NBER WORKING PAPER SERIES

HOW STICKY IS RETIREMENT BEHAVIOR IN THE U.S.? RESPONSES TO CHANGES IN THE FULL RETIREMENT AGE

Manasi Deshpande
Itzik Fadlon
Colin Gray

Working Paper 27190
http://www.nber.org/papers/w27190

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020

We thank John Phillips, Jason Brown, Natalie Lu, Ted Horan, Mark Sarney, Lynn Fisher, and Linda Martin of the Social Security Administration for making this work possible and providing access to data. The authors are grateful to the Ronzetti Initiative for the Study of Labor Markets at the Becker-Friedman Institute for financial support. This research was supported by the U.S. Social Security Administration through grants #RDR18000003 and #DRC12000002 to the National Bureau of Economic Research as part of the SSA Retirement and Disability Research Consortium. The findings and conclusions expressed are solely those of the authors and do not represent the views of the Social Security Administration, any agency of the Federal Government, or the National Bureau of Economic Research. Jesse Kozler provided excellent research assistance.

How Sticky is Retirement Behavior in the U.S.? Responses to Changes in the Full Retirement Age
Manasi Deshpande, Itzik Fadlon, and Colin Gray
NBER Working Paper No. 27190
May 2020
JEL No. H31,H55,J26

ABSTRACT

We study how increases in the Social Security full retirement age (FRA) affect benefit claiming and retirement behavior, and specifically the interaction between these two choices. Using Social Security administrative data, we implement complementary research designs of a traditional cohort analysis and a regression-discontinuity design. We find that while increases in the FRA strongly and immediately shift claiming ages, retirement ages exhibit persistent "stickiness" at the old FRA of 65. We use several strategies to explore the likely mechanisms behind the stickiness in retirement, and we find suggestive evidence of a role for employers in individuals' responses to the FRA.

Manasi Deshpande
Department of Economics
University of Chicago
1126 E 59th St.
Chicago, IL 60637
and NBER
mdeshpande@uchicago.edu

Itzik Fadlon
Department of Economics
University of California, San Diego
9500 Gilman Drive, #0508
La Jolla, CA 92093
and NBER
fadlon@ucsd.edu

Colin Gray
Wayfair
cgray3@wayfair.com

1 Introduction

Social Security benefits are a crucial component of the income of older Americans. The "benchmark" retirement ages within the program are believed to play a key role in shaping American workers' labor supply decisions, as evidenced by the well-known spikes in claiming and retirement at the early and full retirement ages. The increase in the Social Security full retirement age (FRA) through the 1983 Social Security Act Amendments is one of the most prominent changes to retirement policy in the U.S. in recent decades, and several Social Security reform proposals recommend further increases in the FRA in the future (U.S. Congressional Budget Office 2015). The rationale behind these actual and proposed FRA increases is twofold: first, mitigating the fiscal burden on the program by direct reduction in benefits; and, second, inducing later retirement by older American workers whose health capacity to work longer has substantially increased (Coile et al. 2016).

Despite the importance of this reform, there is markedly limited evidence on the effect of the increases in the U.S. Social Security full retirement age on the labor supply and benefit claiming of older Americans. We also know very little about the relationship between households' responses within these two potentially distinct behaviors of claiming benefits and retiring from the labor force. Understanding these responses is important for assessing the degree to which FRA reforms are successful in achieving their goals—of reducing the program's costs and inducing later retirement—and is hence central for the optimal design of Social Security. Moreover, studying the responses to FRA changes is important for understanding what governs Americans' retirement choices, as well as the degree of flexibility in the U.S. labor market for older workers, and it can help shed light on the various suggested explanations for the retirement spikes at the Social Security retirement ages. Indeed, in their well-cited work, Lusdaine, Stock and Wise (1996) write: "The unexplained age-65 spike is important because it limits our ability to predict the effect of potential policy changes, like the planned increase in the Social Security normal retirement age from age 65 to age 67. Would there then be a spike at 67, or would it remain at 65?"

In this paper, we study how American workers respond to the increase in the Social Security FRA, with a particular focus on the potential gap between labor supply and benefit claiming. Specifically, using two complementary research designs, we exploit variations induced by the phased-in FRA increase from the Social Security Act Amendments of 1983 to study age of labor force exit versus retirement benefit claiming across cohorts. Our analysis leverages a long panel of Social Security administrative data stretching back several decades, which provides us with several distinct advantages compared to previous work. First, while previous work has relied on self-reports from survey data, data from the Social Security Administration (SSA) provide accurate and objective administrative measures for both benefit claiming and labor supply. Second, the SSA data provide access to large-scale samples, which allows us to leverage various data-intensive research designs. Third, the time horizon and the detailed information in the datasets we use allow us to investigate possible underlying mechanisms: learning and adjustment costs (as the data cover various pre-reform and post-reform cohorts), geographical location, and the potential role of employers and workplaces.

We begin with a traditional cohort analysis that studies the behavior of American workers around key ages in the Social Security system. We first replicate, for pre-reform cohorts, the well-documented spikes at the Early Eligibility Age (EEA) of 62 and at the old Full Retirement Age (FRA) of 65, in both benefit claiming and labor force exit. We then show that, for the post-reform cohorts whose FRA is 66, the spike in claiming has almost fully moved from the old to the new FRA, whereas the spike in retirement displays "stickiness" at the *old* FRA of 65 with a similar magnitude to the pre-reform cohorts. This phenomenon is robust across different variable definitions and data frequencies, and it is pervasive across subpopulations when we split the data by earnings quartiles and gender.

This descriptive evidence suggests that retirement timing does not shift in lock step with claiming timing when the FRA increases. We next conduct a complementary analysis to identify the causal effect of the FRA increase on claiming and retirement behavior. For cohorts born between 1938 and 1942, the FRA has increased from 65 to 66 in two-month increments per cohort at the January birthdate cutoff. Using a regression-discontinuity design based on this variation, we find that these increases have led to meaningful increases in claiming ages, but have not induced statistically detectable changes in labor supply around the old FRA. This discontinuity result is consistent with our findings from the cohort analysis that claiming age adjusts fully and immediately to increases in the FRA, while retirement behavior exhibits considerable stickiness.

We investigate potential drivers of sticky retirement behavior. Previous studies on Social Security's FRA have been constrained in investigating mechanisms due to limited samples and time horizons and insufficient information on geography and employers. We use the scope and richness of the SSA administrative data to provide new empirical analysis and tests of potential explanations for stickiness in retirement. We first test for evidence of adjustment frictions and learning by studying the responses of individual cohorts. We find limited evidence that later cohorts differentially change their retirement behavior in response to FRA increases, though more cohorts will need to reach retirement age for a definitive conclusion. Specifically, we show that similar-magnitude retirement rates persist at the old FRA for the range of affected cohorts included in our data (spanning 10 birth-years), which is broadly inconsistent with retirement stickiness being primarily driven by delays in adjustment or gradual learning.

We next test whether access to old-age health insurance through Medicare could explain the stickiness in the age 65 retirement spike. We show that workplace-level employee behaviors at the old and new FRA do not vary across firms whose benefit plans, in particular health insurance, cover retirees. Our finding of a limited role for health insurance is consistent with previous literature suggesting that the age 65 Medicare eligibility is unlikely to explain the retirement spike (Lusdaine et al. 1996).

Finally, we investigate the role of geography and firms. We employ a "movers design" across locations and employers. We find some evidence for a role of location in benefit claiming, but not in retirement. For movers across employers, however, we find strong associations with workplace-level behaviors, in both claiming and retirement around the pivotal full retirement ages. This provides suggestive evidence that the workplace could have a meaningful mediating role in responses to FRA policies, through, e.g., retirement norms, demand side factors, employer-based pensions, and the norms related to their withdrawal.

Our paper makes several contributions. First, overall, our findings have both positive and normative implications for retirement behavior and choices in the U.S. The results challenge the common motivation for increases in the FRA: that they will lead to considerable increases in both claiming age and retirement age. Specifically, we have found that the FRA increases were successful in achieving one goal—reducing costs—but have been much less successful in achieving the other stated goal of inducing later retirement. These findings point to meaningful labor market frictions among older American workers. We find evidence for considerable "stickiness" in retirement behavior in the U.S. labor market, and a potential role of the workplace in attenuating the flexibility of adjustment in labor supply choices in response to FRA increases. Normatively, our findings provide evidence on the degree of efficiency of the FRA policy reform and offer some optimal design considerations for future policy changes. Our analysis suggests that if policymakers want to achieve the goal of delaying retirement age, other complementary policy levers, such as employer-based policies and incentives, may be necessary.

