

Does Social Security Eligibility Reduce Living Alone?

Evidence from a Regression Discontinuity Design*

Autonomous Policy Evaluation Project

January 2026

Abstract

Social isolation and loneliness among older adults is recognized as a public health crisis, with living alone being the strongest demographic predictor of loneliness. This paper examines whether Social Security eligibility at age 62—the earliest age at which Americans can claim retirement benefits—affects the probability of living alone. Using American Community Survey microdata from 2016–2022 and a regression discontinuity design, I find that Social Security eligibility *decreases* the probability of living alone by 0.67 percentage points (95% CI: -1.26 to -0.08). This effect is concentrated among unmarried individuals (-2.93 pp, $p < 0.001$) and men (-1.11 pp, $p = 0.01$). A strong first stage confirms that Social Security income receipt increases by 13.4 percentage points at the eligibility threshold. These findings suggest that income security programs may facilitate household consolidation and potentially reduce social isolation among older adults.

Keywords: Social Security, living alone, loneliness, regression discontinuity, retirement

*This paper was generated as part of the Autonomous Policy Evaluation Project (APEP). Data from American Community Survey Public Use Microdata Sample (PUMS), 2016-2022. Replication code available at <https://github.com/dakoyana/auto-policy-evals>.

JEL Codes: H55, J14, J26, I38

1 Introduction

Social isolation and loneliness among older adults has emerged as a major public health concern. In 2023, the U.S. Surgeon General declared loneliness an epidemic, noting that social disconnection increases mortality risk by 26% and is associated with higher rates of dementia, heart disease, stroke, and depression (?). Among the elderly, living alone is one of the strongest predictors of social isolation and loneliness (?).

Approximately 28% of Americans aged 60 and older live alone, a figure that has risen steadily over recent decades (?). Understanding what factors enable or prevent solo living is thus important for policies aimed at addressing the loneliness epidemic. Economic constraints likely play a central role: housing is expensive, and many older adults may live with others out of financial necessity rather than choice. Conversely, income security might enable independence—allowing individuals to maintain solo living arrangements they prefer.

This paper examines how Social Security eligibility affects living arrangements among older Americans. At age 62, individuals become eligible for early Social Security retirement benefits, providing a guaranteed income stream equal to roughly 70% of their full retirement amount. This creates a sharp discontinuity in financial resources that can be exploited for causal identification.

The direction of the effect is theoretically ambiguous. On one hand, Social Security income could *increase* living alone by providing financial independence—allowing individuals to maintain or establish solo households they couldn't otherwise afford. On the other hand, Social Security income could *decrease* living alone by enabling household consolidation: facilitating moves to be near family, affording housing large enough for co-residence, or providing financial stability that makes partnership formation more attractive.

Using American Community Survey (ACS) microdata from 2016–2022 and a regression discontinuity design centered at the age-62 eligibility threshold, I find that Social Security

eligibility *decreases* the probability of living alone by 0.67 percentage points (on a baseline of approximately 19%), a statistically significant effect ($p = 0.026$). This reduction is concentrated among unmarried individuals, where the effect reaches 2.93 percentage points ($p < 0.001$), and among men (-1.11 pp, $p = 0.01$).

The first stage is strong: Social Security income receipt increases discontinuously by 13.4 percentage points at age 62 ($p < 0.001$), confirming that the eligibility threshold creates meaningful variation in claiming behavior. Covariate balance tests show no discontinuities in predetermined characteristics (sex, race, education, disability) at the threshold, supporting the validity of the design.

These findings contribute to several literatures. First, I add to the extensive literature on Social Security's effects on labor supply and retirement timing (???) by documenting effects on a novel outcome: household composition. Second, I contribute to the growing literature on the determinants and consequences of social isolation in aging populations (???). Third, I provide evidence relevant to policy debates about retirement security and its broader social implications.

The finding that Social Security eligibility *reduces* living alone is perhaps surprising. It suggests that for many older Americans approaching retirement age, the constraint on co-residence is not preference for independence but rather financial barriers that income security helps overcome. This has implications for how policymakers think about the social benefits of retirement income programs: beyond direct welfare effects on recipients, these programs may facilitate social connection through household consolidation.

The remainder of this paper is organized as follows. Section 2 provides background on Social Security eligibility rules and living arrangements among older adults. Section 3 describes the data and empirical strategy. Section 4 presents the main results, robustness checks, and heterogeneity analyses. Section 5 discusses mechanisms and implications. Section 6 concludes.

2 Background

2.1 Social Security Eligibility at Age 62

The Social Security retirement program provides monthly benefits to retired workers who have accumulated sufficient work history (generally 40 quarters of covered employment). While full retirement age (FRA) is currently 67 for those born in 1960 or later, individuals can claim early retirement benefits beginning at age 62.

Early claiming results in a permanently reduced benefit: those claiming at 62 receive approximately 70% of what they would receive if they waited until FRA. Despite this actuarial reduction, a substantial fraction of Americans claim at or near age 62. In 2021, approximately 25% of new retired-worker awards went to individuals aged exactly 62 (?).

