Social Security Eligibility and Healthcare Utilization: Evidence from Administrative Data*

Click for Link to Latest Version

October 27, 2021

Abstract

I estimate the impact of Social Security receipt and retirement on healthcare utilization by exploiting the discontinuous increase in claiming and labor market exit at the Early Eligibility Age of 62. Using administrative data on several types of healthcare encounters from New York and California, I find a discontinuous increase in emergency department visits that do not result in hospitalization by 1-2% at this age. I also provide some evidence that this effect is not due to concurrent changes in health insurance status and may instead be attributed to an increase in free time.

1 Introduction

There is a growing body of evidence suggesting that retirement affects health outcomes in both the short- and long-run (Snyder and Evans, 2006; Fitzpatrick and Moore, 2018; Chuard-Keller, 2021; see Fitzpatrick, 2020 for a review). Of these outcomes, healthcare utilization is uniquely important to policy makers in the United States due to its outsized role in domestic spending. However, it is thus far unclear in the literature whether retirement has a meaningful causal effect on healthcare use in the U.S. This is due to in-part to the fact that examinations of this question set in other countries have obtained conflicting results, making it difficult to extrapolate their findings

*I am grateful to my dissertation commitee, Kitt Carpenter, Carolyn Heinrich, Analisa Packham, Lesley Turner, and Michelle Marcus, for their feedback and support on this project. I also thank the applied microeconomics advising group at Vanderbilt for their help comments on earlier versions of my work. This project was supported in-part by a Dissertation Fellowship from the Center for Retirement Research at Boston College and a dissertation improvement grant from the Russell G. Hamilton Graduate Leadership Institute at Vanderbilt University.

to the U.S. context with confidence[1]. Given that the Social Security Administration has the ability to influence financial incentives to encourage earlier/later retirement, obtaining a better understanding of the relationship between retirement and healthcare utilization could allow policymakers to improve social welfare. My study seeks to provide such an understanding by directly estimating the impact of retirement along with Social Security takeup at age 62 on healthcare outcomes in the U.S.

Estimating the causal effect of retirement on healthcare utilization is an empirical challenge for at least two reasons. The primary reason is that it is difficult in retrospective studies to disentangle the effect of retirement on healthcare use from health-related factors that may have also induced people to leave their job. For example, someone who seeks care for a negative health shock may decide to suddenly retire and begin receiving Social Security benefits rather than continuing to work. The second challege to identifying this relationship is that survey data, which are the traditional source of information for retirement research, suffer from at least two major shortcomings: low sample size and self-reporting errors. The former hinders the analysis of relatively rare but serious outcomes, such as inpatient hospitalizations, and the latter can reduce statistical precision and/or lead to misleading estimates.

My study overcomes these obstacles through the combination of two methodological improvements upon prior work. The first is the use of a regression discontinuity design (RDD) in age to account for selection bias into retirement. Specifically, I take advantage of the fact that a significant share of the population begins receiving Social Security payments upon reaching the Early Eligibility Age (EEA) of 62 (Fitzpatrick and Moore, 2018). My empirical strategy, therefore, is to compare healthcare utilization patterns for people just below age 62 to those just above age 62. If I notice a discontinuous change in utilization rates at this threshold, as compared to surrounding ages, I take this as evidence that reaching the EEA, and its downstream effects, play causal roles in determining healthcare usage. The second improvement is my use of administrative data on the near-universe of emergency department (ED) visits and inpatient hospitalizations in New York and California, along with ambulatory surgery (AS) encounters in California alone between 2006-2017. When considered together, these two states account for approximately 18% of the entire

[1]For example, retirement has been shown to reduce healthcare utilization in Austria (Frimmel and Pruckner, 2020) and Denmark (Nielsen, 2019), increase utilization in China (Zhang et al, 2018; Zhou et al., 2021), and produce a null effect in England (Rose, 2020).

country's population as of 2017. These data allow for sample sizes constituting tens of thousands of each type of healthcare encounter per month of age, permitting the identification of relatively small effect sizes that are still of policy-relevant magnitudes. Furthermore, these data are also useful for reducing measurement error since they are produced using hopital billing information rather than self-report and recall.

I estimate that approximately 17% of people claim Social Security and 8% of people retire immediately upon reaching the EEA of 62. Furthermore, reaching the EEA is associated with a discontinuous increase of 1-2% in all-cause ED visits that do not result in hospitalization. Estimates for elective inpatient admissions, admissions from the ED, and ambulatory surgery encounters are statistically insignificant. I am also able to show that changes in health insurance status at retirement do not play a significant role in the effect on ED visits, despite estimating a decrease in employer-sponsored private insurance at age 62 of approximately 4 percentage points. Finally, I show that turning 62 decreases the amount of time people spend working by around 2 hours per week on average, which suggests that shifting time use at retirement may be an important mechanism.

My study provides credible estimates for the impact of Social Security eligibility and retirement on health care use in the U.S. This question has also been examine previously by Gorry et al. (2018) using the Health and Retirement Study for healthcare outcomes. They find that hospitalizations decrease in the years following retirement using an instrumental variable based on year-of-age relative to several Social Security policy thresholds. My study addresses some limitations in the prior work on this subject by bringing to bear information on additional common and expensive forms of healthcare use, such as ED visits without admission and ambulatory surgery encounters, that are not recorded independently in many survey data sets. Furthermore, one may be concerned that research designs exploiting variation in retirement based on year of age may risk comparing individuals who vary across additional dimensions other than retirement and Social Security claiming. The granularity in age and large sample size afforded by the administrative healthcare data allows for comparisons of indivduals within a smaller window around a relevant age threshold using an RDD, circumventing the many of the issues present alternative designs. Lastly, I am able to break down healthcare encounters by diagnosis to further study the mechanisms that connect retirement to healthcare use.

My study's methodological improvements also contribute to the broader quasi-experimental literature on retirement and health. The earliest of these studies in the U.S. also rely mostly on survey data, along with a variety of IV approaches, to study this topic (Charles, 2004; Dave et al., 2008; Neumann, 2008; Coe and Lindeboom, 2008; Coe et al., 2012; Insler, 2014). In contrast, Fitzpatrick and Moore (2018) use U.S. administrative data and an RDD to show that mortality increases by 2% at the EEA. Fitzpatrick (2020) deems this integration of U.S. administrative data into the literature on retirement on health as "the next generation of US studies" due to the ability to answer new questions with rigorous methods that require these data, such as RDDs. My study seeks to be the next step in this literature by bringing to bear administrative data on healthcare utilization.

Taken together, this paper's findings suggest that, at least in the short-run, concurrent retirement and Social Security uptake increase ED visits without admissions because people have more time to spend pursuing healthcare once they stop working. ED visits are unique among healthcare encounters in that they are a mix between urgent and non-urgent cases and visits that do not results in admission are disproportionately comprised of the latter type (Uscher-Pines et al., 2013). This implies that retirement aids people in consuming care that they may have otherwise forgone or deferred, but may still be beneficial for present and/or future health status. This implies that any change in Social Security policy intended to increase the average age of retirement, such as an increase in the Early Eligibility Age, must be weighed against the costs of forgoing desired healthcare for people who would have otherwise retired. Furthermore, these results have implications for evaluating how people make decisions about consuming healthcare versus the opportunity cost in terms for forgone time spent accruing earnings.

2 Background: Social Security Benefits and the Early Eligibility Age

2.1 Eligibility Requirements

The Social Security program is the primary public retirement benefit for aged individuals in the United States, with approximately 50% of the population aged 65+ relying on Social Security for at least half of their household income as of 2015 (Dushi et al., 2017). Furthermore, about 25% of this population receive a full 90% or more of their income through these monthly payments. Most people first become eligible to receive Social Security retirement benefits after working for at least 10 years in qualified employment and reaching the Early Eligibility Age of 62 (SSA, 2021). The EEA was established in 1959 for women and 1961 for men as an alternative to claiming at the age of 65 (McSteen, 1985). However, this policy also made it so that claiming at 62 reduced an individual's monthly payments by 5/9 of 1% for each month between the claiming month and when the individual turns 65, which was rebranded as the Full Retirement Age (FRA) (SSA, 2021). The FRA was eventually increased by two months for each yearly birth cohort after 1937 until the 1943-54 cohorts, where it remained at 66. Starting with the 1955, cohort the FRA continued its increase by two months by year of birth until it stopped at 67 for cohorts born in 1960 or later (Li, 2021). This implies that someone reaching the EEA and claiming Social Security benefits in 2006 would receive approximately 73.33% of the monthly amount they would have earned if they had waited for their FRA. This percent decreased to 72.77% for someone claiming at the EEA in 2017.

People are allowed to apply for Social Security shortly before they turn 62, which means that they can begin to receive benefits within a couple of days of their birthday (SSA, 2021b). Applications must be made at least four months in advance, which means claiming decisions cannot be made entirely spotaneously. People who claim before their FRA and continue work are subject to the Retirement Earnings Test, whereby one dollar of benefits is withheld for every two dollars earned from work above a certain maximum (SSA, 2021c). Social Security is also available to the dependents and survivors of those who qualified based on work history, as decribed above. While spouses and ex-spouses are not beholden to the same work history requirements as the primary claimant, they must still reach age 62 before they are allowed to claim with few exceptions (SSA, 2021b). A subset of other dependents may claim under the primary claimant as well, including children under 18 (or 19 if still in high school) or children of all ages disabled before 22 (SSA, 2021b). Despite these exceptions, only about 15% of individuals claim Social Security benefits of any kind (Disability, Retirement, and/or Survivors Insurance) before 62.

2.2 What Kind of Person Claims and Retires at 62?

Approximately 17% of people first claim Social Security upon turning 62 and, on average, these individuals are more likely to have only a high school education, be of self-reported poor health, to be black, and to have lower personal and household income than those who fully retire at later ages (Rutledge and Wettstein, 2020). Additionally, those claiming at 62 are more likely to have developed work-limiting health conditions by the time they claimed benefits and are more likely to have worked in physically laborious and/or blue-collar jobs (Li et al., 2008; Munnell et al., 2016). It is also the case that individuals who claim at age 62 are more likely to have health insurance policies that are not dependent on employment status (e.g., uninsured, private insurance with retiree coverage, and/or Medicaid) than those who retire at other ages (Rust and Phelan, 1997). However, these are just differences in averages. In fact, those who retire at 62 can be further seperated into two groups of roughly equal size: the advantaged and the disadvantaged (Munnell, et al, 2016). In the advantaged group, approximately 87% have at least some college education, only 10% are blue collar workers, and 59% are in the top quartile of wealth. On the other hand, only 17% of the disadvantaged group have attended at least some college, 77% are blue collar workers, and only 8.3% reside in the top quartile of wealth. This suggests that many people who claim at 62 are not forced into by life circumstances, but because of a desire to retire earlier.

Particularly important for my study are the people who claim within only a couple of months of turning 62 and their retirement decisions. In fact, as of 2006, the majority of those who claim within two months of turning 62 have actually previously retired: approximately 60% of male claimants and 69% of female claimants (Waldron, 2020). Furthermore, approximately 28% of male claimants and 21% of female calimants who claim benefits upon turning 62 also retire concurrently and the remaining 12% of men and 10% of women claimants continue to work (Waldron, 2020). Selection into these subcategories of age 62 claimants is highly correlated with lifetime earnings. Specifically, people in the lower deciles of lifetime earnings are more likely to to have retired prior to claiming at 62, while people in the middle of the distribution are most likely to retire when claiming (Waldron, 2020b).

3 Data and Empirical Strategy

3.1 Administrative Data on Healthcare Utilization

This study's measures of healthcare utilization are based on various sources of administrative data on healthcare encounters from California and New York. I obtained data on the universe of inpatient hospitalizations, emergency department (ED), and ambulatory surgery (AS) visits in California from the state's Office of Statewide Health Planning and Development (OSHPD) for years 2006-2017. Additionally, I obtained data on the near-universe of inpatient and ED encounters in New York state from the Healthcare Cost and Utilization Project (HCUP)'s State Inpatient Database (SID) and State Emergency Department Database (SEDD) between 2006-2017[2]. All of these data are based on the billing information associated with each visit and include a substantial amount of detail about the patient encounter including patient demographics, diagnoses, and patient ZIP codes. Importantly, all of these data contain both the date of encounter and the birthdate of each patient at the month-year level. Combining these pieces of information allows me to calculate a person's age at encounter. However, since I do not have people's exact day of birth, I am only able to determine people's age by calendar month. Therefore, in the month that people turn 62 I cannot distinguish those who have passed their birthday from those that have not yet. I discuss my method for accounting for this measurement error in my empirical strategy (Section 3.4).

3.2 Survey Data for Evaluating Mechanisms

Becoming eligible for Social Security at the EEA has the potential to trigger a wide variety of behavioral responses and changes in life circumstnaces, many of which could foreseeably affect healthcare utilization. The principal source of survey data that I use is the Health and Retirement Study (HRS), which is a longitudinal survey of a representative sample of the U.S. population aged 50+. Since this survey is only conducted in even years, I use the 2006-2018 waves of the HRS to approximate the 2006-2017 time period of my healthcare utilization data. It is from these data that I draw my primary outcome measures for Social Security uptake, retirement, labor force

[2]These data exclude stays in long-term care units of short-term hospitals, Federal hospitals, and free-standing psychiatric hospitals.

participation, health insurance coverage, and heath behaviors[3]. Importantly, the HRS also contains information on month/year of birth and month/year of interview during the relevant sample period. This information allows me to calculate age (in months) at interview in a similar fashion as I do for the administrative data on healthcare utilization. All estimates using the HRS are made at the national level unless otherwise stated.

3.3 Other Data Sets

This study uses a few additional data sets for supplemental analysis. I use the 2007-2018 Current Population Survey - Annual Social and Economic Supplement (CPS-ASEC) for data on Social Security and state pension receipt for the previous year, as well as for data on the national distribution of household income (Flood et al., 2020). I collect data on the median household income by ZIP code from the ACS 5-year estimates, 2010-2014 (Manson et al, 2021). I also make use of natality data collected from historical vital statistics records in combination with population data from the 2010 census 10% sample (Ruggles et al., 2020) to construct population counts by age cohort[4]. Lastly, I make use of the National Health Interview Survey (NHIS), a representative annual survey on health and economic outcomes in the U.S., as a supplement to the HRS. I use the publicly available version from IPUMS, which contains information on date of interview and birth at the month level between 2006-2014 (Blewett et al., 2019).

3.4 Empirical Strategy

I estimate the causal effect of reaching the EEA of 62 on healthcare utilization through the use of an age-based regression discontinuity design (RDD). I utilize the “continuity framework” for identifying the unbiased effect of gaining eligibility at the EEA via RDD (Hahn, Todd, and van de Klaauw, 2001; Cattaneo and Titiunik, 2021). The identifying assumption of this approach is that the potential healthcare outcomes are continuous functions of age at the age 62 cutoff. Another way of viewing this is that, in the absence of treatment, outcomes would have evolved continuously with respect to age rather than jump discontinuously. This is a good approach for establishing the

[3]I discuss the construction of these variables in greater detail in Appendix A

[4]See Appendix A for further discussion and construction of the population denominator.

causal effect of Social Security eligibility as long as people on either side of the age threshold are very similar in almost every way, on average, other than their eligibility status for Social Security.

One potential issue in this approach is if other, non-federal retirement plans also use 62 as a policy-relevant age. While "defined contribution plans" with no age-based eligibility have become the standard for private-sector employees (Dushi et al. 2011), many state and local pensions frequently use age to determine eligibility. Indeed, New York and California both have pension programs for state and local government employees that use age 62 as a relevant threshold for payments[5]. As long as there is not a spike in the claiming of state and local government pensions in these two states at 62 that is comparable to the spike in Social Security claiming, these policies should not significantly influence my results for healthcare utilization. Appendix Figure B.1 compares fractions receiving Social Security and state/local pension income in New York and California by year and by age from 2006-2017 according to the CPS-ASEC. Although Social Security receipt increases dramatically for people aged 62, state and local pension claiming does not increase meaningfully.