Second, we provide evidence on age eligibility for government retirement benefits and older workers' claiming and labor supply choices in the context of the United States, which has distinct institutions from other countries in which this question has been studied. Some recent and concurrent work on responses to variation in the EEA and FRA and potential explanations in countries other than the U.S. includes studies in Austria, Switzerland, and Germany (e.g., Staubli and Zweimüller 2013; Lalive and Staubli 2015; Manoli and Weber 2016; Lalive et al. 2017; Seibold et al. 2019). This work finds strong shifts in labor supply in response to FRA changes with lockstep movements in retirement spikes. In contrast, we find almost no shift in labor supply and retirement spikes in response to FRA changes in the United States.

Finally, in terms of evidence on the United States, the SSA administrative data allow us to study claiming and retirement behaviors with unprecedented accuracy and precision. Our paper is most closely related to the few previous papers that have made important contributions in understanding the effects of FRA changes in the United States. Mastrobuoni (2009) uses the Current Population Survey to estimate the effect of increasing the FRA on the age of labor force exit, and Song and Manchester (2007) use SSA administrative data to study the effect of the FRA increase on benefit claiming. Behaghel and Blau (2012) use the Health and Retirement Study to analyze the effect of the FRA increase in the 1983 Amendments on benefit claiming and labor supply. However, no paper to date has used SSA administrative data to study the effects of the FRA increase on both labor supply and benefit claiming and the relationship between the two. As we discuss above, the scope, richness, and scale of the SSA administrative data offer a number of advantages compared to existing work. The data allow us to study objective measures of claiming and labor supply and separately analyze these behaviors with breakdowns by cohorts and demographics that maintain high degree of precision. We use employer and geographic information and the time horizon to tap into previously unexplored mechanisms of learning, location, and the workplace. Using our data and designs, we provide new analysis and findings about the effects of the Social Security FRA changes, their underlying channels, and novel evidence of stickiness in retirement behavior in the U.S.

The remainder of the paper proceeds as follows. Section 2 describes the datasets and the institutional background. Section 3 presents our analysis of the effects of the increases in the Social Security full retirement age: the cohort analysis in Section 3.1, the regression discontinuity analysis in Section 3.2, and the investigation of potential mechanisms in Section 3.3. Section 4 discusses the implications of our findings and concludes.

2 Data and Institutional Background

2.1 Data and Variable Definitions

We use administrative data from the Social Security Administration's Continuous Work History Sample (CWHS), which is a panel of earnings and other economic data for a 10 percent sample of SSNs. The CWHS allows us to merge multiple files. First, the Summary Earnings Record (SER) contains data on annual earnings in each calendar year taken from Form W-2. Second, the Master Beneficiary Record (MBR) contains the calendar month in which every individual claimed Old Age benefits. The differential data frequency—annual for earnings and monthly for benefit claiming—necessitates adjustments to the particular specifications we use for each of our studied behaviors of labor supply and claiming. To supplement these data sources, we merge in additional demographic data, including date of death, from the Numident file and monthly Old Age payment histories from the Payment History Update System (PHUS) files.

Our main sample consists of primary beneficiaries who were born between 1935 and 1945.

To prevent differential censoring across birth cohorts, we keep beneficiaries who claim Old Age benefits at some point before the age of 70. This accounts for over 99% of primary beneficiaries, since claiming after age 70 is extremely rare.

In the first part of our analysis, we study benefit claiming and retirement by age. We analyze "flow" outcomes, which are defined according to traditional definitions in the literature on retirement in the following way. For benefit claiming, we analyze indicator variables for the age at which claiming began. The variable "claim at age a" assumes the value 1 if the individual had no benefits in the previous period but had positive benefits in the current period, and it assumes the value 0 otherwise. For retirement, the variable "retire at age a" assumes the value 1 if the individual has positive earnings in the current period but no earnings in the next period, and it assumes the value 0 otherwise. We also study an alternative absorbing-state measure, where retirement is defined as the last year in which earnings in the next period are zero but earnings in the current year are positive.

Second, when constructing regression discontinuity estimators we take the following steps. We collapse the data to capture average claiming ages (censored at 70) for different day-of-birth cohorts, where the data frequency allows us to use "monthly"-level ages (e.g., 65 and 2 months). For earnings, we are constrained by the annual frequency of the data. Specifically, earnings at age a (at an "annual" frequency, e.g., 65 rather than 65 and 2 months) are defined as earnings at the end of the year in which affected individuals (above the birthdate cutoff) turned age a.

We use the Detailed Earnings Record (DER) to study movers across workplaces or across geographical locations. For workplaces, we first define each worker's firm in each year as the Employer Identification Number (EIN) in which the worker earned the most in that year. Second, we classify a worker as a "mover" if they switched firms between the ages of 50 and 60 (inclusive). Third, we characterize the behavior of other potential beneficiaries in each firm, to capture leave-out means. To do so, we consider the average behavior of everyone at the firm who is *not* classified as a mover and for whom that was their primary firm in the calendar year that they turned 60. To construct the movers design for locations, which we define as counties, we follow a similar procedure but use a 1% sample of person-year county geocodes that are available in the CWHS.[1]

To study the role of health insurance, we use a supplementary dataset from Form 5500 called "Annual Return/Report of Employee Benefit Plan." This form is filed by employers to report information about their employee benefit plans (if they cover 100 or more participants) as required by the Employee Retirement Income Security Act of 1974 (ERISA) and the Internal Revenue Code. We use publicly available data from the years 1992-1998 and construct an indicator for whether an employer's plan covers retirees, which includes former employees who are receiving group health continuation coverage benefits. We merge these data with our SSA datasets based on EINs.

2.2 Institutional Details

The original Social Security Act of 1935 set the age for receiving full retirement benefits—i.e., the Full Retirement Age (FRA)—at 65. The Social Security Amendments of 1956 enacted the possibility of receiving reduced benefits as early as age 62 for women only. Benefits received prior to age 65 were reduced to account for the longer period over which they would be received. The 1961 amendments extended eligibility for reduced benefits to include men. Sixty-two has been the Early Eligibility Age (EEA) since.

The 1983 Amendments—the core variation for our analysis—phased in a gradual increase in the Full Retirement Age, for which Congress cited improvements in the health of older people and increases in average life expectancy as primary motivations. The amendments specify increases in the FRA from 65 to 67 over a period of 22 years in the following way: 65 for those born in 1937 or earlier (who we call *"pre-reform"* cohorts); increasing by 2 months for every birth cohort from 1938 to 1942 (who we call *"phase-in"* cohorts); 66 for birth cohorts 1943–1954 (who we call *"post-reform"* cohorts); and increasing in 2-month increments per cohort from 1955–1960 until 67 (data not yet available). Our cohort analysis focuses on the "pre-reform" and "post-reform" cohorts so as to use rounded FRAs (given the annual data). Our regression discontinuity analysis leverages the 2-month increments in FRA based on January birthdate thresholds that are used by SSA for determining an individual's FRA.

[1]County geocodes are also reported from Form W-2 that covers only individuals with positive earnings. For each person-year we therefore keep all county codes that correspond to the place of residence rather than the place of employment.

This analysis includes the "phase-in" cohorts with years of birth 1938–1942 as well as the 1937 and 1943 boundary cohorts.

Social Security retirement benefits accrue to individuals whose earnings are subject to Social Security taxes. Generally, to become eligible for retirement benefits, individuals are required to accumulate 40 "credits," which translates to 10 years of work since workers can earn up to 4 credits each year (where, e.g., in 2016, $1,260 in earnings = 1 credit).[2] For this reason, our analysis focuses on the sample of individuals who have 10 years of earnings prior to reaching age 60.

