Framing Social Security Reform:

Behavioral Responses to Changes in the Full Retirement Age

By Luc Behaghel and David M. Blau*

Abstract

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. 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 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. JEL: J26

*Behaghel: Paris School of Economics-INRA and CREST, 48 Boulevard Jourdan, 75014 Paris, France, Luc.behaghel@ens.fr. Blau: Ohio State University and IZA, Arps Hall, 1945 N. High St., Columbus, OH, 43210-1172, 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 authors 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 two referees, 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.

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.[1]

Several explanations for the age 65 spike have been proposed. Some explanations are consistent with the life cycle framework: for instance, a sharp nonlinearity at 65 in the relationship between claiming age and the Social Security benefit, interactions with Defined Benefit (DB) pension plans, which often have financial incentives to retire at 65, or availability of health insurance from Medicare, for which eligibility starts at 65. There are two leading "behavioral" explanations for the age-65 spike that relax the assumptions of farsighted rational behavior and standard preferences. First, the fact that 65 has been the Full Retirement Age (FRA) from the beginning of the OASI program until recently may have caused an endorsement effect. Workers might take this "official" designation as implicit advice from the government about when to retire and claim benefits; if collecting information or deciding by themselves is too difficult or too costly, they may decide to follow this advice. Second, the fact that the FRA is

[1] A 30% spike in the hazard of exit from employment means that 30% of individuals who were working prior to turning 65 left employment with a year after turning 65. Source: 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 often closely related in practice.

used as a reference point in Social Security Administration (SSA) explanations of the SS benefit schedule could influence behavior if workers' preferences are reference-dependent with loss aversion. The intuition is as follows: the salience of the FRA makes it a natural starting point for workers considering when to retire and claim benefits. With standard preferences, for most workers the FRA would not be the optimal choice, and they would move to a different point, specific to each of them, where the marginal rate of substitution between leisure and consumption equals the slope of the benefit schedule. If workers are loss averse, however, some will not deviate from the FRA, as the utility cost of losing benefits or leisure time would be amplified. There is little conclusive evidence to date on these different explanations for the age 65 spike. In particular, most of the evidence on behavioral explanations comes from testing and rejecting other explanations, leaving behavioral economic explanations as the default (Lumsdaine, Stock, and Wise, 1996).

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.[2] The cohorts affected by the reform reached their FRA in 2004-2009, so data on their retirement behavior are 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 by comparing cohorts 1937 and 1943, the increase in the FRA did not change the *slope* of the benefit-claiming-age schedule in the vicinity of the FRA.[3] So there is no economic incentive for someone who, for whatever reason, would have

[2]The 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.

[3]The 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

retired or claimed the benefit at 65 if the FRA had not changed, to instead do so at his 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 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 and the Longitudinal Employer-Household Dynamics data 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.

Our second contribution is to address the more novel question of which groups, defined by observable characteristics, are most responsive to the change in the FRA, i.e. “who is behavioral” (Benjamin, Brown, and Shapiro, 2006) with respect to age of retirement. 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 such as

---

after the FRA is equal to the DRC for his cohort. The loss in expected present discounted value of lifetime benefits from delaying claiming from 65 to 66 is about two percent at a real interest rate of 3%, using U.S. life table mortality rates. There is a similar loss for delaying claiming past age 66.

pension and health insurance coverage; (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 ability respond more strongly to the FRA change.

These results are supportive of a behavioral explanation of the age 65 spike. They do not constitute a full test of which of the leading behavioral explanations – the endorsement effect or reference dependence with loss aversion – is at play.[4] The fact that responsiveness to the FRA increases with cognitive skills indicates that the endorsement effect could not be the full story, as the endorsement effect implies unsophisticated decision making and should be less, not more, prevalent among people with higher cognitive skills. Reference dependence with loss aversion is a form of non-standard preferences and there is no particular reason to expect non-standard preferences to be associated with cognitive ability. Thus our findings do not provide direct support for reference dependence with loss aversion, but are more consistent with that story than with an endorsement effect. One interpretation consistent with the evidence and with both behavioral explanations is that people with lower cognitive skills may be less responsive to the change in the FRA because they do not learn, or learn more slowly about the new "advice" or the new reference point.

