



**THE EFFECT OF EARLY CLAIMING BENEFIT REDUCTION  
ON RETIREMENT RATES**

Damir Cosic and C. Eugene Steuerle

CRR WP 2021-1  
January 2021

Center for Retirement Research at Boston College  
Hovey House  
140 Commonwealth Avenue  
Chestnut Hill, MA 02467  
Tel: 617-552-1762 Fax: 617-552-0191  
<https://crr.bc.edu>

Both of the authors are with the Urban Institute. Damir Cosic is a senior research associate and C. Eugene Steuerle is an Institute fellow and the Richard B. Fisher chair. The research reported herein was derived in whole or in part from research activities performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA, any agency of the federal government, the Urban Institute, or Boston College. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof.

© 2021, Damir Cosic and C. Eugene Steuerle. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

## **About the Center for Retirement Research**

The Center for Retirement Research at Boston College, part of a consortium that includes parallel centers at the National Bureau of Economic Research, the University of Michigan, and the University of Wisconsin-Madison, was established in 1998 through a grant from the U.S. Social Security Administration. The Center's mission is to produce first-class research and forge a strong link between the academic community and decision-makers in the public and private sectors around an issue of critical importance to the nation's future. To achieve this mission, the Center conducts a wide variety of research projects, transmits new findings to a broad audience, trains new scholars, and broadens access to valuable data sources.

Center for Retirement Research at Boston College  
Hovey House  
140 Commonwealth Ave  
Chestnut Hill, MA 02467  
Tel: 617-552-1762 Fax: 617-552-0191  
<https://crr.bc.edu>

*Affiliated Institutions:*  
The Brookings Institution  
Mathematica – Center for Studying Disability Policy  
Syracuse University  
Urban Institute

## Abstract

This paper examines the effect of the increase in the Social Security Full Retirement Age (FRA) and the associated decrease in benefits for early claimants on retirement rates at ages 62 to 65. It uses information on age, sex, and labor force participation from the monthly *Current Population Survey* from 1976 to 2019. Critical components of the analysis include a difference-in-difference framework, comparison of three measures of retirement status, estimation of nonparametric and parametric models, and a test of the assumptions underlying the difference-in-difference approach. Although our model satisfied that test, the results do not guarantee that our specification is valid.

The paper found that:

- The increase in the FRA decreased the retirement rate at age 62 by about 5 percentage points (or 30 percent) for men and by about 2 percentage points (or about 20 percent) for women. The retirement rate at age 63 to 65 was not affected.
- Estimates of the parametric model show that a 1 percent increase in the early claiming reduction decreases the retirement rate by between 0.7 and 0.9 percentage points for men and between 0.2 and 0.4 percentage points for women.

The policy implications of the findings are:

- Our findings can inform evaluations of policy proposals to further increase the Social Security FRA.
- Our findings contribute to the understanding of the effects that the availability and generosity of pension benefits have on retirement decisions.

## Introduction

Even though it is one of the most researched topics in the retirement literature, the extent to which the availability and generosity of pensions, particularly Social Security benefits, influence workers' retirement decisions remains uncertain. Studies have produced a wide range of estimates depending on the methodological approach and data used. Relatively recent estimates of how a reduction in Social Security benefits affects labor force participation at age 62 for the beneficiaries who claim at that age run from a negative impact (Gustman and Steinmeier, 2009) to an impact that is statistically indistinguishable from zero (Blau and Goodstein, 2010) to one that is positive and large (Mastrobuoni, 2009). The main challenge in estimating these effects is identifying the pension effect and isolating it from all other effects that affect retirement behavior, such as the demographic composition of the population, economic conditions, and social norms. This study contributes a fresh perspective to this question by applying a difference-in-difference framework to the latest data.

In particular, we study the effects that an increase in the Social Security Full Retirement Age (FRA),<sup>1</sup> and the associated reduction in benefits for those who claim before the FRA, have had on the retirement behavior of older Americans. Beneficiaries who claim their benefits at the FRA receive the full amount of their primary insurance amount (PIA), which is calculated based on beneficiaries' 35 highest annual earnings. The 1983 Social Security reforms raised the FRA, which had been 65, by two months for each successive birth cohort, beginning with people born in 1938, until it reached 66 for those born in 1943. The reforms further raised the FRA to 67 for people born in 1960 and later, again phasing in the increase by two-month increments beginning with those born in 1955. Benefits can be claimed as early as 62, but early claimants received reduced monthly benefits. The increase in the FRA also reduced monthly Social Security benefits paid to beneficiaries who claim early. For example, a person born in 1937, whose FRA is 65, would receive 80 percent of their full benefits if they claimed at 62. In contrast, a person born in 1943, whose FRA is 66, would receive only 75 percent of their full benefits if they claimed at 62.

Our study exploits this variation in benefit reduction across cohorts to estimate the effect of benefit reduction on labor force participation. We use a difference-in-difference approach,

---

<sup>1</sup> The FRA is also referred to as "Normal Retirement Age" ("Normal Retirement Age," Social Security Administration, accessed August 12, 2020, <https://www.ssa.gov/oact/progdata/nra.html>).

similar to the one used by Song and Manchester (2007), to compare changes in retirement rates at ages 62 to 65 with changes in retirement rates at younger ages at which benefits are not available. This approach allows us to control for secular trends in the retirement rate and to isolate the effects of a Social Security benefits reduction. Because the treatment and control groups include the same birth cohorts, we are also able to control for any cohort-specific effects.

We estimate a range of specifications, which allows us to compare different assumptions and test the sensitivity of our results. In particular, we estimate a parametric and nonparametric difference-in-difference specification. We provide estimates for three age- and cohort-specific measures of retirement status: the *retirement incidence rate*, defined as the share of people born in a given year who retire at a given age; the *retirement hazard rate*, defined as the share of people born in a given year and participating in the labor force before reaching a given age who retire at that age; and the *cumulative retirement rate*, defined as the share of people born in a given year who are retired at a given age. We verify the main identifying assumption in our difference-in-difference design, which is the existence of a common trend in the treatment and control groups, and conduct sensitivity analysis with respect to the control group used. In addition, we try to reproduce estimates by Mastrobuoni (2009) on his original sample and on our sample that includes more recent birth cohorts.

Our preferred measure of retirement status is the incidence rate, for which the estimated effects of the Social Security reform are strong at age 62 for both men and women and statistically insignificant at ages 63 to 65. The retirement hazard rate shows significant and large effects at ages 62 and 65, while those at 63 and 64 are relatively small and not significant in all specifications. We have the least amount of confidence in the estimates that use the cumulative retirement rate, but we present them for comparison. We find that the control age can substantially affect the estimated magnitude, but it rarely changes the sign or statistical significance of the estimates. Our tests of identifying assumptions show that, for most combinations of treatment and control ages, the common trend hypothesis cannot be rejected for incidence and hazard rates. In the case of the cumulative retirement rate, the hypothesis is rejected in most cases for men but not for women.

## **Background**

The decision about when to retire can be dauntingly complex. To make the optimal decision, an individual has to consider a number of choices and the consequences for their potential income streams. This optimization problem in some ways resembles a search for the best move in a chess game, requiring that a player considers not only all possible current moves but also potential future moves and their availability conditional on the current move. An individual's decision to retire not only determines the future stream of income, but, because some retirement decisions are irreversible (e.g. few workers can return to their job after they leave), it also limits the universe of options available in the future. Evaluation of each of these many options requires complex calculations that include estimation of likely future incomes.