In equation (1), I display the primary RDD specification used with the healthcare utilization data. The data are organized into month of age (a), state (s), and encounter year (y) cells. y_{asy} is the outcome of interest, such as the rate of healthcare encounters per 1,000 population [6]. Visits are aggregated in age bins that correspond to the calendar month when people turn a certain age (e.g., the calendar month people turn 62). EEA_a is a dummy that equals 1 for encounters occurring in the month individuals turn 62 or later. $f(admitAge_a, EEA_a)$ is a polynomial function of the running variable that is allowed to vary on either side of the cutoff d_a is a dummy that equals 1 in the month that people turn 62 and 0 otherwise to account for attenuation bias (Dong, 2015). Lastly, $\delta_s \times \sigma_y$ are state-by-year fixed effects to increase the model's explanatory power. All regressions use triangular kernels and standard errors are heteroskedasticity-robust. In specifications estimated using the Calonico, Cattaneo, and Titiunik (CCT) Mean-Squared Error (MSE) optimal bandwidt, I also display p-values derived from their robust bias-correct confidence intervals (Calonico, Cattaneo,

[5]New York's NYSLRS uses 62 as its internal "full retirement age" and California's CALPERS uses 62 as its maximum "normal retirement age".

[6]Specifying the outcome in terms of rates instead of counts controls for sudden changes in population size across age bins. More details on the construction of the denominator can be found in Appendix A.

and Titiunik, 2014).

$$y_{asy} = \alpha + \beta EEA_a + f(admitAge_a, EEA_a) + d_a + \delta_s \times \sigma_y + \varepsilon_{asy} \tag{1}$$

In models using outcomes drived from survey data, I aggregate data only by month of age in order preserve the role of sample weights and estimate equation 2.

$$y_a = \alpha + \beta EEA_a + f(admitAge_a, EEA_a) + d_a + \varepsilon_a \tag{2}$$

4 Conceptual Framework

Initiation of Social Security payments and retirement constitute major life changes that have the potential to impact healthcare utilization in a myriad of ways. Figure 1 is a diagram suggesting several potential channels through which this could occur. First, receiving Social Security benefits, all else equal, constitutes an increase in monthly income of about $1,438 on average as of 2019 (Van de Water and Romig, 2020). Increased lifetime Social Security income has been shown to impact health and healthcare utilization both beneficially (Berman, 2020) and detrimentally (Snyder and Evans, 2006), depending on whether the increased income is accompanied with earlier retirement. Additionally, short-run liquidity shocks induced by the receipt of monthly Social Security checks reduce liquidity and allow people to purchase medical care (Gross et al, forthcoming).

Second, many people decide to retire once they claim Social Security (Fitzpatrick and Moore, 2018). The impact of retirement on healthcare utilization is unclear *a priori* given the large number of potential channels involved. While receiving Social Security increases income, losing wages from exiting the labor force may provide a counterveiling effect. Furthermore, some employers do not provide employer-sponsored health insurance to employees after retirement which may reduce rates of health insurance coverage prior to Medicare eligibility (McArdle et al., 2014). However, one of the biggest changes accompanying retirement is an increase in newly freed time that was formerly spent on work. This time could be spent on a wide variety of behaviors that could either positively or negatively affect healthcare utilization in the short run. For example, spending this time on sedentary activities could decrease health (Koster et al., 2012; Matthew et al., 2012), thereby increasing healthcare utilization, while time spent exercising or eating well could do the opposite. Furthermore, retirement is known to be a time during which people increase their consumption of unhealthy substance such as tobacco or alcohol (Ayyagari, 2016; Chuard-Keller, 2021), again decreasing health. Finally, retirees may simply choose to spend their additional time consuming additional healthcare that is not prompted by a change in underlying health status. In my discussion of mechanisms, I show that this final channel is most likely responsible for the increase in ED visits.

5 Main Results

5.1 First Stages: Effects on Social Security Uptake and Retirement

As discussed in the previous section, there are two primary channels through which eligibility for Social Security at the EEA are likely to affect healthcare utilization outcomes: retirement and benefit receipt. Many of the other potential mechanisms, such as health insurance and time use, would likely flow from at least one of these other two channels. Therefore, I need to show that people do indeed claim Social Security, retire, and exit the labor force discontinuously at 62 between 2006-2018. Figure 2 shows the RD plot for the fraction of people current receiving Social Security using the HRS and displays a positive discontinuity in Social Security receipt at 62 of about 16.6 percentage points[7]. Figure 3 shows how self-reported retirement, labor force participation, and working for pay statuses change for the entire sample at this age. All three outcomes display discontinuous changes at 62 of large magnitudes, though smaller in magnitude than the discontinuity in Social Security receipt[8]. The effects on Social Security uptake and labor market outcomes have been fairly stable over time, as demonstrated in Appendix Figure B.3. Despite some apparent cyclicality, the effects on Social Security claiming and retirement at 62 have hovered around 20% and 10%, respectively, over the two decades prior to 2018.

[7] The outcome is not limited to individuals receiving Social Security Retirement Insurance, and includes individuals receiving Disability and/or Survivors Insurance.

[8] Figure 3 appears to suggest that these effects on labor market outcomes might not show up immediately and could possibly take up to two months to do so. In order to verify if this is truly the case, or just due to sampling variability, I compare the labor force results in the HRS to those in the NHIS between 2006-2014 in Appendix Figure B.2. While both RD plots demonstrate discontinuities at 62, HRS results display a delayed effect, which suggests that it is likely due to sampling error.

Previous research on Social Security eligibility at the EEA has shown that the discontinuities in claiming and labor force outcomes vary considerably by demographic subgroup. In Table 1, I show results from estimating equation (2) with these outcomes across sexes and racial/ethnic categories. Column (1) demonstrates that each demographic subgroup experiences an increase in Social Security receipt at 62 that is significant at at least the 10% level. On the other hand, columns (2-4) show that the effects on labor market outcomes are not as widespread across demographic subgroups. First, unlike in Fitzpatrick and Moore (2018), I find significant increases in retirement, labor force non-participation, and decreases in working for pay among females as well as males. This is likely because labor force attachment is higher in my sample period than their earlier one[9]. Second, I find that changes in labor market outcomes are concentrated among non-hispanic white people and, less robustly, hispanic people.

5.2 Effects on Healthcare Utilization

ED Visits

Figure 4 plots rates of aggregate ED visits per 1,000 population by age in months with linear and quadratic fits by estimating equation (1) with 8-month bandwidths. Both panels (a) and (b) show clear jumps in visits between the month before people turn 62 and the month after people turn 62 that are robust to choice of polynomial. I evaluate these visual assessments by estimating the RD coefficients for each discontinuity under varying polynomial and bandwidth assumptions using equation (1) and display the results in Table 2. The top panel, which displays results for the aggregate sample of both states combined, is consistent with the respective plot in Figure 4. Column (1) indicates that becoming eligible for Social Security at the EEA of 62 is associated with an increase of 0.204 vists per 1,000 people, which is equivalent to an increase of about 1.1% compared to the month before turning 62. This effect is statistically significant at less than the 1% level and is robust to the use of a quadratic polynomial, as shown in column (2). Columns (3) and (4) re-estimate the models using the mean-squarted error (MSE) optimal bandwidth from Calonico, Cattaneo and Titiunik (2014) with 24 months of data on either side of the cutoff. These point

[9] Appendix Figure B.4 demonstrates that this finding is not particular to the HRS by estimating the discontinuity in labor force non-participation by sex using the NHIS. In this data set, the coefficient estimate is actually larger (though not significantly so) among females than males.

estimates are statistically significant at the 1% level, regardless of polynomial choice and inference method. The bottom two panels re-estimates these specifications but for each state seperately. They indicate that the effect is present in both New York and California independently.

Column (1) of Table 3 displays how the effect on ED visits varies across demographic subgroups. Outcomes are specified as the natural log of visits by sex and race/ethnicity since the population denominator cannot be cut by all of these categories. Similar to the effects on Social Security claiming and labor market outcomes, there are significant discontinuities for both males and females at age 62. On the other hand, the positive effect on ED visits appears to be mostly concentrated among non-hispanic white individuals, which is consistent with the effects on labor market outcomes as well. I also break down the aggregate effect by primary diagnosis, which is determined by the physician to be the principal reason why the patient presented to the ED that day. I show these results in Figure 5. There are generally positive or near-zero coefficient estimates across most diagnosis categories, with no category demonstrating an increase that is disproportionate with its share of total visits. Figure 5 also displays a significant negative effect on visits for endocrine and metabolic disorders. Further analysis indicates that this negative effect is driven by decreases in visits for diabetes (β=-0.023, $p<0.01$), disorders of non-thyroid glands (β=-0.025, $p<0.001$), and various metabolic disorders (β=-0.011, $p<0.05$) and is driven by black and hispanic individuals. On the contrary, the effects on the remaining categories of endocrine/metabolic disorder (thyroid problems, nutritional deficiencies) are positive and significant at the 5% level.

Lastly, I examine how the effect on ED visits has changed over the time period of study in Figure 6. The coefficient estimates are mostly positive throughout the entire time period with no visible increases or decreases over time. This is also consistent with the results for Social Security receipt and labor market outcomes which are mostly positive across the relevant time period.

Inpatient Hospitalizations

I turn my attention next to the effect of reaching the EEA on inpatient hospitalizations. I focus on two types of admissions: those originating in the ED and those for elective procedures[10]. Figure 7 shows how these two types of admissions change through the age 62 cutoff. Neither panel (a)

[10] As of 2017, hospitalizations from the ED account for approximately 65% of admissions in New York and California for people aged 62 and 11 months. Elective hospitalizations account for approximately 25% of admissions in the same sample.

nor panel (b) display large discontinuous changes in admissions at 62. Table 4 confirms these null effects across polynomials and bandwidths. However, it is possible that these aggregate results belie significant heterogeneity across demographic groups. Columns (2) and (3) of Table 3 display RD estimates for each type of inpatient admission by subgroup. While coefficients are statistically insignificant in most categories, I estimate decreases in admissions from the ED for black people and people of non-white/hispanic/black races and ethnicities that are significant at the 1% and 10% levels, respectively.

Ambulatory Surgery Encounters

The last form of health care utilization I analyze is ambulatory surgery encounters. Figure 8 shows the RD plot for AS encounters in California and does not suggest an exceptionally large discontinuity at 62. Regression estimates in Table 5 confirm this assessment, displaying insignificant coefficients across all specifications. Furthermore, I do not estimate significant effects within any of the demographic subgroups as displayed in column (4) of Table 3.

5.3 Robustness Checks

I now conduct three evaluations of the robustness of the estimated effect on ED visits. First, I systematically test the robustness of the ED estimates from the top panel of Table 2 to various alternate bandwidths by choice of polynomial and plot the results in Appendix Figure B.5. The coefficient estimates are significant at the 5% level for bandwidths of 4 through 12 months when using a linear fit. Since lower bandwidths tend to suffer less from a variety of undesirable characteristics (e.g., erratic behavior near boundary points, counterintuitive weighting, and overfitting), these findings are reassuring (Calonico and Titiunik, 2021). When using a quadratic fit, the coefficients estimates are significant at the 5% level for every bandwidth except for 4 months where the confidence intervals are relatively large due to the lack of observations.

Next, I present the results from conducting randomization inference with the effect on ED visits as both an alternative inference method and placebo test. I do this in two ways. First, I re-estimate equation (1) with the natural log of aggregate ED visits as the outcome at nearly every month within a 10-year radius of the age 62 cutoff[11]. In total this results in 204 placebo cutoffs. Then, I plot the distribution of these RD estimates in panel (a) of Appendix Figure B.6. The dashed lines indicate the boundaries containing 95% of estimated coefficients and the red line indicates the treatment effect estimate. The second way I conduct randomization inference is by using the t-statistics generated via the CCT robust inference methods instead of just the coefficient estimates. I plot the distribution of t-statistics in panel (b) of Appendix Figure B.6. In both cases, the estimate at age 62 is larger than 95% of estimates. This suggests that the estimated effect of interest is exceptionally large when compared to placebo estimates and is therefore likely to be a true policy effect.

Last, I conduct RD "donut" specifications whereby I exclude months close to the cutoff and re-estimate discontinuity in ED visits per 1,000 population with dummy avariables for those months. I conduct this exercise while bearing in mind that, according to the continuity framework for RDDs, identification of the treatment effect is only identified exactly at the cutoff (Cattaneo and Titiunik, 2021). Since my primary specification already removes the variation from the month in which people turn 62, it is already pushing the assumptions built into this identification strategy. Further distance from the cutoff only reduces the plausibility of identification. Acknowledging this, Table B.1 shows the results of estimating equation (1) with dummies for both one and two additional months from the cutoff. Both the linear and quadratic estimates are robust to excluding an additional month from either side of the cutoff and produce point estimates that are even larger than the base specifications. However, excluding an additional two months renders the results statistically insignificant and negatively signed.

6 Understanding the Healthcare Effects at 62

In Section 5, I articulate the various channels through which Social Security eligibility might impact healthcare utilization. As shown in Figure 1, uptake of Social Security payments and leaving the labor force are the two primary behaviors that are likely to change at age 62. However,

[11] I use the natural log of ED visits instead of the rate per 1,000 population because the vital statistics data do not go back far enough to estimates effects for older individuals. Additionally, for the months to truly be placebo cutoffs, they must not overlap with the effect of the EEA or any other relevant age-based policy. That is to say, I want to avoid "treatment effect contamination" (Calonico and Titiunik, 2021). For this reason, I exclude months within an 8-month radius of age 62 and age 65, as well as the month people turn 65 (and become eligible for Medicare).

each of these behaviors has the potential to cause a variety of cascading effects that may ultimately affect utilization patterns. In this section, I first provide further evidence that Social Security receipt and retirement are driving the effect on ED visits. Then, I show that subsequent changes in health insurance coverage and household income at 62 are not likely to be the main mechanisms at play. Last, I discuss why the effect on ED visits is probably driven by a change in time use.

6.1 Correlation Between Social Security Uptake, Labor Market Outcomes, and the ED Visit Discontinuity

In the previous sections, I have shown that Social Security receipt, retirement, and ED visits all increase discontinuously at age 62. I have also shown how these estimates vary across across demographic subgroups. In Figure 9, I plot the relationship between these effects by each of these subgroups as well as each individual's place in the income distribution[12]. These graphs show that the subgroups with the biggest effects on the first stage outcomes also have the largest effects on ED visits. This further confirms that Social Security receipt, and subsequent retirement, are the primary drivers of the effect on ED visits.

6.2 Health Insurance

Many people lose access to employer-sponsored health insurance when they retire. Furthermore, accessing Social Security income provides additional liquidity that may be used to purchase private health insurance, all else being equal. These changes in health insurance status could, in turn, potentially affect healthcare utilization patterns. I assess whether health insurance status changes discontinuously at age 62 using the HRS in Figure 10. Panel (a) shows that the overall share of respondents reporting having at least one form of health insurance does not change discontinuously at 62, despite the 4.2 percentage point decrease in employer-sponsored private insurance displayed in panel (b)[13]. This decrease is compensated for by a concurrent increase in the share of people reporting other categories of insurance. Appendix Figure B.7 shows how these other categories change independently at 62. Additionally, Appendix Figure B.8 displays statistically

[12] I discuss my method for categorizing observations into income deciles in Appendix A.
[13] This excludes specialty plans such as dental, vision, etc.

insignificant effects on the likelihood of having two or more plans and the total number of plans held.