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 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 III and IV present evidence on the effect of the shift in the FRA. Section V analyzes heterogeneity in the response

[4] By contrast, the fact that the first cohorts exposed to the shift in the FRA responded to the shift does not seem consistent with a "social norms" story, insofar as social norms evolve slowly.

to the FRA. Section VI describes and analyzes the implications of a simple behavioral model of retirement. The final section concludes.

I. Previous Studies

Here, we briefly review existing evidence on explanations for the age 65 spike.

(i) Liquidity Constraint. 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 (Crawford and Lilien, 1981). 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 for the age 65 spike.

(ii) 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, illustrated in Figure 1 for the 1924 birth cohort. Delaying claiming from 62 to 65 resulted in a benefit increase of 6.67% per year, but delaying claiming past 65 resulted in a much smaller increase (1% for individuals who turned 62 in 1981; 3% for individuals who turned 62 from 1982 until 1989). The age-65 spike could be rationalized as a response to a kinked intertemporal budget constraint (Hurd, 1990). However, the 1983 Social Security reforms eliminated the kink, gradually increasing the reward to delaying claiming past the FRA. For individuals who turned 62 in 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.

(iii) 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).[5] DB pensions are much less common today, having been largely supplanted by Defined

[5] In 2005, 59% of DB plans had a normal retirement age of 65. The early retirement age was 55 in 76% of plans that had an early retirement age (U.S. Department of Labor, 2007, Tables 50 and 52).

Contribution (DC) plans.[6] 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 less than it has affected more recent cohorts.[7] The prevalence of DB 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 in the plan.

(iv) 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 (Madrian, 1994). 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.[8] We include controls for employer-provided health insurance and retiree health insurance coverage.

[6] Measured 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. Some DB plans are integrated with Social Security, meaning that the pension benefit depends on the amount of the Social Security benefit. This creates a link between the focal retirement ages in Social Security and DB pensions. We have not been able to find information on how the change in the FRA has affected the normal retirement age in DB plans. The results presented below suggest that the majority of DB plans did not change their normal retirement age.

[7] In 2000, 40% of private sector employees with a pension were covered by DB plans and 75% by DC plans. The total sums to more than 100% because 15% are covered by both types (U.S. Department of Labor, 2003). In the HRS, the corresponding figures for individuals covered by a pension in 2000 were 46% exclusively DB, 48% exclusively DC, and 5% both (authors' tabulations).

[8] Evidence on the role of Medicare in retirement decisions derived from simulations based on structural models is provided in Rust and Phelan (1997), Blau and Gilleskie (2006, 2008), and French and Jones (2011). 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. The results of French and Jones for the 1990s are in between. Some workers are covered by the spouse's health insurance plan, and others purchase a health insurance policy in the market. These possibilities and the option to extend coverage from the employer for up to three years after leaving the firm (so-called COBRA coverage) reduce the potential role of Medicare as an explanation for the age 65 spike.

The most plausible behavioral economic explanations for the age 65 spike are based on the fact that 65 was the Social Security FRA, until recently (Lumsdaine et al., 1996). 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” (Social Security Advisory Board, 2009). 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.[9] Moreover, the distinction between retiring “early” and retiring “late” is explicitly discussed, using the FRA as a reference, in the age-55+ insert sent with the Social Security statement. Clearly, the way the FRA is used in framing benefits invites people to use it as their point of reference.[10] This could result in development of a social norm, or could be taken as advice by agents with limited ability to evaluate the financial implications of alternative claiming and labor force behavior, or could serve as a reference point for individuals whose preferences exhibit loss aversion with reference dependence. Furthermore, advice from non-government sources about when to claim Social Security often uses the FRA as a reference point. For example, the AARP web site begins its advice about when to claim with the statement “If you're healthy and can afford it, you should consider waiting until you reach your full retirement age of 66, or even 70.”[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. Brown et al. (2011) report results from an experiment in which alternative ways of framing the tradeoff between early and

---

[9] See Mastrobuoni (2011) for discussion of the history of the Social Security statement and the impact of its introduction on retirement behavior.

[10] This is true both for cohorts with an integer FRA and for those 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.”