The age-62 threshold is federally uniform and creates a sharp discontinuity in eligibility. Before age 62, individuals cannot receive Social Security retirement benefits regardless of their work history. At exactly 62, eligibility switches on, providing a guaranteed income stream that can substantially change household economic circumstances.

For a worker with average lifetime earnings, the early retirement benefit in 2022 was approximately \$1,500 per month (\$18,000 annually). This represents a meaningful income shock, particularly for those without other retirement income sources.

2.2 Living Arrangements Among Older Adults

Living arrangements in later life are determined by a complex set of factors including health status, marital history, income, housing costs, family availability, and preferences for independence. Among Americans aged 60–64, approximately 18–20% live alone, with the remainder living with spouses, partners, adult children, or others (?).

Living alone is strongly associated with loneliness and social isolation. While living alone

is not synonymous with being lonely—many individuals who live alone maintain robust social networks—the correlation is substantial. Older adults living alone report higher rates of loneliness, depression, and social isolation compared to those in multi-person households (?).

The decision to live alone versus with others involves trade-offs between privacy/independence and companionship/support. Economic factors matter: maintaining a solo household is expensive, requiring the individual to bear full housing costs. For those with limited income, co-residence may be a financial necessity rather than a choice.

2.3 Related Literature

This paper contributes to several distinct literatures that have developed largely in parallel. The first is the extensive literature on Social Security’s effects on labor supply and retirement timing. ? document how the rules governing Social Security benefits affect claiming behavior and work decisions around retirement age. ? estimates a structural model of retirement that highlights Social Security’s central importance in retirement decisions, showing that benefit levels and eligibility rules substantially affect when workers exit the labor force. ? uses quasi-experimental variation from changes in full retirement age to estimate labor supply effects, finding that each year increase in FRA delays retirement by several months. ? examine how workers respond to incentives to delay claiming, finding meaningful responses to both benefit increases and rule changes.

Within this Social Security literature, regression discontinuity designs exploiting age thresholds have been used to study various outcomes. The age-65 Medicare eligibility threshold has been extensively studied for its effects on health insurance coverage, health care utilization, and health outcomes. The age-62 threshold has received less attention, though several studies have examined its effects on labor supply and retirement timing. This paper extends the RDD literature to a novel outcome—living arrangements—that has not been

examined using these methods.

Much less is known about Social Security's effects on non-labor market outcomes. ? examine how Social Security income affects health insurance coverage and find that SS receipt reduces the probability of being uninsured among the elderly. ? study effects on poverty, documenting that SS benefits are the primary income source for most retired households and substantially reduce elder poverty rates. An important related literature has examined how pension income affects elderly living arrangements using historical variation. ? finds that Union Army pensions increased the probability that elderly veterans lived independently rather than with family, providing early causal evidence that public transfers affect household structure. ? extends this analysis to Old Age Assistance programs, showing that income support enabled older nonmarried women to maintain independent households. ? documents the long-run decline in intergenerational coresidence, attributing much of the trend to rising incomes and public pension expansions. This paper contributes to this literature by providing contemporary quasi-experimental evidence using the age-62 eligibility threshold.

A second relevant literature examines the determinants of living alone in later life. ? document the dramatic rise in solo living among the elderly over the 20th century, from under 10% in 1900 to over 25% by the 1990s. This trend has been attributed to rising incomes, expansion of Social Security and pension programs, changes in housing markets, and shifting preferences toward independence. ? examine how economic resources affect living arrangements, finding that increases in Social Security benefits contributed to the rise in elderly living alone by enabling financial independence from family. ? study racial and ethnic differences in solo living, documenting that older Whites are substantially more likely to live alone than Blacks, Hispanics, or Asians, a pattern attributed to differences in family structure, economic resources, and cultural norms around co-residence.

These studies of living arrangement determinants are generally descriptive or rely on correlational evidence. Cross-sectional comparisons of individuals with different income levels

confound the effect of resources with selection: those who live alone may differ from those who co-reside in unobserved ways that also affect income. Longitudinal studies that track individuals over time can control for fixed characteristics but still face concerns about time-varying confounds. The causal effect of income security on living arrangements remains largely unknown.

A third relevant literature examines social isolation and loneliness in aging populations. ? provide a comprehensive review of the consequences of loneliness, documenting associations with depression, cognitive decline, cardiovascular disease, and mortality. ? conduct a meta-analysis finding that social isolation increases mortality risk by 26%, comparable to the risk from obesity or smoking. ? reviews the state of knowledge on social isolation among older adults, emphasizing that it is both common and consequential but underassessed in clinical settings. ? examine the relationship between social isolation, loneliness, and health, emphasizing that the two concepts—objective isolation and subjective loneliness—are distinct though correlated.

Living alone is consistently identified as a strong predictor of social isolation and loneliness. While not all who live alone are lonely, and some who live with others are profoundly lonely, the correlation is substantial. Older adults living alone have fewer daily social interactions, are at higher risk of going days without speaking to anyone, and report higher levels of loneliness on validated scales. Understanding factors that affect living alone is thus directly relevant to the policy challenge of addressing elder loneliness.