As a benchmark, Appendix Figure 1 displays the change in benefit amounts by claiming age, for the different cohorts. It takes into account the reduction rate in benefit amounts if benefits are claimed prior to the FRA (and after the EEA), as well as the Delayed Retirement Credit (DRC) by which benefits increase at a certain percentage if claiming is delayed beyond the FRA and up to age 70. The benefit amount schedule coincides for pre-reform and post-reform cohorts only if claiming occurs at age 70 or later, and it is lower for post-reform cohorts at any other claiming age. As claiming after 70 is rare, this change represents an overall decline in lifetime unearned income. Economic theory would therefore predict that an income effect from this change should postpone retirement (if leisure is a normal good), and in particular, there should be no discrete retirement behavior around the old FRA of 65 driven by the Social Security system itself among post-reform cohorts. We note that the claiming age determines the SSA retirement benefit annuity amount for life. This underscores the importance of also studying changes in benefit claiming from the FRA reform, on top of studying responses in Americans' retirement choices.

3 Empirical Evidence: Responses to Changes in Full Retirement Age

We conduct a series of empirical analyses of the effects of the increases in Social Security's FRA on the benefit claiming and retirement behaviors of American workers. In the first approach, we use a cohort analysis to study the distribution of claiming and retirement

[2] Retirement benefits aim to reflect life-time earnings and are based on a worker's Average Income Monthly Earnings (AIME) over the 35 years in which the worker earned the most.

by age. In the second approach, we employ a regression discontinuity design of 2-month increases in the FRA at the January 2 birthday cutoff. Then, we investigate mechanisms by studying changes in behaviors across a range of individual cohorts and by conducting mover designs with respect to employer and geography.

3.1 Cohort Analysis over the Life Cycle

In the cohort analysis, we plot the benefit claiming and retirement behaviors of older American workers by age for the differentially affected birth cohorts. We test two hypotheses. The first is that cohorts affected by a higher FRA increase their claiming age. Note that an increase in claiming age is not a mechanical effect of the FRA increase, since individuals could respond by claiming early (starting at the EEA). The second hypothesis is that these cohorts also increase their labor supply and retirement age. These hypotheses form the dual motivation for the policy increases in the FRA.

Claiming. Our first outcome of interest is the claiming of SSA retirement benefits, presented in Figure 1. Panel A plots the fraction of the pre-reform and post-reform birth cohorts claiming at certain ages. In this and subsequent graphs, the y-axis captures the behavioral outcome of interest and the x-axis displays individuals' age. Individuals are split into pre-reform cohorts whose FRA is settled at 65 (1935-1937) and post-reform cohorts whose FRA is settled at 66 (1943-1945). In some graphs, we include the phase-in cohorts (1938-1942) whose FRA increases in 2-month increments per cohort from 65 to 66.

Panel A of Figure 1 first replicates the known spikes for pre-reform cohorts in benefit claiming at the EEA, when retirement benefits first become available (at reduced rates), and then at 65, the FRA when eligibility for full benefits begins. Second, the plot for post-reform cohorts clearly shows how claiming behavior has followed the FRA changes: the large spike at the EEA remains whereas the mass of Americans claiming at the FRA shifts almost one-for-one from the old FRA of 65 to the new FRA of 66.

Panel B displays a higher-frequency analysis by leveraging monthly-level claiming information in SSA data. The figure displays claiming rates at "monthly ages" for phase-in and post-reform cohorts. It shows similar first spikes right after the EEA for all cohorts, and second spikes right after each cohort's particular FRA at their monthly level (65 and 2 months for 1938; 65 and 6 months for 1940; 65 and 10 months for 1942; and 66 for 1944, 1946, and 1948), which move in increments similar to the change in FRA across cohorts.

To investigate heterogeneity in claiming behavior, panel C divides the sample by baseline earnings quartiles based on average earnings over ages 25-55. The figure first shows that the patterns of shifts from the old to the new FRA is similar across quartiles, and it also shows that both before and after the reform, the spikes at the FRA are driven by high earners. Panel D splits the sample by gender, and reveals similar responses for men and women with some evidence of larger spikes at the FRA for males.

Together, the panels in Figure 1 indicate that claiming responds immediately and fully to increases in the FRA, and this pattern is strong across observable subgroups.

Retirement. Next, we study the effects of the FRA increases on retirement choices using administrative earnings data by year. Figure 2 studies the retirement behavior of pre-reform and post-reform cohorts in several ways.

First, panel A plots annual retirement flows based on the following traditional measure: an individual is assigned the value 1 if they have positive earnings today and zero earnings in the following period. As a benchmark, the figure shows the well-known spikes for pre-reform cohorts at the EEA and the 65 FRA. For the post-reform cohorts, two patterns stand out. First, consistent with economic theory predictions regarding the induced decline in lifetime unearned income (as manifested in Appendix Figure 1), there is an underlying shift in the distribution toward later retirement. Specifically, there is reduced mass at the EEA of 62 (and at 63) and greater mass at later ages.[3]

The second—and unexpected—pattern that emerges in panel A of Figure 2 is the persistence of the retirement spike at the old FRA of 65 with no spike at the new FRA of 66. More

[3]Could this shift in the distribution reflect a mere time or cohort change? Figure 3 suggests that this is not likely the case. Panel C of Figure 3 plots the cumulative distribution of labor force participation for a range of cohorts. It shows that, particularly in the critical domain of ages 62-66, the distributions for pre-reform cohorts (1935-1937) who all have the same FRA of 65, line up on top of each other, with no apparent "cohort" effect. The shift in the distribution begins with the phase-in cohorts whose FRA increases incrementally (1938-1942). Then, for post-reform cohorts (1943-1945) who all have the same FRA of 66, the distributions again line up on top of each other, with not apparent cohort effect. These patterns are consistent with a modest shift toward later retirement induced by the change in FRA.

so, the magnitude of the spike at 65 is similar across pre-reform and post-reform cohorts. Recall that for post-reform cohorts there is no discontinuous economic incentive or policy anchoring within the Social Security system at the old FRA of 65.

Two alternative measures of retirement, in panels B and C, verify this pattern. Panel B provides a higher-frequency analysis based on annual earnings data and monthly ages at the end of the calendar year. The figure reaches a similar conclusion: a reduction in mass at age 62 and an increase in mass at ages 67 and 68, but surprisingly little change in match at the old FRA at 65. Panel C of Figure 2 uses a similar definition for retirement, but based on the last data period in which such a transition in labor supply applies, which captures retirement in a different manner. The unexpected pattern of a considerable local spike at the old FRA of 65 with no local spike at the new FRA of 66 prevails when we use this measure as well.[4]

To better understand this phenomenon and whether some subgroups may drive the results, we again study heterogeneity by average earnings from ages 25-55 (panel D) and by gender (panel E). Generally, with substantial underlying cross-sectional heterogeneity across them, all different subgroups display similar response patterns of some shifting underlying mass in retirement behavior along with persistent local peaks at age 65 which are comparable in magnitude across pre-reform and post-reform cohorts.[5]

A key takeaway from this analysis is that, while the bunching around the FRA evolves in tandem with increases in the FRA for benefits claiming, changes in labor force exits exhibit a complex pattern. Most puzzling is the apparent stickiness of the age 65 retirement spike, which persists for affected cohorts despite the fact that age 65 is no longer an anchor (in terms of either economic incentives or behavioral inclinations) within the Social Security schedule. We analyze potential mechanisms through which this stickiness operates in Section 3.3.

[4]Unlike our main measure of retirement, this absorbing-state measure is subject to differential censoring across cohorts whose age ranges of available data differ. Therefore, the focus in panel C is not a cross-cohort comparison but rather a within-cohort observation of persistent local spikes at the old FRA.

[5]The overall underlying upward shift in the distribution (from 62 to over 66) seems to be mostly driven by higher earnings individuals.

3.2 Regression Discontinuity

The cohort analysis provided descriptive evidence of the change in the distribution of retirement and claiming behavior after the FRA increase. We next identify the causal effect of the FRA increase on average claiming age and labor supply. We are interested in particular in the causal effect of the FRA increase on retirement, given the evidence of stickiness in retirement behavior from the cohort analysis. We use a regression discontinuity (RD) design that exploits the 2-month increments of FRA increases at the January 2nd birthdate thresholds for the phase-in cohorts in the Social Security Act Amendments of 1983.

The RD design complements the cohort analysis. Because the cohort analysis uses large changes, it is appropriate for identifying aggregate or long run trends. In the presence of frictions, larger changes would more likely induce responses since the costs of misoptimizing are larger. The full-year FRA changes also make it easier from a statistical perspective to pick up responses in labor supply (which is measured annually). In contrast, the RD uses two-month increases in the FRA. The advantage of the RD relative to the cohort analysis is that the causal effect of the FRA is more cleanly identified and it is less likely to be confounded by potential time or cohort effects. However, the RD use of only small increases in the FRA makes it harder to detect effects on claiming and retirement behavior, especially if there are optimization frictions.

