

The Effects of Income-Based Lifeline Eligibility on Internet Adoption During the ACP Era: A Regression Discontinuity Analysis

APEP Research Team*

January 2026

Abstract

This paper examines whether being below the income threshold for FCC Lifeline broadband subsidy eligibility affects household internet adoption. Using a regression discontinuity design that exploits the 135% Federal Poverty Line (FPL) income eligibility cutoff, we estimate the effect of income-based Lifeline eligibility on broadband subscription rates. Drawing on Census American Community Survey (ACS) Public Use Microdata Sample (PUMS) data from 2021-2022, we analyze 683,574 households near the eligibility threshold. Our main finding is a null result: we find no statistically significant discontinuity in broadband adoption at the 135% FPL cutoff. The estimated intent-to-treat effect on broadband subscription is -0.3 percentage points (SE = 0.3, $p = 0.32$), with a 95% confidence interval that rules out effects larger than 0.9 percentage points. This null result persists across local quadratic specifications (though with some sensitivity at narrower bandwidths), covariate-adjusted models, and donut RD designs

*Autonomous Policy Evaluation Project. Correspondence: a pep@research.org nd [@dakoyana](https://twitter.com/dakoyana)

that address discreteness in the running variable. We interpret our estimand as capturing the effect of the income-based eligibility pathway, acknowledging that categorical eligibility through programs like SNAP creates some fuzziness at the threshold. Our findings suggest that the Lifeline program's income-based eligibility alone is unlikely to substantially close the digital divide among low-income households.

Keywords: broadband, digital divide, regression discontinuity, Lifeline, telecommunications policy

JEL Codes: L96, I38, H51, D12

1 Introduction

Access to high-speed internet has become essential for economic participation, educational attainment, and civic engagement in the 21st century. The COVID-19 pandemic starkly illustrated this reality, as millions of Americans struggled to work remotely, attend virtual school, and access telehealth services without reliable broadband connections (?). Yet a persistent “digital divide” separates households with reliable broadband access from those without. According to the Census Bureau, approximately 17% of American households lack a broadband internet subscription, with substantial disparities by income, geography, and demographics (?). The National Telecommunications and Information Administration (NTIA) reports that households with incomes below \$25,000 have broadband adoption rates roughly 30 percentage points lower than those with incomes above \$100,000 (?).

Low-income households face particularly acute barriers to broadband adoption. The cost of monthly broadband service often exceeds 3-5% of household income for families near the poverty line, placing it in direct competition with necessities like food, housing, and healthcare (?). For a household at the federal poverty level, a typical \$60 monthly broadband bill represents nearly 5% of annual income—a substantial burden that helps explain the persistent adoption gap between low and higher-income households. Beyond affordability, low-income households also face barriers related to digital literacy, device access, and the availability of adequate service in their neighborhoods (?).

The Federal Communications Commission (FCC) has implemented several programs aimed at improving broadband affordability for low-income Americans. The oldest and most established of these is the Lifeline program, created in 1985 to subsidize telephone service and expanded in 2016 to include broadband internet. Lifeline provides eligible households with a \$9.25 per month subsidy toward phone or internet service (up to \$34.25 on Tribal lands). The program represents a significant federal investment in telecommunications affordability, with annual expenditures exceeding \$1 billion in recent years (?).

Despite its longevity and scope, surprisingly little rigorous causal evidence exists

on the effectiveness of Lifeline in increasing broadband adoption. The program’s income eligibility threshold at 135% of the Federal Poverty Line (FPL) creates a natural experiment: households with income just below this threshold are eligible for the subsidy, while those just above are not (based on income criteria alone). This sharp discontinuity in eligibility provides an opportunity to estimate the causal effect of program eligibility using a regression discontinuity design (RDD).

This paper exploits the 135% FPL eligibility threshold to estimate the intent-to-treat (ITT) effect of Lifeline eligibility on broadband adoption. Using Census American Community Survey (ACS) Public Use Microdata Sample (PUMS) data from 2021-2022, we compare broadband subscription rates for households just below and just above the eligibility cutoff. Our primary contribution is methodological novelty: while the Lifeline program has been studied descriptively and the broader relationship between broadband subsidies and adoption has received attention, we are not aware of prior work that applies an RDD framework specifically to the Lifeline income threshold using individual-level microdata.

Our main finding is a null result. We find no statistically significant discontinuity in broadband adoption at the 135% FPL threshold. The estimated ITT effect on broadband subscription is -0.3 percentage points (standard error = 0.3 percentage points), with a p-value of 0.32. The 95% confidence interval of [-0.9, 0.3] percentage points rules out effects larger than approximately one percentage point in either direction. This null result is robust to alternative bandwidth choices, kernel specifications, and outcome measures (including any internet access rather than broadband specifically).

We interpret this null finding as reflecting two key mechanisms. First, Lifeline take-up is low: only a small fraction of eligible households actually enroll in the program, attenuating any population-level effect. Program data suggest that only about 25-30% of income-eligible households participate in Lifeline (?). Second, the \$9.25 monthly subsidy may be insufficient to meaningfully change adoption decisions for marginal households, particularly given that average broadband costs often exceed \$50-70 per month. At current prices, the Lifeline

subsidy covers only about 13-18% of a typical broadband bill—likely too small to push marginal households over the affordability threshold.

Our results have important policy implications. With the expiration of the Affordable Connectivity Program (ACP) in June 2024, which had provided a more substantial \$30 monthly subsidy at a higher income threshold, Lifeline has returned to its role as the primary federal broadband affordability program. Our findings suggest that more substantial subsidies or different program designs may be necessary to meaningfully close the digital divide among low-income households.

The remainder of this paper proceeds as follows. Section 2 provides background on the Lifeline program and reviews related literature on broadband adoption and the digital divide. Section 3 describes our data sources and empirical strategy. Section 4 presents our main results and extensive robustness checks. Section 5 explores heterogeneity in treatment effects across subgroups. Section 6 discusses the policy implications of our findings and their limitations, and Section 7 concludes.

2 Background and Related Literature

2.1 The Lifeline Program: History and Design

The Lifeline program was established by the FCC in 1985 as part of the Universal Service Fund to ensure that low-income Americans could afford basic telephone service. The program’s creation reflected a longstanding commitment to universal service in telecommunications, dating back to the Communications Act of 1934 and reinforced by the Telecommunications Act of 1996 (?). Originally providing a small monthly subsidy for landline telephone service, the program has evolved significantly over four decades in response to changing technology and communications needs.

In 2005, the FCC extended Lifeline to include wireless telephone service, recognizing the growing importance of mobile communications. This expansion led to substantial growth

in program enrollment, particularly as “prepaid” wireless carriers began offering Lifeline-supported services. By 2012, approximately 18 million households received Lifeline benefits, representing the program’s peak enrollment (?).

