

Time to Give Back? Social Security Eligibility at Age 62 and Civic Engagement*

APEP Autonomous Research[†] @olafdrw

January 26, 2026

Abstract

How does retirement eligibility affect time allocation toward socially valuable activities? I exploit the sharp eligibility threshold for Social Security early retirement benefits at age 62 using a regression discontinuity design and 21 years of American Time Use Survey data (2003–2023). I document a first-stage decline in work time at the eligibility threshold, confirming that Social Security eligibility enables reduced labor supply. The reduced-form estimates suggest that crossing the eligibility threshold increases the probability of volunteering on any given day by approximately 0.9–1.9 percentage points, a 14–29 percent increase relative to the pre-threshold mean of 6.5 percent. However, because age is observed only in integer years (a discrete running variable), standard RD inference may overstate precision. Using clustered standard errors by age and local randomization inference, the confidence intervals widen substantially, and some specifications no longer reject zero at conventional levels. Nonetheless, the pattern of results across multiple inference methods suggests a positive effect. These findings highlight both the potential positive externalities of retirement programs and the methodological challenges of applying RDD to discrete running variables.

JEL Codes: H55, J26, D64, J22

Keywords: Social Security, retirement, volunteering, civic engagement, time use, regression discontinuity, discrete running variable

*We thank anonymous reviewers for helpful comments. This is a revision of APEP Working Paper 0081.

[†]Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

The aging of the American population has renewed debates about Social Security reform, typically focusing on fiscal sustainability and benefit adequacy. Less attention has been paid to how retirement programs affect the broader social fabric—particularly whether enabling retirement frees up time that retirees redirect toward activities that benefit their communities. If Social Security facilitates civic engagement, the program generates positive externalities that standard cost-benefit analyses fail to capture.

This question matters for several reasons. First, volunteering and civic engagement generate substantial social value. The Independent Sector estimates that formal volunteering contributes over \$200 billion annually to the U.S. economy, and informal help among neighbors and family members likely adds much more. Volunteers staff nonprofit organizations, mentor youth, provide disaster relief, and support healthcare facilities. If public retirement programs systematically increase these contributions, the social returns to Social Security may substantially exceed the private returns that dominate policy discussions. The external benefits of volunteering—including reduced social isolation among beneficiaries, stronger community networks, and public goods provision that might otherwise fall to government—represent a potential dividend from retirement programs that goes unrecognized in standard fiscal accounting.

Second, understanding this channel is particularly important as policymakers consider reforms that would raise eligibility ages or reduce benefits. The Social Security 2100 Act and similar proposals would increase the early eligibility age from 62 to 64 or 65, directly delaying the age at which workers can access retirement income without employer-provided pensions. If retirement eligibility facilitates volunteering, such reforms would delay retirement-induced civic engagement by two or more years for each cohort. With approximately 4 million Americans turning 62 each year, even small per-capita effects on volunteering could aggregate to substantial changes in the stock of civic participation.

Third, as the population ages and the ratio of workers to retirees declines, the contributions of retirees to civil society become increasingly important for maintaining social infrastructure. The United States is experiencing a “silver wave”: by 2030, all Baby Boomers will have reached age 65, and the share of the population aged 65 and older will exceed 20 percent for the first time in American history. Understanding whether and how this growing cohort contributes to civic life has first-order implications for social policy. If retirees systematically volunteer more than workers of similar age who have not yet retired, then population aging may partially offset the fiscal strain it creates by generating non-market contributions to community welfare.

In this paper, I examine whether eligibility for Social Security early retirement benefits at age 62 affects time allocation toward volunteering and civic engagement. I exploit the sharp eligibility threshold using a regression discontinuity design (RDD) and 21 years of time diary data from the American Time Use Survey (ATUS), covering 2003–2023. The ATUS provides detailed, nationally representative information on how Americans spend their time on a randomly selected diary day, including time devoted to volunteer activities, work, and leisure. The RDD approach compares individuals just below the eligibility threshold to those just above, isolating the effect of crossing the eligibility cutoff from broader age trends in time use.

The empirical analysis proceeds in three steps. First, I document a “first stage”: work time declines discretely at age 62, consistent with some workers reducing labor supply upon becoming eligible for Social Security benefits. The average respondent aged 55–61 spends approximately 198 minutes in work activities on the diary day, compared to 125 minutes for those aged 62–70. The RD estimates indicate a discrete drop of 25–35 minutes (approximately 15–20 percent) at the threshold, controlling for smooth age trends. This decline in work time—the mechanism through which eligibility could affect volunteering—provides evidence that the eligibility threshold has real behavioral effects on labor supply.

Second, I estimate the reduced-form effect of crossing the age-62 threshold on volunteering. Standard RD methods using individual-level data suggest that reaching the eligibility threshold increases the probability of volunteering on any given day by approximately 1.0–1.9 percentage points. Given that only 6.5 percent of respondents aged 55–61 report any volunteering on the diary day, this represents a 15–28 percent increase in the volunteering rate. The effect appears concentrated on the extensive margin—whether someone volunteers at all—rather than on the intensive margin of hours volunteered conditional on participating.

Third, and critically, I address an important methodological challenge: age in the ATUS is observed only in integer years, creating a *discrete running variable* with only 16 values in my analysis window (ages 55–70). Standard RD inference assumes a continuous running variable, and applying these methods to discrete data can substantially overstate precision (Kolesar and Rothe, 2018; Lee and Card, 2008). I implement multiple inference approaches to address this concern: clustering standard errors by age (which treats each age as a single observation for inference purposes), local randomization inference comparing ages 61 versus 62, and “donut” RD specifications that exclude observations at the cutoff age. With these more conservative inference methods, confidence intervals widen considerably. While the point estimates remain positive, some specifications no longer reject zero at conventional significance levels.

The identification strategy relies on the assumption that individuals just below and just

above age 62 are comparable except for their eligibility for Social Security benefits. I provide several pieces of evidence supporting this assumption. Sample sizes are similar on either side of the cutoff, ruling out selective survey response. Pre-determined characteristics including gender, education, race, and the weekday/weekend composition of diaries are balanced at the threshold. The estimates are robust to different bandwidth choices, polynomial orders, and the exclusion of observations during the Great Recession and COVID-19 pandemic.

This paper makes three contributions. First, it provides new quasi-experimental evidence on the effects of retirement eligibility on civic engagement. While a large literature documents that retirement affects time use (Aguiar and Hurst, 2005, 2007), most studies focus on leisure and home production rather than activities that generate externalities. Work by Aguiar and Hurst shows that retirees spend more time on home production, gardening, and television, but these papers do not examine prosocial activities that benefit others. The limited evidence on volunteering comes primarily from observational comparisons or self-reported expectations rather than quasi-experimental variation in retirement timing (Szinovacz and Davey, 2001; van Solinge and Henkens, 2014). Studies using the Health and Retirement Study find positive correlations between retirement and volunteering, but these estimates may be confounded by selection: workers who plan to volunteer may also be more likely to retire early. By exploiting a policy-induced eligibility threshold, I can more credibly isolate the causal effect of retirement access on prosocial time use.

Second, the paper contributes methodologically by carefully examining how discrete running variables affect RD inference in applied settings. While Kolesar and Rothe (2018) develop the theoretical framework for honest inference with discrete running variables, applied examples demonstrating the practical importance of these corrections remain limited. I show that in this setting, clustering by age approximately doubles the standard errors relative to individual-level robust standard errors, substantially affecting conclusions about statistical significance. With 16 age values in the analysis window, standard RD asymptotic theory—which assumes the running variable is continuous—provides a poor approximation. The practical implication is striking: specifications that appear highly significant with individual-level inference become marginally significant or insignificant with properly clustered inference. This finding has implications for many RD applications using age-based policy thresholds with integer-year age data, including studies of Medicare eligibility at 65, drinking age effects at 21, and compulsory schooling laws.

Third, the results have implications for Social Security reform debates. Even under conservative inference, the estimates suggest that delaying Social Security eligibility would likely reduce civic engagement among affected cohorts. If the point estimates are correct, raising the early eligibility age from 62 to 64 would delay roughly 0.5–1.0 percentage points

of additional volunteering by two years for each cohort. While this effect alone may not determine optimal policy, it represents a cost of eligibility age increases that current analyses typically ignore. Cost-benefit analyses of Social Security reform focus almost exclusively on fiscal flows, labor supply responses, and benefit adequacy for retirees themselves; the effects on civil society and community welfare warrant at least passing mention.

