

IZA DP No. 5310

Framing Social Security Reform: Behavioral Responses to Changes in the Full Retirement Age

Luc Behaghel
David M. Blau

November 2010

Framing Social Security Reform: Behavioral Responses to Changes in the Full Retirement Age

Luc Behaghel

Paris School of Economics-INRA

David M. Blau

*Ohio State University
and IZA*

Discussion Paper No. 5310
November 2010

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Framing Social Security Reform: Behavioral Responses to Changes in the Full Retirement Age^{*}

We use a US Social Security reform as a quasi-experiment to provide evidence on framing effects in retirement behavior. The reform increased the full retirement age (FRA) from 65 to 66 in two month increments per year of birth for cohorts born from 1938 to 1943. We find strong evidence that the spike in the benefit claiming hazard at 65 moved in lockstep along with the FRA. Results on self-reported retirement and exit from employment are less clear-cut, but go in the same direction. The responsiveness to the new FRA is stronger for people with higher cognitive skills. We interpret the findings as evidence of reference dependence with loss aversion. We develop a simple labor supply model with reference dependence that can explain the results. The model has potentially important implications for framing of future Social Security reforms.

JEL Classification: J26

Keywords: retirement, social security, loss aversion

Corresponding author:

David M. Blau
Department of Economics
Ohio State University
445 Arps Hall, 1945 N. High St.
Columbus, OH 43210-1172
USA
E-mail: blau.12@osu.edu

^{*} The research in this paper was conducted while the second author was a Special Sworn Status researcher of the U.S. Census Bureau at the Triangle Census Research Data Center. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Financial support from NIA grant P30-AG024376 and NSF-ITR grant #0427889 is gratefully acknowledged. We thank Sudipto Banerjee for research assistance. We appreciate helpful comments from seminar participants at MRRC, PSE, and NBER, and from Jonathan Zinman. We are grateful for especially detailed comments from Joe Doyle. The authors are solely responsible for the contents.

NON-TECHNICAL SUMMARY

Some aspects of retirement behavior have proven difficult to explain in the standard life cycle framework used by economists. A key example is the persistence of large spikes in exit from the labor force and take up of Social Security benefits in the US at age 65. In 1983, the US Social Security law was reformed by cutting benefits. However, the reform was framed as an increase in the “Full Retirement Age (FRA)” from 65 to 66, rather than as a benefit cut. We use this quasi-experiment to test a behavioral economic explanation for the age 65 spike, based on reference dependence. We find strong evidence of a shift in the spike in claiming benefits to 66, consistent with reference dependence, with the FRA as the reference point. We go further by addressing the question “who is behavioral?” We find that workers with higher cognitive skill respond more strongly to the FRA change. The most plausible explanation for the evidence we find is reference dependence, which is a form of non standard preferences, as opposed to non standard decision making such as might result from low cognitive skill.

Framing provides a policy maker with a potentially powerful and almost costless tool to influence aggregate labor supply: framing a reform as a change in the reference point magnifies the impact. How to use this knowledge depends on the goals of reform, and this suggests an important avenue for research: how should policy makers take into account loss aversion in designing future reforms?

1. Introduction

Understanding the labor force, benefit take up, and saving decisions of older workers has become increasingly important in today's environment of rapid population aging and large Social Security financial imbalances. The life cycle model provides a powerful framework for modeling retirement decisions, and has been the basis for a substantial amount of informative and policy-relevant research. But some aspects of retirement behavior have proven difficult to explain in the life cycle framework. These include failure to take up employer-provided defined contribution pension plans that provide very favorable terms such as a generous employer match (Madrian and Shea, 2001), and lack of knowledge of pension provisions (Chan and Stevens, 2008).

Another feature of retirement behavior that seems hard to reconcile with a life cycle approach is the persistence of large spikes in exit from the labor force and take up of Social Security benefits (Old Age and Survivors Insurance, or OASI) in the US at age 65. Costa (1998) shows that there was little evidence of a spike in labor force exit in the 1900-1920 period, but a spike at age 65 had emerged by 1940, five years after the establishment of Social Security. The size of the spike was as high as 30% on an annual basis in the mid 1980's (Perrachi and Welch, 1994), and declined to 19% in the 2000's.⁴

Several explanations for the age 65 spike that are consistent with the life cycle framework have been proposed. These include (1) a liquidity constraint that makes it difficult financially for some low income workers to retire before becoming eligible for Social Security benefits, (2) a kink at age 65 in the schedule that determines the Social Security benefit as a function of the age

⁴Source: author's calculations from the Health and Retirement Study, described below. A new spike in the hazard rates of labor force exit and benefit claiming emerged after 1962, when early Social Security entitlement at age 62 became available (Moffitt, 1987). Note that exiting employment and claiming the OASI benefit are distinct choices, although they are closely related in practice. The spikes in

of claiming, (3) the prevalence of age 65 as the normal age of retirement in many Defined Benefit (DB) pension plans, and (4) loss of health insurance coverage as a result of retiring before becoming eligible for Medicare at age 65. The evidence on these explanations (discussed below) is mixed, but overall it suggests that these factors cannot account for more than a small portion of the spike in retirement at 65, either alone or in combination. The evidence is not conclusive, however, because some of these explanations cannot be directly tested.

Other explanations for the age 65 spike have been proposed in the behavioral economics framework, in which the assumptions of farsighted rational behavior and standard preferences are relaxed. 65 was designated as the Full Retirement Age (FRA) from the beginning of the OASI program until recently. Workers might take this “official” designation as implicit advice from the government about when to retire and claim benefits, and as a result retirement at age 65 may have become a social norm. Information provided by the Social Security Administration (SSA) on their web site and in individualized “accounting reports” sent directly to future beneficiaries is explicitly framed with reference to the FRA. To paraphrase: Your FRA is 65. If you continue to work until 65, your monthly benefit will be \$X. If instead you claim at age 62, your benefit will be 80% of X. If you postpone claiming until age 70, your benefit will be 130% of X. This presentation could lead to framing effects at 65, resulting in age 65 becoming a reference point. There is little direct evidence to date on behavioral economic explanations for the age 65 spike. Most of the evidence comes from testing and rejecting other explanations, leaving behavioral economic explanations as the default (Lumsdaine, Stock, and Wise, 1996).

the hazard of claiming benefits at ages 62 and 65 are larger than the labor force exit spikes.

A Social Security reform enacted into law in 1983 increased the FRA from 65 to 66 in two month increments per year of birth for cohorts born from 1938 to 1943.⁵ The cohorts affected by the reform reached their FRA in 2004-2009, so data on their retirement behavior are becoming available now. This provides an unusual opportunity to test both life-cycle and behavioral economic explanations for the age 65 spike. The increase in the FRA is equivalent to a cut in the Social Security benefit: claiming the benefit at any given age results in a lower benefit than if the FRA had not changed. This reduces the expected present discounted value of lifetime benefits (Social Security Wealth), and should cause an increase in the age of retirement if leisure is a normal good. But, as illustrated in Figure 1, the increase in the FRA did not change the *slope* of the benefit-claiming-age schedule in the vicinity of the FRA.⁶ So there is no economic incentive for someone who, for whatever reason, would have retired or claimed the benefit at 65 if the FRA had not changed, to instead do so at his new FRA of 65 and 2 months, or 65 and 4 months, etc. If the spike in retirement or benefit claiming at age 65 shifts across cohorts in parallel with the increase in the FRA, explanations based on the standard life cycle framework would be unable to account for this. This would point toward behavioral economic explanations.

Our first contribution in this paper is to estimate the effect of the increase in the FRA on the hazard of exiting employment and the hazard of claiming the OASI benefit. Several recent

⁵The FRA is 65 for cohorts born before 1938, 65 and 2 months for the 1938 birth cohort, 65 and 4 months for the 1939 birth cohort, etc., and 66 for cohorts born from 1943-1954. A similar stepwise increase from 66 to 67 for cohorts born from 1955 to 1960 was also mandated.

⁶The implied benefit cut is the same at all claiming ages up to and including the FRA except between 62 and 63. The benefit cut implied by a one year increase in the FRA is 5% if the benefit is claimed between 62 and 63, and 6.67% if the benefit is claimed between 63 and 66. Another reform enacted in 1983 gradually increased the Delayed Retirement Credit (DRC), the slope of the benefit-claiming-age profile after the FRA, from 1% for those turning 62 in 1981 to 8% for those turning 62 in 2005. The benefit cut implied by a one year increase in the FRA for an individual who claims the benefit

studies have estimated the effect of the increase in the FRA on the timing of labor force exit or benefit claiming (Blau and Goodstein, 2010; Kopczuck and Song, 2008; Mastrobuoni, 2009; Pingle, 2006; Song and Manchester, 2008), but none have focused specifically on the impact on the spike in retirement at 65. In the first part of the paper, we use data from the Health and Retirement Study (HRS) and the Longitudinal Employer-Household Data (LEHD) to analyze changes in the retirement and claiming hazards across cohorts. We find strong evidence that the spike in the OASI benefit claiming hazard moved in lockstep along with the FRA, consistent with findings from administrative data (Song and Manchester, 2008). Results on self-reported retirement and exit from the labor force are less clear-cut: we find evidence that the spike in labor force exit at 65 decreased substantially and in some cases vanished for cohorts whose FRA increased, but less systematic evidence that new spikes have appeared at the new FRAs. The difference between effects on claiming and labor force exit may indicate that the change in the FRA is less salient for leaving employment than for claiming. Depending on the outcome we examine, the FRA effect can account for 10 to 40% of the initial hazard at age 65.

These results add to the growing literature on evidence “from the field” in behavioral economics. By contrast with other applications of framing effects to labor supply, the FRA offers a particularly favorable setting: the reference is explicitly defined and then exogenously modified by the 1983 reform, so that the behavioral component is clearly identified from cohort discontinuities.⁷

Our second contribution is to address the novel question of “who is behavioral” (Benjamin, Brown, and Shapiro, 2006) with respect to age of retirement. This is important in

after the FRA is equal to the DRC for his cohort.

order to infer which of several potential behavioral economic explanations is most consistent with the results. We define a group of workers as behavioral if the group's claiming and labor force exit behavior closely parallels the shift in the FRA. We define groups in several ways: by (1) socioeconomic characteristics such as gender, education, race, and marital status; (2) job characteristics; (3) cognitive ability; (4) financial literacy and planning horizon; and (5) non-cognitive characteristics such as risk aversion and subjective expectations about longevity. The most consistent finding is that workers with higher cognitive respond more strongly to the FRA change. One might have expected that low cognitive skills would make an individual more likely to interpret the change in the FRA as advice from the SSA. However, the most plausible explanation for the evidence we find is reference dependence, which is a form of non standard preferences, as opposed to non standard decision making such as might result from low cognitive skill. Moreover, the lower FRA effect for workers with low cognitive ability is largely due to the persistence of a large spike at age 65 for this group after the reform. This suggests that people with lower cognitive skills may be slower to learn about the FRA change or to integrate the new reference point in their claiming decision, perhaps instead keeping workers from earlier cohorts as their reference. To summarize, both groups of workers would be "behavioral" in the sense of reacting to a reference point with loss aversion, but workers with lower cognitive skills would be slower to adjust to the new reference points.

The evidence we present is of inherent interest for understanding retirement behavior, but it could also have important implications for future Social Security policy. To illustrate these implications, we develop a simple behavioral model of retirement that incorporates reference

⁷ For a survey of empirical applications, see Della Vigna (2009); for theoretical solutions to the indeterminacy of the reference point, see for instance Koszegi and Rabin (2006).

dependence and loss aversion. The model is used to demonstrate that reference dependence can amplify or dampen the impact on the average age of retirement caused by a reform such as a change in the FRA, depending on how the reform is framed. Simulations of the model for plausible parameter values indicate that the manner of framing a policy reform can have a sizeable effect on its impact on retirement behavior in the presence of reference dependence.