This paper contributes to these literatures by bringing quasi-experimental methods to bear on the question of whether income security programs affect living arrangements. The regression discontinuity design provides credible causal identification of the effect of Social Security eligibility on the probability of living alone, filling a gap in the literature on Social Security's broader effects and on the determinants of elderly living arrangements. I follow best practices for RDD established by ? and ?, including graphical presentation of outcomes

by the running variable, local linear estimation, bandwidth sensitivity checks, and tests for manipulation of the running variable (?). The design is similar in spirit to age-threshold RDD studies in health economics, such as ?'s analysis of Medicare eligibility at age 65. The findings have implications for all three literatures: they document a novel non-labor-market effect of Social Security, provide causal evidence on how economic resources affect living arrangements, and suggest that income support programs may have unintended benefits for social connection.

3 Data and Empirical Strategy

3.1 Data

I use the American Community Survey (ACS) Public Use Microdata Sample (PUMS) for survey years 2016–2022. The ACS is an annual survey of approximately 3.5 million households conducted by the U.S. Census Bureau, providing detailed information on demographics, economic characteristics, and housing.

The key variables for this analysis span several domains. The running variable for the RDD is age (AGEP), which is reported in whole years. While exact age or month of birth would provide more precise measurement, the integer age in ACS still allows for regression discontinuity analysis with appropriate methods for discrete running variables.

The primary outcome variable is constructed from the number of persons in the household (NP). I define living alone as NP equals 1, meaning the respondent is the sole occupant of the household. This is a standard measure of solo living used throughout the literature on elderly living arrangements and provides a clean binary indicator of household composition.

For the first stage of the analysis, I use Social Security income variables (SSP and SSIP) to identify whether respondents receive Social Security retirement benefits or Supplemental

Security Income. Receipt of any Social Security income indicates that the individual has claimed benefits, which is the relevant margin induced by the age-62 eligibility threshold.

Demographic covariates include sex, race (RAC1P), educational attainment (SCHL), marital status (MAR), and disability status (DIS). These variables serve multiple purposes: they are included as controls in the main specification to improve precision, and they are tested for covariate balance at the threshold to assess the validity of the RDD design. All estimates are weighted using person weights (PWGTP) to ensure population-level representativeness.

I restrict the sample to individuals aged 58–66 to focus on ages near the eligibility threshold. I exclude 2020 due to data collection disruptions from the COVID-19 pandemic. To maximize statistical power, I draw from 13 large states (California, Texas, Florida, New York, Illinois, Pennsylvania, Ohio, Georgia, North Carolina, Michigan, Massachusetts, Wisconsin, and West Virginia) that together represent approximately 60% of the U.S. population aged 58–66.

Table 1 presents summary statistics. The final sample includes 1,470,507 person-year observations. At baseline (ages 58–61), approximately 18% live alone, 27% are employed, and 12% receive Social Security income. These patterns shift substantially at age 62 and above.

Table 1: Summary Statistics

	Ages 58–61	Ages 62–66	Full Sample
Living alone	0.183	0.203	0.194
Has SS income	0.120	0.461	0.300
Employed	0.614	0.475	0.540
Married	0.554	0.537	0.545
Female	0.511	0.515	0.513
College degree	0.312	0.298	0.304
Disabled	0.153	0.167	0.161
N (person-years)	671,218	799,289	1,470,507

Notes: Data from ACS PUMS 2016–2022, excluding 2020.

Sample restricted to 13 large states. All statistics are weighted using person weights (PWGTP).

3.2 Empirical Strategy

I employ a regression discontinuity design (RDD) that exploits the sharp eligibility threshold at age 62. The identifying assumption is that individuals just below and just above age 62 are comparable in all respects except for Social Security eligibility, so any discontinuity in outcomes at the threshold can be attributed to eligibility.

The main specification is:

$$Y_i = \alpha + \tau \cdot \mathbf{1}(Age_i \geq 62) + f(Age_i - 62) + X'_i \beta + \gamma_s + \delta_t + \varepsilon_i \quad (1)$$

where Y_i is the outcome (living alone indicator), $\mathbf{1}(Age_i \geq 62)$ is the treatment indicator,

$f(Age_i - 62)$ is a local linear function of the running variable allowed to differ on each side of the cutoff, X_i are covariates (sex, race, education, disability), and γ_s and δ_t are state and year fixed effects. The parameter of interest is τ , the discontinuous change in the outcome at the eligibility threshold.

I use a primary bandwidth of ± 2 years (ages 60–64) and report robustness to alternative bandwidths (± 1 , ± 3 , ± 4). Standard errors are heteroskedasticity-robust (HC1).

An important limitation is that age in the ACS is measured in whole years rather than exact age or month of birth. This creates a discrete running variable that requires careful interpretation. The standard asymptotic theory for RDD relies on continuity of potential outcomes at the cutoff, which formally requires a continuous running variable. With a discrete running variable like integer age, potential outcomes at adjacent integer values (e.g., age 61 vs. 62) may differ for reasons other than the treatment, and no amount of bandwidth shrinkage can isolate the “pure” treatment effect.