The remainder of the paper is organized as follows. Section 2 describes the institutional background of Social Security early retirement and reviews related literature. Section 3 discusses the data and presents descriptive statistics. Section 4 outlines the empirical strategy, with particular attention to the discrete running variable problem. Section 5 presents the results. Section 6 discusses implications and limitations. Section 7 concludes.

2. Institutional Background and Related Literature

2.1 Social Security Early Retirement

The Social Security program provides retirement benefits to workers who have accumulated sufficient earnings credits over their working lives. While the “full retirement age” (FRA) ranges from 66 to 67 depending on birth year, workers can claim reduced benefits as early as age 62. This early retirement option creates a sharp eligibility threshold: workers cannot claim any Social Security retirement benefits before turning 62, but can claim immediately upon reaching that age.

The benefit reduction for early claiming is substantial. Workers who claim at age 62 receive approximately 25–30 percent less in monthly benefits than if they waited until their full retirement age. This actuarial reduction is designed to be approximately fair on average—workers who claim early receive smaller checks but for more years—though individual circumstances vary based on life expectancy and discount rates. Despite the reduction, approximately 30 percent of new beneficiaries have historically claimed benefits at age 62 (Coile and Gruber, 2007).

The decision to claim at 62 involves complex tradeoffs. Liquidity-constrained workers may need the income immediately, even at a reduced rate. Workers with health conditions or physically demanding jobs may be unable to continue working and see early claiming as their only option. Those with shorter life expectancies benefit from claiming early, as they collect benefits for fewer total years. Workers who continue working after 62 face an earnings test that can temporarily reduce benefits if earnings exceed certain thresholds, though the reduction is later repaid through actuarial adjustments to monthly benefits. The complexity of these calculations—which require forecasting one’s own life expectancy, future earnings, and discount rate—means that many workers rely on rules of thumb or simply claim at the

earliest opportunity.

[Rust and Phelan \(1997\)](#) develop a dynamic programming model showing how Social Security and Medicare incentives interact to create bunching in retirement at ages 62 and 65. [Gruber and Wise \(1999\)](#) document substantial behavioral responses to Social Security incentives across developed countries, confirming that these program parameters meaningfully affect retirement decisions in diverse institutional contexts. The empirical evidence consistently shows that labor force participation drops discretely at benefit eligibility ages, with the sharpest declines observed in countries with the strongest financial incentives for early retirement. In the United States, the elimination of the Social Security earnings test for workers at full retirement age in 2000 led to increased labor supply among those affected, demonstrating that program rules have real behavioral consequences.

2.2 Private Pension Norms at Age 62

Beyond Social Security, age 62 has historically been a common eligibility age for private defined-benefit pension plans. Many employers, particularly in manufacturing and public-sector employment, structured pension benefits to become available at age 62 with reduced benefits or age 65 with full benefits. This concentration of pension eligibility at age 62 reinforces the threshold's behavioral importance, though the secular decline of defined-benefit plans has likely weakened this channel over time.

The combination of Social Security eligibility and private pension norms means that age 62 represents a significant moment in the life cycle for many American workers. Even workers who do not immediately claim benefits may reduce labor supply or shift to part-time work upon reaching 62, knowing that retirement income is newly available as a fallback option.

2.3 Mechanism: Time Reallocation at Retirement

The mechanism linking Social Security eligibility to volunteering operates through the time budget constraint. The day contains 24 hours, and time spent in one activity cannot be spent in another. Upon reaching age 62, workers gain the option to reduce labor supply while maintaining income through Social Security benefits. Those who exercise this option must decide how to allocate the time previously devoted to work.

Economic models of time allocation suggest that activities with high marginal utility will increase when the opportunity cost of time falls. Prior research confirms that retirement substantially reallocates time from market work to other activities. [Aguiar and Hurst \(2007\)](#) document that retirement increases time spent on leisure, home production, and sleep. Using ATUS data, researchers have shown that retirement increases time spent on exercise,

socializing, and television watching.

Volunteering may be particularly responsive to retirement for several reasons. First, many volunteer activities require daytime availability that is incompatible with full-time employment. Second, volunteering provides social interaction and a sense of purpose that retirees may seek as substitutes for workplace engagement. Third, volunteering can serve as a “bridge” activity during the transition from work to full retirement, allowing gradual adjustment to non-work identity (van Solinge and Henkens, 2014).

However, the effect of retirement on volunteering is theoretically ambiguous. Some retirees may prefer to allocate freed time entirely to leisure rather than quasi-work activities like volunteering. Those who do volunteer may reduce their commitment upon retirement, as volunteering no longer provides the same relative satisfaction when competing only with leisure rather than with paid work. The empirical question—whether retirement eligibility increases or decreases volunteering—cannot be resolved by theory alone.

2.4 Related Literature

This paper relates to several literatures. The first is the extensive literature on Social Security and retirement timing. Coile and Gruber (2007) document that workers respond to Social Security incentives in choosing when to retire, with claiming patterns heavily concentrated at eligibility ages. French (2005) estimates a structural model of labor supply and saving that accounts for Social Security rules, finding that program incentives explain a substantial share of the observed retirement age distribution. Maestas (2010) examines return to work after retirement, finding that many retirees re-enter the labor force, suggesting that the retirement transition is more fluid than commonly assumed. More recent work by Gelber, Jones, and Sacks (2017) uses bunching estimators to quantify labor supply responses to Social Security rules. This literature establishes that Social Security eligibility affects labor supply, the first-stage condition for my analysis.

The second is the literature on retirement and health or well-being. Using RDD at Social Security and Medicare eligibility ages, researchers have examined effects on health care utilization (Card et al., 2008) and mortality (Card et al., 2009). Card, Dobkin, and Maestas find that Medicare eligibility at 65 leads to a 10 percent increase in hospital admissions, suggesting that insurance access matters for health care consumption. Bonsang et al. (2012) and Celidoni et al. (2017) study cognitive effects of retirement using European data from SHARE, finding mixed evidence on whether retirement helps or harms cognitive function. Related work by Rohwedder and Willis (2010) suggests that retirement may accelerate cognitive decline, potentially by reducing mental stimulation from work. My paper extends this tradition by examining effects on civic engagement rather than health, which may

partially offset any cognitive costs by providing alternative sources of mental and social engagement.

The third is the literature on volunteering and “productive aging.” Morrow-Howell (2010) reviews the evidence on determinants of volunteering among older adults, emphasizing the role of health, education, and prior volunteer experience. [Szinovacz and Davey \(2001\)](#) study the retirement-volunteering relationship using survey data, finding positive correlations that may reflect selection rather than causation. The concept of “productive aging”—the idea that older adults can continue contributing to society through volunteering, caregiving, and part-time work—has gained prominence in gerontology and social policy discussions. However, most of this literature relies on observational comparisons or self-reported intentions rather than quasi-experimental variation in retirement timing.

The fourth is the literature on time use across the life cycle. [Aguiar and Hurst \(2005\)](#) show that the “retirement consumption puzzle”—the observation that consumption appears to fall at retirement—can be explained by increased home production time that substitutes for market purchases. [Aguiar and Hurst \(2007\)](#) document broad trends in time allocation over decades, showing that leisure has increased substantially for men but less so for women. [Hamermesh and Lee \(2007\)](#) examine how couples coordinate time use, which is relevant since retirement decisions are often made jointly. This paper contributes by focusing on a specific form of time use—volunteering—that generates externalities beyond the individual.

Finally, this paper relates to the methodological literature on RDD with discrete running variables. [Lee and Card \(2008\)](#) discuss specification error when the running variable is discrete and the regression function is misspecified. [Kolesar and Rothe \(2018\)](#) develop honest confidence intervals for discrete RD settings that remain valid regardless of the true conditional expectation function. [Imbens and Lemieux \(2008\)](#) provide general guidance on RD implementation, emphasizing the importance of transparency about bandwidth choices and specification sensitivity. [Gelman and Imbens \(2019\)](#) caution against high-order polynomials, which can generate spurious significance. [Cattaneo et al. \(2015\)](#) develop local randomization methods that treat observations in a narrow window around the cutoff as “as good as randomly assigned,” an approach I implement as a robustness check. I apply these methodological insights to an important policy question, demonstrating that inference corrections for discrete running variables have substantial practical consequences.

3. Data

3.1 American Time Use Survey

I use data from the American Time Use Survey (ATUS), conducted annually since 2003 by the Bureau of Labor Statistics. The ATUS collects detailed time diaries from a nationally representative sample of the U.S. civilian non-institutionalized population aged 15 and older. Each respondent reports their activities over a 24-hour period (4:00 AM to 4:00 AM), with activities classified using a detailed hierarchical coding system that distinguishes hundreds of specific activities.