Early studies made their best attempts to model this sort of decision making mathematically by assuming a worker's utility function, constructing their budget constraint, and deriving optimality conditions. Although theoretically sound, these types of models were difficult to translate into empirical models that could identify, with available data, the effect of pensions on retirement decisions. They resulted in a wide range of estimates that depended on researchers' assumptions, from virtually no effect (Gordon and Blinder, 1980) to large effects that explained almost the entire change in retirement rates over time (Boskin, 1977; Boskin and Hurd, 1977). Recognizing this problem, several studies that followed focused on changes in Social Security benefits as the source of identification. Until 1972, when Social Security began providing automatic cost-of-living adjustments, the benefit amount was based on nominal earnings and any adjustment for inflation was set by Congress. These adjustments were usually unpredictable and sometimes large, and some researchers used them as exogenous shocks to individuals' expected future wealth. Hurd and Boskin (1984) found that an unexpectedly large increase in Social Security benefits between 1969 and 1972 could explain most of the decline in the labor force participation of older men during that time. One particularly interesting shock occurred in 1977 when Congress overcorrected an unintentionally large increase in the wage-indexing of benefits enacted in 1972. This created a cohort of Social Security recipients (the so-called notch generation born between 1917 and 1921) for whom the benefit amount, adjusted for inflation, decreased dramatically relative to previous cohorts. Krueger and Pischke (1992) used this anomaly as the source of identification and, with some improvements in methodology, found no significant effect of changes in Social Security wealth on labor supply.

Two more recent studies exemplify these different approaches to the estimation of the effect of changes in Social Security rules on retirement behavior. Blau and Goodstein (2010) take a parametric approach by modeling the retirement decision rule implied by the life-cycle model. Their dataset covers the years 1972, 1977, and 1983, when significant changes to Social Security were made. In contrast, Mastrobuoni (2009) uses a fully nonparametric approach and focuses only on the effects of the 1983 reform. Although Blau and Goodstein found that the 1983 reforms had non-negligible effects on labor supply, their estimates are substantially smaller than Mastrobuoni's. As authors of both papers note, this discrepancy may be due to the capacity of nonparametric specifications to capture noneconomic channels, such as social norms, through which pensions affect retirement behavior. Another advantage of nonparametric models is that they do not require an assumption about the functional form of the explanatory variable. On the other hand, parametric models often provide a useful interpretation of the estimates and allow them to be used in projections of future or counterfactual outcomes.

## **Data and Methods**

We use monthly *Current Population Survey* (CPS)<sup>2</sup> data from 1976 to 2019 to obtain information on labor force participation, age, and gender for a nationally representative sample of people ages 58 to 65. Because the CPS only asks for age at the time of the interview, birth year cannot be determined with certainty. Estimating it by subtracting age from the survey year would misclassify a large share of participants. The likelihood of misclassification is highest for the January sample and the lowest for the December sample. To reduce this likelihood, we follow a suggestion by Mastrobuoni (2009) to use only the December and January samples of CPS and adjust the birth year of the participants interviewed in January by subtracting one year from the difference between the survey year and age.

We constructed three age- and cohort-specific measures of retirement status based on the CPS participants' labor force status. First, we estimated the retirement incidence rate as a weighted share of the population born in a given year who leave the labor force at a given age, which is the sample equivalent of the cohort-specific unconditional probability of retiring at a given age. Second, we estimated the retirement hazard rate as a weighted share of workers born in a given year and participating in the labor force up to a given age who retire at that age, which

---

<sup>2</sup> Data obtained from King et al. (2018).

is the sample equivalent of the cohort-specific conditional probability of retiring at a given age. Third, we estimated the cumulative retirement rate as a weighted share of the population born in a given year who were not in labor force at the time of the survey. Under the assumption that retirement is an absorbing state, this measure is the empirical distribution function of the retirement age, whose value at a particular age equals the sum of the retirement incidence rates at ages younger than or equal to that age. For the incidence and hazard rates, we exploited the panel nature of the monthly CPS to obtain participants' labor status information at two times that are 12 months apart. Because only 50 percent of the sample is repeated at this frequency,<sup>3</sup> the incidence and hazard rates are estimated on a smaller sample, and thus with a lower precision, than the cumulative retirement rate.

Our main goal is to estimate changes in retirement behavior that were caused by the increase in the FRA implemented by the 1983 Social Security reform and isolate it from any secular trends in retirement behavior that might have existed during the observed period. Ages directly affected by the reform are all ages at which Social Security beneficiaries can claim benefits, but we focus on ages at which the affected cohorts would have their benefits reduced due to early claiming: 62 to 65. Our goal, then, is to estimate retirement rates for the 1938 and later cohorts at these ages relative to what they would have been if there had been no change in Social Security rules. The main challenge is that these counterfactual retirement rates are unknown.

One way to overcome this challenge and construct the counterfactual rates is to start with the pre-intervention trend and adjust it by the effects of all other factors that shifted labor force participation over the observed period. If successful, this approach would produce the counterfactual retirement rates that would have prevailed in the absence of Social Security reform. This is the essence of the method used by Mastrobuoni (2009), even though he did not explicitly construct the counterfactual. Because our methods build on his, we start by briefly outlining it.

Mastrobuoni's method can be divided into three steps. First, the dependent variable  $y_i$ , which equals one if person  $i$  of age  $A_i$  born in year  $C_i$  is retired and zero otherwise, is regressed on a set of age and birth year dummies and a vector of individual characteristics and economic outcomes  $X_i$ :

---

<sup>3</sup> U.S. Census Bureau (2006, p. 13-14).



$$y_i = \sum_{a=61}^{65} \mathbb{I}(A_i = a) \left( \lambda_a + \sum_{c \neq 1937} \beta_{a,c} \mathbb{I}(C_i = c) \right) + \delta' X_i + \varepsilon_i \quad 1)$$

The estimate of interest,  $\beta_{ac}$ , is the coefficient on the interaction term between age and birth-year dummies,  $\lambda_a$  represents age fixed effects,  $\delta$  is a vector of parametric estimates of effects of other factors that affect retirement rate, and  $\varepsilon_i$  is the error term. The base birth year omitted from the model is 1937, so the interaction coefficients can be interpreted as the difference of the cumulative probabilities of retirement at age  $a$  between cohorts  $c$  and 1937:

$$\hat{\beta}_{ac} = E[y|C = c, a, X] - E[y|C = 1937, a, X]$$

In the second step, these estimates are combined to calculate the cohort-specific change in the mean retirement age relative to the cohort of 1937. The mean retirement age is defined as the sum of products of each age and the probability of retirement at that age, but Mastrobuoni showed that, under certain assumptions, the cohort-specific change in the retirement age can be calculated by adding the cohort's beta-coefficient estimates for ages 62 to 65. Finally, in the third step, the difference between the post- and pre-intervention trends in the retirement age is calculated as a weighted average of the changes in the mean retirement age, which the author interprets as the difference in retirement age trends caused by the Social Security reform.