The statistically insignificant effect on the likelihood of having health insurance suggests that insurance coverage is not driving the effect on ED visits. In order to provide further evidence that loss of employer-sponsored private insurance cannot explain the effect, Appendix Figure B.9 displays RD plots for both the share and rate of privately insured ED visitors per 1,000 population. While the relative share of patients with private insurance stays constant, I estimate a significant increase in the total number of patients with private insurance at age 62. This implies that the ED results cannot be entirely driven by reductions in private insurance coverage since I estimate an significant increase in ED visits for this population at age 62.

6.3 Income

It is possible that substantial chages in household income upon individuals reaching 62 could affect their healthcare use. I evaluate this by plotting household income by year of age from the CPS-ASEC in Figure 11[14]. Given the coarse granularity of this variable it is difficult to draw firm conclusions, but the the income level at age 62 does not appear to lie significantly outside of the preceding trend.

6.4 Time Use

After eliminating health insurance status and income as the main channels for the effect on ED visits, based on the conceptual framework layed out in Figure 1, the only remaining option is an increase in free time. Since ED visits for non-urgent conditions make up a substantial portion of those visits which do not result in hospitalization (Uscher-Pines et al., 2013), an increase in free time could plausibly result in an increase in ED visits. Previous studies have shown that people shift their time substantially after retiring, particularly toward leisure (Stancanelli and Van Soest, 2016). Additionally, Lucifora and Vigani, (2018) show using the European SHARE data set that retirement causes larger increases in visits to the doctor among people who had more time constraints prior to retirement. Therefore, it is reasonable to suggest that these time-use dynamics

[14] I use the CPS-ASEC instead of the HRS since both surveys ask about income in the past 12 months, instead of current income, but the CPS-ASEC has a larger sample size.

may be present in the U.S. as well and that retirement could increase people's propensity to use healthcare. Using the HRS, I am able to show that at age 62 people spend discontinuously fewer hours per week working. I display these results in Figure 12. If people are spending fewer hours of the day working, then they are able to repurpose that time into other productive activities such as discretionary health care utilization.

7 Conclusion

This study presents evidence that retirement increases more discretionary forms of healthcare utilization in the United States. I accomplish this by combining an RDD, which is a highly internally valid research design, with administrative data on a wide variety of healthcare usage in two of the most populus states in the country. I also show that the most plausible mechanism driving this effect is the increase in free time at retirement and rule out changes in health insurance status as the predominant channel. Taken as a whole, these results contain important implications about the opportunity cost of time for those approaching retirement. Many people who would otherwise use healthcare are prevented from doing so due to the time constraints imposed on them from working. This implies that policies encouraging earlier retirement may help people access more preventative care that could produce health dividends into the future. Alternatively, other policy solutions may be implemented that assist older workers in finding time to use healthcare while they are still working.

Figures

Figure 1: Conceptual Channels Connecting Social Security Eligibility to Healthcare Utilization

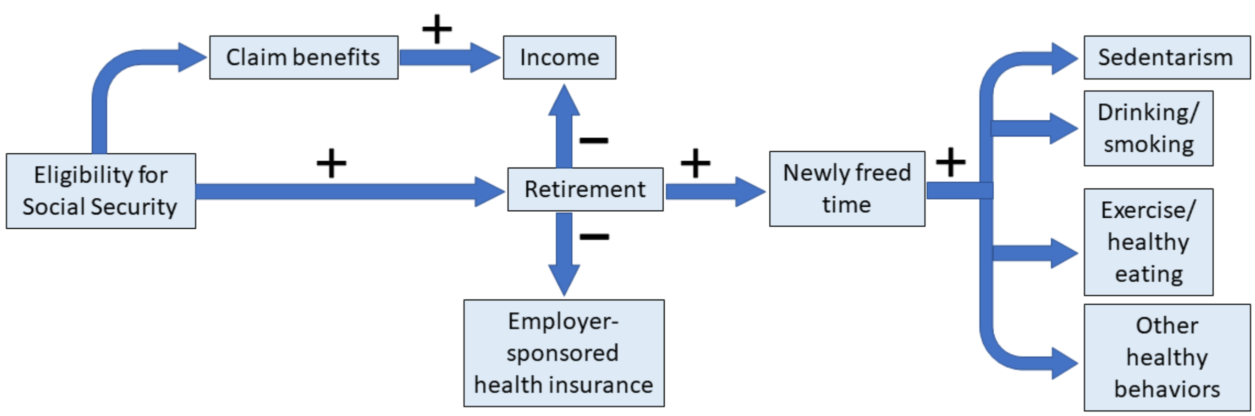

Figure 2: People Join Social Security Discontinuously at 62

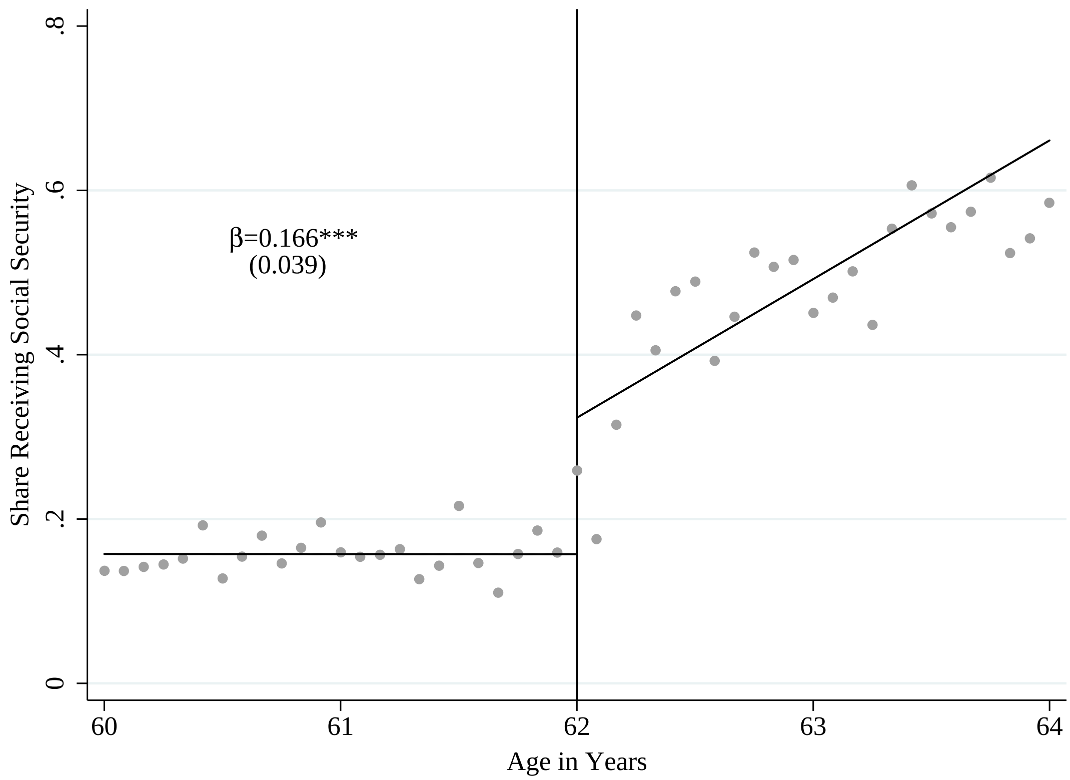

Note: This figure plots the share of the population currently receiving Social Security Retirement, Disability, or Survivor Insurance by age. Linear fit generated by estimating equation (2) with a 24-month bandwidth and triangular kernel. Data are from the Health and Retirement Study waves 2006-2018. Each age bin corresponds to the calendar month that people turned each age (e.g., share receiving income from Social Security in the calendar month people turn 62). Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Figure 3: People Leave Work Discontinuously at 62

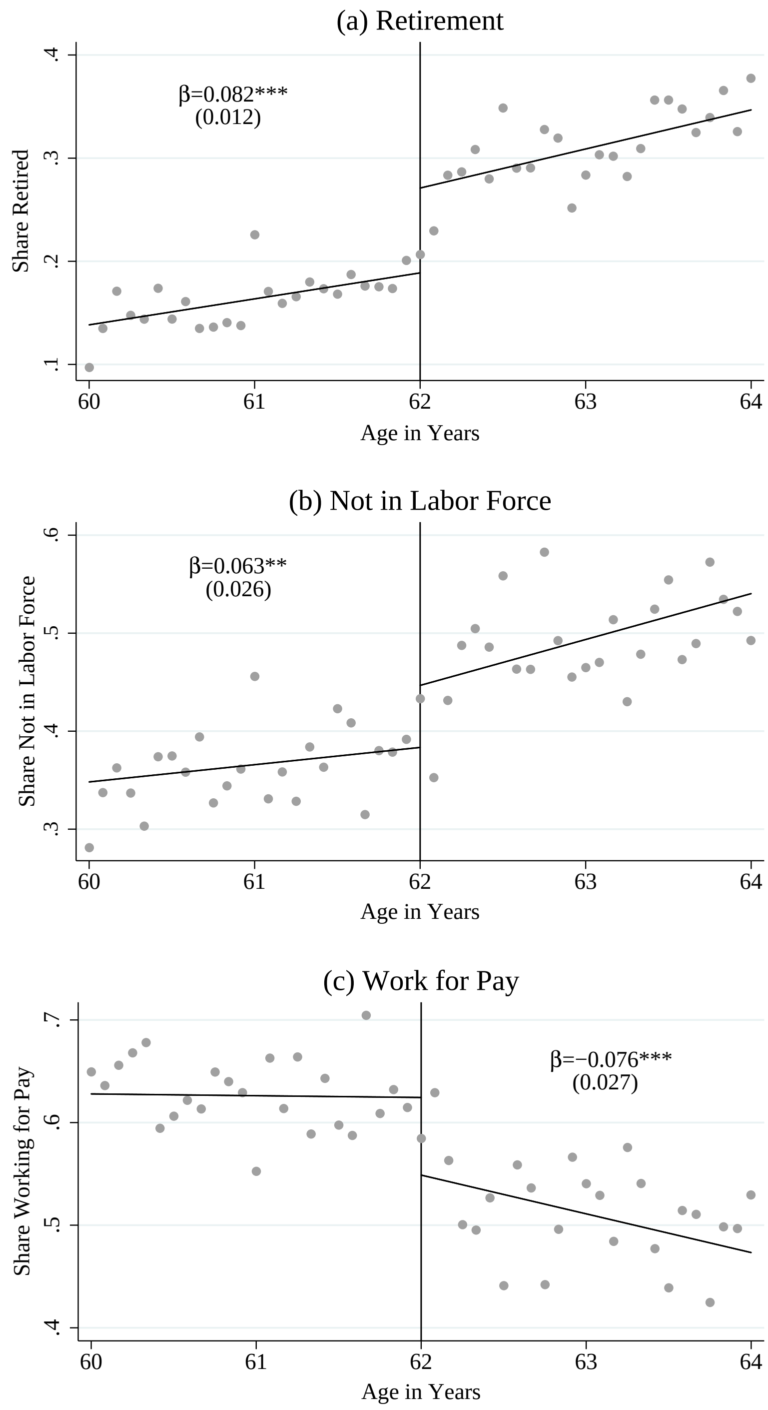

Note: This figure plots the share retired, not in the labor force, and working for pay, by age. Linear fits generated by estimating equation (2) with 24-month bandwidths using triangular kernels. Data are from the Health and Retirement Study waves 2006-2018. Each age bin corresponds to the calendar month that people turned each age (e.g., share retired in the calendar month people turn 62). Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Figure 4: Aggregate ED Visits Increase Visibly at 62

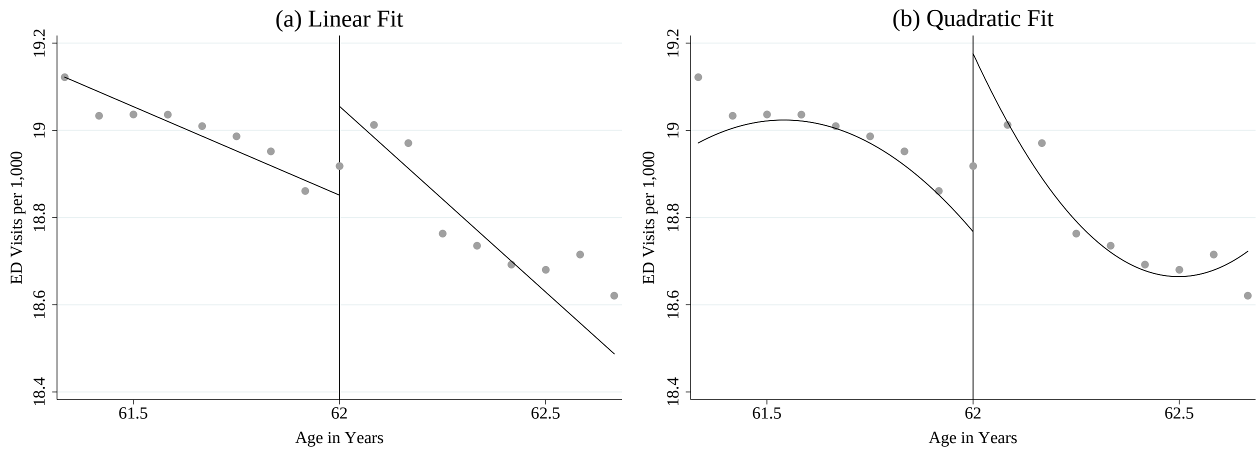

Note: This figure plots ED visits per 1,000 population by age. Data are from HCUP NY SEDD and CA OSHPD ED visit data, years 2006-2017. Each age bin corresponds to the calendar month that people turned each age (e.g., ED visits in the calendar month people turn 62). Polynomial fits generated from estimating equation (1) with 8-month bandwidths and triangular kernels.

Figure 5: Increase in ED Visits Not Driven by Any Particular Diagnoses

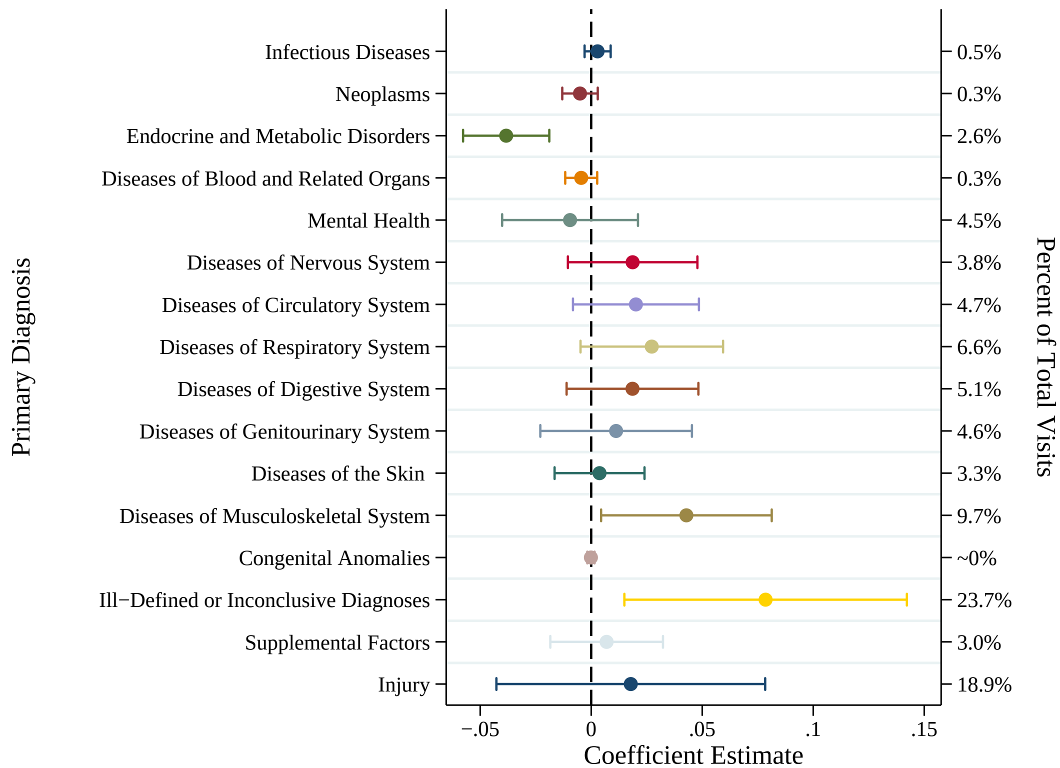

Note: This Figure plots the RD effect on ED visits per 1,000 population by primary diagnosis. Data are from HCUP NY SEDD and CA OSHPD ED visit data, years 2006-2017. Each estimate is made by estimating equation (1) using a linear fit, triangular kernel, and CCT optimal bandwidth. Percent of total visits is calculated based on the month before people turn 62. 95% confidence intervals generated using robust standard errors.