Following ?, I adopt several approaches to address this challenge. First, I focus on local linear specifications rather than higher-order polynomials, which can overfit to the discrete data points. Local linear regression has favorable bias properties and is less sensitive to the exact functional form assumption. Second, I interpret the estimated coefficients as capturing the “jump” in outcomes at the threshold rather than as point estimates of a limit. The treatment effect reflects the average difference between those observed at age 62+ and those observed at age 61 and below, after controlling for the trend in outcomes with age. Third, I report results across multiple bandwidths, recognizing that bandwidth selection is particularly consequential with discrete data. Narrower bandwidths reduce potential bias from nonlinearity but increase variance and can create collinearity problems when only a few mass points remain in the estimation window.

I also consider the possibility of “donut” RDD specifications that exclude observations at exactly age 62. If individuals who claim Social Security immediately at eligibility are sys-

tematically different from those who delay, the age-62 cell may be contaminated by selection effects. In supplementary analyses, I verify that results are similar whether or not the age-62 observations are included in the sample.

For inference, I use heteroskedasticity-robust standard errors (HC1) at the individual level as the primary specification. However, as emphasized by ?, individual-level standard errors can be misleading when the running variable is discrete. With only a handful of age support points in the estimation window (e.g., five ages at bandwidth ± 2), precision is fundamentally limited by the number of support points, not the number of individuals. To address this concern, I supplement the individual-level analysis with a **cell-level approach**: I collapse the data to age \times year cells (30 cells within the ± 2 bandwidth) and run weighted regressions at the cell level with standard errors clustered at the age level. This approach makes the effective variation transparent and provides more conservative inference appropriate for the discrete running variable setting. Following the methodological standards in ? and ?, I report both individual-level and cell-level estimates throughout.

Additionally, to clarify the estimand, I present both **intent-to-treat (ITT)** estimates that capture the effect of SS eligibility, and **fuzzy RDD estimates (LATE)** that use $\text{age} \geq 62$ as an instrument for actual SS receipt. The ITT represents the reduced-form effect of crossing the eligibility threshold, while the LATE (computed as the ratio of reduced form to first stage) represents the causal effect of SS income receipt for compliers—those induced to claim by reaching age 62.

3.3 Identification

The validity of the RDD rests on two key assumptions. First, the treatment (SS eligibility) must change discontinuously at the threshold. This is satisfied by design: age 62 is a bright-line federal rule with no fuzzy margin. Second, other determinants of the outcome must evolve smoothly through the threshold—there should be no other discontinuities in factors

affecting living arrangements precisely at age 62.

I examine several validity checks to assess the credibility of the design. The first and most important check is the first stage: I verify that Social Security income receipt increases discontinuously at age 62. Without a meaningful jump in SS income at the threshold, there would be no policy variation to exploit. A strong first stage is essential for interpreting the reduced-form effects on living arrangements.

The second validity check examines covariate balance. I test whether predetermined covariates—specifically sex, race, education, and disability status—are smooth through the age-62 threshold. Discontinuities in predetermined characteristics would suggest that the individuals just above and below the threshold differ in ways beyond SS eligibility, which would violate the identifying assumption of the RDD.

Third, I conduct a manipulation test by examining whether the density of observations shows unusual bunching at age 62. In some RDD contexts, individuals can manipulate the running variable to sort above or below the threshold. While age cannot be directly manipulated (it is determined by birth date), differential survey response or mortality patterns could create discontinuities in the density. Smooth density through the threshold supports the design's validity.

Finally, I test for discontinuities at placebo cutoffs: ages 59, 60, 61, 63, 64, and 65. Finding effects at false thresholds where no policy change occurs would cast doubt on the attribution of effects to SS eligibility at 62. Conversely, finding effects concentrated at the true threshold strengthens the causal interpretation.

4 Results

4.1 First Stage: Social Security Income Receipt

Figure 1 shows the first stage: the probability of receiving Social Security income by age. There is a clear, large discontinuity at age 62. Before 62, approximately 12–14% of individuals in this age range receive Social Security income (primarily disability benefits). At 62, this jumps to approximately 27%, rising further to 55% by age 65 and 70% by age 66.




Figure 1: First Stage: Social Security Income Receipt by Age

Notes: Each point represents the weighted mean share receiving Social Security income for that age. Points below 62 are shown in blue; points at 62 and above are shown in red. Lines show local linear fits on each side of the cutoff. Data from ACS PUMS 2016–2022 (excluding 2020), 13 large states.

Table 2 presents formal RDD estimates of the first stage. The coefficient on the treatment indicator ($\text{Age} \geq 62$) is 0.134 (SE = 0.003), indicating that SS income receipt increases by 13.4 percentage points at the threshold. This effect is highly significant ($p < 0.001$) and robust to bandwidth choice.