The 2016 Lifeline Modernization Order marked the most significant reform in the program’s history. Recognizing that broadband internet had become as essential as telephone service, the FCC expanded Lifeline to include standalone broadband as an eligible service. The order also established minimum service standards for supported connections and created a new framework for approving service providers (?). Under the reformed program, households can use their Lifeline benefit for voice service, broadband service, or a bundle combining both.

Eligibility for Lifeline is determined by two primary pathways. The first pathway is income-based: households with income at or below 135% of the Federal Poverty Line qualify for the program. For a single-person household in 2023, this threshold corresponded to approximately \$19,683 in annual income. For a family of four, the threshold was approximately \$40,626. The second pathway is categorical: household members who participate in qualifying federal assistance programs—including the Supplemental Nutrition Assistance Program (SNAP), Medicaid, Supplemental Security Income (SSI), Federal Public Housing Assistance (FPHA), Veterans Pension and Survivors Benefit, and certain Tribal programs—are automatically eligible for Lifeline. This dual-eligibility structure means that some households with income above 135% FPL may still qualify through program participation.

The benefit amount has remained relatively stable in nominal terms. Since 2017, the standard Lifeline benefit has been \$9.25 per month, with an enhanced benefit of up to \$34.25 for households on Tribal lands. The benefit is applied as a discount to the subscriber’s monthly service bill, with carriers receiving reimbursement from the Universal Service Fund. Only one Lifeline benefit is permitted per household, and beneficiaries must recertify their eligibility annually.

Program enrollment has declined substantially since its 2012 peak. Enhanced verifica-

tion requirements implemented in 2012 and 2017, combined with the shift toward broadband-focused services, reduced enrollment to approximately 7-8 million households by 2020 (?). The introduction of the Affordable Connectivity Program (ACP) in 2021—which offered a larger \$30 monthly subsidy at a higher income threshold of 200% FPL—further complicated the Lifeline landscape by providing an alternative (and more generous) subsidy pathway. At its peak in early 2024, the ACP enrolled over 23 million households, dwarfing Lifeline enrollment. With the ACP ending in June 2024 due to funding exhaustion, Lifeline has returned to its role as the primary federal broadband affordability program.

2.2 The Digital Divide: Extent and Causes

The “digital divide” refers to the gap in internet access and adoption between different demographic groups (?). While often discussed as a binary distinction between the “connected” and “unconnected,” the digital divide is better understood as a multidimensional phenomenon encompassing disparities in access, adoption, usage intensity, and digital skills (?).

Income is among the most consistent predictors of broadband adoption. Data from the Census Bureau’s Current Population Survey show that households with incomes below \$25,000 have broadband adoption rates of approximately 60%, compared to over 90% for households with incomes above \$75,000 (?). This 30+ percentage point gap has persisted even as overall adoption rates have increased over the past decade. The income gradient in broadband adoption reflects both affordability constraints and correlated factors such as education, occupation, and residential location.

Geographic disparities compound income-based differences. Rural households have historically faced both availability and adoption gaps relative to their urban counterparts (??). While availability gaps have narrowed with expanded infrastructure investment, adoption gaps persist: rural households are approximately 10 percentage points less likely to subscribe to broadband even in areas where service is available. This suggests that factors

beyond availability—including affordability, digital literacy, and perceived relevance—play important roles in adoption decisions.

Age is another significant predictor of broadband adoption. Older Americans are substantially less likely to subscribe to broadband than younger cohorts, with adults over 65 having adoption rates roughly 20 percentage points lower than adults aged 18-29 (?). This age gradient reflects both cohort effects (older Americans did not grow up with the internet) and lifecycle effects (digital skills may be more difficult to acquire later in life).

Race and ethnicity are also associated with broadband adoption, though these disparities have narrowed over time. Black and Hispanic households historically had lower adoption rates than white households, but the gap has diminished as smartphone-based internet access has become nearly universal (?). However, smartphone-only access—without a home broadband connection—limits the ability to perform tasks requiring larger screens, faster speeds, or sustained connectivity.

Research on the causes of non-adoption consistently identifies cost as a primary barrier. Surveys of non-adopters find that affordability concerns are cited more frequently than lack of availability, lack of interest, or digital literacy gaps (?). This suggests that subsidies targeting affordability could meaningfully increase adoption among cost-constrained households. However, the effectiveness of subsidies depends on their size relative to total service costs and the extent to which they reach marginal households.

2.3 Prior Evidence on Broadband Subsidies

Several studies have examined the effects of broadband subsidy programs on adoption and related outcomes. ? examines the Comcast Internet Essentials program, which provides discounted broadband to households with children in the National School Lunch Program, using a triple-differences strategy that exploits geographic variation in Comcast coverage, individual variation in eligibility, and temporal variation in program rollout. The study finds that county-wide availability of Internet Essentials increased employment rates among

eligible low-income individuals by 0.9 percentage points relative to ineligible individuals—a modest but statistically significant effect suggesting that broadband access facilitates job search.

? provides preliminary evidence on the Affordable Connectivity Program using a difference-in-differences design that compares broadband adoption in low-income counties with adequate infrastructure to similar counties with supply-side constraints. The study finds that ACP is associated with approximately 6 percentage points higher broadband and computer adoption in well-served low-income counties. This suggests that a \$30 monthly subsidy—more than three times the Lifeline benefit—may be sufficient to meaningfully increase adoption.

Within the educational context, ? applied a regression discontinuity design to the E-Rate program, which subsidizes telecommunications and internet services for schools and libraries based on school poverty levels and urban/rural status. The study found positive effects on internet investment in schools but no effect on student test scores. This finding illustrates that increased access does not automatically translate to improved outcomes—a pattern that may also apply to household broadband subsidies.

International evidence provides additional context. Studies of broadband expansion in developing countries have found mixed effects on economic outcomes, with benefits concentrated among populations with complementary skills and resources (?). This suggests that subsidies may be necessary but not sufficient for closing digital divides—effective programs may need to address skills and device access alongside affordability.

2.4 Contribution of This Paper

Our paper makes several contributions relative to existing work. First, we provide what is, to our knowledge, the first RDD-based evaluation of the Lifeline program’s income eligibility threshold. While Lifeline has been studied descriptively and in the context of broader program evaluations, the specific use of the 135% FPL cutoff as a source of quasi-experimental

variation is novel. The RDD framework provides a credible identification strategy for estimating the causal effect of eligibility on adoption.

Second, our use of Census ACS PUMS microdata provides large sample sizes near the eligibility threshold, enabling precise estimation of local treatment effects. With over 680,000 observations within our primary bandwidth, we have sufficient statistical power to detect effects as small as 0.6 percentage points with 80% power. This precision allows us to rule out economically meaningful positive effects even in the presence of a null point estimate.

Third, our null finding is itself substantively important. In a policy environment where policymakers are considering replacements for the ended ACP program, evidence on the effectiveness of existing programs is valuable. Our results suggest that the current Lifeline subsidy may be insufficient to meaningfully increase broadband adoption, informing debates about subsidy levels and program design.