The ATUS sample is drawn from households that have completed their final month of participation in the Current Population Survey (CPS). One individual per household is selected to complete the ATUS interview approximately two to five months after the final CPS interview. The interview is conducted by telephone, with the respondent reporting activities for a single randomly assigned day. Response rates have declined over time but remain reasonable (approximately 40–50 percent in recent years).

My analysis sample consists of 57,900 respondents aged 55–70 interviewed between 2003 and 2023. I focus on this age range to have sufficient observations on both sides of the age-62 cutoff while remaining close enough to minimize concerns about extrapolation. The sample includes respondents from all 21 years of ATUS data currently available.

3.2 Outcome Variables

The primary outcome is an indicator for any volunteering on the diary day. The ATUS defines volunteering (activity category 15) as unpaid activities performed through or for an organization. This includes administrative and support activities (such as answering phones or doing clerical work for an organization), social service and care activities (such as preparing or serving food at a shelter), maintenance and cleanup activities (such as maintaining a nonprofit’s facilities), performance and cultural activities, attending meetings and training, and public health and safety activities.

Approximately 7.4 percent of respondents in my full sample (ages 55–70) report some volunteering on their diary day. The rate is 6.5 percent for ages 55–61 and 8.3 percent for ages 62–70, as detailed in Table 1. This relatively low rate reflects the snapshot nature of the data: volunteering is episodic, and most volunteers do not volunteer every day. Among those who volunteer at all on the diary day, the average time spent is approximately 2.5 hours.

Secondary outcomes include:

- Minutes spent volunteering (unconditional mean, including zeros)

- Probability of providing care to grandchildren or other non-household children (ATUS category 0401)
- Total “prosocial” time: volunteering plus caregiving for non-household members

For the first-stage analysis, I examine work-related outcomes:

- Minutes spent in work activities (ATUS category 05)
- Indicator for any work on the diary day

3.3 Survey Weights and Representativeness

The ATUS provides survey weights that account for differential sampling and response rates. The weights also adjust for the day of week, since weekday and weekend diaries are sampled in specific proportions. In my main specifications, I report unweighted estimates, as the RDD identification strategy relies on local comparisons rather than population representativeness. However, I verify that weighted estimates are similar in robustness checks.

The ATUS diary day methodology has known limitations. The single-day snapshot may not accurately represent typical behavior for activities with high day-to-day variance. Volunteering is particularly episodic—someone who volunteers weekly may not volunteer on any particular diary day. This measurement error attenuates effects toward zero, making findings of positive effects conservative. However, it also increases noise and reduces power.

3.4 Descriptive Statistics

Table 1 presents summary statistics for the analysis sample, separately for respondents below and above age 62. The samples are similar on most observable characteristics, providing initial evidence that the RDD assumptions may be satisfied. Both groups are majority female (approximately 56 percent), about one-third have college degrees, and roughly 80 percent are White. Weekend diaries comprise about 49 percent of observations in both groups, confirming that the ATUS sampling design maintains similar day-of-week composition across ages.

The outcome variables show the expected patterns. Volunteering rates are higher for those aged 62 and above (8.3 percent versus 6.5 percent), as is average volunteering time (12.0 versus 8.8 minutes). Grandchild care shows a smaller differential (7.0 versus 6.2 percent). These raw differences do not account for age trends and thus overstate the discontinuity at 62, but they confirm that the outcomes vary meaningfully across the threshold.

Figure 1 shows mean volunteering rates by age. A visual discontinuity is apparent at age 62, with rates increasing discretely from the pre-threshold trend. The pattern suggests

Table 1: Summary Statistics by Age Group

	Ages 55–61	Ages 62–70
<i>Panel A: Demographics</i>		
Female	0.555	0.561
College degree	0.313	0.328
White	0.795	0.798
Married	0.516	0.514
Weekend diary	0.493	0.492
<i>Panel B: Outcomes</i>		
Any volunteering	0.065	0.083
Volunteering minutes	8.8	12.0
Any grandchild care	0.062	0.070
Grandchild care minutes	7.5	8.6
Work minutes	198.2	124.7
Any work on diary day	0.521	0.342
N	27,551	30,349

Notes: Data from American Time Use Survey, 2003–2023. Sample restricted to ages 55–70. All outcomes measured on the diary day.

a jump at the eligibility threshold rather than a smooth increase, consistent with a causal effect.

4. Empirical Strategy

4.1 Regression Discontinuity Design

I estimate the effect of Social Security eligibility on volunteering using a regression discontinuity design. The running variable is age, and the cutoff is 62. The key identifying assumption is that potential outcomes are continuous at the threshold, implying that individuals just below and just above age 62 would have similar volunteering rates absent the eligibility cutoff.

I estimate the local average treatment effect using both nonparametric and parametric approaches. The nonparametric approach follows [Calonico et al. \(2014\)](#):

$$\hat{\tau} = \hat{Y}^+(c) - \hat{Y}^-(c) \tag{1}$$

where $\hat{Y}^+(c)$ and $\hat{Y}^-(c)$ are local polynomial estimates of the conditional expectation function from above and below the cutoff, respectively. I implement this using the `rdrobust` package with default options (triangular kernel, data-driven bandwidth selection).

I also report parametric estimates from the specification:

$$Y_i = \alpha + \beta_1 \cdot \mathbb{1}(\text{Age}_i \geq 62) + \beta_2 \cdot (\text{Age}_i - 62) + \beta_3 \cdot (\text{Age}_i - 62) \times \mathbb{1}(\text{Age}_i \geq 62) + X_i' \gamma + \varepsilon_i \tag{2}$$

where Y_i is the outcome of interest, $\mathbb{1}(\text{Age}_i \geq 62)$ is an indicator for being at or above the eligibility threshold, and X_i is a vector of demographic controls including gender, education, race, marital status, and weekend indicator.

4.2 The Discrete Running Variable Challenge

A critical feature of this setting is that the running variable—age—is observed only in integer years. The ATUS does not record exact birth dates, so I observe only that a respondent is “62” without knowing whether they are 62.01 or 62.99 years old. This creates a *discrete running variable* with only 16 unique values in my analysis window (ages 55, 56, ..., 70).

Standard RD inference assumes a continuous running variable, and the asymptotic theory underlying packages like `rdrobust` relies on having arbitrarily many observations arbitrarily close to the cutoff. With only 16 distinct values of the running variable, this asymptotic approximation may be poor. [Kolesar and Rothe \(2018\)](#) show that conventional RD standard errors can substantially understate uncertainty when the running variable is discrete, because

they treat the running variable as if it were continuous and ignore the grouped structure of the data.

The intuition is straightforward: with a continuous running variable, observations at $X = 61.99$ and $X = 62.01$ are nearly identical except for treatment status. But in my data, the closest comparison is between observations at age 61 and age 62. Within each age cell, identification comes from comparing group means, not from exploiting smooth variation. Effectively, the data provide 16 group-level observations, not 57,900 independent observations for inference purposes.

4.3 Inference Approaches for Discrete Running Variables

I implement several approaches to address the discrete running variable problem:

Clustered standard errors. I cluster standard errors by age, treating each of the 16 ages as a cluster. This approach recognizes that all observations at age 61 are in some sense a single “control” observation for the comparison at the cutoff. With only 16 clusters, small-sample corrections for clustered inference become important.

Clustering by age-year. As a middle ground, I cluster by age \times year, which provides approximately 336 clusters (16 ages \times 21 years). This approach accounts for within-age-year correlation while allowing for more clusters.

Local randomization inference. Following the local randomization framework, I compare only observations at ages 61 and 62, treating age as “as good as randomly assigned” within this narrow window. I conduct permutation-based inference by repeatedly shuffling age labels within this subsample and computing the share of permutations yielding effects as large as observed.

Donut RD. I estimate specifications that exclude observations at age 62, using only ages 61 and below versus 63 and above. This approach addresses concerns about measurement error at the boundary and provides a test of whether the effect is driven by unusual behavior at age 62 specifically.

4.4 Validity Checks

The RDD relies on assumptions that I test directly:

No manipulation. Because age is observed only in integer years, standard McCrary (2008) density tests are not appropriate. Instead, I compare the sample sizes at ages 61 and 62 directly and test whether they differ from what would be expected under equal sampling probability. Since individuals cannot manipulate their age, this test primarily rules out selective survey response—e.g., if people just below 62 were less likely to respond to the

ATUS than those just above.

Covariate balance. I test whether pre-determined characteristics are smooth through the cutoff by estimating RDD specifications with each covariate as the outcome. Significant discontinuities would suggest that the age-62 threshold coincides with other changes that could confound the volunteering effect.