[11] http://www.aarp.org/work/social-security/info-12-2010/top-25-social-security-questions.5.html (question 25). The AARP Social Security calculator contains a bar chart very similar to the one in the Social Security statement, with the FRA as a benchmark.

later claiming were tested. They find some strong effects of framing on the expected age of claiming benefits. Liebman and Luttmer (2011) find that providing information about Social Security in a field experiment caused an increase in labor force participation one year later.

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 plausibly 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, so we cannot rule out information and cognitive ability as explanations for the impact of the FRA. We discuss indirect approaches to assessing the importance of these issues in sections V and VI.

Recent evidence indicates that the increase in the FRA has affected retirement behavior (Mastrobuoni (2009), Blau and Goodstein (2010), Pingle (2006)). These studies do not address the mechanism through which the FRA effect operates, and therefore do 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 which types of workers are more responsive to the new FRA.

II. Data

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 distinguish among alternative behavioral explanations for the age 65 spike.

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 Dynamics (LEHD) files. These files are derived from administrative state Unemployment Insurance (UI) records, with very large sample sizes, and they include information on birth date. These administrative data are used to verify that the trends identified in the HRS are robust. The LEHD is described in the web appendix.

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 a job history at the first interview, employment status at each interview, and the start and end dates of all jobs between interviews, to the nearest month. We construct a monthly employment history. Together with month and year of birth, the employment data are used to compute the age of labor force exit, defined as the age (in months) at which the individual is first out of the labor force for an entire month.[12]

[12]The sample for the labor force exit analysis includes all months during which the individual was

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, so we do not use such cases in the analysis of claiming behavior.[13] The age in months 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. We use only the latter cases in analysis of retirement.

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, household wealth, health insurance and pension coverage, other job characteristics, 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 distinguishing among alternative behavioral explanations for the age 65 spike. 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 (2011), and we follow their approach to construct indicators in three dimensions: Telephone Interview of Cognitive Status (TICS), short term memory, and numeracy (see the web appendix for details).

---

employed for any part of the month. Some people who leave the labor force later return to employment. In such cases there are multiple labor force exits. As a robustness check, we consider that an individual has exited the labor force only if the exit is preceded by an employment spell of at least 3 months, and followed by at least 3 months out of employment.

[13] If 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 provide no information about OASI claiming behavior.

III. Impact of the FRA increase

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 null hypothesis that the increase in the FRA had only a wealth effect against the general "behavioral" prediction that the spike at age 65 should shift along with the FRA for cohorts born after 1937. In the former case we expect an increase in retirement age, but no substantial shift in the spike. As discussed in section III, there are no economic incentives to retire or claim benefits at the new FRA, so in effect we are testing the null hypothesis of rational life cycle behavior against a general unspecified alternative. In Sections V and VI we bring evidence to bear on specific behavioral explanations.

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 post-reform cohorts. The claiming hazard rate is defined as the probability of claiming at a given age conditional on not having claimed previously (claiming is an absorbing state). Age is measured at a bimonthly frequency; e.g. age 65 denotes age 65 0/12 to 65 1/12[14] . In each graph of figure 2, the dotted line depicts the average 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 claiming age (62) and a second larger spike at the FRA (65).[15] About 20% of workers in these cohorts claim at 62, 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). The 1937 cohort is displayed separately, to demonstrate that there were no major shifts in behavior for the last "control" cohort. 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. Very similar patterns appear when

[14] 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 claiming his benefit in April, it is not clear how to classify his claiming status in April without knowing the exact date.

[15] 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 day of the month (see Kopczuk and Song, 2008, for discussion).

men and women are disaggregated (not shown). These results confirm the findings of Song and Manchester (2008), based on 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:

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

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 included (β_a, δ_c).The age 65 coefficient (one of the $\beta's$) 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 and to changes in the estimation sample (columns 2 to 3). Column 4 includes an indicator for whether the individual is subject to the Social Security Earnings Test (SSET) given his age and birth year. This is intended to control for elimination of the SSET in 2000 for individuals at or above the FRA. This reform had a noticeable effect on claiming, but accounting for this policy change does not affect the estimated impact of the FRA. Quantitatively, the estimated impact of the FRA is sizeable: the average 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.