Figure 6: The Effect on ED Visits Remains Stable Over the Time Period

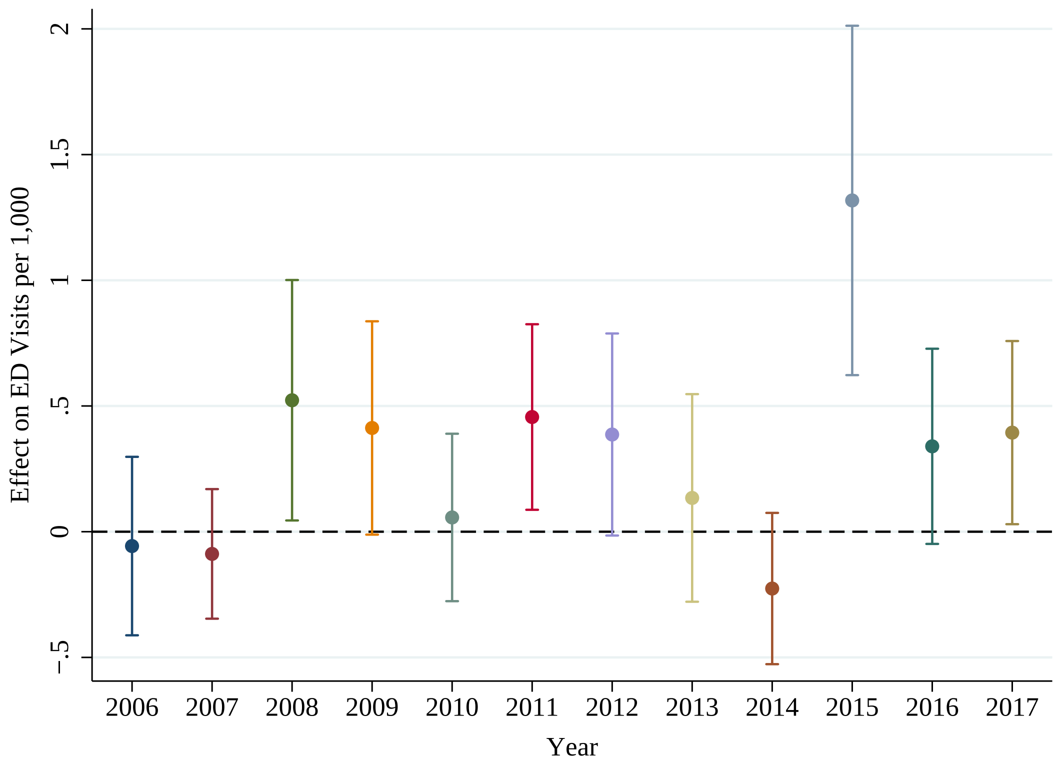

Note: This figure plots the RD effect on ED visits per 1,000 population by year. Data are from HCUP NY SEDD and CA OSHPD ED visit data, years 2006-2017. Each estimate is made by estimating equation (1) using a linear fit, triangular kernel, and CCT optimal bandwidth. 95% confidence intervals generated using robust standard errors.

Figure 7: Aggregate Inpatient Hospitalizations Do Not Change at 62

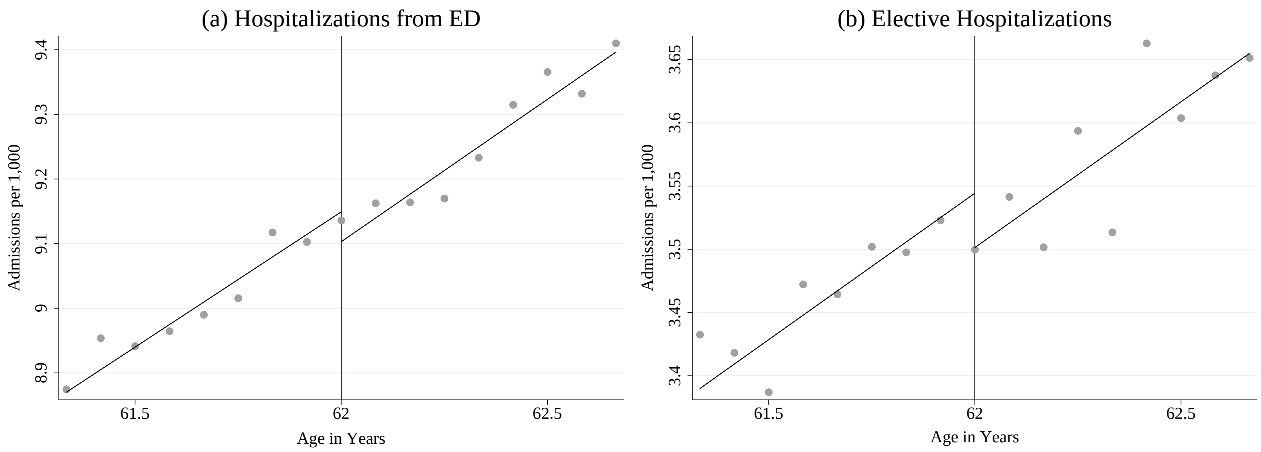

Note: This figure plots inpatient admissions per 1,000 population by age. Data are from HCUP NY SID and CA OSHPD PDD data, years 2006-2017. Each age bin corresponds to the calendar month that people turned each age (e.g., admissions in the calendar month people turn 62). Linear fit generated from estimating equation (1) with linear fit, triangular kernel, and 8-month bandwidth.

Figure 8: Ambulatory Surgery Encounters in CA Do Not Change at 62

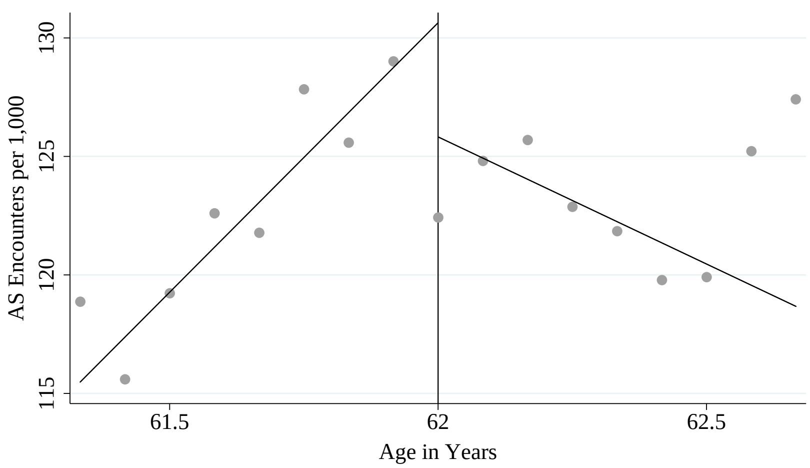

Note: This figure plots ambulatory surgery ancounters per 1,000 population by age. Data are from CA OSHPD AS data, years 2006-2017. Each age bin corresponds to the calendar month that people turned each age (e.g., admissions in the calendar month people turn 62). Linear fit generated from estimating equation (1) with linear fit, triangular kernel, and 8-month bandwidth.

Figure 9: Effect on ED Visits is Correlated with First Stage Outcomes Across Subgroups

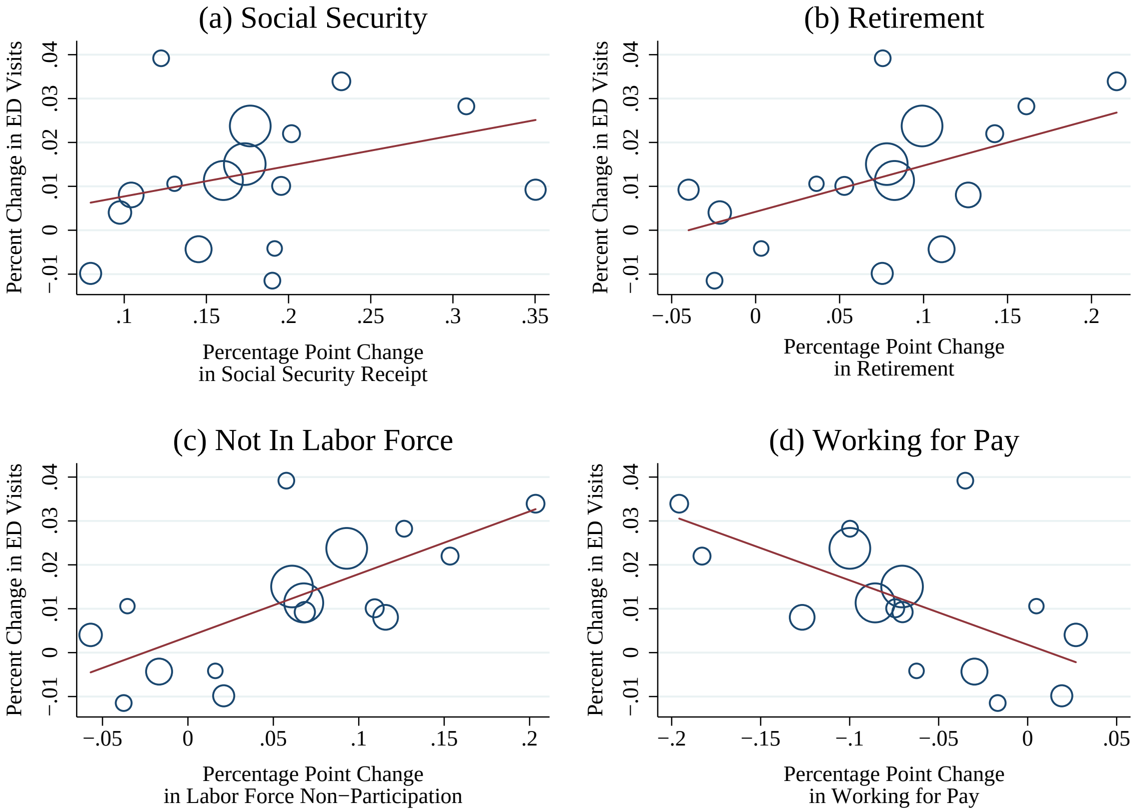

Note: These figures plot the percentage point effect of turning 62 on each first stage outcome against the percent change on ED visits. Points are weighted by the count of ED visits in that subgroup for people aged 61 and 11 months and the size of the circle is directly proportionate to each weight. Each line is a linear fit of the weighted points.

Figure 10: Employer-Sponsored Insurance Decreases at 62, but Overall Rates Stay the Same

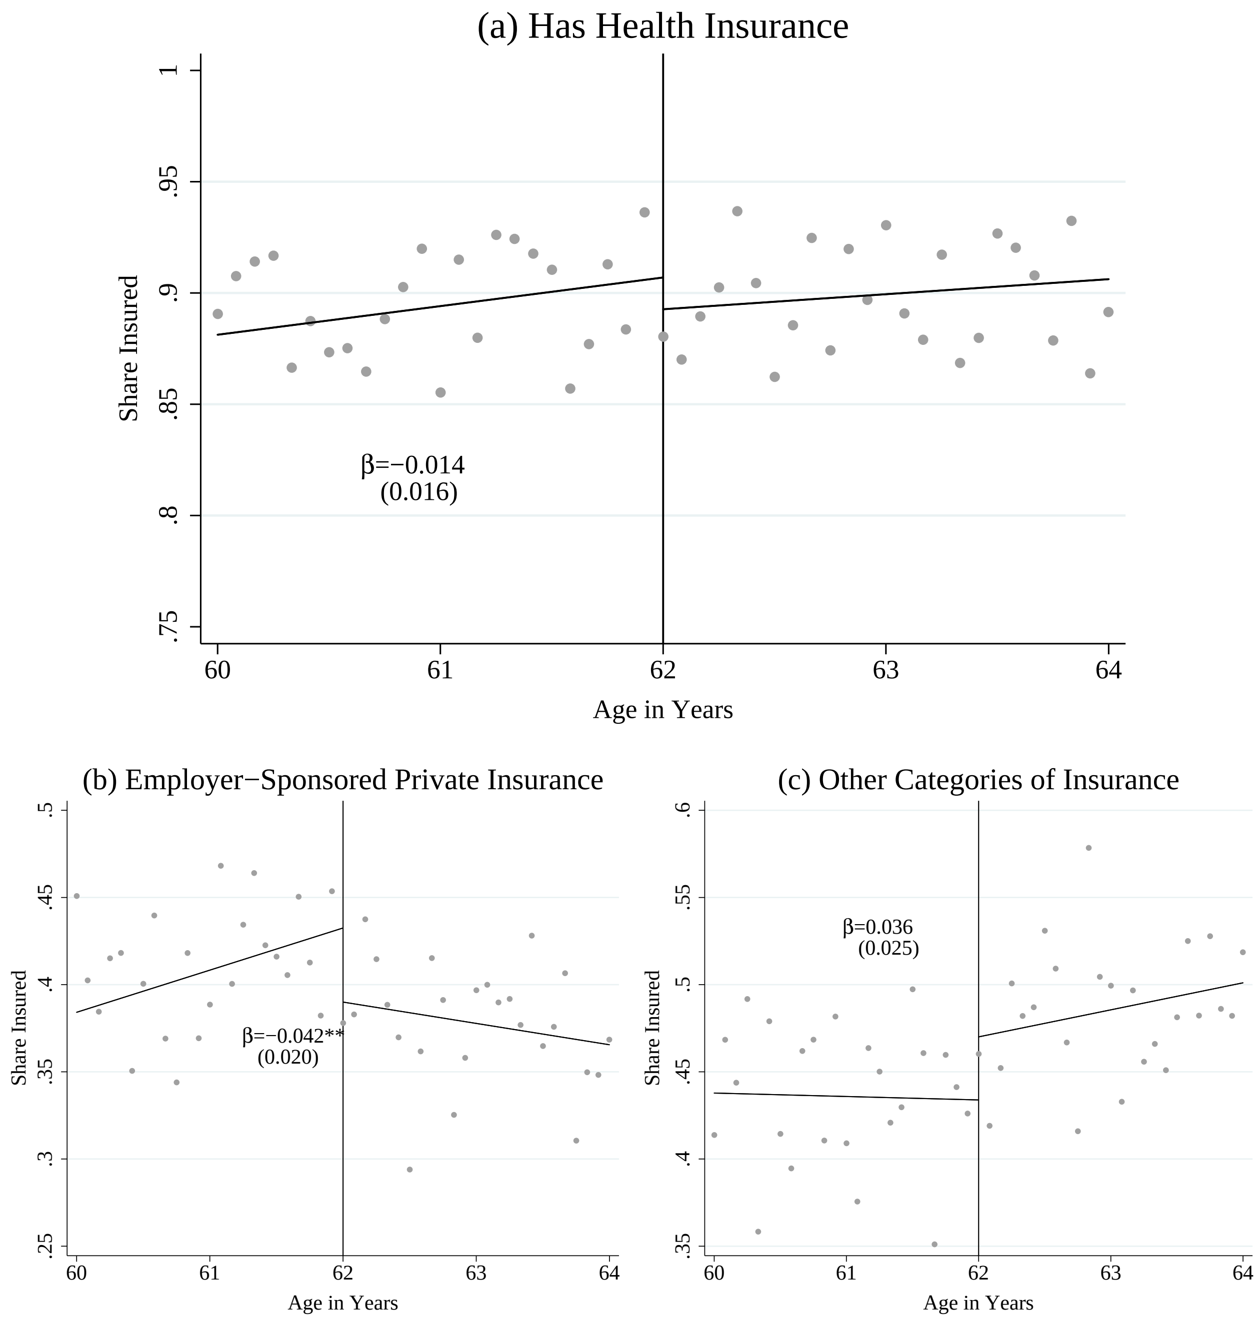

Note: Panel (a) plots the share of share of people with at least one health insurance plan by age. Panel (b) plots the share of people with an employer-sponsored private insurance plan. Panel (c) plots the share of people with at least one form of another type of insurance. Data are from the Health and Retirement Study waves 2006-2018. Linear fits generated by estimating equation (2) with 24-month bandwidth and triangular kernel. Each age bin corresponds to the calendar month that people turned each age (e.g., fraction insured in the calendar month people turn 62). Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Figure 11: Average Household Income Does Not Change Meaningfully at 62