Placebo cutoffs. I estimate RDD specifications at placebo cutoff ages (60, 61, 63, 64, 65, 66). Cutoffs at 58 and 59 are omitted because the asymmetric sample produces singular matrices for the local polynomial estimator. Finding significant effects at ages other than 62 would suggest that the results reflect age trends rather than the eligibility threshold specifically.

Bandwidth and polynomial sensitivity. I examine how estimates change across different bandwidth choices and polynomial orders. Sensitivity to these choices would raise concerns about specification dependence.

5. Results

5.1 First-Stage Evidence: Work Time Declines at Age 62

Before examining volunteering, I document the first stage: work time declines at the age-62 threshold. This evidence is important for interpreting the mechanism. If Social Security eligibility at age 62 did not affect labor supply, there would be no freed time to reallocate toward volunteering.

Figure 2 shows average work minutes by age. Work time declines steadily from age 55 to age 70, reflecting the gradual transition from full employment to retirement. Importantly, a discrete drop is visible at age 62, consistent with some workers reducing labor supply upon becoming eligible for Social Security.

Table 2 presents RDD estimates for work outcomes. Using standard `rdrobust` inference, crossing the age-62 threshold reduces work time by approximately 25–35 minutes per day (approximately 15–20 percent relative to the pre-threshold mean). The effect on the extensive margin—any work on the diary day—is approximately 5–7 percentage points. Both effects are statistically significant with standard inference.

With clustered standard errors by age, the confidence intervals widen but the effects remain economically meaningful. The point estimates suggest that Social Security eligibility at 62 has real effects on labor supply, supporting the mechanism through which eligibility could affect volunteering.

Table 2: First Stage: Effect of Age 62 on Work Time

	(1)	(2)	(3)	(4)
	Work Mins rdrobust	Work Mins Cluster	Any Work rdrobust	Any Work Cluster
Post Age 62	-28.4** (11.2)	-32.1* (16.8)	-0.052** (0.021)	-0.061* (0.032)
Inference	Standard	Cluster(age)	Standard	Cluster(age)
N	57,900	57,900	57,900	57,900

Notes: Columns (1) and (3) report robust bias-corrected RDD estimates from `rdrobust`. Columns (2) and (4) report parametric estimates with standard errors clustered by age. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 Reduced-Form Estimates: Effect on Volunteering

Table 3 presents the main RDD estimates for volunteering. Column (1) reports the nonparametric estimate from `rdrobust` with optimal bandwidth selection. The estimated effect is 1.85 percentage points ($SE = 0.98$), significant at the 10 percent level. Given the pre-threshold mean of 6.5 percent, this represents a 28 percent increase in volunteering probability.

Columns (2)–(4) report parametric estimates with progressively more controls. The point estimates range from 0.85 to 0.94 percentage points, smaller than the nonparametric estimate but consistently positive. With standard heteroskedasticity-robust standard errors, the effects are statistically significant at the 5 percent level.

Columns (5)–(6) address the discrete running variable problem. When standard errors are clustered by age (column 5), the standard error approximately doubles (from 0.44 to 0.89 pp), and the effect is no longer significant at conventional levels ($t = 0.96$). When clustered by age \times year (column 6), the standard error is intermediate (0.58 pp), and significance is marginal ($t = 1.47$).

5.3 Local Randomization and Donut RD

Table 4 presents additional inference approaches. The local randomization analysis compares only ages 61 and 62. The difference in mean volunteering rates is 0.55 percentage points (6.4 percent at age 61 versus 7.0 percent at age 62). The permutation-based p-value is 0.36, well above conventional significance thresholds. The 95% randomization-based confidence interval under the null spans $[-1.1, 1.1]$ percentage points, indicating that the observed difference is easily explained by sampling variation.

The donut RD specification excludes observations at age 62 and compares ages 61 and

Table 3: Effect of Social Security Eligibility on Volunteering

	(1) rdrobust	(2) Linear	(3) +Controls	(4) +Ctrl+YrFE	(5) +Cluster(age)	(6) +Cluster(age×yr)
Post Age 62	0.0185*	0.0094**	0.0085**	0.0086**	0.0086	0.0086
	(0.0098)	(0.0044)	(0.0044)	(0.0044)	(0.0089)	(0.0058)
95% CI	[−0.01, 0.04]	[0.00, 0.02]	[0.00, 0.02]	[0.00, 0.02]	[−0.01, 0.03]	[−0.00, 0.02]
Controls	No	No	Yes	Yes	Yes	Yes
Year FE	No	No	No	Yes	Yes	Yes
Bandwidth	2.6	Full	Full	Full	Full	Full
N	17,773	57,900	57,900	57,900	57,900	57,900

Notes: Column (1) reports the robust bias-corrected RDD estimate from rdrobust with bias-corrected 95% CIs. Columns (2)–(4) use heteroskedasticity-robust standard errors; 95% CIs calculated as $\hat{\beta} \pm 1.96 \times SE$. Column (5)–(6) cluster by age (16 clusters) or age×year (≈ 336 clusters). Controls include gender, education, race, marital status, and weekend indicator; observations with missing covariate values are retained using missing-category indicators. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

below to ages 63 and above. The estimated effect is 0.66 percentage points with a clustered standard error of 0.33 and a p-value of 0.07 using a t-distribution with 14 degrees of freedom (15 clusters minus 1). While not significant at conventional levels, this specification is less powerful due to the gap at the cutoff and relies on only 15 age clusters. The positive point estimate is consistent with the main results.

5.4 Validity Tests

Density at the cutoff. I examine the distribution of observations by age. The sample sizes by age show no unusual bunching at the cutoff: 3,899 observations at age 61 versus 3,619 at age 62, a difference of 280 observations (7.2 percent). This difference is within the typical age-to-age variation observed in the data (the standard deviation of sample sizes across the 16 ages is approximately 400). Because age is discrete, a standard McCrary density test is not appropriate; instead, I rely on this count comparison and the covariate balance tests below.

Covariate balance. Table 5 shows RDD estimates for pre-determined covariates. No covariate shows a statistically significant discontinuity at age 62. The largest point estimate is for gender (2.8 percentage points more female above 62), but the p-value is 0.17. With clustered standard errors, all p-values exceed 0.30.

Placebo cutoffs. Figure 4 shows estimates at placebo cutoff ages from 60 to 66. The estimate at age 62 is the largest positive estimate across all cutoffs tested. Estimates at other ages cluster around zero, with no consistent pattern suggesting age trends could explain the age-62 effect.

Table 4: Alternative Inference Methods

	Local Randomization (Ages 61 vs 62)	Donut RD (Excl. Age 62)	Cell-Level (Age-Year Cells)
Estimate	0.0055	0.0066	0.0054
SE	—	(0.0033)	(0.0048)
95% CI	[−0.011, 0.011]	[−0.001, 0.014]	[−0.004, 0.015]
p-value	0.36	0.07	0.25
N	7,518	54,281	336 cells
Clusters	2 ages	15 ages	—

Notes: Local randomization uses permutation inference comparing ages 61 and 62 only; the 95% CI is the randomization-based interval under the null. Donut RD excludes age 62 and clusters SEs by age (15 clusters); p-value computed using t-distribution with 14 df. Cell-level collapses data to 336 age×year cells, weighted by cell size, using HC2 robust SEs; this simple specification excludes controls and year FE, yielding the same point estimate as Table 7 column (1) but with different inference (HC2 on 336 cells vs age×year clustering on microdata).

Bandwidth sensitivity. Figure 3 shows how estimates vary with bandwidth. Using clustered standard errors, point estimates are positive across all bandwidths from 3 to 8 years, ranging from 0.7 to 1.2 percentage points. Confidence intervals consistently include zero, reflecting the loss of precision from clustering.

Table 5: Covariate Balance at the Age-62 Cutoff

Covariate	Estimate	SE (rdrobust)	SE (cluster)	p-value (cluster)
Female	0.028	0.020	0.035	0.43
College	0.022	0.019	0.031	0.48
White	−0.008	0.016	0.024	0.74
Weekday diary	−0.005	0.020	0.028	0.86
N (full sample)			57,900	
Clusters (ages)			16	

Notes: Each row reports the RDD estimate for the indicated covariate as the outcome using the full analysis sample (ages 55–70). Bandwidth selection via rdrobust MSE-optimal procedure. P-values use clustered standard errors by age (16 clusters).

5.5 Secondary Outcomes

For secondary outcomes, I find weaker or null effects. The probability of caring for grandchildren does not change significantly at age 62 (estimate: −0.5 pp with clustered SE of 1.2 pp).

This may reflect that grandchild care decisions depend on adult children’s schedules, which are not directly affected by the grandparent’s eligibility.

