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

Matthew T. Knowles†

Vanderbilt University

March 20, 2022

Click for Link to Latest Version

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 (ED) visits that do not result in hospitalization by 1-2% at this age. Further analysis demonstrates that this effect is driven by both emergent and nonemergent conditions and is not completely explainable by changes in health insurance status.

*I am grateful to my dissertation commitee, Kitt Carpenter, Carolyn Heinrich, Analisa Packham, and Michelle Marcus, for their feedback and support on this project. I also thank the applied microeconomics advising group at Vanderbilt University, participants at the Southern Economic Association 2021 conference, and seminar attendees at Abt Associates, RTI International, and the Congressional Budget Office for their helpful 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.

†Doctoral Student in Economics, Vanderbilt University, email: matthew.t.knowles@vanderbilt.edu

1 Introduction

The depletion of the Social Security trust fund is a major policy concern in the United States. Many of the current and proposed solutions to this issue, such as increasing the Full Retirement Age, may result in people postponing Social Security benefits and/or retirement. While these efforts may help preserve the trust fund, it's possible that disincentivizing retirement could unintentionally affect social welfare through alternate channels. For example, high quality, quasi-experimental work has demonstrated a causal link between retirement and mortality in the United States (Fitzpatrick and Moore, 2018). However, this effect on mortality may only represent one dimension of the total effect of retirement on health in the U.S. In order to more fully understand this relationship, it is also imperative to examine how retirement impacts other important health outcomes such as healthcare utilization. This is an important gap in the literature on Social Security policy, retirement, and health for several reasons. First, mortality is a particularly extreme measure of health status and it is possible that retirement may also affect people's morbidity. This could be reflected by changes in the utilization of acute care services upon retirement, if healthcare access is kept constant. Second, it is also possible that retirement could improve health by reducing the opportunity cost of time for pursuing treatment for non-emergent conditions. Lastly, unlike mortality, the bulk of any healthcare costs incurred by retirement are placed on other people in the same insurance pools or on the taxpayers, thereby creating an externality. Due to these reasons, studying the relationship between retirement and healthcare utilization in the U.S. is important for determining optimal Social Security policy and for the health and retirement literature as a whole.

While there have been a few studies in European countries and China using quasi-experimental designs to study the effect of retirement on healthcare utilization, it is difficult to extrapolate their results to the U.S. context due to cultural and institutional differences (Frimmel and Pruckner, 2020; Nielson, 2019; Zhang et al., 2018; Zhou et al., 2021; Rose, 2020). Indeed, I am aware of only one other study that has attempted to study this reseach question in the U.S. Gorry et al., (2018) use an instrumental variable design based on year-of-age relative to various Social Security policy thresholds, along with data from the Health and Retirement Study, to estimate a decrease in hospitalizations following retirement. My study improves upon the limitations of prior work in several ways. First, I take advantage of the fact that a significant share of the population begin to receive Social Security payments and leave work immediately upon reaching the Early Eligibility Age (EEA) of 62 in a natural experiment to analyze the causal effects of retirement (Fitzpatrick and Moore, 2018). I do this using a regression discontinuity design (RDD) that examines age-related trends in healthcare outcomes for people on each side of the age 62 threshold. The principal advantage of the RDD, compared to alternative research designs, is its high internal validity and transparent identifying assumptions. To the best of my knowledge, mine is the first study to use this approach to study retirement and healthcare outcomes in the United States.

Second, I am able to overcome the small sample sizes inherent to survey data by using administrative data on healthcare encounters from New York and California between 2006-2017. These data, which contain observations on tens of thousands of healthcare encounters per month of age, allow me to examine relatively rare but costly healthcare episodes (i.e., ED visits, inpatient hospitalizations, and ambulatory surgery encounters) with statistical precision. Third, these administrative data are not suscptible to the self-report error inherent to survey data since they are based on hospital billing records. This reduces the possibility of measurement error contributing to misleading estimates. Fourth, ED visits and ambulatory surgery encounters have not previously been examined by prior work on health and retirement in the United States. The examination of ED visits is particularly important since the ED is used as a source of both acute and primary care. Therefore, the frequency of ED visits may change in response to retirement not only due to changes in health status, but also due to increases in free time to consume more discretionary forms of healthcare. Fifth, the administrative data contain detailed information on each healthcare encounter, such as primary diagnosis, that is unavailable in common survey data sets. I am able to use this information to determine whether any effects on ED visits are being driven by emergent versus nonemergent conditions.

I estimate that, across the entire U.S., at least 17% of people claim Social Security retirement benefits and approximately 8% of people retire immediately upon reaching the EEA of 62. I also show that turning 62 decreases the amount of time people spend working by around 2 hours per week on average. Furthermore, reaching the EEA is associated with a discontinuous increase of 1-2% in all-cause ED visits that do not result in hospitalization in New York and California. Estimates for discontinuities in elective inpatient admissions, admissions from the ED, and ambulatory surgery encounters are statistically insignificant. Furthermore, the increase in ED visits occurs across a wide variety of both emergent and nonemergent conditions. I am also able to show that changes in health insurance status and income at retirement do not play a significant role in the effect on ED visits, despite estimating a decrease in employer-sponsored private insurance coverage at age 62 of approximately 4 percentage points. Taken together, these results suggest that the increase in ED visits is potentially driven by both a decrease in health status and a decrease in the opportunity cost of receiving care upon retirement.

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.

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, 2021a). 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, 2021b). 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, 2021a). 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).

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, 2020a). 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 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 substances such as tobacco or alcohol (Ayyagari, 2016; Chuard-Keller, 2021), again decreasing health. Additionally, free time spent outside of work could be directed to more injury-prone activities such as home maintainence or driving. 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 in Section 6, I evaluate many of these possible channels to determine which are most likely to be driving the effects on ED visits.

4 Data and Empirical Strategy

4.1 Administrative Data on Healthcare Utilization

This study's measures of healthcare utilization are based on various sources of administrative data from California and New York. The California data were provided by the Department of Healthcare Acces and Information (HCAI)[1] and the New York data were provided by the Health-

[1]Formerly called the Office of Statewide Health Planning and Development (OSHPD).

care Cost an Utilization Project (HCUP)[2]. The primary outcome I study is emergency department (ED) visits without hospitalization that occurred between 2006-2017 in either state. ED visits are well as suited as an outcome in this study because people receive care in the ED for a wide variety of reasons. While the prototypical ED visit is for an emergent condition that could not be appropriately treated in an alternative setting, a large share of patients are seen in the ED for nonemergent conditions or for primary care treatable conditions (Johnston et al., 2017). Furthermore, as laid out in Section 3, retirement could affect healthcare utilization through several channels such as changes in health status and/or time use. Given the wide variety of conditions and reasons for which people receive care in the ED, and the range of plausible mechanisms through which Social Security eligibility could affect healthcare utilization, it is reasonable that the likelihood of visiting the ED might change discontinuously when reach the EEA. In addition to ED visits, I also examine changes in the rate of inpatient hospitalizations in both states between 2006-2017. Unlike most survey data sets, I am able to separate out inpatient hospitalization as elective or admitted via the ED. Lastly, I examine ambulatory surgery (AS) encounters between 2006-2017 from California only due to to data limitations.

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 4.4).

[2]The New York data are from HCUP's State Emergency Department Database (SEDD) and State Inpatient Database (SID). The data exclude stays in long-term care units of short-term hospitals, Federal hospitals, and free-standing psychiatric hospitals.

4.2 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 circumstances, 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 participation, and health insurance coverage. 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 in order to maximize statistical power. However, I also include robustness checks for certain outcomes in which I narrow analyses only to the census divisions that include NY and CA . This is the most granular level of geography that is present in the publicly-available HRS files.

4.3 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 also make use of natality data collected from historical vital statistics records in combination with population data from the 2010 census 10% sample to construct population counts by age cohort[3] (Ruggles et al., 2020). 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]See Appendix A for further discussion and construction of the population denominator.

4.4 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 rely upon 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 in the absence of treatment. 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 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. This research design can be viewed as a "fuzzy" RDD, rather than "sharp", for several reasons. First, not everyone who turns 62 is eligible to receive retirement benefits. In fact, the SSA estimated that in 2010, about 4% of people aged 62-84 that year would never go on to claim retirement benefits, principally due to a lack of qualifying work history (Whitman et al., 2011). Furthermore, eligible individuals are neither forced to claim benefits nor retire when they turn 62. Therefore, any causal effects on healthcare utilization may attributed to the "compliers", or those who claim benefits and/or retire upon their 62nd birthday.

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[4]. 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

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

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 [5]. 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. Main specifications use triangular kernels and standard errors are heteroskedasticity-robust. In specifications estimated using the Calonico, Cattaneo, and Titiunik (CCT) Mean-Squared Error (MSE) optimal bandwidth, I also display p-values derived from their robust bias-correct confidence intervals (Calonico, Cattaneo, 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.

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

A notable threat to identification of causal effects on healthcare outcomes using this research design is if people systematically migrate away from NY and/or CA when the turn 62. I can test this question directly by plotting in Appendix Figure B.2 the estimated population of people living in these states by quarter of age using the 10% sample of the 2010 Decennial Census. I also fit linear polynomials on either side of the cutoff in order to estimate any change in population. The RD estimate, as displayed in Figure B.2, is statistically insignificant which suggests that aggregate population does not change at the cutoff.

[5]Specifying the outcome in terms of rates instead of counts controls for sudden changes in population size across age bins. This is particularly important when disaggregating analysis by year.

5 Main Results

5.1 Effects on Social Security Uptake and Labor Market Outcomes

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 that could affect healthcare utilization, such as changes in health insurance or changes in income, would likely flow from at least one of these channels. Below, I use the HRS to estimate the share of people who claim Social Security and retirement upon turning 62 during my sample period. I also use information from the HRS to quantify the substitution in time use away from working. I 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, by age, who self-report receiving some type of Social Security income using the HRS full national sample. The figure displays a visible positive discontinuity in Social Security receipt at 62 as well as a point estimate of about 16.6 percentage points that is statistically significant at $p < .01$[6]. Figure 3 shows RD plots and point estimates of discontinuities in various measures of retirement and labor force participation at 62, also for the entire U.S. These measures include the average number of hours people report working per week in order to quantify the substitution in time use that occurs upon retirement. All four outcomes display discontinuous changes at 62 of large magnitudes. The estimates in Panels (a) - (c) display an 8.2 percentage point increase in share retired, a 6.3 percentage point increase in labor force non-participation, and a 7.6 percentage point decreaese in the share working for pay[7]. Panel D displays that people report working almost two hours less per week, on average, when they turn 62[8]. This discontinuity can be viewed as an "intent-to-treat" effect. since only about 8 percent of people retire discontinuously at this age. If one assumes that the entire effect on work-related time use

[6] The outcome is not limited to individuals receiving Social Security Retirement Insurance, and includes individuals receiving Disability and/or Survivors Insurance. This explains the non-zero levels of Social Security uptake prior to 62.

[7] 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.3. While both RD plots demonstrate discontinuities at 62, HRS results display a delayed effect, which suggests that it is likely due to sampling variation.

[8] Respondents who do not report working any job or spending time on any businesses are coded as working zero hours per week.

is driven by the effect on retirement, then a back-of-the-envelope calculation suggests that people who retire at 62 work approximately 23.9 hours less per week on average[9]. The effects on Social Security uptake and labor market outcomes have been fairly stable over time, as demonstrated in Appendix Figure B.4. 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.

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[10]. Second, I find that changes in labor market outcomes are concentrated among non-hispanic white people and, less robustly, hispanic people.

As I discuss in Section 4.2, the limited sample size of the HRS necessitates the use of the national sample to maximize statistial precision when estimating treatment effects. However, since the healthcare data I use in this study are from NY and CA, it is useful to provide additional evidence that the social security and labor market effects observed in the national sample are also present for these states specifically. Using the publicly-available HRS files, I can limit the sample to the Pacific and Middle Atlantic Census Divisions. According the US Census Bureau, people living in NY and CA comprise approximately 60% of the population aged 60-64 in these divisions (U.S. Census Bureau, Population Division, 2020)[11]. Appendix Figure B.6 shows the results for Social Security uptake and labor market outcomes in these two Census Divisions alone.The effects

[9] $\frac{\textit{Reduced Form}}{\textit{First Stage}} = \frac{-1.959}{0.08} = -24.49$

[10] Appendix Figure B.5 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.

