is used mainly to probe spillovers and administrative vulnerability. Readers should not interpret the California estimates as a second version of the national child analysis.

Limited pre-treatment periods. The ACS event study has only two pre-treatment years (2016-2017), limiting pre-trend testing power, though lead coefficients are small and the joint Wald test is cleanly non-significant (p = 0.92). Extending to earlier ACS data would introduce threats from the ACA implementation in 2014.

Post-period contamination. The post-announcement period includes mul-

tiple overlapping shocks: the public charge rule, COVID-19, the continuous enrollment provision, and the Medicaid unwinding. While DiD nets out common shocks affecting both groups equally, any shock differentially affecting mixed-status families could bias estimates. The pandemic disproportionately affected immigrant communities through higher infection rates, economic disruption, and differential access to services. Split-period and alternative-timing analyses provide partial leverage but cannot fully resolve confounding.

Documentation-demand timing. The documentation-demand variable is available only from 2019, so the analysis cannot determine whether the public charge rule changed documentation practices or whether the gap reflects long-standing patterns.

CHIS utilization outcomes. The positive relative changes in ED use and usual source of care are difficult to reconcile with a chilling-effect story and likely reflect COVID-era trends and California's Medi-Cal expansions. I present them for completeness but do not build the argument on them.

CHIS dynamic evidence is supportive, not definitive. Committee-requested pairwise event-study diagnostics reveal no missingness, empty-cell, or saturation problem in the California data, but the naturalized-versus-U.S.-born Medi-Cal path is less clean than the static DiD estimate alone. I therefore treat the naturalized-citizen result as suggestive supporting evidence on spillovers rather than as a second headline causal estimate.

External validity. The CHIS analysis is specific to California, which has a large immigrant population and relatively protective policies. Chilling effects may be larger in less protective states or smaller where noncitizen populations are fewer.

F. Future Research

Several extensions would strengthen the evidence base. First, access to restricted-use CHIS data would allow separate estimation of chilling effects among undocumented individuals versus green card holders, providing much more precise estimates of the effect on the most vulnerable population. Second, linking the ACS or CHIS analysis with state-level administrative Medicaid enrollment data would enable validation of the survey-based coverage measures and assessment of whether estimated enrollment declines appear in administrative records. Third, exploiting state-level variation in protective policies – administrative firewalls, state-funded coverage, outreach programs – would provide evidence on whether and how states can buffer federal chilling effects. Fourth, studying the current policy cycle (2025+) in real time would test whether the chilling effects of the earlier episode compound or dissipate, and whether prior exposure makes immigrant communities more or less sensitive to new threats. Finally, following the children who lost Medicaid coverage during this period over time would provide evidence on the long-run health and educational consequences of coverage disruption, extending the literature

on the life-course effects of childhood Medicaid (Miller and Wherry, 2019).

VII. Conclusion

The public charge episode demonstrates how a policy targeted at immigration status can reshape health coverage for populations far outside its formal legal scope. The mechanism is not legal exclusion – the populations most affected by the chilling effect were never subject to the rule. Rather, the mechanism is fear: fear amplified by policy complexity, transmitted through immigrant networks, and facilitated by an administrative environment that scrutinizes immigrant applicants more intensively than native-born citizens.

In this paper's preferred design, low-income U.S. citizen children in mixed-status families experienced a 2.0 percentage point decline in Medicaid enrollment after the public charge rule's announcement, an effect that persisted through the rule's reversal and deepened during the Medicaid unwinding. California provides supporting evidence that chilling fears spilled over to naturalized citizens – a group with no direct legal vulnerability to the rule – and that immigrant groups face a markedly more burdensome documentation environment when seeking health coverage. The California adult results are not as decisive as the national ACS child analysis, but they clarify the spillover and administrative channels through which the rule likely operated.

These findings carry implications for both immigration and health policy. For immigration policy, they show that the collateral effects of restrictive rules on health coverage are substantial, measurable, and persistent. For health policy, they show that enrollment infrastructure and administrative design can either amplify or attenuate the chilling effects of policy shocks. For both domains, they underscore a fundamental asymmetry: it is easier to generate fear than to undo it.

The public charge rule has been formally reversed, but its effects have not been fully erased. Any future reinstatement of similar restrictions should be evaluated not only as immigration policy but also as child health policy and Medicaid policy – because that is what, in practice, it turned out to be.

References

Acevedo-Garcia, D., Joshi, P. K., Ruskin, E., Walters, A. N., & Sofer, N. (2021). Restoring an inclusionary safety net for children in immigrant families: A review of three social policies. *Health Affairs*, 40(7), 1099–1107.

Alsan, M., & Yang, C. S. (2024). Fear and the safety net: Evidence from Secure Communities. *Review of Economics and Statistics*, 106(6), 1427–1441.

Artiga, S., Garfield, R., & Damico, A. (2019). Estimated impacts of final public charge inadmissibility rule on immigrants and Medicaid coverage. *KFF Issue Brief*.

Artiga, S., Tolbert, J., & Rudowitz, R. (2019). Impact of shifting immigration policy on Medicaid enrollment and utilization of care among health center patients. *KFF Issue Brief*.

Barofsky, J., Vargas, A., Rodriguez, D., & Barrows, A. (2020). Spreading fear: The announcement of the public charge rule reduced enrollment in child safety-net programs. *Health Affairs*, 39(10), 1752–1761.

Bernstein, H., McTarnaghan, S., & Gonzalez, D. (2019). Safety net access in the context of the public charge rule. *Urban Institute Report*.

Bhargava, S., & Manoli, D. (2015). Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment. *American Economic Review*, 105(11), 3489–3529.

Bustamante, A. V., Felix-Beltran, L., Nwadiuko, J., & Ortega, A. N. (2022). Avoiding Medicaid enrollment after the reversal of the changes in the public charge rule among Latino and Asian immigrants. *Health Services Research*, 57(Suppl 2), 195–203.

Capps, R., Greenberg, M., Fix, M., & Zong, J. (2018). Gauging the impact of DHS' proposed public-charge rule on U.S. immigration. *Migration Policy Institute Report*.

Cohodes, S. R., Grossman, D. S., Kleiner, S. A., & Lovenheim, M. F. (2016). The effect of child health insurance access on schooling: Evidence from public insurance expansions. *Journal of Human Resources*, 51(3), 727–759.

Currie, J. (2006). The take-up of social benefits. In A. Auerbach, D. Card, & J. Quigley (Eds.), *Public Policy and the Income Distribution* (pp. 80–148). Russell Sage Foundation.

Currie, J., & Gruber, J. (1996). Health insurance eligibility, utilization of medical care, and child health. *Quarterly Journal of Economics*, 111(2), 431–466.

Kravitz, C., Auchincloss, A. H., Chaparro, M. P., Argibay, S., Eastus, A., & Langellier, B. A. (2024). ICE detainer requests were associated with lower Medicaid and SNAP enrollment among eligible adults, 2011–19. *Health Affairs*, 43(9), 1244–1253.

Fix, M. E., & Capps, R. (2020). Anticipated "chilling effects" of the public-charge rule are real: Census data reflect steep decline in benefits use by immigrant families. *Migration Policy Institute Commentary*.

Finkelstein, A., & Notowidigdo, M. J. (2019). Take-up and targeting: Experimental evidence from SNAP. *Quarterly Journal of Economics*, 134(3), 1505-1556.

FWD.me. (2024). Immigration reform can keep millions of mixed-status families together. *FWD.me Blog*.

Choi, S. W., Agbese, E., Kim, G., Makhlouf, M. D., & Leslie, D. L. (2023). Uninsured immigrants in the United States significantly delayed the initiation of prenatal care after the changes to the public charge rule. *Public Health*, 225, 1–7.

Gangopadhyaya, A., & Braga, B. (2022). Immigrant families faced multiple barriers to safety net programs in 2021. *Urban Institute Report*.

Giuntella, O., & Lonsky, J. (2020). The effects of DACA on health insurance, access to care, and health outcomes. *Journal of Health Economics*, 72, 102320.

