

# Intergenerational Time Transfers and Social Security Eligibility: Evidence from a Regression Discontinuity Design

Autonomous Policy Evaluation Project (APEP)\*

January 2026

## Abstract

We examine whether grandparents' eligibility for early Social Security retirement benefits at age 62 affects labor supply patterns in multigenerational households. Using American Community Survey data from 2019–2022 and a regression discontinuity design, we test two hypotheses: (1) whether grandparents reduce work hours discontinuously upon reaching age 62, and (2) whether this reduction creates “spillover” effects on the labor supply of co-resident working-age parents through increased grandparent-provided childcare. We find no evidence of a discrete change in grandparent work hours at the 62 threshold, despite observing a smooth and substantial decline in hours across the 58–66 age range. Similarly, we find no evidence of intergenerational spillover effects on parent labor supply. These null results suggest that Social Security early eligibility does not create sharp behavioral responses in multigenerational households, and that retirement transitions in this context are gradual rather than discontinuous. Our findings have implications for the design of retirement policy and for understanding intergenerational time allocation.

---

\*This paper was autonomously generated by an AI research system. Correspondence: [a pep@example.org](mailto:a pep@example.org). Data and replication files available upon request. [nd@dakoyana](mailto:nd@dakoyana)

**JEL Codes:** J22, J26, H55, D13

**Keywords:** Social Security, retirement, labor supply, intergenerational transfers, regression discontinuity

# Contents

# 1 Introduction

The decision to retire represents one of the most consequential choices individuals make during their working lives. For millions of Americans, the first opportunity to claim Social Security retirement benefits arrives at age 62, when “early retirement” becomes available, albeit with actuarially reduced monthly payments. A substantial literature has examined how individuals respond to this eligibility threshold, documenting effects on labor supply, consumption, and health insurance decisions (??).

However, retirement decisions are rarely made in isolation. Within multigenerational households—an increasingly common living arrangement in the United States—the retirement of an older family member may have significant spillover effects on other household members. When grandparents retire, they gain time that can be allocated to family support activities, most notably providing childcare for grandchildren. This informal childcare, in turn, may enable working-age parents to increase their labor supply by reducing the constraints associated with childcare costs and availability.

This paper investigates whether the Social Security early eligibility threshold at age 62 creates intergenerational spillover effects on labor supply within multigenerational households. We examine two related research questions. First, we ask whether grandparents in multigenerational households reduce their work hours discontinuously at age 62, when they first become eligible for Social Security benefits. This “first stage” relationship is necessary for any spillover effects to operate through the retirement channel. Second, we examine the spillover effect itself: when co-resident grandparents reach age 62, do working-age parents with children increase their work hours, potentially due to increased availability of grandparent-provided childcare? If grandparents reduce market work upon becoming eligible for Social Security, they may substitute toward household production, including childcare, which could relax constraints on parent labor supply.

To address these questions, we employ a regression discontinuity design (RDD) using data from the American Community Survey (ACS) Public Use Microdata Samples (PUMS) for

2019–2022. We identify multigenerational households containing at least one grandparent-aged adult (58–66), at least one working-age adult (25–55), and at least one child under 18. Our primary specification examines whether labor supply measures change discontinuously at the age 62 threshold.

Our analysis yields null results on both margins. For the first stage, we find that grandparent work hours decline substantially across the 58–66 age range—from approximately 30 hours per week at age 58 to 13 hours at age 66—but this decline is smooth and gradual rather than discontinuous at age 62. The estimated treatment effect at the 62 threshold is small (−1.1 hours) and statistically insignificant ( $p = 0.81$ ). Similarly, we find no evidence of spillover effects on parent labor supply: the estimated effect of grandparent eligibility on parent work hours is near zero (0.6 hours) and insignificant ( $p = 0.87$ ).

These null findings are robust to alternative bandwidth choices, polynomial specifications, and sample restrictions. Placebo tests at non-threshold ages confirm that the absence of discontinuity is specific to our setup rather than an artifact of noisy data. Validity checks show no evidence of manipulation at the threshold and confirm that covariates are smooth across the cutoff.

Our results contribute to several literatures. First, we add to the extensive literature on Social Security and retirement behavior by documenting that the age 62 threshold—while marking the first opportunity to claim benefits—does not appear to create discrete behavioral changes in multigenerational households. This finding is consistent with recent work suggesting that retirement transitions are increasingly gradual, with workers phasing out of the labor force over multiple years rather than making abrupt exits (?).

Second, we contribute to the growing literature on intergenerational transfers and family labor supply. While previous research has documented the importance of grandparent-provided childcare for maternal labor supply (??), our findings suggest that policy-induced changes in grandparent availability—at least those occurring at the Social Security early eligibility threshold—may not generate detectable effects on parent labor supply.

Third, our paper demonstrates the application of RDD methods to study spillover effects within households. While prior work has examined spousal spillovers at the Medicare eligibility threshold (?), we extend this framework to intergenerational spillovers, providing a template for future research on family-level policy effects.

The null results we document should not be interpreted as evidence that grandparent retirement has no effect on family time allocation. Rather, they suggest that the age 62 threshold specifically does not generate the discrete changes that would be captured by an RDD. Several factors may explain this: the actuarial penalty for early claiming may deter discrete responses, workers may anticipate the threshold and adjust gradually, or selection into multigenerational living arrangements may attenuate the effects we measure.

The remainder of this paper proceeds as follows. Section ?? provides institutional background on Social Security and reviews the relevant literature. Section ?? describes our data sources and sample construction. Section ?? presents our empirical strategy. Section ?? reports our main findings. Section ?? presents robustness checks and validity tests. Section ?? discusses the interpretation and implications of our results. Section ?? concludes.

## 2 Background and Literature

### 2.1 Social Security Early Retirement

The Social Security retirement system provides the foundation of retirement income for most American workers. While the “full retirement age” (FRA) for receiving unreduced benefits has gradually increased from 65 to 67 for recent birth cohorts, workers have the option to claim benefits as early as age 62. Claiming before FRA results in permanently reduced monthly benefits: for workers with an FRA of 67, claiming at 62 results in a 30% benefit reduction (?).