Our RD design uses the standard estimating equation of the following form:

$$Y_i = \alpha + \beta \mathbb{1}\{\text{DOB}_i > 0\} + \gamma \text{DOB}_i + \delta(\text{DOB}_i \times \mathbb{1}\{\text{DOB}_i > 0\}) + \varepsilon_i. \quad (1)$$

Here, Y_i is an outcome such as claiming age or earnings for individual i. The running variable is DOB_i, which is the individual's date of birth (in days) relative to January 2nd, the date at which the FRA increases by 2 months. The indicator $\mathbb{1}\{\text{DOB}_i > 0\}$ is turned on for individuals whose date of birth is after January 2nd. The coefficient of interest is β, which gives the discontinuity in the outcome variable around the cutoff. We interpret this coefficient as the causal effect of the 2-month increase in the FRA on benefit claiming or labor supply.

Figure 4 provides a visual representation of behaviors around the cutoff, and the corresponding regressions that estimate the RD coefficient β are provided in Table 1. Panels A-B of Figure 4 and panel A of Table 1 analyze the claiming behavior. First, the monthly claiming age displays a meaningful increase: a 2-month increase in FRA leads to .84-month average increase (=.07×12) in the age at claiming. Second, narrowing the analysis to the vicinity of the actual monthly FRAs reveals sharp responses. For every cohort stacked in our RD design, we study how claiming at the new FRA (attributed to those just above the cutoff) changes around the treatment discontinuity. The propensity to claim benefits at the exact vicinity of the new FRA—within a 2-month range—increases by 12 pp (see panel B of Figure 4 and column 2 in Table 1.A). This underscores changes in average claiming age that are exactly at the locality of the new FRA.

We turn next to studying wage earnings. Consistent with our findings in the previous section and in contrast to claiming, there is much less movement in labor supply than in claiming. Since the labor supply data are annual but the FRA increases are in two-month increments, we study the continuous measure of earnings at the calendar year at which just treated individuals turned a given ("annual") age, and we focus on the anchoring ages relevant to the policy change, i.e., ages 65 and 66. Panels C-D of Figure 4 and panel B of Table 1 show that there are no detectable responses in labor supply. The estimates allow us to reject earnings increases of more than \$292 for age 65 and more than \$312 for age 66. A useful benchmark to gauge these magnitudes is the expected increase if earnings were to respond as flexibly as claiming, i.e., by a .84-month equivalent. Based on baseline annual earnings (using the constant terms in Table 1.B), these amount to \$1,154 for age 65 and \$955 for age 66, which are much larger than the upper bounds of our estimates.

The cohort analysis and RD analysis present similar stories of responses in claiming and retirement to increases in the FRA. While benefit claiming closely moves with the exact FRA, labor supply displays much less movement and, in particular, it displays no change around the old FRA. We next attempt to shed light on the underlying forces that can drive this result.

3.3 Potential Mechanisms

Having established the stickiness of retirement behavior relative to claiming behavior, we investigate potential mechanisms behind this phenomenon. We are specifically interested in trying to explain the responses, or lack thereof, around the critical new and old full retirement ages (65 and 66) in both claiming and retirement.

Learning and Delayed Adjustment. A natural question is whether delays in adjustments and gradual learning can account for the persistent spike at the old FRA. Since the FRA increases were announced in 1983, affected workers received notice of the change decades in advance of retirement. If they were fully informed and forward looking, then learning could not explain the results. Still, it could be that learning and its corresponding impact begin when changes actually take place and the policy comes into effect (rather than in anticipation of it); that is, as cohorts mature into the new FRA in practice. In that case, if individuals can more easily postpone their claiming (e.g., through the use of private savings), but more time is needed for changes in the retirement decision (which could also depend on the the employer side), then delays in adjustment may provide some explanation. To test this hypothesis, we look separately at the various post-reform cohorts in panel F of Figure 2. Generally inconsistent with this potential explanation, the peaks at the old 65 FRA seem similarly persistent across post-reform cohorts up to the last cohort in our data for which we can observe behaviors at age 65 (those born in 1948, where the first cohort whose FRA has changed is 1938).

Medicare Eligibility. Another candidate explanation for the persistent retirement spike at 65 could be eligibility for health insurance through Medicare . Using data from the Form 5500 on retiree health coverage linked by EIN to our CWHS data, we run regressions of workplace-level employee behaviors on benefit plan characteristics among matched workplaces. The analysis and additional details are provided in Appendix Table 1. The outcomes include SSA benefit claiming and retirement behaviors, among pre-reform cohorts and using the difference between post-reform and pre-reform cohorts, at both the old FRA and the new FRA. The right-hand side variables include an indicator for whether the employer's benefit plan covers retirees and a range of other plan characteristics. In these regressions, we find no evidence of any partial correlation between the provision of retiree benefits and employee behaviors. Our findings are consistent with Lusdaine et al. (1996), who compare employees with and without employee-provided retiree health insurance and find that high age-65 retirement hazard rates cannot be explained by Medicare eligibility.

Geography and Employer Roles. Our data allow us to analyze the role of two additional potential channels that are likely candidates: geographical location and employers. First, recent studies have highlighted the importance of one's location in determining many important aspects over the life-cycle, including educational attainment, future earnings and income, intergenerational mobility, healthcare utilization, and mortality (see, e.g., Chetty and Hendren 2018a, 2018b and Finkelstein et al. 2016). It is therefore natural to extend these findings to the investigation of benefit claiming and retirement choices. Second, directly related to retirement choices, Stock and Wise (1990) have found an important role for employer-provided pensions in individuals' retirement behavior. Therefore, the extent to which employees respond to the FRA, and to changes in it, could be mediated by their workplaces. Workplaces could be viewed as a "composite" of channels that could potentially encompass workplace retirement norms, demand side factors (such as effective upper age bounds for employment from the employer's standpoint), employer-based pensions and retirement savings accounts, and norms related to their withdrawal.

To investigate the potential role of these two channels, we employ a "movers" design that has become prevalent in empirical work to study mechanisms that could drive observed variation in behaviors across individuals.[6] Our ability to conduct this analysis, as compared to work using survey data, is attributable to the richness of the administrative data (which include linkages across individuals, locations, and employers) and the scale of the data (which allows us to identify a sufficient number of movers from the specific post-reform cohorts).

[6] For example, among others, Finkelstein et al. (2016) use migration of Medicare patients to study the drivers of geographic variation in U.S. healthcare utilization, and Chetty et al. (2014) use employee switches across firms to study the role of employer contributions to retirement savings accounts on Danish individuals' overall retirement savings.

We use the following specification:

$$Y_{iod} = \alpha_o + \beta \text{StayerPre}_d + \gamma \text{StayerDelta}_d + \varepsilon_{iod}. \quad (2)$$

Here, Y_{iod} is an outcome like claiming or retiring at 65 for individual i from a post-reform cohort who moved from origin o to destination d between the ages of 50 and 60, inclusive. The α_o are origin fixed effects, StayerPre_d is the outcome variable (e.g., fraction retiring at 65) for stayers in unit d, and StayerDelta_d is the change in the outcome variable for stayers in unit d before and after the FRA increase (that is, the difference in behavior across pre-reform and post-reform cohorts).

In an alternative specification, we replace the origin fixed effects with right-hand side variables similar to those for destination but for origin along with a constant term:

$$Y_{iod} = \alpha + \beta_0 \text{StayerPre}_o + \gamma_0 \text{StayerDelta}_o + \beta_1 \text{StayerPre}_d + \gamma_1 \text{StayerDelta}_d + \varepsilon_{iod}. \quad (3)$$

The alternative specification provides some useful benchmarks for considering magnitudes. The constant term provides a baseline that captures the average behavior of movers from post-reform cohorts whose origin workplaces display 0 "intensity" level. For example, in the analysis of claiming at age 65, this would translate to mean claiming rate at 65 of movers from workplaces where no employee starts claiming at 65 and where post-reform and pre-reform cohorts do not display differential claiming behavior at that age. The regression coefficients on the destination variables then correspondingly measure how an individual's own behavior varies when moving from "0 intensity" origin (where no one displays that behavior) to a "1 intensity" destination (where everyone displays that behavior).