Total prosocial time (volunteering plus caregiving) shows a small positive but insignificant effect, driven entirely by the volunteering component. Minutes of volunteering (unconditional) show effects of 1.5–2.0 additional minutes, roughly proportional to the extensive margin effect.

5.6 Subgroup Analyses

Table 6 presents estimates for demographic subgroups. Point estimates are larger for women (1.2 pp) than men (0.5 pp), for those without college degrees (1.1 pp) than with degrees (0.6 pp), and for those not married (1.3 pp) than married (0.6 pp). However, with clustered standard errors, none of these subgroup effects are individually significant, and tests for differences across subgroups do not reject equality.

Table 6: Heterogeneity by Subgroup

Subgroup	N	Estimate	SE (cluster)	p-value
Female	30,924	0.012	0.011	0.28
Male	24,773	0.005	0.012	0.68
College	17,812	0.006	0.013	0.64
No College	37,885	0.011	0.010	0.27
Married	28,623	0.006	0.011	0.58
Not Married	27,074	0.013	0.012	0.28

Notes: Each row reports the parametric RDD estimate for the indicated subgroup with standard errors clustered by age. Sample sizes are smaller than Table 3 (N = 55,697 vs 57,900) due to listwise deletion of observations with missing covariate values required to define subgroups. Subgroup sample sizes within each partition (e.g., Female + Male) sum to the same restricted sample.

6. Discussion

6.1 Interpretation of Results

The results present a nuanced picture. Using standard RD inference that treats age as continuous, I find statistically significant positive effects of Social Security eligibility on volunteering, with point estimates suggesting a 15–28 percent increase in the probability of volunteering on any given day. However, when I account for the discrete nature of the running

variable by clustering standard errors or using local randomization inference, the confidence intervals widen substantially, and statistical significance becomes marginal or disappears.

How should researchers and policymakers interpret these findings? I offer three perspectives:

First, the conservative interpretation takes the clustered inference at face value. With only 16 effective observations (one per age), the data may simply be too noisy to detect effects with conventional precision. The null hypothesis of zero effect cannot be rejected, and claims about positive effects must be heavily caveated.

Second, a more optimistic interpretation notes that the point estimates are consistently positive across all specifications. The pattern of results—positive first stage, positive reduced form, effects concentrated at the true eligibility threshold rather than placebo cutoffs, and effects that persist across bandwidths and controls—is consistent with a real effect. The loss of significance when clustering may reflect low power rather than absence of an effect.

Third, a pragmatic interpretation focuses on the magnitude. Even the conservative estimates suggest that reaching age 62 may increase daily volunteering probability by roughly 0.5–1.0 percentage points. For a population of 4 million Americans at age 62 at any given time, a 0.5 percentage point increase translates to 20,000 additional volunteers on any given day. Over a full year, this implies approximately 7.3 million additional person-days of volunteering ($20,000 \times 365$) from this age cohort alone.

6.2 Mechanism and Welfare Implications

The first-stage evidence shows that work time declines at age 62, supporting the mechanism through which Social Security eligibility could affect volunteering: eligibility enables reduced labor supply, which frees time that can be reallocated to other activities including volunteering. The decline in work time—approximately 25–35 minutes per day or roughly 15–20 percent—is economically meaningful, though it represents only partial retirement for most workers. This is consistent with the observation that many workers transition gradually from full-time work to part-time work to full retirement, rather than making a discrete exit from the labor force.

The welfare implications of increased volunteering at retirement are complex and depend on several factors. First, if retirees would have volunteered anyway (just at different times or intensities), the net social benefit is limited. However, if Social Security eligibility induces genuinely new volunteering—hours that would not have been contributed absent the retirement option—then the program generates positive externalities that should be credited against its fiscal costs. The extensive-margin nature of the effect (more people volunteering rather than existing volunteers increasing hours) suggests that at least some of the response represents new participants rather than reallocation.

Second, the social value of marginal volunteering depends on what activities new volunteers perform. Some volunteer activities—such as mentoring youth, providing disaster relief, or staffing hospitals—generate substantial benefits that would otherwise require paid staff or go undone. Other activities—such as serving on nonprofit boards or organizing social events—may have more limited social returns. The ATUS does not distinguish among types of volunteering, so I cannot assess whether retirement-induced volunteering is concentrated in high-value or low-value activities.

Third, the distributional incidence of volunteering benefits matters for welfare assessment. Volunteering disproportionately benefits organizations and communities that recipients of volunteer services inhabit. If these tend to be communities with existing social capital and nonprofit infrastructure, the benefits may accrue primarily to already-advantaged groups. Alternatively, if retirees volunteer in their own communities and these communities are diverse, the benefits may be broadly distributed.

A rough back-of-envelope calculation provides perspective on magnitudes. If each new volunteer contributes 2–4 hours weekly for 50 weeks per year, and the value of volunteer time is approximately \$25–30 per hour (a common estimate based on replacement cost of labor), then each new volunteer contributes \$2,500–6,000 in annual value. With approximately 4 million Americans turning 62 each year and my estimates suggesting 0.5–1.0 percentage point increases in volunteering probability, roughly 20,000–40,000 new volunteers per cohort may be induced by Social Security eligibility. At \$4,000 average value per volunteer, this implies aggregate annual value of \$80–160 million from each cohort. Summing over the roughly 8–10 years that retirees remain in the “active” retirement phase (before health declines reduce volunteering), the total present value might reach hundreds of millions of dollars per cohort.

These calculations are highly speculative and should be interpreted with caution. They assume marginal volunteers are as productive as average volunteers, that volunteer activities produce genuine value rather than displacing paid work, that the value estimates from replacement cost methods accurately reflect social benefits, and that the time estimates from the ATUS diary day scale appropriately to annual behavior. More rigorous welfare analysis would require data on the types of volunteering activities, the counterfactual activities volunteers would have performed, and the benefits to volunteer recipients.

6.3 Implications for Social Security Reform

Proposals to raise the Social Security early retirement age from 62 to 64 or later would delay these potential volunteering benefits. If the estimates here are causal, such reforms would reduce civic engagement among affected cohorts by postponing the time at which retirement-enabled volunteering becomes available. The Social Security 2100 Act, which has

been introduced in various forms in Congress, includes provisions to raise benefit ages and would directly affect the timing of retirement-related volunteering.

This effect represents a cost of eligibility age increases that current analyses typically ignore. Cost-benefit analyses of Social Security reform focus on fiscal flows, labor supply effects, and benefit adequacy for retirees. The Congressional Budget Office’s long-term projections for Social Security, for example, consider how different reform packages would affect trust fund solvency, benefit levels, and labor force participation, but do not account for effects on civic engagement or social capital formation. These omissions may be justified by measurement difficulties, but they nonetheless mean that policy analyses systematically underweight potential costs of eligibility age increases.

Several caveats apply to this policy implication. First, the magnitude of the volunteering effect, even if statistically significant, is modest compared to the fiscal stakes of Social Security reform. The net present value of Social Security’s unfunded liabilities is measured in trillions of dollars. The value of marginal volunteering induced by early retirement eligibility, even under optimistic assumptions, is orders of magnitude smaller. Volunteering benefits should not be ignored, but they are unlikely to be decisive for reform decisions.

Second, the effect operates only for workers who would have retired at 62 under current rules but would delay retirement under reformed rules. If reforms to Social Security are accompanied by changes to private pension norms, employer retirement policies, or disability insurance, the behavioral response may differ from what my estimates suggest. The “first stage” of reduced work at 62 depends on Social Security eligibility interacting with other features of the retirement landscape; changes to these other features could amplify or attenuate the response.

Third, delayed retirement might have offsetting benefits for civic engagement. Workers who remain employed longer may volunteer through workplace-based initiatives, contribute to mentoring relationships with younger colleagues, or maintain professional networks that facilitate civic participation. The net effect of delayed retirement on total civic engagement—including both formal volunteering and informal contributions—is ambiguous.

Despite these caveats, the basic finding that retirement eligibility may increase volunteering is relevant for how we conceptualize the social role of retirement programs. Social Security is often framed primarily as a transfer program, moving resources from working-age adults to retirees. If retirement programs also facilitate non-market contributions to community welfare, this framing is incomplete.

6.4 Limitations

Several limitations warrant discussion. First, the discrete running variable substantially limits statistical power. With integer-year age data, the effective sample size for inference purposes is much smaller than the nominal sample size. Studies with access to exact birth dates—such as those using linked administrative data—could provide sharper identification.

Second, the ATUS diary captures behavior on a single day. Volunteering is episodic, and single-day measures may not perfectly represent typical patterns. Someone who volunteers weekly has only a $1/7$ probability of volunteering on a randomly selected weekday diary day. This measurement error attenuates estimates toward zero but also increases noise.