Table 2: First Stage: Effect of Age-62 Threshold on SS Income Receipt

Bandwidth	± 1	± 2	± 3	± 4
Treatment (Age ≥ 62)	—	0.134*** (0.003)	0.115*** (0.002)	0.101*** (0.002)
Observations	327,660	824,709	1,152,715	1,470,507

Notes: Robust standard errors in parentheses. *** $p < 0.01$. All specifications include local linear control for age (allowed to differ on each side of cutoff), state and year fixed effects, and controls for sex, race, education, and disability. Estimates weighted by person weights. Bandwidth ± 1 estimate omitted due to collinearity with discrete running variable.

4.2 Main Result: Living Alone

Figure 2 shows the main outcome: the probability of living alone by age. The pattern shows a steady upward trend across ages, from about 17% at age 58 to about 21% at age 66. However, there is a visible discontinuity at age 62: the rate of living alone appears to be slightly *lower* just after 62 than would be predicted by extrapolating the pre-62 trend.




Figure 2: Main Outcome: Living Alone by Age

Notes: Each point represents the weighted mean share living alone for that age. Points below 62 are shown in blue; points at 62 and above are shown in red. Lines show local linear fits on each side of the cutoff. Living alone is defined as $NP = 1$ (household size of one person).

Table 3 presents the main RDD estimates. The preferred specification (bandwidth ± 2) yields a coefficient of -0.0067 ($SE = 0.003$), indicating that Social Security eligibility *decreases* the probability of living alone by 0.67 percentage points. This effect is statistically significant at the 5% level ($p = 0.026$).

The 95% confidence interval is $[-1.26, -0.08]$ percentage points, allowing us to rule out zero effect as well as large positive effects. The point estimate represents a relative decline of approximately 3.5% from the baseline rate of 19%.

Table 3: Main Results: Effect of SS Eligibility on Living Alone

	Bandwidth			
	± 1	± 2	± 3	± 4
Treatment (Age ≥ 62)	—	-0.0067** (0.003)	-0.0028 (0.002)	-0.0039** (0.002)
95% CI		[-0.013, -0.001]	[-0.007, 0.001]	[-0.007, -0.001]
<i>p</i> -value		0.026	0.191	0.029
Observations	327,660	824,709	1,152,715	1,470,507

Notes: Robust standard errors in parentheses. ** $p < 0.05$. All specifications include local linear control for age, state and year fixed effects, and demographic controls. Baseline living alone rate ≈ 0.19 .

4.2.1 Cell-Level Inference and Fuzzy RDD Estimates

As discussed in the Methods section, the discrete running variable raises concerns about individual-level inference. Table 4 presents results using the cell-level approach, which collapses data to 30 age \times year cells and clusters standard errors at the age level. This approach makes the effective variation transparent: with only 5 age support points in the ± 2 bandwidth, inference is fundamentally limited by this small number of clusters.

Reassuringly, the cell-level estimates are nearly identical to the individual-level estimates. The first stage coefficient is 0.134 (SE = 0.007, $p < 0.001$), and the main outcome coefficient is -0.0067 (SE = 0.001, $p < 0.001$). The clustered standard errors are actually smaller than the individual-level robust standard errors in this case, reflecting the high precision of the cell means. The key finding that SS eligibility reduces living alone remains robust to this more conservative inference approach.

Table 4: Cell-Level Inference: Age \times Year Collapsed Estimates

	First Stage (SS Receipt)	Living Alone (Main Outcome)
Treatment (Age \geq 62)	0.134*** (0.007)	-0.0067*** (0.001)
95% CI	[0.119, 0.148]	[-0.009, -0.005]
<i>p</i> -value	<0.001	<0.001
Age \times Year cells	30	30
Age clusters	5	5

Notes: Data collapsed to age \times year cells. Standard errors clustered at age level. Bandwidth ± 2 . *** $p < 0.01$.

Table 5 presents the fuzzy RDD estimates that interpret the treatment as actual SS receipt rather than eligibility. Using age ≥ 62 as an instrument for SS receipt, the LATE is -0.050 (SE = 0.023, $p = 0.026$). This implies that among compliers—those induced to claim SS by reaching age 62—receiving Social Security income reduces the probability of living alone by approximately 5 percentage points. The larger LATE compared to the ITT reflects the fact that only about 13% of individuals are induced to claim at the 62 threshold (the first stage).

Table 5: Fuzzy RDD: Effect of SS Receipt on Living Alone

	First Stage	Reduced Form	2SLS (LATE)
Coefficient	0.134*** (0.003)	-0.0067** (0.003)	-0.050** (0.023)
<i>p</i> -value	<0.001	0.026	0.026
Observations	824,709	824,709	824,709

Notes: Fuzzy RDD using $\text{Age} \geq 62$ as instrument for SS income receipt. First stage: effect of eligibility on SS receipt. Reduced form: ITT effect of eligibility on living alone. 2SLS: LATE of SS receipt on living alone. SE for LATE computed via delta method. Bandwidth ± 2 . *** $p < 0.01$, ** $p < 0.05$.