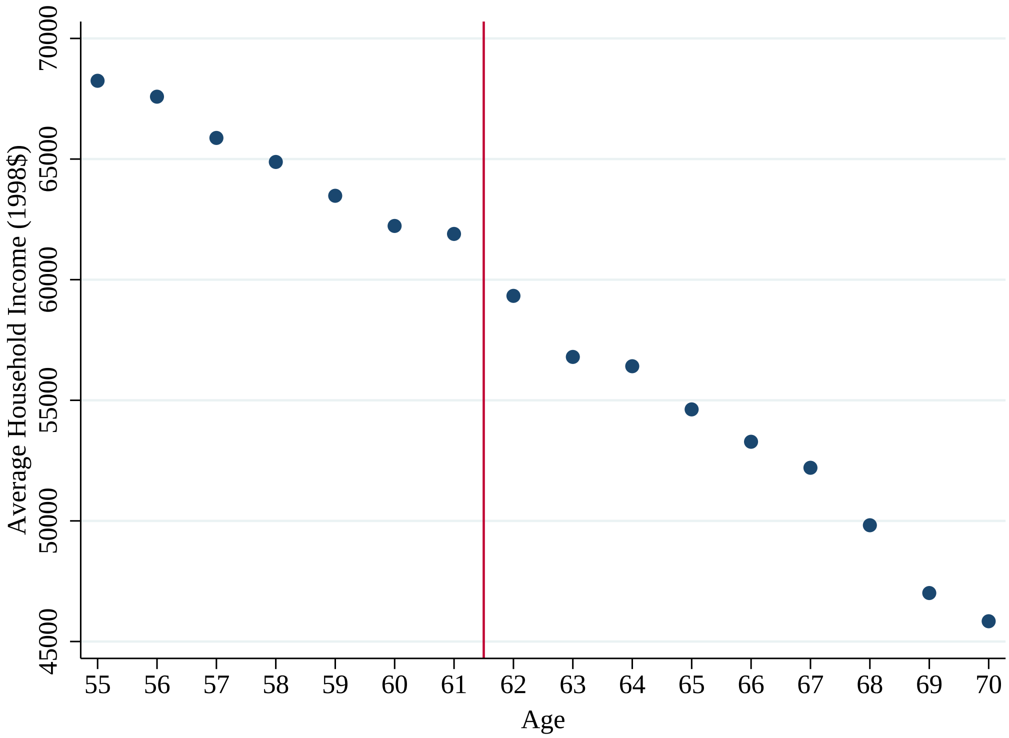

Note: This figure plots average household income by year of age from the CPS-ASEC years 2006-2017.

Figure 12: People Spend Discontinuously Fewer Hours per Week Working at 62

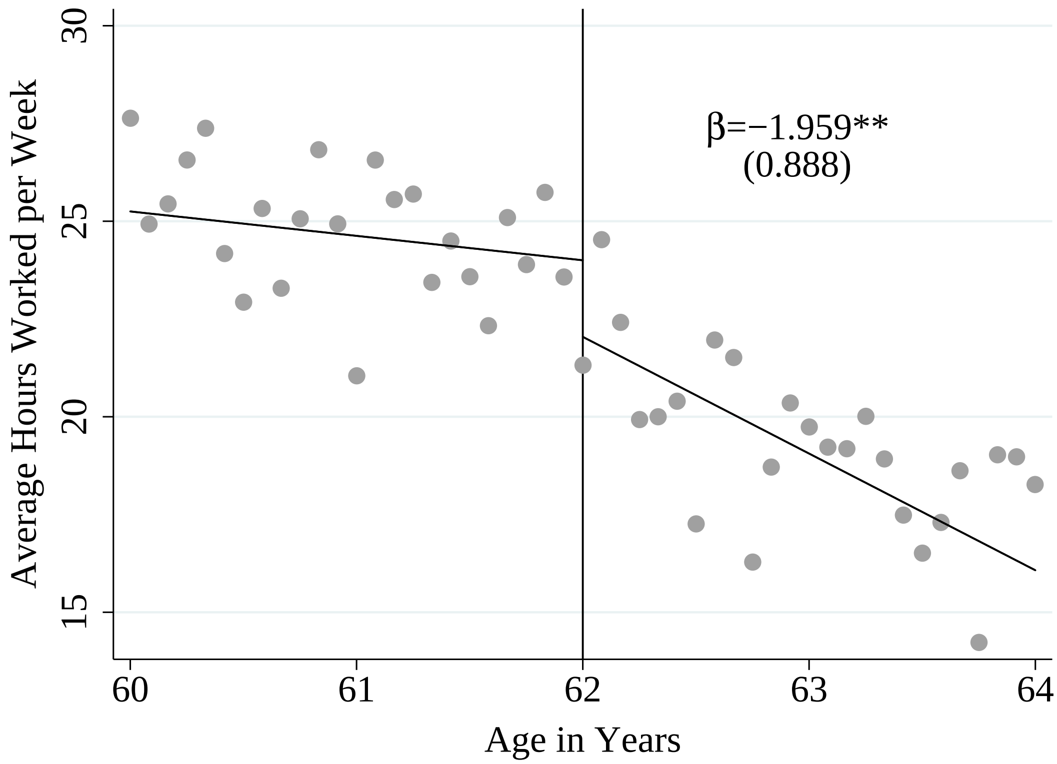

Note: This figure plots average hours of work per week by age. Data are from the Health and Retirement Study waves 2006-2018. Each age bin corresponds to the calendar month that people turned each age (e.g., admissions in the calendar month people turn 62). Linear fit generated from estimating equation (2) with linear fit, triangular kernel, and 24-month bandwidth. Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Tables

Table 1: How the Effects on Social Security and Labor Market Outcomes Vary by Sex and Race

Outcome:	(1) **Social Security**	(2) **Retired**	(3) **Not in Labor Force**	(4) **Work for Pay**
Sex:				
Male	0.159***	0.082***	0.067**	-0.085***
	(0.050)	(0.022)	(0.033)	(0.032)
Female	0.173***	0.078***	0.061	-0.071*
	(0.041)	(0.028)	(0.038)	(0.041)
Race:				
White	0.176***	0.099***	0.093***	-0.099***
	(0.041)	(0.015)	(0.031)	(0.032)
Hispanic	0.145***	0.111***	-0.017	-0.030
	(0.058)	(0.039)	(0.072)	(0.075)
Black	0.097*	-0.021	-0.057	0.027
	(0.052)	(0.043)	(0.045)	(0.042)
Other Races/Ethnicities	0.350***	-0.040	0.068	-0.070
	(0.095)	(0.063)	(0.089)	(0.086)

Note: Regressions are estimated using equation (2) with the Health and Retirement Study waves 2006-2018. All regressions use triangular kernel, linear polynomial, and 24-month bandwidths. Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Table 2: Regression Estimates of the Effect on ED Visits per 1,000 Population

	(1)	(2)	(3)	(4)
Sample: Both States				
Aged 62+	0.204***	0.409***	0.292***	0.397***
	(0.064)	(0.108)	(0.074)	(0.105)
			[0.001]	[0.001]
BW	8	8	5.5	8.1
Sample: New York Only				
Aged 62+	0.234**	0.550***	0.346***	0.548***
	(0.096)	(0.162)	(0.107)	(0.162)
			[0.001]	[0.002]
BW	8	8	6.1	7.7
Sample: California Only				
Aged 62+	0.174**	0.267*	0.213**	0.232*
	(0.086)	(0.147)	(0.102)	(0.130)
			[0.121]	[0.157]
BW	8	8	5.6	9.8
Poly. Deg.	1	2	1	2
CCT?			X	X

Note: Regressions are estimated using equation (1) using NY HCUP SEDD and CA OSHPD ED data for years 2006-2017 with ED visits per 1,000 population as the outcome. All regressions use triangular kernels. Columns 1 and 3 estimate the model using a linear fit and columns 2 and 4 estimate the model with a quadratic fit. Columns 1 and 2 use a set bandwidth of 8 months and columns 3 and 4 use CCT optimal bandwidths. The sample range for CCT optimal bandwidth calculations is 24 months on either side of the cutoff. Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01. Brackets contain p-values from bias-corrected confidence intervals.

Table 3: How RD Estimates for Healthcare Encounters Vary by Sex and Race

Outcome:	(1) **ED Visits**	(2) **Inpatient Admissions - ED**	(3) **Inpatient Admissions - Elective**	(4) **AS Encounters**
Sex:				
Male	0.011**	-0.001	-0.012	0.015
	(0.005)	(0.009)	(0.010)	(0.013)
	[0.025]	[0.926]	[0.207]	[0.141]
Female	0.015**	-0.004	-0.016	-0.005
	(0.006)	(0.008)	(0.011)	(0.008)
	[0.033]	[0.491]	[0.151]	[0.726]
Race:				
White	0.024***	0.012	-0.010	0.010
	(0.006)	(0.010)	(0.010)	(0.010)
	[0.001]	[0.195]	[0.334]	[0.246]
Hispanic	-0.004	0.005	-0.006	-0.009
	(0.009)	(0.011)	(0.031)	(0.013)
	[0.753]	[0.658]	[0.930]	[0.407]
Black	0.004	-0.037***	0.002	0.000
	(0.010)	(0.013)	(0.032)	(0.022)
	[0.574]	[0.006]	[0.817]	[0.824]
Other Races/Ethnicities	0.009	-0.026*	-0.031	-0.012
	(0.008)	(0.015)	(0.029)	(0.014)
	[0.335]	[0.111]	[0.432]	[0.628]

Note: Regressions are estimated using equation (1) with triangular kernel, linear polynomial, and CCT optimal bandwidths. The sample range for CCT optimal bandwidth calculations is 24 months on either side of the cutoff. All outcomes are specified as the natural logs of various types of healthcare encounters. ED visit data are taken from NY HCUP SEDD and CA OSHPD ED data. Inpatient hospitalization data are taken from NY HCUP SID and and CA OSHPD PDD. Ambulatory surgery data are taken from CA OSHPD AS data. All data from years 2006-2017. Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01. Brackets contain p-values from bias-corrected confidence intervals.

Table 4: Regression Estimates of the Effect on Inpatient Hospitalizations per 1,000 Population

	(1)	(2)	(3)	(4)
Sample: Hospitalizations from the ED				
Aged 62+	-0.047	-0.022	-0.034	-0.027
	(0.047)	(0.081)	(0.057)	(0.079)
			[0.589]	[0.798]
BW	8	8	5.5	8.2
Sample: Elective Hospitalizations				
Aged 62+	-0.040	0.002	-0.038	0.010
	(0.026)	(0.004)	(0.025)	(0.042)
			[0.160]	[0.705]
BW	8	8	8.7	7.4
Poly. Deg.	1	2	1	2
CCT?			X	X

Note: Regressions come from estimating equation (1) using using NY HCUP SID and CA OSHPD PDD data for years 2006-2017. Outcomes are inpatient admissions per 1,000 population. All regressions use triangular kernels. The sample range for optimal bandwidth calculations is 24 months on either side of the cutoff. Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01. Brackets contain p-values from bias-corrected confidence intervals.

Table 5: Regression Estimates of the Effect on Ambulatory Surgery per 1,000 Population

	(1)	(2)	(3)	(4)
Sample: Aggregate				
Aged 62+	-0.029	0.014	0.035	0.137
	(0.060)	(0.104)	(0.074)	(0.105)
			[0.444]	[0.172]
BW	8	8	5.4	7.9
Poly. Deg.	1	2	1	2
CCT?			X	X

Note: Regressions come from estimating equation (1) using CA OSHPD AS data for years 2006-2017. Outcomes are ambulatory surgery encounters per 1,000 population. All regressions use triangular kernels. The sample range for optimal bandwidth calculations is 24 months on either side of the cutoff. Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01. Brackets contain p-values from bias-corrected confidence intervals.

Appendix A: Variable Construction

Population Denominator

I model my construction of the population denominator on the method developed in Arenberg et al. (2020). In this method, I approximate the population for each state-by-year-by-age-by-calendar month cell by combining historical vital statistics data on births by month between 1941-1958 with adjustments from state population estimates from the 10% 2010 Decennial U.S. Census (Ruggles et al., 2020). Specifically, the approximation for each cell's population value is written in equation (A.1):

$$\hat{pop}_{asym} = births_{asym} * \frac{pop2010_{asyq}}{births_{asyq}} \tag{A.1}$$

$\hat{pop}_{asym}$ is an estimate of the cohort size of people aged a months old in state s in year y and were born on month m. $births_{asym}$ is the number of people of age (in months) a in state s on year y ithat were born in calendar month m. Births alone is an sufficient measure of population size between 2006-2017 since people may have either died or moved between when they were born and this time period. Therefore, I adjust monthly birth counts each quarter by $\frac{pop2010_{asq}}{births_{asyq}}$, the ratio of the 2010 population of each age cohort by state[15] to $births_{asyq}$, which is the number of births aggregated to quarter-of-birth instead of month-of-birth. Since my outcomes are aggregated by state, age, and year, I must take the sum of $\hat{pop}_{asym}$across calendar months of birth to calculate the denominator for regressions. Thus, my final denominator is displayed in equation (A.2), where M is a set containing all months of birth for people of age a in state s and year y. The population estimate for each age, state, and year cell can be interpreted as the total number of people in a given state and year that are ever a given age (e.g., the total number of people who are ever 62 and 1 month old in NY in 2006).

$$\bar{pop}_{asy} = \sum_{m \in M} \hat{pop}_{asym} \tag{A.2}$$

Appendix Figure A.1 shows the smoothness of the population estimates within a two-year radius of the age 62 cutoff. Each dot represents the mean population for each month of age across

[15] Estimated by multiplying the 10% Census sample's population count by 10.

years and states. This figure indicates that my method does not produce "jumpy" population estimates that would inappropriately affect RD estimates.