For this design and specifications to carry a causal interpretation, the identifying assumption is that, conditional on their origin, individuals who move to destinations with differential intensities in behavior do not systematically differ in their underlying potential outcomes. Our analysis is motivated by this possible causal interpretation, but, unlike other applications of mover designs, our context does not allow us to test the necessary condition of parallel pre-trends. This is because our outcomes of interest are one-time events that lead into absorbing states (e.g., claiming and retirement), so these behaviors cannot be studied in the pre-move periods. Establishing the causality of our mechanism analysis and additional work that aims to identify the role of location and employers in claiming and retirement choices could be a fruitful route for future research. Our goal here is merely to investigate whether the patterns are consistent with a role for employers or geography. Our analysis essentially studies the degree to which an individual's behavior is associated with the degree of "treatment" intensity induced by the move, in terms of the outcome of interest, when the individual relocate to a higher- or lower-intensity destination unit as compared to their origin unit.

Table 2 presents the results, with movers across geography in Panel A and with movers across employers in Panel B. We find evidence for a potential role for geographic location in claiming behavior, specifically, in terms of claiming at the new FRA of 66. We find no evidence for a role of geography in retirement behavior. In terms of employers, we find patterns that are consistent with a role for workplaces in mediating Americans' responses to the FRA, in both benefit claiming and retirement choices. Across specifications, ages, and outcomes, there are strong associations between the mean behavior of employees in the workplace to which an individual switches and the individual's own claiming and retirement behavior.

Consider the retirement behavior at age 65 of movers from post-reform cohorts (whose FRA is 66). Column 6 of Table 2.B suggests that moving from a 0 to 1 intensity workplace in terms of employees' propensity to retire at 65 is associated with a 8.4 pp higher propensity to retire at the old FRA for an individual from a post-reform cohort. The estimates also suggest that relative to someone who moves to a workplace where the average retirement age increased following the FRA change, someone who moves to a workplace where the average retirement age did not increase is 7.4 pp more likely to retire at age 65. These coefficients are larger than 100% of the baseline level (captured by the constant term) of a 6.4 pp retirement rate at the old FRA of 65. Overall, the results suggest a potentially important role for employers in shaping and mediating retirement decisions as they relate to Social Security's full retirement age.

4 Discussion and Conclusion

Using Social Security administrative data, we study how increases in the full retirement age (FRA) affect benefit claiming and retirement behavior. We find that claiming behavior closely follows the FRA, but that retirement behavior displays considerable "stickiness" that persists over time.

This finding has two main implications. First, from a positive perspective, it points to meaningful frictions in the labor market and in Americans' retirement choices which, by comparison to contexts of previous work (e.g., Gelber et al. 2020 on the earnings test), display persistence throughout a rather long period of post-reform cohorts. As we find that some degree of stickiness in retirement remains over time, explanations such as adjustment costs, learning, or necessary time to prepare (e.g., via personal plans or negotiations with employers) seem unlikely in our context.

Second, from a policy standpoint, our results suggest that merely increasing the FRA is not likely to efficiently achieve the stated goal of inducing later retirement by Americans. That is, we find that changing the FRA in isolation is not enough to move people away from the age 65 spike in retirement. If policymakers want to encourage delayed retirement, they may need to combine the FRA increase with other policy levers.

The diverging patterns in the responses across claiming and retirement underscores the importance of analyzing these behaviors separately moving forward, in contrast to earlier research that has assumed their concordance. Moreover, both of these implications call for investigating potential underlying mechanisms that may drive retirement stickiness. This would allow for better understanding of the labor market for older Americans and what governs their retirement choices, and could also point to the policy tools needed in combination with future increases in the FRA for inducing later retirement. Our initial findings point to a potentially important role for employers, which can include retirement norms, demand side factors, as well as employer-based pensions and norms in initiating their withdrawal. Further analysis that identifies the causal channels by which employers may affect responses to the FRA could be an important route for future research.

References

Behaghel, L. and D.M. Blau, "Framing Social Security Reform: Behavioral Responses to Changes in the Full Retirement Age," *American Economic Journal: Economic Policy*, 2012, *4* (4), 41–67.

Chetty, R., J.N. Friedman, T. H. Nielsen, and T. Olsen, "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," *Quarterly Journal of Economics*, 2014, *129* (3), 1141–1219.

Chetty, Raj and Nathaniel Hendren, "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects," *Quarterly Journal of Economics*, 2018, *133* (3), 1107–1162.

_ and _ , "The Impacts of Neighborhoods on Intergenerational Mobility II: County-Level Estimates," *Quarterly Journal of Economics*, 2018, *133* (3), 1163–1228.

Coile, C., K. Milligan, and D.A. Wise, *Health Capacity to Work at Older Ages: Evidence from the United States*, University of Chicago Press, 2016.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams, "Sources of Geographic Variation in Health Care: Evidence from Patient Migration," *Quarterly Journal of Economics*, 2016, *131* (4), 1681–1726.

Gelber, Alexander M., Damon Jones, and Daniel W. Sacks, "Estimating Adjustment Frictions Using Nonlinear Budget Sets: Method and Evidence from the Earnings Test," *American Economic Journal: Applied Economics*, January 2020, *12* (1), 1–31.

Lalive, Rafael and Stefan Staubli, "How Does Raising Womens Full Retirement Age Affect Labor Supply, Income, and Mortality," *NBER working paper*, 2015, *w18660.*

_ , Arvind Magesan, and Stefan Staubli, "Raising the Full Retirement Age: Defaults vs Incentives," Technical Report, Working paper 2017.

Lusdaine, R.L., J.H. Stock, and D.A. Wise, *Why Are Retirement Rates So High at Age 65?*, University of Chicago Press, 1996.

Manoli, Dayanand S and Andrea Weber, "The Effects of the Early Retirement Age on Retirement Decisions," *NBER working paper*, 2016, *w22561.*

Mastrobuoni, G., "Labor Supply Effects of the Recent Social Security Benefit Cuts: Empirical Estimates Using Cohort Discontinuities," *Journal of Public Economics*, 2009, *93* (11-12), 1224–1233.

Seibold, Arthur et al., "Reference Points for Retirement Behavior: Evidence from German Pension Discontinuities," Technical Report, CESifo Group Munich 2019.

Song, J. and J. Manchester, "Have People Delayed Claiming Retirement Benefits? Responses to Changes in Social Security Rules," *Social Security Bulletin*, 2007, *67* (2).

Staubli, Stefan and Josef Zweimüller, "Does Raising the Early Retirement Age Increase Employment of Older Workers?," *Journal of Public Economics*, 2013, *108*, 17–32.

Stock, James H and David A Wise, "Pensions, the Option Value of Work, and Retirement," *Econometrica*, 1990, *58* (5), 1151–1180.

U.S. Congressional Budget Office, "Social Security Policy Options," Report, Congressional Budget Office 2015.