4.3 Validity Checks

Covariate Balance. Table 6 tests for discontinuities in predetermined covariates at age 62. None of the four covariates (sex, race, education, disability) show statistically significant discontinuities at the threshold. This supports the identifying assumption that individuals just above and below the threshold are comparable.

Table 6: Covariate Balance at Age-62 Threshold

Covariate	Coefficient	SE	<i>p</i> -value
Female	0.004	0.004	0.273
Race	0.031	0.019	0.091
Education	-0.017	0.031	0.579
Disability	0.003	0.003	0.404

Notes: Each row reports the RDD estimate of the discontinuity in the listed covariate at age 62. Bandwidth ± 2 . None significant at the 5% level.

Placebo Cutoffs. Table 7 tests for discontinuities in living alone at false cutoffs. Of six placebo tests (ages 59, 60, 61, 63, 64, 65), only age 60 shows a marginally significant effect ($p = 0.045$). Given six tests, one significant result at the 5% level is not unexpected by chance alone. The true eligibility threshold at 62 shows a larger and more robust effect than any placebo.

Table 7: Placebo Tests: Discontinuities at False Cutoffs

Cutoff Age	Coefficient	SE	<i>p</i> -value
59	—	—	—
60	−0.006	0.003	0.045*
61	0.006	0.003	0.060
62 (True)	−0.007	0.003	0.026**
63	0.001	0.003	0.746
64	0.005	0.003	0.090
65	−0.004	0.003	0.234

Notes: Each row tests for a discontinuity in living alone at the listed age cutoff. Bandwidth ± 2 around each cutoff. Age 59 omitted due to sample boundary. * $p < 0.05$, ** $p < 0.05$.

Manipulation Test. Figure ?? (Panel D of the main results figure) shows the density of observations by age. There is no evidence of unusual bunching or discontinuity in the density at age 62. This is expected: age is determined by birth date and cannot be manipulated.

4.4 Heterogeneity

Table 8 examines heterogeneity in the treatment effect across subgroups. The most striking finding is the difference by marital status. Among unmarried individuals, SS eligibility reduces living alone by 2.93 percentage points ($p < 0.001$)—more than four times the overall effect. Among married individuals, the effect is essentially zero (0.23 pp, $p = 0.06$).

This pattern makes sense: married individuals rarely live alone (by definition, they live with their spouse), so there is little scope for SS eligibility to affect their living arrangements.

For unmarried individuals, by contrast, the relevant margin is between living alone versus with other non-spouse household members (adult children, siblings, roommates, etc.).

The effect is also larger for men (-1.11 pp, $p = 0.01$) than for women (-0.27 pp, $p = 0.52$). This may reflect gender differences in social networks or housing circumstances. The effect does not differ significantly by education or disability status.

Table 8: Heterogeneity in Treatment Effects

Subgroup	Coefficient	SE	p-value	N
<i>By Sex</i>				
Male	-0.011^{**}	0.004	0.010	403,841
Female	-0.003	0.004	0.519	420,868
<i>By Marital Status</i>				
Married	0.002	0.001	0.061	452,879
Unmarried	-0.029^{***}	0.007	<0.001	371,830
<i>By Disability</i>				
Disabled	-0.012	0.008	0.141	130,082
Not Disabled	-0.006^*	0.003	0.064	694,627
<i>By Education</i>				
College+	-0.007	0.005	0.195	254,417
No College	-0.007^*	0.004	0.070	570,292

Notes: Each row reports the RDD estimate for the indicated subgroup. Bandwidth ± 2 . *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

4.5 Secondary Outcomes

Table 9 examines effects on secondary outcomes. Employment shows a significant discontinuous decline at age 62 (-3.6 pp, $p < 0.001$), consistent with prior literature on SS and retirement. The marriage rate shows a marginally significant decline (-0.64 pp, $p = 0.086$), though this should be interpreted cautiously given the discrete age measurement.

Table 9: Effects on Secondary Outcomes

Outcome	Coefficient	SE	<i>p</i> -value
Employment	-0.036^{***}	0.004	<0.001
Marriage	-0.006^*	0.004	0.086

Notes: Bandwidth ± 2 . *** $p < 0.01$, * $p < 0.10$.

5 Discussion

5.1 Interpretation

The main finding is that Social Security eligibility at age 62 *reduces* the probability of living alone by approximately 0.7 percentage points overall, with effects concentrated among unmarried individuals (2.9 pp) and men (1.1 pp). This suggests that income security enables household consolidation rather than independence. While the effect may seem modest in percentage point terms, it represents a meaningful 3.5% relative decline from the baseline rate of living alone and implies substantial changes in living arrangements for the subpopulations most affected by the policy.

Several mechanisms could explain why Social Security eligibility reduces rather than increases living alone. The first mechanism involves geographic mobility. Social Security in-

come may enable individuals to relocate to be near family members, facilitating co-residence that would otherwise be impractical. Without guaranteed income, older adults may be “stuck” in their pre-retirement location near former workplaces, separated from family by geographic distance. The income security provided by SS benefits could finance a move and provide the economic stability needed to establish a new residence near relatives. This mechanism is particularly plausible for individuals who maintained employment-driven residences during their working years but whose families live elsewhere.