Figure A.1: The Population Denominator is Smooth Through the Cutoff

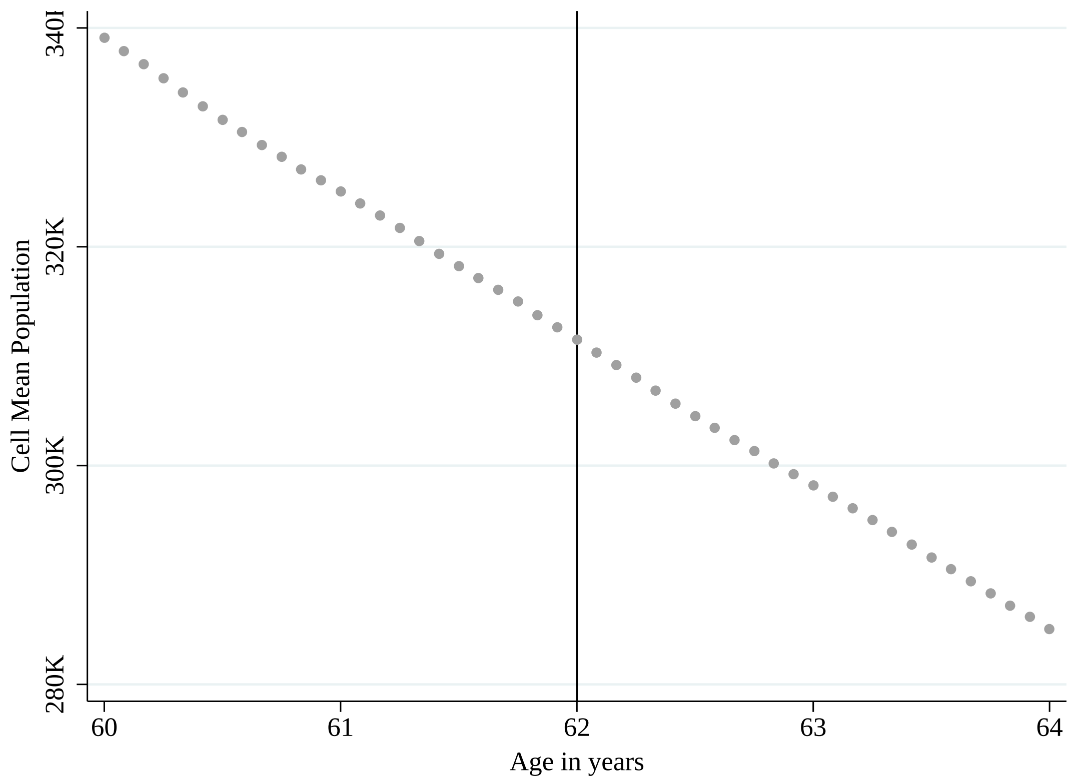

Note: This figure plots the estimated population counts by month of age. Estimates are derived from from 1941-1958 historical vital statistics data combined with the publicly available 10% extract of the 2010 U.S. Census.

Income Deciles

I define household income deciles in separate ways for the healthcare utilization data and for the HRS household survey data. The healthcare utilization data do not have information on individuals' incomes, but do report ZIP code of residence. First, I obtain data on the median household income of each ZIP code in the country as estimated using the ACS 2010-2014 five-year estimates (Manson et al., 2021). Then, I rank these ZIP codes into deciles and assign the ZIP codes in NY and CA to their appropriate deciles in the national distribution. I report the ceiling of each decile in the left column of Appendix Figure A.2. On the other hand, not only does the HRS have detailed information on household income, the longitudinal structure also allows me to determine income levels during from previous years. The benefit of this is that it prevent the treatment at 62 from affecting a household's income classification. I classify individuals in the HRS according their "lifetime" household income by, first, assigning them an income value equal to their earliest observable age before 61 ("measurement age). Then, for each year in the sample, I compare that income value to the national distribution of incomes from the 2006-2017 CPS-ASEC at the individual's measurement age to account for changes in the income distribution over time. I use this comparison to assign each individual a household income decile. To give an example of these deciles, I list the ceilings for household income deciles when pooling together the 2006-2017 CPS-ASEC years for people aged 55. Implementing these methods, Appendix Figure A.3 estimates how the effects on Social Security receipt and labor market outcomes vary by decile and Appendix Figure A.4 shows the corresponding figure for ED visits.

Figure A.2: ZIP Code and Household Income Decile Ceilings

	Median HH Income (zip codes)	HH Income at Age 55
Decile: 1	$30,536	$10,000
2	$36,250	$22,500
3	$40,587	$34,018
4	$44,309	$47,240
5	$48,302	$60,000
6	$52,462	$76,000
7	$57,813	$97,573
8	$65,484	$124,634
9	$79,688	$170,000
10	$250,001	$1,111,500

Note: The left-hand side column is produced using the ACS 2010-2014 five-year estimates. The right-hand side column is produced using the 2006-2017 CPS-ASEC.

Figure A.3: Effect on First Stages by Household Income Decile

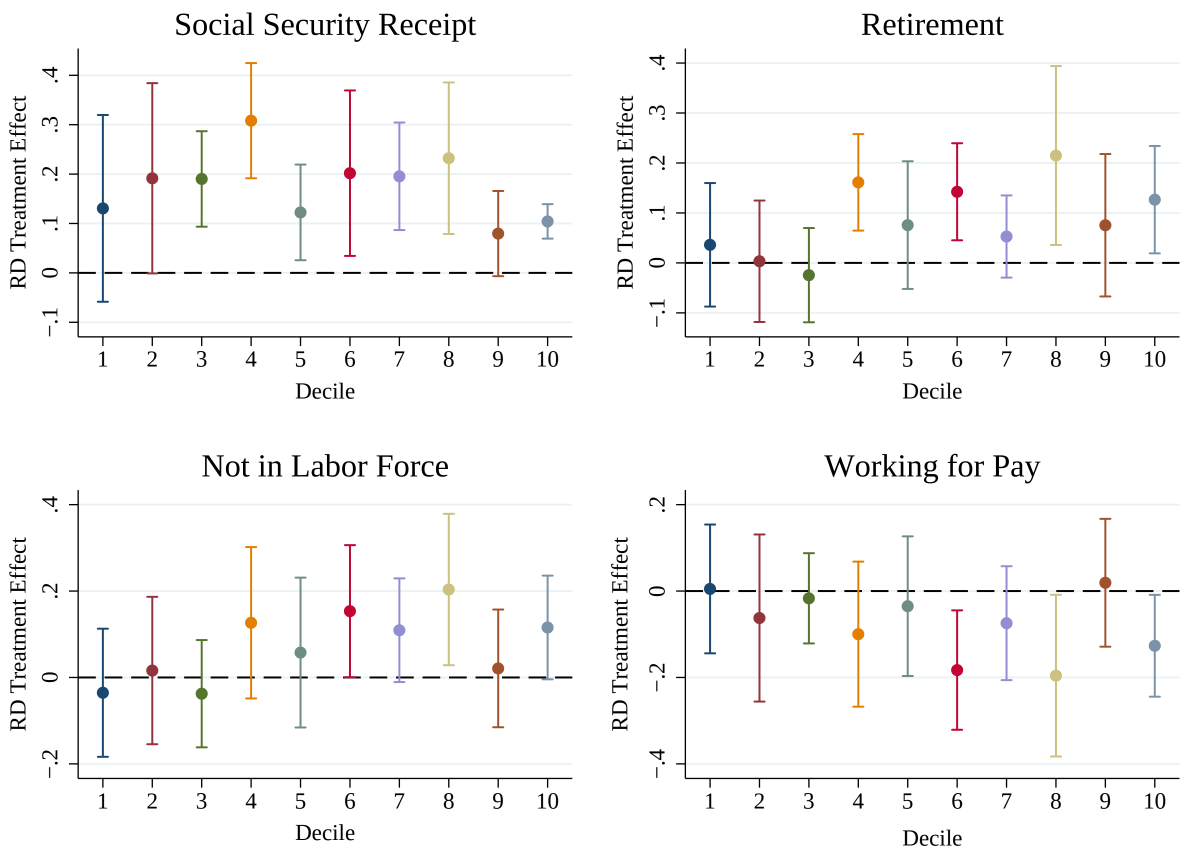

Note: This figure plots RD estimates for Social Security receipt, share retired, not in the labor force, and working for pay, by household income decile. Linear fits generated by estimating equation (2) with 24-month bandwidths using triangular kernels. Data are from the Health and Retirement Study waves 2006-2018. 95% confidence intervals are estimated using robust standard errors.

Figure A.4: Effect on ED Visits by Decile of Median Household Income by ZIP Code of Residence

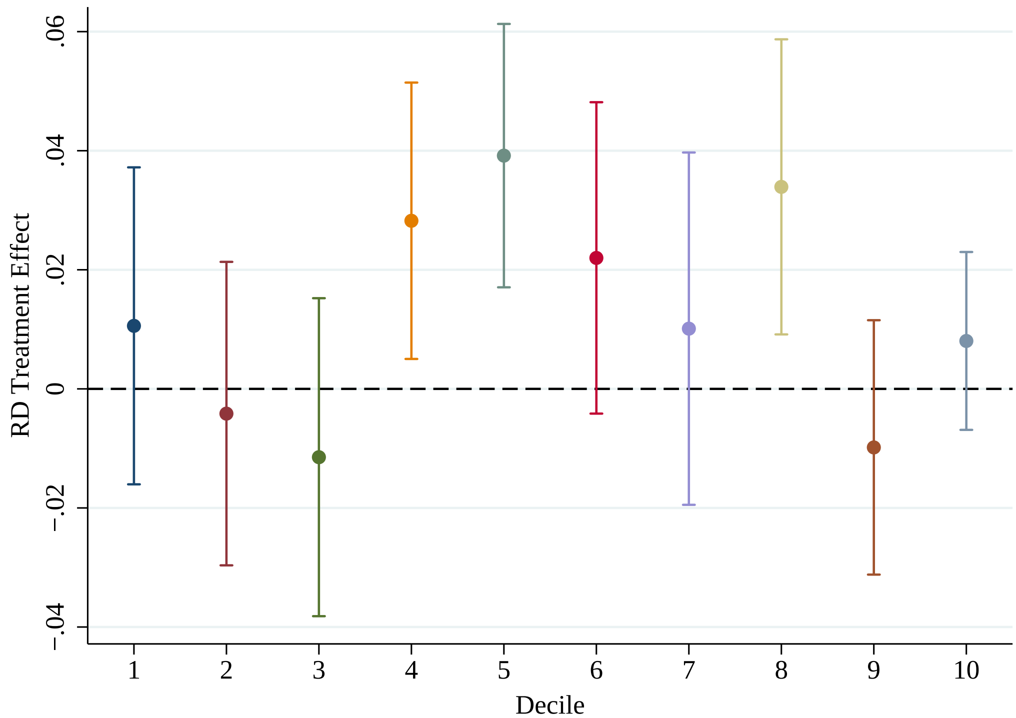

Note: This figure plots RD estimates for Social Security receipt, share retired, not in the labor force, and working for pay, by decile of median household income by ZIP code. Linear fits generated by estimating equation (1) with triangular kernel and CCT optimal bandwidth. The sample range for optimal bandwidth calculations is 24 months on either side of the cutoff. Data are from NY HCUP SEDD and CA OSHPD ED data, years 2006-2017. 95% confidence intervals are estimated using robust standard errors.

Appendix B: Supplemental Results

Figure B.1: Social Security Claiming Increases Meaningfully in NY and CA at 62, State and Local Pension Claiming Do Not

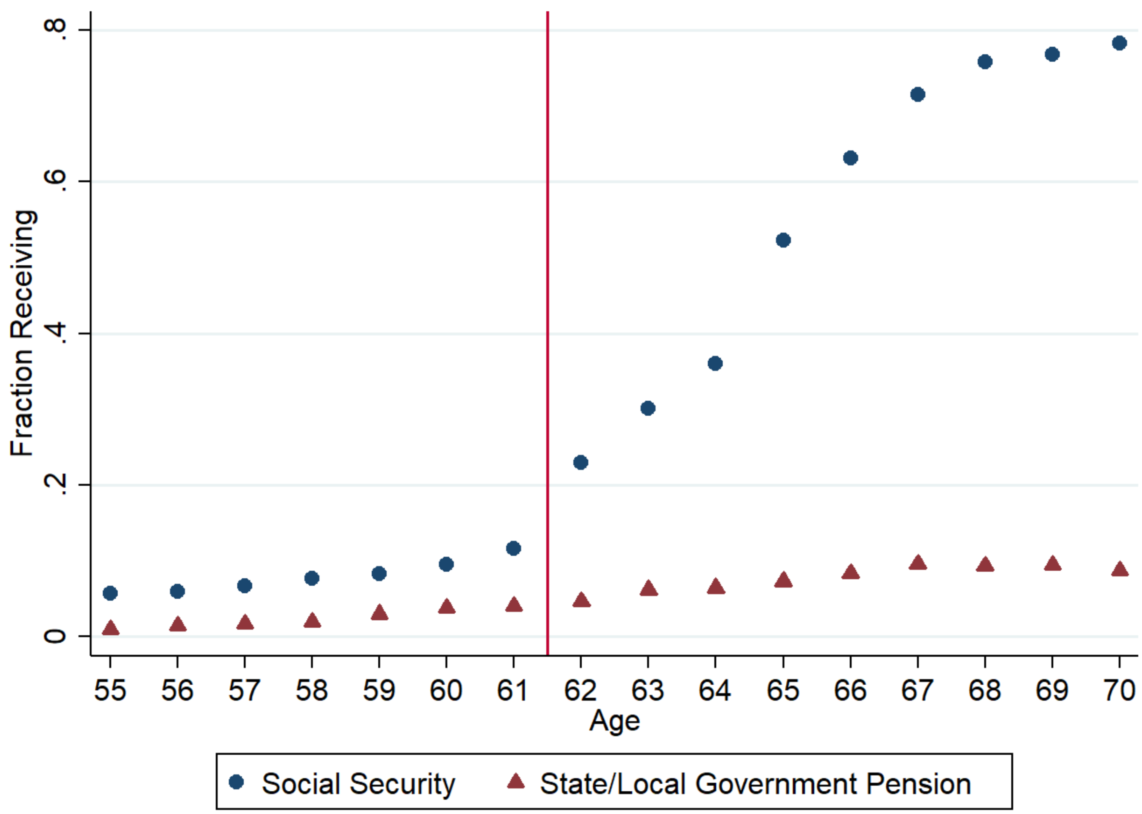

Note: This figure plots rates of Social Security and State/Local Government Pension receipt in NY and CA by age in 2006-2017. Data are taken from the CPS-ASEC 2007-2018 waves, which asks about people's sources of retirement income from the previous year. Since the CPS-ASEC is conducted every March and asks about the previous year, ages are calculated by subtracting one year from the respondent's current age in years.