The next section of the paper briefly reviews previous findings and places our contribution in context. The following section describes the data, and sections 4 and 5 present evidence on the effect of the shift in the FRA. Section 6 analyzes the question of “who is behavioral?” Section 7 describes and analyzes the implications of a simple behavioral model of retirement. The final section concludes.

2. Background and Previous Studies

Here, we briefly review evidence on explanations for the age 65 spike that are consistent with the standard life cycle model.

(1) Liquidity Constraint. Until 1962, eligibility for OASI began at age 65. Low income workers who saved little during their working years could face a liquidity constraint that makes it difficult to finance consumption during retirement before receiving the OASI benefit. This could explain the prevalence of retirement at the earliest age of eligibility. However, the earliest age of eligibility for OASI was changed to 62 in 1962, but the spike at age 65 remained and even grew, thus providing evidence against a liquidity constraint explanation. It is possible that liquidity constraints were an important source of the spike before 1962 and were gradually supplanted by other factors. And a liquidity constraint is a plausible explanation for the spike at age 62.

(2) Nonlinear Budget Set. Until the 1990s there was a sharp kink at the FRA in the schedule that determines the Social Security benefit as a function of the age of claiming. Delaying claiming from 62 to 65 resulted in an approximately actuarially fair benefit increase (6.67% per year), but delaying claiming past 65 resulted in a much less than actuarially fair increase (1% up to 1981; 3% from 1981 until 1989). The age-65 spike could be rationalized as a response to a kinked intertemporal budget constraint. However, the 1983 Social Security reforms eliminated the kink, gradually increasing the reward to delaying claiming past the FRA from 3% for cohorts reaching their FRA in 1989 to 8% for cohorts reaching their FRA in 2008. By the mid-2000's there was no longer a kink in the benefit-claiming age schedule, and today there is even a slight convex kink, yet the spike in retirement at age 65 for cohorts with an FRA of 65 persisted, suggesting that the budget constraint kink was not the main factor behind the age 65 spike in retirement, at least in recent years.

(3) Defined Benefit Pensions. DB pension plans often have a normal retirement age of 65, and these plans usually have very strong incentives to retire by the normal age (conditional on not having retired at the earliest age of eligibility, which is typically quite attractive as well). DB pensions are much less common today, having been largely supplanted by Defined Contribution (DC) plans.⁸ DC plans do not have any incentives to retire at 65 or any other particular age. However, the switch from DB to DC plans affected the cohorts reaching their mid 60's in the 2000s much less than it has affected more recent cohorts. The prevalence of DB

⁸Measured in terms of annual contributions by employers, DB plans accounted for 60% of total private sector employer contributions in 1980, and only 13% in 2000 (Poterba, Venti, and Wise, 2007). DB plans remain prevalent in the public sector.

pension coverage in these older cohorts is consistent with the persistence of the age 65 spike. We control for DB pension coverage and, for those with DB plans, the normal retirement age.

(4) Health Insurance. Workers who would lose their employer-provided health insurance upon retiring might prefer to postpone retirement until Medicare eligibility at age 65 in order to avoid being uninsured, leading to a spike in retirement at 65. The age of eligibility for Medicare has been 65 since the program was introduced in 1965, so there is no direct evidence on this explanation, although the age-65 spike was present before 1965. It is possible that Medicare has replaced the other explanations for the age-65 spike as they have become less relevant, but this is difficult to determine because of the lack of variation in the Medicare eligibility age.⁹

The most plausible behavioral economic explanation for the age 65 spike stems from the fact that 65 was the Social Security FRA until recently. The FRA is not presented as a norm – the SSA presents things in a balanced way: “If you retire early, you may not have enough income to enjoy the years ahead of you. Likewise, if you retire late, you’ll have a larger income, but fewer years to enjoy it. Everyone needs to find the right balance based on his or her own circumstances.” (SSA, 2008a). However, in personalized Social Security statements the FRA is explicitly used as a reference in a bar chart illustrating benefits as a function of claiming age.¹⁰ Moreover, the distinction between retiring “early” (that is, earlier than the FRA) and retiring

⁹Evidence on the role of Medicare in retirement decisions derived from simulations based on structural models is provided in Rust and Phelan (1997) and Blau and Gilleskie (2006, 2008). Rust and Phelan conclude that Medicare was an important determinant of retirement timing in the 1970s, while Blau and Gilleskie conclude that it was much less important in the 1990s.

¹⁰ See Mastrobuoni (2010) for discussion of the history of the Social Security statement and the impact of its introduction on retirement behavior.

“late” (later than the FRA) is explicitly discussed (SSA, 2008b; see appendix). Clearly, the way the FRA is used in framing benefits invites people to use it as their point of reference.¹¹

There is little direct evidence on the explanatory power of specific behavioral economic explanations for the age 65 spike. Featherstonehaugh and Ross (1999) and Liebman and Luttmer (2009) pose hypothetical questions to survey respondents to gauge the importance of framing effects in the presentation of information about Social Security. Framing effects appear to matter at some ages and not at other ages, but the scenarios are hypothetical. The hypothetical experiment we have in mind would cut Social Security benefits by a given amount, and frame the cut in alternative ways: (1) as a neutral across-the-board cut irrespective of the age at which the benefit is claimed, (2) as a cut in the benefit available at a particular reference age, holding constant the slope of the benefit-claiming-age profile, and (3) as an increase in the age at which a given reference benefit level is available, again holding the slope constant. In each scenario, individuals would be perfectly well informed about the cut, and capable of determining their optimal response, given their preferences. This would eliminate lack of information and limited cognitive ability as confounders, so any differences in responses across scenarios could be attributed to reference dependence. The actual quasi-experiment induced by the reform did hold the benefit-claiming-age profile roughly constant (c.f. Figure 1), and used framing option 3 (an increase in the FRA). We cannot compare results for different framing options and we cannot rule out information and cognitive ability as confounders. We discuss indirect approaches to assessing the importance of these issues in sections 5 and 6.

¹¹ This is true for cohorts with non-integer FRAs. For instance, for workers born in 1939, the statement reads: “The earliest age at which you can receive an unreduced retirement benefit is 65 and 4 months.”

Recent evidence indicates that the increase in the FRA has affected retirement behavior. Mastrobuoni (2009) analyzed the age of exit from the labor force of the “treated” cohorts, comparing them to earlier cohorts not affected by the increase in the FRA. His estimates indicate that each two month increase in the FRA caused a one month increase in retirement age. Blau and Goodstein (2010) and Pingle (2006) also estimate that the increase in the FRA has caused an increase in employment at older ages. The evidence in these three studies does not address the mechanism through which the FRA effect operates, and therefore does not shed light on the question of *why* the FRA has affected retirement behavior. The evidence is consistent with a wealth effect that would alter behavior at all ages, but it does not address the question of whether there is a shift in the spike at age 65.

In contrast, recent evidence on Social Security claiming clearly suggests a behavioral economic interpretation. There is no economic incentive to claim the OASI benefit at the FRA, yet that is precisely what the treated cohorts have done. Song and Manchester (2008) and Kopczuck and Song (2008) use administrative data to show that the increase in the FRA has caused the spike in claiming at age 65 to shift almost completely to the new FRA for the affected cohorts. This finding is difficult to explain in the life cycle framework. The key unanswered questions that we seek to address are whether there is similar evidence for employment; if so, how can we explain it; and who is “behavioral.”

3. Data

Previous studies of the impact of the FRA have used data from the Current Population Survey (Blau and Goodstein, 2010; Mastrobuoni, 2009), the Survey of Income and Program

Participation (Pingle, 2006), and SSA administrative data (Song and Manchester, 2008; Kopczuk and Song, 2008). These data sources provide large sample sizes, which may be needed to detect the effects of small changes in the FRA (two months per birth cohort). We use data from the Health and Retirement Study (HRS), which provides a rich set of potential explanatory variables to study behavioral aspects of retirement. These variables are unique to the HRS and offer the opportunity to get closer to *explaining* behavioral aspects of retirement as well as documenting their existence.

The main disadvantage of the HRS is the relatively small samples available to study retirement and claiming behavior at the FRA. Roughly three quarters of workers retire and claim benefits before reaching the FRA, so despite sample sizes of about 1,000 respondents per year of birth, the effective sample size is closer to 250 per cohort. This provides about 1,125 observations on the treated cohorts, and 1,750 for the control cohorts. Therefore, we also use data from the Longitudinal Employer-Household Data (LEHD) files. These files are derived from administrative state Unemployment Insurance (UI) records, with very large sample sizes. These administrative data are used only for descriptive purposes to verify that the trends identified in the HRS are robust, since they contain no information on any of the explanatory variables of interest. The LEHD is described in Section 4.

The HRS is a biennial survey of a sample of households containing individuals over the age of 50, and their spouses. The survey began in 1992 with birth cohorts 1931-1941, and new cohorts were added in 1998 and 2004. We use the 1992-2008 waves. The analysis sample is birth cohorts 1931-1942, since these cohorts had reached their FRA as of 2008, while later cohorts had

not. The 1931-37 cohorts are the controls (FRA=65.0) and the 1938-42 cohorts are the treated cohorts (FRA>65.0).

The HRS records employment status at each interview, and in addition asks for information on the start and end dates of all jobs between interviews, to the nearest month. We construct a monthly employment history.¹² Together with month and year of birth, the employment data are used to construct the month and year of labor force exit.¹³

The HRS contains self-reported information on the month and year in which the respondent first received a Social Security benefit payment. In some cases, the reported date is before the respondent turned 62, indicating that he or she first received some type of Social Security benefit other than OASI, such as disability or dependent benefits. In these cases it is not possible to identify when the respondent claimed the OASI benefit.¹⁴ The month and year of OASI claiming is the second outcome of interest.

A third outcome of interest is the self-reported month and year of retirement. This “subjective” measure is frequently used as an indicator of retirement, and it is of interest to determine whether “retirement” and “employment” differ with respect to the FRA. However, there are many longitudinally inconsistent self-reports of retirement age, so only about half the sample has a self-reported retirement age that is reasonably consistent across waves.

¹²The initial interview collected information on the start date of the job held at the interview date, start and end dates of up to two other jobs before the first interview (longest and last), and the number of additional jobs that lasted five or more years.

¹³Some people who leave the labor force later return to employment. In such cases there are multiple labor force exit dates. As a robustness check, we use labor force exits that are preceded by an employment spell of at least 3 months, and followed by at least 3 months out of employment.

¹⁴If an individual reaches his FRA while receiving Social Security Disability Insurance, the benefit is switched to OASI, but this is purely an administrative adjustment, so it does not provide any information about claiming behavior. Most SSDI recipients never leave the SSDI rolls, and therefore

The monthly record is merged with permanent characteristics such as race, gender, ethnicity, and education, and with time-varying measures recorded at the survey dates, including health status, the wage rate, job characteristics, household wealth, health insurance, pension coverage, and marital status. If there was a change in one of these time-varying variables between waves, we assume the change occurred midway between the waves.

The HRS contains several variables that are useful in inferring “who is behavioral”. These include measures of cognitive ability, risk aversion, self control, and financial planning horizon. Cognitive ability measures have been explored by McArdle, Smith and Willis (2009), and we follow their approach to construct indicators in three dimensions: Telephone Interview of Cognitive Status (TICS), short term memory, and numeracy (see data appendix for details).

4. Impact of the FRA increase

HRS Results

In this section, we describe evidence from the HRS on how the change in the FRA has modified the timing of OASI benefit claiming, labor force exit, and self-reported retirement. The goal is to test the “behavioral” prediction that the spike at age 65 should shift along with the FRA for cohorts born after 1937, against the null hypothesis that the increase in the FRA had only a wealth effect. In the latter case we expect an increase in retirement age, but no substantial shift in the spike.