[11] The Pacific Census Division consists of Alaska, California, Hawaii, Oregon, and Washington. The Middle Atlantic Census Division consists of New Jersey, New York, and Pennsylvania.

are even larger in magnitude and significance than those for the entire sample for each outcome, which suggests that these effects are likely to hold in NY and CA independently.

5.2 Effects on Healthcare Utilization

ED Visits

Figure 4 plots rates of total 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 5. 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 estimates are statistically significant at the 1% level, regardless of polynomial choice and inference method. Lastly, columns (5) and (6) estimate alternative models using uniforn kernels with CCT-selected bandwidths. These point estimates are of similar magnitudes and statistical significance to those in columns (3) and (4), although the bandwidth selection is notably smaller for the linear model in column (5). The bottom two panels seperately re-estimate these specifications for each state. Although the estimates for New York are generally larger and more statistically significant than those for California, the results indicate that the effect on ED visits is present in both states. These findings suggest that the discontinuity at 62 may be generalizable to other states as well since they do not appear to be driven's by a single state's policies.

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. Lastly, I examine how the effect on ED visits has changed over the time period of study in Figure 5. 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[12]. Figure 6 shows how these two types of admissions change through the age 62 cutoff. Neither panel (a) 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 7 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.

[12]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.

5.3 Robustness Checks

I now conduct two evaluations of the robustness of the estimated effect on ED visits. First, I systematically test the robustness of the aggregate ED estimates with triangular kernel to various alternate bandwidths by choice of polynomial and plot the results in Appendix Figure B.7. The coefficient estimates are significant at the 5% level for bandwidths of 4 through 12 months when using a linear fit. Since smaller 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 interval is relatively large due to the lack of observations.

Next, I present the results from re-estimating the effect on ED visits at a variety of placebo cutoffs. 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[13]. In total this results in 204 placebo cutoffs. Then, I plot the distribution of these RD coefficient estimates in panel (a) of Appendix Figure B.8. 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 this placebo test 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.8. 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.

6 Understanding the Healthcare Effects at 62

In Section 3, 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 labor

[13] 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).

force departure are the two primary behaviors that are likely to change at age 62. However, each of these behaviors has the potential to cause a variety of cascading effects that may ultimately affect utilization patterns. Understanding the relative contibutions of these mechanisms in important for policy making and is worth further analysis. In this section, I first provide some evidence about types of ED visits that change disontinuouslyat 62. 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.

6.1 Emergent versus Nonemergent Care

Stratifying the effect on ED visits by primary diagnosis can aid in evaluating which mechanisms might underpin the aggregate effect[14]. In Figure 8, I display the point estimates and 95% confidence intervals from re-estimating equation (1) by International Classification of Diseases (ICD) category. 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 8 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 other hand, the effects on the remaining categories of endocrine/metabolic disorder (thyroid problems, nutritional deficiencies) are positive and significant at the 5% level in aggregate.

Given the diffuse nature of the change in ED visits across types of conditions, it may aid interpretation to sort visits according to their level of urgency. In particular, visits for emergent conditions may be more likely to be driven by sudden changes in health status than visits for nonemergent conditions. One method for classifying visits is to use an algorithm developed by researchers at New York University (NYU) (Billings et al., 2000). These researchers used extensive internal records on ED visits from six hospitals in the Bronx, NY and categorized each visit into one of the following categories: (1) nonemergent; (2) emergent, primary care treatable; (3) emergent, ED care needed but preventable; or (4) emergent, ED care needed, not preventable. The authors then separated out visits due to injury, mental health, alcohol use, or substance use

[14] Primary diagnosis is defined as the principal reason why the patient showed up to the ED that day, as determined by the assigned physician

and placed them in their own designated, or "carved out", categories. Using these classifications, the researchers assigned to each primary diagnosis code (except for those carved out) a vector of four probability weights, one for each category, corresponding to the share of visits that fell into the given category[15]. All diagnosis codes that were not present among the researcher-categorized visits are designated as "uncategorized". In my analysis, I use an updated version of this NYU algorithm that supplemented the algorithm with ICD-9 and ICD-10 codes added since the original study (Johnston et al., 2017).

I estimate the effect of turning 62 on ED visits within each of these aforementioned categories and display the point estimates and 95% confidence intervals in Figure 10. Each outcome is constructed by summing the probability weight for each category within each age-state-year bin and the treating the sum as the "expected number of visits" for that category. These results indicate that the effect on ED visits at 62 is being driven by emergent conditions, including unavoidable ones, as well as non-emergent conditions. Estimates for each of the carved out categories, as well as for the remaining unclassifiable visits, are insignificantly different from zero. Further analysis also reveals the specific diagnosis categories driving the effects for each of these categories. The increase in emergent, unavoidable visits is mostly driven by diagnoses of signs, symptoms, and ill-defined conditions ($\beta = 0.042, p < 0.01$), though also somewhat by genitourinary issues ($\beta = 0.014, p < 0.05$) and digestive issues ($\beta = 0.010, p < 0.10$). Although there is no aggregate increase in avoidable emergent visits, this belies a substantial decrease in endocrine-related visits ($\beta = -0.035, p < 0.001$) that is counterbalanced by a statistically insignificant by large increase in circulatory and respiratory visits. The aggregate increase in primary care treatable emergent visits is driven by a wide variety of diagnoses, largest of which is an increase in respiratory issues ($\beta = 0.039, p < 0.001$) mostly related to COPD. COPD-related conditions are important drivers of mortality at 62 as well (Fitzpatrick and Moore, 2018). There is also a statistically significant decresease in endocrine-related visits in this category ($\beta = -0.008, p < 0.01$). Lastly, the increase in nonemergent visits is nearly entirely driven by musculoskeletal issues ($\beta = 0.046, p < 0.01$).

However, it is possible that these estimated effects on both emergent and nonemergent visits are just artefacts from the aggregation scheme of the probability weights. In Appendix Figure B.9,

[15] For example, if 50% of visits for ICD-9 code X were emergent, but primary care treatable and 50% of visits were nonemergent, then those two categories would receive weights of .5 in each of those two categories and the other categories would receive weights of 0.

I display two alternate approaches if assigning diagnoses to categories. The first approach involves assigning each non-carved out visit as either "emergent" or "nonemergent" depending on whether the sum of emergent categories' weights summed to more than .5 or whether the nonemergent weight is larger the .5. If both the sum of the emergent weights equal .5, then the visit is categories as ambiguous. The second method is similar to this, but classifies visits as "high" or "low" probability emergent versus non-emergent depending on if the weights summed between .75 and 1 or .5 and .75, respectively. Reassuringly, both methods produce similar results as the main method.

Taken together, the findings displayed in Figures 8 and 9 suggest that the change in ED visits is driven both by changes in health status and, potentially, the opportunity cost of time. Assuming the cost of care stays relatively constant after turning 62, an increase in unavoidable emergent visits where ED care is needed is a strong indicator that health status is discontinuously changing. On the other hand, an increase in nonemergent visits upon turning 62 suggests that something about receiving Social Security payments and/or retirement increases the propensity to seek out care for existing issues. Given the discontinuous increase in time spent not working at this age, a possible explanation for the effect on nonemergent visits is an increase in the opportunity cost of time spent receiving healthcare. A piece of evidence in favor of this explanation is the discontinuous drop in ED visits for diabetes at 62. Diabetes is a challenging condition to manage properly while at work and workers with diabetes often sustain higher-than-optimal glucose counts as a result (Ruston et al., 2013). It is entirely possible that an increase in time to administer self-care caused a decrease in ED visits for diabetes-related episodes, particularly since a large share of the drop is among emergent by preventable conditions. Thus, if retiring at 62 can increase people's time to take care of their diabetes at home, it is possible that they also have more time to pursue other forms of treatment outside of the house.

6.2 Health Insurance

The interpretations for the results discussed in Section 6.1 hinge upon there being no concurrent change in the financial cost of receiving care. 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)[16]. This decrease is compensated for by a concurrent increase in the share of people reporting other categories of insurance. Appendix Figure B.10 shows how these other categories change independently at 62. Additionally, Appendix Figure B.11 displays statistically 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. However, the relatively small sample size in the HRS may underpower these models to detect small but meaningful effects on health insurance status. In order to provide further evidence that loss of employer-sponsored private insurance cannot explain the effect, Appendix Figure B.12 displays RD plots for both the share and rate of privately insured ED visitors per 1,000 population. If the effect on ED visits is driven entirely by formerly insured individuals seeking care at the ED, one would expect to see a net decrease in privaly insured visits at 62. Instead, my findings sugges the opposite. 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 also possible that substantial chages in household income upon individuals reaching 62 could have affected their healthcare use. I evaluate this by plotting household income by year of age from the CPS-ASEC in Figure 11[17]. Given the coarse granularity of this variable it is difficult to draw firm conclusions, but the income level at age 62 does not appear to lie significantly outside

[16]This excludes specialty plans such as dental, vision, etc.

[17]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.

of the preceding trend.

7 Conclusion

This study presents evidence that claiming Social Security benefits and retiring at the Early Eligibility Age of 62 increases healthcare utilization for both emergent and nonemergent conditions. I am able to provide more comprehensive and credible evidence on this question than previous studies due to my use of administrative data on healthcare encounters and an internally valid regression discontinuity design. I also show that these healthcare effects cannot be entirely accounted for by concurrent losses of employer-sponsored health insurance. This paper's findings that both nonemergent and emergent healthcare utilization increase upon retirement suggests a complementary pair of policy recommendations. First, since early claiming of benefits and retirement appear to decrease health status and burden the Social Security trust fund, it may be beneficial to both population health and public finances to reduce policy incentives to do so. On the other hand, retirement also appears to increase episodes of nonemergent healthcare utilization that could protect the elderly's future health status. The available evidence suggests that this is most likely due to a reduction in the opportunity cost of time that occurs when people are no longer working. Therefore, encouraging public and firm policies to protect workers' ability to seek out primary care during the workday may serve to both increase long-run health and disincentivize early retirements.