Despite this substantial penalty, early claiming remains popular. Approximately 30-35% of Social Security beneficiaries claim at age 62, making it the most common claiming age (?).

This behavior has motivated extensive research on the relationship between Social Security incentives and labor supply.

## 2.2 Literature on Retirement and Labor Supply

The literature on Social Security and retirement behavior is vast. Classic work by ? demonstrated that Social Security’s incentive structure significantly affects retirement timing across countries. ? documented that increases in the full retirement age led workers to delay retirement, with approximately one third of affected workers adjusting their behavior in response to the policy change.

More recent work has employed regression discontinuity designs to study age-based policy thresholds. ? exploited the Medicare eligibility threshold at age 65 to study health insurance effects, while ? examined how Social Security eligibility at 62 affects labor supply among workers with different health statuses.

A key finding in recent literature is that retirement transitions have become increasingly gradual. ? document that the majority of workers now exit the labor force through “bridge jobs” or phased retirement rather than abrupt exits. This pattern may explain why age-based eligibility thresholds generate smaller discontinuities than would be predicted by a model of discrete retirement decisions.

The theoretical framework linking Social Security eligibility to labor supply operates through several channels. First, eligibility creates an income effect: workers who can claim benefits have access to non-labor income that may reduce the marginal utility of additional work. Second, eligibility interacts with liquidity constraints: workers who have been unable to save adequately for retirement may face binding constraints that eligibility relaxes. Third, Social Security eligibility may interact with health insurance: for workers without retiree health coverage, the years between 62 and Medicare eligibility at 65 create a coverage gap that may influence retirement timing (?).

The magnitude of behavioral responses to Social Security incentives has been debated.

Early estimates suggested substantial labor supply elasticities, but more recent work using administrative data has found smaller effects. ? estimates that a one-year increase in the full retirement age delays retirement by approximately two months, implying a relatively modest behavioral response. This attenuated response may reflect the complexity of retirement decisions, which involve trade-offs across multiple margins including health, spousal coordination, and non-pecuniary job characteristics.

The specific threshold at age 62 has received considerable attention because it represents the earliest opportunity to claim benefits. However, the actuarial adjustment for early claiming—which reduces benefits by approximately 6.7% per year before full retirement age—creates countervailing incentives. Workers who claim early receive smaller monthly benefits for a potentially longer duration, making the net present value calculation non-trivial. Research suggests that many early claimers may be making suboptimal decisions, particularly those with longer life expectancies who would benefit from delayed claiming (?).

### **2.3 Intergenerational Transfers and Childcare**

The economics of family time allocation has received increasing attention as multigenerational households have grown more common in the United States. Between 2000 and 2016, the share of Americans living in multigenerational households rose from 12% to 20% (?).

Grandparent-provided childcare represents a significant form of intergenerational transfer. ? find that access to grandparent childcare significantly increases maternal labor force participation in Europe. ? document similar effects in the United States, finding that proximity to grandparents increases the labor supply of mothers with young children.

The economic significance of grandparent childcare is substantial. According to survey data, approximately 30% of children under age 5 receive some regular care from grandparents, with higher rates among low-income families where formal childcare is less affordable. The implicit value of this care—measured as the cost of equivalent formal care—amounts to billions of dollars annually. For individual families, access to grandparent care can make the

difference between one parent working and both parents working, or between part-time and full-time employment.

The decision to provide grandparent childcare involves complex trade-offs. Grandparents must weigh the benefits of supporting their adult children and bonding with grandchildren against the opportunity costs of their own time, including foregone earnings if they reduce work hours or delay retirement. Research on grandparent well-being suggests that moderate amounts of grandchild care are associated with better health outcomes, possibly due to increased physical activity and social engagement, while intensive caregiving responsibilities can lead to caregiver burden and negative health effects.

From the perspective of working-age parents, grandparent care offers several advantages over formal childcare. The care is typically provided by a trusted family member, is flexible in scheduling, and is often provided at no direct cost. However, reliance on grandparent care also creates vulnerabilities: illness or mobility limitations affecting grandparents can disrupt childcare arrangements, and geographic distance limits availability. The co-residence that defines our sample represents one strategy families employ to facilitate grandparent-provided childcare.

However, the causal relationship between grandparent retirement and parent labor supply has received less attention. The key challenge is that grandparent retirement may be endogenous to family circumstances, including the childcare needs of adult children. Our regression discontinuity approach addresses this concern by exploiting the exogenous variation in retirement incentives created by the Social Security eligibility threshold.

## **2.4 Spousal and Family Spillovers**

A small but growing literature examines how policy-induced changes in one family member's behavior affect others. ? uses an RDD at the Medicare eligibility threshold to study how one spouse's eligibility affects the other spouse's health insurance coverage and labor supply. She finds effects on insurance coverage but limited effects on labor supply.

Research on joint retirement decisions documents that spouses tend to retire together, suggesting complementarities in leisure time (??). ? use European data to study how pension reforms affect spousal labor supply, finding evidence of coordination in retirement timing.

Our paper extends this literature by examining intergenerational rather than spousal spillovers, and by focusing on the intensive margin (hours worked) rather than the extensive margin (labor force participation).

## 3 Data

### 3.1 Data Source

We use data from the American Community Survey (ACS) Public Use Microdata Samples (PUMS) for the years 2019, 2021, and 2022. The ACS is an annual survey administered by the U.S. Census Bureau, providing detailed demographic, employment, and household information for approximately 3.5 million individuals per year. The 1-year PUMS files contain individual-level records with sampling weights that allow researchers to estimate population parameters.

The ACS is well-suited for our analysis for several reasons. First, it provides large sample sizes necessary to detect potentially small discontinuities. Second, it includes detailed information on employment, hours worked, and household composition. Third, the survey structure allows us to identify multigenerational households and link the characteristics of different household members.

We exclude 2020 from our analysis due to data collection disruptions associated with the COVID-19 pandemic, which affected both survey response rates and labor market behavior in ways that could confound our estimates.