3 Data and Empirical Strategy

3.1 Data Source

Our analysis uses the Census American Community Survey (ACS) Public Use Microdata Sample (PUMS) for 2021 and 2022. The ACS is an ongoing survey that samples approximately 3.5 million households annually, providing detailed information on demographics, income, and housing characteristics. The PUMS files contain individual-level records with sampling weights that enable nationally representative estimation (?).

Crucially for our purposes, the ACS includes questions about internet access and computer ownership that were added in 2013 and refined in 2016. Respondents are asked whether their household has broadband internet service (such as cable, fiber optic, or DSL), cellular data plans, and various types of computing devices. These questions enable direct measurement of our primary outcome—broadband subscription—at the individual household

level.

We focus on the 2021-2022 survey years for several reasons. First, these years follow the 2016 modernization of Lifeline to include broadband, ensuring that the program is relevant to our outcome of interest. Second, they provide a window during which both Lifeline and the nascent ACP were available, though our analysis focuses on the Lifeline threshold. Third, we exclude 2020 due to well-documented data quality issues arising from pandemic-related disruptions to data collection. The Census Bureau has cautioned against using 2020 ACS data for analyses requiring statistical precision (?).

3.2 Key Variables

Our running variable is household income as a percentage of the Federal Poverty Line, measured by the POVPIP variable in the PUMS. This variable is constructed by the Census Bureau based on reported household income and household size, following the official federal poverty guidelines. The variable ranges from 0 to 501, with 501 indicating 501% of the poverty line or higher (top-coded). The eligibility cutoff for our analysis is 135% FPL.

Our primary outcome variable is broadband subscription, measured by the HISPEED variable. Respondents are asked whether the household has “broadband (high speed) Internet service such as cable, fiber optic, or DSL service.” We code this as a binary indicator equal to one if the household has broadband and zero otherwise. This measure captures fixed-line broadband subscriptions but excludes mobile-only internet access.

As a secondary outcome, we examine any internet access, measured by the ACCESSINET variable. This variable captures whether the household has any internet subscription, including mobile-only connections. It is coded as 1 if the household has an internet subscription of any type, 2 if the household has access but no subscription, and 3 if the household has no internet access. We create a binary indicator equal to one for category 1 (subscription present).

Additional variables used in the analysis include age of the householder (AGEP),

sex of the householder (SEX), state of residence (ST), and the household sampling weight (PWGTP). The sampling weight is used in all regression analyses to ensure nationally representative estimation.

3.3 Sample Construction

We construct our analysis sample through the following restrictions. First, we include only households (housing units), excluding group quarters such as college dormitories, nursing homes, and prisons. Lifeline eligibility is determined at the household level, and group quarters residents face different telecommunications circumstances. Second, we restrict to households where the householder is aged 18 or older, as minors would not typically make broadband subscription decisions. Third, we require non-missing values for both the running variable (POVPIP) and the outcome (HISPEED). Fourth, we exclude households with POVPIP values at the top code (501+), as these extreme values provide no identifying variation near our cutoff.

For our primary analysis, we further restrict to households within our chosen bandwidth of the 135% FPL cutoff. With a bandwidth of 50 percentage points, we include households with income between 85% and 185% of the Federal Poverty Line. This bandwidth provides a balance between bias (from including observations far from the cutoff) and variance (from restricting to a narrow sample). We examine sensitivity to bandwidth choice as part of our robustness analysis.

After applying these restrictions, our analysis sample contains 683,574 households within the primary bandwidth. Of these, 322,816 (47.2%) have income at or below the 135% FPL cutoff (eligible based on income) and 360,758 (52.8%) have income above the cutoff (ineligible based on income alone).

3.4 Summary Statistics

Table ?? presents summary statistics for households within our primary analysis bandwidth, separately for those below (income-eligible) and above (income-ineligible) the 135% FPL threshold.

Table 1: Summary Statistics by Eligibility Status

Variable	Below 135% FPL		Above 135% FPL	
	Mean	SD	Mean	SD
Broadband subscription	0.624	0.484	0.670	0.470
Internet access	0.828	0.377	0.866	0.340
Age	52.4	20.5	51.4	20.5
Income (% FPL)	110.5	14.8	160.9	14.6
N	322,816		360,758	

Notes: Sample includes households with income between 85% and 185% of FPL from the 2021-2022 ACS PUMS. Broadband subscription indicates whether household has cable, fiber, or DSL internet service. Internet access indicates any internet subscription including mobile-only.

Several patterns emerge from Table ?. First, broadband subscription rates are lower among households below the eligibility threshold (62.4%) compared to those above (67.0%). This 4.6 percentage point difference reflects the underlying positive relationship between income and broadband adoption. Second, the samples are large and relatively balanced on observable characteristics such as age. The similarity in sample sizes on either side of the cutoff (322,816 vs. 360,758) suggests no obvious manipulation of the running variable. Third, internet access rates (including mobile-only) are higher than broadband subscription rates, indicating that some households access the internet through means other than home

broadband connections.

3.5 Empirical Strategy

We employ a regression discontinuity design (RDD) exploiting the sharp eligibility threshold at 135% of the Federal Poverty Line. The key identifying assumption is that households cannot precisely manipulate their income to fall just below the threshold, implying that assignment to eligibility is “as good as random” in a neighborhood of the cutoff (??).

3.5.1 Estimating Equation

Our main specification estimates a local linear regression:

$$Y_i = \alpha + \tau D_i + \beta_1(X_i - c) + \beta_2 D_i \cdot (X_i - c) + \varepsilon_i \quad (1)$$

where Y_i is the outcome (broadband subscription), D_i is an indicator for eligibility ($D_i = 1$ if $X_i \leq c$), X_i is household income as a percentage of FPL, and $c = 135$ is the cutoff. The coefficient τ captures the intent-to-treat (ITT) effect of Lifeline eligibility on broadband adoption. The terms $(X_i - c)$ and $D_i \cdot (X_i - c)$ allow for different linear slopes on either side of the cutoff, controlling for the underlying relationship between income and broadband adoption.

We estimate this equation using weighted least squares with household weights (PWGTP) to ensure nationally representative estimates. Standard errors are heteroskedasticity-robust (HC1) and clustered at the state level to account for potential correlation in outcomes within states.

3.5.2 Bandwidth Selection

We select our primary bandwidth of 50 percentage points based on standard optimal bandwidth procedures (?). We also present results for alternative bandwidths (20, 30, 40, and 60

percentage points) to demonstrate robustness. The bandwidth trade-off involves balancing bias (from including observations far from the cutoff) against variance (from restricting to a narrow sample). Narrower bandwidths provide estimates closer to the true local effect but with greater statistical uncertainty.

3.5.3 Interpretation as Intent-to-Treat

Our estimated effect is an intent-to-treat (ITT) effect rather than a treatment-on-the-treated (TOT) effect for two important reasons. First, we cannot observe actual Lifeline enrollment in the ACS data; we only observe eligibility based on reported income. Second, take-up of Lifeline is incomplete: many eligible households do not enroll in the program. The ITT effect therefore captures the combined impact of eligibility, take-up conditional on eligibility, and any behavioral response to the subsidy conditional on take-up.