Figures

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

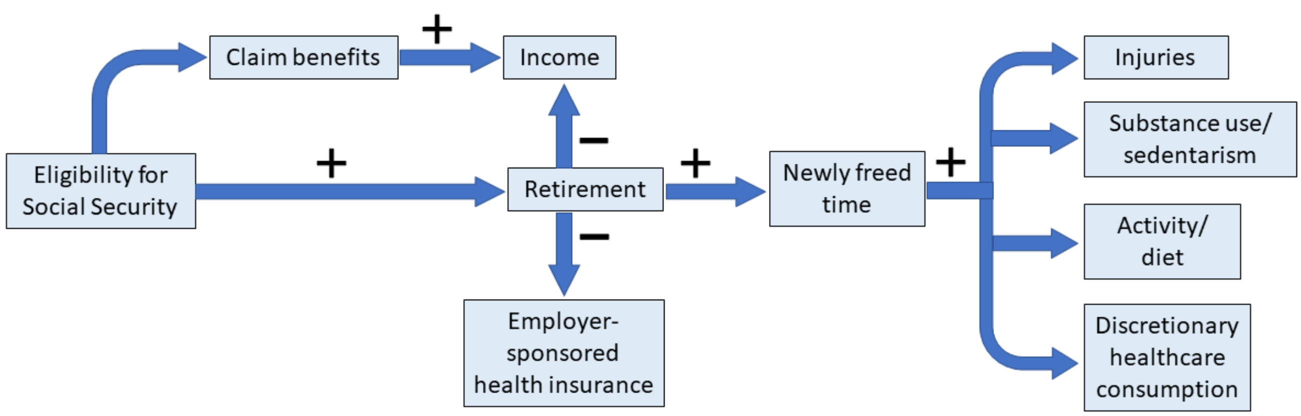

Figure 2: People Join Social Security Discontinuously at 62

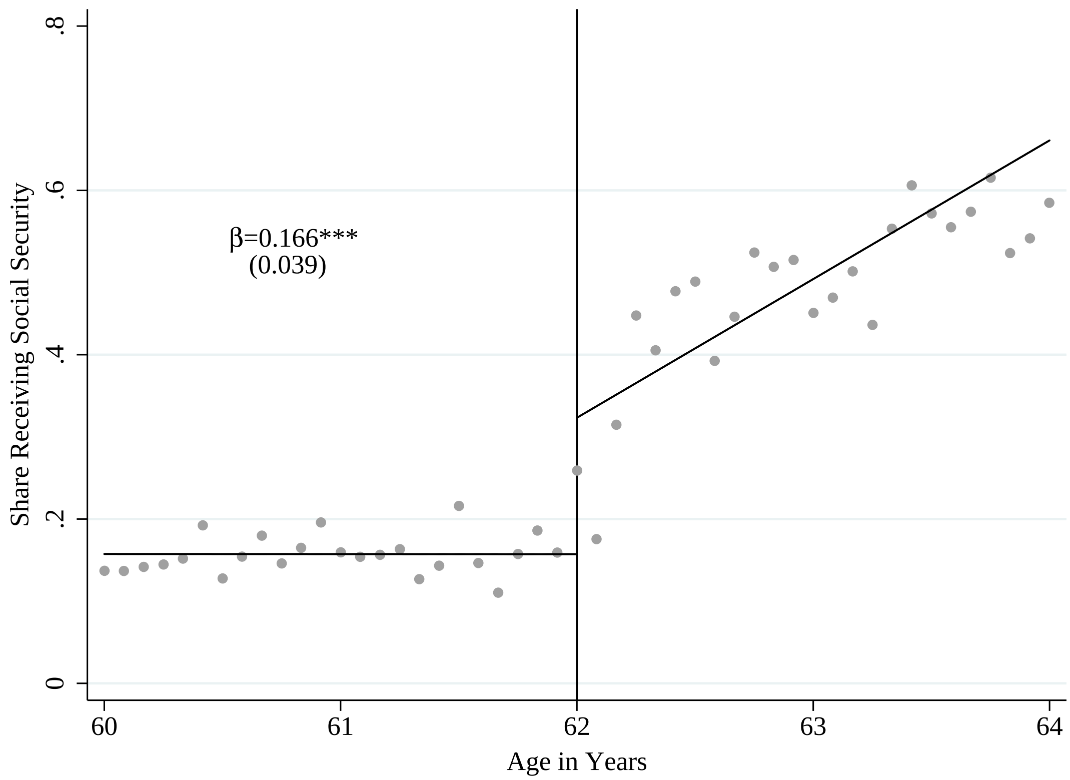

Note: This figure plots the share of the U.S. population currently receiving Social Security Retirement, Disability, or Survivor Insurance by age. Linear fit generated by estimating equation (2) with a 24-month bandwidth and triangular kernel. Data are from the Health and Retirement Study waves 2006-2018. 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

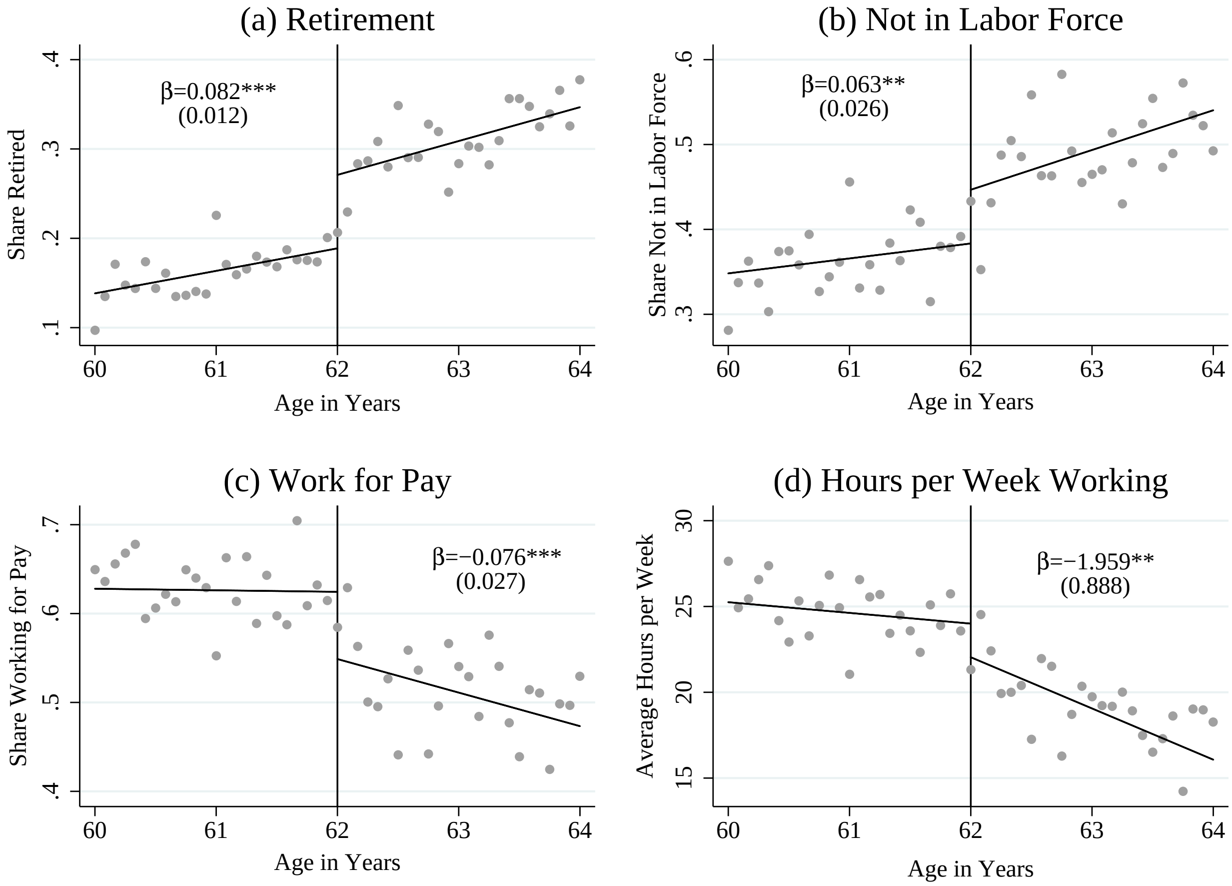

Note: This figure plots the share retired, not in the labor force, and working for pay, by age. It also plots the average number of hours people report working by age. Linear fits are generated by estimating equation (2) with 24-month bandwidths using triangular kernels. Data are from the Health and 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

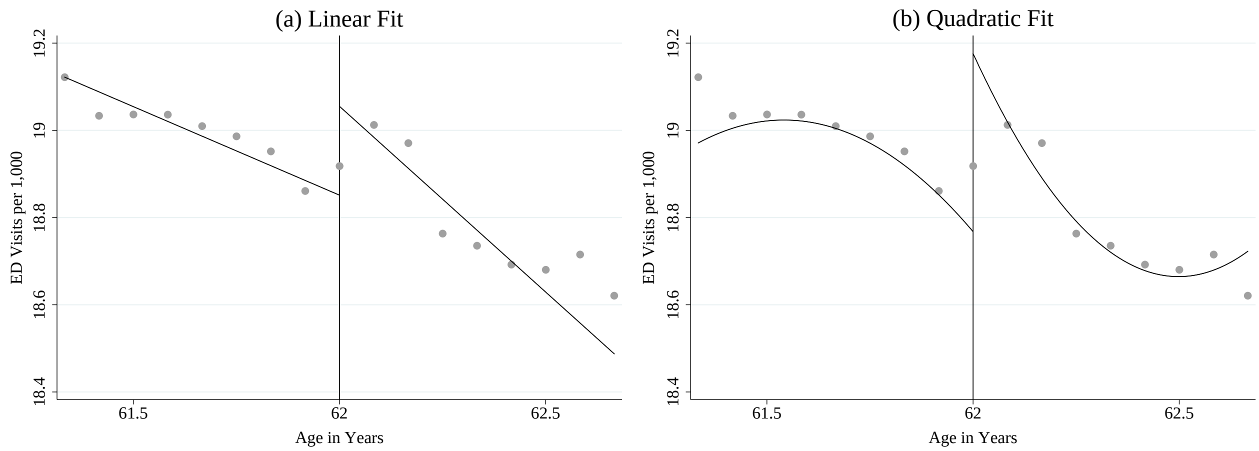

Note: This figure plots ED visits per 1,000 population by age. Data are from HCUP NY SEDD and CA OSHPD ED visit data, years 2006-2017. Each age bin corresponds to the calendar month that people turned each age (e.g., ED visits in the calendar month people turn 62). Polynomial fits generated from estimating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

Figure 5: The Effect on ED Visits Remains Stable Over the Time Period

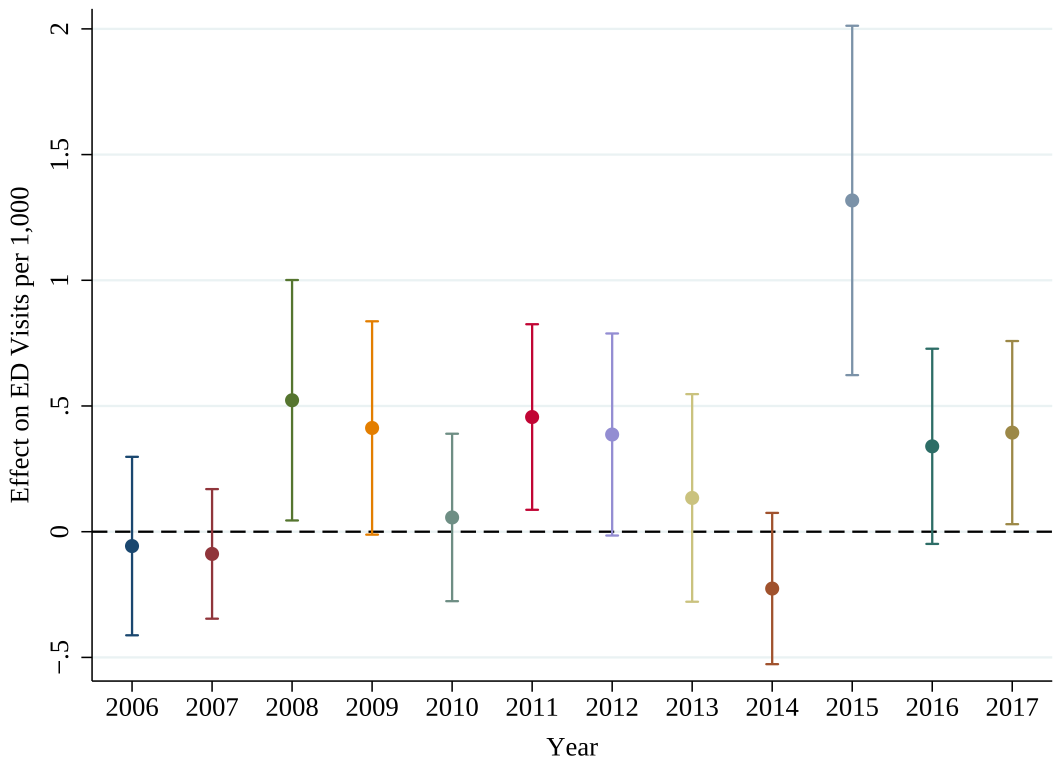

Note: This figure plots the RD effect on ED visits per 1,000 population by year. Data are from HCUP NY SEDD and CA OSHPD ED visit data, years 2006-2017. Each estimate is made by estimating equation (1) using a linear fit, triangular kernel, and CCT optimal bandwidth. 95% confidence intervals generated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

Figure 6: Aggregate Inpatient Hospitalizations Do Not Change at 62

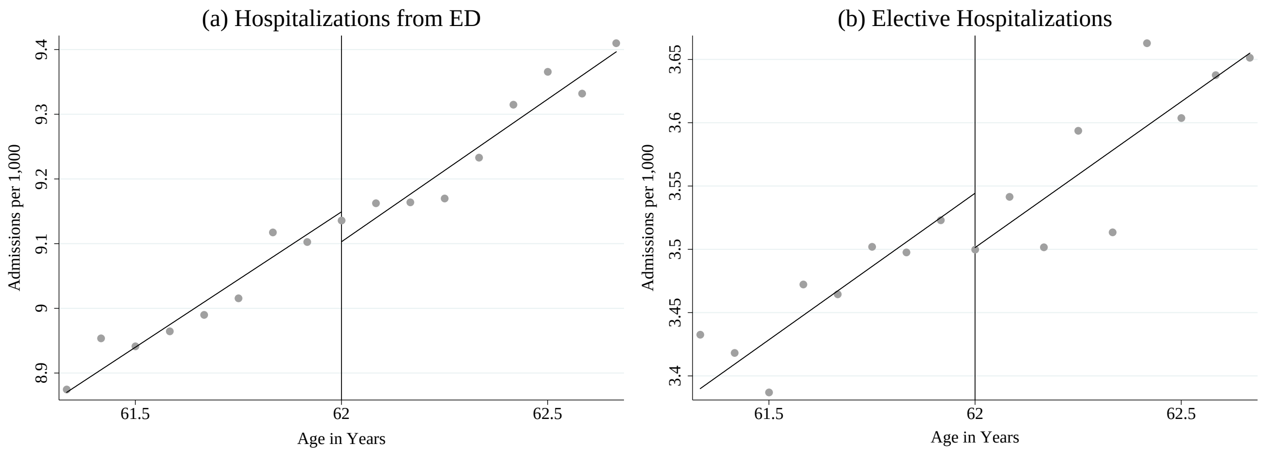

Note: This figure plots inpatient admissions per 1,000 population by age. Data are from HCUP NY SID and CA OSHPD PDD data, years 2006-2017. Each age bin corresponds to the calendar month that people turned each age (e.g., admissions in the calendar month people turn 62). Linear fit generated from estimating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