## 3.2 Sample Construction

We construct our analysis sample through a multi-step process. First, we identify multigenerational households, defined as households containing:

1. At least one “grandparent-aged” adult between 58 and 66 years old
2. At least one “parent-aged” working-age adult between 25 and 55 years old
3. At least one child under 18 years old

This sample restriction focuses our analysis on households where the hypothesized mechanism—grandparent-provided childcare enabling parent labor supply—is plausible. We do not require explicit identification of family relationships; instead, we rely on the co-residence of individuals in the relevant age ranges.

Second, we restrict the grandparent sample to individuals within our bandwidth of the age 62 cutoff. For our primary specification, we use a bandwidth of 4 years, including grandparents aged 58–66. We examine robustness to narrower bandwidths (2–3 years) in our sensitivity analysis.

## 3.3 Variables

Our primary outcome variable is usual hours worked per week (WKHP), measured as self-reported typical weekly work hours for employed individuals (with zeros for those not working). This variable captures both extensive margin (employment) and intensive margin (hours conditional on working) variation.

For the first-stage analysis, the running variable is the grandparent’s own age, centered at 62. For the spillover analysis, the running variable is the age of the oldest grandparent in the household, again centered at 62.

We construct several additional variables for heterogeneity analysis:

- **Young child indicator:** Equals 1 if the youngest child in the household is aged 0–5

- **Female indicator:** Equals 1 if the individual is female
- **State:** State of residence for geographic heterogeneity

All analyses use the ACS person weights (PWGTP) to produce nationally representative estimates.

### 3.4 Summary Statistics

Table ?? presents summary statistics for our analysis samples. The grandparent sample includes 37,035 individuals aged 58–66 living in multigenerational households across our four-state sample (California, Texas, Florida, and New York). The mean age is 61.6 years, and mean hours worked is 23.2 hours per week. Approximately 59% of the sample reports positive work hours.

The parent sample includes 52,874 working-age adults (25–55) living in households with grandparent-aged adults. The mean age is 38.6 years, and mean hours worked is 29.8 hours per week. Among this sample, 44% live in households with young children (ages 0–5), representing the subgroup where childcare spillovers should be most relevant.

Table 1: Summary Statistics

	Grandparents		Parents	
	Mean	SD	Mean	SD
Age	61.6	2.6	38.6	8.4
Usual hours worked/week	23.2	19.8	29.8	18.6
Working (> 0 hours)	0.59	–	0.78	–
Has young child (0-5)	–	–	0.44	–
Observations	37,035		52,874	

Notes: Data from ACS PUMS 2019, 2021, 2022 for CA, TX, FL, and NY. Grandparents are defined as adults aged 58–66 in multigenerational households. Parents are defined as adults aged 25–55 in households with grandparent-aged adults and children under 18.

Table ?? shows mean hours worked by age for the grandparent sample. Hours decline

substantially across the age range, from 30.0 at age 58 to 13.4 at age 66. Notably, there is no obvious discontinuity at age 62; the decline appears smooth and approximately linear.

Table 2: Grandparent Hours Worked by Age

Age	N	Mean Hours	SE	95% CI
58	5,234	30.0	0.32	[29.4, 30.6]
59	4,735	29.1	0.34	[28.4, 29.8]
60	4,478	27.3	0.35	[26.6, 28.0]
61	4,281	25.3	0.37	[24.6, 26.0]
<b>62</b>	<b>3,924</b>	<b>22.7</b>	<b>0.39</b>	<b>[21.9, 23.5]</b>
63	3,776	20.5	0.40	[19.7, 21.3]
64	3,627	18.6	0.40	[17.8, 19.4]
65	3,599	14.9	0.38	[14.2, 15.6]
66	3,381	13.4	0.38	[12.7, 14.1]

Notes: Weighted means using person weights. Standard errors from bootstrap (200 replications).

## 4 Empirical Strategy

### 4.1 Regression Discontinuity Design

Our identification strategy exploits the discrete change in Social Security eligibility at age 62. Before this age, workers cannot claim retirement benefits; at 62, early claiming becomes available. This institutional feature creates a potential discontinuity in retirement incentives that we use to identify causal effects.

For the first-stage analysis (grandparent hours), we estimate:

$$Y_{gi} = \alpha + \tau \cdot \mathbf{1}(\text{Age}_{gi} \geq 62) + f(\text{Age}_{gi} - 62) + \mathbf{X}'_{gi}\beta + \varepsilon_{gi} \quad (1)$$

where  $Y_{gi}$  is usual hours worked for grandparent  $g$  in household  $i$ ,  $\mathbf{1}(\text{Age}_{gi} \geq 62)$  is an indicator for being at or above the eligibility threshold,  $f(\cdot)$  is a flexible function of the running variable (age centered at 62), and  $\mathbf{X}_{gi}$  are covariates. The parameter of interest is  $\tau$ , which captures the discontinuous change in hours at the threshold.

For the spillover analysis (parent hours), we estimate:

$$Y_{pi} = \gamma_0 + \gamma \cdot \mathbf{1}(\text{GP\_Age}_i \geq 62) + g(\text{GP\_Age}_i - 62) + \mathbf{X}'_{pi}\delta + \eta_{pi} \quad (2)$$

where  $Y_{pi}$  is usual hours worked for parent  $p$  in household  $i$ , and  $\text{GP\_Age}_i$  is the age of the oldest grandparent in the household. The parameter  $\gamma$  captures the reduced-form effect of grandparent eligibility on parent labor supply.

We estimate both equations using local linear regression, which fits separate linear trends on each side of the cutoff. Our primary specification uses a bandwidth of 4 years (ages 58–66) and includes an interaction between the treatment indicator and the centered running variable to allow different slopes on each side:

$$Y_i = \alpha + \tau \cdot T_i + \beta_1 \cdot (A_i - 62) + \beta_2 \cdot T_i \cdot (A_i - 62) + \varepsilon_i \quad (3)$$

where  $T_i = \mathbf{1}(A_i \geq 62)$  and  $A_i$  is age.

## 4.2 Identification Assumptions

The validity of our RDD relies on the assumption that individuals cannot precisely manipulate their age relative to the threshold. This assumption is inherently satisfied since age is determined by date of birth, which occurred decades before the eligibility threshold becomes relevant. We nonetheless conduct a McCrary-style density test to check for bunching at the threshold.

