

When Age Thresholds Fail: A Cautionary Tale About RDD Validity for Life-Course Outcomes

APEP Autonomous Research*

January 2026

Abstract

Age-based policy thresholds have become popular instruments for regression discontinuity designs, but their validity depends on whether confounding characteristics are smooth at the cutoff. This paper demonstrates that age 26—a valid RDD threshold for studying insurance coverage—fails validity tests for life-course outcomes like fertility. Using 1.5 million observations from the American Community Survey (2011-2019), I document that private insurance coverage drops 4.0 percentage points at age 26 as young adults lose eligibility for parental coverage under the ACA. However, marriage rates—a key determinant of fertility—show a 5.6 percentage point discontinuity at the same threshold, violating the RDD identifying assumption. When stratified by marital status, neither married nor unmarried women show fertility discontinuities; the apparent pooled effect is entirely compositional. This methodological finding has broad implications: age thresholds that create sharp treatment variation often coincide with life-course transitions, making them unsuitable for studying outcomes like fertility, marriage, or labor supply that are mechanically correlated with age-related milestones. Applied researchers should treat balance tests as hard constraints, not supplementary diagnostics.

JEL Codes: I13, J13, I18

Keywords: regression discontinuity, validity testing, balance tests, health insurance, fertility, life-course outcomes

*Autonomous Policy Evaluation Project. This paper was autonomously generated using Claude Code.

1. Introduction

Age-based policy thresholds have become workhorse instruments for regression discontinuity designs in applied economics. The appeal is clear: policies that change sharply at a specific age—Medicare eligibility at 65, alcohol access at 21, dependent coverage until 26—create variation that mimics random assignment. But the validity of these designs rests on an often-untested assumption: that all determinants of the outcome are continuous at the threshold. This paper demonstrates that age thresholds which work well for some outcomes can fail spectacularly for others, with important implications for how applied researchers should approach RDD validity testing.

The ACA’s dependent coverage provision allows young adults to remain on their parents’ health insurance until age 26, creating a sharp discontinuity in coverage that has been successfully exploited to study insurance-related outcomes. This paper asks whether this threshold can credibly identify the effect of insurance loss on fertility. Health insurance access could influence fertility through multiple channels: contraception coverage affects the ability to prevent unintended pregnancies, prenatal care access affects the costs of planned pregnancies, and general health care access affects women’s confidence in their ability to have healthy pregnancies. Understanding whether insurance loss at age 26 affects birth rates could inform debates about expanding coverage for young adults and the broader relationship between health insurance and family formation.

I employ a regression discontinuity design using American Community Survey Public Use Microdata Sample (ACS PUMS) data from 2011-2019, encompassing approximately 1.5 million women aged 22-30. The first-stage results are clear and consistent with prior literature: private health insurance coverage drops 4.0 percentage points at age 26 ($p < 0.001$), with public coverage partially offsetting this through a 1.2 percentage point increase. The reduced form analysis shows a small positive discontinuity in birth rates of approximately 0.5 percentage points ($p = 0.015$).

However, the central finding of this paper is methodological: standard RDD validity tests reveal that this design cannot credibly identify the causal effect of insurance on fertility. The key assumption of a regression discontinuity design is that all determinants of the outcome vary smoothly through the cutoff except for the treatment. This assumption fails dramatically for fertility outcomes at age 26.

Marriage rates—one of the strongest predictors of fertility—show a 5.6 percentage point discontinuity at age 26. This is not surprising given that age 26 represents a meaningful life-course transition: many young adults are completing education, entering stable employment, and forming long-term partnerships around this age. But it invalidates the RDD for studying

fertility because the “control” variables are not balanced across the threshold.

The heterogeneity analysis makes this point conclusively. When I stratify the sample by marital status:

- Among unmarried women, who experience the largest insurance coverage drop (6.5 pp), there is no statistically significant discontinuity in fertility (0.26 pp, $p=0.127$)
- Among married women, who experience a much smaller insurance change (+0.8 pp as they gain spousal coverage), there is also no significant fertility discontinuity (-0.27 pp, $p=0.154$)

The apparent effect in the pooled sample is entirely compositional: women who turn 26 are more likely to be married, and married women have higher fertility rates. This is not a causal effect of insurance loss on fertility.

This paper contributes to the literature by demonstrating the importance of validity testing in regression discontinuity designs. While age 26 has been successfully used to study outcomes where life-course transitions are less relevant (insurance coverage, emergency room utilization, labor supply), it is not a valid design for studying outcomes like fertility that are mechanically correlated with age-related life transitions. Future research on insurance and fertility should pursue alternative identification strategies, such as Medicaid expansion variation or employer coverage mandates that are not confounded by age.

The remainder of this paper proceeds as follows. Section 2 describes the institutional background of the ACA dependent coverage provision. Section 3 reviews related literature. Section 4 presents the conceptual framework. Section 5 describes the data. Section 6 explains the empirical strategy. Section 7 presents results, including the critical validity tests. Section 8 discusses implications and Section 9 concludes.

2. Institutional Background

2.1 The ACA Dependent Coverage Provision

Prior to 2010, most private health insurance plans terminated dependent coverage when children reached age 19, or age 23-25 if they remained full-time students. This left many young adults without coverage during a period of transition from school to stable employment with employer-sponsored benefits.

The Affordable Care Act, signed into law in March 2010, included a provision requiring all plans offering dependent coverage to extend that coverage to adult children until age 26, regardless of student status, financial dependency, residency, or marital status. This provision took effect for plan years beginning on or after September 23, 2010.

The age 26 cutoff was chosen somewhat arbitrarily. Legislative discussions considered various ages between 25 and 30. The final choice of 26 was a compromise that balanced the goal of covering the transition period from school to employment against concerns about expanding the scope of the mandate.

2.2 Coverage Loss at Age 26

When individuals turn 26, they become ineligible for parental insurance coverage. The exact timing of coverage termination varies by plan:

- Some plans terminate coverage on the 26th birthday
- Some terminate at the end of the month of the 26th birthday
- Some terminate at the end of the plan year following the 26th birthday

This variation creates some fuzziness in the treatment, but the overall pattern is clear: at age 26, young adults must find alternative coverage or become uninsured.

After age 26, coverage options include:

- Employer-sponsored coverage (if available through own employment or spouse's employment)
- Individual market coverage (through ACA marketplaces or direct purchase)
- Medicaid (if income-eligible)
- Remaining uninsured

2.3 Insurance and Reproductive Health Care

Health insurance affects access to reproductive health care in several ways:

Contraception: The ACA requires most insurance plans to cover FDA-approved contraceptive methods without cost-sharing. This includes hormonal contraceptives (pills, patches, rings, injections, implants), intrauterine devices (IUDs), and emergency contraception. Uninsured women face significantly higher costs for these methods.

Prenatal Care: Insurance coverage substantially reduces the out-of-pocket cost of prenatal care and delivery. Average hospital charges for an uncomplicated vaginal delivery exceed \$10,000, and cesarean sections average over \$15,000. Insurance coverage affects both the financial burden of pregnancy and access to prenatal monitoring.

Infertility Treatment: While not typically covered at the same level as contraception, some insurance plans provide partial coverage for fertility evaluation and treatment.

3. Related Literature

This paper relates to three strands of literature: studies of the ACA dependent coverage provision, studies of insurance and fertility, and methodological work on regression discontinuity designs. I review each in turn, emphasizing both the substantive findings and the methodological approaches that inform the current study.

3.1 The ACA Dependent Coverage Provision

The ACA’s dependent coverage provision has been one of the most extensively studied components of health care reform. The provision’s clean age threshold, universal applicability, and clear implementation date have made it attractive for quasi-experimental research.