Turning to labor force exit behavior, 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 4 in table 2 use the same specifications as in 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 mean monthly hazard of labor force exit at age 65 is 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. As previously shown by Baker and Benjamin (1999) for the case of Canada, the effect of Social Security reform on benefit claiming behavior can significantly differ from the effect on labor force participation and retirement behavior. It is common to work after claiming and to move in and out of the labor force, but claiming is a one-time event. There may also be more measurement error in the reports of the timing of labor force exit, as it may not be as salient as claiming. [16]

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.

As noted above, we also used administrative earnings data from the LEHD in order to verify that the labor force participation results from the HRS are robust. The results from analysis of the LEHD are generally quite consistent with findings from the HRS. These results are discussed in the web appendix. 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 the HRS, but the LEHD has very large samples and

[16] Results based on a smaller sample that eliminates temporary withdrawals from the labor force (three months or less) gave very similar results. Note that the monthly hazard rate of labor force exit is much smaller than the monthly hazard of Social Security claiming. For example, at age 65 the latter is equal to about 0.4 for claiming but only 0.02 for employment for the 1931-36 birth cohorts (compare Figures 2 and 3). Some of this difference could be due to measurement error in constructing a monthly labor force history from retrospective between-wave questions. However, claiming can occur before, after, or at the same time as labor force exit, so there is no reason why the levels of the two hazards should be of the same order of magnitude.

administrative data. The combination of findings from the two sources gives us some confidence that labor supply decisions are affected, though clearly not by as much as claiming decisions.

An important question is whether the changes in labor force and retirement behavior that we find 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 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). A direct comparison of this elasticity to our results is not possible, but some calculations suggest that the wealth effect is too small to account for more than a minor part of the estimated effect of the FRA on the age 65 spikes[17].

Finally, as noted in section II, it has been difficult to rule out Medicare as a cause of the age 65 spike in labor force exit. If Medicare is an important cause of the age 65 spike, then the spike should not disappear entirely when the FRA moves away from 65. The evidence presented in this section shows that the age-65 spikes in labor force exit and self-reported retirement gradually, if irregularly, disappeared as successive post-1937 birth cohorts reached age 65. However, these results pertain to the entire population, while the availability of Medicare at 65 may be important only for the subpopulation without retiree benefits from their employer-provided health insurance plan. When we limit the analysis to this subpopulation, the results are very close to those of table 1, although they are much less precisely estimated because of the smaller sample. Overall, the results suggest at most a minor role for Medicare in explaining the age 65 spike.

[17] The largest estimate of the elasticity of labor force exit with respect to SSW presented by Coile and Gruber (2007) is 0.16. Samwick (1998) estimates the elasticity to be approximately zero. Manoli, Mullen, and Wagner (2009) using Austrian administrative data estimate an elasticity of 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%.

IV. Distinguishing among behavioral economic explanations

Results from section IV provide strong evidence that OASI benefit claiming behavior has been influenced by the increase in the FRA, and weaker evidence that the same is true of labor force participation. As argued in section II, this finding leaves behavioral factors as likely explanations. However, as stressed by Kahneman (1999) and Featherstonehaugh and Ross (1999), loss aversion is only one of the potential behavioral explanations: the implicit advice or endorsement of the FRA by the SSA as a "normal" retirement age is another one. A third possible explanation is that the FRA has become a "social norm", because it is followed by the majority. Conceptually, these three leading explanations are different. Loss aversion is a form of non-standard preferences, whereas allowing one's behavior to be governed by advice or by social norms are forms of non-standard decision making.[18] Discriminating empirically between these three explanations is hard, even in an experimental context where one can, as in Brown et al. (2011), manipulate the framing. The conceptual difference between non-standard decision making and non-standard preferences however suggests an indirect test, by asking 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 (see the discussions by Burtless, 2006; Liebman et al., 2009; 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, 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:

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

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

[18] Advice and social norms are close, but they differ by who sets the norm and by their degree of inertia: a given authority may change its advice with immediate effects whereas social norms are likely to change slowly.

different response to the FRA in the population characterized by the *Type* variable. The type variable is also interacted with cohort and age dummies, and the full sets of main effects from Tables 1-3 are included as well. Equation (2) basically tests whether the FRA has heterogeneous effects by a set of observable characteristics. We group the type variables into 3 broad categories: socioeconomic; pension and job characteristics; and cognition and behavior.