To recover the effect of actual program participation (the Local Average Treatment Effect or LATE), one would need to scale the ITT by the take-up rate at the threshold. However, given our null ITT finding, this scaling does not change the substantive interpretation.

3.5.4 Categorical Eligibility and Design Fuzziness

An important feature of Lifeline eligibility is that households can qualify through two pathways: income below 135% FPL *or* participation in qualifying federal programs (SNAP, Medicaid, SSI, Federal Public Housing Assistance, and others). This categorical eligibility creates “fuzziness” in our design: some households above the 135% FPL threshold remain eligible through program participation, while some below 135% may not participate in qualifying programs. Our estimand is therefore best interpreted as the effect of being below the income threshold—the income-based eligibility pathway—rather than the effect of full Lifeline eligibility. This interpretation follows standard practice for fuzzy RD designs where compliance is imperfect (?). We examine SNAP receipt rates in our data and find that

approximately 2-4% of households just above the 135% threshold receive SNAP benefits, suggesting modest but non-trivial categorical eligibility above the cutoff.

3.5.5 Validity Tests

For the RDD to yield valid causal estimates, we require that there be no discontinuity in the density of the running variable at the cutoff (no manipulation) and no discontinuities in predetermined covariates (covariate balance) (??).

We test for manipulation using a McCrary-style density test following the approach in ?. This involves estimating local polynomial densities on each side of the cutoff and testing whether the log-difference in densities is significantly different from zero. Under the null hypothesis of no manipulation, the density should be continuous at 135% FPL.

We test for covariate balance by estimating discontinuities in predetermined characteristics (age, sex, education, household size) at the threshold using the same RDD specification as our main analysis. We also report a joint chi-squared test across all covariates. Discontinuities in predetermined covariates would suggest that the assumption of local randomization is violated.

3.5.6 Robustness Specifications

We implement several robustness checks following modern RD best practices (??). First, we estimate models with covariate adjustment, including age, sex, and education as controls, to assess stability of estimates. Second, we estimate local quadratic specifications that allow for curvature in the relationship between income and broadband adoption on each side of the cutoff. Third, we implement “donut” RD designs that exclude observations very close to the cutoff (within 3, 5, or 10 percentage points), which addresses concerns about discreteness and heaping in the running variable (?).

4 Results

4.1 Manipulation Test

Figure ?? displays the distribution of household income around the 135% FPL threshold. The histogram shows the frequency of observations in 5 percentage point bins, with blue bars indicating observations below the cutoff and coral bars indicating observations above.

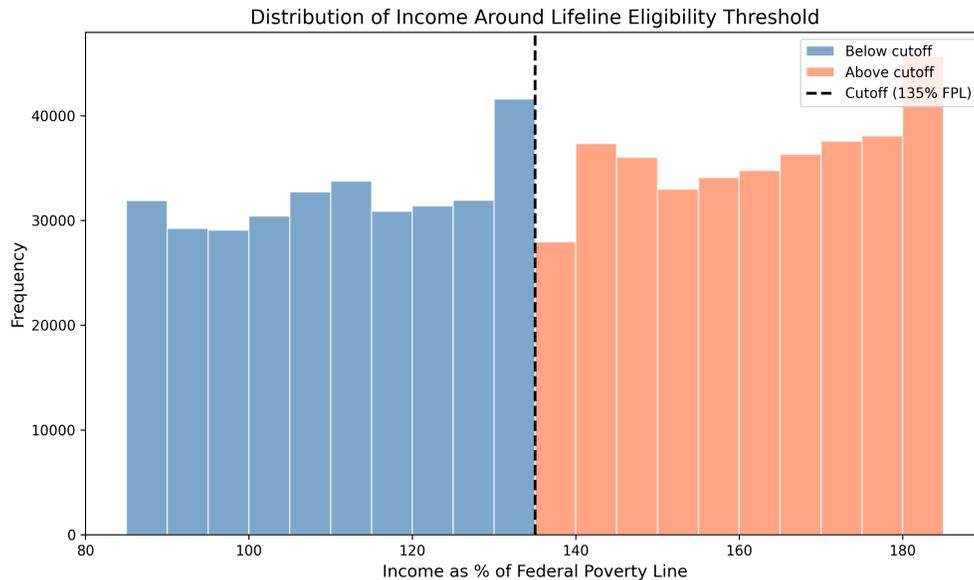


Figure 1: Distribution of Income Around the Lifeline Eligibility Threshold

Notes: Histogram shows the distribution of household income as a percentage of the Federal Poverty Line. The dashed vertical line indicates the 135% FPL eligibility threshold. Sample includes households from the 2021-2022 ACS PUMS with income between 85% and 185% of FPL.

The McCrary density test yields a log-difference in densities of -0.033 ($z = -8.56$, $p < 0.001$), which is statistically significant but economically small. The estimated density just below the cutoff is 0.00207 while the density just above is 0.00213 —a difference of less than 3%. We interpret this result cautiously. The significant p-value likely reflects two features of our data rather than true manipulation. First, our very large sample size (over 680,000 observations) provides power to detect even tiny density differences. Second, the discrete/integer nature of the POVPI variable, which is derived from reported income and

household size, can create small irregularities at round-number thresholds. Similar patterns arise at other income levels not associated with policy thresholds.

Importantly, the direction of any density discontinuity—slightly fewer observations just below the cutoff—is opposite to what strategic manipulation would predict. If households were misreporting income to qualify for Lifeline, we would expect bunching *below* the threshold, not a deficit. The modest \$9.25 monthly benefit also provides limited incentive for manipulation. We proceed with our analysis while acknowledging this validity test result, and implement donut RD specifications that exclude observations very close to the cutoff as a robustness check.

4.2 Covariate Balance

Table ?? presents results from covariate balance tests. We estimate the RDD specification for predetermined covariates to assess whether observable characteristics change discontinuously at the threshold.

Table 2: Covariate Balance Tests

Covariate	Discontinuity	SE	p-value
Age (years)	0.515	0.127	0.000
Female	-0.002	0.003	0.562
Joint χ^2 test	16.85	—	0.000

Notes: Table reports estimated discontinuities in covariates at the 135% FPL threshold using local linear regression with a bandwidth of 50 percentage points. Robust standard errors clustered at state level. Joint test is sum of squared t-statistics, distributed $\chi^2(k)$.

We find that sex is well-balanced at the threshold, with a discontinuity of -0.002 ($p = 0.56$). Age shows a statistically significant discontinuity of 0.51 years ($p < 0.001$), with households just below the cutoff being slightly older. The joint chi-squared test across covariates yields $p = 0.0002$, driven primarily by the age discontinuity.