A second mechanism centers on housing affordability. Social Security income may make it affordable to maintain larger housing that accommodates other household members. A one-bedroom apartment sufficient for solo living may be replaced by a two-bedroom unit that allows an adult child or sibling to move in. Alternatively, pooling Social Security income with family members may enable households that would otherwise be economically impossible—neither party could afford independent housing alone, but together they can afford a shared residence. This income-pooling mechanism is consistent with the strong effects observed among lower-income populations and the unmarried.

Third, Social Security income might enable partnership formation that reduces solo living. The marginal decline in marriage rates observed at age 62 (Table 6) is small and imprecise, but unmarried cohabitation may increase without being detected in formal marriage statistics. Having guaranteed income could make an individual a more attractive partner, or could provide the financial security needed to take the risk of forming a new household with a romantic partner. This mechanism would be particularly relevant for previously divorced or never-married individuals who face both economic and social barriers to partnership later in life.

A fourth mechanism involves supporting others. Social Security income may enable older adults to support adult children or other family members who then move in, transforming a one-person household into a multi-person household. An adult child facing economic

hardship might move in with a parent who can now afford to house them thanks to SS income. This “reverse co-residence”—where the older adult supports the younger relative rather than vice versa—is increasingly common in an era of rising housing costs and economic insecurity among younger generations.

The heterogeneity results help adjudicate between these mechanisms. The concentration of effects among the unmarried is particularly informative. Married individuals rarely live alone by definition (they live with their spouse), so there is limited scope for SS eligibility to affect their living arrangements through any of the mechanisms described above. For unmarried individuals, by contrast, all four mechanisms could operate: they could move to be near family, afford larger housing, form partnerships, or support relatives who move in. The large effect size for the unmarried (2.9 percentage points, implying that roughly 22% of induced SS claimers transition out of solo living) suggests these mechanisms are quantitatively important for this population.

5.2 Magnitude

Is a 0.7 percentage point reduction in living alone economically meaningful? On a baseline of 19%, this represents a 3.5% relative decline. Given that the first stage shows SS income receipt increases by 13.4 pp at the threshold, a rough “Wald” estimate would attribute the entire living alone reduction to SS income: $-0.007/0.134 = -0.05$. This implies that among those induced to claim SS at 62, approximately 5% experience a change from living alone to living with others.

For the unmarried subgroup, the implied effect is larger: $-0.029/0.134 = -0.22$. Among unmarried individuals induced to claim SS by reaching age 62, approximately 22% transition out of solo living. This is a substantial effect.

5.3 Limitations

Several limitations of this analysis warrant discussion. Understanding these limitations is important for interpreting the findings and for designing future research to address remaining questions.

The first limitation concerns the discrete running variable. Age in the ACS is measured in whole years rather than exact age or month of birth. This creates several challenges for the RDD methodology. Standard RDD theory relies on a continuous running variable that allows the analyst to compare observations arbitrarily close to the threshold. With a discrete running variable, comparisons are necessarily between different integer values (e.g., age 61 vs. age 62), and any differences between these groups could reflect either the treatment effect or other factors that change between adjacent ages. I address this limitation by using local linear specifications that are less sensitive to functional form assumptions, by reporting results across multiple bandwidths, and by following the methodological recommendations of ? for inference with discrete running variables. However, the fundamental challenge remains, and readers should interpret the estimates as capturing the jump at the threshold rather than as point estimates of a well-defined treatment effect at a limit.

A second limitation involves first stage timing. Not everyone eligible at 62 claims Social Security immediately. In fact, a substantial fraction of individuals delay claiming until age 65 or later, hoping to receive higher monthly benefits by waiting. The RDD identifies an intent-to-treat effect of eligibility rather than the effect of actual claiming. Because many individuals who become eligible at 62 do not claim immediately, the reduced-form effect of eligibility on living arrangements may substantially underestimate what would happen if all eligible individuals claimed. Conversely, the reduced-form effect may capture anticipation effects: individuals may begin adjusting their living arrangements before actually claiming, in anticipation of future benefits. Disentangling these timing effects would require linking to administrative claiming data, which is beyond the scope of the current analysis.

Third, living arrangements are a slow-moving outcome. Unlike variables that can change rapidly in response to policy (such as labor supply, which can adjust within weeks), living arrangements typically change infrequently and with considerable lag. Moving to a new residence involves significant transaction costs—finding housing, relocating belongings, establishing new routines. Partnership formation or dissolution similarly involves extended processes rather than instantaneous changes. The annual cross-sectional ACS data may miss effects that accumulate over longer time horizons. If the full effect of Social Security eligibility on living arrangements takes 2–3 years to materialize, the estimates based on the jump at age 62 would capture only the initial portion of this effect. The findings should therefore be interpreted as lower bounds on the eventual effect if adjustment is gradual.