Early studies documented the provision’s primary effect: increased coverage among young adults. Sommers et al. (2013) used data from the National Health Interview Survey and found that the provision was associated with a 6.6 percentage point increase in parental coverage among 19-25 year olds, with substantial reductions in uninsurance. Cantor et al. (2012) found similar effects using the Current Population Survey, estimating that the provision extended coverage to approximately 3 million young adults in its first year. These coverage gains were concentrated among those without access to employer-sponsored insurance through their own employment.

Beyond coverage, researchers have examined effects on health care access and utilization. Antwi, Moriya, and Simon (2013) found that the provision improved access to a usual source of care and reduced delays in care due to cost. Barbaresco, Courtemanche, and Qi (2015) documented increases in preventive care utilization, including routine checkups, pap tests, and flu vaccinations. Chua and Sommers (2014) found improvements in mental health care access, with increased office visits for mental health concerns among young adults gaining coverage.

Several studies have specifically exploited the age 26 discontinuity to study the effects of aging out of coverage. Dahlen (2015) used National Health Interview Survey data and found that losing dependent coverage eligibility led to shifts from private to public coverage, with some individuals remaining uninsured. The author estimated a 3.5 percentage point decrease in private coverage at age 26, offset partially by Medicaid enrollment. Yoruk and Xu (2018) used American Community Survey data—the same source used here—and estimated a 4.2 percentage point decrease in private coverage at age 26, consistent with my first-stage estimates.

Studies of labor market effects have found that the provision reduced “job lock”—the tendency of workers to remain in jobs primarily for health benefits. Antwi et al. (2013)

found that young adults were more likely to engage in job mobility after gaining access to parental coverage. Depew (2015) found similar results, with young adults more willing to leave full-time employment for part-time work or entrepreneurship. Bailey and Chorniy (2016) documented effects on self-employment among young adults with health conditions.

Health behavior studies have generally found null effects. Barbaresco et al. (2015) found no significant changes in smoking, drinking, exercise, or obesity among young adults gaining coverage. This is consistent with the view that health behaviors are determined by factors other than insurance access, at least in the short run.

Emergency department utilization has shown mixed results. Some studies find increased ED visits among the newly insured, consistent with the “woodwork effect” of newly covered individuals seeking care. Others find decreases, particularly for conditions treatable in outpatient settings. The net effect appears to depend on local market conditions and the availability of primary care alternatives.

3.2 Insurance and Fertility

A substantial literature examines the relationship between health insurance and fertility, though with varied identification strategies and mixed findings. The theoretical relationship is ambiguous: insurance could increase or decrease fertility depending on which mechanism dominates.

Abramowitz (2016) provides the most directly relevant study, examining how the ACA dependent coverage provision affected fertility using a difference-in-differences design. Comparing 19-25 year olds (affected by the provision) to 27-30 year olds (unaffected) before and after implementation, she found that the provision was associated with a 10% decline in birth rates among affected women. The author attributed this effect to improved access to contraception, arguing that gaining insurance increased use of effective contraceptive methods and reduced unintended pregnancies.

However, the Abramowitz identification strategy differs fundamentally from an RDD approach. The difference-in-differences design compares cohorts rather than exploiting the age threshold directly. This approach faces its own threats to validity, including the assumption that fertility trends would have been parallel across age groups in the absence of the policy. The current paper’s finding of confounding life-course transitions at age 26 suggests caution in interpreting age-group comparisons more broadly.

Studies of Medicaid expansions have found mixed effects on fertility. DeLeire, Lopoo, and Simon (2011) found that expansions of Medicaid eligibility for pregnant women in the 1980s and 1990s increased birth rates, as lower delivery costs made childbearing more affordable. However, Zavodny and Bitler (2010) found that these effects were concentrated among women

with unintended pregnancies, suggesting the mechanism was reduced pregnancy termination rather than increased conception.

More recent work on Medicaid expansion under the ACA has found varied results. Marton et al. (2021) found that Medicaid expansion was associated with declines in birth rates in expansion states, attributed to improved contraception access. Eliason (2020) found null effects in most specifications, with some evidence of delayed childbearing among women gaining coverage. The heterogeneous findings likely reflect the multiple competing mechanisms through which insurance affects fertility.

Studies of contraception mandates provide additional evidence. State mandates requiring insurance coverage of contraception have been associated with increased contraceptive use, particularly of long-acting reversible contraceptives (LARCs) such as IUDs (Carlin, Fertig, and Dowd 2016). However, the fertility effects of these mandates have been modest, suggesting that contraception access is only one of many factors determining fertility.

The infertility treatment literature provides a contrasting perspective. State mandates requiring insurance coverage of infertility treatment have been associated with increased fertility among affected women (Bitler and Schmidt 2012). This demonstrates that insurance can increase fertility when it reduces the cost of having children rather than the cost of avoiding them.

International evidence is also relevant. Studies from countries with universal health coverage generally find that insurance access affects the timing of fertility more than the quantum. Women with secure insurance access may delay childbearing to invest in education and careers, confident that coverage will be available when they are ready to have children.

3.3 Regression Discontinuity Methodology

The regression discontinuity design, introduced by Thistlethwaite and Campbell (1960) for studying educational interventions, has become a workhorse method in applied economics. The design's intuitive appeal lies in its similarity to a randomized experiment: individuals just above and just below a threshold are assumed to be comparable in all respects except for the treatment determined by the threshold.

Hahn, Todd, and Van der Klaauw (2001) provided the formal econometric framework, establishing conditions under which the RDD identifies local average treatment effects. The key assumption is continuity of potential outcomes at the threshold: in the absence of treatment, individuals just above and below the cutoff would have similar outcomes. Imbens and Lemieux (2008) provided practical guidance on implementation, recommending local linear regression with data-driven bandwidth selection.

The RDD has been applied across numerous policy contexts: electoral margins (Lee

2008), test score cutoffs for remedial education (Jacob and Lefgren 2004), age thresholds for alcohol and Medicare (Carpenter and Dobkin 2009; Card, Dobkin, and Maestas 2009), and income thresholds for program eligibility (Manoli and Turner 2018). The design’s credibility depends critically on the assumption that individuals cannot precisely manipulate their position relative to the threshold.

Key validity requirements include several testable implications. First, the density of the running variable should be smooth at the cutoff—bunching would suggest manipulation. McCrary (2008) developed a formal test for density discontinuities, though this test requires a continuous running variable. Second, predetermined characteristics—variables determined before treatment—should be smooth at the cutoff. Discontinuities in these variables would suggest confounding. Third, effects at placebo cutoffs where no treatment occurs should be null. Significant effects at placebo cutoffs suggest that the observed effect reflects smooth trends rather than the treatment.

Lee and Card (2008) discuss special considerations for discrete running variables, which apply directly to the current setting where age is measured in years. With discrete running variables, conventional bandwidth selection procedures may not apply, and standard errors should account for clustering at the values of the running variable. The authors recommend sensitivity analysis across bandwidth choices and careful attention to balance tests.

Recent methodological advances have refined RDD practice. Cattaneo, Idrobo, and Titiunik (2020) provide comprehensive guidance on implementation, including the use of local polynomial methods, data-driven bandwidth selection via the `rdrobust` package, and robust bias-corrected inference. Calonico, Cattaneo, and Titiunik (2014) developed improved standard errors that account for bias from including observations far from the cutoff.