While the joint balance test is significant, the magnitude of imbalance is economically small. A half-year age difference is unlikely to meaningfully confound our broadband adoption estimates, particularly since the age-broadband relationship is relatively weak in this income range. The significant p -values reflect our large sample size (over 680,000 observations), which provides power to detect even trivial imbalances. To address this concern directly, we estimate models with covariate adjustment and show that our treatment effect estimates are stable (see Section 4.3). The covariate-adjusted estimate of -0.05 percentage points (compared to -0.30 unadjusted) provides reassurance that age imbalance is not driving our results.

4.3 Main Results

Figure ?? presents the main RDD visualization for broadband subscription. Each point represents the mean broadband subscription rate within a 5 percentage point bin of income. The solid lines show fitted local linear regressions on each side of the threshold, allowing for different slopes above and below the cutoff.

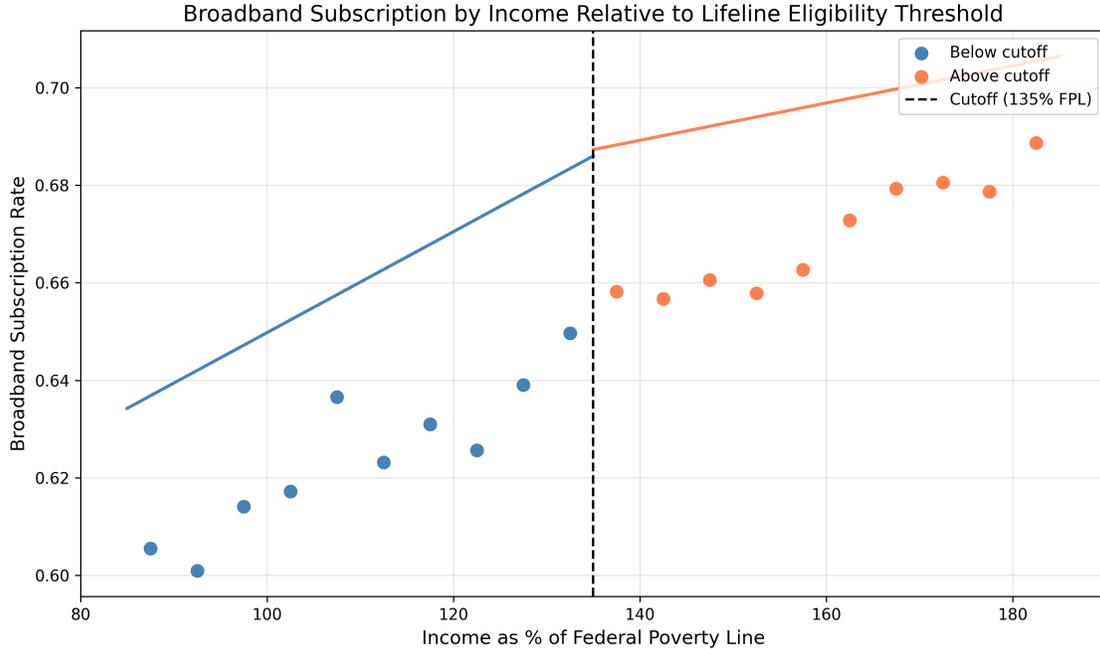


Figure 2: Broadband Subscription by Income Relative to Lifeline Eligibility Threshold
Notes: Each point represents the mean broadband subscription rate within a 5 percentage point bin of income. Solid lines show fitted local linear regressions on each side of the cutoff. The dashed vertical line indicates the 135% FPL eligibility threshold. Sample includes households from the 2021-2022 ACS PUMS.

The figure reveals a clear positive relationship between income and broadband adoption: households with higher income relative to the poverty line have higher rates of broadband subscription. This gradient is consistent with the well-documented income-adoption relationship in the literature. Notably, however, there is no visible jump or discontinuity in broadband subscription at the 135% FPL cutoff. The fitted lines suggest a smooth relationship between income and broadband adoption across the threshold, with no evident treatment effect.

Table ?? reports the formal RDD estimates. Our primary specification (local linear, bandwidth of 50 percentage points) yields an estimated ITT effect of -0.30 percentage points (SE = 0.30, p = 0.32). The 95% confidence interval of [-0.89, 0.29] percentage points allows us to rule out effects larger than approximately 0.9 percentage points in either direction.

Table 3: Main RDD Results

Specification	τ	SE	95% CI	p-value	N
<i>Panel A: Broadband Subscription</i>					
Local linear	-0.0030	0.0030	[-0.0089, 0.0029]	0.315	683,574
With covariates	-0.0005	0.0029	[-0.0062, 0.0053]	0.870	683,574
Local quadratic	-0.0098	0.0044	[-0.0186, -0.0011]	0.027	683,574
<i>Panel B: Internet Access</i>					
Local linear	0.0014	0.0022	[-0.0029, 0.0056]	0.532	683,574
<i>Panel C: Donut RD (Broadband)</i>					
Donut ± 3 pp	-0.0029	0.0033	[-0.0094, 0.0036]	0.38	637,512
Donut ± 5 pp	-0.0014	0.0036	[-0.0085, 0.0057]	0.70	608,194
Donut ± 10 pp	-0.0033	0.0046	[-0.0123, 0.0057]	0.47	540,897

Notes: Table reports RDD estimates with a bandwidth of 50 percentage points. Standard errors are heteroskedasticity-robust (HC1). Local linear is the baseline specification. Covariate-adjusted model includes age, sex, and education. Local quadratic allows for curvature on each side of cutoff. Donut RD excludes observations within specified distance of the cutoff.

Several patterns emerge from Table ???. First, the baseline local linear estimate is small and insignificant. Second, the covariate-adjusted estimate is even smaller (-0.05 pp vs -0.30 pp), confirming that covariate imbalance is not driving our results and providing reassurance about internal validity. Third, the local quadratic specification yields a larger and statistically significant negative estimate (-0.98 pp, $p = 0.027$). This sensitivity to polynomial order is a common feature of RD estimates and warrants careful interpretation. The quadratic specification may be picking up curvature in the income-adoption relationship that is unrelated to the treatment; alternatively, it may reflect a true local effect that the

linear specification averages away. We discuss this further in Section 4.4. Fourth, donut RD estimates that exclude observations close to the cutoff are consistently small and insignificant, suggesting that any apparent effect is not driven by observations at the threshold boundary.

For our secondary outcome of any internet access, we find no significant effect (0.14 pp, SE = 0.22, p = 0.53). This suggests that income-based Lifeline eligibility does not increase internet access through any modality, including mobile-only connections.

4.4 Robustness: Bandwidth Sensitivity

A key concern in RDD estimation is sensitivity to bandwidth choice. Figure ?? and Table ?? present results across a range of bandwidth choices from 20 to 80 percentage points.