Figure 7: Ambulatory Surgery Encounters in CA Do Not Change at 62

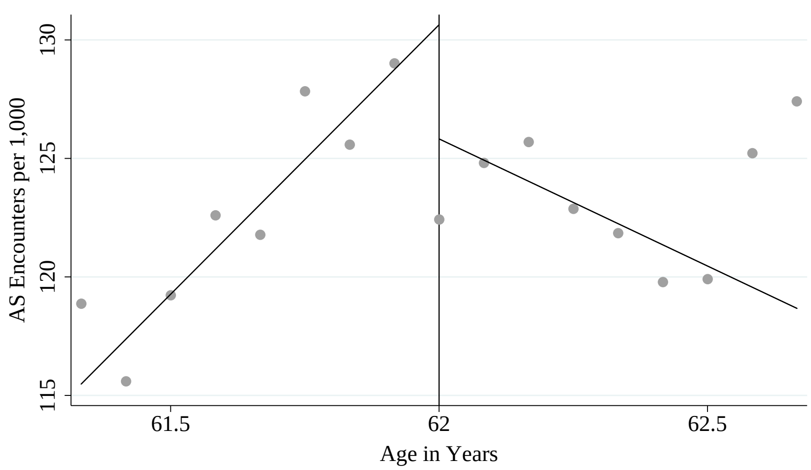

Note: This figure plots ambulatory surgery ancounters per 1,000 population by age. Data are from CA OSHPD AS data, years 2006-2017. Each age bin corresponds to the calendar month that people turned each age (e.g., admissions in the calendar month people turn 62). Linear fit generated from estimating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

Figure 8: Effect on ED Visits Not Driven by Any Particular Diagnoses

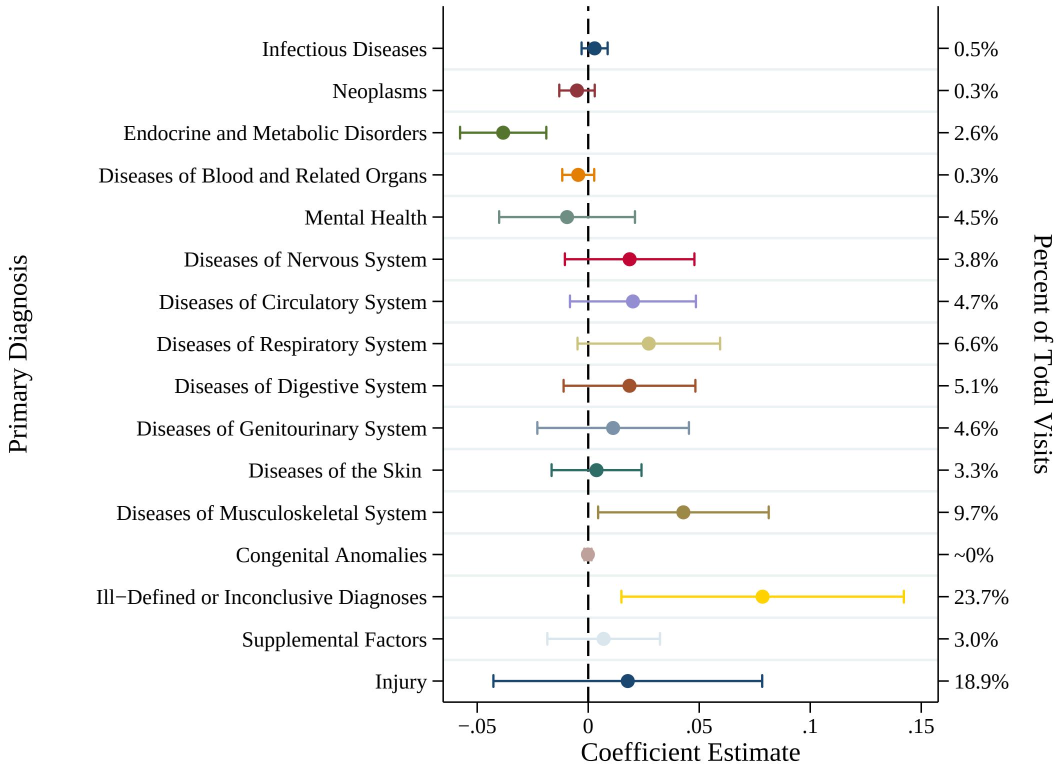

Note: This figure plots the RD effect on ED visits per 1,000 population by primary diagnosis. Data are from HCUP NY SEDD and CA OSHPD ED visit data, years 2006-2017. Each estimate is made by estimating equation (1) using a linear fit, triangular kernel, and CCT optimal bandwidth. Percent of total visits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

Figure 9: Effect on ED Visits Drive by Emergent and Nonemergent Conditions

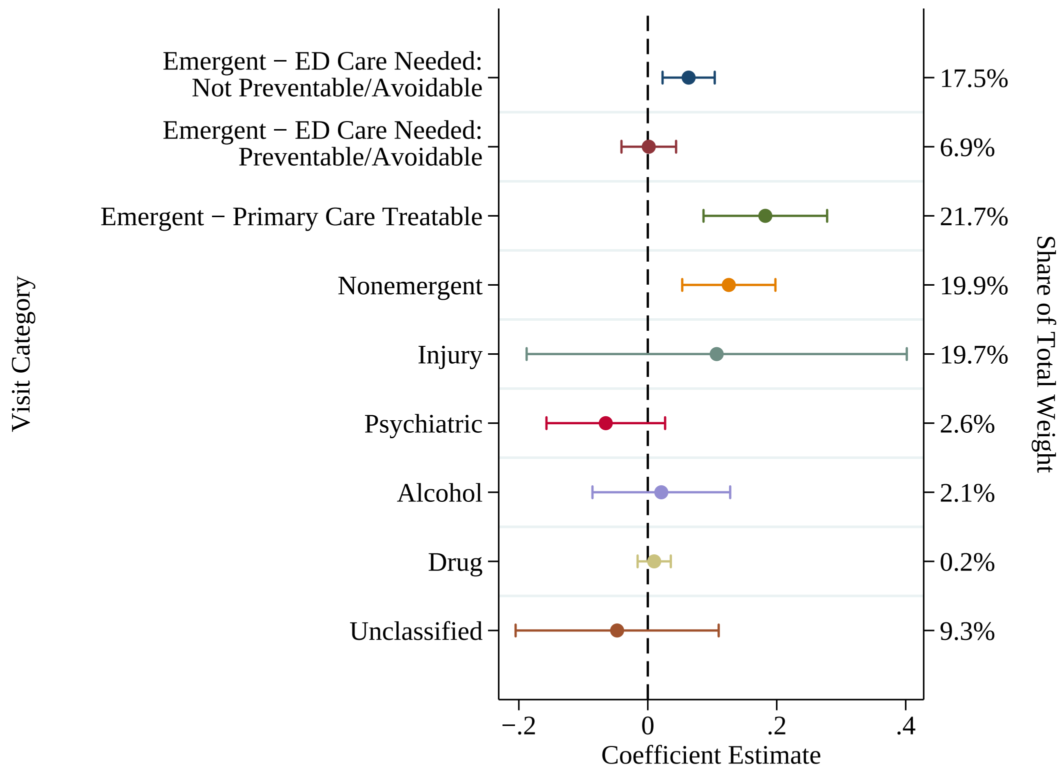

Note: This figure plots the RD effect on ED visits per 1,000 population according to NYU algorithm classification. The outcome for each category of visit is the sum of the probability weight across all visits in each age-state-year bin. Data are from HCUP NY SEDD and CA OSHPD ED visit data, years 2006-2017. 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 10: Employer-Sponsored Insurance Decreases at 62, but Overall Rates Stay the Same

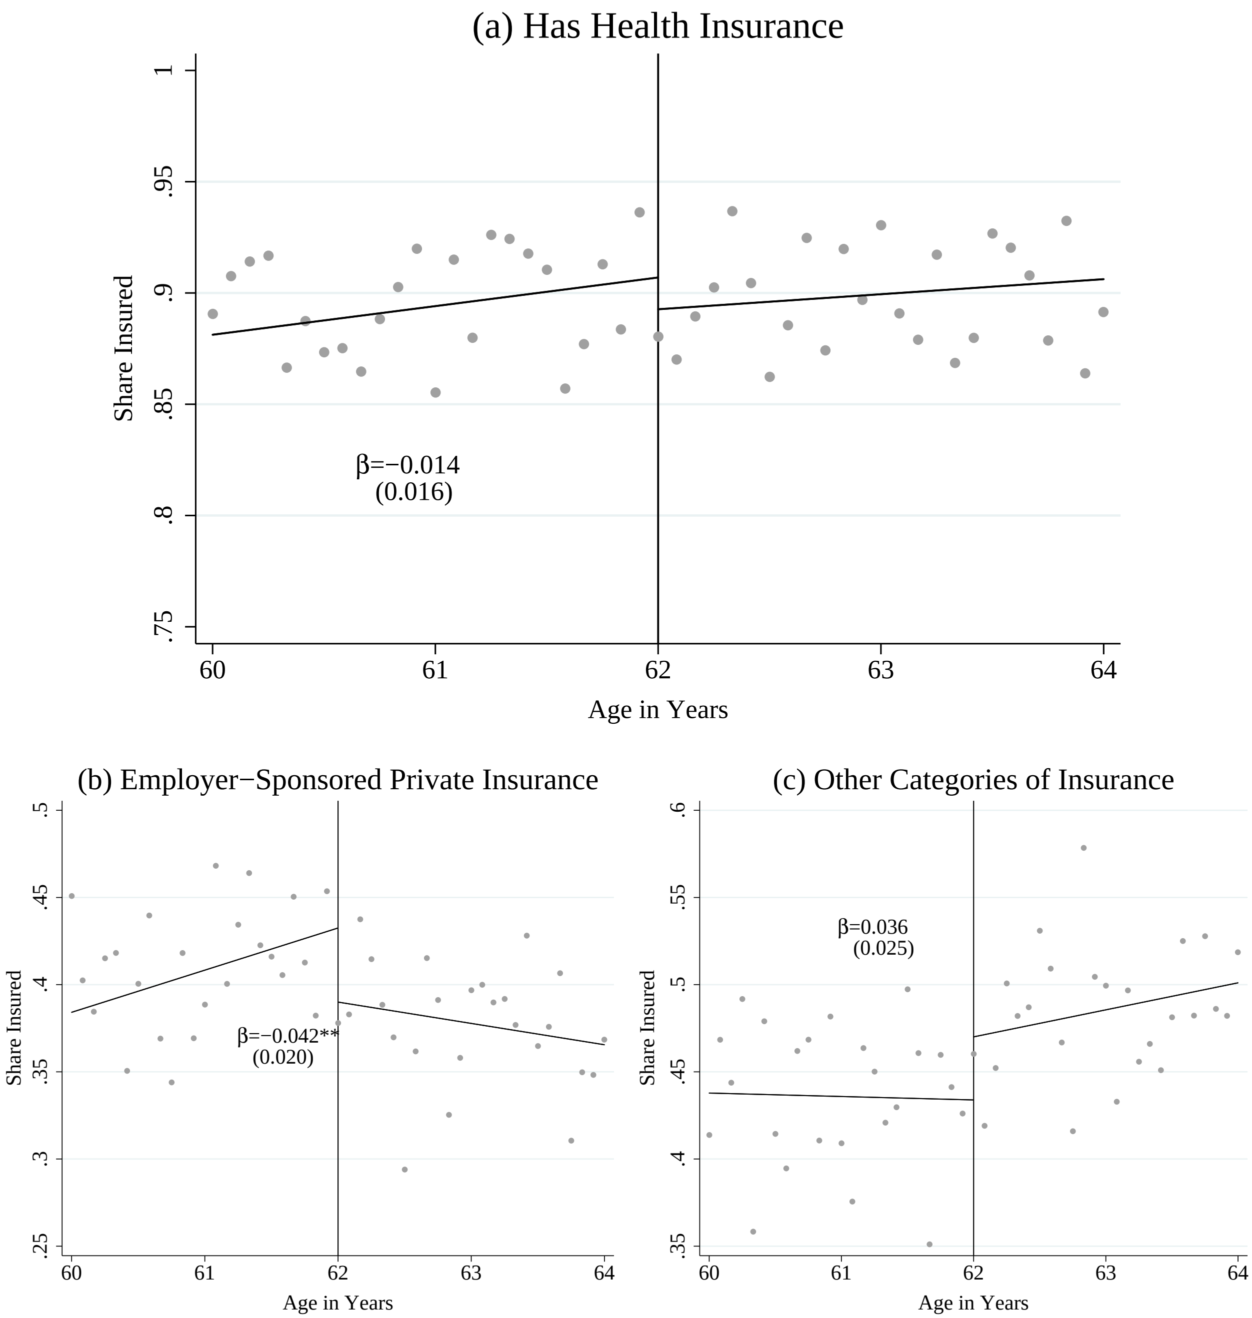

Note: Panel (a) plots the share of share of people with at least one health insurance plan by age. Panel (b) plots the share of people with an employer-sponsored private insurance plan. Panel (c) plots the share of people with at least one form of another type of insurance. Data are from the Health and 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