We start with graphical evidence on the timing of OASI benefit claiming across cohorts, pooling men and women. Figure 2 displays average monthly claiming hazard rates for pre and

provide no information about OASI claiming behavior.

post-reform cohorts. The claiming hazard rate is defined as the probability of claiming at a given age, conditional on not having claimed previously. Age is measured at a bimonthly frequency; e.g. age 65 denotes age 65 0/12 to 65 1/12¹⁵. In each graph of figure 1, the dotted line depicts the claiming hazard for workers born between 1931 and 1936. As shown by previous studies, there is a first spike in the hazard at or just after the early retirement age (62) and a second larger spike at the FRA (65).¹⁶ About 20% of workers claim at the early retirement age, and 30% of those who have not claimed before 65 claim at 65. For each cohort, the vertical lines indicate age 62, age 65, and the FRA (if different from age 65). There is clear evidence that the spike in the claiming hazard moves in lockstep along with the FRA. The spike at age 65 does not completely disappear for the treated cohorts, but it becomes progressively smaller across cohorts. Very similar patterns appear when men and women are disaggregated (not shown). These results confirm the findings of Song and Manchester (2008), using administrative data.

Regression analysis is useful here to summarize the graphical evidence and to quantify the impact of the FRA. We adopt the following difference-in-difference specification:

$$P_{iac} = \theta FRA_{iac} + x_{iac} \gamma + \beta_a + \delta_c + \varepsilon_{iac} \quad (4.1)$$

where P_{iac} is an indicator variable equal to one if individual i born in cohort c claims at age a (in months), conditional on not having claimed previously. FRA is the indicator variable for age a being his FRA, x_{iac} is a set of individual controls, and full sets of cohort and age dummies are

¹⁵ One reason for measuring age at a bimonthly frequency is to make the graph easier to read. It is also useful because there is some arbitrariness in measuring the age at which an event occurs. If an individual reports leaving his job in April, it is not clear whether to classify his employment status in April as employed or not employed, without knowing the exact date.

¹⁶ The spike at age 62 is slightly after age 62 because benefits are payable beginning in the first month in which a person is 62 throughout the whole month, unless the person was born on the 1st or 2nd of the month (see Kopczuk and Song, 2008, for discussion).

included (β_a, δ_c) . The age 65 coefficient (one of the β 's) captures the part of the spike that is not explained by the fact that 65 is the FRA for cohorts up to 1937. The parameter of interest θ is identified by the interaction of age and cohort, under the assumption that the control variables capture any non-FRA-related motives to claim at the FRA. Results are shown in table 1. The four columns differ by the estimation sample or the controls. Column 1 has no controls other than age and cohort effects, and restricts the analysis to ages 64 to 65 11/12. Reaching the FRA increases the claiming hazard by 14 percentage points. The effect is statistically highly significant, and robust to the inclusion of controls (socioeconomic characteristics, pension and job characteristics, measures of cognitive ability, planning horizon and risk aversion) and to changes in the estimation sample (columns 2 to 4). Column 4 drops the 1938 cohort since the graphical analysis suggests that the effect of the FRA change might have been smaller for that cohort (with a higher persistence of the old FRA spike), possibly due to a learning effect. This however makes little difference in the estimation. Quantitatively, the estimated impact of the FRA is sizeable: the claiming hazard at age 65 is around 30% for cohorts born between 1931 and 1937; more than 40% (14/30) of the claims occurring at that age for the control cohorts can therefore be explained by the fact that 65 is their FRA.

Equation (4.1) relies on the identifying assumption that any change in the shape of the claiming hazard can be attributed to the FRA increase. Other changes over the period may also have affected the timing of claiming decisions¹⁷. In particular, the Social Security earnings test for people who have reached their full retirement age was eliminated in 2000. The rule before

¹⁷ As noted above, the 1983 reform increased the Delayed Retirement Credit (DRC) as well as the FRA. An increase in the reward to claiming the benefit after the FRA could affect claiming behavior (and

2000 was the following: workers who have not reached the FRA have their earnings reduced by \$1 for every \$2 earned beyond the earnings test threshold; workers who have reached the FRA have their benefits reduced by \$1 for every \$3 earned beyond the earnings test threshold (see Manchester and Song, 2008). Although this reduction is compensated by increased benefits at older ages (making it more or less neutral in terms of SS wealth), it may be felt as a disincentive to claim benefits by those who intend to continue working (Friedberg, 2000; Gruber and Orszag, 2003; Haider and Loughran, 2008). In 2000, the earnings test was eliminated for workers who have reached the FRA. The earnings test was unchanged before the FRA. The introduction of this discontinuity may generate a spike in the claiming hazard around the FRA, if some workers delay claiming SS benefits until they reach the FRA, in order to avoid the earnings test.

Fortunately, the cohorts impacted by the increase in the FRA and the removal of the earnings test do not fully overlap, making it is possible to separately identify the two effects (see Song and Manchester, 2008, for details). It is therefore possible to extend equation (4.1) to separately identify the impact of the earnings test removal:

$$P_{iac} = \theta FRA_{iac} + \lambda ETR_{iac} + x_{iac} \gamma + \beta_a + \delta_c + \varepsilon_{iac}, \quad (4.2)$$

where ETR_{iac} is an indicator for the month in which the earnings test ceases to apply to individual i . The results are given in column (5). The coefficient on the FRA remains unchanged. The impact of the earnings test removal is positive and significant, as expected, suggesting that part of the age 65 spike for some cohorts was due to the removal of the earnings test. However, the impact of reaching the FRA itself is much larger.

does, according to Blau and Goodstein, 2010, and Pingle, 2006). But the DRC is cohort-specific, so its effects are absorbed by cohort dummies.

As shown by Coile et al. (2002) and Baker and Benjamin (1999), benefit claiming behavior can significantly differ from labor force participation and retirement behavior. Before interpreting the shift in the claiming spike from a behavioral perspective, it is therefore important to check whether retirement and employment patterns are similar to claiming patterns. Figure 3 shows that the increase in the FRA resulted in a progressive fall in the age 65 spike in the labor force exit hazard. However, there is no systematic evidence of new spikes at the FRA for cohorts born after 1937.

Columns 1 to 5 in table 2 use the same specifications as table 1. The impact of the FRA is smaller than for claiming, but is positive and significantly different from zero. For cohorts born before 1937, the monthly hazard of labor force exit at age 65 is about 4.6%. Roughly 20% of this spike ($0.9/4.6$) can be explained by the fact that 65 is the FRA for these cohorts.¹⁸ The FRA effect is again robust to inclusion of a control for elimination of the earnings test.

Figure 4 presents evidence on the monthly hazard of entry to self-reported retirement, comparable to figures 2 and 3. The results are in between. There is some reasonably strong evidence of a shift in the spike for cohorts born after 1939, consistent with an effect of the FRA on retirement decisions. However, the spike at the old FRA (age 65) persists for some of these cohorts. The regression results (table 3) are very imprecise. The point estimate implies a smaller FRA impact than for claiming and exit from employment. The mean monthly retirement hazard for the control cohorts at age 65 is around 13%. 10% of this spike ($1.1/13$) can be accounted for by the fact that 65 is their FRA. The results are robust to controlling for the elimination of the earnings test, but in this case the effect of the earnings test is larger than the effect of the FRA.

¹⁸ Results based on a smaller sample that eliminates temporary withdrawals from the labor force (three months or less) gave very similar results.

LEHD results

In this section, we describe evidence from the LEHD on how the change in the FRA has affected the timing of labor force exit. The LEHD Infrastructure File system is based on state Unemployment Insurance (UI) administrative files, with data available from 31 states covering about 80% of the U.S. work force for the years 1990-2004, although the period covered varies by state (Abowd, Haltiwanger, and Lane, 2004). Employers covered by UI file a quarterly report for each individual who received any covered earnings in the quarter. UI covers about 96% of private non-farm wage-salary employment, with lower coverage of agricultural and government workers, and no coverage of the unincorporated self-employed. The UI records contain the individual's Social Security number, and an identification number and quarterly earnings for each employer from which he has any covered earnings during the quarter. These data are merged by the Census Bureau with the Census Personal Characteristics File, which contains the exact date of birth, place of birth, sex, and a measure of race/ethnicity.¹⁹

The main advantage of the LEHD for our analysis is the very large sample size. The main disadvantages are absence of information on hours of work, and lack of data after the third quarter of 2004. Without data on hours of work, we must use changes in quarterly earnings to infer changes in employment status. The absence of data beyond the third quarter of 2004 means

¹⁹An extensive discussion of the construction and the content of these files is provided in Abowd *et al.* (2006). We use a subsample of the full LEHD files, consisting of workers who were employed at any employer at which a member of the 1990-2001 panels of the Survey of Income and Program Participation (SIPP) worked. This is a very large subsample of the full LEHD, somewhat skewed toward large firms. See Blau and Shvydko (in press) for a description of the subsample. In order to reduce the number of quarterly observations from the hundreds of millions to the more manageable level of several

that only two of the “treated” birth cohorts (1938 and 1939) can be used to analyze changes in labor force behavior around age 65..

The main outcome of interest is a binary indicator of zero earnings in a given quarter, conditional on positive earnings in the previous quarter, which we interpret as the hazard of labor force exit. We aggregate the data into cells defined by month and year of birth (January 1931 through December 1942), calendar quarter (1994:Q1 through 2004:Q3, although not all quarters are represented for all cohorts), and age in quarters (62.0, 62.25, 62.5,..., 66.0), and conduct the analysis using cell means, with no loss of information. Data are available for 2,400 cells, with roughly 1,000 individual observations per cell on average. Men and women are pooled and gender and state dummies (fractions, after aggregation) are included in the regression models.

The specification used in column 1 of table 4 is the same as in tables 2-4. Age and birth year dummies are included, so that the effect of the FRA variable is identified from the interaction of age and birth cohort.²⁰ The results are consistent with reference dependence, with a positive FRA effect on the exit hazard, significantly different from zero. Given the fact that the data only has two post-reform cohorts, we check a new specification in column 2. The FRA indicator is interacted with cohort dummies: so, the coefficient on the FRA*(coh=1938) variable can be interpreted as the average impact of the FRA for cohort 1938, compared to individuals of the same age in the 1931-37 cohorts, and accounting for cohort trends. The results are

million, we use data from only three states. Census Bureau guidelines prevent us from identifying the three states.

²⁰ If a worker leaves employment at his FRA, his earnings will be positive in the quarter in which his FRA falls (unless he quits on the first day of the quarter), so zero earnings in the first calendar quarter after the quarter in which he reaches his FRA is the only reliable measure of exit at the FRA. Depending on birth month within a quarter, some cases do not provide any evidence on the impact of the change in the FRA. A more detailed discussion of which cases contribute to identification is available from the authors.

unexpected for the 1938 cohort: reaching the FRA *decreases* the hazard of exit. The effect for cohort 1939 is small and statistically insignificant. The specification in column 2 has the advantage of being less restrictive than in column 1, as it does not force the magnitude of the FRA effects to be the same for all cohorts, but it also uses less information: in column 1, the spike at the FRA observed in the 1931-37 cohorts and its decline afterwards contribute to the identification of the FRA coefficient, whereas the estimate in column 2 only tests for the emergence of spikes at the new FRAs.²¹

Summary

Overall, combining the information on labor force transitions from the LEHD and the HRS as well as from self-reported retirement age from the HRS provides only mixed evidence that the labor supply decisions of workers have been affected by the change in the FRA in a manner consistent with a behavioral interpretation. Limited statistical power and measurement error are issues with each of these data sources. However, their combination gives us some confidence that labor supply decisions are affected, but only in a limited way.