Figure B.2: The Discontinuity in Labor Force Participation Occurs Immediately at 62

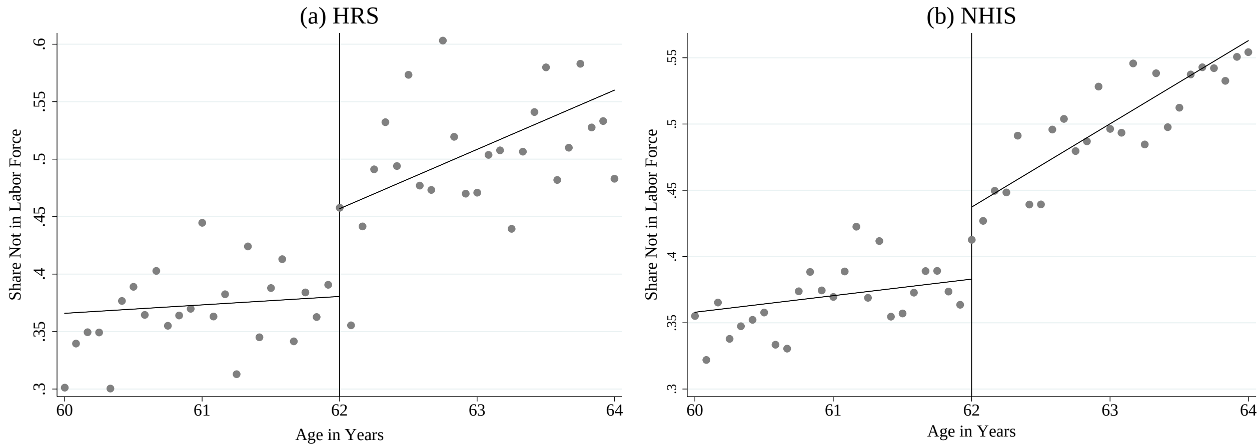

Note: This figure plots the share of the population currently not in the labor force. Linear fit obtained form estimating equation (2) with a 24-month bandwidth and triangular kernel. Data are from the Health and Retirement Study waves 2006-2014 and National Health Interview Survey waves 2006-2014. Each age bin corresponds to the calendar month that people turned each age (e.g., share working in the calendar month people turn 62).

Figure B.3: Effect on Social Security and Labor Market Outcomes by Year

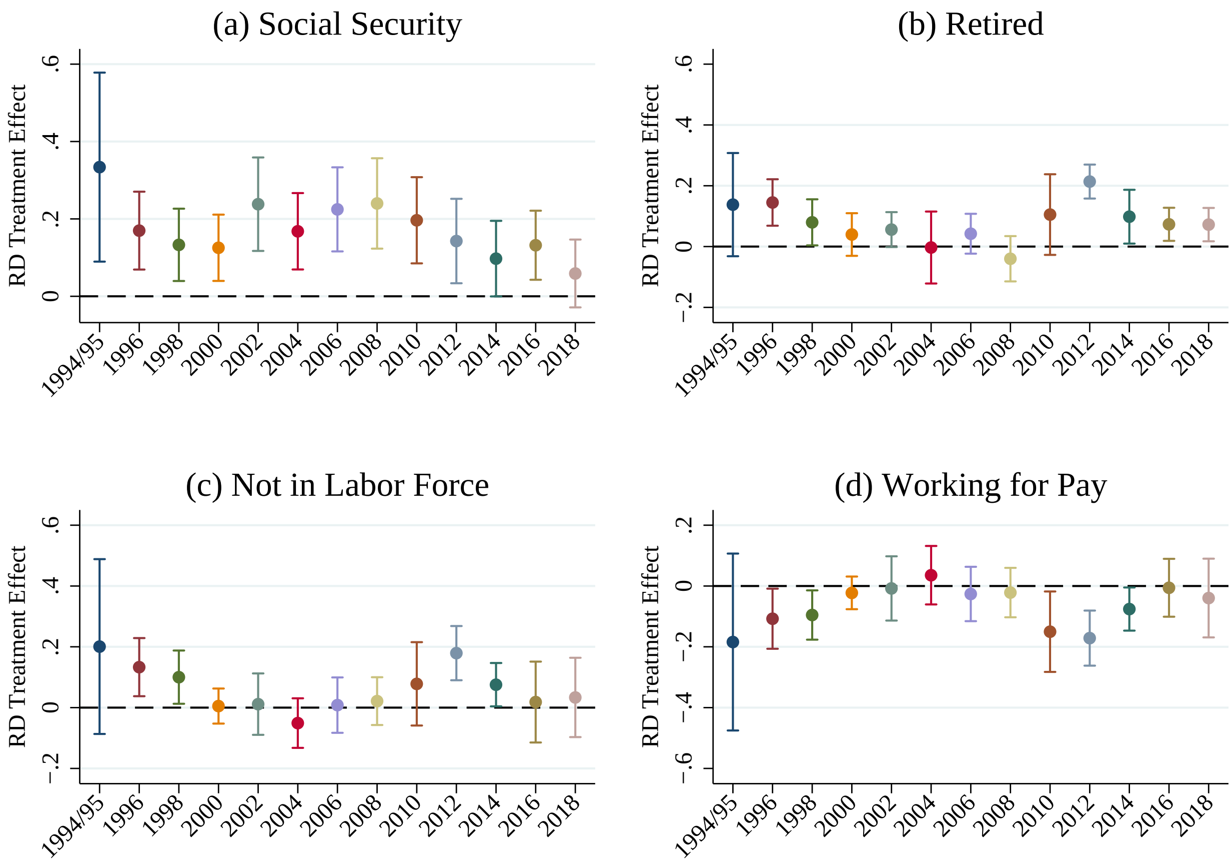

Note: This figure plots the effect of turning 62 on Social Security receipt and labor market outcomes by year by estimating β using equation (2). Data are from the Health and Retirement Study waves 2006-2018. Each specification controls for a dummy for the month people turn 62. Confidence intervals are derived from robust standard errors.

Figure B.4: The Discontinuity in Female Labor Force Participation is Also Present in the NHIS

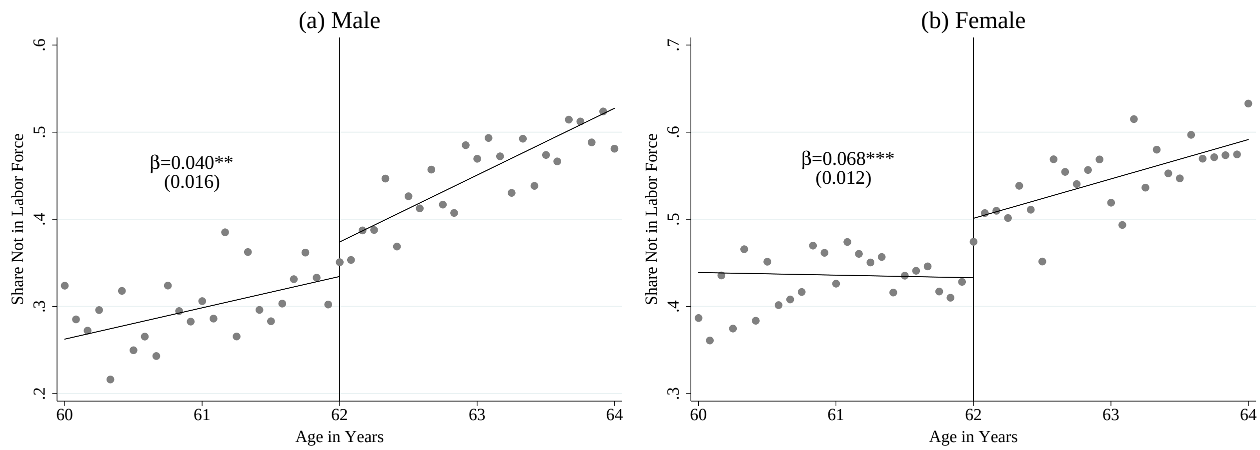

Note: This figure plots the share of the population currently not in the labor force. Linear fit obtained from estimating equation (2) with a 24-month bandwidth and triangular kernel. Data are from the National Health Interview Survey waves 2006-2014. Each age bin corresponds to the calendar month that people turned each age (e.g., share working in the calendar month people turn 62).

Figure B.5: Estimated Effect on ED Visits is Robust to Many Bandwidths, Particularly Narrow Ones

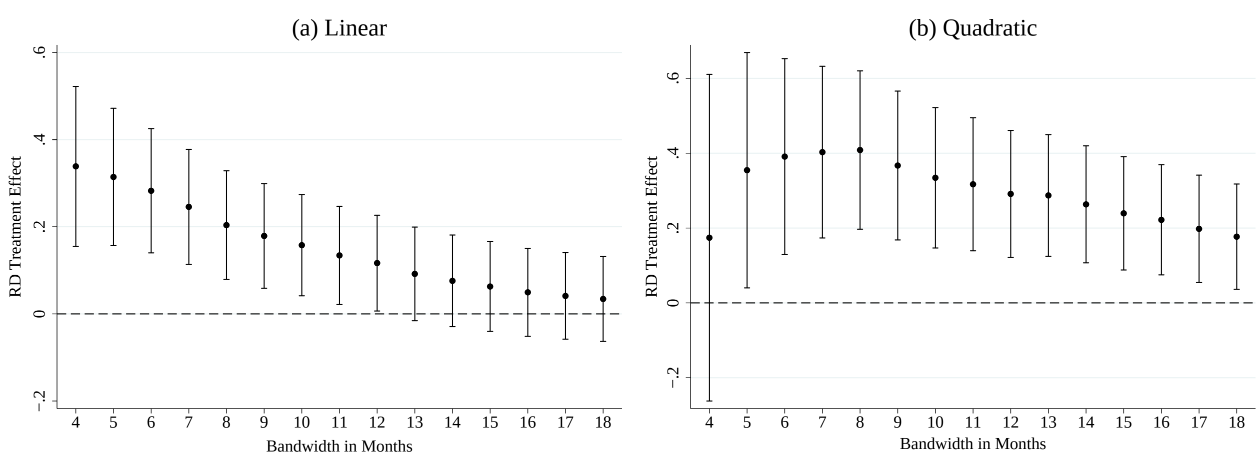

Note: This figure plots RD estimates using equation (1) by bandwidth in months and polynomial fit. Data come from NY HCUP SEDD and CA OSHPD ED for years 2006-2017. Outcomes are ED visits per 1,000 population. All regressions estimated with triangular kernels. 95% confidence intervals are calculated using robust standard errors.

Figure B.6: The Effect on ED Visits and Robust t-Statistic are Larger in Magnitude than 95% of Placebo Estimates

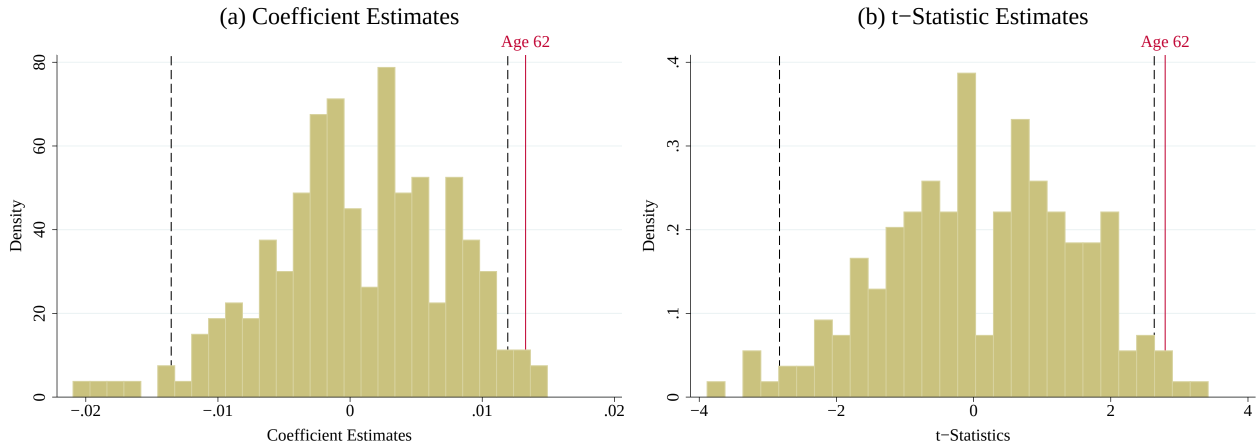

Note: This figure plots the distribution of RD coefficients and robust t-statistics, estimated by equation (1), at the age 62 cutoff and a set of placebo age cutoffs within a +/- 10-year radius. Data come from NY HCUP SEDD and CA OSHPD ED data for years 2006-2017. The outcomes are the natural log of aggregate ED visits. All regressions estimated with linear fits, CCT optimal bandwidths, and triangular kernels. The sample range for CCT optimal bandwidth calculations is 24 months on either side of the cutoff. I do not include coefficients using placebo cutoffs within 8-month radiuses of age 62 or age 65, or a cutoff at 65 itself, in order to avoid treatment effect contamination (Calonico and Titiunik, 2021). 95% of the estimated coefficients and t-statistics fall within the dotted black lines, and the solid red line indicates the estimate at age 62.

Figure B.7: Effect on Insurance Coverage by Type of Payer

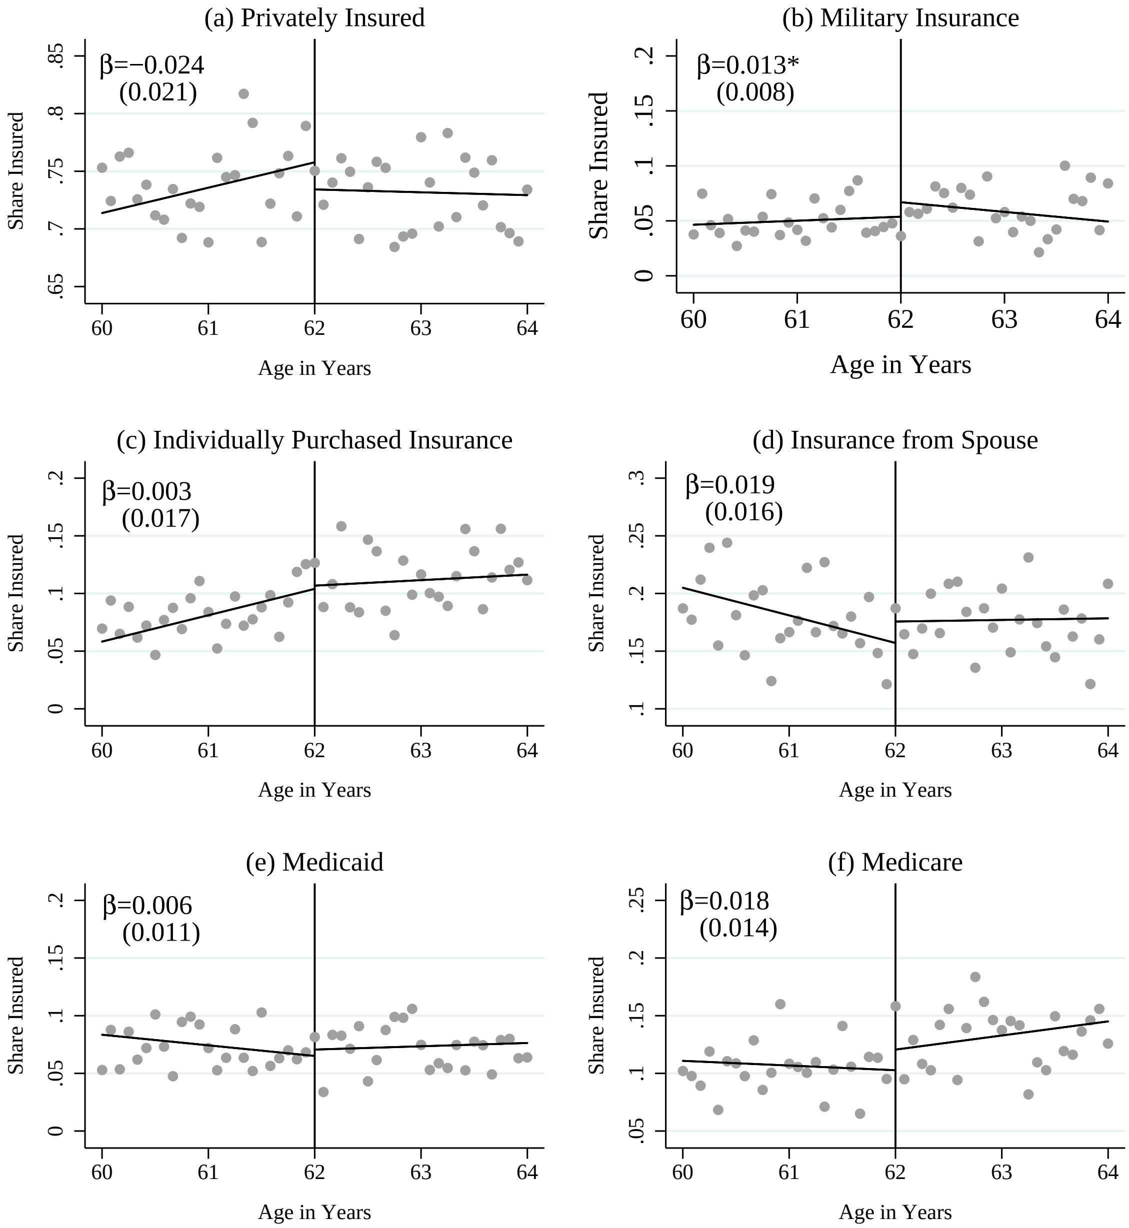

Note: This figure plots the rates of of insurance coverage by age and type of payer. Data are from the Health and Retirement Study waves 2006-2018. Linear fits are generated by estimating equation (2) with 24-month bandwidths and triangular kernel. Each age bin corresponds to the calendar month that people turned each age (e.g., share retired in the calendar month people turn 62). Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Figure B.8: Effect on Number of Insurance Policies