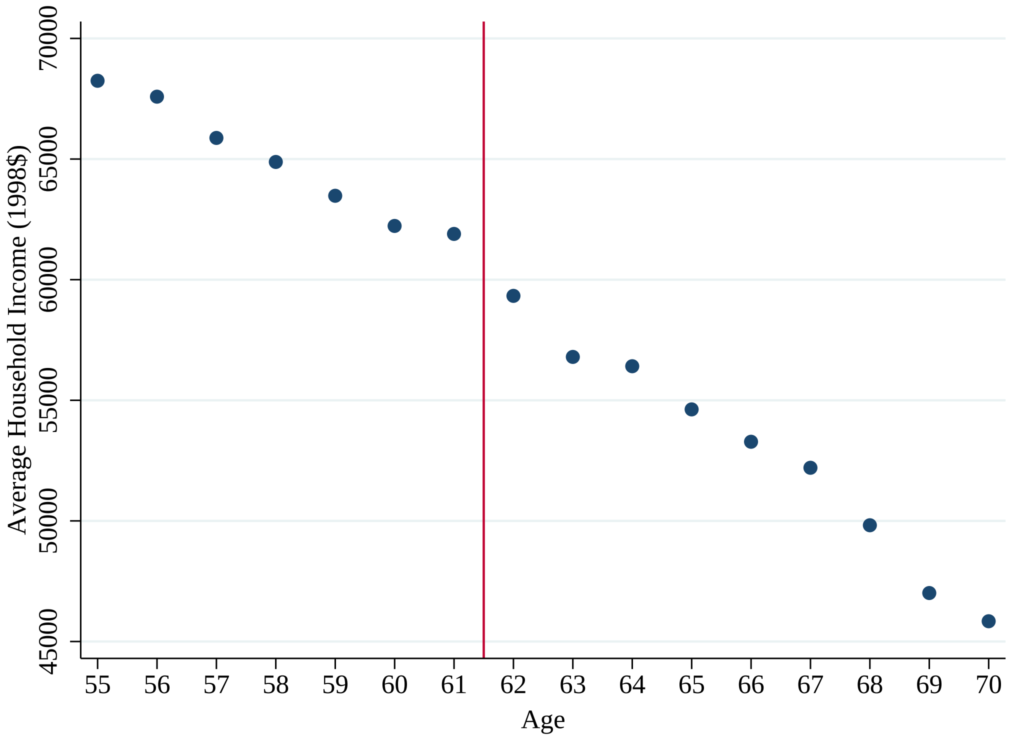

Note: This figure plots average household income by year of age from the CPS-ASEC years 2006-2017.

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)	(5)	(6)
Sample: Both States						
Aged 62+	0.204***	0.409***	0.292***	0.397***	0.357***	0.315***
	(0.064)	(0.108)	(0.074)	(0.105)	(0.096)	(0.109)
			[0.001]	[0.001]	[0.000]	[0.005]
BW	8	8	5.5	8.1	3.9	8.3
Sample: New York Only						
Aged 62+	0.234**	0.550***	0.346***	0.548***	0.400***	0.569***
	(0.096)	(0.162)	(0.107)	(0.162)	(0.129)	(0.176)
			[0.001]	[0.002]	[0.002]	[0.001]
BW	8	8	6.1	7.7	4.5	7.8
Sample: California Only						
Aged 62+	0.174**	0.267*	0.213**	0.232*	0.208**	0.216
	(0.086)	(0.147)	(0.102)	(0.130)	(0.101)	(0.135)
			[0.121]	[0.157]	[0.082]	[0.187]
BW	8	8	5.6	9.8	5.6	9.9
Poly. Deg.	1	2	1	2	1	2
CCT?			X	X	X	X
Kernel	Triangular	Triangular	Triangular	Triangular	Uniform	Uniform

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. The odd-numbered columns estimate the model using a linear fit and the even-numbered columns estimate the model with a quadratic fit. Columns (1)-(2) use a set bandwidth of 8 months while columns (3)-(6) use CCT optimal bandwidths. Columns (1)-(4) use triangular kernels and columns (5)-(6) use uniform kernels. 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

	(1)	(2)	(3)	(4)
Outcome:	**ED Visits**	**Inpatient Admissions - ED**	**Inpatient Admissions - Elective**	**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 insufficient 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 and quarter-of-birth[18] 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

[18] 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

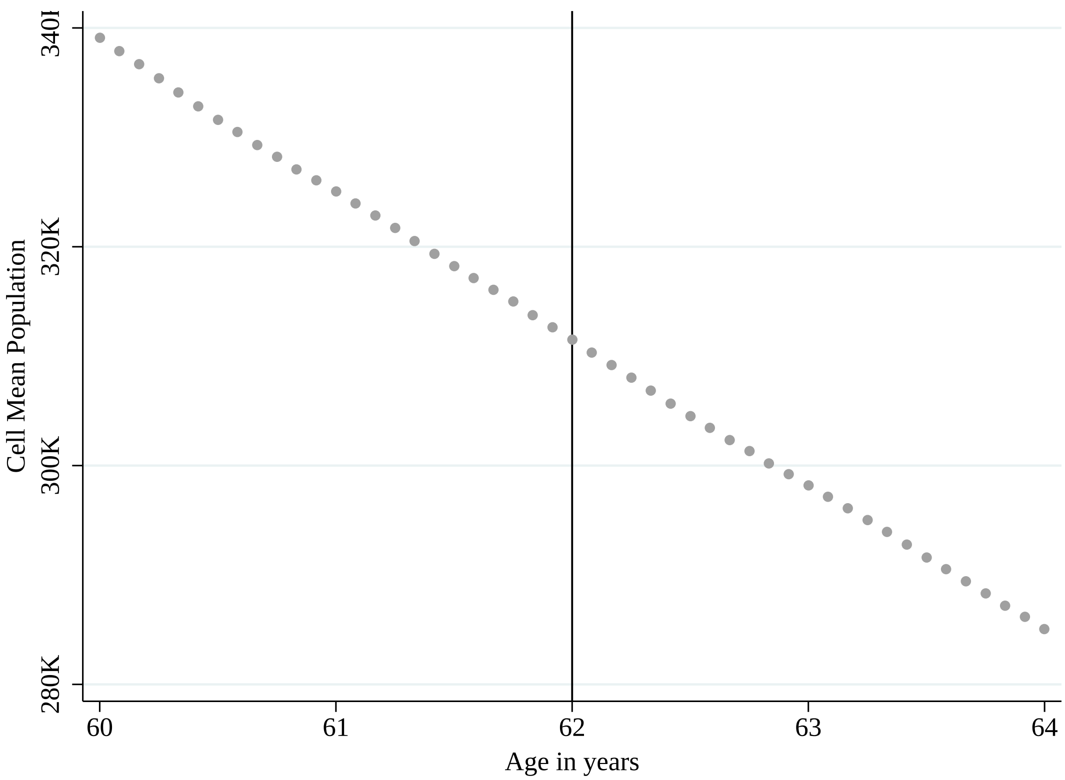

Note: This figure plots the estimated population counts by month of age. Estimates are derived from from 1941-1958 U.S. vital statistics data combined with the publicly available 10% extract of the 2010 U.S. Census.

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

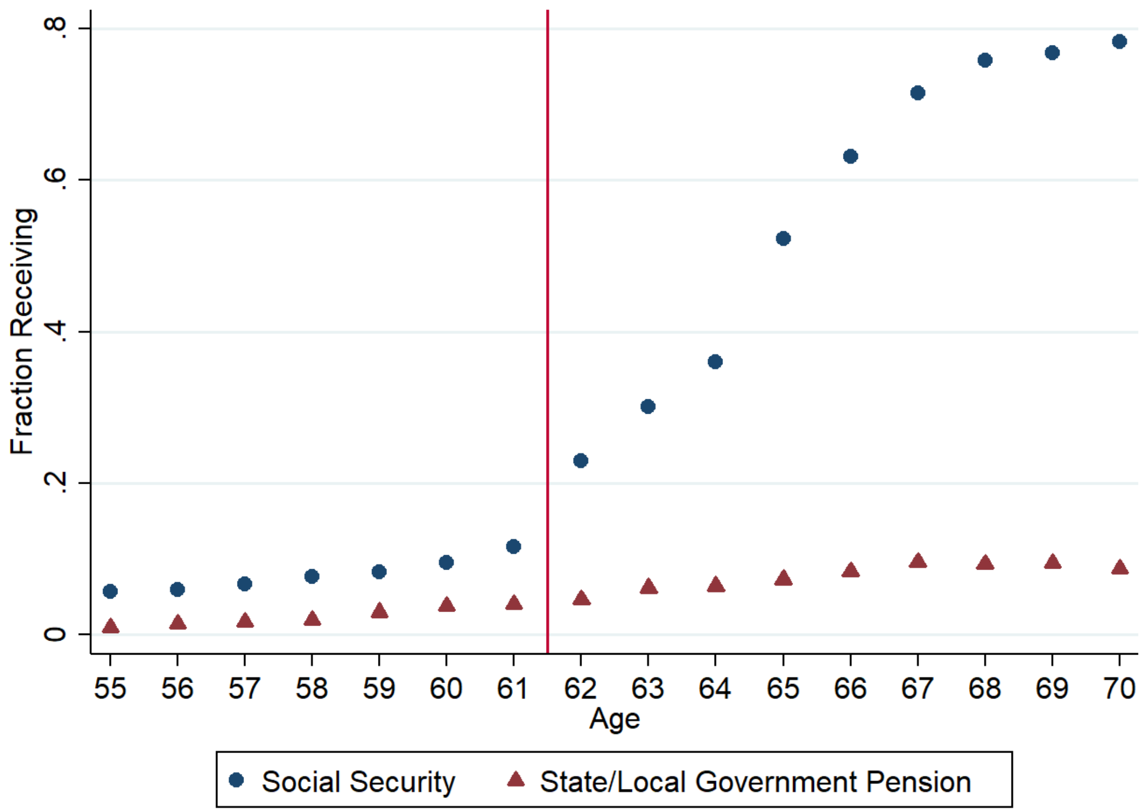

Note: This figure plots rates of Social Security and State/Local Government Pension receipt in NY and CA by age in 2006-2017. Data are taken from the CPS-ASEC 2007-2018 waves, which asks about people's sources of retirement income from the previous year. Since the CPS-ASEC is conducted every March and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

Figure B.2: The Aggregate Population in NY and CA Does Not Change Discontinuously at 62