An important question is whether the magnitudes of the changes in claiming, labor force, and retirement behavior could be explained solely by wealth effects as a response to the benefit cut implied by the increase in the FRA. We cannot directly address this question because we do

²¹ We constructed graphs like those in Figures 2-4 and estimated models with non-parametric specifications (like those in Mastrobuoni, but for hazards rather than levels). The results were similar to the HRS results in Figure 3, showing clear evidence of a decline in the hazard at 65, but less clear evidence of increases at other ages. We omit these results for brevity, but they are available on request. We also analyzed exit from employment in the HRS data using the same cohorts as in the LEHD data and interacting the FRA indicator with indicators for the 1938 and 1939 cohorts; estimates of the FRA impact go in the same direction: 0.5 (0.5) when pooling all sources of identification like in column 1 of table 4; -0.9 (0.8) and 1.6 (1.2) for the FRA indicator interacted with birth cohorts 1938 and 1939, respectively.

not estimate the wealth effect; rather, as in Mastrobuoni (2009), we estimate the total effect, including the wealth effect and any “behavioral” effects. Several papers report estimates of the elasticity of the hazard of labor force exit with respect to Social Security Wealth (SSW): (1) Coile and Gruber (2007): the largest effects they find are 0.16 when evaluated at mean SSW, and .075 when evaluated at median SSW. (2) Samwick (1998): approximately zero. (3) Day, Mullen, and Wagner (2009) using Austrian administrative data: 0.40. The change in SSW wealth implied by a change in the FRA from 65 to 66 is 6.67%. If we take the largest elasticity estimate, 0.4, this would imply a 2.7% (not percentage point) decrease in the hazard of LF exit. Using the largest estimate from Coile and Gruber, 0.16, implies a 1.07% decline in the hazard. These numbers cannot be compared directly to the results reported in Figures 2-4, but the visual impression from these figures is of effects much larger than 1-3%.

5. Who is behavioral?

Results from section 4 provide strong evidence that OASI benefit claiming behavior has followed the increase in the FRA, and weaker evidence that the same is true of labor force participation. As argued in section 2, this finding leaves behavioral factors as likely explanations. Three leading candidates are reference dependence with loss aversion, “advice” from the SS administration, and “social norms”. An indirect way to discriminate between alternative behavioral explanations is to ask a simpler, descriptive question: which types of workers respond most strongly to the FRA shift? This is interesting per se – indeed, the recent retirement literature has stressed the fact that aggregate retirement behavior may hide considerable heterogeneity (e.g.

This contrasts with the claiming results, where the FRA impact is positive (and significant) in all specifications, as can be expected from figure 2.

see the discussions by Burtless, 2004; Liebman et al., 2008; and the empirical applications in Coile et al., 2002, and Chan and Stevens, 2008). It may also shed light on the most likely behavioral mechanism. For instance, if workers with lower cognitive skills respond more to the FRA, this would point toward non-standard decision making (bounded rationality, for instance), making an “advice” or social norm explanation plausible.

A simple way to look at this question is to compare the FRA impacts across subpopulations. The corresponding regression model is:

$$P_{iac} = \theta_1 FRA_{iac} + \theta_2 FRA_{iac} \times Type_{iac} + \beta_{1a} + \delta_{1c} + \gamma_{1a} \times Type_{iac} + \zeta_{1c} \times Type_{iac} + \varepsilon_{iac}, \quad (5.1)$$

where *Type* is an indicator variable that splits the population in two (for instance, *Type* is 1 for individuals with higher numeracy, 0 otherwise). The estimate of θ_2 reveals whether there is a different response to the FRA in the population characterized by the *Type* variable. The type variable is also interacted with cohort and age dummies.

Table 5 displays estimates for OASI benefit claiming, using the HRS data. 18 stratification dimensions are considered separately; they can be grouped into 3 broad categories: socioeconomic; pension and job characteristics; and cognition and behavior. The parameter of interest θ_2 is reported as the coefficient on the “FRA*interaction term” line. White workers tend to respond more than non-whites (+8 percentage points), although the difference is not statistically significant. Differences in the impact of the FRA along other socioeconomic dimensions – sex, marital status, education, and health – are small and/or statistically insignificant. In the second panel, the type of pension and the level of wealth appear to matter: the FRA impact is significantly lower among workers whose current job provides a defined benefit pension and significantly higher among workers holding above median wealth. Measures

of cognitive ability shown in the last panel matter too: having a high TICS score increases the response to the FRA by 9.6 pp; better memory and higher numeracy scores increase it by 5.7 pp (not statistically significant) and 12.2 pp., respectively. Overall, this suggests that wealthier and more cognitively skilled workers are “more behavioral” in the sense of following the FRA more closely. This seems to go against a bounded rationality explanation whereby workers with lower cognitive skills would follow the FRA as a default solution, or consider it as advice from the SSA. It therefore tends to point in the direction of non-standard preferences – like reference dependence and loss aversion – even though this still begs the question of why workers with higher cognitive skills would more strongly display such non-standard preferences.²² The impact of defined benefit pension coverage could reflect the fact that these pension plans maintain a normal age of 65, reducing the salience of the change in the FRA.

A limitation of model (5.1) is that it does not allow us to determine which of the many interaction effects reported separately in table 5 are the most important. Interacting all 18 characteristics with the FRA indicator, the set of cohort dummies and the set of age dummies is not feasible due to limited sample size and multicollinearity. A more parsimonious specification introduces interactions with the most promising dimensions of heterogeneity, based on table 5: pension characteristics, wealth, cognitive measures, and the basic socioeconomic variables. The various specifications are given in table 6. Multicollinearity pushes the standard errors up, so that few interaction terms remain significant. In column 1, where nine dimensions of heterogeneous impact are considered simultaneously, the only significant effects are due to holding a defined

²² Benjamin, Brown, and Shapiro (2006) find that small-stakes risk aversion and short run discounting are less common among those with higher cognitive ability. These findings might conflict with our results, to the extent that non-standard preferences share a common component across domains.

benefit pension, and to high memory. The wealth effect shrinks somewhat, suggesting that it was picking up in part the effect of cognitive ability. Combining the cognitive measures into a single index confirms the positive impact of cognitive skills (column 2).²³ Dropping the DB variable (only available for a subsample of workers) does not significantly alter the results (column 3).

Of course, it may still be the case that these interaction effects are driven by unobserved sources of heterogeneity. However, as noted above, a plausible causal interpretation of the DB interaction effect is that the presence of a DB pension reduces the salience of the SS FRA, in particular when the DB pension plan maintains a normal retirement age at 65. Accordingly, when we restrict the sample to DB holders we find that the responsiveness to the SS FRA is lower for those with a DB plan that has a normal retirement age at 65.²⁴ By contrast, the negative impact of low cognitive skills may seem harder to interpret causally. We first check that it is not capturing the impact of stressful or otherwise demanding job characteristics (column 4), and that it is not due to a lower sensitivity to the removal of the SS earnings test (column 5).²⁵ Checking for non linear effects, we find that most of the cognition effect is due to the lowest quartile, with smaller differences among the upper three quartiles (column 6). Furthermore, the effect is weaker for cohorts 1940+ compared to 1938-39 (column 7). This can be traced to the persistence of a large age-65 spike for people with lower cognitive skills. As show by figure 5 (replicating figure 2 by cognitive skill group), the age profiles of the hazard rate of claiming are very similar for people with higher and lower cognitive skills born in 1937. Spikes at the new FRAs appear for cohorts

Their paper does not analyze reference dependence. Also, we do not find significant differences by levels of risk aversion.

²³ See the Appendix for details on the cognition index.

²⁴ The coefficient on FRA*(DB NRA=65) is -.095 (marginally significant with a standard error of .059).

born after 1938, but, for people with lower cognitive skills, the spike at the FRA is somewhat smaller and the spike at age 65 persists. This suggests that people with lower cognitive skills may be slower to learn about the change in the FRA and adapt it into their decision making. Instead, they may use workers from earlier cohorts as a reference.^{26 27}

6. Implications of the results for framing Social Security reform

We interpret the empirical results presented above as suggesting that reference dependence is a factor in claiming and retirement decisions. In this section, we introduce reference dependence in a lifetime labor supply model in order to draw out its implications for framing of Social Security reforms. Specifically, we derive conditions under which reference dependence leads to a greater increase in employment in response to a benefit cut than would be predicted by the wealth effect alone. The model echoes the way the SS administration frames the retirement decision, as a tradeoff between income and “years to enjoy it”.

The set up is as simple as possible. Workers choose their optimal retirement and claiming age (assumed to be the same for simplicity),²⁸ by trading off years of leisure l against lifetime consumption c . The age at death (T) is fixed and known, so choosing retirement age R is

²⁵ Note that there is no significant difference across the two cognitive groups in the response to the removal of the SS earnings test.

²⁶ This section has focused on claiming behavior, for which there is strong and robust evidence of responsiveness to the FRA. Analysis of employment exit and self-reported retirement indicates that there is little heterogeneity that can be detected in the FRA effect on these outcomes (results available from the authors).

²⁷ Other explanations are possible. One would be that the two groups take the FRA as the claiming age recommended by the SSA, but only workers with higher cognitive skills read their statements and learn about the new FRA. We consider this explanation as less plausible given the care taken by the SSA not to imply any advice in their phrasing of the leisure / consumption tradeoff.

²⁸ This assumption implies that we ignore the lower bound on claiming age. In the simulations described below, we account for it.

equivalent to choosing lifetime leisure: $l = T - R$ (for convenience, we assume life begins at labor force entry). The budget constraint is: $c = k + wR$, where $k + wR$ is a linear approximation (in the vicinity of the FRA) of the lifetime income derived from retiring at age R ; k (initial wealth) and w (annual compensation, including the wage and the increment to the Social Security benefit resulting from an additional year of work) are fixed parameters in this approximation. This yields the standard static labor supply model, interpreted as a model of lifetime labor supply:

$$\begin{aligned} & \max_{c,l} U(l, c) \\ \text{s.t. } & c = k + w(T - l). \end{aligned} \tag{6.1}$$

The SS rules and statements suggest a specific age, the FRA, as a reference. Let c_{FRA} and l_{FRA} denote the levels of consumption and leisure from retiring and claiming at the FRA. Workers may experience loss aversion with respect to either leisure or consumption or both: they may be reluctant to reduce the number of “years to enjoy retirement” below the number implied by retiring at their FRA, and they may be reluctant to consume less than the level implied by retiring at the FRA. Following Tversky and Kahneman (1991), we incorporate reference dependence in a two-good model with no uncertainty by specifying the payoff from choice (c, l) as:

$$U_{FRA}(l, c) = aR_1(l) + R_2(c), \tag{6.2}$$

with

$$\begin{aligned} R_1(l) &= \lambda_1(u(l) - u(l_{FRA})) \quad \text{if } l \leq l_{FRA} \\ R_1(l) &= u(l) - u(l_{FRA}) \quad \text{if } l > l_{FRA} \end{aligned}$$

and

$$\begin{aligned} R_2(c) &= \lambda_2(v(c) - v(c_{FRA})) \quad \text{if } c \leq c_{FRA} \\ R_2(c) &= v(c) - v(c_{FRA}) \quad \text{if } c > c_{FRA}. \end{aligned}$$

λ_1 and λ_2 are the coefficients of loss aversion, with $\lambda_1, \lambda_2 > 1$ if there is loss aversion, and $\lambda_1 = \lambda_2 = 1$ otherwise. $u(\cdot)$ and $v(\cdot)$ are increasing and concave utility subfunctions. Lastly,

$a > 0$ is an individual-specific parameter that allows for heterogeneity in the preference for leisure.²⁹

This specification captures an asymmetry in preferences with regard to losses and gains around the FRA reference. Starting from the reference set by the FRA, the marginal utility of increasing consumption by one dollar is $v'(c_{FRA})$ whereas the utility loss from decreasing consumption by one dollar is $\lambda_2 v'(c_{FRA})$. Similarly, increasing leisure time by one day increases utility by $au'(l_{FRA})$, whereas reducing it by one day decreases utility by $a\lambda_1 u'(l_{FRA})$. Both dimensions of loss aversion increase the likelihood that the FRA is the optimal retirement age. It is straightforward to show that the solution to problem 6.1 can be characterized by two critical

values of a : individuals with low preference for leisure ($a < \frac{w}{\lambda_1} \frac{v'(c_{FRA})}{u'(l_{FRA})}$) retire after the full

retirement age; those with high preference for leisure ($a > w\lambda_2 \frac{v'(c_{FRA})}{u'(l_{FRA})}$) retire before the full

retirement age; and workers with intermediate preferences for leisure

($a \in \left[\frac{w}{\lambda_1} \frac{v'(c_{FRA})}{u'(l_{FRA})}, w\lambda_2 \frac{v'(c_{FRA})}{u'(l_{FRA})} \right]$) retire exactly at the FRA. These three cases are illustrated in

figure 6, which plots U_{FRA} as a function of l , after substituting for consumption from the budget constraint. U_{FRA} has a kink at l_{FRA} . This kink generates a mass point at the FRA in the distribution of retirement ages. Let F denote the c.d.f. of a , and P_{FRA} denote the fraction of workers retiring at

²⁹ We introduce a as the only source of heterogeneity, and derive the distribution of retirement ages from the distribution of a . One could introduce other sources of heterogeneity, either in preferences – λ_1 , λ_2 , $u(\cdot)$ and $v(\cdot)$ may vary across individuals – or in budget constraints – variations in k , w or T . However, these other sources of heterogeneity have similar implications for the retirement age distribution: as long as they are continuously distributed (so that they do not generate a kink in

the FRA. Then:

$$P_{FRA} = F\left(w\lambda_2 \frac{v'(c_{FRA})}{u'(l_{FRA})}\right) - F\left(\frac{w}{\lambda_1} \frac{v'(c_{FRA})}{u'(l_{FRA})}\right). \quad (6.3)$$

The fact that P_{FRA} is strictly positive if $\lambda_1 > 1$ or $\lambda_2 > 1$ shows that loss aversion in either of the two dimensions (loss of leisure or loss of benefits) is enough to generate the spike. However, these two dimensions have opposite impacts on the rest of the retirement age distribution.