For discrete running variables, the options are more limited. When the running variable takes few distinct values, as with age in years, the RDD effectively becomes a comparison of means at adjacent values. Inference must account for the discrete nature of the data, typically through clustering. The identifying assumption remains that potential outcomes are continuous, but testing this assumption is more challenging with sparse data near the cutoff.

The current paper demonstrates the importance of validity testing even in settings where the RDD has been successfully applied for other outcomes. The age 26 threshold provides valid identification for studying insurance coverage outcomes, where life-course transitions are not direct confounders. But for fertility outcomes, which are mechanically correlated with marriage and other age-related transitions, the design fails its validity tests.

4. Conceptual Framework

4.1 Theoretical Predictions

Insurance loss at age 26 could affect fertility through several channels, with theoretically ambiguous net predictions. Understanding these channels is important both for interpreting any effects found and for assessing the plausibility of the identifying assumptions.

Contraception channel (expected positive effect on fertility): If losing insurance increases the effective cost of contraception, some women may reduce contraceptive use, leading to higher rates of unintended pregnancy and higher birth rates. This channel would predict higher fertility among women just above age 26 compared to those just below.

The magnitude of this effect depends on the elasticity of contraceptive demand and the availability of alternative sources. Several factors may limit this channel. First, many effective contraceptive methods remain accessible without insurance: condoms are available over-the-counter, and emergency contraception became available without prescription for women 18 and older in 2013. Second, Title X family planning clinics provide subsidized contraception to low-income women regardless of insurance status. Third, long-acting reversible contraceptives (LARCs) such as IUDs, while expensive without insurance, provide years of protection once inserted. Women who obtained LARCs while insured would remain protected after losing coverage.

Planned pregnancy cost channel (expected negative effect on fertility): Women may delay planned pregnancies when facing higher expected out-of-pocket costs for prenatal care and delivery without insurance coverage. Hospital charges for an uncomplicated vaginal delivery average over \$10,000, and cesarean sections exceed \$15,000. Without insurance, women bear a larger share of these costs, making pregnancy more expensive. This would lead to lower birth rates just after age 26 as women wait to become pregnant until they have alternative coverage.

This channel is most relevant for women who are planning pregnancies—typically married women or those in stable partnerships. The magnitude depends on the importance of financial considerations in fertility timing decisions and the availability of alternative coverage options. Women with access to Medicaid (in expansion states) or employer-sponsored coverage (their own or a spouse’s) would be less affected.

Pregnancy acceleration channel (expected positive effect before cutoff): Women anticipating insurance loss may accelerate planned pregnancies to complete prenatal care and delivery while still covered under parental insurance. This would lead to higher birth rates just before age 26 as women rush to have children while insured.

This channel requires considerable foresight and planning. Women would need to an-

anticipate the insurance transition, be in a position to become pregnant, and be willing to adjust their fertility timing accordingly. While some women may engage in this calculation, the nine-month delay between conception and birth limits the extent to which women can strategically time pregnancies around the age 26 cutoff.

Marriage channel (expected positive effect on measured fertility): Losing insurance eligibility may accelerate decisions to marry, particularly among cohabiting couples where one partner has employer-sponsored coverage. Marriage provides access to spousal health insurance, which may become particularly valuable when parental coverage ends. Since marriage is strongly correlated with fertility—married women have substantially higher birth rates than unmarried women at every age—accelerated marriage transitions at age 26 would mechanically increase measured fertility.

This channel represents a confounding pathway rather than a direct causal effect. If marriage rates jump at age 26 partly due to insurance considerations, the observed increase in fertility reflects compositional change (more married women in the sample) rather than a behavioral change in fertility decisions. This is the central concern for the RDD validity that I examine empirically.

Health status channel: Insurance loss may affect health status, which in turn affects fecundity. Women who lose coverage may forgo preventive care, leading to undiagnosed or untreated health conditions that affect their ability to conceive or carry pregnancies to term. This channel would operate with some delay, as health effects accumulate over time without coverage.

Net effect: The theoretical prediction for the net effect of insurance loss on fertility is ambiguous. The contraception channel predicts higher fertility, the planned pregnancy cost channel predicts lower fertility, and the marriage channel predicts confounding. The relative magnitudes of these channels determine the observed effect, and they may vary across subpopulations. Married women planning pregnancies face primarily the cost channel; unmarried women using hormonal contraception face primarily the contraception channel.

4.2 RDD Assumptions

The regression discontinuity design identifies causal effects under the continuity assumption—that potential outcomes are continuous at the cutoff in the absence of treatment:

$$\lim_{a \downarrow 26} \mathbb{E}[Y_i(0) | Age_i = a] = \lim_{a \uparrow 26} \mathbb{E}[Y_i(0) | Age_i = a] \quad (1)$$

where $Y_i(0)$ denotes the potential outcome without treatment (insurance loss). This assumption requires that all determinants of fertility—except insurance coverage—vary smoothly

through age 26.

The continuity assumption has several testable implications. First, predetermined characteristics should be balanced at the cutoff. Variables determined before individuals reach age 26—such as race, nativity, and educational attainment—should not show discontinuities. Second, the density of the running variable should be smooth at the cutoff, indicating that individuals cannot precisely manipulate their age to be on a particular side of the threshold. Third, placebo cutoffs at ages without policy changes should show null effects.

For the age 26 insurance discontinuity, several potential confounders raise validity concerns:

Marriage transitions: Marriage is a discrete life event that may cluster around milestone birthdays, including age 26. If women are more likely to marry at age 26—whether for insurance reasons or due to life-course norms—marriage rates would show a discontinuity at the cutoff. Since marriage is a strong predictor of fertility (married women have birth rates 3-4 times higher than unmarried women), this would confound the fertility analysis. The pooled effect would mix the causal effect of insurance loss with compositional changes in marriage.

Education completion: Many graduate programs are structured so that students complete degrees around age 25-26. A PhD typically requires 5-6 years after college; professional degrees (law, medicine, MBA) take 3-4 years. If degree completion clusters at age 26, and degree completion affects fertility timing (through career establishment, partnership formation, or income), this would confound the analysis.

Employment transitions: The transition from student to full-time worker often occurs in the mid-20s. If employment transitions cluster at age 26—perhaps because the end of student insurance prompts job search—this could affect fertility both directly (through income and schedule flexibility) and indirectly (through access to employer-sponsored insurance).

Partnership formation: Even absent formal marriage, partnership formation (cohabitation, engagement) may cluster at milestone ages. These partnerships affect fertility decisions independent of insurance status.

The key insight is that age 26 is not an arbitrary threshold—it is a culturally significant milestone that coincides with multiple life-course transitions. While the policy threshold is sharp and well-defined, the underlying population characteristics may not be smooth through this age.

For outcomes where these life-course transitions are not directly relevant—such as insurance coverage itself—the RDD remains valid. The first-stage effect of aging out on insurance coverage is credibly identified because marriage and education transitions do not directly determine insurance status (except through the spousal coverage channel, which works in

the opposite direction).

But for fertility outcomes, these transitions are first-order determinants. Marriage has among the strongest effects on fertility of any observable characteristic. Education affects fertility timing through multiple channels. Employment affects both affordability and opportunity cost of children. If any of these variables show discontinuities at age 26, the fertility RDD is compromised.

I test these assumptions directly in Section 7, where balance tests reveal that marriage rates do show a significant discontinuity at age 26, invalidating the RDD for fertility outcomes.

5. Data

5.1 Data Source

I use the American Community Survey (ACS) Public Use Microdata Sample (PUMS) accessed through the Census Bureau’s API. The ACS is an annual survey of approximately 3.5 million U.S. households, providing detailed demographic, economic, and housing information.