Gonzalez, D., Bernstein, H., & Guelespe, D. (2025). Immigrant families express worry as they prepare for policy changes. *Urban Institute Brief*.

Gonzalez, D., Bernstein, H., & Karpman, M. (2020). Amid confusion over the public charge rule, immigrant families continued avoiding public benefits in 2019. *Urban Institute Report*.

Hacker, J. S. (2004). Privatizing risk without privatizing the welfare state: The hidden politics of social policy retrenchment in the United States. *American Political Science Review*, 98(2), 243–260.

Hacker, J. S., & Pierson, P. (2019). Policy feedback in an age of polarization. *Annals of the American Academy of Political and Social Science*, 685(1), 8–28.

Haro-Ramos, A. Y., Axeen, S., Gorman, A., Schneberk, T., & Ro, A. (2026). Self-pay emergency department visits by undocumented patients after 2018 public charge announcement. *JAMA Network Open*, 9(1), e2555081.

Heinrich, C. J. (2016). The bite of administrative burden: A theoretical and empirical investigation. *Journal of Public Administration Research and Theory*, 26(3), 403–420.

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

Kandula, N. R., Grogan, C. M., Rathouz, P. J., & Lauderdale, D. S. (2004). The unintended impact of welfare reform on the Medicaid enrollment of eligible immigrants. *Health Services Research*, 39(5), 1509–1526.

Kaushal, N., & Kaestner, R. (2005). Welfare reform and health insurance of immigrants. *Health Services Research*, 40(3), 697–722.

Tolbert, J., & Corallo, B. (2024). An examination of Medicaid renewal outcomes and enrollment changes at the end of the unwinding. *KFF Issue Brief*.

Miller, S., & Wherry, L. R. (2019). The long-term effects of early life Medicaid coverage. *Journal of Human Resources*, 54(3), 785–824.

Miller, S., Wherry, L. R., & Aldana, G. (2022). Covering undocumented immigrants: The effects of a large-scale prenatal care intervention. NBER Working Paper 30299.

Moffitt, R. (1983). An economic model of welfare stigma. *American Economic Review*, 73(5), 1023–1035.

Moynihan, D. P., Herd, P., & Harvey, H. (2015). Administrative burden: Learning, psychological, and compliance costs in citizen-state interactions. *Journal of Public Administration Research and Theory*, 25(1), 43–69.

Ortega, A. N., McKenna, R., Pintor, J. K., et al. (2018). Health care access and physical and behavioral health among undocumented Latinos in California. *Medical Care*, 56(11), 919–926.

Perreira, K. M., & Pedroza, J. M. (2019). Policies of exclusion: Implications for the health of immigrants and their children. *Annual Review of Public Health*, 40, 147–166.

Pourat, N., Wallace, S. P., Hadler, M. W., & Ponce, N. (2014). Assessing health care services used by California's undocumented immigrant population in 2010. *Health Affairs*, 33(5), 840–847.

Twersky, S. E. (2022). Do state laws reduce uptake of Medicaid/CHIP by U.S. citizen children in immigrant families: Evaluating evidence for a chilling effect. *International Journal for Equity in Health*, 21, 50.

Vargas, E. D. (2015). Immigration enforcement and mixed-status families: The effects of risk of deportation on Medicaid use. *Children and Youth Services Review*, 57, 83–89.

Watson, T. (2014). Inside the refrigerator: Immigration enforcement and chilling effects in Medicaid participation. *American Economic Journal: Economic Policy*, 6(3), 313–338.

Wherry, L. R., & Meyer, B. D. (2016). Saving teens: Using a policy discontinuity to estimate the effects of Medicaid eligibility. *Journal of Human Resources*, 51(3), 556–588.

Tables and Figures

Table 1: ACS Descriptive Statistics, Below 200% FPL

Variable	Mixed-Status	Citizen-Only	Std. Diff.
N	220,387	1,604,764	—
Age (mean)	8.24	9.04	-0.15
Female	0.489	0.489	0.00
Below 200% FPL	0.583	0.335	0.51
Medicaid (pre-period)	0.602	0.337	0.55
Uninsured (pre-period)	0.071	0.037	0.15
Any coverage (pre-period)	0.929	0.963	-0.15

Notes: Pre-period (2016-2018) weighted means. Standardized differences are reported to characterize level differences between treatment and control groups.

Table 2: CHIS Descriptive Statistics by Citizenship Group

Variable	U.S.-Born	Naturalized	Noncitizens
N	168,525	35,656	16,744
Medi-Cal	0.224	0.254	0.057
ED visit	0.205	0.169	0.169
Usual source	0.857	0.862	0.715
Forgone care	0.103	0.064	0.079
Asked SSN/citizenship	0.000	0.073	0.117
Age (mean)	46.5	53.3	42.2
Female	0.511	0.524	0.497
Below 200% FPL	0.257	0.363	0.570

Notes: CHIS 2015-2024 pooled public-use files.

Table 3: Main ACS Estimates

Sample	Outcome	Beta	SE	p	N
Below 200% FPL	Medicaid	**-0.020**	**0.005**	**<0.001**	**1,622,582**
Below 200% FPL	Any public coverage	**-0.020**	**0.004**	**<0.001**	**1,622,582**
Below 200% FPL	Uninsured	+0.008	0.005	0.113	1,622,582
Full sample	Medicaid	-0.005	0.005	0.340	5,256,322
Full sample	Uninsured	+0.008	0.003	0.012	5,256,322
Full sample	Any coverage	-0.008	0.003	0.012	5,256,322

Notes: The below-200% FPL Medicaid specification is the paper's preferred ACS design. Full-sample estimates are secondary and should be interpreted with caution because the uninsured event study fails the joint pre-trend test.

Table 4: Preferred CHIS Multi-Group Difference-in-Differences

Outcome	Noncitizen x Post	p	Naturalized x Post	p	N
Medi-Cal	-0.000	0.979	**-0.013**	**0.030**	220,925
ED visit	+0.055	<0.001	+0.050	<0.001	220,925
Usual source	+0.054	<0.001	+0.007	0.368	220,925
Forgone care	-0.004	0.616	-0.012	0.027	220,925

Notes: Reference group is U.S.-born citizens. Robust standard errors are the primary CHIS inference. The naturalized-citizen Medi-Cal estimate is the paper's preferred California spillover result.

Table 5: Administrative Vulnerability, CHIS Documentation Demands (Immigrant-Only Universe)

Specification	Estimate	SE	p	N
(1) Noncitizen vs. naturalized, year FE only	**+0.092**	**0.008**	**<0.001**	33,464
(2) + age, sex, below-200% FPL	**+0.055**	**0.008**	**<10-^{1}0**	33,464
(3) + race/ethnicity dummies	**+0.057**	**0.009**	**<10-^{1}0**	33,427
Reference-group mean (naturalized)	0.105	—	—	24,085

Notes: CHIS public-use adult sample, 2019–2024, restricted to respondents not born in the United States — the survey universe for the AL19V2 documentation-demand item. Each row is a weighted linear-probability regression of being asked for a Social Security number or citizenship documentation when seeking coverage on a noncitizen indicator (naturalized = reference), with year fixed effects and the listed controls. Heteroskedasticity-robust standard errors. The earlier inclusive-universe specification compared noncitizens and naturalized citizens against a U.S.-born baseline of structural zeros (the question is filtered to immigrants); that comparison mechanically inflates the gradient and is not reported here. The corrected within-immigrant-universe gradient is 5–9 percentage points depending on covariate set: noncitizens face that much more documentation scrutiny than naturalized citizens at the same enrollment-system entry points.