Starting from a situation without loss aversion ($\lambda_1 = \lambda_2 = 1$), an increase in λ_1 attracts workers who would otherwise work longer toward the FRA, thus reducing the average retirement age (see the model appendix for details). By contrast, an increase in λ_2 attracts workers who would otherwise retire earlier toward the FRA, thus increasing the average retirement age. Overall, the impact of reference dependence and loss aversion on the average retirement age is ambiguous *a priori*.

Impact of the 1983 reform

The 1983 reform, as framed by the SSA, can easily be incorporated into the model as a change in the reference age. In order to maintain the same level of benefits, workers born after 1937 must delay retirement by (FRA - 65 years). For instance, in order to receive 100% of the primary insurance amount, workers born in 1937 must claim at age 65, whereas workers born in 1943 must delay claiming until age 66. In the model's notation, the reform is such that $\Delta c_{FRA} = 0$ and $\Delta l_{FRA} < 0$. All other things equal, this has two effects on the retirement age distribution:

First, the spike in the retirement hazard shifts to the new FRA. Second, the probability of retiring

preferences or the budget constraints), they cannot by themselves account for a spike in the retirement

before the FRA increases, whereas the probability of retiring after the FRA decreases. The combined effect is an increase in the average retirement age. We now ask whether a different framing of the reform would have yielded different results. Specifically, how does $dE(R)/dk$, the response of the average retirement age R to a given shift in the intercept of the benefit schedule dk , vary with the way the reform is framed? In all cases, we have

$$E(R) = T - E(l) = T - \int l^*(a) f(a) da, \quad (6.4)$$

where f is the density of a , and $l^*(a)$ is the level of leisure chosen by a worker given his preference for leisure. The quantity we are interested in is

$$\frac{dE(R)}{dk} = - \int \frac{dl^*(a)}{dk} f(a) da. \quad (7.5)$$

The first framing option we consider is neutral: your benefit schedule is lower than the schedule of your older peers, without reference to a specific age (see Figure 1). In this case the response of $l^*(a)$ to the reform is simply given by differentiating the standard first order conditions for an interior solution, yielding the standard wealth effect (see the model appendix).³⁰

Under the other two framing options, a reference point is explicitly given so that the utility of a worker with preference for leisure a is described by equations 6.2 above, and the optimal retirement age is characterized by solving the first order conditions for the cases in

which $a < \frac{w}{\lambda_1} \frac{v'(c_{FRA})}{u'(l_{FRA})}$ and $a > w\lambda_2 \frac{v'(c_{FRA})}{u'(l_{FRA})}$, and by setting $l = l_{FRA}$ if

hazard.

³⁰ This only holds if there was no loss aversion and framing before the reform. Otherwise, it is unclear how a neutral framing of the reform would be perceived: would it cancel the initial reference? If

$a \in \left[\frac{w}{\lambda_1} \frac{v'(c_{FRA})}{u'(l_{FRA})}; w\lambda_2 \frac{v'(c_{FRA})}{u'(l_{FRA})} \right]$. The two framing options differ by the fact that only the third

framing option implies a change in the reference point.

In the second framing option, workers have the same reference point after the reform (l_{FRA}). However, SSA tells them that there is a cut in benefits for claiming at the FRA. In other words, they still perceive that they meet a target if they retire at 65 after the reform, but the target is now lower. This arises if the reform is framed as a change in the PIA (the benefit amount available if claimed at the FRA) with no change in the FRA.

Finally, the third framing option is the one actually mandated by the reform: a change in the FRA with no mention of a benefit cut. In this case, things remain as in the 2nd framing option for workers with low and high preference for leisure. Things do however differ for workers with intermediate preferences for leisure. The condition for claiming at the FRA is the same as under the second framing option. However, the FRA itself changes, with $dl_{FRA} = \frac{1}{w} dk$.

The model appendix shows that $\left[\frac{dE(R)}{dk} \right]_3 > \left[\frac{dE(R)}{dk} \right]_2$, where the subscript indicates the framing option. In the presence of reference dependence, the reform has a stronger impact if it is framed as a change in the reference point than if it is framed as an equivalent change in the benefit at the reference point. The comparison with $\left[\frac{dE(R)}{dk} \right]_1$ (neutral framing) is less immediate; it depends on the parameters. However, empirical estimates suggest that

not, and if workers keep the initial reference point, the first framing option would be equivalent to the second option, described below.

$\left[\frac{dE(R)}{dk} \right]_3 > \left[\frac{dE(R)}{dk} \right]_1$ ³¹. Indeed, Mastrobuoni (2009) finds that the average response to a 1 year increase in the FRA (under the 3rd framing option, which is how it was framed by SSA) is a .5 year increase in the average retirement age. For workers at the FRA, loss aversion implies a 1 for 1 response. In this model, the .5 response must be a weighted average of 1 for people at the FRA and δ (the average response for people above and below the FRA). This implies that $\delta < 1$. More precisely, the magnifying effect under the 3rd framing option is positively correlated with the share of the population clustered at the FRA. Assuming that the response is roughly constant for other workers, we have indeed

$$\left[\frac{dE(R)}{dk} \right]_3 \approx -\frac{1}{w} [(1 - P_{FRA})\delta + P_{FRA}] = -\frac{1}{w} [\delta + P_{FRA}(1 - \delta)]. \quad (6.6)$$

In sum, compared to a situation without a reference point ($P_{FRA} = 0$), reference dependence magnifies the impact of a reform if the reform is expressed as a change in the reference point. This magnifying effect increases with the share of the population initially clustered at the reference point. Figure 7 summarizes the main lesson of the model using simulated data. We use CRRA specifications for u and v and assume that leisure preference a follows a log-normal distribution with mean 1. We set $w=1$ and $k=0$, and simulate the model for various values of the risk aversion coefficients in u and v and the variance of the log-normal distribution. We select only those simulations that yield plausible distributions of the retirement

³¹ In our simulations discussed below, $\left[\frac{dE(R)}{dk} \right]_1$ and $\left[\frac{dE(R)}{dk} \right]_2$ are roughly similar.

age.³² We then take averages across the corresponding simulations.³³ The graphs show the simulated impact of loss aversion (x-axis) on (i) the average retirement age before the reform (top row); (ii) the share of workers retiring exactly at the FRA before the reform (middle row); and (iii) the impact of the reform on the average retirement age (bottom row). The latter impact is analyzed under two framing options: the reform is framed as an increase in the FRA (dark full line, “framing option number 3”) or the reform is framed as a cut in the PIA (dashed grey lines, “framing option number 2”). The first framing option (neutral framing) corresponds to the point with $\lambda_1 = \lambda_2 = 1$ on the x-axis. At that point, there is no loss aversion, and framing does not matter (assuming, as discussed above, that there was no framing effect before the reform: even though workers would display loss aversion if there was a reference, when no reference is given everything happens as if people were not loss averse). The three columns correspond to different assumptions on loss aversion: in column 1, workers suffer only from loss aversion with regard to leisure time below their reference; in column 2, workers suffer only from loss aversion with regard to benefits below their reference; in columns 3, the two types of loss aversion are combined.

Looking at the first row illustrates the fact that the impact of loss aversion on the average retirement age is, in general, indeterminate. As noted above, aversion with regard to loss of leisure reduces the average retirement age (top graph on the left), but aversion with regard to loss of benefits increases it (middle graph on the top). The combined effect is, in general, almost 0 (top graph on the right). By contrast, the two types of loss aversion both cause an increase in the size of the spike at the FRA (middle row). Finally, the bottom row shows the crucial role played

³² Specifically, we drop simulations with sets of parameters that lead to less than 72% of workers retiring before the FRA, and more than 20% retiring after the FRA, in the absence of loss aversion ($\lambda_1 = \lambda_2 = 1$).

³³ Results using each set of parameters individually are very similar.

by framing in the presence of loss aversion. In the absence of framing effects (which corresponds to $\lambda_1 = \lambda_2 = 1$), the reform – modeled here as a 1-year increase in the FRA or a 6.7% benefit cut (as in the 1983 reform) – increases the average retirement age by less than one month (this is represented by the point $\lambda_1 = \lambda_2 = 1$ on the horizontal axis in the bottom three graphs). In the presence of strong loss aversion, framing the reform as an increase in the FRA strongly magnifies the impact, to about 4 months if $\lambda_1 = \lambda_2 = 3$ (bottom right graph, solid line). By contrast, the impact of the reform is about the same as in the no-loss-aversion case if the reform is framed as a cut in the PIA (bottom right graph, dashed line). The loss aversion parameters λ_1 and λ_2 have strong implications for the impact of the SS reform on retirement age. Even though they are hard to estimate, the simulations suggest that the amplifying effect is roughly proportional to the share of workers clustered at the FRA. This motivated our empirical attempt above at quantifying the magnitude of the spike at the FRA.

Is loss aversion enough to explain the unexpectedly strong impact of the 1983 reform (Mastrobuoni, 2009)? Our analysis suggests that reference dependence matters, by shifting the workers clustered at the old reference point toward the new reference point. However, as noted above, the share of workers who have not retired or claimed benefits by the age of 65 is small, so the impact on the average retirement age remains modest.³⁴ Moreover, Mastrobuoni's results show that the 1983 reform also strongly affected the retirement distribution at ages 62 to 64, suggesting that while loss aversion has probably magnified the impact of the reform, it may not fully explain the large effect of the reform.

³⁴ For instance, if 10% of workers initially retired at their FRA, the amplifying effect due to these workers would be less than .1 years (10% of workers postponing retirement by 1 year).

7. Conclusion

This paper has used the 1983 Social Security reform as a quasi-experiment to provide evidence on framing effects in retirement behavior. From a methodological perspective, the FRA is particularly well-suited to study reference dependence: by contrast with other applications of reference dependence to labor supply analysis, the reference is explicitly defined and then exogenously modified by the 1983 reform. The FRA impact is unambiguously identified by cohort discontinuities. Although one cannot fully rule out alternative behavioral explanations such as social norms or reliance on SSA “advice”, the latter explanations seem at odds with the fact that workers with higher cognitive ability respond more to the FRA change. Responsiveness to the FRA does not seem to be due to unsophisticated decision making.