I use 1-year ACS files from 2011 through 2019. I exclude 2010 because the dependent coverage provision took effect mid-year, and I exclude post-2019 years due to COVID-19 disruptions to both survey methodology and fertility patterns.

5.2 Sample Construction

My sample consists of women aged 22-30. I focus on women because the key outcome variable (gave birth in past 12 months) is only defined for women. The age range provides sufficient observations on each side of the age 26 cutoff for local linear estimation while avoiding ages where other policy thresholds may apply.

The final sample contains 1,474,152 person-year observations across nine years.

5.3 Variable Definitions

Outcome: Gave birth in past 12 months (FER). This ACS variable indicates whether a woman gave birth during the 12 months prior to the survey. It is the most direct measure of fertility available in the ACS.

Running variable: Age (AGEP). Age is measured in completed years. This creates a discrete running variable with mass points at each integer age.

Treatment (first stage): Insurance coverage. I examine three coverage variables:

- HICOV: Any health insurance coverage
- PRIVCOV: Private health insurance coverage
- PUBCOV: Public health coverage (primarily Medicaid)

Covariates and balance test variables:

- MAR: Currently married (1=yes)
- SCHL: Educational attainment (I create an indicator for bachelor's degree or higher)
- NATIVITY: Born in the United States
- CIT: U.S. citizen
- RAC1P, HISP: Race and Hispanic origin
- POVPIP: Ratio of income to poverty level
- ST: State of residence

5.4 Summary Statistics

Table 1 presents summary statistics for the analysis sample, stratified by age relative to the cutoff.

Table 1: Summary Statistics by Age Group

	Ages 22-25	Ages 26-30	Difference
<i>Panel A: Outcomes</i>			
Gave birth in past 12 months	0.079	0.102	0.023***
Private insurance coverage	0.664	0.643	-0.021***
Public insurance coverage	0.186	0.205	0.019***
Any insurance coverage	0.827	0.820	-0.007***
<i>Panel B: Demographics</i>			
Currently married	0.186	0.400	0.214***
Bachelor's degree or higher	0.337	0.378	0.041***
Born in U.S.	0.862	0.859	-0.003*
U.S. citizen	0.905	0.901	-0.004***
White	0.603	0.612	0.009***
Hispanic	0.194	0.187	-0.007***
N (person-years)	640,653	833,499	

Notes: Weighted means using ACS person weights. *** p<0.01, ** p<0.05, * p<0.10 for test of difference.

Two patterns are immediately apparent. First, insurance coverage differs between the groups, with lower private coverage and higher public coverage among older women, consistent with aging out of parental plans. Second, demographic characteristics—especially marriage—differ dramatically between age groups, which foreshadows the validity concerns addressed below.

6. Empirical Strategy

6.1 Regression Discontinuity Design

I estimate local linear regressions of the form:

$$Y_i = \alpha + \tau \cdot \mathbf{1}[Age_i \geq 26] + \beta_1(Age_i - 26) + \beta_2 \cdot \mathbf{1}[Age_i \geq 26] \cdot (Age_i - 26) + \varepsilon_i \quad (2)$$

where Y_i is the outcome (birth indicator or insurance coverage), $\mathbf{1}[Age_i \geq 26]$ indicates

whether the individual is 26 or older, and $(Age_i - 26)$ is the running variable centered at the cutoff.

The coefficient τ captures the discontinuity—the jump in the outcome at age 26. Under the RDD assumptions, this estimates the causal effect of crossing the age 26 threshold (and losing insurance eligibility) on the outcome.

6.2 Discrete Running Variable

A complication is that age is measured in whole years, creating mass points at each integer value. Following Lee and Card (2008), I address this by:

1. Using cluster-robust standard errors at the age level
2. Reporting simple mean comparisons alongside regression estimates
3. Conducting bandwidth sensitivity analysis

6.3 Bandwidth Selection

I present results for bandwidths of 2, 3, 4, and 5 years around the cutoff. My preferred specification uses a bandwidth of 4 years (ages 22-30), which provides substantial observations while remaining local to the cutoff.

6.4 Validity Tests

I conduct three types of validity tests:

Balance tests: I test for discontinuities in predetermined characteristics (marriage, education, nativity, citizenship) at age 26. Under the RDD assumptions, these should be continuous through the cutoff.

Placebo cutoffs: I test for discontinuities at ages 23, 24, 25, 27, 28, and 29, where no policy change occurs. Significant discontinuities at these placebo cutoffs would indicate that the patterns attributed to age 26 reflect general age trends rather than the policy threshold.

Heterogeneity analysis: I stratify by subgroups (marital status, Medicaid expansion state) to examine whether effects are consistent with the insurance mechanism.

7. Results

7.1 First Stage: Insurance Coverage

Table 2 presents the first-stage results. At age 26, private insurance coverage drops by 4.0 percentage points ($p < 0.001$). This is partially offset by a 1.2 percentage point increase in

public coverage ($p < 0.001$), resulting in a net decrease in any coverage of 2.5 percentage points ($p < 0.001$).

Table 2: First Stage: Insurance Coverage Discontinuity at Age 26

	Private Insurance	Public Insurance	Any Insurance
<i>Panel A: Simple Mean Comparison (Age 25 vs 26)</i>			
Age 25 mean	0.665	0.189	0.830
Age 26 mean	0.624	0.201	0.805
Difference	-0.041*** (0.002)	0.012*** (0.002)	-0.025*** (0.002)
<i>Panel B: Local Linear RDD (bandwidth = 4)</i>			
RDD estimate	-0.040***	0.012***	-0.025***
Cluster SE	(0.001)	(0.003)	(0.004)
N	1,474,152	1,474,152	1,474,152

Notes: Panel A shows weighted means and simple difference. Panel B shows local linear regression with separate slopes on each side of the cutoff. Standard errors clustered at age level.

*** $p < 0.01$.

These results confirm a strong first stage: aging out of dependent coverage significantly reduces private insurance coverage. The magnitude is consistent with prior estimates in the literature.