We take issue with two assumptions underlying these estimates. First, the calculation of the change in the mean retirement age in the second step assumes that the change in the retirement probabilities at ages 61 or younger and 66 or older relative to the 1937 cohort is zero. However, this assumption almost certainly does not hold because there are clear trends in retirement probabilities at ages 61 and younger (see figures 1 and 2). This would bias the estimate of the change in the retirement age based only on retirement probabilities at ages 62 to 65. To test this assumption, we replicate the author's estimates with an expanded age range. If the assumption holds, we expect little sensitivity in the estimates; otherwise, the calculation of changes in the mean retirement age requires summation over a wider range of possible retirement ages. Because our results indicate that the latter is likely the case, we eschew the mean retirement age as a measure of retirement status and focus on retirement rates, which are more easily separable by age.

The second assumption that might not hold is the assumption that the counterfactual retirement rate can be constructed based on the pre-intervention trend and observed factors that affect the retirement rate. For this assumption to hold, one would have to control for all the factors whose effect on the retirement rate changed over the observed period. Considering the current state of understanding of trends in the retirement rate, however, it is unlikely that we are able to do that. Because of the limited knowledge of the processes that shape labor force participation at older ages, it is likely that our list of factors that affect it is incomplete. Moreover, even some factors on the list, including life expectancy, income and wealth inequality, and social norms other than those related to retirement age (e.g. social norms about women in labor force), affect retirement rate in a way that is difficult to estimate parametrically. We address this issue by using a difference-in-difference design that allows us to construct counterfactuals based on a control group that was not affected by the changes in Social Security rules.

In addition, we hope to improve on Mastrobuoni's estimates by using more recent data. At the time of his study, data were available for only four single-year birth cohorts that were affected by the reform (people born between 1938 and 1941), which might not be sufficient to identify changes in trends with the required precision. To test the sensitivity of trend estimates, we replicate the estimation on an expanded dataset that includes surveys from January 1975 to January 2019, which adds the 1942 to 1953 cohorts. We replicate only the specification that the author calls "restricted" and does not include additional controls.

Our difference-in-difference framework allows us to compare pre- and post-intervention trends in retirement rates between the treatment and control groups. In an ideal case, the two groups would have the same composition and the treatment would be randomly assigned. This would allow us to assume that the post-intervention trend in the control group is the same as the counterfactual post-intervention trend in the treatment group. Consequently, the difference in the changes in outcomes between the treatment and control groups represents the average treatment effect. Most designs, including ours, do not rise to this high standard, which is why we test our assumptions and the sensitivity of our estimates.

Our treated groups consist of individuals ages 62 to 65 who are eligible for Social Security benefits; the control groups include individuals ages 58 to 61 who are not yet eligible for the benefits. We make two strong assumptions. First, we assume that retirement behavior at

pre-eligibility ages is not affected by changes in the benefit amount. Second, we assume that pre-intervention trends in the retirement rate are the same in the control and treatment groups and, consequently, that in the absence of an intervention they would have remained the same in the post-intervention period. The first assumption is difficult to test, but we draw some confidence from the Social Security Board of Trustees who make the same assumption in their annual report (Office of the Chief Actuary, 2019, p. 63). We test the second assumption by estimating a model augmented by a linear trend.

Our design is similar to the one developed by Song and Manchester (2007) who used it to estimate the effect of changes in Social Security rules on the rate of claiming Social Security benefits. Their treatment group included individuals ages 62 to 70 and the control group consisted of individuals ages 71 and 72 who were not affected by changes in the Social Security rules. Their design implies the same assumptions as ours: that claiming behavior is not affected at ages for which benefits did not change, and the claiming trends in the treatment and control groups would have been the same in the absence of the benefits change.

Unlike Song and Manchester, our design uses birth cohorts to define pre- and post-intervention periods, rather than calendar years. This approach has a couple of advantages. First, it allows us to control for cohort effects, which ensures that, aside from age, demographic characteristics of the treatment and control groups are the same.<sup>4</sup> Second, because the treatment is defined with respect to birth year rather than calendar year, intervention timing varies with age. For example, the treatment starts in 2000 for 62-year-olds, in 2001 for 63-year-olds, etc. Using birth cohort as the time dimension removes this dependency and avoids the complications associated with staggered treatment (see Goodman-Bacon, 2019).

The choice of the control group has to reconcile two opposing requirements. To ensure similarity with the treatment group, the control group ages should be as close as possible to the treatment group ages. But to ensure that a change in Social Security rules does not affect the retirement decisions of people in the control group, their ages should be sufficiently distant from the eligibility ages. There is no way for us to determine the optimal trade-off between these two requirements, which is why we provide estimates for a range of control ages (58 to 61) and conduct a sensitivity analysis.

---

<sup>4</sup> This is not entirely true because of differences in mortality rates between population groups. However, it is a good approximation because our treatment and control ages are relatively close.

Rather than using individual-level observations, we aggregate them by age and birth year. In this, we follow a suggestion by Angrist and Pischke (2009) who caution about the risk of underestimating standard errors and consequently overstating the precision of estimates, when individual-level residuals are correlated at the group level. In our case, the main source of such correlation is the treatment, which varies by age and cohort. Denoting the average retirement rate for cohort  $c$  at age  $a$  with  $y_{ac}$  and treatment ages with  $a_T$ , which include ages 62 to 65, our nonparametric specification is:

$$y_{ac} = \lambda_a + \gamma_c + \beta_c^{a_T} \mathbb{I}(c \geq 1938 \wedge a = a_T) + \varepsilon_{ac} \quad 2)$$

The  $\beta_c^{a_T}$  coefficients represent the cohort-specific effects of the Social Security reform on the retirement rates at treatment age  $a_T$ . We estimate this equation separately for each treatment age, and calculate the mean  $\hat{\beta}^{a_T}$  as the mean effect of the reduction in benefits on the retirement rate for age  $a_T$ .

In a parametric version of the model, we define the treatment  $D_{ac}$  as the percentage increase in the benefit reduction for cohort  $c$  born in 1938 or later at age  $a$  relative to the benefit reduction for the 1937 cohort at the same age. Figure 3 shows values of  $D_{ac}$ , and the following equation represents the econometric specification:

$$y_{ac} = \lambda_a + \gamma_c + \beta D_{ac} + \varepsilon_{ac} \quad 3)$$

The coefficient  $\beta$  can be interpreted as the percentage-point reduction in the retirement rate due to a one-percent increase in the benefit reduction.

A common problem with using ordinary least squares (OLS) to estimate this type of model is the serial correlation among predictors, and in particular in the treatment variable. As Bertrand, Duflo, and Mullainathan (2004) point out, a predictor with positive serial correlation can cause a substantial underestimation of standard error. To deal with this problem, we use the randomization inference, an approach suggested by the authors. We generate 200 random permutations of the treatment across ages and cohorts for each gender, then estimate the above equations for each permutation. We then sort the estimates by size and use the 5th and 195th as the interval in which 95 percent of estimates would fall under the null hypothesis that reducing early claiming benefits has no effects on retirement behavior. If the point estimate falls outside of the interval, the null hypothesis can be rejected with 95 percent certainty.