Figure 1: Raw Medicaid Enrollment Trends, Below 200% FPL

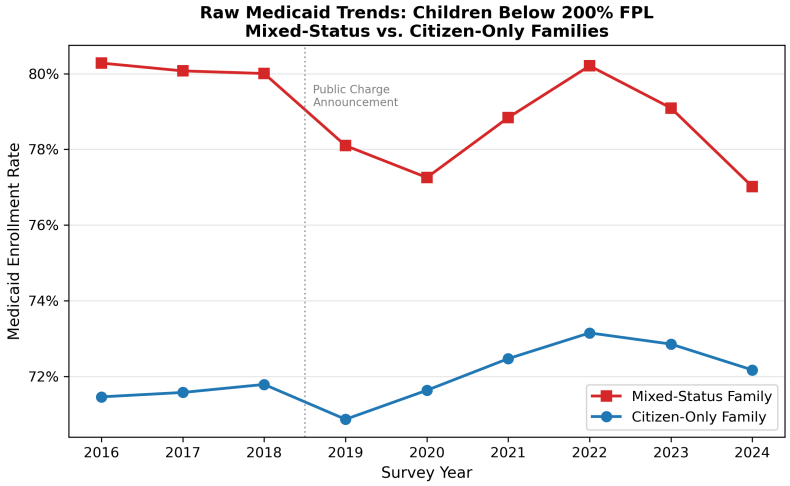

Figure 10: Raw Medicaid Trends Below200Fpl

Note: This figure shows raw trends for the raw Medicaid Trends Below200Fpl. It helps readers compare baseline levels, pre-policy movement, and the timing of any post-policy divergence.

Weighted mean Medicaid enrollment rates for U.S. citizen children in mixed-status versus citizen-only families, below 200% FPL. Pre-announcement trends are parallel. A visible gap opens after the September 2018 announcement and persists through 2024.

Figure 2: Event Study, Medicaid Enrollment, Below 200% FPL

Lead coefficients are small and jointly insignificant ($p = 0.92$). The post-period path becomes negative after the announcement, deepens around implementation, and remains negative through 2024.

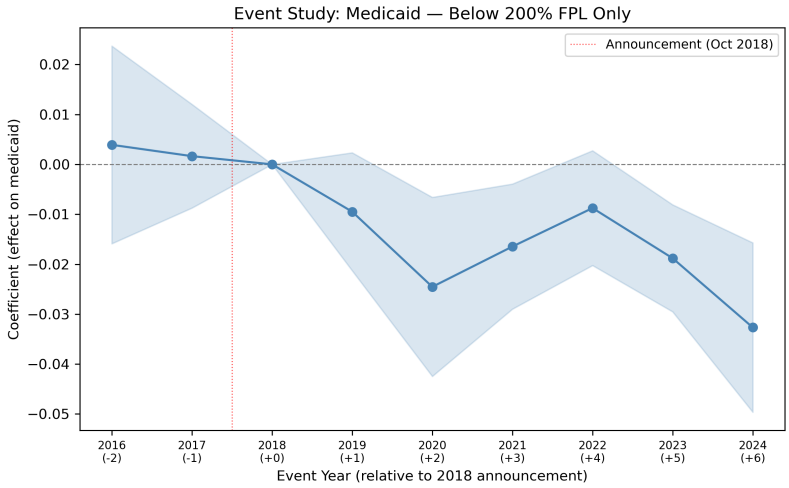

Figure 11: Event Study Medicaid Below200Fpl

Note: This figure plots event-time estimates for the event Study Medicaid Below200Fpl. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

Figure 3: CHIS Event-Study Panels

Four-panel event studies for Medi-Cal enrollment, ED visits, usual source of care, and forgone care. Joint Wald tests of pre-period coefficients are non-significant for all four outcomes.

Appendix

Appendix A1: ACS Full-Sample Secondary Results

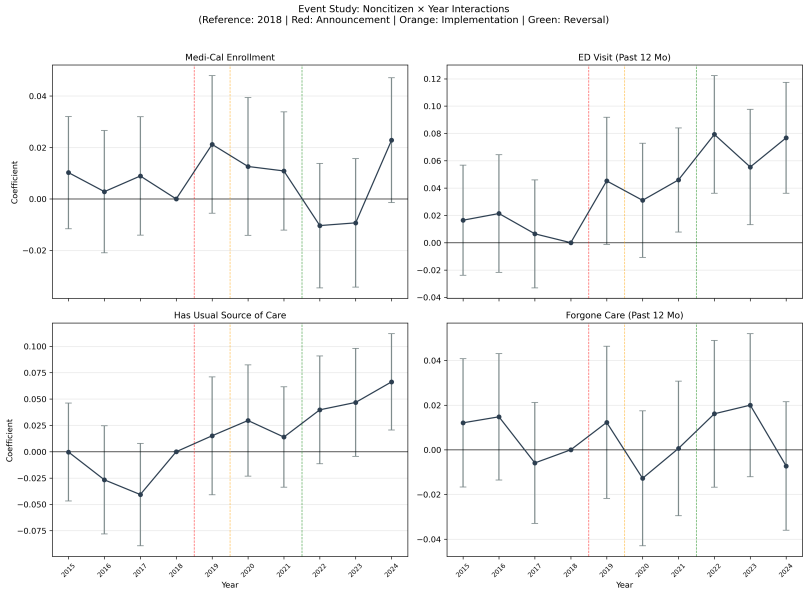

Figure 12: Event Study Combined

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

Outcome	Beta	SE	p	N
Medicaid	-0.005	0.005	0.340	5,256,322
Uninsured	+0.008	0.003	0.012	5,256,322
Any coverage	-0.008	0.003	0.012	5,256,322
Triple-difference Medicaid	-0.003	0.007	0.703	5,256,322

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.

Appendix A2: Joint Wald Tests of Pre-Period Event-Study Coefficients

Analysis	Outcome	Chi-square	p
ACS below 200% FPL	Medicaid	0.16	0.92
ACS full sample	Uninsured	7.41	0.025
CHIS	Medi-Cal	1.11	0.77
CHIS	ED visit	1.20	0.75
CHIS	Usual source	4.13	0.25
CHIS	Forgone care	3.34	0.34

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.

Appendix A3: Unwinding Decomposition, Below 200% FPL

Period	Outcome	Beta	SE	p	N
Full post-period (2019-2024)	Medicaid	**-0.020**	**0.005**	**<0.001**	**1,622,582**
Pre-unwinding (2019-2022)	Medicaid	-0.017	0.005	0.001	1,622,582
Unwinding period (2023-2024)	Medicaid	-0.028	0.008	<0.001	1,622,582

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.

Appendix A4: County-Enforcement Heterogeneity DiD — Restricted-CHIS Adults

Outcome	Coefficient	Beta	SE	p
Medi-Cal	Naturalized × Post (low-enf)	+0.006	0.012	0.629
	Green-card × Post (low-enf)	+0.039	0.025	0.124
	No-green-card × Post (low-enf)	+0.125	0.022	<0.001
	High-enf × Naturalized × Post	–0.002	0.008	0.785
	High-enf × Green-card × Post	–0.032	0.018	0.074
	High-enf × No-green-card × Post	**–0.045**	**0.019**	**0.021**
Usual source of care	Naturalized × Post (low-enf)	+0.001	0.011	0.930
	Green-card × Post (low-enf)	+0.076	0.024	0.002
	No-green-card × Post (low-enf)	+0.130	0.024	<0.001
	High-enf × Naturalized × Post	+0.009	0.009	0.331
	High-enf × Green-card × Post	–0.060	0.017	<0.001
	High-enf × No-green-card × Post	**–0.084**	**0.019**	**<0.001**
ED visit (past year)	Naturalized × Post (low-enf)	+0.052	0.011	<0.001
	Green-card × Post (low-enf)	+0.066	0.022	0.004
	No-green-card × Post (low-enf)	+0.077	0.020	<0.001
	High-enf × Naturalized × Post	–0.004	0.008	0.601
	High-enf × Green-card × Post	–0.021	0.014	0.135

Outcome	Coefficient	Beta	SE	p
	High-enf × No-green-card × Post	**–0.032**	**0.016**	**0.049**
Forgone care	Naturalized × Post (low-enf)	–0.012	0.008	0.138
	Green-card × Post (low-enf)	–0.007	0.011	0.533
	No-green-card × Post (low-enf)	–0.001	0.017	0.933
	High-enf × Naturalized × Post	+0.001	0.006	0.927
	High-enf × Green-card × Post	–0.016	0.008	0.058
	High-enf × No-green-card × Post	+0.019	0.013	0.162

Notes: Restricted-CHIS adult sample, 2015–2024 (N = 220,925). U.S.-born citizens omitted. Low-enforcement baseline corresponds to the San Francisco AOR (48 California counties); the high-enforcement extra coefficients identify the differential post-period response in the Los Angeles + San Diego AORs (10 counties). High-enforcement is an AOR-level above-median per-capita arrest indicator constructed from the Deportation Data Project (CC-0). 2015–2021 county-year values are back-filled with the AOR's 2022 value because public AOR-level enforcement data begin in fiscal year 2022; substantive interpretation is therefore a Southern-versus-Northern California enforcement-intensity contrast on the level of the post-2018 response, not a time-varying enforcement-shock instrument. Variance estimated via 80 CHIS replicate jackknife weights. Output table: analysis/tables/exhibit_chis_county_enforcement_did.csv and data/dac_submission/DAC260336748_20260424_third_ticket_output/adult_county_heterogeneity_doc4.csv.