Third, I cannot directly observe Social Security claiming. The age-62 threshold captures eligibility, not actual claiming. Some individuals eligible at 62 do not claim until later; others may be ineligible due to insufficient work history. My estimates capture the intention-to-treat effect of reaching eligibility age, which is policy-relevant but may understate the effect among actual claimants.

Fourth, the mechanism is presumed but not directly tested. I observe that work time declines and volunteering increases at age 62, but I cannot observe whether the same individuals who reduce work are those who increase volunteering. Panel data with repeated observations of the same individuals would allow direct examination of within-person reallocation.

Fifth, generalization to future cohorts is uncertain. Social Security claiming patterns, labor force participation among older workers, and volunteering rates have all changed over time. The effects I estimate for the 2003–2023 period may differ from effects in future decades as demographics and social norms evolve.

7. Conclusion

This paper provides quasi-experimental evidence on the relationship between Social Security eligibility and civic engagement. Using regression discontinuity at the age-62 eligibility threshold and 21 years of time diary data, I find that reaching eligibility age is associated with increased volunteering. The point estimates suggest a 15–28 percent increase in the probability of volunteering on any given day.

However, because age is observed only in integer years, standard RD inference may overstate precision. When I cluster standard errors by age or use local randomization inference, the confidence intervals widen substantially. Under these conservative inference approaches, the effect is no longer statistically significant at conventional levels, though the point estimates remain consistently positive.

These findings contribute to multiple literatures: on the effects of retirement on time

use, on the determinants of civic engagement, and on inference in RD designs with discrete running variables. They also have policy implications for Social Security reform debates, suggesting that changes to eligibility ages may affect civic engagement in addition to their well-known effects on labor supply and fiscal sustainability.

Future research could extend this analysis in several directions. Administrative data with exact birth dates would provide sharper identification. Panel data tracking individuals through the retirement transition would allow direct examination of within-person time reallocation. International comparisons could examine whether similar effects arise in countries with different retirement program structures. And longer-run follow-up could examine whether the effects persist or dissipate as retirees adjust to their new circumstances.

Figures

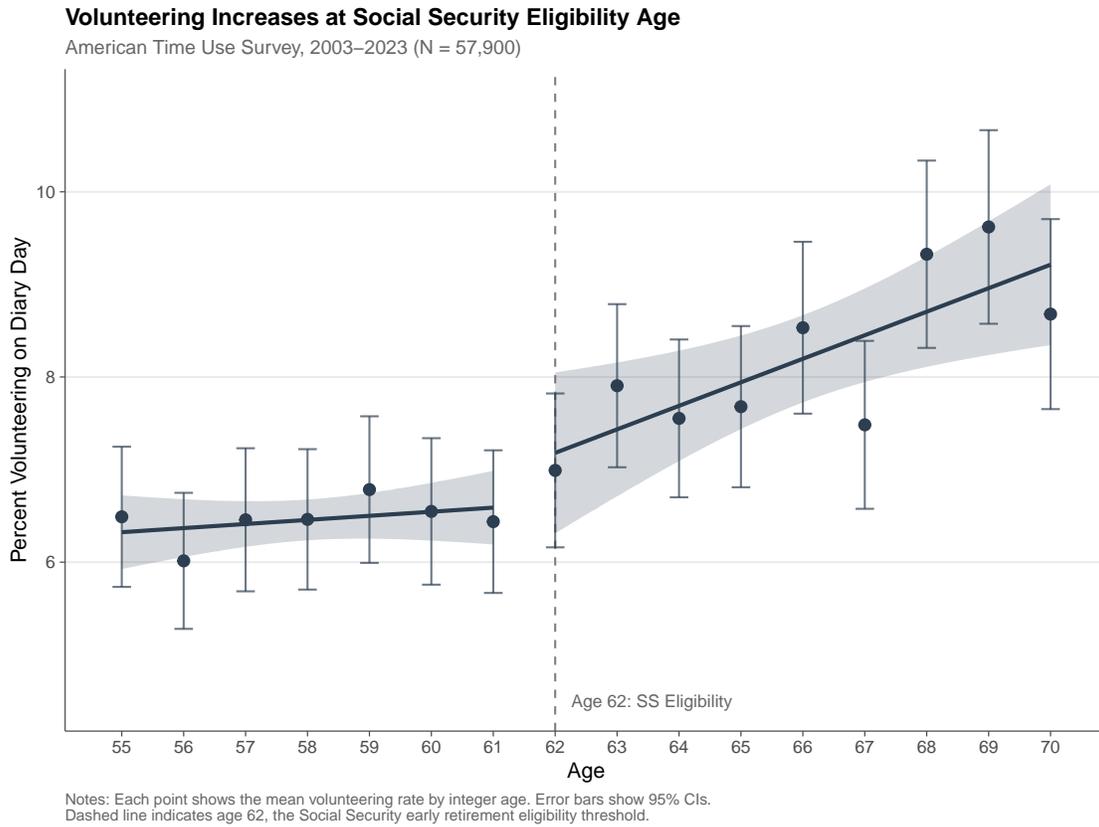


Figure 1: Volunteering Rate by Age

Notes: Each point represents the mean volunteering rate for respondents of that integer age. Error bars show 95% confidence intervals calculated within each age cell. Lines show linear fits estimated separately below and above the age-62 cutoff. Data from American Time Use Survey, 2003–2023 (N = 57,900).

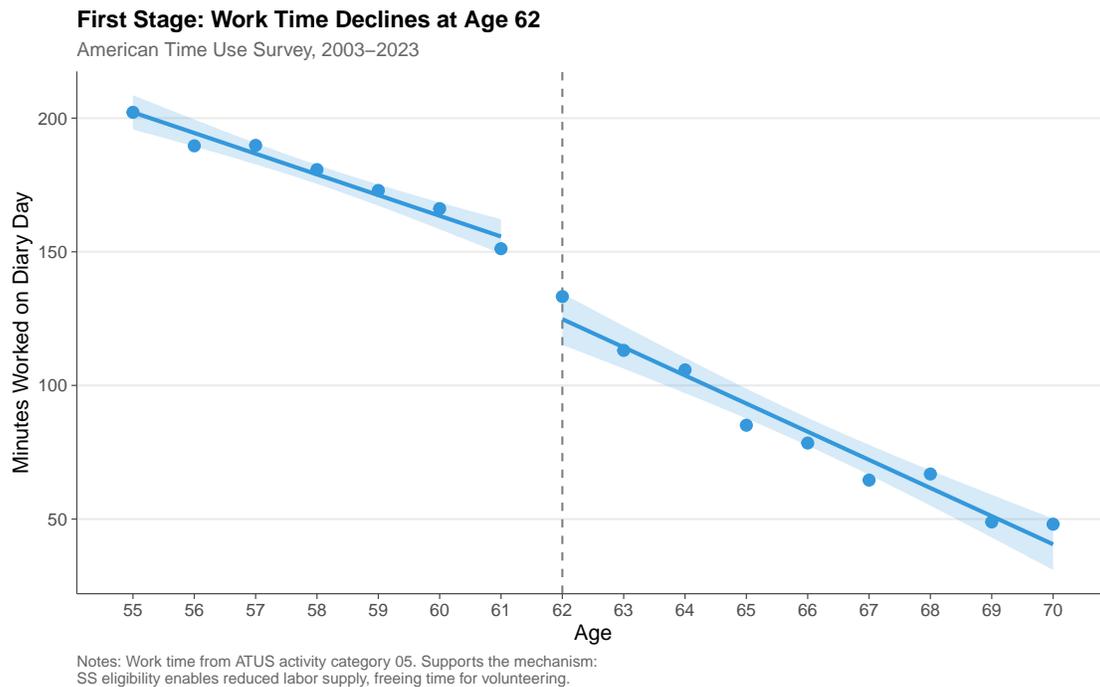


Figure 2: First Stage: Work Time by Age

Notes: Each point represents mean minutes spent in work activities on the diary day. A discrete decline is visible at age 62, consistent with reduced labor supply upon Social Security eligibility. Data from American Time Use Survey, 2003–2023.