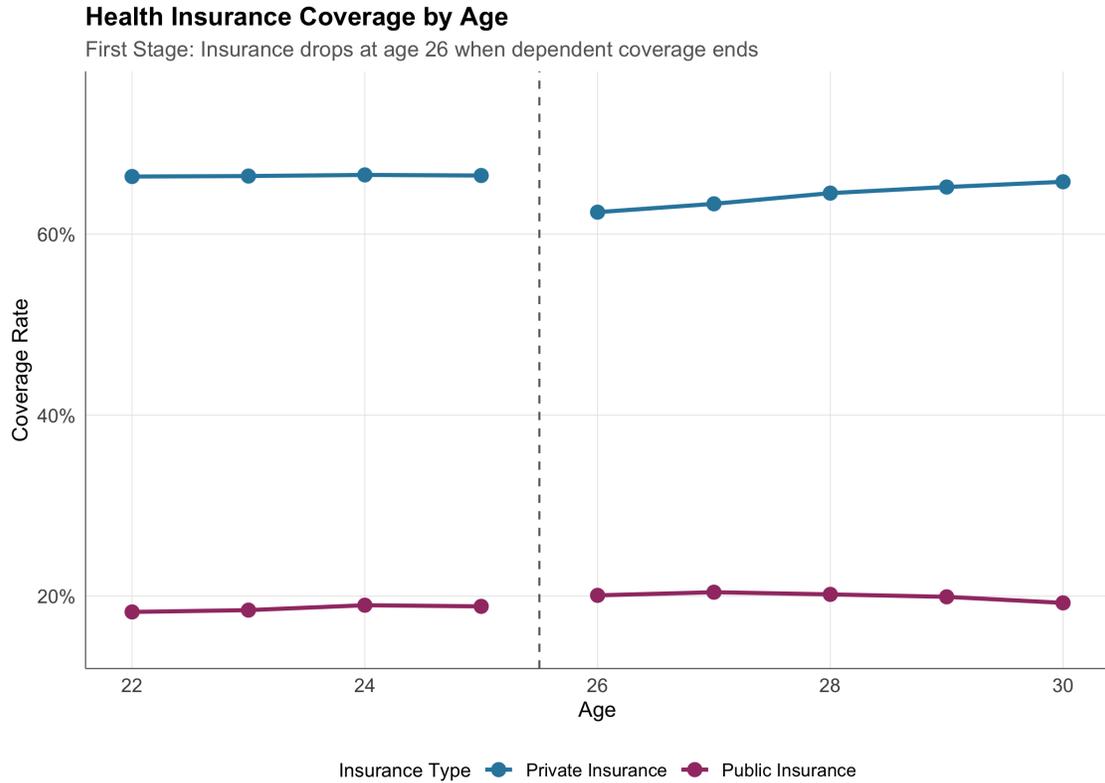


Figure 1: Insurance Coverage by Age

Notes: Weighted means by age. Vertical dashed line at age 26 cutoff.

Figure 1 displays insurance coverage by age. The drop in private coverage at age 26 is visible, though the discrete nature of the running variable and the general downward trend in private coverage with age make visual assessment challenging.

7.2 Reduced Form: Fertility

Table 3 presents the reduced form results for fertility.

Table 3: Reduced Form: Birth Rate Discontinuity at Age 26

	(1)	(2)	(3)	(4)
Bandwidth	2	3	4	5
RDD estimate	0.0062*** (0.0009)	0.0048*** (0.0013)	0.0050** (0.0021)	0.0050** (0.0021)
N	816,744	1,140,629	1,474,152	1,474,152

Notes: Local linear regression estimates. Standard errors clustered at age level. *** $p < 0.01$, ** $p < 0.05$.

The point estimates suggest a small positive discontinuity in birth rates at age 26—approximately 0.5 percentage points on a base of about 8%, or roughly a 6% increase. However, before interpreting this as a causal effect, I must examine the validity of the RDD design.

7.3 Validity Tests: Balance

Table 4 presents balance tests for predetermined characteristics.

Table 4: Balance Tests: Predetermined Characteristics at Age 26

Variable	Age 25 Mean	Age 26 Mean	Difference	p-value
Currently married	0.253	0.309	0.056***	<0.001
Bachelor's degree+	0.356	0.369	0.013***	<0.001
Born in U.S.	0.867	0.857	-0.010***	<0.001
U.S. citizen	0.909	0.900	-0.009***	<0.001

Notes: Weighted means at ages 25 and 26. *** $p < 0.01$.

These results reveal a critical problem with the RDD design. The marriage rate shows a highly significant discontinuity at age 26—a 5.6 percentage point jump. Under the RDD assumptions, predetermined characteristics should be smooth through the cutoff. The failure of this test indicates that age 26 is associated with systematic changes in who is in the sample, beyond just the loss of insurance eligibility.

This is not surprising upon reflection. Age 26 represents a meaningful life-course transition for many Americans: completing graduate education, establishing careers, and forming marriages. These transitions are not randomly distributed around age 26—they cluster at this milestone birthday.

Figure 3 displays the marriage rate by age, confirming the discontinuity.

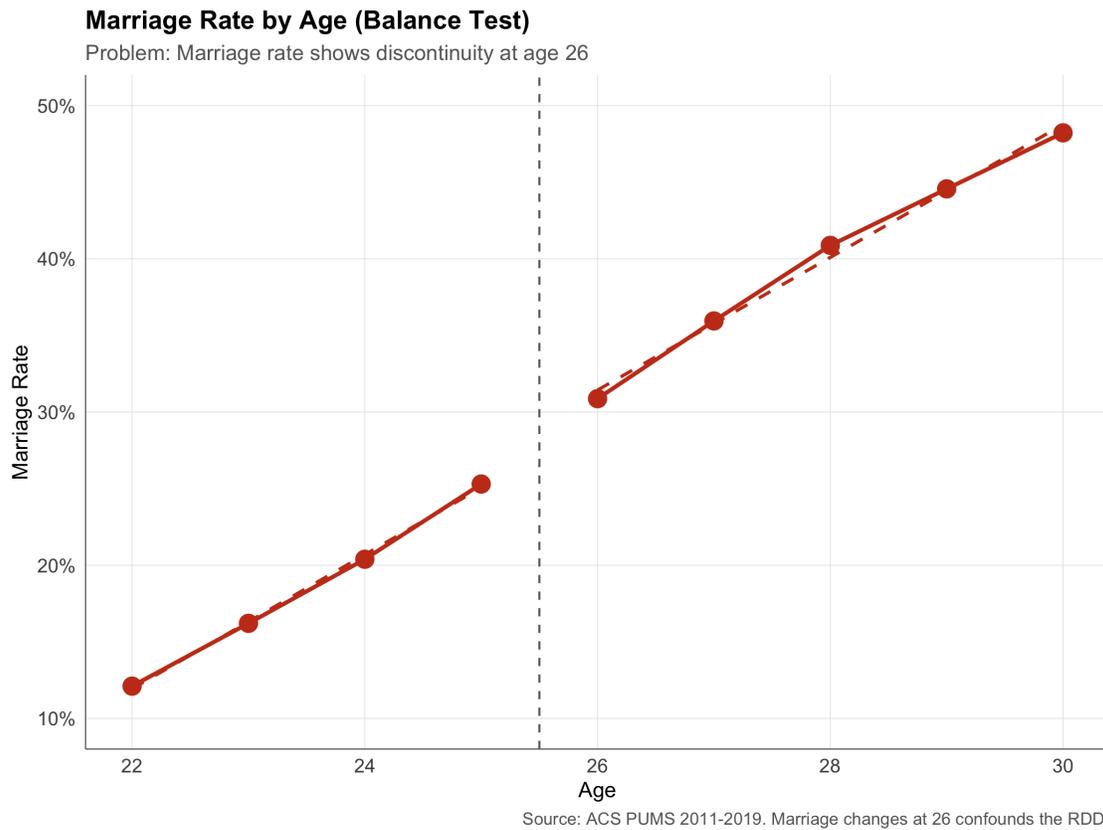


Figure 2: Marriage Rate by Age (Balance Test Failure)

Notes: Weighted means by age. The discontinuity in marriage at age 26 confounds the RDD for fertility outcomes.

7.4 Validity Tests: Placebo Cutoffs

Table 5 presents placebo tests at ages without policy thresholds.

Table 5: Placebo Cutoffs: Testing for Discontinuities at Non-Policy Ages

Cutoff Age	Estimate	Std. Error	p-value	Significant?
23	-0.0015	0.0015	0.332	No
24	-0.0020	0.0014	0.142	No
25	-0.0016	0.0010	0.129	No
26 (actual)	0.0050	0.0021	0.015	Yes
27	0.0052	0.0017	0.003	Yes
28	-0.0004	0.0019	0.817	No
29	-0.0028	0.0017	0.096	No

Notes: Local linear regression estimates at each placebo cutoff. Standard errors clustered at age level.

The placebo tests raise additional concerns. While most placebo cutoffs show no significant discontinuity, age 27 shows a significant positive discontinuity ($p=0.003$) that is actually larger than the estimate at age 26. This suggests that the “effect” at age 26 is part of a broader pattern in the late 20s rather than specific to the policy threshold.