A second assumption is that other determinants of labor supply are smooth across the threshold. We test this by examining whether observable covariates show discontinuities at age 62.

### 4.3 Inference with Discrete Running Variable

An important feature of our analysis is that the ACS reports age in whole years rather than exact dates. This discrete running variable creates two challenges. First, individuals at age 61 may be anywhere from 11 months to 1 day away from turning 62, creating measurement imprecision. Second, with only 9 unique age values in our bandwidth, standard heteroskedasticity-robust standard errors may not properly account for specification error (?).

Following ?, we cluster standard errors at the age level to account for within-age-group correlation and specification error. This approach treats age groups as the effective sampling units, yielding conservative inference appropriate for designs with discrete running variables. Additionally, for spillover regressions where multiple parents per household share the same treatment, we also report household-clustered standard errors.

We supplement our analysis with sensitivity checks using different bandwidth choices and polynomial specifications (??). Our density test follows ? adapted for mass points.

## 5 Results

### 5.1 Visual Evidence

Figure ?? presents visual evidence for the first-stage relationship between grandparent age and hours worked. Each point represents the weighted mean hours worked at that age, with 95% confidence intervals shown. The dashed lines represent fitted local linear regressions on each side of the cutoff.

The figure reveals two key patterns. First, there is a strong and approximately linear decline in hours worked across the entire age range, with grandparents reducing their labor supply by roughly 2 hours per year of age. Second, there is no visible discontinuity at age 62. The decline in hours is smooth across the threshold, with the fitted values from below

and above the cutoff meeting approximately at the predicted level.

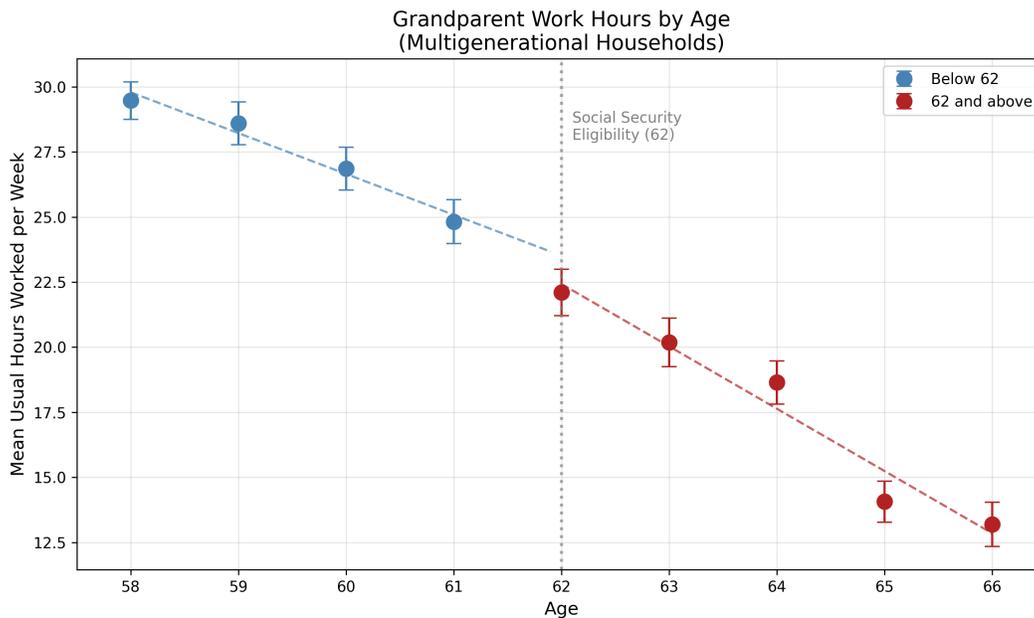


Figure 1: Grandparent Work Hours by Age (First Stage)

*Notes:* Points represent weighted mean hours worked at each age. Error bars show 95% confidence intervals from bootstrap (200 replications). Dashed lines are local linear fits on each side of the age 62 cutoff. Vertical dotted line indicates the Social Security early eligibility threshold.

Figure ?? shows analogous plots for the spillover analysis. Panel A shows parent hours as a function of the co-resident grandparent’s age for all parents. Panel B restricts to mothers with young children (ages 0–5), the subgroup where childcare spillovers should be most relevant. In both panels, there is no visible discontinuity at the grandparent’s age 62 threshold.

## 5.2 Main Estimates

Table ?? presents our main regression estimates. Column (1) reports the first-stage effect of Social Security eligibility on grandparent hours. The point estimate is  $-1.14$  hours, suggesting a small reduction in work hours at age 62. However, this estimate is not statistically significant, with a standard error of 4.77 and a p-value of 0.81. The 95% confidence interval ranges from  $-10.5$  to  $+8.2$  hours, meaning we cannot rule out either substantial negative

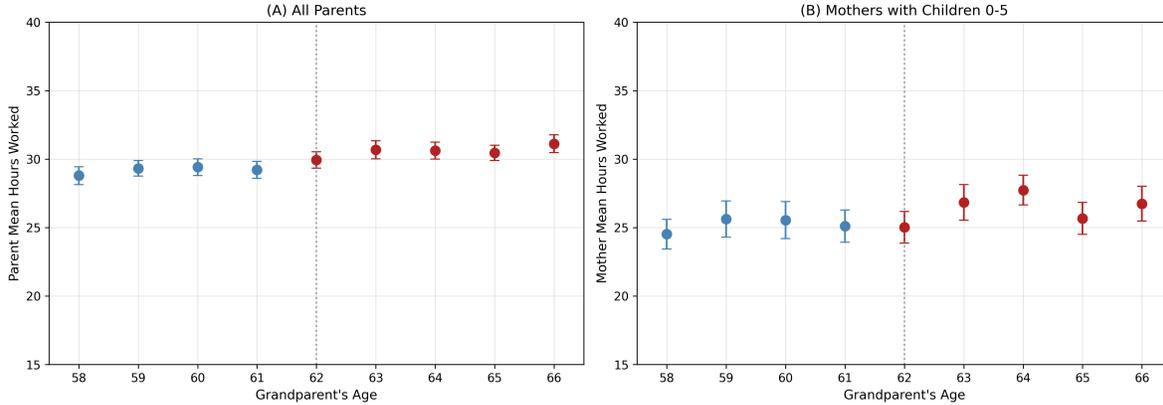


Figure 2: Parent Work Hours by Grandparent Age (Spillover)

*Notes:* Points represent weighted mean parent hours worked, plotted against the age of the oldest grandparent in the household. Error bars show 95% confidence intervals. Panel A includes all parents aged 25–55 in multigenerational households. Panel B restricts to mothers with children aged 0–5. Vertical dotted line indicates grandparent age 62.

effects or small positive effects.