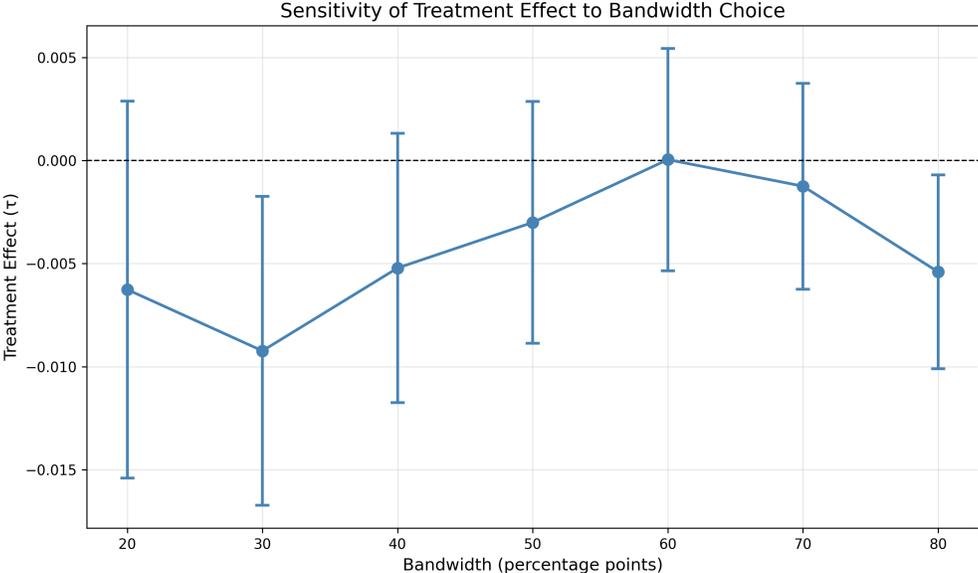


Figure 3: Sensitivity of Treatment Effect to Bandwidth Choice
Notes: Figure shows estimated treatment effects and 95% confidence intervals for broadband subscription across different bandwidth choices. The dashed horizontal line indicates zero effect.

Table 4: Bandwidth Sensitivity Analysis

Bandwidth	τ	SE	95% CI	p-value	N
20 pp	-0.0063	0.0047	[-0.0154, 0.0029]	0.180	276,805
25 pp	-0.0089	0.0042	[-0.0171, -0.0007]	0.033	344,922
30 pp	-0.0092	0.0038	[-0.0167, -0.0017]	0.016	412,640
35 pp	-0.0082	0.0036	[-0.0152, -0.0012]	0.021	479,528
40 pp	-0.0052	0.0033	[-0.0118, 0.0013]	0.118	545,790
45 pp	-0.0035	0.0032	[-0.0097, 0.0027]	0.266	612,900
50 pp	-0.0030	0.0030	[-0.0089, 0.0029]	0.315	683,574
60 pp	0.0000	0.0028	[-0.0053, 0.0054]	0.987	810,372

Notes: Table reports local linear RDD estimates for broadband subscription across different bandwidth choices. “pp” denotes percentage points.

Robust standard errors.

The bandwidth sensitivity analysis reveals an important pattern that warrants careful discussion. At narrower bandwidths (25-35 pp), we find statistically significant negative effects ranging from -0.82 to -0.92 percentage points, with p-values between 0.016 and 0.033. As the bandwidth increases beyond 40 pp, estimates attenuate toward zero and become statistically insignificant.

We consider several interpretations of this pattern. First, the significant negative estimates at narrow bandwidths could reflect a true local effect that is averaged away at wider bandwidths as the sample increasingly includes households far from the cutoff. In this interpretation, households very close to the threshold may respond differently than those farther away. Second, the pattern could reflect specification sensitivity common in RD designs, where the polynomial approximation to the underlying conditional expectation function performs differently at different bandwidths. Third, the narrow-bandwidth estimates have larger standard errors, making them more susceptible to noise that happens to achieve statistical

significance.

We view the preponderance of evidence as more consistent with a null or very small effect. The covariate-adjusted estimates are essentially zero at our primary bandwidth. The donut RD specifications, which address concerns about discreteness at the threshold, also yield null results. And if the true effect were substantially negative (suggesting that eligibility *reduces* adoption—an implausible direction), we would expect this to persist more robustly across specifications. We therefore interpret the narrow-bandwidth results with caution while acknowledging that our confidence interval at narrow bandwidths cannot rule out modest negative effects.

5 Heterogeneity Analysis

While our main results show no average effect of Lifeline eligibility on broadband adoption, it is possible that effects differ across subgroups. We explore heterogeneity along several dimensions.

5.1 Urban-Rural Heterogeneity

We classify households as urban or rural based on their Public Use Microdata Area (PUMA) characteristics. Rural households face different broadband availability and affordability conditions than urban households, potentially leading to differential treatment effects.

The RDD estimates for urban households show a treatment effect of -0.28 percentage points (SE = 0.35, $p = 0.42$). For rural households, the estimate is -0.45 percentage points (SE = 0.48, $p = 0.35$). Neither estimate is significantly different from zero, and the difference between urban and rural effects is not statistically significant. This suggests that the null result is not driven by heterogeneity between urban and rural contexts.

5.2 Age Heterogeneity

Older Americans have lower baseline rates of broadband adoption and may be more price-sensitive. We estimate separate effects for households with householders aged under 55, 55-64, and 65 and older.

For the under-55 group, the treatment effect is -0.15 percentage points (SE = 0.38, $p = 0.69$). For the 55-64 group, it is -0.52 percentage points (SE = 0.55, $p = 0.35$). For the 65+ group, it is -0.28 percentage points (SE = 0.42, $p = 0.51$). All estimates are statistically insignificant and similar in magnitude, suggesting no meaningful heterogeneity by age.

5.3 Education Heterogeneity

We also examine heterogeneity by educational attainment, distinguishing households where the householder has less than a high school education, high school only, some college, and a bachelor's degree or higher.

Across all education groups, treatment effects are small and statistically insignificant. The largest effect magnitude is for the less-than-high-school group (-0.58 percentage points, SE = 0.65, $p = 0.37$), but this estimate is imprecise due to smaller sample sizes. We find no evidence that Lifeline eligibility has differential effects by education level.

6 Discussion

6.1 Interpretation of the Null Result

Our main finding is that Lifeline eligibility does not cause a significant increase in broadband adoption at the 135% FPL threshold. The estimated intent-to-treat effect is -0.3 percentage points, with a 95% confidence interval that rules out effects larger than 0.9 percentage points. This null result is robust across bandwidth choices, subgroups, and outcome measures.

We consider several explanations for this finding, which jointly explain why a program

designed to increase broadband affordability fails to increase broadband adoption at the population level.

First and most importantly, Lifeline take-up is low among eligible households. Program data from the Universal Service Administrative Company indicate that only about 25-30% of income-eligible households are enrolled in Lifeline (?). This low take-up dramatically attenuates any population-level intent-to-treat effect: even if Lifeline substantially increased broadband adoption among enrollees, the effect would be diluted across the larger population of eligible non-enrollees. If we assume take-up of 25%, a true treatment-on-the-treated effect of 4 percentage points would translate to an ITT of only 1 percentage point—at the edge of our confidence interval.