Supplementary Appendix

Fear as a Barrier to Coverage: The Chilling Effect of the Public Charge Rule on Children's Medicaid Enrollment

---

Appendix Exhibit A1: Descriptive Statistics — ACS Sample (Below 200% FPL)

Variable	Mixed-Status	Citizen-Only	Std. Diff.
N	220,387	1,604,764	—
Age (mean)	8.24	9.04	–0.15
Female	0.489	0.489	0.00
Below 200% FPL	0.583	0.335	0.51
Medicaid (pre-period)	0.602	0.337	0.55
Uninsured (pre-period)	0.071	0.037	0.15
Any coverage (pre-period)	0.929	0.963	–0.15

Notes: Pre-period (2016–2018) weighted means. Standardized differences > |0.10| flagged.

Appendix Exhibit A2: Descriptive Statistics — CHIS Sample

Variable	US-Born	Naturalized	Noncitizens
N	168,525	35,656	16,744
Medi-Cal	0.224	0.254	0.057
ED visit	0.205	0.169	0.169
Usual source	0.857	0.862	0.715
Forgone care	0.103	0.064	0.079
Asked SSN/cit	0.000	0.073	0.117
Age (mean)	46.5	53.3	42.2
Female	0.511	0.524	0.497
Below 200% FPL	0.257	0.363	0.570

Notes: This table documents the source files, scripts, variables, or data inputs used in the analysis. It is included to make the construction of the analytic evidence reproducible.

Appendix Exhibit A3: ACS Full-Sample Results (All Income Levels)

Specification	Outcome	β	SE	p	N
Basic DiD	Medicaid	–0.005	0.005	0.34	5,256,322
Basic DiD	Uninsured	+0.008	0.003	0.012	5,256,322
Basic DiD	Any coverage	–0.008	0.003	0.012	5,256,322
Triple-DiD	Medicaid	–0.003	0.007	0.70	5,256,322
Drop COVID (20–21)	Medicaid	–0.003	0.006	0.62	4,203,852
COVID interaction	Medicaid	–0.003	0.006	0.58	5,256,322

Notes: Full-sample results presented as secondary specifications. The attenuated Medicaid effect (–0.5 pp, p = 0.34) is consistent with dilution from higher-income families for whom Medicaid is not the relevant coverage margin. The significant uninsured result (+0.8 pp, p = 0.012) should be interpreted with caution: the joint Wald test rejects the null for uninsured pre-period coefficients (p = 0.025; see Exhibit A5), though the pre-trend direction (positive 2016 coefficient) biases estimates toward zero rather than away.

Appendix Exhibit A4: Event-Study Coefficients — Medicaid, Full Sample

Year	β	SE	p
2016 (–2)	+0.007	0.008	0.34
2017 (–1)	+0.005	0.005	0.33
2018 (ref)	0	—	—
2019 (+1)	–0.002	0.005	0.66
2020 (+2)	–0.011	0.004	0.016
2021 (+3)	+0.003	0.007	0.63
2022 (+4)	+0.008	0.004	0.027
2023 (+5)	+0.004	0.003	0.22
2024 (+6)	–0.006	0.007	0.38

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.

Appendix Exhibit A5: Joint Wald Tests of Pre-Period Event-Study Coefficients

Analysis	Outcome	N pre-periods	Chi-square	p
ACS (below 200% FPL)	Medicaid	2	0.16	0.92
ACS (full sample)	Uninsured	2	7.41	0.025
CHIS	Medi-Cal	3	1.11	0.77
CHIS	ED visit	3	1.20	0.75
CHIS	Usual source	3	4.13	0.25
CHIS	Forgone care	3	3.34	0.34

Notes: Joint Wald tests use the estimated covariance matrix of the pre-period event-study coefficients rather than an approximation that ignores covariance across leads. The primary specification (below 200% FPL, Medicaid) shows no evidence of pre-policy divergence (p = 0.92). The full-sample uninsured pre-trend failure (p = 0.025) is driven by a positive 2016 coefficient (opposite direction from the hypothesized treatment effect), which biases the DiD estimate toward zero. All CHIS outcomes pass the joint Wald test. The ACS event study has only two pre-treatment periods (2016, 2017), limiting the power of parallel trends testing; this is acknowledged as a limitation.

Appendix Exhibit A6: Placebo Tests — ACS

Test	Outcome	β	SE	p	N
Citizen-only × High NC × Post	Medicaid	–0.004	0.005	0.42	4,655,075
Citizen-only × High NC × Post	Uninsured	+0.002	0.002	0.39	4,655,075
Citizen-only × High NC × Post	Any coverage	–0.002	0.002	0.39	4,655,075

Notes: Placebo test using citizen-only families comparing high vs. low noncitizen-share states. All null, confirming the treatment effect is specific to mixed-status families.

Appendix Exhibit A7: COVID-19 Sensitivity — CHIS

Outcome	Specification	β	SE	p	N
Medi-Cal	Drop COVID	–0.003	0.006	0.61	146,774
ED visit	Drop COVID	+0.064	0.011	<0.001	146,774
Medi-Cal	COVID interaction	–0.003	0.006	0.58	185,269
ED visit	COVID interaction	+0.064	0.011	<0.001	185,269

Notes: Secondary two-group CHIS comparisons restricted to noncitizens versus U.S.-born citizens. The preferred California specification in the main text is the multi-group model reported in Exhibit 3.

Appendix Exhibit A8: BRR Variance Estimation — CHIS

Outcome	Coefficient	β	BRR SE	Robust SE	BRR p
Medi-Cal	Noncitizen × Post	–0.000	0.001	0.006	0.85
Medi-Cal	Naturalized × Post	–0.013	0.001	0.006	<0.001
ED visit	Noncitizen × Post	+0.055	0.001	0.010	<0.001
ED visit	Naturalized × Post	+0.050	0.001	0.008	<0.001
Usual source	Noncitizen × Post	+0.054	0.002	0.012	<0.001
Usual source	Naturalized × Post	+0.007	0.001	0.008	<0.001
Forgone care	Noncitizen × Post	–0.004	0.001	0.007	<0.001
Forgone care	Naturalized × Post	–0.012	0.001	0.005	<0.001

Notes: BRR SE from 80 CHIS replicate weights for the preferred multi-group CHIS model in Exhibit 3. The BRR SEs are substantially smaller than heteroskedasticity-robust SEs because BRR captures survey sampling variance only, while the robust SEs also absorb misspecification and heteroskedasticity concerns. I therefore retain the robust SEs as the primary CHIS inference.

Appendix Exhibit A9: Alternative Treatment Timing (Post = 2020, Implementation)

Analysis	Outcome	β	SE	p	N
ACS	Medicaid	–0.004	0.003	0.21	5,256,322
ACS	Uninsured	+0.005	0.002	0.035	5,256,322
CHIS	Medi-Cal	–0.004	0.006	0.48	185,269
CHIS	ED visit	+0.050	0.010	<0.001	185,269

Notes: CHIS rows report the secondary two-group noncitizen-versus-U.S.-born comparison. The preferred California specification in the main text is the multi-group model reported in Exhibit 3.