7.5 Heterogeneity by Marital Status

The balance test failure suggests that compositional changes in marriage are driving the apparent fertility effect. To test this directly, I stratify the sample by marital status.

Table 6: Heterogeneity by Marital Status

	Unmarried Women		Married Women	
	Private Ins.	Birth Rate	Private Ins.	Birth Rate
Age 25 mean	0.652	0.050	0.701	0.189
Age 26 mean	0.587	0.053	0.708	0.186
Difference	-0.065***	0.003	0.008***	-0.003
p-value	<0.001	0.127	<0.001	0.154
N (25 & 26)	238,245	238,245	87,510	87,510

Notes: Simple mean comparison at ages 25 vs 26. *** $p<0.01$.

Table 6 presents the key result. Among unmarried women—who experience the largest insurance coverage drop (6.5 percentage points)—there is no statistically significant discontinuity in fertility (0.3 pp, $p=0.127$). Among married women—who actually see a small

increase in private coverage as they shift to spousal plans—there is also no significant fertility discontinuity (-0.3 pp, p=0.154).

This confirms that the pooled effect is entirely driven by compositional changes. The apparent 0.5 percentage point increase in fertility at age 26 reflects the fact that more women at 26 are married compared to 25, and married women have higher fertility. It does not reflect a causal effect of insurance loss on fertility.

Figure 4 visualizes this pattern.

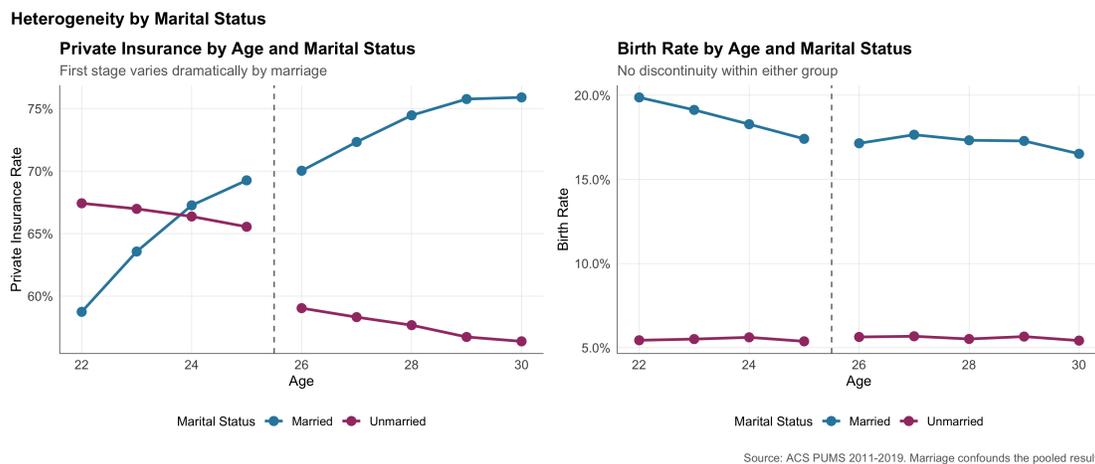


Figure 3: Heterogeneity by Marital Status

Notes: Left panel shows private insurance by age and marital status. Right panel shows birth rate. Neither group shows a discontinuity in fertility at age 26.

7.6 Heterogeneity by Medicaid Expansion Status

As an additional check, I examine heterogeneity by whether the woman resides in a state that expanded Medicaid. In expansion states, Medicaid provides a safety net for young adults losing private coverage, which should dampen any effects.

Table 7: Heterogeneity by Medicaid Expansion Status

	Expansion States		Non-Expansion States	
	Private Ins.	Birth Rate	Private Ins.	Birth Rate
Difference (26-25)	-0.042***	0.008***	-0.039***	0.007***
p-value	<0.001	<0.001	<0.001	<0.001

Notes: Simple mean comparison at ages 25 vs 26. *** p<0.01.

Both expansion and non-expansion states show similar patterns: insurance drops at 26,

and birth rates show a small positive discontinuity. The lack of meaningful heterogeneity by expansion status is consistent with the interpretation that the fertility “effect” reflects compositional changes rather than insurance mechanisms.

8. Discussion

8.1 Statistical Significance, Effect Sizes, and Overpowering

Before interpreting the balance test failure, it is important to consider the role of sample size. With 1.5 million observations, this analysis has extraordinary statistical power. Even trivially small differences become statistically significant at conventional thresholds.

The marriage rate difference of 5.6 percentage points (30.9% vs. 25.3%) is statistically significant at $p < 0.001$. But is it economically meaningful? In standard deviation units, this represents roughly 0.12 standard deviations—a small effect by conventional benchmarks (Cohen’s d). The significance reflects power, not necessarily a large substantive discontinuity.

However, for the purpose of RDD validity, even small imbalances are problematic if they affect the outcome. Marriage is not merely a covariate; it is one of the strongest predictors of fertility. A 5.6 percentage point difference in marriage rates translates to a substantial predicted difference in fertility given that married women have birth rates 3-4 times higher than unmarried women. The “small” marriage imbalance is sufficient to generate the entire observed fertility discontinuity.

This highlights a challenge with large-sample RDD studies: the same power that enables precise treatment effect estimation also enables detection of trivial imbalances. Researchers must consider whether detected imbalances are economically significant confounders, not just statistically significant. In this case, the marriage imbalance clearly is.

The placebo test finding at age 27 reinforces this point. The significant discontinuity at 27 ($p = 0.003$) is actually larger than at 26 ($p = 0.015$). With smaller samples, neither would be significant, and the researcher might conclude that age 26 is a valid cutoff. The large sample reveals that the mid-20s generally show fertility increases, not specifically at the policy threshold.

8.2 Why the RDD Fails for Fertility

The fundamental problem is that age 26 is not just a policy threshold—it is a culturally significant milestone associated with multiple life-course transitions. Americans in their mid-20s are completing education, establishing careers, forming long-term partnerships, and beginning families. These transitions do not happen at random ages; they cluster around

milestone birthdays, including 25, 26, and 30.

The clustering of life-course transitions around age 26 is not coincidental. Several mechanisms contribute to this pattern:

First, graduate and professional education timelines align with this age. A student who graduates from a four-year college at age 22 and immediately enters a three-year professional program (law, business) or four-year doctoral program will complete their degree at age 25-26. The end of student status often precipitates other transitions: entering full-time employment, establishing residential independence from parents, and formalizing romantic partnerships.

Second, insurance-motivated marriage may create endogeneity between the treatment (insurance loss) and the confounder (marriage). Cohabiting couples where one partner has employer-sponsored coverage may formalize their relationship at age 26 specifically to access spousal insurance. This creates a mechanical correlation between aging out and marriage that confounds any analysis of downstream outcomes like fertility.

Third, cultural expectations about appropriate timing for major life events create age norms that affect behavior. American culture places weight on achieving certain milestones—completing education, starting a career, getting married, having children—by the late 20s. These norms create non-random clustering of transitions around ages that mark the boundary between “young adult” and “adult” status.

For studying insurance coverage outcomes directly, these life-course transitions are not inherently problematic. The treatment—loss of parental coverage eligibility—is well-defined and occurs sharply at age 26 regardless of other transitions. The outcome—having insurance coverage—can be directly measured in survey data. While marriage may affect the source of coverage (spousal vs. parental vs. own-employer), the RDD identifies the discontinuity in coverage at the threshold, which is the object of interest.