Figure 1: Social Security Retirement Benefit Claiming

A. Aggregate Annual Rates

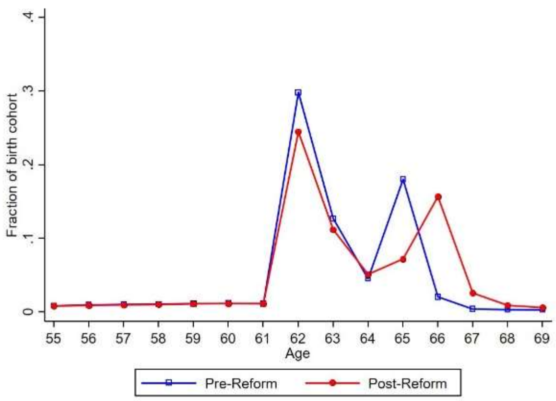

B. Monthly Rates by Cohort

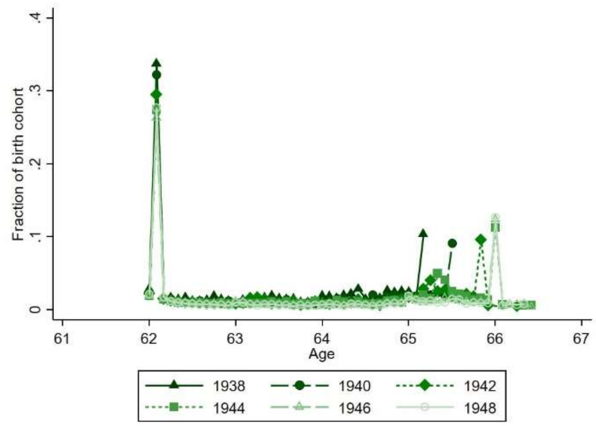

C. Claiming by Earnings Quartile

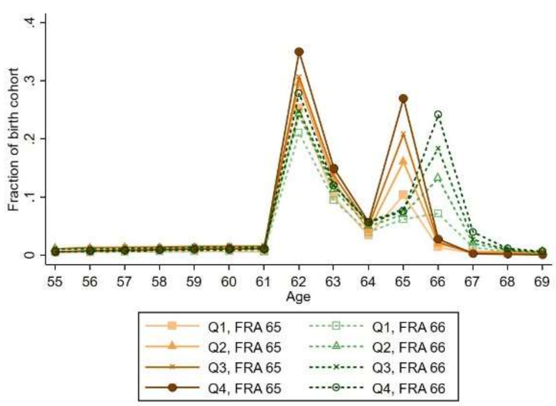

D. Claiming by Gender

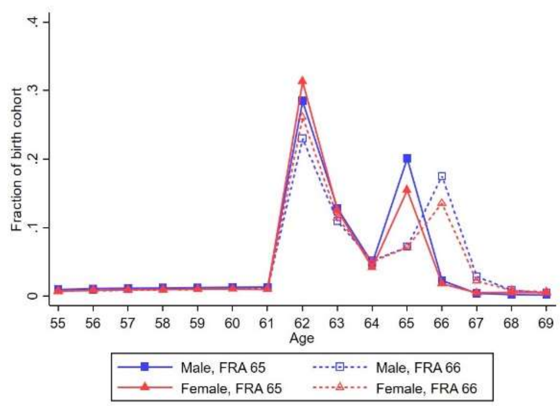

Notes: These figures plot claiming of Social Security retirement benefits by age, for different pre-reform and post-reform cohorts. Claiming at a given age is defined as starting to collect Social Security benefits at that age; that is, it is an indicator variable that is assigned the value 1 if an individual has positive benefits in a given period and no benefits in the previous period. Panel A aggregates behavior by age at the annual level. Panel B breaks it down to ages at the monthly level, and it also provides a breakdown by specific birth cohorts. Months prior to age 62 are dropped due to small cell sizes (and the corresponding data reporting rules). Panel C splits the sample by earnings quartiles based on an average over ages 25-55, and panel D splits the sample by gender.

Figure 2: Retirement Rates

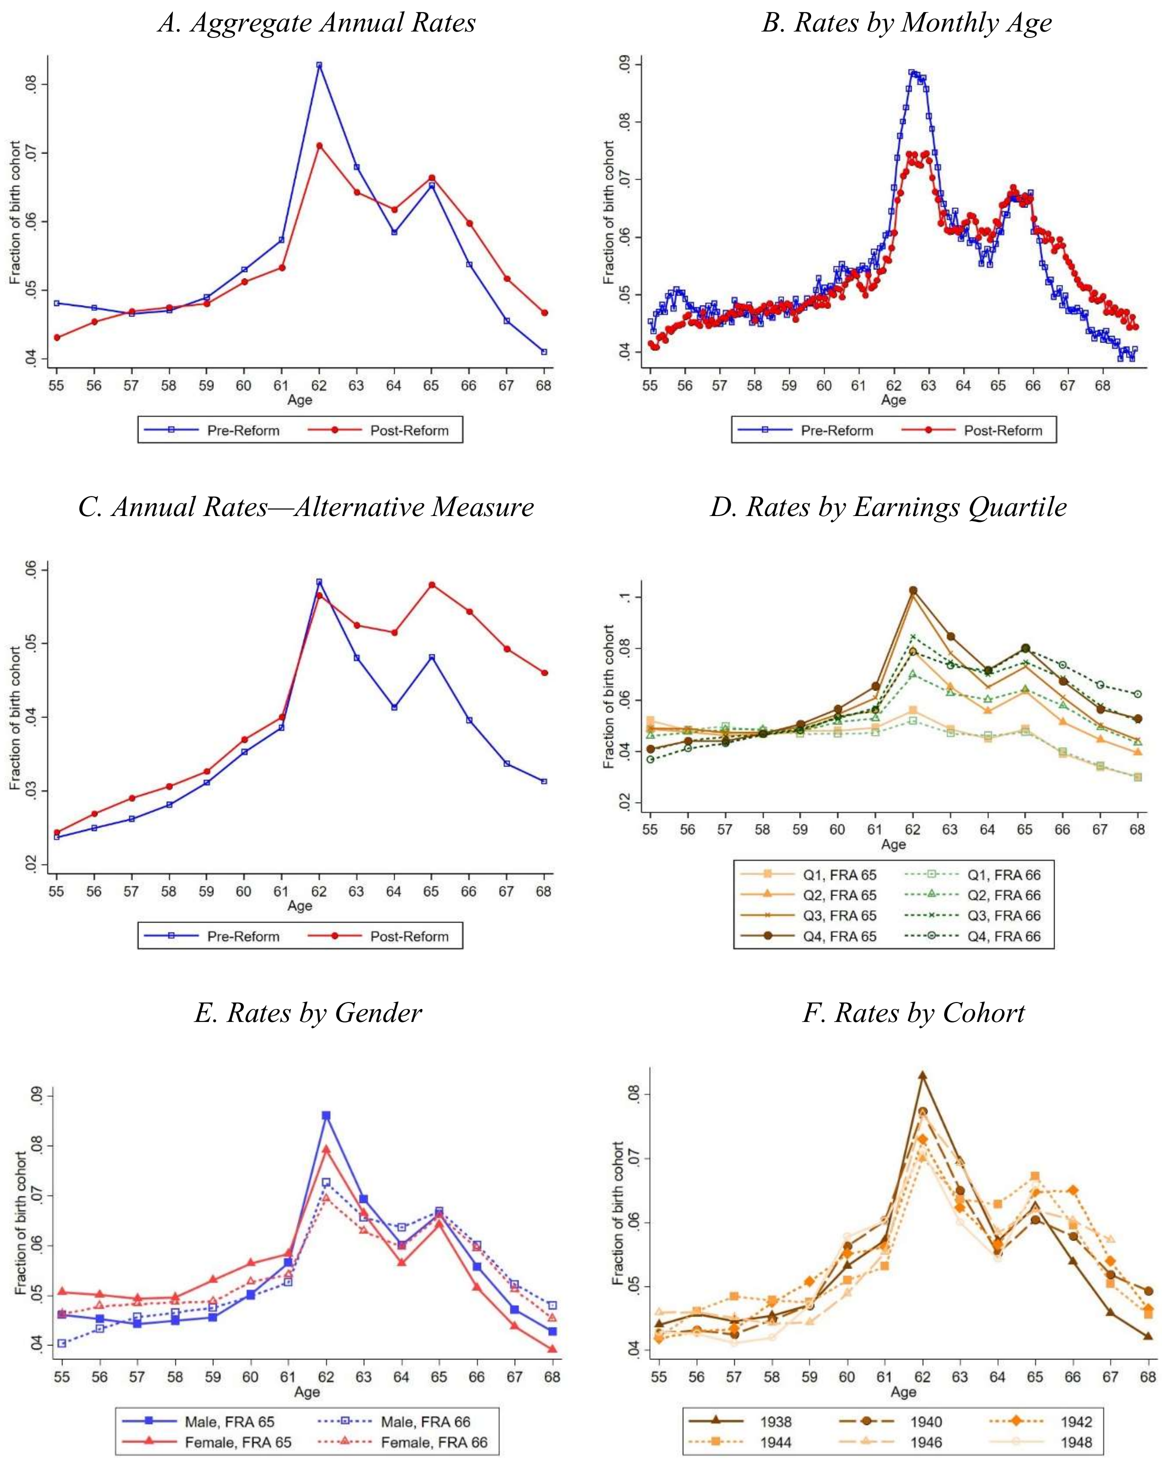

Notes: These figures plot retirement rates by age, for different pre-reform and post-reform cohorts. Retirement at a given age is defined by an indicator variable that is assigned the value 1 if an individual has positive earnings in a current period but no earnings in the next period. Panel A plots retirement rates using this measure, and panel C uses an alternative definition based on an indicator variable that is assigned the value 1 if a given age is the last data period in which an individual has positive earnings in a current period but no earnings in the next period. Panel B exploits exact dates of birth to construct monthly-level ages based on an individual's age at the end of the calendar year. Panel D splits the sample by earnings quartiles based on an average over ages 25-55, and panel E splits the sample by gender. Panel F plots retirement rates by age for a range of pre-reform and post-reform cohorts.