Appendix Exhibit A10: Heterogeneity by Poverty Level — CHIS

Subgroup	Outcome	β	SE	p	N
Below 200% FPL	Medi-Cal	–0.001	0.009	0.87	—
Above 200% FPL	Medi-Cal	–0.013	0.008	0.11	—
Below 200% FPL	ED visit	+0.065	0.014	<0.001	—
Above 200% FPL	ED visit	+0.048	0.014	<0.001	—

Notes: Secondary two-group CHIS comparisons restricted to noncitizens versus U.S.-born citizens. The preferred California specification in the main text is the multi-group model reported in Exhibit 3.

Appendix Exhibit A11: Unwinding Decomposition — Full Sample

Period	Outcome	β	SE	p	N
Full post (2019–2024)	Medicaid	–0.005	0.005	0.340	5,256,322
Pre-unwinding (2019–2022)	Medicaid	–0.004	0.004	0.282	5,256,322
Unwinding (2023–2024)	Medicaid	–0.005	0.007	0.463	5,256,322
Pre-unwinding (2019–2022)	Uninsured	+0.007	0.005	0.178	1,622,582
Unwinding (2023–2024)	Uninsured	+0.010	0.005	0.069	1,622,582

Notes: Full-sample and uninsured unwinding decomposition. Consistent with the primary specification, effects are attenuated in the full sample. The below-200% FPL uninsured result shows a marginally significant unwinding-period effect (p = 0.069), directionally consistent with the Medicaid findings.

Appendix Exhibit A12: Earlier Treatment Timing Sensitivity

Timing	Sample	Outcome	β	SE	p
Baseline: post=2019+ (NPRM)	**Below 200% FPL**	**Medicaid**	**–0.020**	**0.005**	**<0.001**
Alt A: post=2018+ (Reg Agenda)	Below 200% FPL	Medicaid	–0.018	0.007	0.012
Alt A: post=2018+ (Reg Agenda)	Below 200% FPL	Uninsured	+0.009	0.005	0.085
Alt A: post=2018+ (Reg Agenda)	Full sample	Medicaid	–0.007	0.006	0.296
Alt B: post=2017+ (Leaked EO)	Below 200% FPL	Medicaid	–0.017	0.010	0.092
Alt B: post=2017+ (Leaked EO)	Below 200% FPL	Uninsured	+0.008	0.005	0.146
Alt B: post=2017+ (Leaked EO)	Full sample	Medicaid	–0.007	0.008	0.356

Notes: Sensitivity to alternative treatment timing reflecting the public charge rule's emergence before the October 2018 NPRM. A leaked draft executive order on public benefits was reported in January 2017 (Washington Post); DHS placed the rule on the Unified Regulatory Agenda in October 2017. Alt A (post=2018+) remains significant for below-200% FPL Medicaid, suggesting anticipation effects. Alt B (post=2017+) is directionally consistent but attenuated and marginally significant, consistent with gradual onset. The event study centered on 2016 (Figure A11) shows a flat 2017 coefficient, with effects emerging in 2018 and deepening from 2020.

Appendix Exhibit A13: Event-Study Coefficients — Medicaid, Below 200% FPL (Reference: 2016)

Year	Event Year	β	SE	p
2016 (ref)	0	0	—	—
2017	+1	–0.002	0.008	0.78
2018	+2	–0.004	0.010	0.70
2019	+3	–0.013	0.012	0.26
2020	+4	–0.028	0.016	0.075
2021	+5	–0.020	0.007	0.007
2022	+6	–0.013	0.007	0.070
2023	+7	–0.023	0.009	0.012
2024	+8	–0.037	0.017	0.033

Notes: Event study with 2016 as reference year (earliest available) to examine whether effects began before the NPRM. The 2017 coefficient is near zero (–0.2 pp, p = 0.78), suggesting no measurable chilling from the leaked EO at the population level. Effects emerge gradually from 2018 and deepen sharply from 2020. This trajectory is consistent with chilling intensifying as the policy moved from leaked draft to formal rulemaking to implementation.

Appendix Exhibit A14: Pairwise CHIS Event-Study Audit

Panel A. Sample support and saturation

Comparison	Treated N	Control N	Missingness in modeled variables	Empty year-by-group cells	Fully saturated event study
Noncitizens vs. U.S.-born citizens	16,744	168,525	No	No	Yes
Naturalized vs. U.S.-born citizens	35,656	168,525	No	No	Yes

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

Panel B. Medi-Cal pairwise results

Comparison	Pairwise DiD β	SE	p	Joint pre-trend p	Average post coefficient
Noncitizens vs. U.S.-born citizens	–0.001	0.006	0.789	0.946	–0.000
Naturalized vs. U.S.-born citizens	–0.013	0.006	0.030	0.171	–0.030

Notes: Committee-requested pairwise CHIS audit. Panel A summarizes the missingness and support checks using /analysis/tables/comment48_missingness_audit.csv and /analysis/tables/comment48_cell_support.csv. Panel B reports the pairwise DiD and event-study summary for Medi-Cal from /analysis/tables/comment48_pairwise_did.csv and /analysis/tables/comment48_event_study_summary.csv. Both pairwise comparisons estimate all expected year interactions. The naturalized-versus-U.S.-born event study does not reject the joint pre-trend test, but one lead coefficient is modestly negative, so I treat that California result as supporting rather than definitive evidence.

Appendix Exhibit A15: Pairwise CHIS Medi-Cal Event Studies

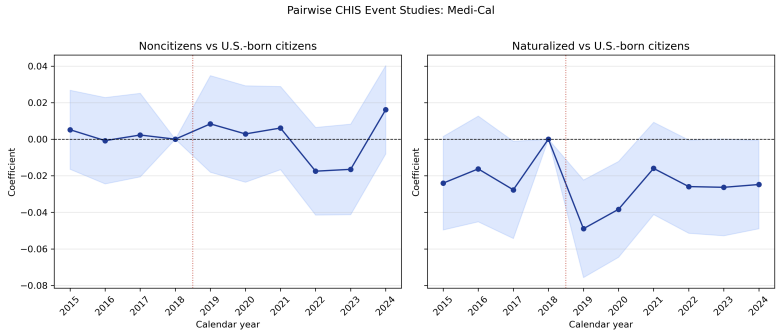

Figure 13: Pairwise CHIS event studies

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

Notes: Pairwise event-study estimates comparing noncitizens with U.S.-born citizens and naturalized citizens with U.S.-born citizens. Reference year is 2018. The noncitizen series remains close to zero throughout. The naturalized series is negative in most post-period years, but one pre-period coefficient is also modestly negative, reinforcing a cautious interpretation of the California spillover result.

Appendix Exhibit A16: Triple-Difference Calibration Check — High Noncitizen-Share States

Coefficient	β	SE	p
Mixed-status × High NC share	+0.057	0.014	<0.001
High NC share × Post	–0.004	0.005	0.437
Mixed-status × High NC share × Post	–0.003	0.007	0.703

Notes: Triple-difference calibration check using the full ACS child sample. The null triple interaction shows that the preferred result is not strengthened by additionally interacting exposure with high state-year noncitizen share. I therefore treat this as a weak secondary specification rather than a core identification result.

Appendix Exhibit A17: ACS State-by-Year Fixed-Effects Event-Study Summary

Analysis	Average post coefficient	SE	p	Joint pre-trend p	N
Medicaid, below 200% FPL	–0.019	0.004	<0.001	0.922	1,622,582

Notes: Separate ACS event-study sensitivity with state-by-year fixed effects, estimated on the preferred below-200% FPL sample. This specification absorbs any state-specific shocks common to mixed-status and citizen-only children within the same year. The average post coefficient is the mean of the 2019–2024 event-study coefficients from state_year_fe_event_study_below200fpl.csv.