Framing effects have been well documented in the related domain of pension plan choice. Our results indicate that they exist in benefit claiming and retirement as well. Given that around 3 workers out of 4 have already claimed SS benefits before the FRA, the aggregate impact of loss aversion in the context of the 1983 reform remains modest. However, the mechanisms at play are quite general and have potentially important implications for framing of future reforms, in the same way as findings on savings and pension plan decisions have led to “behavioral institutional design” recommendations (see, e.g., Benartzi and Thaler, 2004). Indeed, as shown by a simple extension of the standard labor supply model, framing provides the decision maker with a potentially powerful and almost costless tool to influence aggregate labor supply: framing a reform as a change in the reference point magnifies the impact, whereas framing it as a benefit cut dampens the response. How to use this knowledge depends on the goals of reform, and this

suggests an important avenue for research: how should policy makers take into account loss aversion in designing future reforms?

References

- Abowd, John M., Bryce E. Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L. McKinney, Marc Roemer, and Simon Woodcock (2006). "The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators," LEHD Technical Paper 2006-01, January.
- Abowd, John M., John Haltiwanger, and Julia Lane (2004). "Integrated Longitudinal Employee-Employer Data for the United States," LEHD Technical Paper 2004-02, May.
- Baker, Michael, and Dwayne Benjamin (1999). "Early Retirement Provisions and the Labor Force Behavior of Older Men: Evidence from Canada," *Journal of Labor Economics* 17 (4), Part 1, October: 724-756.
- Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro (2006). "Who is 'Behavioral'? Cognitive Ability and Anomalous Preferences." Harvard University, mimeo, 2006.
- Benartzi, Shlomo and Richard Thaler. 2004. "Save More Tomorrow (TM): Using Behavioral Economics to Increase Employee Saving." *Journal of Political Economy*, 112(S1): S164-S187.
- Blau, David M., and Donna Gilleskie (2008). "The Role of Retiree Health Insurance in the Employment Behavior of Older Men" *International Economic Review*, 49 (2): 475-514.
- Blau, David M., and Donna Gilleskie (2006). "Health Insurance and Retirement of Married Couples" *Journal of Applied Econometrics*, 21, Nov.: 935-953.
- Blau, David M., and Ryan Goodstein (2010). "Can Social Security Explain Trends in Labor Force Participation of Older Men in the United States?" *Journal of Human Resources*. 45 (2): 326-63.
- Blau, David M., and Tetyana Shvydko (in press). "Labor Market Rigidities and the Employment Behavior of Older Workers," *Industrial and Labor Relations Review*.
- Burtless, G. (2004) "Social Norms, Rules of Thumbs, and Retirement: Evidence for Rationality in Retirement Planning", mimeo, The Brookings Institution, October 2004.
- Chan, Sewin, and Ann Huff Stevens (2008). "What You Don't Know Can't Help You. Pension Knowledge and Retirement Decision-Making," *Review of Economics and Statistics*, May.
- Coile, Courtney C., Peter Diamond, Jonathan Gruber, and Alain Jouten (2002). "Delays in Claiming Social Security Benefits," *Journal of Public Economics*, 84 (3), July: 357-385.
- Coile, Courtney, and Jonathan Gruber (2007). "Future Social Security Entitlements and the Retirement Decision," *Review of Economics and Statistics* 89 (2): 234-46.

Costa, Dora (1998). The Evolution of Retirement: An American Economic History, 1880-1990, Chicago: University of Chicago Press.

Della Vigna, Stefano (2009) "Psychology and Economics: Evidence from The Field", *Journal of Economic Literature*, June, Vol. 47, pp. 315-372.

Featherstonehaugh, David, and Lee Ross (1999). "Framing Effects and Income Flow Preferences in Decisions About Social Security," in Henry Aaron (ed.) Behavioral Dimensions of Retirement Economics Washington DC, Brookings Institution Press and Russell Sage Foundation: 187-209.

Friedberg, Leora (2000). "The Labor Supply Effects of the Social Security Earnings Test," *Review of Economics and Statistics* 82 (1): 48-63.

Gruber, Jonathan, and Peter Orszag (2003). "Does the Social Security Earnings Test Affect Labor Supply and Benefits Receipt?" *National Tax Journal* 56 (4): 755-73.

Haider, Steven J., and David S. Loughran (2008). "The Effect of the Social Security Earnings Test on Male Labor Supply: New Evidence from Survey and Administrative Data," *Journal of Human Resources* 43 (1):57-87.

Kopczuk, Wojciech, and Jae Song (2008). "Stylized Facts and Incentive Effects Related to Claiming of Retirement Benefits Based on Social Security Administration Data," Michigan Retirement Research Center Working Paper 2008-200, October.

Koszegi Botond, and Matthew Rabin (2006) "A Model of Reference-Dependent Preferences", *Quarterly Journal of Economics*, 122(4), 1133-1165.

Liebman J., E. Luttmer and D. Seif (2008), "Labor Supply Responses to the Social Security Tax-Benefit Link", Mimeo, July 2008.

Liebman, Jeffrey B., and Erzo F.P Luttmer (2009). "The Perception of Social Security Incentives for Labor Supply and Retirement: The Median Voter Knows a Lot More Than You'd Think," Working paper, Harvard University, April.

Lumsdaine, Robin L., James H. Stock, and David A. Wise (1996). "Why Are Retirement Rates So High at Age 65?" in David A. Wise (ed.) Advances in the Economics of Aging Chicago: University of Chicago Press: 61-82.

Madrian, Brigitte C., and Dennis F. Shea. (2001). "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior," *Quarterly Journal of Economics* 116:1149-87.

- Manoli, Dayand, Kathleen Mullen, and Mathis Wagner (2009). “Risk Aversion and Retirement Decisions: Using Policy Variation to Identify and Estimate a Structural Model of Retirement,” Working Paper, UCLA.
- Mastrobuoni, Giovanni (2009). “Labor Supply Effects of the Recent Social Security Benefit Cuts: Empirical Estimates Using Cohort Discontinuities,” *Journal of Public Economics* 93 (11-12), December: 1224-33.
- Mastrobuoni, Giovanni (2010). “The Role of Information in Retirement Behavior: Evidence Based on the Stepwise Introduction of the Social Security Statement,” Working Paper, Collegio Carlo Alberto, January.
- McArdle, John, James Smith and Robert Willis (2009). “Cognition and Economic Outcomes in the Health and Retirement Survey”, *NBER working paper*, 15266, August.
- Moffitt, Robert A. (1987). “Life Cycle Labor Supply and Social Security: A Time Series Analysis.” In *Work, Health, and Income among the Elderly*, ed. Gary Burtless, 183–220. Washington, D.C.: The Brookings Institution.
- Perrachi, Franco, and Finis Welch (1994). “Trends in Labor Force Transitions of Older Men and Women,” *Journal of Labor Economics* 12 (2): 210-42.
- Poterba, James, Steven Venti, and David A. Wise (2007). “The changing landscape of pensions in the United States,” NBER Working Paper 13381, September.
- Pingle, Jonathan F. (2006). “Social Security’s Delayed Retirement Credit and the Labor Supply of Older Men,” Working Paper, Board of Governors of the Federal Reserve System, June.
- Rust John, and Christopher Phelan. (1997). “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets,” *Econometrica* 65: 781-831.
- Samwick, Andrew A. (1998). “New Evidence on Pensions, Social Security, and the Timing of Retirement,” *Journal of Public Economics* 70: 207-36.
- Social Security Administration (2008a) <http://www.ssa.gov/mystatement/currentstatement.pdf>
- Social Security Administration (2008b) <http://www.ssa.gov/mystatement/insert55+.pdf>
- Song, Jae, and Joyce Manchester (2008). “Have People Delayed Claiming Social Security Benefits? Responses to Changes in Social Security Rules,” Congressional Budget Office, Working Paper 2008-04, May.
- Tversky, Amos, and Daniel Kahneman, 1991. “Loss Aversion in Riskless Choice: A Reference Dependent Model”, *Quarterly Journal of Economics* 106(4): 1039-1061.

Data appendix

The main features of the data are described in section 4. This appendix describes in more detail the variables used to differentiate cognitive and behavioral types. Most variables are derived from the raw HRS files.

Cognitive variables:

High TICS score. The TICS variable provides an objective measure of the respondent's memory and ability to think quickly. This is the sum of points (1 for each correct answer) for a set of questions the respondent has answered; it ranges from 0 to 10. The variable was only collected in waves 3-8 of the HRS. We use a dummy for high TICS scores (above 9).

High episodic memory. As part of the interview the interviewer asks the respondent to listen and then recall a list of words. This variable indicates how many words the respondent can recall immediately after hearing them. In waves 1 and 2, the list included 20 words. In subsequent waves, the list included 10 words. The score for waves 1 and 2 was divided by 2 for comparability. A second variable indicates how many words the respondent can recall some time after hearing the list, i.e., after answering a few other questions. We average the two scores. We then use a dummy for high memory: a score higher than 5 out of 10.

High numeracy. The numeracy variable is the number of correct answers to the following three questions: Correctly answer the question "If 5 people all have the winning numbers in the lottery and the prize is two million dollars, how much will each of them get?"; "If the chance of getting a disease is 10 percent, how many people out of 1,000 would be expected to get the disease?"; "Let's say you have \$200 in a savings account. The account earns 10 percent interest per year. How much would you have in the account at the end of two years?". These questions were only asked during waves 6 to 8 of the HRS. We generate an indicator variable for high numeracy, indicating at least two correct answers out of three.

We construct a cognitive index equal to a weighted sum of the TICS, memory and numeracy scores (putting a weight 3.3 times as large on numeracy, to compensate for the smaller range of values). The index can go from 0 to 30.

Behavioral variables

High self control. We create a self-control score using a series of 5 questions such as "I feel that I can do as I please. (Often? Sometimes? Not often? Never?)." The maximum score is 19, and we generate an indicator for high self control that equals 1 for people with a score higher than 6.

The next two variables come from the RAND release of the HRS:

Long financial planning horizon. The "long financial planning horizon" dummy is equal to 1 for those who respond to plan at least for the next few years.

Risk averse. The “risk averse” dummy is equal to 1 for workers classified as the most risk averse (out of four categories).

Working condition variables

These are from the following questions: “I’ll read some statements that are true for some people’s jobs but not for other people’s jobs. Thinking of your job, please tell how often these statements are true.”

- My job requires lifting heavy loads.
- My job requires lots of physical effort.
- My job requires stooping, kneeling, or crouching.
- My job requires good eyesight.
- My job involves a lot of stress.

Stressful job is an indicator variable for responding “all or almost all of the time” or “most of the time” to the last item.

Physically demanding job is an indicator variable for being in the top half of the sample for a synthetic variable summing answers to the first three items.

Model Appendix

Here we show the results stated in section 6.

Starting from a situation without loss aversion ($\lambda_1 = \lambda_2 = 1$), an increase in λ_1 attracts workers who would otherwise work longer toward the FRA, thus reducing the average retirement age: $dP_{FRA-} / d\lambda_1 = 0$, $dP_{FRA} / d\lambda_1 > 0$ and $dP_{FRA+} / d\lambda_1 < 0$. An increase in λ_2 attracts workers who would otherwise retire earlier toward the FRA, thus increasing the average retirement age: $dP_{FRA-} / d\lambda_2 < 0$, $dP_{FRA} / d\lambda_2 > 0$ and $dP_{FRA+} / d\lambda_2 = 0$.

Impact of the 1983 reform: the probability of retiring before the FRA (P_{FRA-}) increases, whereas the probability of retiring after the FRA (P_{FRA+}) decreases. This can be directly seen from

$$P_{FRA-} = 1 - F\left(w\lambda_2 \frac{v'(c_{FRA})}{u'(l_{FRA})}\right); \quad (\text{A.1})$$

$$P_{FRA+} = F\left(\frac{w}{\lambda_1} \frac{v'(c_{FRA})}{u'(l_{FRA})}\right), \quad (\text{A.2})$$

noting that $d \frac{v'(c_{FRA})}{u'(l_{FRA})} = -\frac{v'(c_{FRA}) \times u''(l_{FRA})}{(u'(l_{FRA}))^2} dl_{FRA} < 0$.