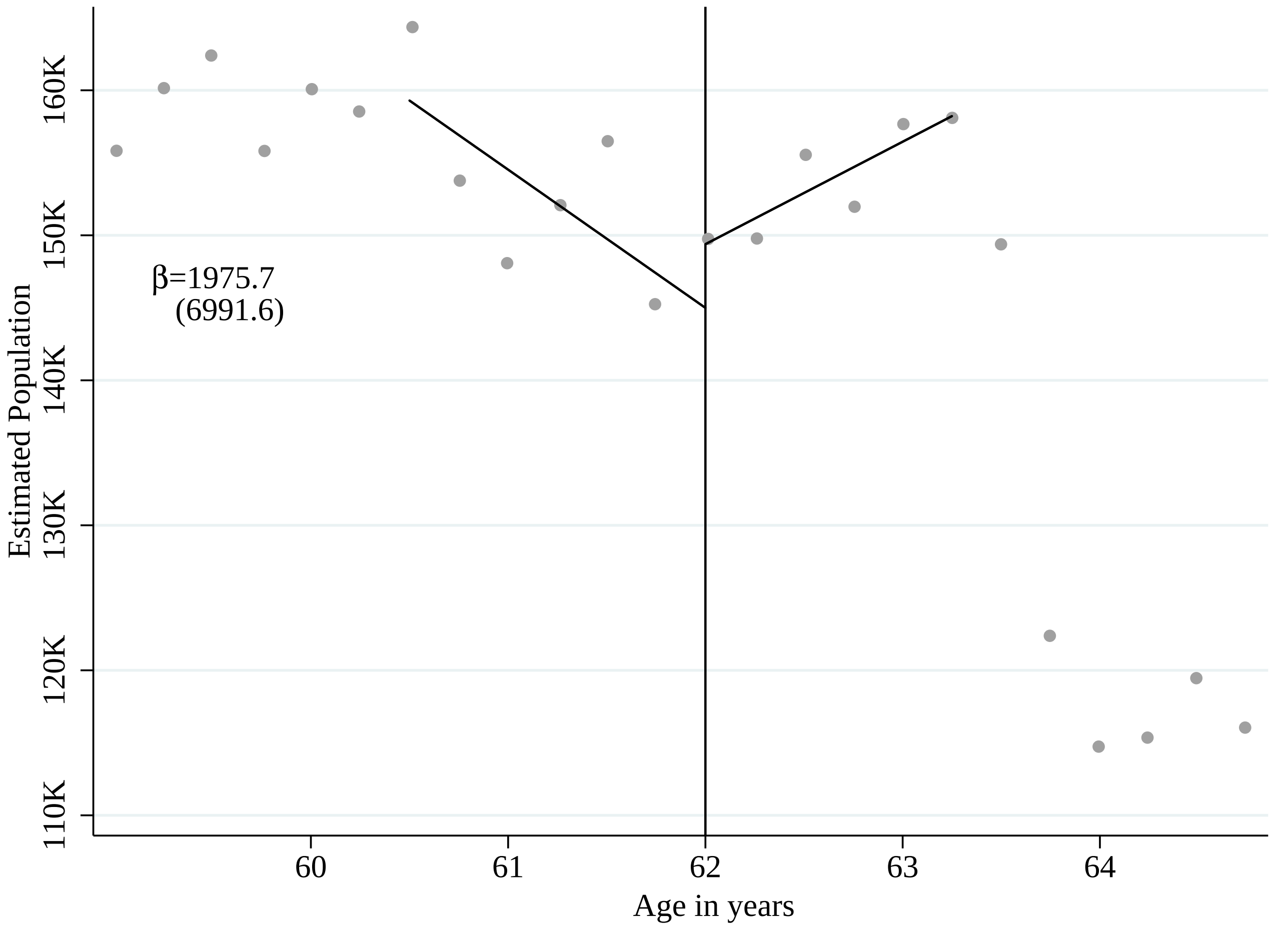

Note: This figure plots the estimated population size in NY and CA by quarter of age using the 10% sample of the 2010 Decennial Census. Population counts for each age bin are multiplied by 10 to account for the 10% random sample. Linear fits are estimated using triangular kernels and 6-quarter (18-month) bandwidths.

Figure B.3: The Discontinuity in Labor Force Participation Occurs Immediately at 62

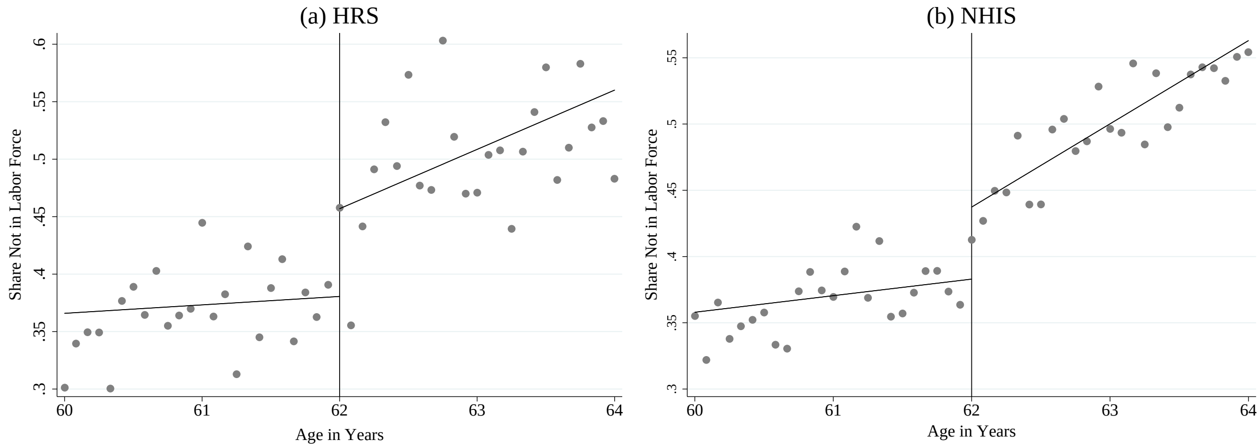

Note: This figure plots the share of the population currently not in the labor force. Linear fit obtained form estimating equation (2) with a 24-month bandwidth and triangular kernel. Data are from the Health and Retirement Study waves 2006-2014 and National Health Interview Survey waves 2006-2014. Each 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.4: Effect on Social Security and Labor Market Outcomes by Year

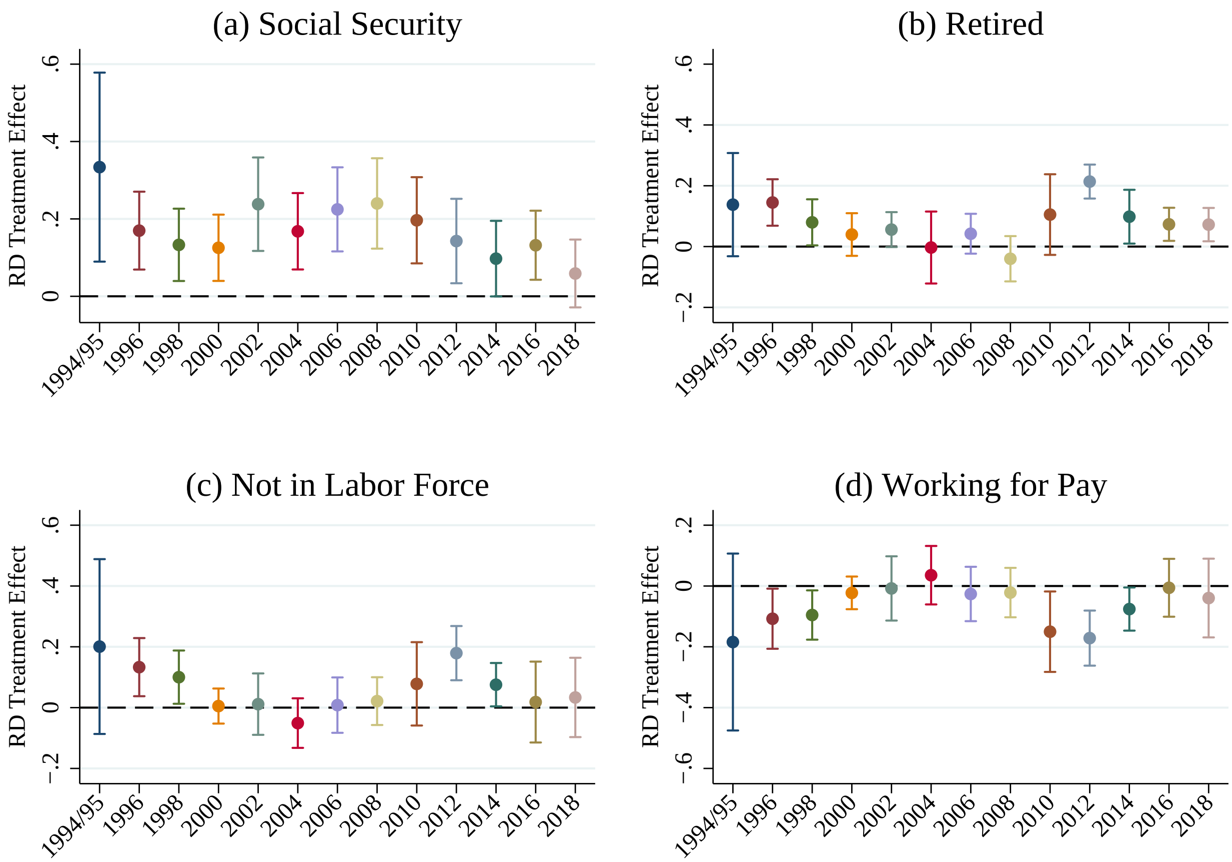

Note: This figure plots the effect of turning 62 on Social Security receipt and labor market outcomes by year by estimating β using equation (2). Data are from the Health and Retirement Study waves 1994/95-2018. Each specification controls for a dummy for the month people turn 62. 95% confidence intervals are derived from robust standard errors.

Figure B.5: The Discontinuity in Female Labor Force Participation is Also Present in the NHIS

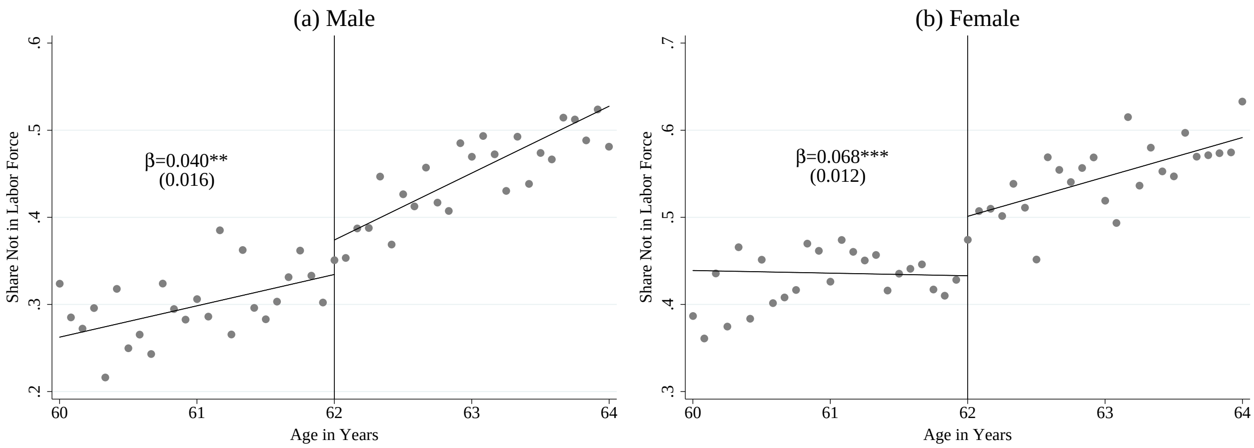

Note: This figure plots the share of the population currently not in the labor force. Linear fit obtained from estimating equation (2) with a 24-month bandwidth and triangular kernel. Data are from the National Health Interview Survey waves 2006-2014. Each age bin corresponds to the calendar month that people turned each age (e.g., share working in the calendar month people turn 62). Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Figure B.6: The Effects on Social Security Uptake and Labor Market Outcomes: Pacific and Middle Atlantic Census Divisions