Appendix Exhibit A18: ACS Event Study with State-by-Year Fixed Effects

Notes: State-by-year fixed-effects event-study estimates for Medicaid enrollment among children below 200% FPL. Reference year is 2018. Pre-period coefficients remain near zero, while the post-period path turns negative in 2020 and remains negative through 2024. This supports the interpretation that the main ACS result is not driven by broad state-level shocks that affect both family types equally within a year.

Appendix Exhibit A19: County-Enforcement Heterogeneity DiD — Restricted-CHIS Adults

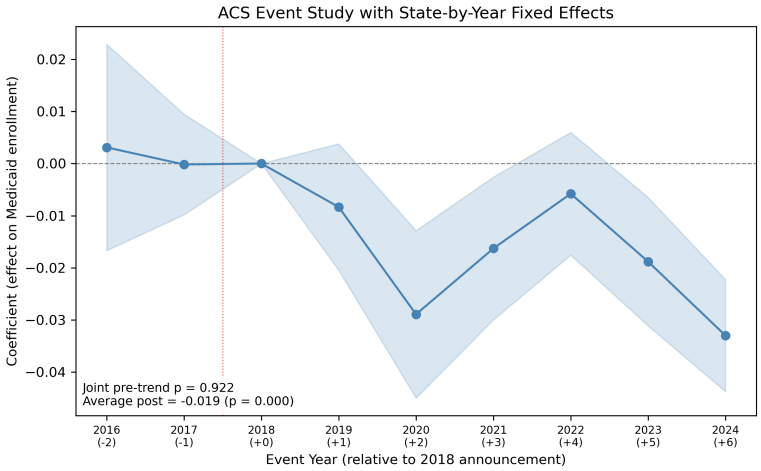

Figure 14: State-by-year FE event study

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

Outcome	Coefficient	β	SE	p
Medi-Cal	Naturalized × Post (low-enf baseline)	+0.006	0.012	0.629
Medi-Cal	Green-card × Post (low-enf baseline)	+0.039	0.025	0.124
Medi-Cal	No-green-card × Post (low-enf baseline)	+0.125	0.022	<0.001
Medi-Cal	High-enf × Naturalized × Post	–0.002	0.008	0.785
Medi-Cal	High-enf × Green-card × Post	–0.032	0.018	0.074
Medi-Cal	**High-enf × No-green-card × Post**	**–0.045**	**0.019**	**0.021**
Usual source	Naturalized × Post (low-enf baseline)	+0.001	0.011	0.930

Outcome	Coefficient	β	SE	p
Usual source	Green-card × Post (low-enf baseline)	+0.076	0.024	0.002
Usual source	No-green-card × Post (low-enf baseline)	+0.130	0.024	<0.001
Usual source	High-enf × Naturalized × Post	+0.009	0.009	0.331
Usual source	High-enf × Green-card × Post	–0.060	0.017	<0.001
Usual source	**High-enf × No-green-card × Post**	**–0.084**	**0.019**	**<0.001**
ED visit	Naturalized × Post (low-enf baseline)	+0.052	0.011	<0.001
ED visit	Green-card × Post (low-enf baseline)	+0.066	0.022	0.004
ED visit	No-green-card × Post (low-enf baseline)	+0.077	0.020	<0.001
ED visit	High-enf × Naturalized × Post	–0.004	0.008	0.601
ED visit	High-enf × Green-card × Post	–0.021	0.014	0.135
ED visit	**High-enf × No-green-card × Post**	**–0.032**	**0.016**	**0.049**
Forgone care	Naturalized × Post (low-enf baseline)	–0.012	0.008	0.138
Forgone care	Green-card × Post (low-enf baseline)	–0.007	0.011	0.533
Forgone care	No-green-card × Post (low-enf baseline)	–0.001	0.017	0.933
Forgone care	High-enf × Naturalized × Post	+0.001	0.006	0.927

Outcome	Coefficient	β	SE	p
Forgone care	High-enf × Green-card × Post	–0.016	0.008	0.058
Forgone care	High-enf × No-green-card × Post	+0.019	0.013	0.162

Notes: Restricted-CHIS adult sample, 2015–2024 (N = 220,925; sample matches all four outcomes). Specification: four-level CITIZEN1 multi-group DiD with high-enforcement-county interactions, controlling for age, sex, race, below-200% FPL, survey-year fixed effects, and CITIZEN1 fixed effects. U.S.-born citizens omitted. The low-enforcement baseline corresponds to the San Francisco AOR (48 California counties); high-enforcement extra coefficients identify the differential post-period response in the Los Angeles + San Diego AORs (10 counties; 100 county-years, ~17% of the merged county-year set; 480 SF-AOR county-years, ~83%). High-enforcement is an AOR-level above-median per-capita arrest indicator constructed from the Deportation Data Project (CC-0; data accessed 2026-04-24). Because public AOR-level DDP data begin in fiscal year 2022, 2015–2021 county-years use the AOR's 2022 value as the earliest publicly observable enforcement-intensity rank, so the indicator is essentially time-invariant within the AOR-pair partition during the 2018–2024 paper window. Substantive interpretation is a Southern-versus-Northern California enforcement-intensity contrast on the level of the post-2018 response, not a within-AOR county-level gradient and not a time-varying enforcement-shock instrument. The within-noncitizen-group post coefficients (low-enforcement baseline) are dominated by California's full-scope Medi-Cal expansions for undocumented adults during the paper window (January 2020 ages 19–25, May 2022 ages 50+, January 2024 ages 26–49); the high-enforcement-extra interactions identify a consistent differential — adults without green cards in higher-enforcement AORs gained less coverage and access during the expansion calendar than adults without green cards in the lower-enforcement AOR — most clearly on the access margin (usual source of care, –8.4 pp, p < 10^{-4}). Variance estimated via 80 CHIS replicate jackknife weights using the design-supplied multiplier. Output files are archived with the analysis tables and DAC third-ticket output.

Appendix Figures

Figure A1. Uninsured Event Study

Figure A2. Any Coverage Event Study

Figure A3. Medicaid Event Study

Figure A4. Placebo Event Study

Figure A5. ACS Covariate Balance

Figure A6. CHIS Covariate Balance

Figure A7. Medicaid Race/Ethnicity Heterogeneity

Figure A8. Uninsured Race/Ethnicity Heterogeneity

Figure A9. CHIS Event-Study Panels

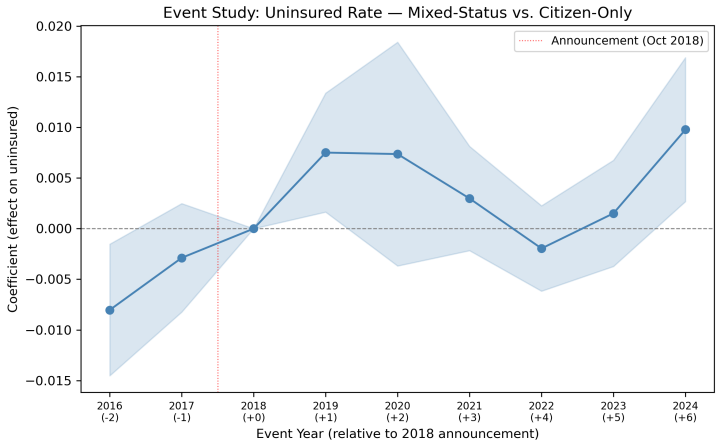

Note: This figure plots event-time estimates for the event study uninsured. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