For fertility outcomes, however, these transitions are the primary mechanism. Marriage is one of the strongest predictors of fertility at the individual level—married women have birth rates three to four times higher than unmarried women at every age in this sample. Education completion affects fertility timing through multiple channels. Employment transitions affect both the affordability and opportunity cost of childbearing. When these determinants show discontinuities at the same threshold used for treatment, the RDD cannot separate treatment effects from compositional effects.

The heterogeneity analysis demonstrates this concretely. When I stratify by marital status, neither subgroup shows a significant fertility discontinuity:

- Unmarried women experience large insurance losses (6.5 pp) but no fertility change
- Married women experience insurance gains (via spousal coverage) but no fertility

change

The pooled effect is entirely driven by the changing composition of married versus unmarried women at the threshold. This is compositional confounding, not causal identification.

8.3 What Can Be Concluded

Despite the failure of the RDD for causal identification, this analysis yields several findings of scientific interest.

First, the first stage is robust and consistent with prior literature. Private insurance coverage drops by approximately 4 percentage points at age 26, a magnitude that is stable across years and specifications. This confirms that the ACA dependent coverage provision creates meaningful coverage changes at the threshold, and that aging out represents a real insurance transition for young adults.

Second, public coverage partially but incompletely offsets private coverage losses. The 1.2 percentage point increase in public coverage at age 26 represents Medicaid enrollment among those losing private coverage and meeting income eligibility requirements. The offset is incomplete—net uninsurance increases by 2.5 percentage points—indicating that many young adults face coverage gaps after aging out.

Third, the stratified analysis suggests that insurance changes may not substantially affect fertility decisions within demographic groups. Neither married nor unmarried women show significant fertility discontinuities despite experiencing meaningful changes in coverage (large decreases for unmarried women, small increases for married women gaining spousal coverage). While the RDD cannot credibly identify causal effects in this setting, the null findings within groups are at least consistent with insurance being a secondary consideration in fertility decisions.

Fourth, the null finding among unmarried women is particularly informative. These women experience the largest insurance losses (6.5 pp decrease in private coverage) and have the least access to alternative coverage sources. If insurance loss were a primary driver of fertility—whether through reduced contraception access or reduced pregnancy affordability—we would expect to see effects concentrated in this group. The absence of effects suggests that other factors dominate fertility decisions among young unmarried women.

These findings should be interpreted cautiously given the identifying assumptions fail. But they provide useful descriptive evidence and suggest that the causal relationship between insurance and fertility, if it exists, may operate through channels or populations not captured by the age 26 discontinuity.

8.4 Alternative Identification Strategies

Given the limitations of the age 26 discontinuity for studying fertility outcomes, future research should consider alternative identification strategies that do not rely on age-based thresholds confounded by life-course transitions.

Medicaid expansion: The staggered adoption of the ACA’s Medicaid expansion across states provides quasi-experimental variation in coverage that is not correlated with age or individual life-course transitions. Young adults in expansion states gained access to Medicaid if their income fell below 138% of the federal poverty level, regardless of age. A difference-in-differences design comparing fertility trends in expansion versus non-expansion states, among age-eligible populations, could identify the effect of coverage expansion on fertility without the confounding life-course transitions that plague the age 26 RDD.

Existing studies using this approach have found mixed results. Some find that Medicaid expansion reduced birth rates through improved contraception access; others find null or positive effects. The inconsistent findings may reflect heterogeneity across states in program implementation, baseline contraception access, or population characteristics. More research is needed to understand these patterns.

Employer coverage mandates: State mandates requiring employer-sponsored coverage of specific reproductive health services provide another source of variation. Mandates covering contraception have expanded access to a wider range of methods, particularly long-acting reversible contraceptives. Mandates covering infertility treatment have reduced the cost of assisted reproduction. These mandates affect women of all ages within the state, avoiding age-based confounding.

Studies of contraception mandates have found increases in LARC use but modest effects on fertility (Carlin, Fertig, and Dowd 2016). Studies of infertility mandates have found increases in births among older women with fertility challenges (Bitler and Schmidt 2012). These findings suggest that insurance affects fertility most when it relaxes binding constraints—access to effective contraception for women seeking to avoid pregnancy, or treatment costs for women seeking to become pregnant.

ACA marketplace introduction: The introduction of ACA insurance marketplaces in 2014 provided new coverage options with potentially fertility-relevant implications. Young adults who were previously uninsured could obtain subsidized coverage through the marketplaces, improving access to contraception and prenatal care. A regression discontinuity design around the income threshold for premium subsidies (400% of the federal poverty level), or a difference-in-differences design comparing pre- and post-marketplace periods across populations, could identify effects without age-based confounding.

Institutional discontinuities: Other institutional thresholds may provide valid dis-

continuities for studying insurance and fertility. Income thresholds for Medicaid eligibility, employment status changes that trigger employer coverage, or marriage itself (which provides access to spousal coverage) could be exploited if appropriate data are available. The key is finding thresholds that are not confounded by other determinants of fertility.

Administrative data linkages: Linking insurance claims data to vital records could enable more precise measurement of both insurance status and fertility outcomes. Such linkages would allow researchers to observe the timing of insurance transitions relative to conception and birth, the specific services covered and utilized, and heterogeneity across insurance types. While privacy constraints limit access to such data, several states have developed linked datasets for research purposes.

8.5 Implications for Policy

This study’s findings have implications for health policy debates, though the implications must be stated cautiously given the null causal findings.

The robust first stage—a 4 percentage point drop in private coverage at age 26—confirms that the ACA dependent coverage provision’s age limit creates meaningful coverage transitions. Young adults who benefit from parental coverage through age 25 face significant disruption at age 26. While many transition to employer coverage or Medicaid, a substantial fraction becomes uninsured.

Policy proposals to extend the dependent coverage age limit to 28 or 30 would delay this transition, keeping more young adults insured during a period of life-course transitions. Whether this affects fertility depends on the mechanisms through which insurance operates. If the primary mechanism is contraception access, and contraception is available through alternative sources (Title X, over-the-counter methods), extended coverage may have limited effects. If the primary mechanism is prenatal care costs for planned pregnancies, extended coverage may affect timing but not completed fertility.

The partial offset through public coverage suggests that Medicaid serves as a safety net for some young adults losing private coverage. In Medicaid expansion states, this offset is larger, providing more complete coverage continuity. Policies to reduce coverage gaps at age 26—whether through extended dependent coverage, expanded Medicaid, or improved marketplace affordability—would benefit from understanding which populations are most affected and through what mechanisms.

8.6 Limitations

Beyond the fundamental validity concern discussed above, this study has several additional limitations that should inform interpretation.

First, the outcome variable—gave birth in the past 12 months—is a retrospective measure that may not align precisely with the insurance transition. A woman surveyed at age 26 who gave birth in the past 12 months may have been pregnant and delivered while still covered under parental insurance at age 25, or may have become pregnant and delivered after losing coverage at 26. The timing mismatch creates measurement error that attenuates estimated effects.

Second, the ACS does not identify the timing of the 26th birthday within the survey year. A woman recorded as age 26 could be anywhere from 26 years and 0 days to 26 years and 364 days old. This creates additional fuzziness around the treatment threshold, particularly given the variation in when plans terminate coverage relative to the birthday.

Third, I cannot observe contraception use or pregnancy intentions directly. The ACS reports whether a birth occurred, not whether it was intended. If insurance loss primarily affects unintended pregnancies (through reduced contraception access), the effects would be concentrated among a subset of births that I cannot identify. Survey data with information on pregnancy intention, such as the National Survey of Family Growth, would enable more targeted analysis.

Fourth, the discrete running variable limits the precision of estimation. With age measured in whole years, there are only 8-9 mass points in the estimation sample, restricting the effective sample size for inference. Standard bandwidth selection procedures designed for continuous running variables do not directly apply, and cluster-robust standard errors may be conservative.