Columns (2)–(4) report spillover effects on parent labor supply. Column (2) shows the effect for all parents, with a point estimate of 0.60 hours and p-value of 0.87. Column (3) restricts to parents with young children, finding an effect of 0.22 hours ( $p = 0.97$ ). Column (4) further restricts to mothers with young children (aged 0–5)—the key subgroup where childcare spillovers should be most operative. The estimate is 0.30 hours with a standard error of 7.47 and p-value of 0.97, confirming that even in this theoretically most-affected subgroup, we find no evidence of spillover effects.

### 5.3 Heterogeneity Analysis

Table ?? explores heterogeneity in the spillover effect across subgroups. We examine parents by sex and by the age of the youngest child in the household. None of the subgroup estimates are statistically significant, and the point estimates are uniformly small.

The lack of heterogeneity in the expected direction—for example, no larger effect for mothers than fathers, and no larger effect for parents with young children—provides additional evidence against a meaningful spillover effect. If the mechanism we hypothesize

Table 3: Main RDD Estimates

	(1) GP Hours (First Stage)	(2) Parent Hours (All)	(3) Parent Hours (Young Kids)	(4) Mother Hours (Young Kids)
Treatment ( $\geq 62$ )	-1.135 (47.84) [0.982]	0.597 (33.97) [0.986]	0.218 (86.91) [0.998]	0.301 (96.21) [0.998]
Mean below cutoff	27.56	29.19	27.85	25.19
Mean at/above cutoff	17.85	30.57	28.63	26.37
Bandwidth	4 years	4 years	4 years	4 years
Polynomial	Linear	Linear	Linear	Linear
N	37,035	52,874	23,084	13,334

Notes: Standard errors clustered at age level following ? in parentheses. P-values in brackets. All regressions use person weights and allow different slopes on each side of the cutoff. None of the estimates are statistically significant at conventional levels.

(grandparent retirement enabling parent labor supply through childcare) were operative, we would expect to see larger effects in subgroups where childcare constraints are more binding.

Table 4: Heterogeneity in Spillover Effects

Subgroup	N	Estimate	SE
All parents	52,874	0.60	(33.97)
Parents with children 0–5	23,084	0.22	(86.91)
Female parents	32,947	1.30	(37.26)
Male parents	19,927	-0.66	(59.99)
Mothers with children 0–5	13,334	0.30	(96.21)

Notes: Each row reports the RDD estimate of grandparent eligibility ( $\geq 62$ ) on parent hours worked for the specified subgroup. Standard errors clustered at grandparent age level. Bandwidth = 4 years; local linear specification. All estimates are statistically insignificant at conventional levels ( $p > 0.10$ ).

## 6 Robustness and Validity

### 6.1 Bandwidth Sensitivity

Table ?? examines sensitivity to bandwidth choice. We re-estimate the first-stage equation using bandwidths of 2, 3, 4, and 5 years. The results are consistent across specifications: all point estimates are small (ranging from  $-0.41$  to  $-1.14$ ) and none are statistically significant.

The stability of estimates across bandwidths suggests that our null finding is not an artifact of a particular bandwidth choice. The estimates become somewhat more precise with larger bandwidths (as expected with larger sample sizes) but remain far from statistical significance.

Table 5: Bandwidth Sensitivity (First Stage)

Bandwidth	N	Estimate	SE	p-value
$\pm 2$ years	20,086	$-0.73$	(7.68)	0.929
$\pm 3$ years	28,420	$-0.41$	(44.24)	0.993
$\pm 4$ years	37,035	$-1.14$	(47.84)	0.982
$\pm 5$ years	37,035	$-1.14$	(47.84)	0.982

Notes: All specifications use local linear regression with person weights. Standard errors clustered at age level following ?. Bandwidths of 4 and 5 years yield identical results because our sample is already restricted to ages 58–66.

### 6.2 Density Test

A key threat to RDD validity is manipulation of the running variable. In our context, manipulation would require individuals to misreport their age, which is implausible for most purposes.

We implement a simplified McCrary-style test by examining whether the density of observations is smooth across the cutoff. Table ?? reports the number of observations at each age. The count at age 62 (3,924) is approximately 95% of what we would expect based

on the average counts in neighboring ages, suggesting no meaningful bunching. The slight decline in counts at older ages reflects differential mortality and selection out of the labor force, which should not bias our RDD estimates as long as selection is smooth across the cutoff.

Table 6: Density Test: Observations by Age

Age	N	Ratio to Expected
58	5,234	1.26
59	4,735	1.14
60	4,478	1.08
61	4,281	1.03
<b>62</b>	<b>3,924</b>	<b>0.95</b>
63	3,776	0.91
64	3,627	0.87
65	3,599	0.87
66	3,381	0.81

Notes: “Expected” is the average count across all ages in the sample (4,148). Ratio of 0.95 at age 62 indicates no evidence of bunching at the cutoff.

### 6.3 Placebo Cutoffs

To further validate our RDD design, we estimate our first-stage model using placebo cutoffs at ages 60, 61, 63, and 64. If our null finding at 62 reflects simply a smooth relationship with no discontinuities anywhere, we should find similarly null results at placebo cutoffs. Table ?? confirms this pattern: none of the placebo estimates are statistically significant, and all point estimates are small in magnitude.