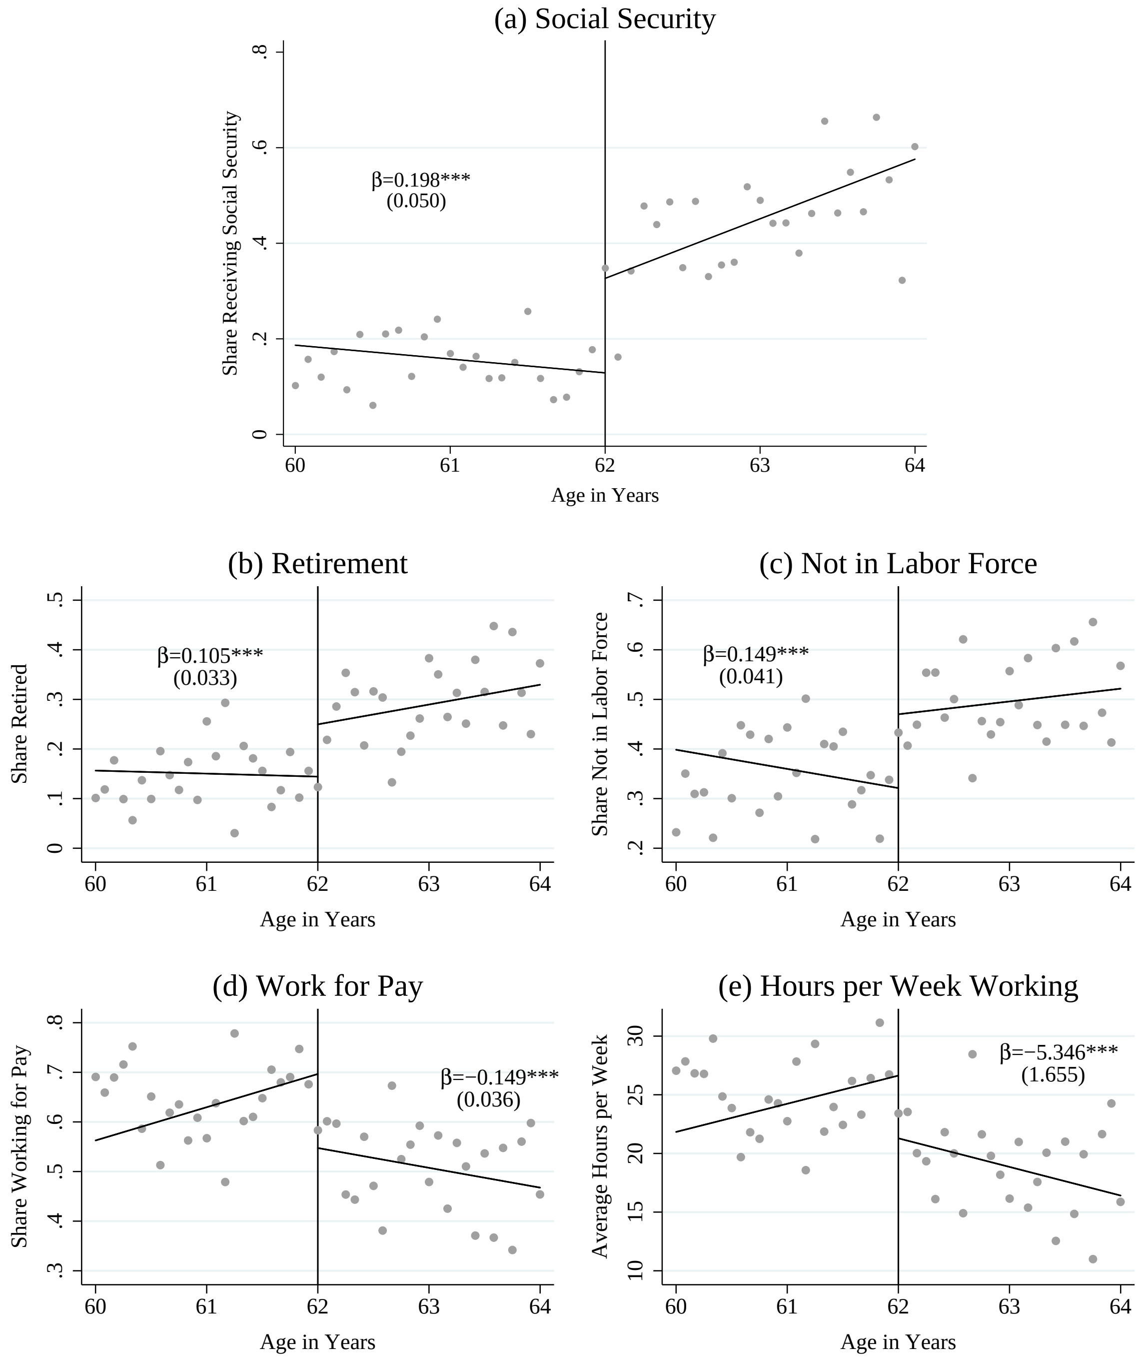

Note: These figures plot the discontinuity in Social Security receipt and labor outcomes for people in the Pacific and Middle Atlantic Census Divisions. Linear fit obtained from estimating equation (2) with a 24-month bandwidth and triangular kernel. Data are from the Health and Retirements Study waves 2006-2018. Each age bin corresponds to the calendar month that people turned each age (e.g., share working in the calendar month people turn 62). Parentheses contain robust standard errors where * p <.1, ** p<.05, *** p<.01.

Figure B.7: Estimated Effect on ED Visits is Robust to Many Bandwidths, Particularly Narrow Ones

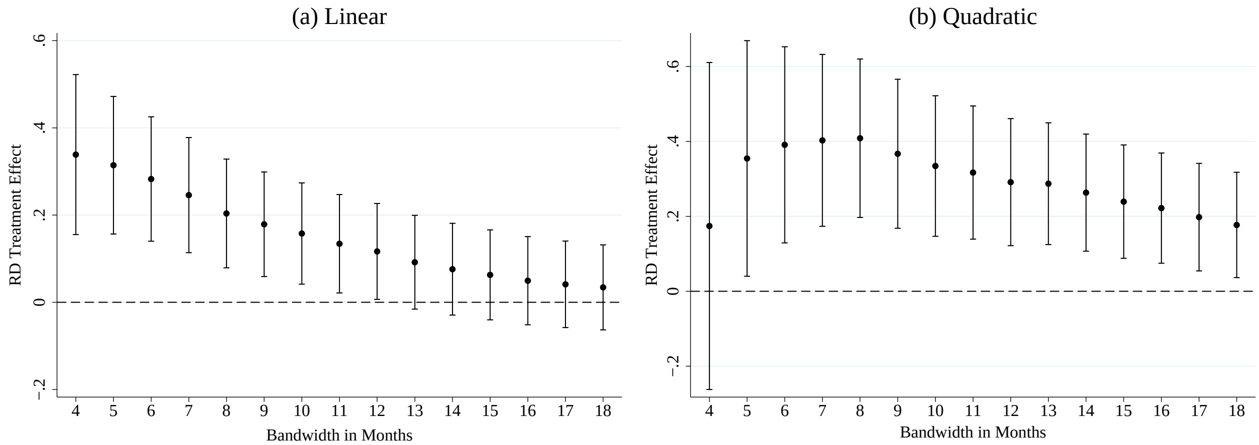

Note: This figure plots RD estimates using equation (1) by bandwidth in months and polynomial fit. Data come from NY HCUP SEDD and CA OSHPD ED for years 2006-2017. Outcomes are ED visits per 1,000 population. All regressions estimated with triangular kernels. 95% confidence intervals are calculated using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

Figure B.8: The Effect on ED Visits and Robust t-Statistic are Larger in Magnitude than 95% of Placebo Estimates

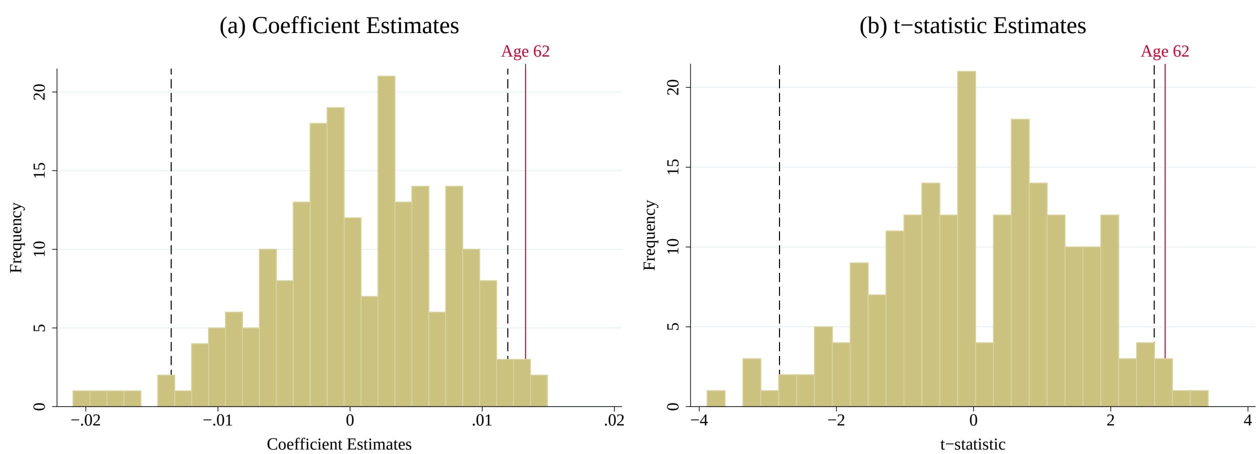

Note: This figure plots the distribution of RD coefficients and robust t-statistics, estimated by equation (1), at the age 62 cutoff and a set of placebo age cutoffs within a +/- 10-year radius. Data come from NY HCUP SEDD and CA OSHPD ED data for years 2006-2017. The outcomes are the natural log of 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.9: Effect on Emergent and Nonemergent Visits Is Robust To Alternative Classification Methods

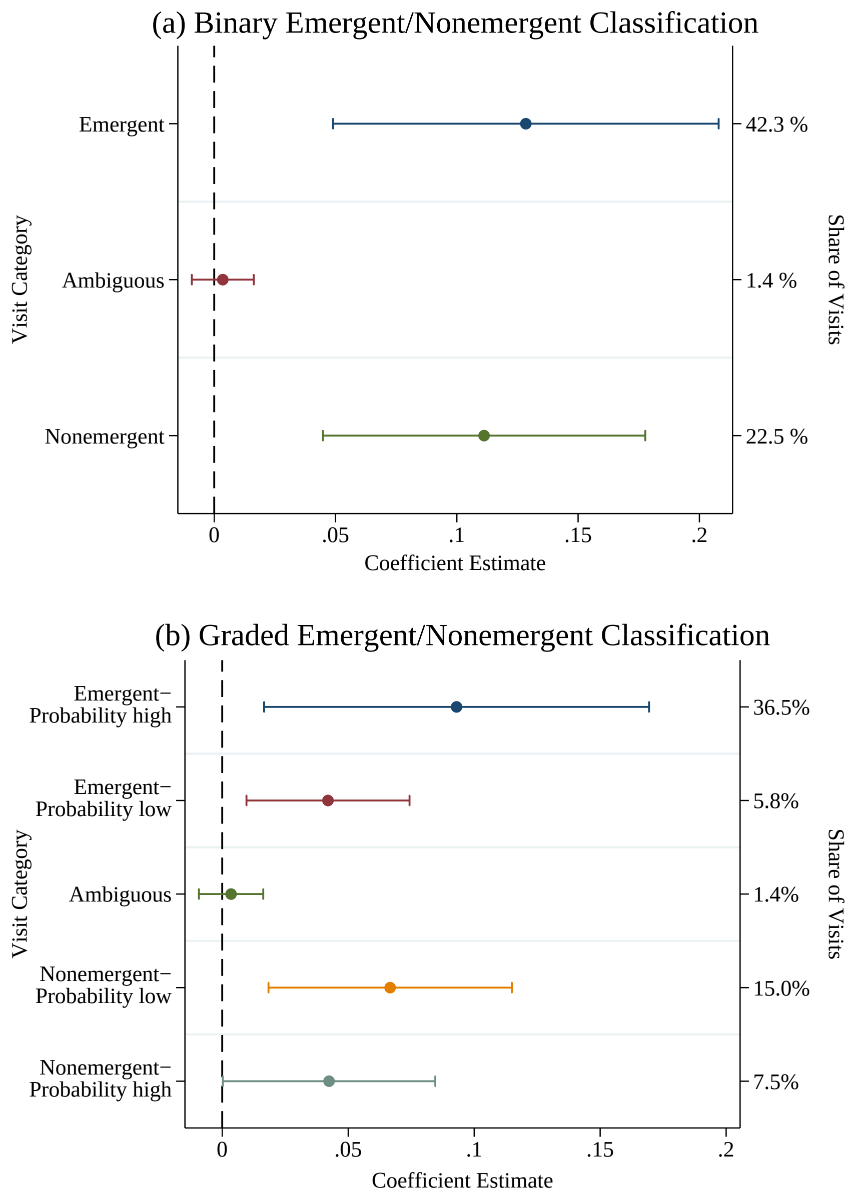

Note: This figure plots the RD effect on ED visits per 1,000 population as classified by the NYU algorithm via Johnston et al. (2017). Panel (a) classifies visits as emergent if the probability weights in each emergent category sum to greater than .5 and classifies visits as nonemergent if the same holds true for the nonemergent probability weight. Panel (b) does the same, but adds gradation of low and high probabilities for each category depending on whether the probabilitity weights sum between .5 to .75 or .75 to 1. Cases in which the probability weights for emergent categories and nonemergent categories are both .5 are classified as ambiguous. 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 B.10: Effect on Insurance Coverage by Type of Payer