Table 4 shows selected estimates from several specifications of equation (2) that include alternative combinations of the type variables. The parameters of interest are the coefficient on the "FRA*interaction term," (θ_2). For instance, -12.5 in column (1) on the "FRA*defined benefits" line means that workers with a defined benefit pension are 12.5 percentage points less responsive to the FRA than are other individuals, other things equal. Wealth is associated with higher responsiveness to the FRA (but the coefficient estimate is not significantly different from zero). Higher cognitive ability is associated with a larger response to the FRA, with a statistically significant effect of the memory index. There are no significant differential effects by socio-demographic dimensions such as race, gender, marital status, or education. Having used the online SSA calculator to compute future benefits or having the benefit calculated by the SSA has no statistically significant impact on the responsiveness to the FRA.

Columns 2 to 6 of Table 4 explore the robustness of the positive interaction effect on cognitive skills. Combining the cognitive measures into a single index yields a highly significant effect (column 2).[19] Dropping the DB variable (only available for a subsample of workers) does not significantly alter the results (column 3). Adding FRA interactions with stressful and physically demanding job indicators, health, subjective life expectancy, health insurance coverage, financial planning horizon, and risk aversion has very little impact on the FRA-cognition coefficient estimate (column 4). None of these additional FRA interaction coefficients are significantly different from zero. The FRA-cognition coefficient estimate is also robust to inclusion of an FRA-earnings test interaction (column 5). 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). By contrast, there is no evidence of nonlinear effects of wealth.

[19] See the web appendix for details on the cognition index.

A potential problem with the results in Table 4 is that the response to the FRA may be lower for workers with low cognitive skills in absolute (percentage points) term because the age 65 spike is initially lower for them, although the responses may be more similar in relative terms. Figure 5 (replicating figure 2 by cognitive skill group) shows that this is not the case: 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 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 last fact 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.[20] [21] The alternative behavioral explanations – advice and social norms – imply unsophisticated decision making, which would predict a *negative* FRA*cognition effect.

Of course, it may still be the case that these interaction effects are driven by unobserved sources of heterogeneity. While we are confident that the 1983 reform allows us to identify the causal effect of the FRA on different groups, the sources of variation that we use do not allow us to say what causes the differences in response. 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

[20] 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).

[21] 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. Another possible explanation is that the loss in expected lifetime utility from claiming at the FRA is small, because the present discounted value of lifetime benefits is intended to be approximately invariant to the age of claiming. If gathering the information needed to make a well-informed decision is costly (Mastrobuoni, 2011) then it could be optimal to simply follow the path of least resistance by claiming at a salient age such as the FRA. And if Social Security wealth provides a smaller share of retirement income for smart and/or wealthy individuals, this could make the utility cost of non-optimal claiming especially small. We cannot evaluate this explanation directly. Coile et al. (2002) find that in many circumstances it is optimal to delay claiming until the FRA but not beyond the FRA. However, their analysis was based on the rules for the 1930 birth cohort, for whom the Delayed Retirement Credit was only 3%.

responsiveness to the SS FRA is lower for those with a DB plan that has a normal retirement age at 65.[22]

V. Implications of the results for framing Social Security reform

We interpret the empirical results presented above as suggesting that reference dependence with loss aversion is the behavioral explanation for the age 65 spike in claiming and retirement decisions that is most consistent with the evidence. Admittedly, we have no direct evidence to support this explanation; rather we have at least some evidence against the alternatives. 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 SSA 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,[23] 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 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) to 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:

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

[23] We assume that retirement and claiming occur at the same for age, for simplicity. This assumption allows us to avoid introducing dynamics, which would be required to deal with minimum and maximum claiming ages and the possibility of a liquidity constraint. This assumption implies that we ignore the lower bound on claiming age. In the simulations described below, we do account for the lower bound on the claiming age. In practice, the majority of individuals claim at the same age at which they retire (Coile et al., 2002).

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

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:

(4) $$U_{FRA}(l,c) = aR_1(l) + R_2(c),$$

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

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

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(.)* and *v(.)* are increasing and concave utility subfunctions. Lastly, $a > 0$ is an individual-specific parameter that allows for heterogeneity in the preference for leisure.[24]

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

[24] 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(.)* and *v(.)* 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 preferences or the budget constraints), they cannot by themselves account for a spike in the retirement hazard.