One of the main concerns in any difference-in-difference design is the validity of the common trend assumption. For equations (2) and (3) to be a valid model for estimation of the treatment effect, it must be the case that the retirement rates for the control and treated groups were following the same trend before the intervention. Only then can we assume that they would have followed the same trend after the intervention in the absence of the treatment and conclude that the difference in post-intervention trends was caused by the intervention. To test this assumption, we estimate equation (4), which expands equation (2) by adding an interaction between a linear trend  $t$  and the treatment-age dummy variable.

$$y_{ac} = \lambda_a + \gamma_c + \beta_c \mathbb{I}(c \geq 1938 \wedge a = a_T) + \theta \mathbb{I}(a = a_T)t + \varepsilon_{ac} \quad 4)$$

A rejection of the null hypothesis that coefficient  $\theta$  is zero would indicate that the common trend is not a valid assumption.<sup>5</sup> Related to the common-trend assumption is the assumption about the functional form of the dependent variable (Kahn-Lang and Lang, 2020). For example, two groups could have a common trend in the retirement incidence rate but diverging trends in the hazard rate. This is why we estimate the above models for all three measures of retirement status.

A potential problem with the cumulative retirement rate as a measure of retirement status at a given age is that it depends not only on retirement status at that age but also on retirement statuses at all younger ages. For example, the cumulative retirement rate at 62 equals the cumulative rate at 61 plus the incidence rate at 62. An increase in the cumulative rate at 61 will automatically raise the cumulative retirement rate at 62 even if the rate at which 62-year-olds exit the labor force has not changed. It is not clear that this a desirable feature of an age-specific measure of retirement status, but this feature may also adversely affect our difference-in-difference design by effectively eliminating the comparison with the control group. The estimator of the average treatment effect on the cumulative retirement rate at 62 that uses age 61 as the control group effectively calculates the difference between pre- and post-intervention differences in the cumulative retirement rates at ages 62 and 61. But, because the cumulative rate at 62 additively includes the cumulative rate at 61, the cumulative rates at age 61 cancel each other for each cohort, and the remaining terms represent the difference in the retirement rate at

---

<sup>5</sup> As Bilinski and Hatfield (2020) point out, a failure to reject, on the other hand, does not guarantee that the assumption is valid.

age 62 between post- and pre-intervention cohorts, effectively collapsing the difference-in-difference design into an interrupted-time-series design.<sup>6</sup> Because the estimator relies only on the difference between pre- and post-intervention trends and not on the comparison with the control group, it requires a stronger assumption that the pre-intervention trend would have remained the same in the post-intervention period in the absence of the intervention. This is why we prefer the other two measures over the cumulative retirement rate. However, we present results based on all three measures to demonstrate differences in how the choice of a measure affects estimates and to allow comparison with other studies. The above criticism would not apply if we applied a non-homomorphic transformation, such as the logarithmic or logit functions, to the cumulative retirement rate. We estimated the above models based on these two transformations too, but tests rejected the common trend hypothesis and we do not report those results in the paper.

## Results

We start by testing Mastrobuoni’s results (2009) for sensitivity to birth cohorts included in the sample and to ages included in the estimation of average retirement age. Table 1 contains the estimates of the change in the average retirement age for cohorts born in 1938 and later. This is the final result of the three-stage estimation that starts with an estimation of equation (1), the summation of  $\beta_{a,c}$  by cohort for ages 62 to 65 to obtain cohort-specific changes in the average retirement age, and averaging these changes to obtain the change in trends in the average retirement age for cohorts born in 1938 and later. The first two columns show results of our replication of Mastrobuoni’s estimates for the restricted case without controlling for individual characteristics; the last two columns show the estimates obtained by using the same method but on an expanded dataset that includes more recent cohorts. We were not able to reproduce the author’s results exactly, but considering that all our intermediate results are relatively close to the author’s, we are reasonably confident that we faithfully reproduced all the key aspects of his method.<sup>7</sup> Our estimates on the extended sample show that results are sensitive to the range of

---

<sup>6</sup> If we denote cumulative and incidence rates for age  $a$  and a given cohort as  $y_a$  and  $i_a$ , then  $y_{62} = y_{61} + i_{62}$ . If we add superscripts 0 and 1 to denote pre- and post-intervention periods, the average treatment effect at 62 with control age 61 is  $\Delta = (y_{62}^1 - y_{62}^0) - (y_{61}^1 - y_{61}^0) = (y_{61}^1 + i_{62}^1 - y_{61}^0 + i_{62}^0) - (y_{61}^1 - y_{61}^0) = i_{62}^1 - i_{62}^0$ .

<sup>7</sup> Our estimated change in the trend for women is 0.65 and not statistically significant, comparing to the author’s 0.88 and statistically significant at the 95 percent level (Mastrobuoni, 2009, table 5, panel A, column 5). Our

birth cohorts included in the sample. The inclusion of more recent cohorts significantly changed the estimates but not in the expected direction. The extended sample includes birth cohorts that were increasingly more affected by the Social Security reforms as the early claiming reduction increased for later cohorts. For example, a person born in 1938 who claimed at age 62 would have received 1 percent lower benefits than a person born a year earlier who claimed at the same age, while someone born in 1943 claiming at the same age would receive six percent lower benefits. Assuming that the main effect on the retirement age comes from this reduction in benefits, we expected that including more recent cohorts would increase the mean effect on the retirement age. Instead, the estimates for both men and women significantly decreased when we included more recent cohorts. We interpret this result as evidence that data on only three post-intervention cohorts were insufficient for estimation of a change in trend.

Our second sensitivity test looks at the second step in Mastrobuoni's estimation that estimates a change in the mean retirement age. Under the assumption that there were no changes in retirement rates at ages 61 or younger and ages 66 or older, the author calculates changes in the mean cohort-specific retirement age based only on retirement rates at ages 62 to 65. Those assumptions produce the change in the retirement age trend discussed above and shown again in first two columns of table 2. Under a relaxed assumption that includes ages 55 to 65 in the estimation, the resulting changes in the retirement age trend, which are shown in the last two columns of the same table, are more than 50 percent larger. This indicates that the assumption of no change in retirement rates at ages younger than 62 and older than 65 might be too strong.

We now turn to our results, but before we show the estimates of the models described in the previous section we start with the most intuitive way of estimating treatment effects: as the difference between the post- and pre-intervention outcomes for the treated group minus the difference between the post- and pre-intervention outcomes for the control group. These results based on retirement incidence rates at treated ages 62 to 65 and control ages 58 to 61 are shown in table 3. The "Pre" and "Post" columns show the average retirement rate for the 1920 to 1937 cohorts and the 1938 to 1955 cohorts, respectively, and the "Post-Pre" column shows the difference between these two columns indicating the change in retirement rate after the intervention. According to this column, the average retirement rate for men changed

---

estimate for men, 1.09, is slightly higher than 0.94 by the author and statistically significant at the same level (Mastrobuoni, 2009, table 5, panel B, column 5).

substantially after the intervention only at age 62, when it fell by 5.1 percentage points, or one-third, from 14.7 percent to 9.6 percent. Changes at other ages were less than one percentage point. The results are different for women, whose retirement rates increased by one percentage point at ages 58, 59, 61, and 64, decreased by one percentage point at age 62, and remained relatively constant at other ages.