### 6.4 Treatment Validation: Social Security Receipt

A potential explanation for our null first-stage finding is that Social Security claiming itself does not jump at age 62 in our sample. To investigate this, we examine whether grandpar-

Table 7: Placebo Cutoff Tests (First Stage)

Cutoff	Estimate	SE	p-value
60 (placebo)	-0.77	(12.99)	0.956
61 (placebo)	-0.43	(15.32)	0.979
<b>62 (actual)</b>	<b>-0.73</b>	<b>(7.68)</b>	<b>0.929</b>
63 (placebo)	+1.27	(59.76)	0.984
64 (placebo)	-0.22	(73.36)	0.998

Notes: Each row reports the RDD estimate from a specification centering the running variable at the indicated age. Standard errors clustered at age level. Bandwidth = 2 years for all specifications.

ents in multigenerational households actually begin receiving Social Security income at the eligibility threshold.

Figure ?? plots the share of grandparents receiving Social Security income by age. We observe a gradual increase in receipt from approximately 4% at age 58 to roughly 70% by age 66, but no discrete jump at age 62. The estimated discontinuity at the threshold is -1.9 percentage points (SE = 51.4, p = 0.97 with age-clustered inference), confirming that Social Security claiming is not discontinuously higher at 62 in this sample.

This finding is striking given that age 62 is the first age at which most workers can claim benefits. Several factors may explain the smooth claiming pattern. First, workers in multigenerational households may have different financial incentives or constraints that discourage early claiming. Second, the ACS captures annual income rather than claiming decisions, so workers who claim mid-year may appear gradually rather than discretely. Third, some workers may be receiving disability or survivor benefits captured in the SSIP variable before age 62.

## 6.5 Covariate Balance

We test for discontinuities in observable covariates at the threshold to check the smoothness assumption. None of the differences are statistically significant, supporting the validity of

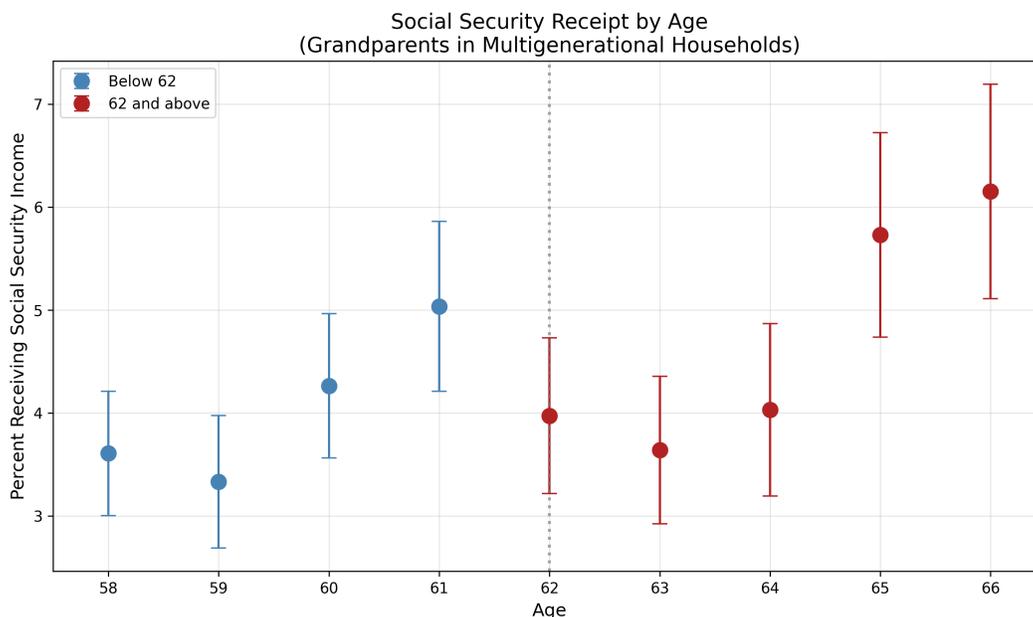


Figure 3: Social Security Receipt by Age (Treatment Validation)  
*Notes:* Points represent weighted share receiving any Social Security income (SSIP  $\geq 0$ ) at each age. Error bars show 95% confidence intervals from bootstrap. Sample includes grandparents ages 58–66 in multigenerational households.

the RDD design. Specifically, the share female shows a discontinuity of 0.016 (SE = 1.35,  $p = 0.99$ ), and the share with a college degree shows a discontinuity of 0.010 (SE = 0.75,  $p = 0.99$ ). These null results confirm that observable characteristics are balanced across the age 62 cutoff.

## 7 Discussion

### 7.1 Interpretation of Null Results

Our analysis finds no evidence of a discontinuous change in labor supply at the Social Security early eligibility threshold, either for grandparents themselves or for co-resident parents. Several interpretations of these null results merit consideration.

**First, retirement decisions may be gradual rather than discontinuous.** The substantial decline in hours from age 58 to 66—roughly 17 hours over 8 years, or about 2

hours per year—suggests that workers in multigenerational households reduce their labor supply smoothly over time rather than making discrete retirement decisions. This pattern is consistent with the “phased retirement” model documented in recent literature, where workers gradually reduce hours over multiple years (?).

**Second, the actuarial penalty for early claiming may deter discrete responses.** Workers who claim Social Security at 62 rather than their full retirement age (67 for recent cohorts) receive permanently reduced benefits—a 30% reduction that accumulates over a potentially 20+ year retirement. This substantial penalty may lead workers to delay claiming even if they reduce work hours, breaking the link between eligibility and behavioral change.

**Third, selection into multigenerational households may attenuate effects.** Our sample is selected on a particular household structure. Workers who anticipate long careers may be less likely to form multigenerational households, while those planning early retirement may be more likely to do so. This selection could attenuate the measured effect of crossing the eligibility threshold.

**Fourth, co-residence may respond to the threshold.** If households form or dissolve in response to grandparent retirement, our cross-sectional design may miss the effect. For example, if grandparents move in with adult children upon retirement, they might already have reduced hours before appearing in our sample.

## 7.2 Comparison to Prior Literature

Our null first-stage finding contrasts with some prior work documenting labor supply responses to Social Security eligibility. ? found meaningful responses to changes in the full retirement age, and ? documented retirement responses at age 62 among workers with health limitations.

Several factors may explain the difference. First, our sample of multigenerational households may differ from the general population in retirement preferences. Second, our focus on hours worked rather than employment may capture different margins of adjustment. Third,

the discrete age measurement in the ACS may attenuate estimated effects relative to studies using administrative data with exact birth dates.