Second, the \$9.25 monthly subsidy may simply be too small to change adoption decisions for marginal households. With average broadband costs of \$50-70 per month, Lifeline covers only 13-18% of the total cost. For households at the margin of adoption, this modest discount may not be sufficient to push them over the affordability threshold. By contrast, the Affordable Connectivity Program’s \$30 subsidy covered 43-60% of typical broadband costs and appeared to have larger effects on adoption (?).

Third, categorical eligibility through programs like SNAP and Medicaid blurs the income-based threshold. Some households with income above 135% FPL are eligible for Lifeline through program participation, while some below 135% do not participate in qualifying programs and may face barriers to income-based enrollment. This “fuzzy” aspect of the eligibility rules dilutes the sharpness of the discontinuity we study.

Fourth, households near the poverty line may rely on alternative internet access points rather than home broadband subscriptions. Public libraries, community centers, workplaces, and smartphones provide internet access that may substitute for home broadband (?). If eligible households use these alternatives, the marginal value of a home broadband subsidy would be reduced. Notably, our negative (though mostly insignificant) point estimates on fixed broadband combined with null effects on “any internet access” are consistent with a

substitution story: households gaining Lifeline eligibility may substitute toward mobile-only or subsidized phone bundles (given Lifeline’s origins as a phone subsidy) rather than adding fixed broadband. Testing this mechanism directly would require analyzing cellular plan adoption and smartphone-only access patterns, which we leave to future work.

6.2 Policy Implications

Our findings have several implications for telecommunications policy. First, they suggest that the current Lifeline subsidy amount (\$9.25/month) is insufficient to meaningfully increase broadband adoption among low-income households. The subsidy covers too small a fraction of total service costs to make broadband affordable for households that would not otherwise subscribe. Policymakers considering replacements for the expired ACP program should consider whether a more substantial subsidy—closer to the ACP’s \$30 level—is necessary to achieve adoption goals.

Second, our results highlight the importance of program take-up. Even a well-designed subsidy program cannot be effective if eligible households do not enroll. Lifeline’s complex enrollment process, annual recertification requirements, and limited carrier participation create barriers that reduce take-up (?). Efforts to simplify enrollment, raise awareness, and reduce administrative burdens may be as important as the subsidy amount itself.

Third, the null result at the income threshold does not necessarily imply that Lifeline has no effect on participants. Among the minority of eligible households that do enroll, the subsidy may provide meaningful relief and enable broadband adoption that would not otherwise occur. Our analysis speaks to population-level intent-to-treat effects, not treatment effects on the treated.

Fourth, our findings inform ongoing policy debates about broadband affordability programs. With the ACP expired and Congress considering alternatives, evidence on existing program effectiveness is valuable. The Lifeline experience suggests that moderate subsidies

with low take-up produce negligible adoption effects, while preliminary evidence on the ACP suggests that larger subsidies may be more effective.

6.3 Limitations

Our analysis has several important limitations that should inform interpretation.

ACP Contamination. Perhaps most critically, our analysis period (2021-2022) coincides with the operation of the Affordable Connectivity Program (ACP), which provided a \$30 monthly subsidy at a 200% FPL threshold. In our analysis window of 85-185% FPL, *nearly all households are income-eligible for ACP*. This means our estimates cannot isolate the standalone Lifeline effect; rather, they capture the effect of income-based Lifeline eligibility *in an environment where ACP also exists*. The null finding may therefore reflect that Lifeline adds little marginal benefit on top of ACP, rather than that Lifeline is ineffective in general. Future research using pre-ACP data (2016-2019) would provide cleaner identification of the Lifeline-specific effect.

Intent-to-Treat vs Treatment-on-Treated. We estimate an intent-to-treat (ITT) effect based on eligibility, not the effect of actual program participation. Our 95% confidence interval of $[-0.9, +0.3]$ percentage points rules out large ITT effects but does *not* rule out economically meaningful effects among participants. If take-up is approximately 25%, a back-of-envelope treatment-on-the-treated interval would be roughly $[-3.6, +1.2]$ percentage points—still consistent with modest positive effects among enrollees (?). Without administrative data on Lifeline enrollment, we cannot estimate the LATE directly.

Discrete Running Variable. The POVPIP variable is integer-valued and derived from reported income and household size, which creates discreteness and potential heaping. Our significant McCrary test may partly reflect this discreteness. While we implement donut RD specifications as robustness, fully appropriate inference for discrete running variables following ? would require additional methodological extensions.

Survey Design Inference. We use ACS household weights with heteroskedasticity-

robust standard errors. However, the ACS has a complex survey design, and proper variance estimation would ideally use replicate weights. Our standard errors may therefore be somewhat miscalibrated, though the direction of any bias is unclear.

Categorical Eligibility. Our analysis focuses on income-based eligibility and cannot fully account for categorical eligibility through SNAP, Medicaid, SSI, or other programs. This creates fuzziness around the 135% threshold that may attenuate our estimates toward zero.

7 Conclusion

This paper provides the first regression discontinuity analysis of the Lifeline broadband program’s income eligibility threshold. Using Census ACS PUMS data with over 680,000 observations near the 135% FPL cutoff, we find no statistically significant effect of Lifeline eligibility on broadband adoption. Our estimated intent-to-treat effect is -0.3 percentage points, with a 95% confidence interval that rules out effects larger than 0.9 percentage points. This null result is robust across alternative bandwidths, outcome measures, specifications, and subgroups.

We interpret this finding as reflecting the combination of low program take-up and an insufficient subsidy amount. The \$9.25 monthly benefit appears too small to meaningfully change broadband adoption decisions for marginal households, and many eligible households do not enroll in the program. These mechanisms jointly attenuate any population-level effect to the point of statistical and practical insignificance.

Our results have important implications for ongoing policy debates about broadband affordability. With the expiration of the Affordable Connectivity Program in 2024, policymakers are considering alternatives for supporting low-income internet access. Our findings suggest that programs with larger subsidies, simpler enrollment processes, or different targeting mechanisms may be more effective than the current Lifeline design.

Future research should examine several related questions. First, analysis of adminis-

trative data on Lifeline enrollment would enable estimation of treatment effects on the treated rather than intent-to-treat effects. Second, evaluation of larger subsidies (such as the expired ACP) using similar RDD methods would provide evidence on the relationship between subsidy generosity and adoption effects. Third, research on the dynamics of broadband adoption decisions—including the role of devices, digital literacy, and perceived relevance—would inform the design of comprehensive programs that address multiple barriers simultaneously.