The average treatment effect for each treated age can be calculated with respect to each control age as the difference in the values in the “Post-Pre” column between the treated and control ages. For example, the treatment effect for 62-year-old men based on control age 61 equals the difference between -5.1 and -0.3, or -4.8. Based on other control ages, the treatment effect for men at this age ranges from -5.4 to -4.5. The treatment effect for men ages 63 to 65 is less than one percentage point. For women, the treatment effect is also the strongest at age 62, but it is smaller than for men of the same age. Depending on control age, it ranges from -2.2 to -1.3 percentage points. At age 63, the treatment effect is about -1 percentage point based on control ages 58, 59, and 61, and only -0.3 percentage points based on control age 60. The treatment effect for age 64 varies from zero to 0.9 percentage points, and for age 65 from -0.4 to 0.5 percentage points.

Even though these calculations, under our identifying assumptions, provide unbiased and intuitive estimates of the magnitude of the changes in retirement rates, they do not allow statistical inference about the estimates’ precision. For that, we use OLS regression with randomization inference. Our estimates of equation (2) for treatment ages 62 to 65 and control ages 58 to 61 are presented in tables 4, 5 and 6. Estimates in table 4 use the retirement incidence rate, those in table 5 use the retirement hazard rate, and those in table 6 use the cumulative retirement rate. Each table contains estimates for men and women. Each row represents estimates for one combination of the treatment and control age. The “treatment effect” column contains point estimates of the effect. The “placebo interval” column shows the interval that was generated by the randomization inference process, which estimated equation (2) 200 times, each time with a permuted treatment variable. The interval, which was constructed by using the 5<sup>th</sup>- and 195<sup>th</sup>-largest estimates as its endpoints, denotes the range of values in which estimates would land 95 percent of the time if the treatment was placebo.

The magnitude of the estimated effects shown in table 4 is similar to those obtained by simple differencing shown in table 3, but the placebo interval shows that the treatment effect for



men is statistically significant only at age 62 (for all control ages). For all other treatment ages, the point estimate lies within the placebo interval indicating that it was statistically indistinguishable from a placebo. The estimated effect on the retirement rate at age 62 represents a 30 to 35 percent reduction relative to pre-1938 cohorts. The women's retirement rate at age 62 decreased by between 1.7 and 2.5 percentage points, or around 20 percent, after the change in the early claiming benefit reduction, and the estimates are statistically significant for all control ages. The treatment effect on women's retirement rate is also significant at age 63 for control ages 59 and 61.

The estimated effects on the retirement hazard rate, which are shown in table 5, are statistically significant at all treatment ages for both men and women, at least for some control ages, but they are largest and most consistently significant for ages 62 and 65. For men, the retirement hazard rate at 62 decreased by between 8.1 and 9.8 percentage points, depending on the control age. The relative magnitude of these changes is similar to the magnitude obtained for the incidence retirement rate, around 30 percent. At 63, the effect is significant only for control age 59. The estimates at 64 are statistically significant for all control ages and range from -5.7 to -4.2 percentage points. For age 65, the estimated effects are large and significant, ranging from -8.7 to 6.4 percentage points. For women, the retirement hazard rate decreased by between 5 and 7.2 percentage points at age 62; between 3.5 and 6.8 percentage points at age 63; and between 5.6 and 8.7 percentage points at 65. At 64, the estimated effect was around -3 percentage points and statistically significant for control ages 58 and 59, and statistically insignificant for control ages 60 and 61.

The estimates based on the hazard and incidence rate are consistent at age 62 but they diverge for older ages. Part of the reason for these differences is that the hazard rate is a composite measure that depends on the retirement incidence rate in its numerator and the retirement survival function in the denominator. Even if the retirement incidence rate at a given age does not change, a change in incidence rates at younger ages would change the survival function and therefore the denominator of the hazard rate. For example, table 4 shows that the effect on the incidence rate was significant at age 62 but not at 63. This reduction in the incidence rate at age 62 decreased the cumulative rate at the same age and increased the retirement survival function, which is the denominator of the hazard rate at age 63. Even though there was no effect on the incidence rate at age 63, its decrease at 62 can cause the hazard rate at

63 to decrease. This dependence of the hazard rate on retirement status over a range of ages is one reason why we prefer the incidence rate over it; the other reason is the higher volatility of the hazard rate.

Results obtained by using the cumulative retirement rate, which are shown in table 6, are fundamentally different from the results obtained by using the other two functional forms. This is especially clear from the estimates of changes in women's cumulative retirement rate, which are not statistically significant at any age, even though changes in both the incidence and hazard rates were statistically significant at age 62 for all control ages. One reason for this is the additivity of cumulative retirement rate and its elimination of the control group described in the previous section. As table 3 shows, the decrease in the incidence rate for 62-year-old women was small, only 1.1 percentage point relative to the pre-intervention rate. But because there was a commensurable increase in incidence rates at all but one control age, the change relative to the control group was 2 percentage points. In the case of the cumulative retirement rate, the only comparison being made is in the incidence rate between pre- and post-intervention periods, and this change was relatively small. The other reason is the steady downward trend of women's cumulative retirement rates, which creates a large variation in the sample resulting in wide placebo intervals that are much wider than those estimated for men. In contrast, there is much less variation in the incidence and hazard rates, and consequently placebo intervals estimated with these two measures for women are in most cases narrower than those for men.

We now examine the validity of the assumption that pre-intervention trends in control and treatment groups were parallel by estimating equation (4) and testing the hypothesis that  $\theta$  equals zero. Tables 7 contains p-values for the test of the common trend assumption for the three functional forms, treatment ages 62 to 65 and control ages 56 to 61. For each functional form of the dependent variable and each combination of control and treatment ages, the value in the cell is the p-value of a test that the coefficient  $\theta$  in equation (4) is zero. When the p-value is less than 5, the common trend hypothesis can be rejected at the 95 percent significance level. For men, the retirement incidence rate passes the test for all combinations of treatment and control age with p-values well above the 5-percent threshold. The retirement hazard rate does not perform as well, having the common trend hypothesis rejected for two combinations of treatment and control ages, but the p-values for other combinations are mostly high. The cumulative retirement rate, however, fails the test for most cases. The test fails to reject the common-trend hypothesis

only for treatment age 62 and all control ages, and treatment age 63 and control age 61. For all other combinations, the p-value is less than 5. The test results are different for women, for whom the cumulative retirement rate yields the highest p-values and the test failed to reject the common-trend hypothesis for any combination of treatment and control ages. The outcome is similar for the incidence rate, although the p-values in a couple of cases are close to 5 percent. The common trend hypothesis was rejected only in the case of the retirement hazard rate in four combinations of treatment and control ages.

Finally, we show estimates of the parametric model (3) based on the retirement incidence rate. The estimates of the parameter  $\beta$  in equation (3), which are shown in table 9, represent the average effect on the retirement incidence rate of a 1 percent increase in the benefit reduction. To calculate the total effect for a particular age and cohort, this estimate must be multiplied by the benefit reduction for that age and cohort. For example, the estimated effect for 62-year-old men with control age 61 is a 0.75 percentage point decrease in the retirement rate per 1 percent of increase in benefit reduction. Because beneficiaries born in 1943 and claiming at 62 would have their benefits cut by 6.25 percent more than their counterparts born in 1937, men's retirement rate decreased by 4.7 percentage points due to the benefit reduction.