Under framing option 1: $\frac{dl^*(a)}{dk} = \frac{au''(l) + w^2v''(W - wl)}{wv''(W - wl)}$ so that

$$\left[\frac{dE(R)}{dk}\right]_1 = -\int \frac{au''(l) + w^2v''(W - wl)}{wv''(W - wl)} f(a) da, \text{ where } W = k + wT,$$

Under framing options 2 and 3 the optimal retirement age is

$$l^* = \begin{cases} l & \text{st } a\lambda_1 u'(l) - wv'(W - wl) = 0 & \text{for } a < \underline{a} \\ l_{FRA} & & \text{for } a \in [\underline{a}; \bar{a}] \\ l & \text{st } au'(l) - w\lambda_2 v'(W - wl) = 0 & \text{for } a < \bar{a} \end{cases}$$

with $\underline{a} \equiv \frac{wv'(W - wl_{FRA})}{a\lambda_1 u'(l_{FRA})}$ and $\bar{a} \equiv \frac{w\lambda_2 v'(W - wl_{FRA})}{au'(l_{FRA})}$. In the second framing option, due to the

reduction in benefits at the FRA, the thresholds for being below, at or above the FRA become

$$\underline{a}' \equiv \frac{wv'(W + dW - wl_{FRA})}{a\lambda_1 u'(l_{FRA})} \text{ and } \bar{a}' \equiv \frac{w\lambda_2 v'(W + dW - wl_{FRA})}{au'(l_{FRA})}. \text{ However, for a marginal change } dk,$$

the probability that $a \in [\underline{a}'; \underline{a}]$ or $a \in [\bar{a}'; \bar{a}]$ is 0, so that we can neglect these marginal workers

when computing the impact of the reform on the expected retirement age. We are left with three

categories of workers. Those with low preference for leisure respond to the benefit cut

$$\text{by } \frac{dl^*(a)}{dk} = \frac{a\lambda_1 u''(l) + w^2 v''(W - wl)}{wv''(W - wl)}, \text{ and those with high preference for leisure respond}$$

$$\text{by } \frac{dl^*(a)}{dk} = \frac{au''(l) + w^2 v''(W - wl)}{w\lambda_2 v''(W - wl)}. \text{ However, workers with intermediate preference for leisure}$$

do *not* respond to the benefit cut: setting aside the marginal workers, they are such that

$a \in [\underline{a}'; \bar{a}']$ and they stay at the FRA. So, the impact of the reform under the second framing

option is

$$\left[\frac{dE(R)}{dk} \right]_2 = - \int_{-\infty}^{\underline{a}} \frac{a\lambda_1 u''(l) + w^2 v''(W - wl)}{wv''(W - wl)} f(a) da - \int_{\bar{a}}^{+\infty} \frac{au''(l) + w^2 \lambda_2 v''(W - wl)}{w\lambda_2 v''(W - wl)} f(a) da.$$

In the third framing option, workers with low and high preference for leisure respond as in the second option, but workers with intermediate preferences for leisure also respond. The condition for claiming at the FRA is the same as under the second framing option: $au'(l_{FRA}) = w\lambda_2 v'(W - wl_{FRA})$.

However, the FRA itself changes, with $dl_{FRA} = \frac{1}{w} dk$. Consequently, these workers delay

claiming by $\frac{1}{w} dk$. The impact of the reform on the mean retirement age in this case is

$$\left[\frac{dE(R)}{dk} \right]_3 = \left[\begin{array}{l} - \int_{-\infty}^a \frac{a\lambda_1 u''(l) + w^2 v''(W - wl)}{wv''(W - wl)} f(a) da \\ - \int_a^{+\infty} \frac{a u''(l) + w^2 \lambda_2 v''(W - wl)}{w\lambda_2 v''(W - wl)} f(a) da - \frac{1}{w} P_{FRA} \end{array} \right].$$



Thinking of retiring?

www.socialsecurity.gov

Some things to consider

Retirement can have more than one meaning these days. It can mean that you have applied for Social Security retirement benefits ... or that you are no longer working. Or it can mean that you have chosen to receive Social Security while still working, either full or part-time. All of these choices are available to you. And your retirement decisions can have very real effects on your ability to maintain a comfortable retirement.

If you retire early, you may not have enough income to enjoy the years ahead of you. Likewise, if you retire late, you'll have a larger income, but fewer years to enjoy it. Everyone needs to try to find the right balance, based on his or her own circumstances.

We hope the following information will help you as you plan for your future retirement and consider your retirement options.


It's so easy to apply online for benefits

The easiest way to apply for Social Security retirement benefits is to go online at www.socialsecurity.gov/applyforbenefits. If you do not have access to the Internet, you can call 1-800-772-1213 (TTY number, 1-800-325-0778) between 7 a.m. and 7 p.m., Monday through Friday, to apply by phone. You also can apply at any Social Security office. To avoid having to wait, call first to make an appointment.

What is the best option for you?

Everyone's situation is different. That is why Social Security has created several retirement planners to help you decide what would be best for you and your family. Social Security has a new online calculator that can provide immediate and accurate retirement benefit estimates to help you plan for your retirement.

The online Retirement Estimator is a convenient, secure and quick financial planning tool. It uses your own earnings record information, thereby eliminating any need to manually key in years of earnings information. The estimator also will let you create "what if" scenarios. You can, for example, change your "stop work" date or expected future earnings to create and compare different retirement options. To use the Retirement Estimator, go to our website at www.socialsecurity.gov/estimator.



There is one more thing you should remember as you crunch the numbers for your retirement. You may need your income to be sufficient for a long time, because people are living longer than ever before, and generally, women tend to live longer than men. For example:

- The typical 65-year-old today will live to age 83;
- One in four 65-year-olds will live to age 90; and
- One in 10 65-year-olds will live to age 95.

Once you decide on the best age for you to actually retire, remember to complete your application *three months before* the month in which you want retirement benefits to begin.

Don't forget Medicare

Even if you don't plan to receive monthly benefits, you should sign up for Medicare *three months before* reaching age 65. Otherwise, your Medicare medical insurance, as well as prescription drug coverage, could be delayed and you could be charged higher premiums. For more information about eligibility and costs, visit www.medicare.gov.

Receiving benefits while you work

When you reach your full retirement age, you can work and earn as much as you want and still receive your full Social Security benefit payment. If you are younger than full retirement age and if your earnings exceed certain dollar amounts, some of your benefit payments during the year will be withheld.

This does not mean you must try to limit your earnings. If we withhold some of your benefits because you continue to work, we will pay you a higher monthly benefit amount when you reach your full retirement age. In other words, if you would like to work and earn more than the exempt amount, you should know that it will not, on average, reduce the total value of lifetime benefits you receive from Social Security—and may actually increase them.

Here is how this works: after you reach full retirement age, we will recalculate your benefit amount to give you credit for any months in which you did not receive some benefit because of your earnings. In addition, as long as you continue to work, we will check your record every year to see whether the additional earnings will increase your monthly benefit.

Many people can continue to work and still receive retirement benefits. If you want more information on how earnings affect your retirement benefits, ask for *How Work Affects Your Benefits* (Publication No. 05-10069), which has current annual and monthly earnings limits, and is available on our website.

Retirement age considerations

Full retirement age

For persons born during the years 1943-1954, the full retirement age is 66. If you were not born in this period, you can find your full retirement age on page 2 of your *Social Security Statement*.

Retiring early

If you've earned 40 credits (credits are explained on page 2 of your *Statement*), you can start receiving Social Security benefits at 62 or at any month between 62 and full retirement age. However, your benefits will be reduced based on the number of months you receive benefits before you reach full retirement age.

If your full retirement age is 66, benefits will be reduced:
25 percent at age 62;
20 percent at age 63;
13 ½ percent at age 64; or
6 ½ percent at age 65.

Delaying retirement

You may decide to wait beyond your full retirement age before

choosing to receive benefits. If so, your benefit will be increased by a certain percentage for each month you don't receive benefits between your full retirement age and age 70. This table shows the rate your benefits increase if you delay retiring.

Year of birth	Yearly increase rate
1937 - 1938	6.5%
1939 - 1940	7.0%
1941 - 1942	7.5%
1943 or later	8.0%

Rules that may affect your survivor

If you are married and die before your spouse, he or she may be eligible for a benefit based on your work record. If you start benefits before your full retirement age, we cannot pay your surviving spouse a full benefit from your record. Also, if you wait until after your full retirement age to begin benefits, the surviving spouse benefits based on your record will be higher.

Need more information?

You can find answers to frequently asked questions about Social Security, learn about factors that could affect your benefits and much more by visiting Social Security online at www.socialsecurity.gov.

If you do not have access to the Internet, you can get information about Social Security by calling 1-800-772-1213 (1-800-325-0778 for the deaf or hard of hearing) or by visiting a local Social Security office.

Other useful websites

www.mymoney.gov

This website contains calculators for financial planning and information on money-related matters, such as retirement planning and starting a small business.

www.dol.gov/ebsa/pdf/nearretirement.pdf

Have you determined how much money you will need in retirement? There are many tools available to help you, such as the *Taking the Mystery Out of Retirement Planning Workbook* available at this link.

www.sec.gov/investor/seniors.shtml


Are you looking for information about the investment options available to you as you enter retirement? The Securities and Exchange Commission has a wealth of information on different investment products and topics available at this website.

www.usa.gov/topics/seniors.shtml

This website has a variety of resources for seniors on topics including retirement planning, housing and health.


www.socialsecurity.gov

Figure 1: Relationship between Social Security Benefit Claiming Age and Benefit level as a percent of the Primary Insurance Amount for Two Birth Cohorts




Source: Authors' calculation from Social Security rules.

Figure 2: SS Benefit Claiming Hazard




Notes: The graphs show the average monthly claiming hazard rates for pre and post-reform cohorts. The claiming hazard rate is defined as the probability of claiming at a given age, conditional on not having claimed previously. Age is measured at a bimonthly frequency; e.g. age 65 denotes age 65 0/12 to 65 1/12. In each graph, the dotted line depicts the claiming hazard for workers born between 1931 and 1936. For each cohort, the vertical lines indicate age 62, age 65, and the FRA (if different from age 65).

Figure 3: Hazard of Exit from Employment




Notes: see notes to figure 2.

Figure 4: Hazard of Retirement




Notes: see notes to figure 2. To keep the same scale on the vertical axis for 1942 as for other birth cohorts, the hazard rate at age 790 (65 10/12) has been arbitrarily set at .2. The observed value is .6 (over only 5 individuals).