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 3 given preferences characterized by equation 4 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 the FRA. Then:

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

The fact that P_{FRA} is strictly positive if either $\lambda_1 > 1$ or $\lambda_2 > 1$ shows that loss aversion in either of the two dimensions can 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 web 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*.

A. 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). 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 before the FRA goes up, while 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*, holding the slope (*w*) constant, vary with the way the reform is framed? In all cases, we have

$$E(R) = T - E(l) = T - \int l^*(a) f(a) da \, , \tag{6}$$

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

$$\frac{dE(R)}{dk} = -\int \frac{dl^*(a)}{dk} f(a) da \, ; \tag{7}$$

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. 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, which yields the standard wealth effect (see the web appendix).[25]

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 consumption target if they retire at 65 after the reform,

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

but the target is now lower. This arises if the reform is framed as a change in the Primary Insurance Amount (the benefit amount available if claimed at the FRA) with no change in the FRA itself.

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 web 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 on the average age of retirement and claiming if it is framed as a change in the reference point than if it is framed as an equivalent change in the benefit at an unchanged 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 $\left[\frac{dE(R)}{dk}\right]_3 > \left[\frac{dE(R)}{dk}\right]_1$. 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 0.5 year increase in the average retirement age. In the framework of our model, the .5 response must be a weighted average of 1 for people at the old FRA (who will move to the new FRA, a one-for-one effect) and δ (the average response for people above and below the FRA). This implies that $\delta < 1$. More precisely, the magnifying effect caused by framing 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 (δ is a constant), we have

$$\text{(8)} \qquad \left[\frac{dE(R)}{dk}\right]_3 \approx -\frac{1}{w}\left[(1-P_{FRA})\delta + P_{FRA}\right] = -\frac{1}{w}\left[(\delta + P_{FRA}(1-\delta)\right]$$

which is increasing in P_{FRA}.

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.[26]

Is loss aversion enough to explain the unexpectedly strong impact of the 1983 reform found by Mastrobuoni (2009) and others? 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 relatively small, so the impact on the average retirement age should be relatively modest.[27] 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.

VI. 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: in 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

[26] We parameterized the model and simulated the impact of a benefit cut under various assumptions about loss aversion, framing, and parameter values. The results indicate that framing the cut as an increase in the FRA amplifies its effect on the mean retirement age substantially, while neutral framing and framing as a cut in the PIA lead to much smaller effects. The simulation results are described in the web appendix.

[27] 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).

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 is probably 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, we showed that in 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

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.

Brown, Jeffrey, Arie Kapteyn, Olivia S. Mitchell. 2011. "Framing Effects and Expected Social Security Claiming Behavior," NBER Working Paper 17018, May.

Burtless, Gary. 2006 "Social Norms, Rules of Thumbs, and Retirement: Evidence for Rationality in Retirement Planning", in K.W. Schaie and L.L. Carstensen eds. Social Structures, Self-Regulation, and Aging, New York: Springer.

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 Jousten. 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.

Crawford, Vincent P., and David M. Lilien. 1981. "Social Security and the Retirement Decision," *Quarterly Journal of Economics* 96 (3): 505-29.

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.

French, Eric, and John Bailey Jones. 2011. "The Effects of Health Insurance and Self-Insurance on Retirement Behavior, *Econometrica* 79 (3): 693-732.

Hurd, Michael. 1990. "Research on the Elderly: Economic Status, Retirement, and Consumption and Saving," *Journal of Economic Literature* 28 (2): 565-637.

Kahneman, Daniel. 1999. Comment on "Framing Effects and Income Flow Preferences in Decisions About Social Security," by Featherstonehaugh and Ross, in Henry Aaron ed. Behavioral Dimensions of Retirement Economics Washington DC, Brookings Institution Press and Russell Sage Foundation: 210-213.

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.

Liebman J., E. Luttmer and D. Seif. 2009. "Labor Supply Responses Marginal Social Security Benefits: Evidence from Discontinuities", *Journal of Public Economics* 93 (11-12): 1208-23.

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.

Liebman, Jeffrey B., and Erzo F.P Luttmer. 2011. "Would People Behave Differently if They Better Understood Social Security? Evidence from a Field Experiment" NBER Working Paper 17287, August.

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. 1994. “The Effect of Health Insurance on Retirement,” *Brookings Papers on Economic Activity* (1): 181-252.