The sign and significance of estimates is almost identical to the nonparametric estimates. The treatment effect is negative and statistically significant for men and women only at age 62. The estimates lie within the placebo interval for all other age combinations except for treatment age 64 and control age 60 for women. To compare these estimates to those from the nonparametric model, we can calculate the average increase in the early claiming reduction across the 1938 to 1955 cohorts and multiply it by the parametric estimate of the effect of a 1 percent increase. For example, the average increase in the benefit reduction for claiming at 62 for these cohorts was 5.4 percent. Multiplying it by -0.75 estimate for 62-year-old men yields an average incidence rate change of -4.1 percentage points. The equivalent non-parametric estimate is -4.8 percentage points. Taking other control ages into account, the estimates range from -5.1 to -3.9 for the parametric model and -5.8 to -4.8 for the nonparametric model. For women of the same age, these ranges are -2.3 to -1.2 in the parametric case, and -2.7 to -1.7 in the nonparametric case. The proximity of the estimates based on two different models provides some assurance that the models are not misspecified.

## Conclusion

A potential future increase in the Social Security FRA is a commonly discussed policy in the academic literature and among policy makers. Our findings have important implications for its evaluation. In addition to reducing Social Security outlays, increasing the FRA has the potential to increase labor force participation among older workers, which, in turn, could help grow the economy and increase tax revenues. Modeling these effects hinges on the assumption about the impact that such a policy would have on labor force participation. This paper contributes a new perspective on this relationship by studying the effect on retirement behavior of the 1983 Social Security reform, which raised the FRA. In addition to providing a set of estimates based on the latest data, we also examine different methodological choices and their consequences for the analysis. We adopt a difference-in-difference estimation as our framework and, within it, explore three functional forms, compare parametric and nonparametric model specifications, and examine how the choice of a control group affects the estimates.

We found that functional form is a particularly important choice that affects conclusions in fundamental ways. The analysis based on the incidence retirement rate indicates that the reform reduced retirement rates only at age 62 for both men and women. Results based on the hazard rate are also symmetrical with respect to gender, but the effects are significant for all four affected ages that we analyzed. In contrast, when we use cumulative retirement rate as a measure, the effects were large and significant for men at all four ages, but insignificant for women. Based on our literature review, the most commonly used measures are the cumulative retirement rate and the retirement hazard rate, but there is little guidance for selecting one or the other. However, our analysis suggests that the incidence rate provides more useful insight into changes in retirement behavior. It measures the share of the population retiring at a single age and is not affected by changes at other ages, as is the case with the other two measures. In addition, using the incidence rate yields the most precise estimates (because of narrow placebo intervals and low variation of estimates across control ages) and the highest p-values in our test of the common-trend hypothesis. Finally, we provide parametric estimates of the effect of an increase in FRA on retirement rates, which are more suitable for use in projections and simulations.


## References

- Angrist, Joshua David and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics*. Princeton, NJ: Princeton University Press.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan. 2004. "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics* 119(1): 249-275.
- Bilinski, Alyssa and Laura A. Hatfield. 2020. "Nothing to See Here? Non-Inferiority Approaches to Parallel Trends and Other Model Assumptions." *ArXiv:1805.03273 [Stat]* January.
- Blau, David M. and Ryan M. Goodstein. 2010. "Can Social Security Explain Trends in Labor Force Participation of Older Men in the United States?" *Journal of Human Resources* 45(2): 32-63.
- Boskin, Michael J. 1977. "Social Security and Retirement Decisions." *Economic Inquiry* 15(1): 1-25.
- Boskin, Michael J. and Michael D. Hurd. 1977. "The Effect of Social Security on Early Retirement." Working Paper 204. Cambridge, MA: National Bureau of Economic Research.
- Goodman-Bacon, Andrew. 2018. "Difference-in-Differences with Variation in Treatment Timing." Working Paper 25018. Cambridge, MA: National Bureau of Economic Research.
- Gordon, Roger H. and Alan S. Blinder. 1980. "Market Wages, Reservation Wages, and Retirement Decisions." *Journal of Public Economics* 14(2): 277-308.
- Gustman, Alan L. and Thomas Steinmeier. 2009. "How Changes in Social Security Affect Recent Retirement Trends." *Research on Aging* 31(2): 261-290.
- Hurd, Michael D. and Michael J. Boskin. 1984. "The Effect of Social Security on Retirement in the Early 1970s." *Quarterly Journal of Economics* 99(4): 767-790.
- Kahn-Lang, Ariella and Kevin Lang. 2020. "The Promise and Pitfalls of Differences-in-Differences: Reflections on 16 and Pregnant and Other Applications." *Journal of Business & Economic Statistics* 38(3): 613-620.
- King, Miriam, Steven Ruggles, Sarah Flood, and Robert J. Warren. 2018. "Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [Dataset]." Minneapolis, MN: IPUMS. Available at: <https://doi.org/10.18128/D030.V6.0>

- Mastrobuoni, Giovanni. 2009. "Labor Supply Effects of the Recent Social Security Benefit Cuts: Empirical Estimates Using Cohort Discontinuities." *Journal of Public Economics* 93(11): 1224-1233.
- Office of the Chief Actuary. 2019. "The Long-Range Economic Assumptions for the 2019 Trustees Report." Baltimore, MD: U.S. Social Security Administration.
- Song, Jae and Joyce Manchester. 2007. "Have People Delayed Claiming Retirement Benefits- Responses to Changes in Social Security Rules?" *Social Security Bulletin* 67(2).
- U.S. Census Bureau. 2006. "Design and Methodology." CPS TP66. Current Population Survey Technical Paper. Washington, DC.

## Figures


Figure 1. *Cumulative Retirement Rates for Men by Age and Birth Year*



Notes: Lines represent predicted values from regressions of retirement rates on birth-year dummies. The vertical line represents the 1938 birth year.

Source: *Current Population Survey* (1978-2019).

Figure 2. *Cumulative Retirement Rates for Women by Age and Birth Year*




Notes: Lines represent predicted values from regressions of retirement rates on birth-year dummies. The vertical line represents the 1938 birth year.

Source: *Current Population Survey* (1978-2019).



Figure 3. Increase in Early Claiming Benefit Reduction Due to the Social Security Reform



Notes: Increase in benefit reduction at each age is calculated relative to early claiming benefit reduction for recipients born in 1937. For example, a person born in 1937 who claimed at 62 would receive 80 percent of their full PIA. A person born in 1943 claiming at the same age would receive only 75 percent of their full PIA, or 6.25 percent less, which represents the percentage increase in reduction.

## Tables

Table 1. *Estimates of the Changes in Retirement Age Trends are Sensitive to Included Birth Cohorts*

<b>Birth cohorts 1928-1941</b>		<b>1928-1953</b>	
<b>Men</b>	<b>Women</b>	<b>Men</b>	<b>Women</b>
1.087**	0.650	0.556*	0.387
(0.520)	(0.474)	(0.320)	(0.291)

Notes: Standard errors in parentheses, \*\*  $p < 0.05$ , \*  $p < 0.1$ . Estimation using methods by Mastrobuoni (2009) based on two samples that include different ranges of birth cohorts. The results are calculated based on estimates of  $\beta_{a,c}$ , which are first summed by cohort for ages 62 to 65 to obtain the cohort-specific change in the average retirement age. These changes in the average retirement age are then averaged to obtain the change in trends for cohorts born in 1938 and later.