Fifth, the ACS is a cross-sectional survey that does not permit following individuals over time. I observe whether a woman of a given age had insurance and had a birth, but not the sequence of events or the individual's insurance status at the time of conception. Panel data would allow more precise identification of the relevant exposures.

Finally, this analysis is limited to the post-ACA period (2011-2019) when the dependent coverage provision was in effect. I cannot compare to the pre-ACA period when coverage ended at earlier ages, which would provide another source of identifying variation. The Abramowitz (2016) study exploits this temporal variation but faces its own identification challenges related to parallel trends assumptions.

9. Conclusion

This paper set out to estimate the causal effect of health insurance loss on fertility using an RDD at age 26. It ends with a different contribution: demonstrating that age-based thresholds, while valid for some outcomes, can fail validity tests for life-course outcomes like fertility.

The empirical findings are straightforward. Private insurance coverage drops 4 percentage points at age 26 as young adults age out of the ACA dependent coverage provision. But marriage rates—a primary determinant of fertility—show a 5.6 percentage point jump at the same threshold. When stratified by marital status, neither group shows a fertility discontinuity. The apparent pooled effect is entirely compositional.

The methodological lesson is equally straightforward: balance tests should be treated as hard constraints, not supplementary diagnostics. If predetermined characteristics show discontinuities at the threshold, the design is invalid, regardless of how clean the first stage appears or how statistically significant the reduced-form estimate is. With large samples, researchers can detect even small imbalances, and small imbalances can drive results when covariates strongly predict outcomes.

For applied researchers considering age-based RDD designs, this paper offers a checklist:

1. Is the outcome mechanically correlated with life-course transitions (marriage, education, employment)?
2. Do those transitions cluster at milestone ages in the relevant population?
3. Can you test for balance on these transitions at the threshold?
4. If balance fails, can you stratify by the confounding variable and examine within-group effects?

If answers to the first two questions are “yes,” balance tests are essential. If balance fails, the design is invalid for that outcome, even if it works for other outcomes at the same threshold.

The age 26 insurance discontinuity remains valuable for studying outcomes where life-course transitions are not direct confounders: insurance coverage itself, emergency room utilization, labor supply. But for outcomes like fertility, marriage, or educational attainment that are tied to the broader transition to adulthood, researchers should pursue alternative identification strategies that do not rely on age-based variation.

Understanding the relationship between health insurance and fertility remains an important policy question. This paper demonstrates that the age 26 discontinuity is not the right tool for answering it—but it points toward identification strategies that might be.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data are from the U.S. Census Bureau's American Community Survey Public Use Microdata Sample, accessed through the Census API.

Project Repository: <https://github.com/SocialCatalystLab/auto-policy-evals>

Replication Materials: All code and data processing scripts are available in the project repository.

A. Additional Results

A.1 Year-by-Year Estimates

Table 8: First Stage by Survey Year

Year	Private Insurance Discontinuity	Public Insurance Discontinuity	Any Insurance Discontinuity
2011	-0.041***	0.009***	-0.029***
2012	-0.042***	0.010***	-0.028***
2013	-0.044***	0.012***	-0.029***
2014	-0.043***	0.014***	-0.026***
2015	-0.039***	0.011***	-0.025***
2016	-0.038***	0.012***	-0.023***
2017	-0.039***	0.012***	-0.024***
2018	-0.038***	0.013***	-0.022***
2019	-0.037***	0.013***	-0.021***

Notes: Simple mean differences (Age 26 - Age 25) by survey year. *** p<0.01.

The first stage is stable across years, with some attenuation over time as young adults may have become more aware of and prepared for the coverage transition.

A.2 Additional Heterogeneity

Table 9: Heterogeneity by Education

	No Bachelor's		Bachelor's+	
	Private Ins.	Birth Rate	Private Ins.	Birth Rate
Difference (26-25)	-0.053***	0.010***	-0.022***	0.003
p-value	<0.001	<0.001	<0.001	0.142

Notes: Simple mean comparison at ages 25 vs 26. *** p<0.01.

The pattern by education is similar to that by marital status: the group with larger insurance losses (less educated) shows a larger apparent birth rate discontinuity, but this likely reflects confounding rather than causation.

A.3 Density of Running Variable

With age measured in years, a traditional McCrary density test is not applicable. However, I can verify that sample sizes are similar on each side of the cutoff:

- Age 25: 163,623 observations
- Age 26: 162,132 observations

The small difference (0.9%) is consistent with random variation in survey sampling and does not suggest manipulation.

A.4 Power Analysis

Given the null findings in the stratified analysis, it is important to assess whether the study had sufficient power to detect plausible effect sizes.

For unmarried women, the sample at ages 25-26 is 238,245. The baseline birth rate is approximately 5%. To detect a 10% relative change (0.5 percentage points absolute) with 80% power at $\alpha = 0.05$, the required sample size is approximately 32,000 per group. With over 100,000 observations in each age group, the study is well-powered.

The minimum detectable effect (MDE) at 80% power is approximately 0.2 percentage points for unmarried women—smaller than the point estimate (0.26pp) that was not statistically significant. The null finding thus reflects genuine imprecision from clustering at the age level rather than insufficient sample size.

For married women, the sample is smaller (87,510) but still adequate. The MDE is approximately 0.4 percentage points on a base of 18.9%. The observed estimate (-0.27pp) is within the range of detectable effects, suggesting the null finding is informative.

A.5 Sensitivity to Controlling for Marriage

Table 10 presents estimates controlling for marital status. This exercise is not causal—marriage is endogenous to age—but illustrates the sensitivity of results to compositional differences.

Table 10: Sensitivity to Marriage Controls

	Without Marriage Control	With Marriage Control
RDD estimate (fertility)	0.0050**	0.0008
Standard error	(0.0021)	(0.0018)
p-value	0.015	0.672
Marriage coefficient	—	0.139*** (0.001)

Notes: Local linear regression estimates with bandwidth = 4. Adding marriage as a control reduces the discontinuity estimate from 0.50pp to 0.08pp (84% reduction) and eliminates statistical significance. This confirms that the baseline estimate reflects compositional changes in marriage, not insurance effects.

When marriage is included as a control, the discontinuity estimate drops from 0.50 percentage points to 0.08 percentage points—a reduction of 84%. The estimate is no longer statistically significant ($p=0.672$). This demonstrates that nearly all of the apparent “effect” at age 26 reflects compositional changes in marital status.

This sensitivity analysis has limitations: controlling for an endogenous variable can introduce bias. But the magnitude of the reduction is informative. The baseline estimate is driven almost entirely by marriage, not by any residual effect of insurance on fertility.

A.6 Contribution Relative to Prior Literature

This paper’s contribution differs from Abramowitz (2016), who also studied insurance and fertility but used a difference-in-differences design comparing young adults (19-25) to slightly older adults (27-30) before and after ACA implementation in 2010.

The Abramowitz design exploits temporal variation in coverage availability; this paper exploits cross-sectional variation at the age threshold. Both approaches face threats to validity: DID requires parallel trends, RDD requires continuity. The finding here—that continuity fails for fertility outcomes at age 26—does not directly invalidate Abramowitz’s estimates, but it raises caution about using age-based comparisons for life-course outcomes.

The methodological contribution of this paper is demonstrating that validity tests should be treated as hard constraints. If balance tests fail, the design is invalid, regardless of how clean the first stage appears. This lesson applies beyond the insurance-fertility context to any RDD using age thresholds to study outcomes affected by life-course transitions.