Fourth, external validity is limited by the sample construction. To maximize statistical power, I restrict the analysis to 13 large states that represent approximately 60% of the U.S. population aged 58–66. Effects may differ in smaller or more rural states not included in the analysis. In particular, housing markets vary substantially across states: the financial trade-offs between living alone and co-residence depend on local housing costs, which range from very high in states like California and New York to much lower in states like West Virginia. The effect of Social Security income on living arrangements may be larger in high-cost states where housing affordability is more binding, or smaller in low-cost states where even modest pre-SS income could sustain independent living. Generalizing the findings to the full U.S. population should be done cautiously.

Finally, the ACS does not contain measures of actual loneliness or social isolation. Living alone is used as a proxy for loneliness based on the well-documented correlation between the two, but it is an imperfect measure. Some individuals who live alone maintain rich social networks and experience little loneliness; others who live with family members feel profoundly isolated. The finding that Social Security eligibility reduces living alone does not directly demonstrate that it reduces loneliness. Future research linking RDD estimates to survey measures of subjective loneliness would strengthen the connection between this analysis and

the policy goal of reducing elder isolation.

5.4 Policy Implications

These findings have several policy implications:

First, retirement income programs may have unintended social benefits beyond direct welfare effects on recipients. By enabling household consolidation, Social Security may reduce social isolation among older adults—a recognized public health concern.

Second, proposals to raise the early eligibility age (currently 62) should consider effects on living arrangements. Delaying eligibility would delay the household consolidation effects documented here, potentially prolonging social isolation for some older adults.

Third, policies aimed at reducing elder loneliness might consider the role of income constraints. The finding that financial resources affect living arrangements suggests that income support could be a lever for addressing social isolation.

6 Conclusion

This paper provides the first quasi-experimental evidence on how Social Security eligibility affects living arrangements among older Americans. Using a regression discontinuity design centered at the age-62 early eligibility threshold, I find that Social Security eligibility *reduces* the probability of living alone by 0.67 percentage points, with effects concentrated among unmarried individuals (2.9 pp) and men (1.1 pp). A strong first stage confirms that Social Security income receipt increases by 13.4 percentage points at the eligibility threshold, and validity checks including covariate balance tests, placebo cutoffs, and manipulation tests support the credibility of the research design.

These findings suggest that income security enables household consolidation rather than

independence. The theoretical prediction was ambiguous: Social Security income could either increase living alone (by enabling financial independence) or decrease it (by enabling co-residence). The empirical results decisively favor the latter interpretation. For many older Americans approaching retirement age, the barrier to co-residence is not a preference for independence but rather financial constraints that Social Security helps relieve. This has important implications for how we understand the relationship between economic resources and social outcomes in later life.

The results also contribute to ongoing policy debates about Social Security reform. Proposals to raise the early eligibility age from 62 to 64 or higher are often evaluated in terms of effects on labor supply, retirement timing, and fiscal sustainability. This paper suggests that such proposals would also delay the household consolidation effects documented here, potentially prolonging social isolation for some older adults during the additional years before eligibility. While this should not be the sole consideration in evaluating eligibility age changes, it represents a previously unrecognized cost that should enter the policy calculus.

More broadly, as policymakers grapple with the loneliness epidemic among older adults, these results highlight an underappreciated channel through which retirement income programs may facilitate social connection. The U.S. Surgeon General’s 2023 advisory on loneliness emphasized the need for interventions that address social isolation, but focused primarily on community programs, mental health services, and social infrastructure. The findings of this paper suggest that income support programs—already a cornerstone of U.S. social policy—may have unintended benefits for social connection by enabling household formation. Understanding these broader effects is important for comprehensive evaluation of social insurance programs and for designing holistic approaches to addressing elder isolation.

Finally, future research should examine whether these effects persist over longer time horizons, whether they extend to other income thresholds (such as full retirement age at 67), and whether similar dynamics operate in other countries with different Social Security

systems. Linking administrative data on Social Security claiming to longitudinal surveys with richer measures of social contact and loneliness would allow researchers to more directly assess the welfare implications of the household composition changes documented here. As populations age worldwide, understanding the interplay between retirement systems and social outcomes will become increasingly important for policy.

A Additional Figures




Figure 3: Summary of Main Results

Notes: Panel A shows the first stage (SS income receipt by age). Panel B shows the main outcome (living alone by age). Panel C shows employment by age. Panel D shows the density of observations by age (manipulation test). Points below age 62 in blue; points at 62 and above in red.

B Robustness to Bandwidth

Table 10: Robustness to Alternative Bandwidths

Bandwidth	Living Alone			SS Income (First Stage)		
	Coef	SE	<i>p</i>	Coef	SE	<i>p</i>
±2	-0.007	0.003	0.026	0.134	0.003	<0.001
±3	-0.003	0.002	0.191	0.115	0.002	<0.001
±4	-0.004	0.002	0.029	0.101	0.002	<0.001

Notes: All specifications include local linear control for age, state and year FE, and demographic controls.