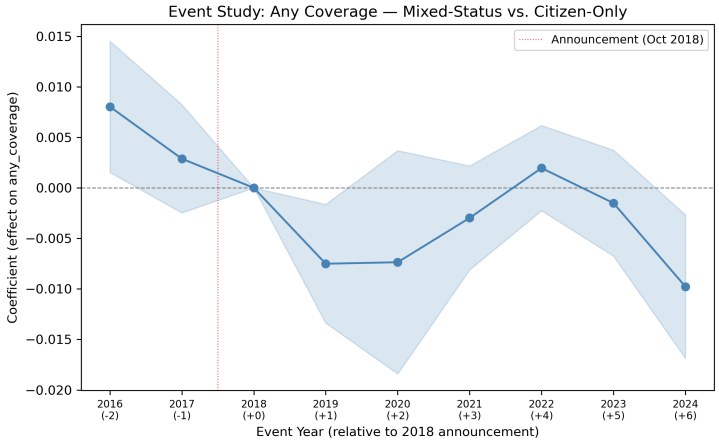

Note: This figure plots event-time estimates for the event study any coverage. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

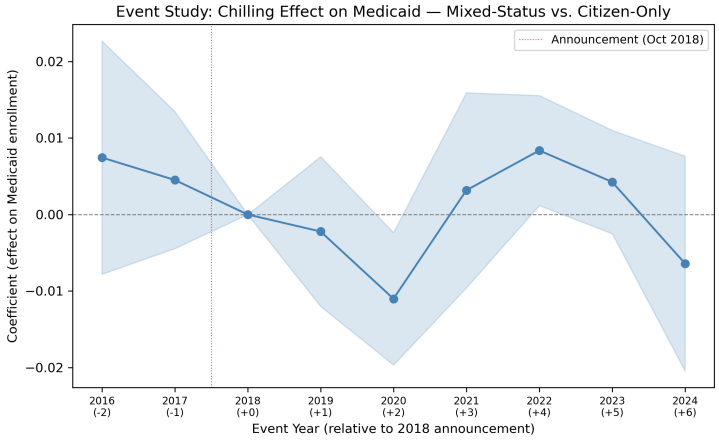

Note: This figure plots event-time estimates for the event study medicaid. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

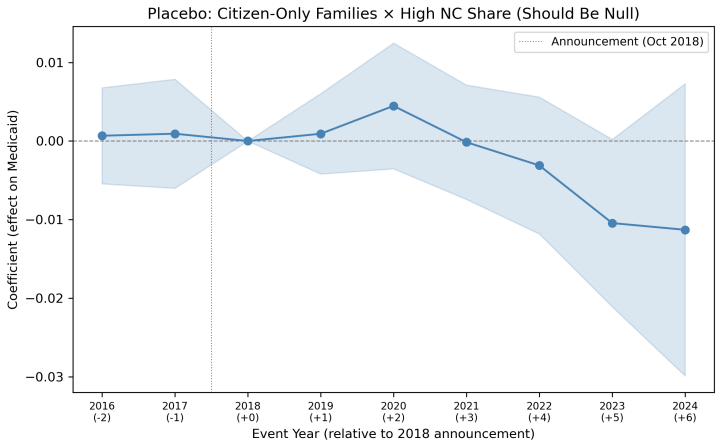

Note: This figure plots event-time estimates for the event study placebo medicaid. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

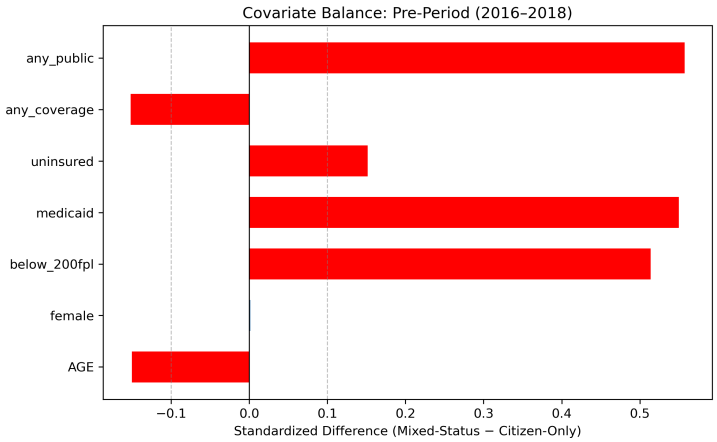

Note: This figure presents the covariate balance. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

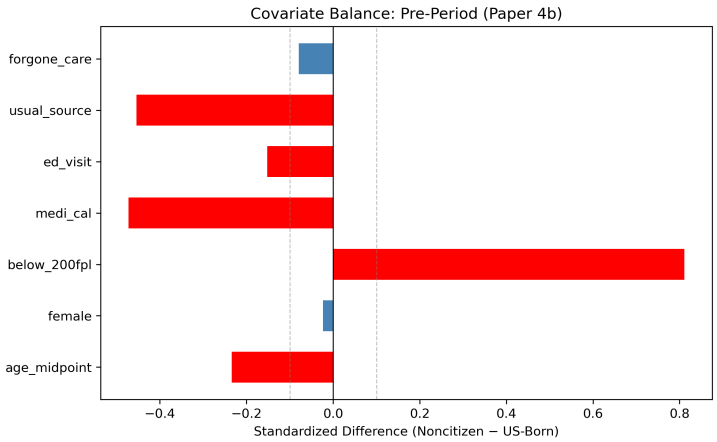

Note: This figure presents the covariate balance 4b. It is included to make the empirical design, sample structure, or headline result easier to read alongside the surrounding text.

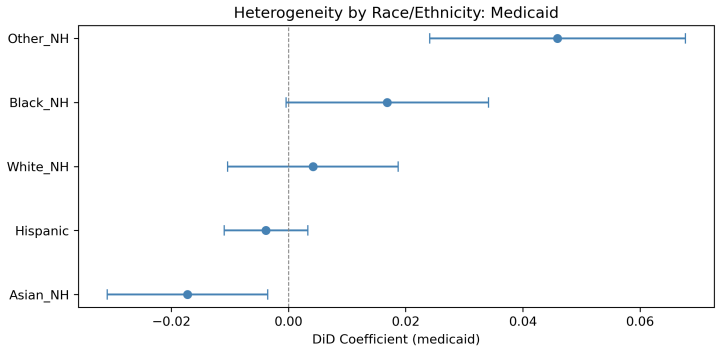

Note: This figure compares estimates across groups or specifications for the heterogeneity race medicaid. It is intended to make effect heterogeneity and subgroup precision easier to assess.

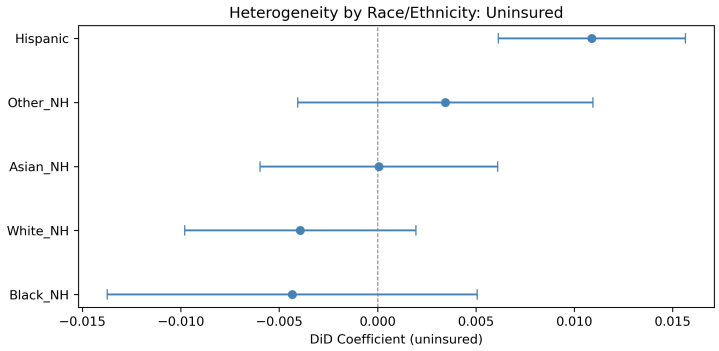

Note: This figure compares estimates across groups or specifications for the heterogeneity race uninsured. It is intended to make effect heterogeneity and subgroup precision easier to assess.