Figure 3: Cumulative Distributions of Claiming and Labor Supply by Cohort

A. Claiming and Participation

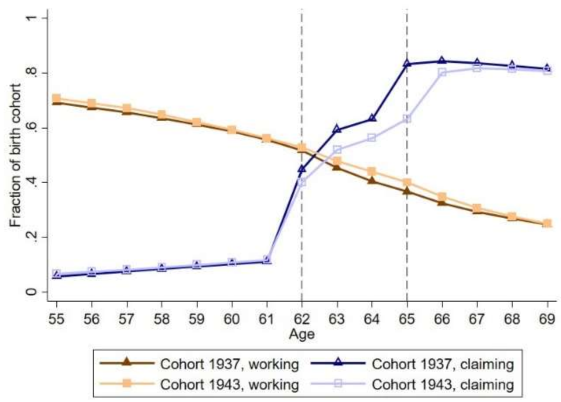

B. Claiming

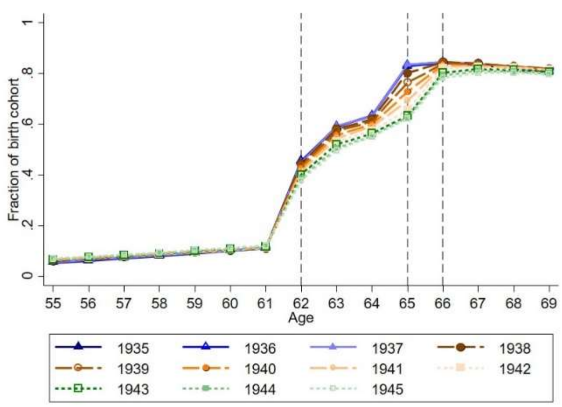

C. Participation

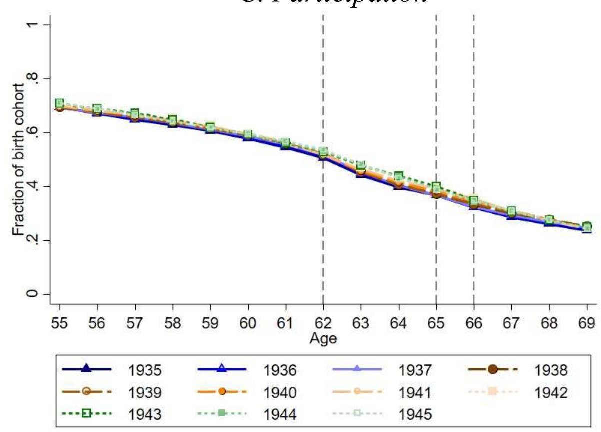

Notes: These figures plot cumulative distributions for Social Security benefit claiming and labor force participation by age. Panel A compares a pre-reform cohort (1937 with FRA 65) to a post-reform cohort (1943 with FRA 66). Panels B and C extend the analysis to encompass a range of cohorts—pre-reform cohorts (1935-1937 with FRA 65), phase-in cohorts (1938-1942 with increasing FRA at 2-month increments per cohort), and post-reform cohorts (1943-1945 with FRA 66)—for claiming and labor force participation, respectively.

Figure 4: Regression Discontinuity

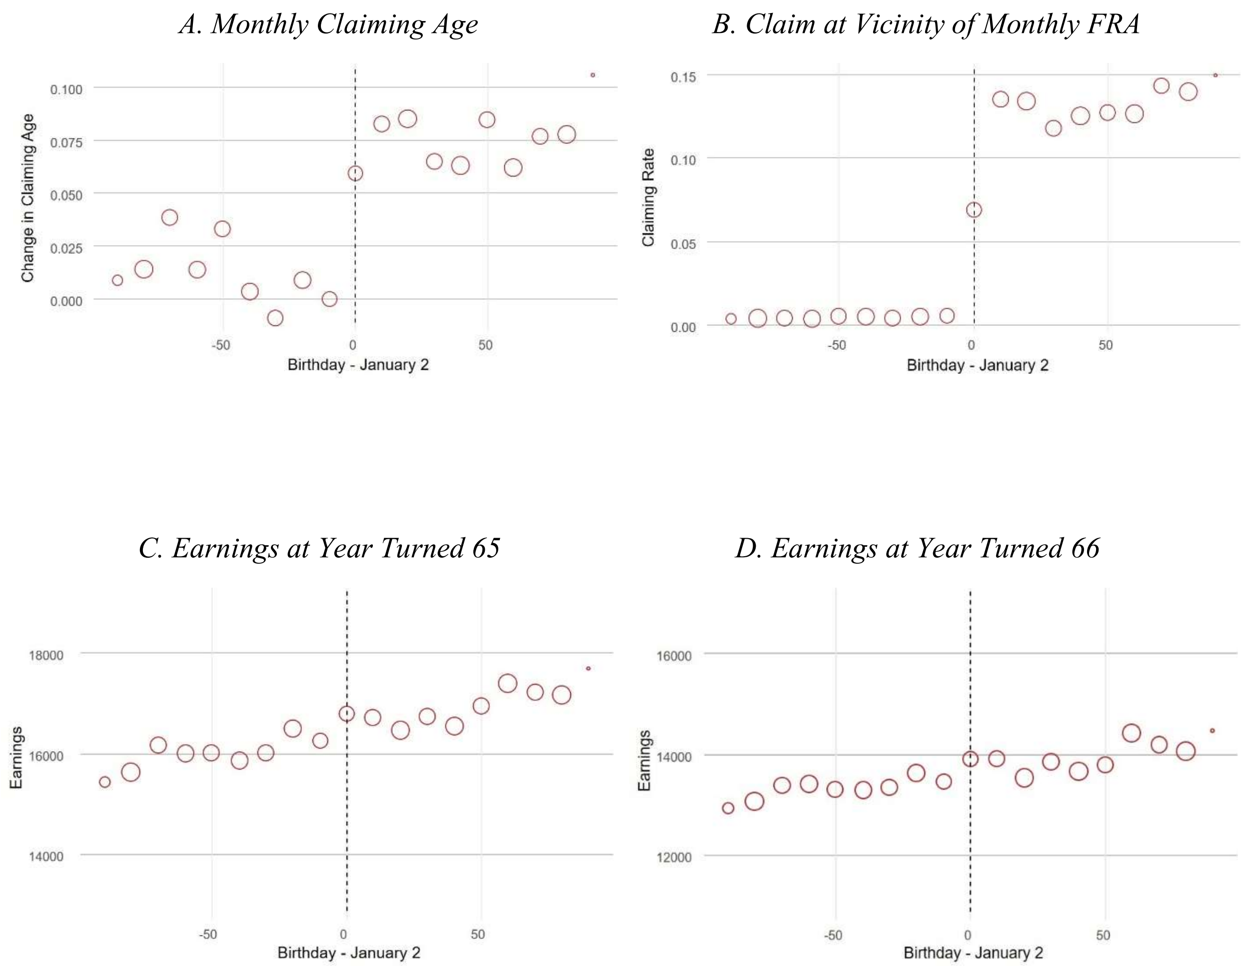

Notes: These figures leverage the 2-month increments in FRA around the January 2nd birthdate cutoff to study its effects on Social Security benefit claiming and earnings using a Regression Discontinuity (RD) design. Panels A-B display claiming outcomes, and panels C-D display labor supply outcomes. Earnings at a given age are reported for the calendar year at which affected individuals above the cutoff reached that age.