Table 2. *Estimates of the Changes in Retirement Age Trends are Sensitive to Included Ages*

<b>Ages 62-65</b>		<b>Ages 55-65</b>	
<b>Men</b>	<b>Women</b>	<b>Men</b>	<b>Women</b>
1.087**	0.650	1.670**	1.084
(0.520)	(0.474)	(0.691)	(0.663)

Notes: Standard errors in parentheses, \*\*  $p < 0.05$ , \*  $p < 0.1$ . Estimation using methods by Mastrobuoni (2009) based on two who estimated the change in cohort-specific average retirement age as the sum of  $\beta_{a,c}$  coefficients for ages 62 to 65, under the assumption that these coefficients for ages outside of the range are zero. The estimates for the 55-65 age range relax this assumption and include ages 55 to 61.

Table 3. Average Retirement Incidence Rates and Average Treatment Effects

Age	Pre	Post	Post - Pre	Average Treatment Effect			
				61	60	59	58
<i>Men</i>							
58	4.8	5.0	0.2				
59	4.8	5.1	0.3				
60	7.1	6.6	-0.6				
61	6.6	6.3	-0.3				
62	14.7	9.6	-5.1	-4.8	-4.5	-5.4	-5.2
63	8.7	8.6	-0.1	0.2	0.5	-0.4	-0.2
64	7.7	7.3	-0.4	-0.2	0.1	-0.8	-0.6
65	9.8	9.5	-0.3	0.0	0.3	-0.6	-0.5
<i>Women</i>							
58	4.7	5.8	1.1				
59	5.3	6.4	1.1				
60	6.5	6.7	0.2				
61	6.2	7.2	1.0				
62	9.5	8.4	-1.1	-2.1	-1.3	-2.2	-2.2
63	7.2	7.0	-0.1	-1.1	-0.3	-1.2	-1.2
64	6.2	7.3	1.1	0.1	0.9	0.0	0.0
65	7.5	8.2	0.6	-0.4	0.5	-0.4	-0.4

Notes: Column "Pre" shows the average retirement rate for 1920 to 1937 cohorts, column "Post" for cohorts 1938 to 1955, and column "Post-Pre" shows the difference between first two columns. The four columns under "Treatment Effect" show the difference between the values in the "Post-Pre" column for the treatment age in the intersecting row and control age in the intersecting column. E.g. value -5.6 for treatment age 62 and control age 61 is the difference of -6.6 in the "Post-Pre" column at age 62 and -1.0 at age 61. Sample contains cohorts 1920 to 1955.

Table 4. *Nonparametric Estimates of the Treatment Effect on the Retirement Incidence Rate*

Age	Men		Women	
	Control age	Treatment Effect Placebo Interval	Treatment Effect Placebo Interval	
62	58	-5.51 [-2.70, 2.49]	-2.22 [-0.66, 1.49]	
	59	-5.77 [-3.00, 2.07]	-2.45 [-0.98, 1.57]	
	60	-4.76 [-2.99, 1.89]	-1.69 [-0.81, 1.51]	
	61	-4.82 [-2.90, 1.87]	-2.69 [-0.68, 1.56]	
63	58	-0.16 [-3.36, 1.90]	-0.98 [-1.04, 1.75]	
	59	-0.40 [-3.06, 1.57]	-1.25 [-0.84, 1.49]	
	60	0.32 [-2.86, 1.44]	0.02 [-0.95, 1.42]	
	61	0.30 [-3.23, 1.64]	-1.26 [-1.01, 1.73]	
64	58	-0.64 [-1.73, 1.75]	0.36 [-0.62, 1.58]	
	59	-0.77 [-1.59, 1.73]	-0.09 [-1.30, 1.30]	
	60	0.00 [-1.99, 1.31]	1.04 [-1.03, 1.18]	
	61	-0.07 [-1.58, 1.70]	0.22 [-1.07, 1.73]	
65	58	-0.56 [-2.12, 5.12]	-0.38 [-1.05, 3.56]	
	59	-0.63 [-2.13, 4.16]	-0.47 [-1.06, 4.19]	
	60	0.52 [-1.63, 4.80]	0.43 [-1.68, 3.92]	
	61	0.22 [-2.10, 5.15]	-0.47 [-0.91, 4.33]	

Notes: The average treatment effect was estimated as the average of  $\beta$  coefficients in equation (2). The placebo interval represents the interval between 2.5 and 97.5 percent of the distribution of estimates generated by randomization inference. A point estimate outside of that interval indicates a statistically significant effect at the 95 percent level. The sample includes people born between 1920 and 1955.

Table 5. *Nonparametric Estimates of the Treatment Effect on the Retirement Hazard Rate*

Age	Men		Women		
	Control age	Treatment Effect	Placebo Interval	Treatment Effect	Placebo Interval
62	58	-9.31	[-3.22, 3.14]	-7.24	[-2.03, 2.97]
	59	-9.82	[-3.37, 2.85]	-6.62	[-1.79, 3.01]
	60	-9.17	[-3.46, 2.63]	-5.01	[-1.73, 2.79]
	61	-8.12	[-2.93, 2.61]	-6.47	[-1.66, 2.93]
63	58	-3.64	[-4.94, 2.59]	-6.78	[-1.92, 3.53]
	59	-4.47	[-4.39, 3.00]	-6.27	[-1.90, 2.83]
	60	-4.39	[-5.32, 2.22]	-3.48	[-2.17, 2.62]
	61	-3.18	[-4.93, 2.84]	-5.64	[-2.29, 3.53]
64	58	-5.40	[-2.38, 2.33]	-2.98	[-1.82, 2.56]
	59	-4.47	[-2.31, 2.30]	-3.06	[-2.78, 2.38]
	60	-5.66	[-3.23, 1.93]	-0.10	[-2.81, 1.56]
	61	-4.23	[-2.35, 2.22]	-1.92	[-2.59, 2.56]
65	58	-8.65	[-3.10, 3.39]	-8.69	[-2.53, 3.41]
	59	-6.97	[-3.02, 3.04]	-8.28	[-3.49, 2.96]
	60	-7.01	[-2.61, 3.22]	-5.60	[-2.80, 3.04]
	61	-6.42	[-2.78, 2.98]	-7.02	[-2.84, 3.32]

Notes: The average treatment effect was estimated as the average of  $\beta$  coefficients in equation (2). The placebo interval represents the interval between 2.5 and 97.5 percent of the distribution of estimates generated by randomization inference. A point estimate outside of that interval indicates a statistically significant effect at the 95 percent level. The sample includes people born between 1920 and 1955.