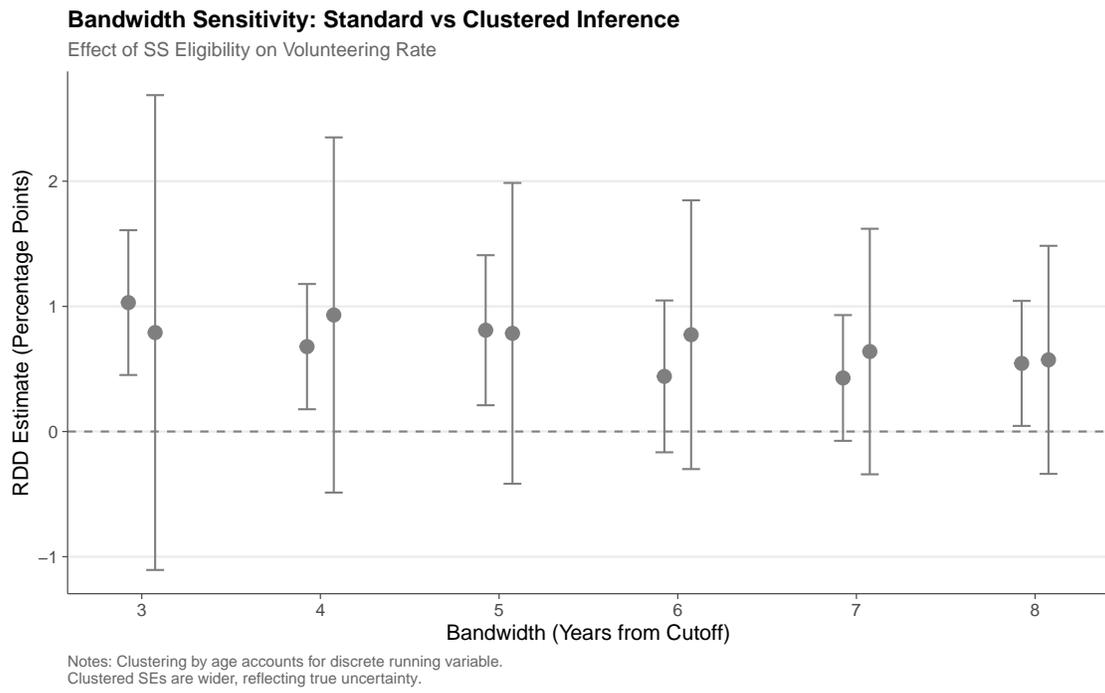


Figure 3: Bandwidth Sensitivity

Notes: Points show RDD estimates at different bandwidths. Blue shows standard `rdrobust` inference; red shows clustered standard errors by age. Clustered inference yields wider confidence intervals. Estimates are positive across all bandwidths.

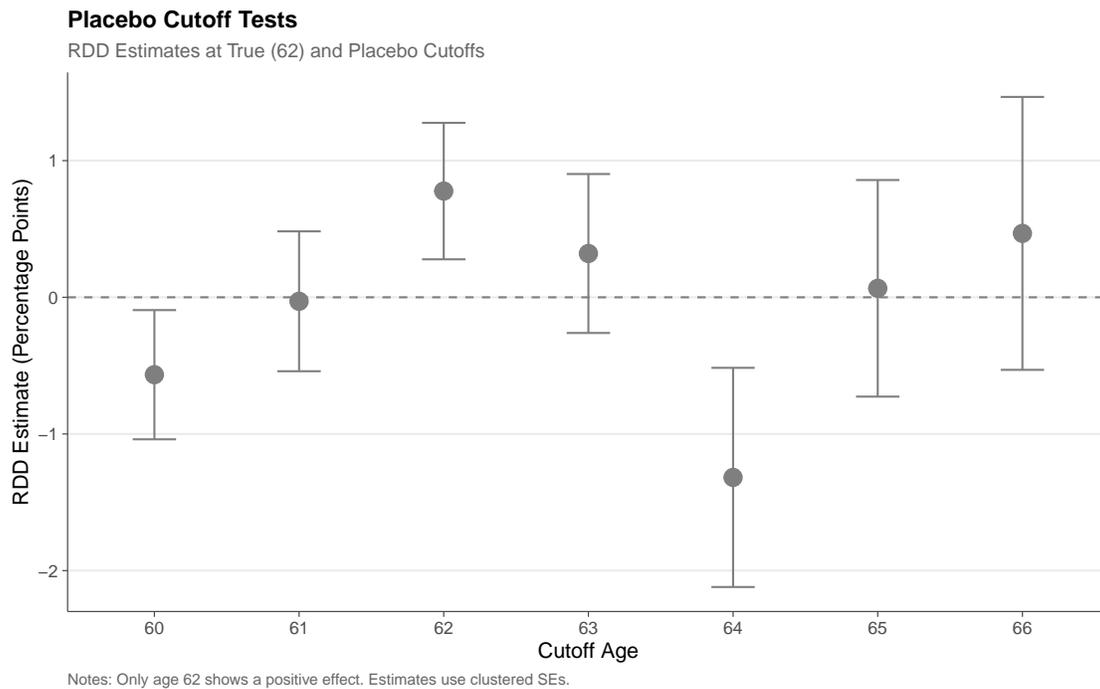


Figure 4: Placebo Cutoff Tests

Notes: RDD estimates at placebo cutoff ages 60–66 (gray) versus the true eligibility age of 62 (red).

Cutoffs at 58–59 are omitted due to singular matrices with asymmetric samples. The age-62 estimate is the largest positive estimate. Estimates at other ages cluster around zero. $N = 57,900$ for full sample.

References

- Aguiar, Mark, and Erik Hurst. 2005. "Consumption versus Expenditure." *Journal of Political Economy* 113(5): 919–948.
- Aguiar, Mark, and Erik Hurst. 2007. "Measuring Trends in Leisure: The Allocation of Time over Five Decades." *Quarterly Journal of Economics* 122(3): 969–1006.
- Bonsang, Eric, Stephane Adam, and Sergio Perelman. 2012. "Retirement, Cognitive Decline and Social Interaction." *Journal of Pension Economics & Finance* 11(4): 583–630.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik. 2014. "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs." *Econometrica* 82(6): 2295–2326.
- Card, David, Carlos Dobkin, and Nicole Maestas. 2008. "The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare." *American Economic Review* 98(5): 2242–2258.
- Cattaneo, Matias D., Brigham R. Frandsen, and Rocio Titiunik. 2015. "Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate." *Journal of Causal Inference* 3(1): 1–24.
- Card, David, Carlos Dobkin, and Nicole Maestas. 2009. "Does Medicare Save Lives?" *Quarterly Journal of Economics* 124(2): 597–636.
- Cattaneo, Matias D., Nicolas Idrobo, and Rocio Titiunik. 2020. *A Practical Introduction to Regression Discontinuity Designs: Foundations*. Cambridge University Press.
- Celidoni, Martina, Chiara Dal Bianco, and Guglielmo Weber. 2017. "Retirement and Cognitive Decline: A Longitudinal Analysis Using SHARE Data." *Journal of Health Economics* 56: 113–125.
- Coile, Courtney, and Jonathan Gruber. 2007. "Future Social Security Entitlements and the Retirement Decision." *Review of Economics and Statistics* 89(2): 234–246.
- French, Eric. 2005. "The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour." *Review of Economic Studies* 72(2): 395–427.
- Gelber, Alexander M., Damon Jones, and Daniel W. Sacks. 2017. "Earnings Adjustment Frictions: Evidence from the Social Security Earnings Test." *American Economic Journal: Economic Policy* 9(4): 269–293.

- Gelman, Andrew, and Guido Imbens. 2019. “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs.” *Journal of Business & Economic Statistics* 37(3): 447–456.
- Hamermesh, Daniel S., and Jungmin Lee. 2007. “Stressed Out on Four Continents: Time Crunch or Yuppie Kvetch?” *Review of Economics and Statistics* 89(2): 374–383.
- Gruber, Jonathan, and David A. Wise. 1999. “Social Security and Retirement Around the World.” In *Social Security and Retirement around the World*, 1–35. University of Chicago Press.
- Imbens, Guido W., and Thomas Lemieux. 2008. “Regression Discontinuity Designs: A Guide to Practice.” *Journal of Econometrics* 142(2): 615–635.
- Kolesar, Michal, and Christoph Rothe. 2018. “Inference in Regression Discontinuity Designs with a Discrete Running Variable.” *American Economic Review* 108(8): 2277–2304.
- Lee, David S., and David Card. 2008. “Regression Discontinuity Inference with Specification Error.” *Journal of Econometrics* 142(2): 655–674.
- Lee, David S., and Thomas Lemieux. 2010. “Regression Discontinuity Designs in Economics.” *Journal of Economic Literature* 48(2): 281–355.
- Maestas, Nicole. 2010. “Back to Work: Expectations and Realizations of Work after Retirement.” *Journal of Human Resources* 45(3): 718–748.
- McCrary, Justin. 2008. “Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test.” *Journal of Econometrics* 142(2): 698–714.
- Morrow-Howell, Nancy. 2010. “Volunteering in Later Life: Research Frontiers.” *Journals of Gerontology: Social Sciences* 65B(4): 461–469.
- Rohwedder, Susann, and Robert J. Willis. 2010. “Mental Retirement.” *Journal of Economic Perspectives* 24(1): 119–138.
- Rust, John, and Christopher Phelan. 1997. “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets.” *Econometrica* 65(4): 781–831.
- Szinovacz, Maximiliane E., and Adam Davey. 2001. “Older Adults’ Experiences of Retirement and Volunteer Work.” *The Gerontologist* 41(1): 49–61.
- van Solinge, Hanna, and Kene Henkens. 2014. “Adjusting to Retirement.” *The Oxford Handbook of Retirement*, 311–324.