Table 1: Regression Discontinuity

Panel A: Claiming

	Claiming Monthly Age (1)	Claiming at Vicinity of Monthly FRA (2)
RD	0.070***	0.120***
	(0.012)	(0.003)
Constant	63.56***	0.005***
	(0.009)	(0.0003)
Number of observations	507,751	507,751

Panel B: Earnings

	Earnings in Year Turn 65 (1)	Earnings in Year Turn 66 (2)
RD	-51.363	-19.417
	(175.140)	(169.113)
Constant	16,487.15***	13,640.81***
	(128.888)	(126.344)
Number of observations	507,743	507,740

Notes: These tables leverage the 2-months increments in FRA around the January 2nd birthdate cutoff to study its effects on claiming (panel A) and labor supply (panel B). Earnings at a given age are reported for the calendar year at which treated individuals above the cutoff reached that age. The tables display regression discontinuity estimates using specification (1) with a ninety-day bandwidth. Standard errors clustered at the level of days relative to the cutoff are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01

Table 2: Potential Mechanisms—Geographical Location and Employers

Panel A: Moving across Counties

	Claiming				*Retirement*			
	Age 65		*Age 66*		*Age 65*		*Age 66*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Origin rate		0.049*		0.139		0.019		0.040
		(0.029)		(0.086)		(0.036)		(0.035)
Destination rate	0.067**	0.066**	0.270***	0.309***	0.017	0.046	-0.012	-0.009
	(0.030)	(0.027)	(0.094)	(0.086)	(0.043)	(0.039)	(0.037)	(0.033)
Origin change in rates		0.041		0.067***		0.018		0.027
		(0.026)		(0.022)		(0.029)		(0.027)
Destination change in rates	0.022	0.027	0.126***	0.134***	-0.014	-0.001	0.006	0.003
	(0.025)	(0.023)	(0.021)	(0.019)	(0.027)	(0.024)	(0.029)	(0.026)
Constant		0.099***		0.108***		0.077***		0.073***
		(0.006)		(0.006)		(0.004)		(0.004)
Number of observations	22,447	21,883	22,447	21,883	20,178	19,669	20,178	19,669

Panel B: Moving across Employers

	Claiming				*Retirement*			
	Age 65		*Age 66*		*Age 65*		*Age 66*	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Origin rate		0.012		0.149***		0.029		0.030
		(0.019)		(0.051)		(0.023)		(0.030)
Destination rate	0.086***	0.121***	0.237***	0.276***	0.060***	0.084***	0.093***	0.094***
	(0.017)	(0.021)	(0.039)	(0.041)	(0.023)	(0.023)	(0.033)	(0.032)
Origin change in rates		-0.0005		0.083***		0.013		0.009
		(0.016)		(0.020)		(0.018)		(0.027)
Destination change in rates	0.073***	0.107***	0.167***	0.205***	0.055***	0.074***	0.051**	0.056***
	(0.016)	(0.017)	(0.014)	(0.019)	(0.019)	(0.018)	(0.020)	(0.020)
Constant		0.086***		0.091***		0.064***		0.060***
		(0.006)		(0.006)		(0.002)		(0.002)
Number of observations	65,982	45,121	65,982	45,121	61,260	42,287	61,260	42,287

Notes: These tables display regression estimates based on specification (2) which leverages movers designs across geographical locations (panel A) and employers (panel B). The equation's right-hand side variables, indicated in the left-most column, are defined per specification and with respect to the studied outcome. For example, when studying the outcome "claiming at age 65" for our sample of movers, the variable "destination rate" refers to the rate at which individuals in the destination unit start claiming at age 65. Baseline "rates" are calculated using pre-reform cohorts, and "change in rates" are calculated using the difference between the behavior of post-reform cohorts and pre-reform cohorts. All rates are calculated using non-movers only and are measured in percentage points. For each outcome, we run two specifications which correspond to the two respective columns: the first specification studies variation in the destination and includes fixed effects for origin; and the second specification replaces these fixed effects with continuous measures for origin rates and a constant term. Standard errors clustered at the origin level are reported in parentheses. *p<0.1; **p<0.05; ***p<0.01

Appendix Figure 1: Social Security Benefit Annuities by Claiming Age

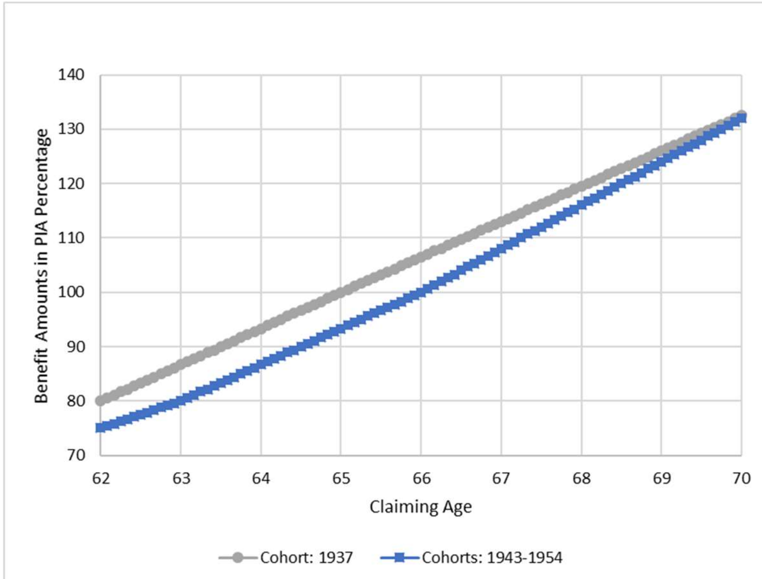

Note: This figure plots the annuity amount, in terms of percent out of the Primary Insurance Amount (PIA), against claiming age. It provides the schedule for pre-reform cohorts whose FRA is 65 (those born in 1937 or earlier) and for the post-reform cohorts whose FRA is 66 (cohorts 1943-1954).

Appendix Table 1: Workplace-Level Employee Behaviors—Correlation with Benefit Plan Characteristics

Panel A: Claiming

	Pre-Cohort Age 65	Diff. in Post-Cohort and Pre-Cohort Age 65	Pre-Cohort Age 66	Diff. in Post-Cohort and Pre-Cohort Age 66
	(1)	(2)	(3)	(4)
Retiree benefits	-0.0122	0.00771	-0.000221	0.00328
	(0.00981)	(0.0169)	(0.00208)	(0.0118)
Constant	0.290***	-0.174***	0.00639***	0.179***
	(0.0117)	(0.0188)	(0.00223)	(0.0131)
Number of EINs	11,656	6,320	11,656	6,320

Panel B: Retirement

	Pre-Cohort Age 65	Diff. in Post-Cohort and Pre-Cohort Age 65	Pre-Cohort Age 66	Diff. in Post-Cohort and Pre-Cohort Age 66
	(1)	(2)	(3)	(4)
Retiree benefits	0.000376	0.00660	0.00767	0.00203
	(0.00567)	(0.0109)	(0.00510)	(0.00991)
Constant	0.0724***	-0.0176	0.0473***	0.0197*
	(0.00687)	(0.0120)	(0.00610)	(0.0113)
Number of EINs	11,656	6,320	11,656	6,320

Notes: This table reports the correlation between our workplace-level employee behaviors and an indicator for whether the employer's benefit plans cover retirees. This indicator is based on Form 5500 called "Annual Return/Report of Employee Benefit Plan (With 100 or more participants)" from the years 1992-1998. The U.S. Department of Labor, Internal Revenue Service, and the Pension Benefit Guaranty Corporation jointly developed the Form 5500 Series for employee benefit plans to utilize in order to satisfy annual reporting requirements under the Employee Retirement Income Security Act of 1974 (ERISA) and the Internal Revenue Code. ERISA is a federal law that sets minimum standards for most voluntarily established retirement and health plans in private industry to provide protection for individuals in these plans, and it requires plans to provide participants with plan information including important information about plan features and funding (see www.dol.gov for more details). Our indicator for provision of retiree benefits is based on 'Box 7b' which documents the number of "Retired or separated participants receiving benefits." The instructions for filling this box indicate: "7b. Inactive participants receiving benefits are any individuals who are retired or separated from employment covered by the plan and who are receiving benefits under the plan. This includes former employees who are receiving group health continuation coverage benefits pursuant to Part 6 of ERISA who are covered by the employee welfare benefit plan." The table reports the coefficients from regressions that include as controls other plan characteristics reported on From 5500, specifically, whether the employer offers a defined benefit or defined contribution pension plan, whether some of those covered by the plan are members of a collective bargaining unit, the number of employees and the number of covered employees, the plan's average of net assets, average annual amounts of overall contributions to the plan and of employer contributions to the plan, and whether a plan was terminated. The sample includes all EINs for which there was a match across our SSA data and the publicly available datasets from Form 5500. *p<0.1; **p<0.05; ***p<0.01