Table 6. *Nonparametric Estimates of the Treatment Effect on the Cumulative Retirement Rate*

Age	Men			Women	
	Control age	Treatment Effect	Placebo Interval	Treatment Effect	Placebo Interval
62	58	-8.68	[-1.90, 2.83]	-1.10	[-6.51, 3.13]
	59	-8.14	[-2.09, 3.64]	-1.92	[-5.65, 3.43]
	60	-7.00	[-1.43, 3.18]	-1.02	[-6.23, 4.32]
	61	-5.89	[-1.46, 2.31]	-1.59	[-5.50, 3.35]
63	58	-10.01	[-3.16, 3.27]	-0.07	[-6.08, 2.64]
	59	-9.47	[-3.44, 4.02]	-0.89	[-6.16, 1.49]
	60	-8.33	[-2.70, 3.78]	0.00	[-5.99, 2.99]
	61	-7.23	[-2.75, 3.03]	-0.56	[-5.50, 4.40]
64	58	-10.74	[-3.65, 3.57]	0.15	[-7.43, 3.42]
	59	-10.16	[-3.78, 3.53]	-0.74	[-6.46, 3.65]
	60	-9.33	[-3.25, 3.90]	0.08	[-6.21, 3.26]
	61	-8.10	[-3.30, 2.79]	-0.48	[-6.36, 5.40]
65	58	-11.79	[-3.64, 3.22]	-0.69	[-7.05, 2.56]
	59	-11.20	[-4.28, 4.01]	-1.56	[-6.36, 3.94]
	60	-10.43	[-3.94, 4.00]	-0.74	[-7.04, 5.24]
	61	-9.31	[-3.84, 3.21]	-1.39	[-6.97, 4.12]

Notes: The average treatment effect was estimated as the average of  $\beta$  coefficients in equation (2). The placebo interval represents the interval between 2.5 and 97.5 percent of the distribution of estimates generated by randomization inference. A point estimate outside of that interval indicates a statistically significant effect at the 95 percent level. The sample includes people born between 1920 and 1955.

Table 7. *P-values for a Common Trend Test, Men*

Treatment age	Men				Women			
	Control ages				Control ages			
	58	59	60	61	58	59	60	61
<b>Retirement incidence rate</b>								
62	78.79	69.65	91.13	96.98	20.68	8.13	91.14	36.67
63	10.52	31.54	25.99	15.18	65.89	5.86	55.59	38.87
64	51.87	61.98	50.61	50.51	92.88	6.60	63.87	57.32
65	73.58	81.21	50.81	43.14	82.91	7.81	77.87	69.72
<b>Retirement hazard rate</b>								
62	43.43	42.04	73.25	97.19	5.27	97.35	61.99	23.30
63	4.45	32.74	18.40	11.79	81.51	3.40	80.25	1.46
64	22.91	11.70	5.49	13.10	44.82	31.55	73.97	77.30
65	17.14	8.60	2.26	7.46	0.10	9.65	1.60	17.66
<b>Cumulative retirement rate</b>								
62	57.16	74.89	9.70	16.43	39.88	19.14	97.38	98.74
63	1.05	0.76	0.51	17.61	42.18	23.01	92.32	93.82
64	0.01	0.01	0.00	0.66	29.51	17.55	78.35	75.97
65	0.00	0.00	0.00	0.00	82.53	68.18	21.94	48.31

Notes: P-values in percentages for a test of the null-hypothesis that pre-intervention trend is the same for control and treated ages. Control ages are 58 to 61. Treated ages are 62 to 65. Pre-intervention birth cohorts are 1937 or earlier; post-intervention birth cohorts are 1938 or later. For each combination of treated and control ages, a measure of retirement status was regressed on birth cohort dummies interacted with a dummy for treated age and post-intervention cohorts, age dummies, cohort dummies, and linear trend interacted with a dummy for treated ages. A value less than five indicates rejection of the common trend hypothesis.

Table 9. *Parametric Estimates of the Treatment Effect on Retirement Incidence Rate*

Age	Men			Women		
	Control age	Treatment Effect	Placebo Interval	Treatment Effect	Placebo Interval	
62	58	-0.88	[-0.42, 0.40]	-0.31	[-0.25, 0.19]	
	59	-0.94	[-0.42, 0.39]	-0.35	[-0.29, 0.25]	
	60	-0.72	[-0.49, 0.50]	-0.23	[-0.20, 0.20]	
	61	-0.75	[-0.40, 0.45]	-0.42	[-0.18, 0.23]	
63	58	-0.09	[-0.34, 0.29]	-0.07	[-0.26, 0.22]	
	59	-0.14	[-0.26, 0.24]	-0.11	[-0.19, 0.18]	
	60	0.01	[-0.23, 0.27]	0.04	[-0.21, 0.18]	
	61	-0.01	[-0.32, 0.29]	-0.14	[-0.21, 0.27]	
64	58	-0.14	[-0.30, 0.33]	0.10	[-0.18, 0.19]	
	59	-0.18	[-0.25, 0.30]	0.03	[-0.20, 0.23]	
	60	-0.02	[-0.33, 0.28]	0.20	[-0.16, 0.16]	
	61	-0.04	[-0.27, 0.33]	0.04	[-0.27, 0.23]	
65	58	-0.11	[-0.31, 0.33]	0.04	[-0.25, 0.21]	
	59	-0.15	[-0.36, 0.34]	0.01	[-0.26, 0.28]	
	60	0.08	[-0.30, 0.30]	0.15	[-0.33, 0.28]	
	61	0.03	[-0.35, 0.33]	-0.03	[-0.27, 0.28]	

Notes: The average treatment effect is the point estimate of  $\beta$  in equation (3). The placebo interval represents the interval between 2.5 and 97.5 percent of the distribution of estimates generated by randomization inference. The sample includes people born between 1920 and 1955.



RECENT WORKING PAPERS FROM THE  
CENTER FOR RETIREMENT RESEARCH AT BOSTON COLLEGE

**Financial Security at Older Ages**

*Barbara A. Butrica and Stipica Mudrazija, December 2020*

**Do People Work Longer When They Live Longer?**

*Damir Cosic, Aaron R. Williams, and C. Eugene Steuerle, December 2020*

**Do State and Local Government Employees Save Outside of Their Defined Benefit Plans When They Need To?**

*Laura D. Quinby and Geoffrey T. Sanzenbacher, November 2020*

**How Much Taxes Will Retirees Owe on Their Retirement Income?**

*Anqi Chen and Alicia H. Munnell, November 2020*

**A Behavioral Economics Assessment of SSDI Earnings Reporting Documents**

*Denise Hoffman, Jonah Deutsch, and Britta Seifert, November 2020*

**How Accurate Are Retirees' Assessments of Their Retirement Risk?**

*Wenliang Hou, July 2020*

**Is Nontraditional Work at Older Ages Associated with Better Retirement Security?**

*Matthew S. Rutledge and Gal Wettstein, July 2020*

**New Insights on Self-Employment of Older Adults in the United States**

*Joelle Abramowitz, July 2020*

**What Jobs Do Employers Want Older Workers to Do?**

*Alicia H. Munnell, Gal Wettstein, and Abigail N. Walters, June 2020*

**Technological Innovation and Labor Income Risk**

*Leonid Kogan, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Jae Song, June 2020*

**Are Older Nontraditional Workers Able to Find Health and Retirement Coverage?**

*Matthew S. Rutledge, March 2020*

**Employer Perceptions of Older Workers – Surveys from 2019 and 2006**

*Alicia H. Munnell and Gal Wettstein, March 2020*

**Are Homeownership Patterns Stable Enough to Tap Home Equity?**

*Alicia H. Munnell, Abigail N. Walters, Anek Belbase, and Wenliang Hou, January 2020*

*All working papers are available on the Center for Retirement Research website  
(<https://crr.bc.edu>) and can be requested by e-mail ([crr@bc.edu](mailto:crr@bc.edu)) or phone (617-552-1762).*