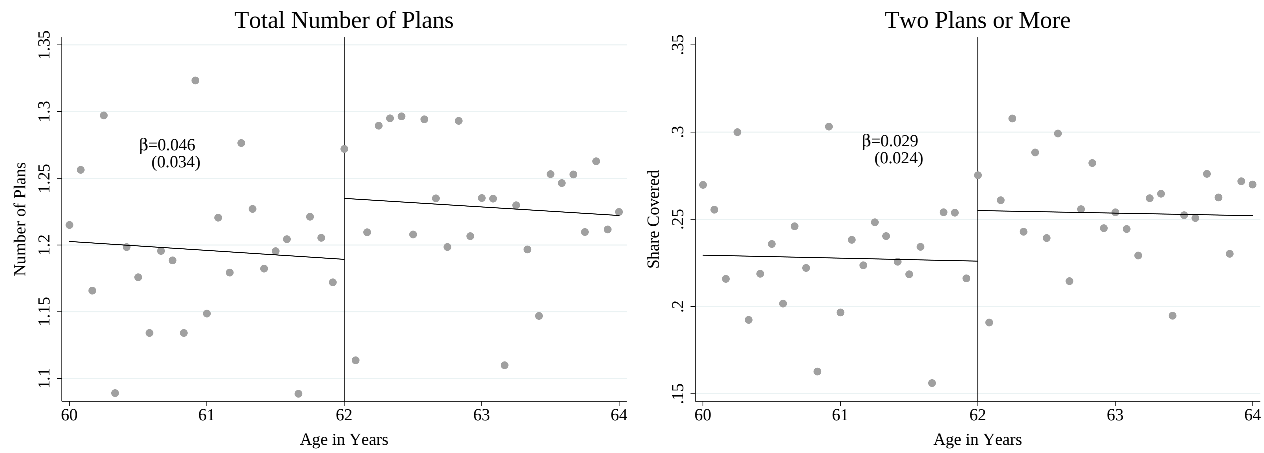

Note: This figure plots the rates of of insurance coverage by age. Data are from the Health and Retirement Study waves 2006-2018. Linear fits are generated by estimating equation (2) with 24-month bandwidths and triangular kernel. Each age bin corresponds to the calendar month that people turned each age (e.g., share retired in the calendar month people turn 62). Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Figure B.9: Effect on Share and Rate of ED Visits Among Privately Insured

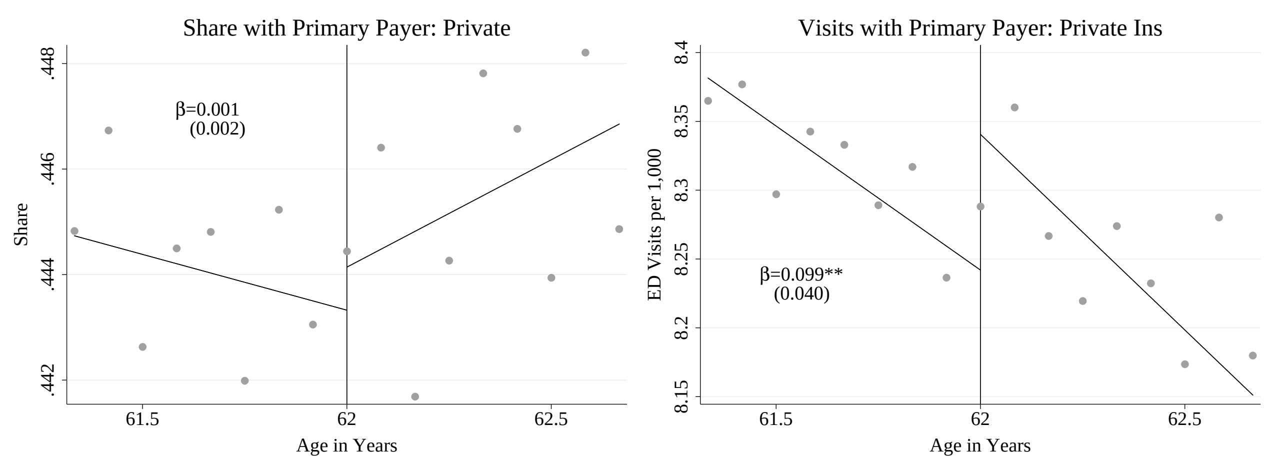

Note: This figure plots the estimated discontinuities in the share and rate per 1,000 population of privately insured patients. Data come from NY HCUP SEDD and CA OSHPD ED data for years 2006-2017. Linear fits are generated by estimating equation (1) with 8-month bandwidths and triangular kernel. Each age bin corresponds to the calendar month that people turned each age (e.g., share retired in the calendar month people turn 62). Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Table B.1: RD Donut Specifications for Effect on ED Visits

	One-Month Radius		Two-Month Radius	
	(1)	(2)	(3)	(4)
Aged 62+	0.257***	0.505***	-0.063	-0.085
	(0.073)	(0.105)	(0.069)	(0.092)
	[0.003]	[0.000]	[0.518]	[0.520]
BW	5.7	8.1	6.4	10.5
Poly. Deg.	1	2	1	2
CCT?	X	X	X	X

Note: Regressions come from estimating equation (1) using data from NY HCUP SEDD and CA OSHPD ED data for years 2006-2017 for years 2006-2017. Outcomes are ED visits per 1,000 population. All regressions use triangular kernels. Each model controls for state-by-year fixed effects a dummy for the month people turn 62. The sample range for optimal bandwidth calculations is 24 months on either side of the cutoff. Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01. Brackets contain p-values from bias-corrected confidence intervals.

References

[1] Berman, J. (2020). Can Income Buy Health? Evidence from Social Security Benefit Discontinuities and Medicare Claims. Working Paper.

[2] Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014). Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs. *Econometrica 82 (6)*, 2295-2326.

[3] Cattaneo, M. D. & Titiunik, R. (2021). Regression Discontinuity Designs. Working Paper.

[4] Charles, K. K. (2004). Is Retirement Depressing?: Labor Force Inactivity and Psychological Well Being in Later Life

[5] Chuard-Keller, P. (2021). With Booze, you Lose: The Mortality Effects of Early Retirement. Working Paper.

[6] Coe, N. B. & Lindeboom, M. (2008). Does Retirement Kill You? The Evidence from Early Retirement Windows. CentER Discussion Paper Series No. 2008-93. Tilburg, ND: Tilburg University. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1295315

[7] Coe, N. B., von Gaudecker, H. M., Lindeboom, M., & Maurer, J. (2012). The Effect of Retirement on Cognitive Functioning. *Health Economics, 21 (8)*, pp. 913-927.

[8] Coe, N. B. & Zamarro, G. (2015). Does retirement impact health care utilization. CESR-Schaeffer Working Paper No. 2015-032.

[9] Dave, D., Rashad, I., & Spasojevic, J. (2008). The Effect of Retirement on Physical and Mental Health Outcomes. *Southern Economic Journal, 75 (2)*. pp. 497-523.

[10] Dong, Y. (2014). Regression Discontinuity Applications with Rounding Errors in the Running Variable. *Journal of Applied Econometrics, 30 (3)*, pp. 422-446.

[11] Dushi, I., Iams, H. M., and Lichtenstein, J. (2011). Assessment of retirement plan coverage by firm size using W-2 tax records. Social Security Bulletin, 71 (2) (2011), pp. 53-65.

[12] Dushi, I., Iams, H. M., & Trenkamp, B. (2017). The Importance of Social Security Benefits to the Income of the Aged Population. Social Security Bulletin, 77 (2) (2017).

[13] 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*, pp. 121-137.

[14] Fitpatrick, M. D. (2020). Does Working Longer Enhance Old Age? Wharton Pension Research Council Working Paper No 2020-20.

[15] Frimmel, M. D. & Pruckner, G. J. (2020). Retirement and healthcare utilization. *Journal of Public Economics, 184*, 104146.

[16] Sarah Flood, Miriam King, Renae Rodgers, Steven Ruggles and J. Robert Warren. Integrated Public Use Microdata Series, Current Population Survey: Version 7.0 [dataset]. Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D030.V7.0

[17] Gorry, A., Gorry, D., & Slavov, S. N. (2018). Does retirement improve health and life satisfaction? *Health Economics, 27 (13):* 2067-2086.

[18] Gustman, A. L. & Steinmeier, T. L. (2005). The social security early entitlemnt age in a structural model of retirement and wealth. *Journal of Public Economics, 89 (2-3),* pp. 441-463. doi:10.1016/j.jpubeco.2004.03.007.

[19] Gross, T., Layton T. J., & Prinz, D. (forthcoming). The Liquidity Sensitivity of Healthcare Consumption. *American Economic Review: Insights*.

[20] Hahn, J., Todd, P., & van der Klaauw, W. (2001). Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design. *Econometrica, 69 (1)*, pp. 201-209.

[21] Insler, M. (2014). The Health Consequences of Retirement. *Journal of Human Resources, 49 (1)*, pp. 195-233.

[22] Koster, A., Caserotti, P. *et al.* (2012). Association of sedentary time with mortality independent of moderate to vigorous physical activity. *PLoS One, 7 (6)*, Article e37696.

[23] Lee, D. S. & Lemieux, T. (2010). Regression Discontinuity Designs in Economics. *Journal of Economics Literature, 48 (2)*, 281-355.

[24] Lucifora, C. & Vigani, D. (2018). Health care utilization at retirement: The role of the opportunity cost of time. *Health Economics, 27 (12)*, pp. 2030-2050.

[25] Matthews. C E., George, S. M., Moore, S. C., *et al.* (2012). Amount of time spent in sedentary behaviors and cause-specific mortality in US adults. *American Journal of Clinical Nutrition, 95*, pp. 437-445.

[26] McSteen, M. A. (1985). Fifty Years of Social Security. Social Security Bulletin, 48 (8), pp. 36-44.

[27] Munnell, A. H. (2015). The Average Retirement Age - An Update. *Issue in Brief* 15-4. Chestnut Hill, MA: Center for Retirement Research at Boston College.

[28] Munnell, A. H., Sanzenbacher, G. T., Webb, A., & Gillis, C. M. (2016). Are Early Claimers Making a Mistake? CRR WP 2016-5.

[29] Li, X., Hurd, M., & Loughran, D. S. (2008). The Characteristics of Social Security Benefits Who Claim Benefits at the Early Eligibility Age. AARP Public Polbicy Institute Research Report No. 2008-19. AARP, Washington DC.

[30] Li, Z. (2021). The Social Security Retirement Age. Congressional Research Service R44670.

[31] Lynn A. Blewett, Julia A. Rivera Drew, Miriam L. King and Kari C.W. Williams. IPUMS Health Surveys: National Health Interview Survey, Version 6.4 [dataset]. Minneapolis, MN: IPUMS, 2019. https://doi.org/10.18128/D070.V6.4

[32] Steven Manson, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles. IPUMS National Historical Geographic Information System: Version 16.0 [dataset]. Minneapolis, MN: IPUMS. 2021. http://doi.org/10.18128/D050.V16.0

[33] Neuman K. (2008). Quit Your Job and Get Healthier? The Effect of Retirement on Health. *Journal of Labor Research, 29 (2)*, pp. 177-201.

[34] Nielsen, N. F. (2019). Sick of retirement? *Journal of Health Economics, 65*, pp.133-152.

[35] Rose, L. (2020). Retirement and health: Evidence from England. *Journal of Health Economics, 73*, 102352.

[36] Ruggles, S., Flood, S., Goeken, R., Grover, J., Meyer, E., Pacas, J., & Sobek, M. IPUMS USA: Version 10.0 [dataset]. Minneapolis, MN: IPUMS, 2020. https://doi.org/10.18128/D010.V10.0

[37] Rutledge, M. S. & Wettstein, G. (2020). Is Nontraditional Work At Older Ages Associated with Better Retirement Security? CRR WP 2020-13.

[38] Rust, J. & Phelan, C. (1997). How Social Security and Medicare Affect Retirement Behavior In a World of Incomplete Markets. *Econometrica, 65 (4)*, pp. 781-831.

[39] Snyder, S. E. & Evans, W. E. (2006). The effect of income on mortality: Evidence from the Social Security Notch. *Review of Economics and Statistics*, *88* (3): 482-495.

[40] Social Security Administration (SSA). (2020). Social security beneficiaries by benefit level. Social Security. https://www.ssa.gov/oact/progdata/benefitlevel.html?type=ra.

[41] Social Security Administration (SSA). (2021a). Population Profile: Never Beneficiaries, Aged 60 or Older, 2021. Retrieved from: https://www.ssa.gov/policy/docs/population-profiles/never-beneficiaries.html

[42] Social Security Administration. (2021b). Retirement Benefits. Publication No. 05-10035.

[43] Social Security Administration. (2021c). How Work Affects Your Benefits. Publication No. 05-10069.

[44] Stancanelli, E. & Van Soest, A. (2016). Partners' leisure time truly together upon retirement. *IZA Journal of Labor Policy, 5(1)*, pp. 1-19.

[45] Uscher-Pines, L., Pines, J., Kellermann, A., Gillen, E., & Mehrotra, A. (2013). Deciding to Visit the Emergency Department for Non-Urgent Conditions: A Systematic Review of the Literature. *American Journal of Managed Care, 19 (1)*, pp. 47-59.

[46] Van de Water, P. N. & Romig, K. (2020). Social Security Benefits are Modest. Center of Budget and Policy Priorities and Policy: Policy Futures.

[47] Waldron, H. (2020a). Trends in Working and Claiming Behavior at Social Security's Early Eligibility Age by Sex. ORES Working Paper Series No. 114. Washington, DC: SSA, Office of Retirement and Disability Policy, Office of Research, Evaluation, and Statistics.

[48] Waldron, H. (2020b). Working and Claiming Behavior at Social Security's Early Eligibility Age Among Men by Lifetime Earnings Decile. ORES Working Paper Series No. 115. Washington, DC: SSA, Office of Retirement and Disability Policy, Office of Research, Evaluation, and Statistics.

[49] World Bank. (2021). Life expectanc at birth, total (years). *World Development Indicators*. The World Bank Group, https://databank.worldbank.org/reports.aspx?source=world-development-indicators.

[50] Zhang, Y., Salm, M., & van Soest, A V. (2018). The effect of retirement on healthcare utilization: Evidence from China. *Journal of Health Economics, 62*, pp. 165-177.

[51] Zhou, Q., Eggleston, K., & Liu, G. G. (2021). Healthcare utilization at retirement in China. *Health Economics*, early view.