Figure 5: SS Benefit Claiming Hazard
Workers with higher vs. lower cognitive skills



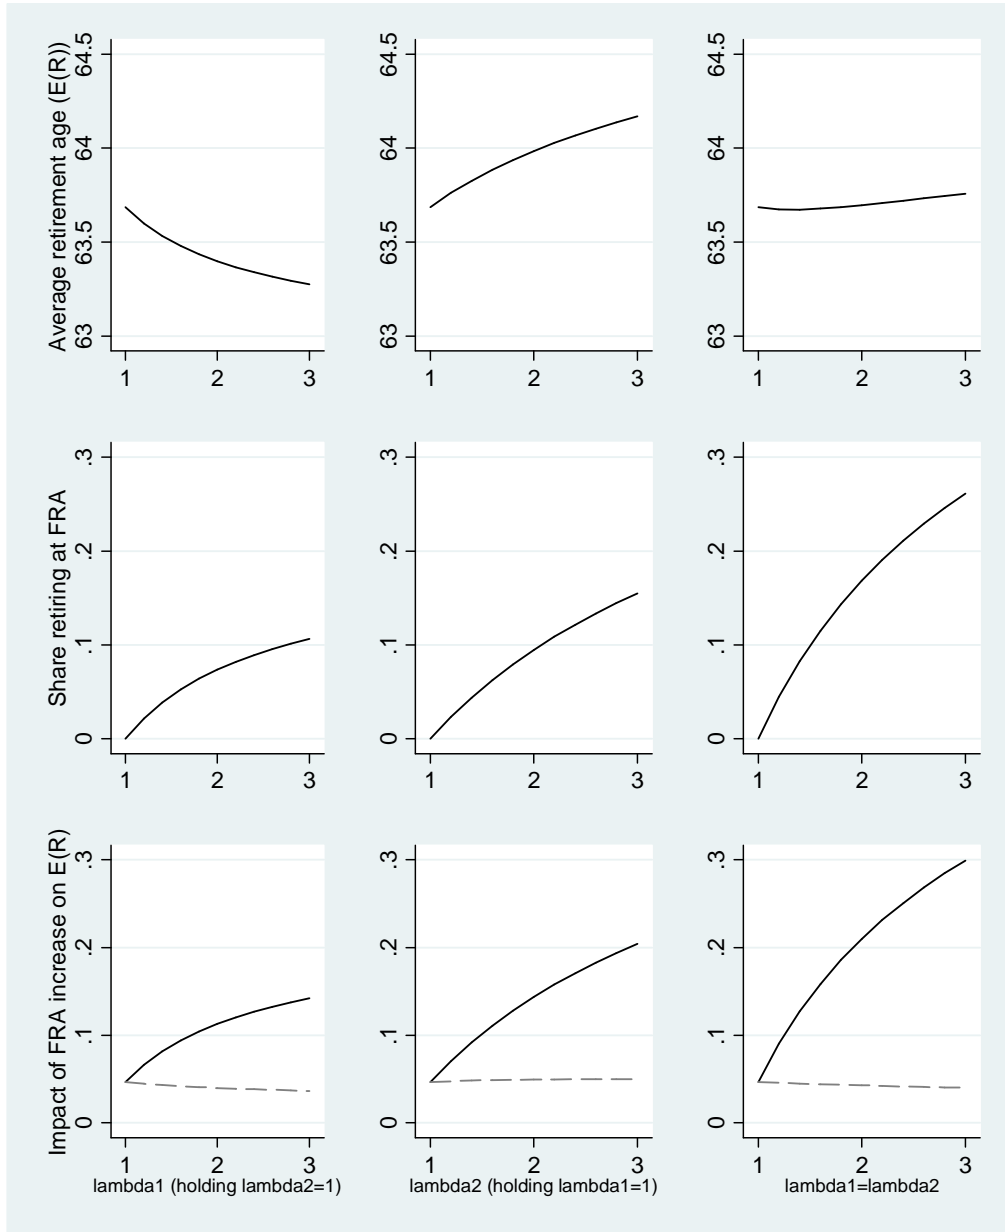
Notes: see notes to figure 2. People with a cognitive index above the median are in black dashed lines, people with a cognitive index below the median are in gray full lines.

Figure 6: Heterogeneous preference for leisure and optimal retirement age



Notes: see discussion in the text.

Figure 7: Simulated Effects of Framing with Different Degrees of Loss Aversion



Note: The graphs show the simulated impact of loss aversion (x-axis) on (i) the average retirement age before the 1983 reform (top row); (ii) the share of workers retiring exactly at the FRA before the 1983 reform (middle row); (iii) the impact of 1983 reform on the average retirement age (bottom row). The latter impact is analyzed under two framing options: the reform is framed as an increase in the FRA (dark full line) or the reform is framed as a cut of the PIA (dashed grey lines).

The three columns correspond to different assumptions on loss aversion: in column 1, values above 1 on the x-axis ($\lambda_1 > 1$) mean that workers suffer only from aversion with regard to cuts in leisure time below their reference; in column 2, values above 1 on the x-axis ($\lambda_2 > 1$) mean that workers suffer only from aversion with regard to cuts in benefits below their reference; in columns 3, the two types of loss aversion are combined. In all graphs, 1 on the x-axis ($\lambda_1 = \lambda_2 = 1$) is the case with no loss aversion. The slope of the curves in each graph shows the impact of increasing loss aversion.

Table 1: Impact of the FRA on OASI Benefit Claiming Hazard

	Claiming social security (OASI) benefits				
	(1)	(2)	(3)	(4)	(5)
FRA	13.8*** (1.9)	13.6*** (1.9)	13.3*** (1.9)	13.7*** (2.1)	13.6*** (1.9)
SS earnings test removal					3.2*** (1.0)
Controls	No	Yes	Yes	Yes	Yes
Cohorts	1932-41	1932-41	1932-41	1932-37; 1939-41	1932-41
Age range	64-66	64-66	62-66	64-66	64-66
N	25801	25801	89348	23570	25801
R ²	0.146	0.154	0.162	0.152	0.155

Notes: FRA is a dummy variable equal to 1 if the current month is the FRA and zero otherwise. The coefficients measure the percentage point increase in the SS benefit claiming monthly hazard at the FRA.

See equation (4.1) for the specification. The dependent variable is a monthly dummy for claiming social benefits. The models were estimated by OLS with standard errors clustered by individuals. Each regression includes a full set of monthly age dummies and birth cohort dummies. Controls in columns (2)-(5): measures of cognitive capability, planning horizon, risk aversion, pension and job characteristics, health and socio-demographic characteristics.

Sample: HRS waves 1992-2008. Age range and birth cohorts included in the regression differ by column. *, **, and *** indicate that the coefficient estimate is significantly different from zero at the 10%, 5%, and 1% level, respectively.

Table 2: Impact of the FRA on the Hazard of Exit from Employment

	Exit from Employment				
	(1)	(2)	(3)	(4)	(5)
FRA	1.0*** (0.4)	0.9** (0.4)	0.9** (0.4)	1.3*** (0.4)	1.0*** (0.4)
SS earnings test removal					0.5* (0.3)
Controls	No	Yes	Yes	Yes	Yes
Cohorts	1932-41	1932-41	1932-41	1932-37; 1939-41	1932-41
Age range	64-66	64-66	62-66	64-66	64-66
N	68952	68952	152753	62662	68952
R ²	0.016	0.047	0.054	0.047	0.047

Notes: see notes to Table 1.

Table 3: Impact of the FRA on the Hazard of Retirement

	Retiring				
	(1)	(2)	(3)	(4)	(5)
FRA	1.5 (1.6)	1.1 (1.6)	1.1 (1.6)	1.8 (2.0)	1.1 (1.6)
SS earnings test removal					2.2** (1.1)
Controls	No	Yes	Yes	Yes	Yes
Cohorts	1932-41	1932-41	1932-41	1932-37; 1939-41	1932-41
Age range	64-66	64-66	62-66	64-66	64-66
N	16387	16387	46737	14801	16387
R ²	0.056	0.082	0.085	0.083	0.083

Notes: see notes to Table 1.

Table 4: Impact of the FRA on the Hazard of Exit from Employment (LEHD data)

	Exit from Employment	
	(1)	(2)
FRA*(coh=1938)		-2.8*** (0.7)
FRA*(coh=1939)		-0.1 (1.6)
FRA	1.2** (0.5)	
Cohorts	1931-42	1931-42
Age range	62-66	62-66
N	2256	2256
R ²	0.97	0.97

Notes: The dependent variable is a dummy for the first quarter with 0 earnings, interpreted as exit from employment. Data are means over cells defined by the month of birth, observed in a given quarter. Each regression includes a full set of monthly age dummies and birth cohort dummies. Controls: gender and state dummies (fraction, after aggregation). *, **, and *** indicate that the coefficient estimate is significantly different from zero at the 10%, 5%, and 1% level, respectively. Sample: LEHD from three states.

Table 5: Differential impact of FRA on SS Benefit Claiming Hazard, by type of worker

		Interaction term: sociodemographics					
		White	Hispan	Woman	Married	High-school dropout	Bad health
FRA		6.6 (4.4)	13.4*** (2.1)	13.3*** (2.5)	14.0*** (3.4)	13.9*** (2.2)	14.1*** (2.2)
FRA*interaction term		8.0 (4.9)	0.1 (6.0)	0.3 (4.0)	-0.7 (4.1)	-2.9 (5.0)	-3.9 (4.8)
N		88229	88228	88229	88011	88106	88074
R ²		0.14	0.14	0.14	0.14	0.14	0.14
		Interaction terms: pension and job characteristics					
		Covered by health insurance from current or private job	Covered by spouse's health insurance	Pension available from the current job	Pension type is defined benefits	Above median wealth	Above median wage
FRA		12.3*** (3.2)	14.2*** (2.1)	13.2*** (3.5)	22.1*** (3.7)	5.7* (3.1)	16.7*** (2.6)
FRA*interaction term		2.3 (4.0)	-4.7 (6.3)	3.6 (4.4)	-11.8** (5.3)	12.1*** (4.0)	-3.9 (5.8)
N		87899	87014	66676	38787	88115	59566
R ²		0.15	0.14	0.12	0.09	0.14	0.11
		Interaction term: cognition and behavior					
		High TICS score	High episodic memory	High numeracy	High self control	Long financial planning horizon	Risk averse
FRA		5.2 (5.0)	10.7*** (2.8)	8.0** (3.1)	6.8 (7.0)	11.5*** (3.6)	16.7*** (3.0)
FRA*interaction term		9.6* (5.5)	5.7 (3.9)	12.2*** (4.3)	-1.6 (9.1)	3.2 (4.3)	-4.8 (4.1)
N		86127	87105	69532	15394	86053	84070
R ²		0.14	0.14	0.14	0.14	0.14	0.14

Notes: see notes to Table 1. See equation (5.1) for the model specification. Each regression includes a full set of monthly age dummies and birth cohort dummies interacted with the interaction variable (one different interaction variable per regression).

Table 6: Differential impact of FRA on SS Benefit Claiming Hazard, by type of worker

	Claiming social security (OASI) benefits						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
FRA*high wealth	9.6 (6.4)						
FRA*defined benefits	-13.3** (5.8)	-13.1** (5.8)					
FRA*high numeracy	2.2 (6.1)						
FRA*high memory	16.6*** (6.1)						
FRA*high TICS	15.8 (9.7)						
FRA*cognition index		2.1** (1.0)	2.2*** (0.8)	2.4*** (0.9)	2.2*** (0.8)		
FRA*years of education	0.5 (1.0)	1.0 (1.0)	0.4 (0.8)	0.4 (1.0)	0.4 (0.8)	0.4 (0.8)	0.9 (1.0)
FRA*white	4.6 (7.7)	11.1 (7.5)	2.6 (6.3)	-2.0 (6.9)	2.8 (6.3)	2.4 (6.3)	-2.1 (8.2)
FRA*woman	7.6 (6.8)	10.6 (6.5)	0.5 (4.6)	3.2 (5.2)	0.5 (4.6)	-0.7 (4.6)	-0.7 (5.7)
FRA*married	-0.9 (6.9)	-0.3 (7.0)	-3.1 (4.9)	0.2 (5.4)	-3.1 (4.9)	-3.4 (4.9)	0.5 (6.2)
FRA*(2nd quartile cognition index)						16.9*** (6.5)	12.8 (8.3)
FRA*(3rd quartile cognition index)						15.1** (6.6)	7.8 (7.8)
FRA*(4th quartile cognition index)						24.9*** (6.9)	19.3** (8.2)
FRA*stressful job				-6.8 (5.2)			
FRA*physically demanding job				0.6 (4.7)			
SS earnings test removal*cognition index					-0.5 (0.3)		
N	32709	32709	69892	54856	69892	69892	57904
R ²	0.11	0.10	0.15	0.11	0.15	0.15	0.15

Notes: see notes to Table 5. Each regression includes a full set of monthly age dummies and birth cohort dummies interacted with the different interaction variables, as well as a direct FRA effect. Column (7) excludes workers born in 1938 and 1939.