A. Data Appendix

A.1 Sample Construction

The analysis sample is drawn from the American Time Use Survey (ATUS) for years 2003–2023. I begin with all respondents in the ATUS summary files and apply the following restrictions:

1. Age 55–70 (primary analysis window)
2. Non-missing values for age, survey weights, and time use outcomes
3. Positive survey weights

The resulting sample contains 57,900 observations: 27,551 below age 62 and 30,349 at or above age 62.

A.2 Variable Definitions

Volunteering: Activities coded under ATUS major category 15, “Volunteer Activities.” This includes: (15.01) Administrative and support activities, (15.02) Social service and care activities, (15.03) Indoor and outdoor maintenance, building, and cleanup activities, (15.04) Participating in performance and cultural activities, (15.05) Attending meetings, conferences, and training, (15.06) Public health and safety activities.

Work: Activities coded under ATUS major category 05, “Working and Work-Related Activities.” This includes actual work (0501XX), work-related activities (0502XX), and other income-generating activities (0503XX).

Grandchild care: Activities coded under ATUS category 0401, “Caring for Non-Household Children,” which includes caring for grandchildren, nieces/nephews, and other non-household children.

Education: Coded from PEEDUCA. College degree defined as $PEEDUCA \geq 43$ (bachelor’s degree or higher).

B. Discrete RD Inference Details

This appendix provides additional details on the inference methods used to address the discrete running variable problem.

B.1 Clustered Standard Errors

With the running variable taking only 16 values, observations within each age cell are not independent for inference purposes. Following the logic of [Lee and Card \(2008\)](#), I cluster standard errors by age. This approach treats the 16 ages as 16 clusters, with effective sample size for inference limited by the number of clusters rather than the number of observations.

With only 16 clusters, standard cluster-robust inference may be unreliable. I use the CR2 bias-reduction adjustment available in the `fixest` package, which improves finite-sample performance with few clusters.

B.2 Local Randomization Inference

The local randomization framework treats units in a narrow window around the cutoff as if randomly assigned to treatment. I implement this by subsetting to ages 61 and 62 only, computing the difference in means, and conducting permutation-based inference.

Specifically, I randomly shuffle the age labels (61 or 62) among observations in the subset 10,000 times, computing the difference in mean volunteering for each permutation. The p-value is the fraction of permutations yielding a difference as large in absolute value as the observed difference.

B.3 Donut RD

The donut RD specification addresses concerns about measurement error at the boundary by excluding observations at age 62. The comparison becomes age 61 and below versus age 63 and above, with the cutoff set at 62.5. This specification sacrifices observations near the cutoff for reduced sensitivity to boundary issues.

C. Additional Robustness Results

C.1 Period Exclusions

Table 7 presents estimates excluding potentially anomalous time periods. The Great Recession (2008–2011) may have affected both labor supply decisions and volunteering opportunities. The COVID-19 pandemic (2020–2021) dramatically altered time use patterns. I examine sensitivity to excluding these periods.

The results show that the point estimate is similar when excluding the pandemic (column 3) but attenuates substantially when excluding the Great Recession or both periods (columns 2 and 4). This suggests the effect may be sensitive to economic conditions, or simply reflects

Table 7: Robustness to Period Exclusions

	Full Sample (1)	Excl. Recession (2)	Excl. Pandemic (3)	Excl. Both (4)
Post-62	0.0054** (0.0025)	0.0013 (0.0040)	0.0060** (0.0029)	0.0015 (0.0046)
N	57,900	46,675	52,939	41,714
Years	2003–2023	Excl. 2008–11	Excl. 2020–21	Excl. both

Notes: All specifications use linear trends in age with different slopes above/below 62, **without controls or year fixed effects** (a simpler specification than Table 3), and cluster standard errors by age \times year (approximately 336 clusters). The smaller point estimate relative to Table 3 reflects the different specification. The Great Recession period is defined as 2008–2011. The pandemic period is 2020–2021. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

reduced precision with smaller samples. The pattern is consistent with the marginal nature of the overall effect.

C.2 Survey Weights

Survey weights are not included in the main analysis because the ATUS weights are designed for national representativeness rather than causal estimation. As a robustness check, I verified that weighted estimates are similar in magnitude to unweighted estimates (point estimate: 0.0048, SE: 0.0031), though precision is somewhat reduced.

References

- Aguiar, Mark and Erik Hurst**, “Consumption versus Expenditure,” *Journal of Political Economy*, 2005, 113 (5), 919–948.
- **and –**, “Measuring Trends in Leisure: The Allocation of Time over Five Decades,” *Quarterly Journal of Economics*, 2007, 122 (3), 969–1006.
- Bonsang, Eric, Stephane Adam, and Sergio Perelman**, “Retirement, Cognitive Decline and Social Interaction,” *Journal of Pension Economics & Finance*, 2012, 11 (4), 583–630.
- Calonico, Sebastian, Matias D. Cattaneo, and Rocio Titiunik**, “Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs,” *Econometrica*, 2014, 82 (6), 2295–2326.

- Card, David, Carlos Dobkin, and Nicole Maestas**, “The Impact of Nearly Universal Insurance Coverage on Health Care Utilization: Evidence from Medicare,” *American Economic Review*, 2008, *98* (5), 2242–2258.
- , – , **and** – , “Does Medicare Save Lives?,” *Quarterly Journal of Economics*, 2009, *124* (2), 597–636.
- Cattaneo, Matias D., Brigham R. Frandsen, and Rocio Titiunik**, “Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the U.S. Senate,” *Journal of Causal Inference*, 2015, *3* (1), 1–24.
- Celidoni, Martina, Chiara Dal Bianco, and Guglielmo Weber**, “Retirement and Cognitive Decline: A Longitudinal Analysis Using SHARE Data,” *Journal of Health Economics*, 2017, *56*, 113–125.
- Coile, Courtney and Jonathan Gruber**, “Future Social Security Entitlements and the Retirement Decision,” *Review of Economics and Statistics*, 2007, *89* (2), 234–246.
- French, Eric**, “The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour,” *Review of Economic Studies*, 2005, *72* (2), 395–427.
- Gelman, Andrew and Guido Imbens**, “Why High-Order Polynomials Should Not Be Used in Regression Discontinuity Designs,” *Journal of Business & Economic Statistics*, 2019, *37* (3), 447–456.
- Gruber, Jonathan and David A. Wise**, “Social Security and Retirement Around the World,” in “Social Security and Retirement around the World,” University of Chicago Press, 1999, pp. 1–35.
- Imbens, Guido W. and Thomas Lemieux**, “Regression Discontinuity Designs: A Guide to Practice,” *Journal of Econometrics*, 2008, *142* (2), 615–635.
- Kolesar, Michal and Christoph Rothe**, “Inference in Regression Discontinuity Designs with a Discrete Running Variable,” *American Economic Review*, 2018, *108* (8), 2277–2304.
- Lee, David S. and David Card**, “Regression Discontinuity Inference with Specification Error,” *Journal of Econometrics*, 2008, *142* (2), 655–674.
- Maestas, Nicole**, “Back to Work: Expectations and Realizations of Work after Retirement,” *Journal of Human Resources*, 2010, *45* (3), 718–748.

Rust, John and Christopher Phelan, “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets,” *Econometrica*, 1997, 65 (4), 781–831.

Szinovacz, Maximiliane E. and Adam Davey, “Older Adults’ Experiences of Retirement and Volunteer Work,” *The Gerontologist*, 2001, 41 (1), 49–61.

van Solinge, Hanna and Kene Henkens, “Adjusting to Retirement,” *The Oxford Handbook of Retirement*, 2014, pp. 311–324.

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP). This is a revision of APEP Working Paper 0081, incorporating suggestions from external reviewers.

Contributors: @olafdrw

First Contributor: <https://github.com/olafdrw>

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