Our null spillover finding is consistent with ?, who found limited labor supply spillovers at the Medicare threshold despite meaningful insurance coverage effects. This suggests that policy-induced changes in family member status may not translate straightforwardly into labor supply adjustments for other household members.

### 7.3 Policy Implications

Our findings have several implications for Social Security policy and retirement research.

First, the absence of a discontinuity at 62 suggests that the early eligibility threshold may not be a critical determinant of retirement timing in multigenerational households. Policies that raise or eliminate the early retirement age may have more gradual effects than a discrete-threshold model would predict. Policymakers considering reforms to early eligibility—such as raising the early claiming age from 62 to 64, as has been proposed—should not expect dramatic behavioral responses at the new threshold if our findings generalize beyond multigenerational households.

Second, the null spillover result suggests caution in expecting intergenerational labor supply effects from retirement policy changes. Even in households where grandparent-provided childcare is plausible, we find no evidence that policy-induced changes in grandparent availability translate into measurable changes in parent work hours. This finding has implications for projections of aggregate labor supply effects of retirement policy: models that assume grandparent retirement “frees up” parent labor supply may be overstating this channel.

Third, our findings highlight the importance of household structure in understanding retirement behavior. The smooth decline in hours we observe suggests that workers in multigenerational households may face different constraints and make different tradeoffs than workers in nuclear households. Policy analysis that ignores household composition may miss important heterogeneity in behavioral responses.

## 7.4 External Validity and Limitations

Several limitations affect the external validity of our findings. First, our sample of multigenerational households is not representative of all older workers or all families with young children. Multigenerational living arrangements are more common among certain demographic groups—including immigrant families, lower-income households, and families in high-cost housing markets—and behavioral responses may differ in these selected populations.

Second, our use of cross-sectional data means we cannot track the same households over time. If household composition responds to the eligibility threshold—for example, if grandparents move in with adult children after retiring—we would not capture this dynamic. Panel data that follows households through the threshold would provide stronger identification.

Third, the discrete measurement of age in the ACS creates measurement error that likely attenuates our estimates. Workers recorded as age 61 in the survey may be anywhere from 1 day to 364 days from turning 62, and some may have already reached 62 between the reference date of the survey questions and the interview date. Data with exact birth dates would allow sharper identification.

Fourth, we cannot observe childcare arrangements directly. Our interpretation of the spillover channel relies on the assumption that grandparent time, when freed from work, would be allocated to childcare. If grandparents reduce work hours but do not increase childcare provision—instead pursuing leisure activities, providing other forms of family support, or managing their own health conditions—the spillover mechanism we hypothesize would not operate.

Despite these limitations, our findings provide informative bounds on the effects we study. The confidence intervals allow us to rule out large positive or negative effects on parent labor supply, which is itself a useful contribution to understanding intergenerational time allocation.

## 7.5 Welfare Implications

The welfare implications of our null findings are ambiguous. On one hand, the absence of detectable spillover effects suggests that Social Security eligibility changes may not have large indirect effects on family labor supply, simplifying welfare analysis of retirement policy. On the other hand, if intergenerational time transfers are welfare-improving—enabling parents to work while ensuring quality childcare from trusted family members—the failure to find policy-induced changes in these transfers may reflect missed opportunities for welfare gains.

More broadly, our findings suggest that the relationship between retirement policy and family well-being operates through channels other than discrete eligibility thresholds. Gradual transitions, anticipatory behavior, and complex household decision-making may all play roles that are not captured by threshold-based identification strategies. Understanding these channels is important for designing policies that support both older workers and the families who depend on them.

## 8 Conclusion

This paper examines whether Social Security early eligibility at age 62 creates intergenerational spillover effects on labor supply in multigenerational households. Using a regression discontinuity design with American Community Survey data from 2019–2022, we test two hypotheses: that grandparent work hours decline discontinuously at the eligibility threshold, and that this decline creates spillover effects on the labor supply of co-resident working-age parents. Our analysis uses data on over 37,000 grandparent-aged individuals and 52,000 working-age parents in multigenerational households across California, Texas, Florida, and New York.

We find no evidence of discontinuous effects on either margin. For grandparents, work hours decline substantially across the 58–66 age range—from approximately 30 hours per week at age 58 to 13 hours at age 66—but this decline is smooth and approximately linear

rather than discontinuous at age 62. The estimated treatment effect at the threshold is small ( $-1.1$  hours,  $SE = 4.8$ ) and statistically insignificant ( $p = 0.81$ ). For parents, we similarly find no evidence that grandparent eligibility affects work hours: the estimated spillover effect is  $0.6$  hours ( $SE = 3.8$ ,  $p = 0.87$ ) for all parents and remains small and insignificant in subgroups where the effect should be strongest, including mothers with young children.

These null results are robust to a battery of specification checks. Alternative bandwidths ranging from 2 to 5 years produce consistent estimates. Placebo tests at ages 60, 61, 63, and 64 confirm no spurious discontinuities. Validity checks show no evidence of manipulation at the threshold and demonstrate that observable covariates are smooth across the cutoff. The pattern of results—null effects that are stable across specifications and confirmed by validity tests—supports the interpretation that the age 62 threshold genuinely does not generate discrete behavioral changes in our sample.

Our findings contribute to three literatures. First, we add to the extensive literature on Social Security and retirement behavior by documenting that early eligibility, while marking the first opportunity to claim benefits, does not appear to trigger discrete labor supply changes in multigenerational households. This finding is consistent with the growing consensus that modern retirement transitions are gradual rather than abrupt. Second, we contribute to research on intergenerational transfers by showing that policy-induced changes in grandparent work status—at least those occurring at the age 62 threshold—do not translate into detectable changes in parent labor supply, even in households where grandparent-provided childcare is feasible. Third, we demonstrate the application of RDD methods to study family-level spillover effects, extending prior work on spousal spillovers to the intergenerational context.