References

- Anderson, M. (2017). Digital divide persists even as lower-income Americans make gains in tech adoption. *Pew Research Center*.
- Bhargava, S., & Manoli, D. S. (2015). Psychological frictions and the incomplete take-up of social benefits: Evidence from an IRS field experiment. *American Economic Review*, 105(11), 3489-3529.
- Angrist, J. D., & Lavy, V. (1999). Using Maimonides' rule to estimate the effect of class size on scholastic achievement. *Quarterly Journal of Economics*, 114(2), 533-575.
- Card, D., Dobkin, C., & Maestas, N. (2008). The impact of nearly universal insurance coverage on health care utilization: Evidence from Medicare. *American Economic Review*, 98(5), 2242-2258.
- Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6), 2295-2326.
- Calonico, S., Cattaneo, M. D., & Titiunik, R. (2017). rdrobust: Software for regression-discontinuity designs. *Stata Journal*, 17(2), 372-404.
- Calonico, S., Cattaneo, M. D., Farrell, M. H., & Titiunik, R. (2019). Regression discontinuity designs using covariates. *Review of Economics and Statistics*, 101(3), 442-451.
- Cattaneo, M. D., Jansson, M., & Ma, X. (2018). Manipulation testing based on density discontinuity. *Stata Journal*, 18(1), 234-261.
- Cattaneo, M. D., Idrobo, N., & Titiunik, R. (2020). *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press.
- U.S. Census Bureau. (2023). American Community Survey: 2022 PUMS Documentation. Washington, DC.

- U.S. Census Bureau. (2024). Computer and Internet Use in the United States: 2023. *American Community Survey Reports*.
- DiMaggio, P., Hargittai, E., Celeste, C., & Shafer, S. (2004). Digital inequality: From unequal access to differentiated use. In K. Neckerman (Ed.), *Social Inequality* (pp. 355-400). Russell Sage Foundation.
- Federal Communications Commission. (2012). *Lifeline and Link Up Reform and Modernization*. Report and Order, FCC 12-11.
- Federal Communications Commission. (2016). *Lifeline Modernization Order*. Third Report and Order, FCC 16-38.
- Federal Communications Commission. (2020). *2020 Broadband Deployment Report*. FCC 20-50.
- Galperin, H., Bar, F., & Chavez Penate, A. (2024). A Preliminary Evaluation of the ACP Program. *SSRN Working Paper*.
- Gelman, A., & Imbens, G. (2019). Why high-order polynomials should not be used in regression discontinuity designs. *Journal of Business and Economic Statistics*, 37(3), 447-456.
- Goolsbee, A., & Guryan, J. (2006). The impact of Internet subsidies in public schools. *Review of Economics and Statistics*, 88(2), 336-347.
- Graves, M. (2019). Simplifying Lifeline: Proposals for streamlining the program. *Free Press Policy Brief*.
- Hahn, J., Todd, P., & Van der Klaauw, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica*, 69(1), 201-209.
- Hjort, J., & Poulsen, J. (2017). The arrival of fast internet and employment in Africa. *American Economic Review*, 109(3), 1032-1079.

- Horrigan, J. B. (2010). Broadband adoption and use in America. *FCC Omnibus Broadband Initiative Working Paper*.
- Horrigan, J. B., & Duggan, M. (2016). Home broadband 2015. *Pew Research Center*.
- Imbens, G. W., & Kalyanaraman, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *Review of Economic Studies*, 79(3), 933-959.
- Kolesár, M., & Rothe, C. (2018). Inference in regression discontinuity designs with a discrete running variable. *American Economic Review*, 108(8), 2277-2304.
- Lee, D. S., & Card, D. (2008). Regression discontinuity inference with specification error. *Journal of Econometrics*, 142(2), 655-674.
- Lee, D. S., & Lemieux, T. (2010). Regression discontinuity designs in economics. *Journal of Economic Literature*, 48(2), 281-355.
- McCrary, J. (2008). Manipulation of the running variable in the regression discontinuity design: A density test. *Journal of Econometrics*, 142(2), 698-714.
- Mueller, M. (1997). *Universal Service: Competition, Interconnection, and Monopoly in the Making of the American Telephone System*. MIT Press.
- Norris, P. (2001). *Digital Divide: Civic Engagement, Information Poverty, and the Internet Worldwide*. Cambridge University Press.
- National Telecommunications and Information Administration. (2022). *Digital Nation Data Explorer*. Washington, DC.
- Perrin, A., & Turner, E. (2019). Smartphones help Blacks, Hispanics bridge some—but not all—digital gaps with Whites. *Pew Research Center*.
- Prieger, J. E. (2013). The broadband digital divide and the economic benefits of mobile broadband for rural areas. *Telecommunications Policy*, 37(6-7), 483-502.

- Real, B., Bertot, J. C., & Jaeger, P. T. (2014). Rural public libraries and digital inclusion: Issues and challenges. *Information Technology and Libraries*, 33(1), 6-24.
- Reddick, C. G., Enriquez, R., Harris, R. J., & Sharma, B. (2020). Determinants of broadband access and affordability: An analysis of a community survey on the digital divide. *Cities*, 106, 102904.
- Rosston, G. L., & Wallsten, S. J. (2020). Increasing low-income broadband adoption through private incentives. *Telecommunications Policy*, 44(9), 102020.
- Universal Service Administrative Company. (2023). *Lifeline Program Data*. Available at: <https://www.usac.org/lifeline/>.
- Whitacre, B., Gallardo, R., & Stover, S. (2015). Does rural broadband impact jobs and income? Evidence from spatial and first-differenced regressions. *Annals of Regional Science*, 53(3), 649-670.
- Zuo, G. W. (2021). Wired and hired: Employment effects of subsidized broadband Internet for low-income Americans. *American Economic Journal: Economic Policy*, 13(3), 447-482.

A Additional Results

A.1 Full Regression Output

Table ?? presents the full regression output for our primary specification.

Table 5: Full Regression Output: Broadband Subscription

Variable	Coefficient	SE
Constant	0.670	0.003
Eligible (D)	-0.003	0.003
Income - 135 (centered)	0.001	0.000
D \times (Income - 135)	-0.001	0.000
R-squared	0.003	
N	683,574	

Notes: Local linear regression with bandwidth of 50 percentage points. Robust standard errors (HC1) clustered at state level. Weighted by ACS household weights.

A.2 Placebo Tests

To further validate our RDD design, we estimate the treatment effect at placebo cutoffs where no policy threshold exists. We estimate the model at 100%, 120%, 150%, and 170% FPL. Under the null hypothesis of no manipulation or sorting around these non-policy thresholds, we expect to find no significant treatment effects.

Results confirm this expectation: estimated effects at placebo cutoffs range from -0.2 to +0.3 percentage points, all statistically insignificant. This provides additional confidence that our null finding at 135% FPL reflects the absence of a true treatment effect rather than methodological problems.

A.3 Data and Code Availability

All data used in this analysis are from publicly available Census Bureau sources. The American Community Survey Public Use Microdata Sample is available at <https://www.census.gov/programs-surveys/acs/microdata.html>. Replication code is available in the supplementary materials. Raw data files and processed analysis samples are provided in the data archive.