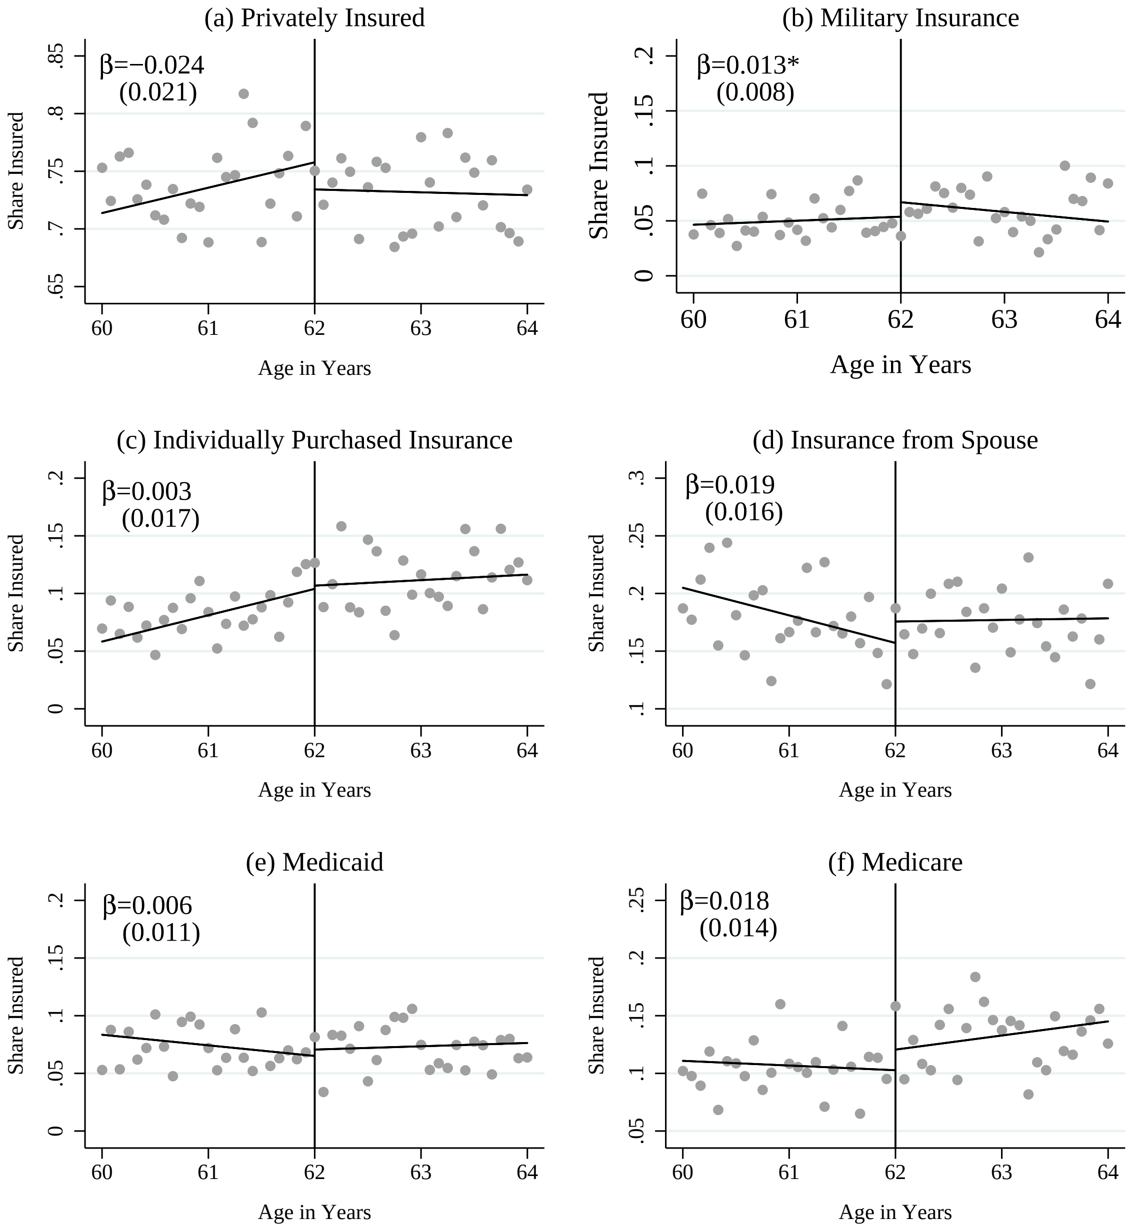

Note: This figure plots the rates of of insurance coverage by age and type of payer. Data are from the Health and Retirement Study waves 2006-2018. Linear fits are generated by estimating equation (2) with 24-month bandwidths and triangular kernel. Each age bin corresponds to the calendar month that 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.11: Effect on Number of Insurance Policies

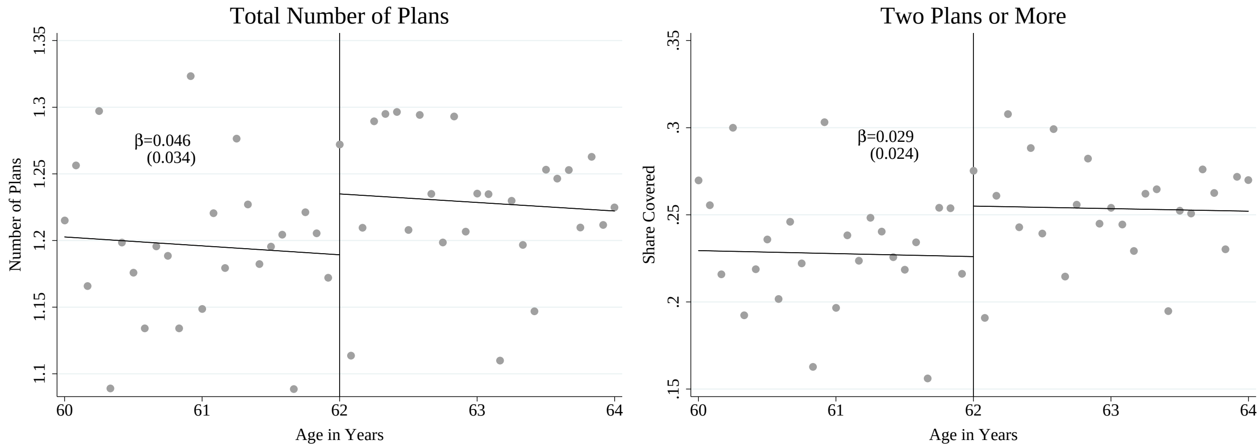

Note: This figure plots the rates of of insurance coverage by age. Data are from the Health and Retirement Study waves 2006-2018. Linear fits are generated by estimating equation (2) with 24-month bandwidths and triangular kernel. Each age bin corresponds to the calendar month that people turned each 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.12: Effect on Share and Rate of ED Visits Among Privately Insured

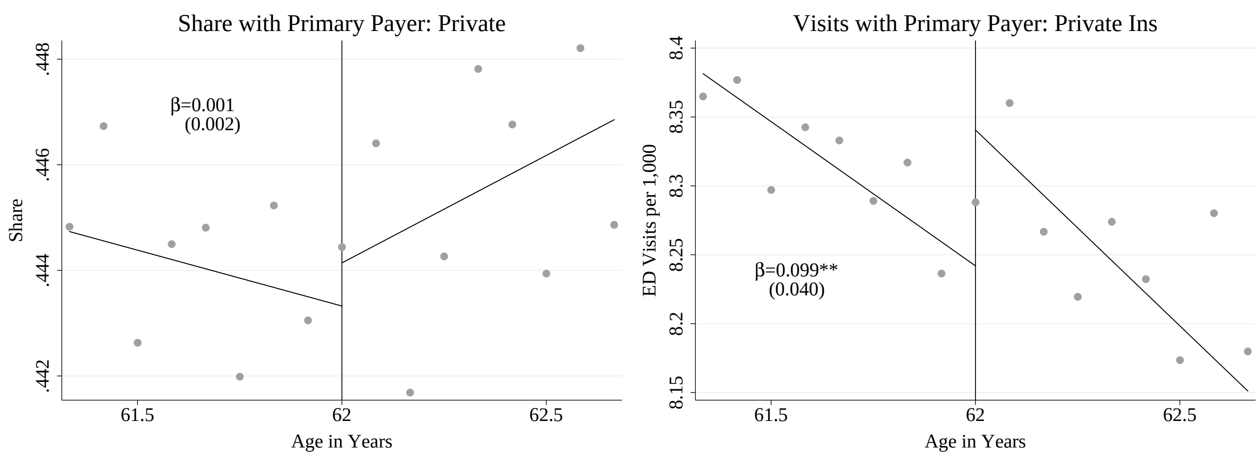

Note: This figure plots the estimated discontinuities in the share and rate per 1,000 population of privately insured patients. Data come from NY HCUP SEDD and CA OSHPD ED data for years 2006-2017. Linear fits are generated by estimating equation (1) with 8-month bandwidths and triangular kernel. Each 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.

References

[1] Arenberg, S., Neller, S., & Stripling, S. (2020). The Impct of Youth Medicaid Eligibility on Adult Incarceration. Working Paper.

[2] Ayyagari, P. (2016). The impact of retirement on smoking behavior. *Eastern Economic Journal 43*, 270-287.

[3] 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"

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

[5] Billings, J., N, Parikh, and T. Mijanovich. 2000. "Emergency Department Use in New York City: A Substitute for Primary Care?" *Issue Brief Commonwealth Fund.*

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

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

[8] Charles, K. K. (2004). Is Retirement Depressing?: Labor Force Inactivity and Psychological Well Being in Later Life. *Research in Labor Economics 23*, 269-299.

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

[10] 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

[11] 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.

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

[13] 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.

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

[15] 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.

[16] 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).

[17] 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.

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

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

[20] 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

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

[22] 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.

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

[24] 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.

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

[26] Johnston, K. J., Allen, L., Melanson, T. A., & Pitts, S. R. (2017). A “Patch” to the NYU Emergency Department Visit Algorithm,. *Health Services Research, 52 (4)*, pp. 1264-1276. doi: 10.1111/1475-6773.12638.

[27] 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.

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

[29] 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.

[30] 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.

[31] McArdle, F., Neuman, T., & Huang, J. (2014). Retiree health benefits at the crossroads. Menlo Park, CA: Kaiser Family Foundation.

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

[33] 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.

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

[35] 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.

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

[37] 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

[38] 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

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

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

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

[42] 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

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

[44] 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.

[45] Ruston, A., Smith, A., & Fernando, B. (2013). Diabetes in the Workplace - Diabetic's Perceptions and Experiences of Managing Their Disease at Work: A Qualitative Study. *BMC Public Health 13, 386.*

[46] 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.

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

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

[49] Social Security Administration (SSA). (2021b). Annual Statistical Supplement to the Social Security Bulletin, 2021. Retrieved from: https://www.ssa.gov/policy/docs/statcomps/supplement/2021/supplement21.pdf

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

[51] United States Census Bureau, Population Division. (2020). Annual Estimates of the Resident Population for Selected Age Groups by Sex for Alaska, California, Hawaii, New Jersey, New York, Oregon, Pennsylvania, Washington: April 1, 2010 to July 1, 2019.

[52] 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.

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

[54] 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.

[55] 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.

[56] Whitman, K., Reznik, G. L., & Shoffner, D. (2011). Who Never Receives Social Security Benefits? Social Security Bulletin, 71 (2) (2011), pp. 17-24.

[57] 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.

[58] 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.

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