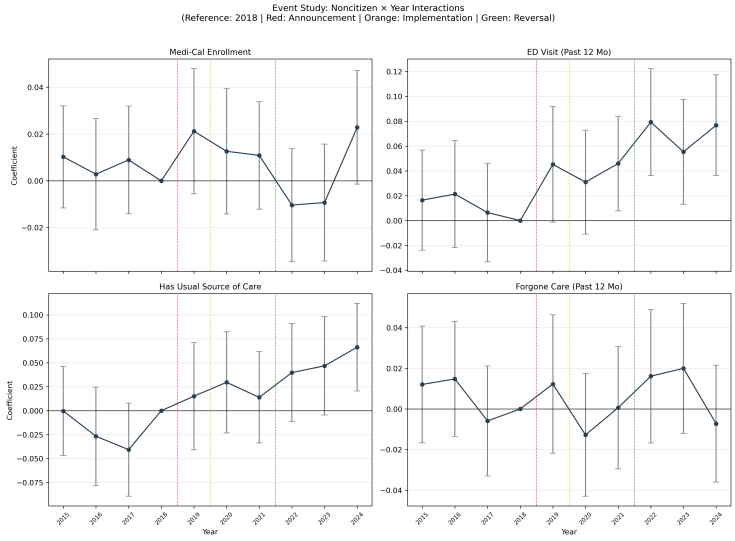

Note: This figure plots event-time estimates for the event study combined. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

Figure A10. Naturalized-Citizen Placebo

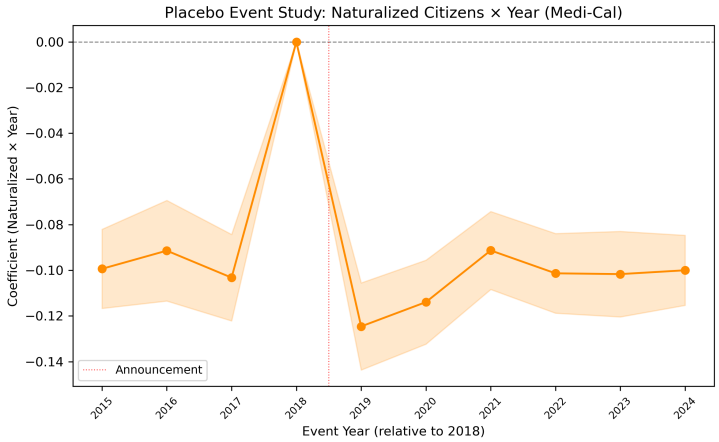

Note: This figure plots event-time estimates for the event study placebo nat medi cal. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

Figure A11. Early-Timing Medicaid Event Study

Figure A12. Raw Outcome Trends

Figure A13. Multi-Outcome Event Study: Levels

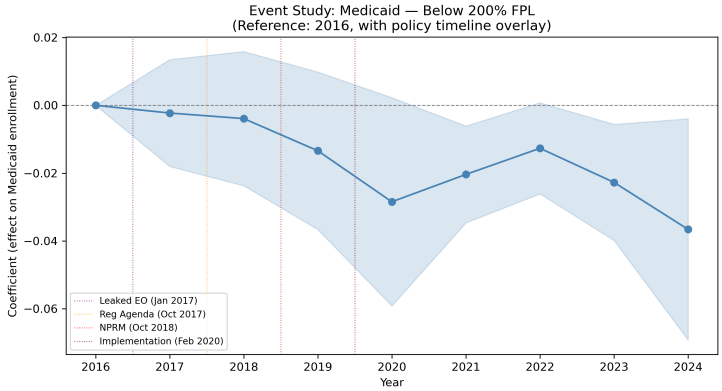

Note: This figure plots event-time estimates for the event study early timing ref2016. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

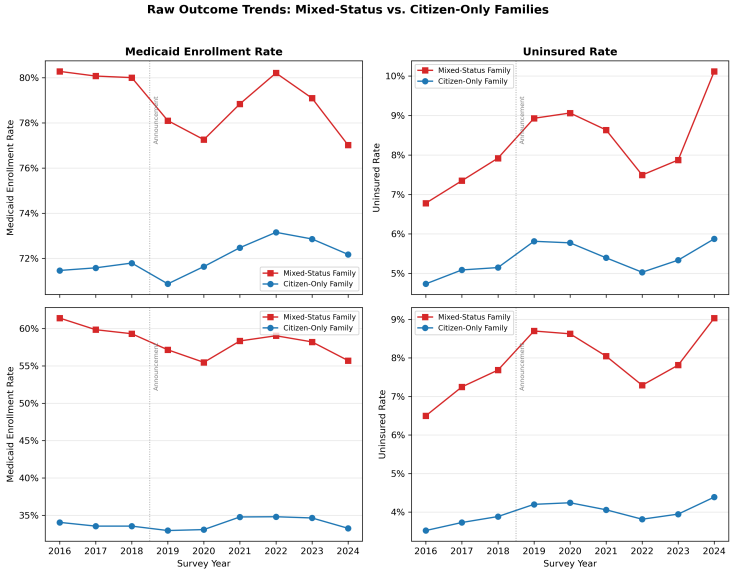

Note: This figure shows raw trends for the raw outcome trends. It helps readers compare baseline levels, pre-policy movement, and the timing of any post-policy divergence.

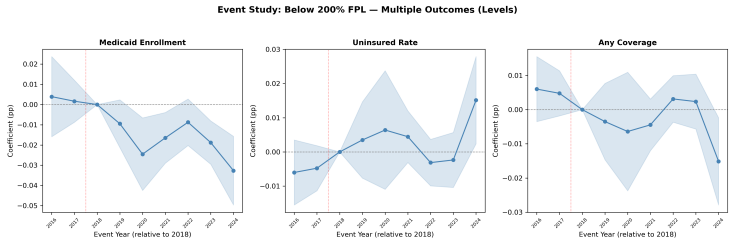

Note: This figure plots event-time estimates for the event study multi outcome levels. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

Figure A14. Multi-Outcome Event Study: Percent Change

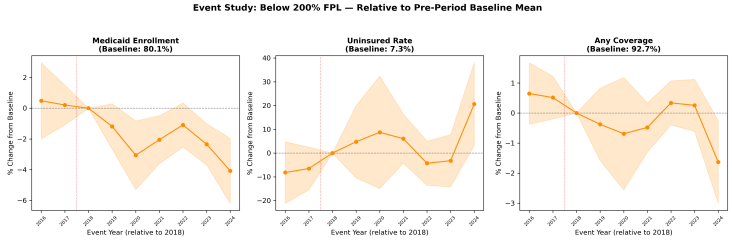

Note: This figure plots event-time estimates for the event study multi outcome pct change. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

Figure A15. Pairwise CHIS Medi-Cal Event Studies

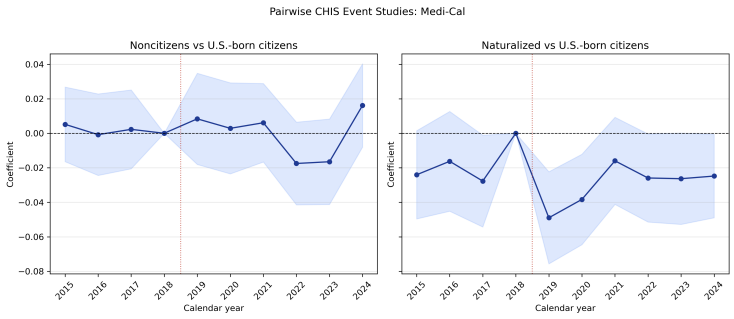

Note: This figure plots event-time estimates for the comment48 pairwise event study medi cal. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.

Figure A16. State-Year FE Medicaid Event Study

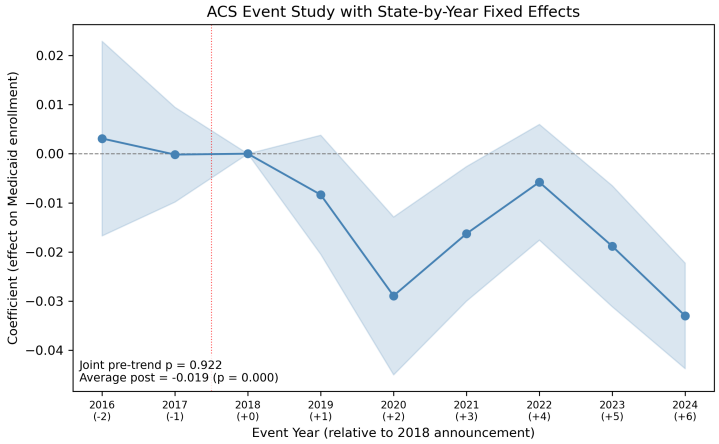

Note: This figure plots event-time estimates for the event study medicaid state year fe. Points show period-specific effects relative to the omitted reference period, with uncertainty intervals where reported.