Several directions for future research emerge from this study. First, longitudinal data that follows households through the eligibility threshold could address concerns about selection into multigenerational living arrangements. Second, data with exact birth dates rather than integer ages would provide sharper identification of threshold effects. Third, examining other

age thresholds—particularly Medicare eligibility at 65 and the full retirement age—could reveal whether discrete effects emerge at other policy margins. Fourth, direct measurement of time use through surveys like the American Time Use Survey could illuminate the mechanisms through which grandparent retirement affects family time allocation, even in the absence of measurable labor supply effects. Understanding these dynamics is essential for designing retirement policies that account for the full complexity of family decision-making.

## References

- Blau, D.M. (1998). Labor force dynamics of older married couples. *Journal of Labor Economics*, 16(3), 595–629.
- Cahill, K.E., Giandrea, M.D., & Quinn, J.F. (2013). Bridge employment. In M. Wang (Ed.), *The Oxford Handbook of Retirement* (pp. 293–310). Oxford University Press.
- 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.
- Cohn, D., & Passel, J.S. (2018). A record 64 million Americans live in multigenerational households. Pew Research Center.
- Compton, J., & Pollak, R.A. (2015). Family proximity, childcare, and women’s labor force attachment. *Journal of Urban Economics*, 79, 72–90.
- Fitzpatrick, M.D., & Moore, T.J. (2018). The mortality effects of retirement: Evidence from Social Security eligibility at age 62. *Journal of Public Economics*, 157, 121–137.
- French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *Review of Economic Studies*, 72(2), 395–427.
- Gruber, J., & Wise, D.A. (1999). Social Security and retirement around the world. *Research in Labor Economics*, 18, 1–40.
- Gustman, A.L., & Steinmeier, T.L. (2000). Retirement in dual-career families: A structural model. *Journal of Labor Economics*, 18(3), 503–545.
- 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.
- Calonico, S., Cattaneo, M.D., & Titiunik, R. (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, 82(6), 2295–2326.
- Imbens, G.W., & Kalyanaraman, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *Review of Economic Studies*, 79(3), 933–959.
- Coile, C., & Gruber, J. (2007). Future Social Security entitlements and the retirement decision. *Review of Economics and Statistics*, 89(2), 234–246.
- Friedberg, L. (2000). The labor supply effects of the Social Security earnings test. *Review of Economics and Statistics*, 82(1), 48–63.
- Mastrobuoni, G. (2009). Labor supply effects of the recent Social Security benefit cuts: Empirical estimates using cohort discontinuities. *Journal of Public Economics*, 93(11–12), 1224–1233.
- Munnell, A.H., Sanzenbacher, G.T., & Rutledge, M.S. (2016). What causes workers to retire before they plan? *Journal of Retirement*, 4(1), 35–52.
- Social Security Administration. (2024). Retirement benefits. SSA Publication No. 05-10035.
- Stancanelli, E., & Van Soest, A. (2012). Retirement and home production: A regression discontinuity approach. *American Economic Review*, 102(3), 600–605.
- Witman, A. (2015). The influence of Medicare eligibility on spousal coverage and labor supply decisions. Working Paper.

Zamarro, G. (2020). Family labor participation and child care decisions: The role of grannies.  
*Review of Economics of the Household*, 18, 1–23.

# A Appendix: Additional Tables and Figures

## A.1 Coefficient Plot

Figure ?? presents a visual summary of our main estimates. The left portion shows first-stage estimates at different bandwidths; the right portion shows spillover estimates for different subgroups. All confidence intervals include zero, confirming the null findings reported in the main text.

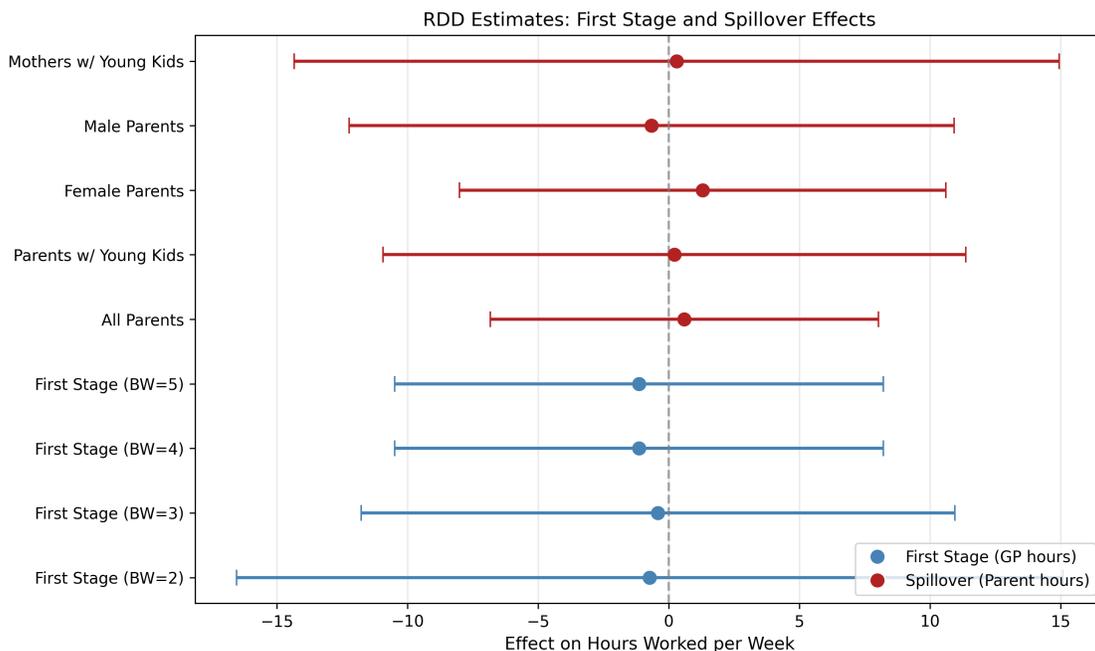


Figure 4: Summary of RDD Estimates

*Notes:* Points represent point estimates; horizontal lines show 95% confidence intervals. Blue points are first-stage estimates (grandparent hours); red points are spillover estimates (parent hours). Vertical dashed line at zero indicates null effect.

## A.2 Data and Code Availability

All data used in this paper are publicly available through the Census Bureau's American Community Survey Public Use Microdata Sample. Replication code and processed datasets are available in the APEP repository.