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. 2011. “The Role of Information in Retirement Behavior: Evidence Based on the Social Security Statement,” *Journal of Public Economics* 95 (7-8): 913-25.

McArdle, John, James Smith and Robert Willis. 2011. “Cognition and Economic Outcomes in the Health and Retirement Survey,” in David Wise (ed.) Explorations in the Economics of Aging, Chicago: University of Chicago Press for the NBER: 209-33.

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 Advisory Board. 2009. “The Social Security Statement: How It Can Be Improved,” http://www.ssab.gov/documents/Pre-Publication_Version_The_Social_Security_Statement_web.pdf, May, accessed 11/02/2011.

Song, Jae, and Joyce Manchester. 2008. “Have People Delayed Claiming Social Security Benefits? Responses to Changes in Social Security Rules, *Social Security Bulletin* 67 (2): 1-23.

Tversky, Amos, and Daniel Kahneman. 1991. “Loss Aversion in Riskless Choice: A Reference Dependent Model”, *Quarterly Journal of Economics* 106(4): 1039-1061.

U.S. Department of Labor. 2003. “National Compensation Survey: Employee Benefits in Private Industry in the United States, 2000,” U.S. Bureau of Labor Statistics, January, Bulletin 2555.

U.S. Department of Labor. 2007. “National Compensation Survey: Employee Benefits in Private Industry in the United States, 2005,” U.S. Bureau of Labor Statistics, May, Bulletin 2589.

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

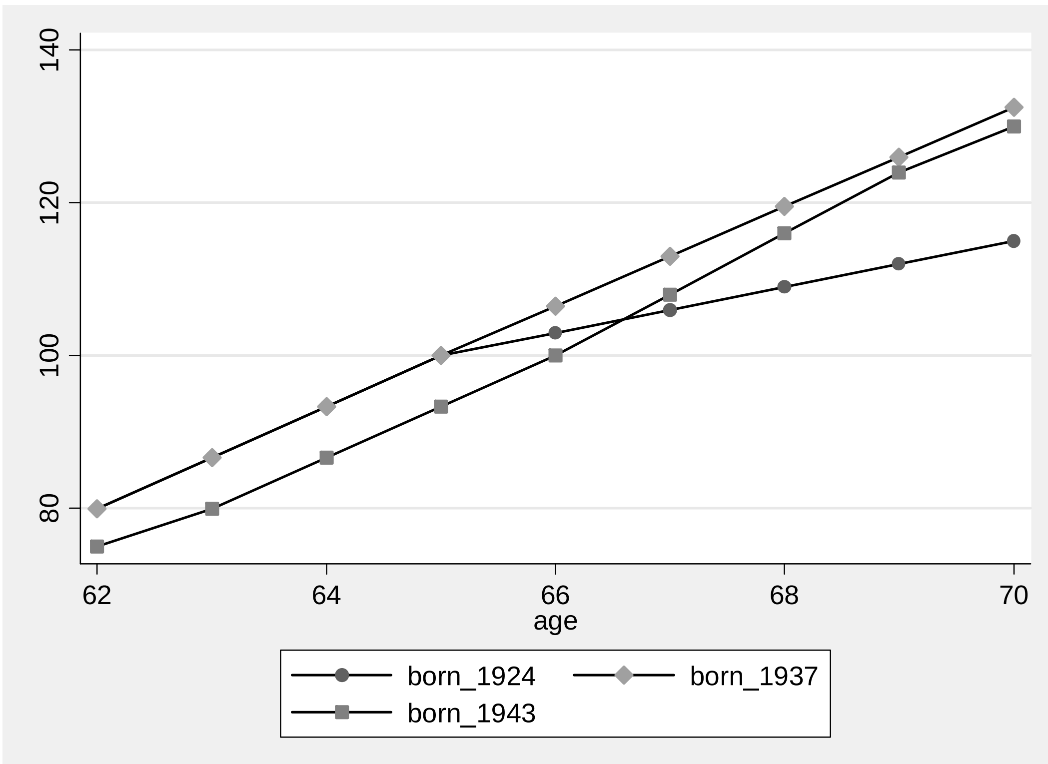

Source: Authors' calculation from Social Security rules. Primary Insurance Amount is the benefit amount when claimed at the Full Retirement Age.

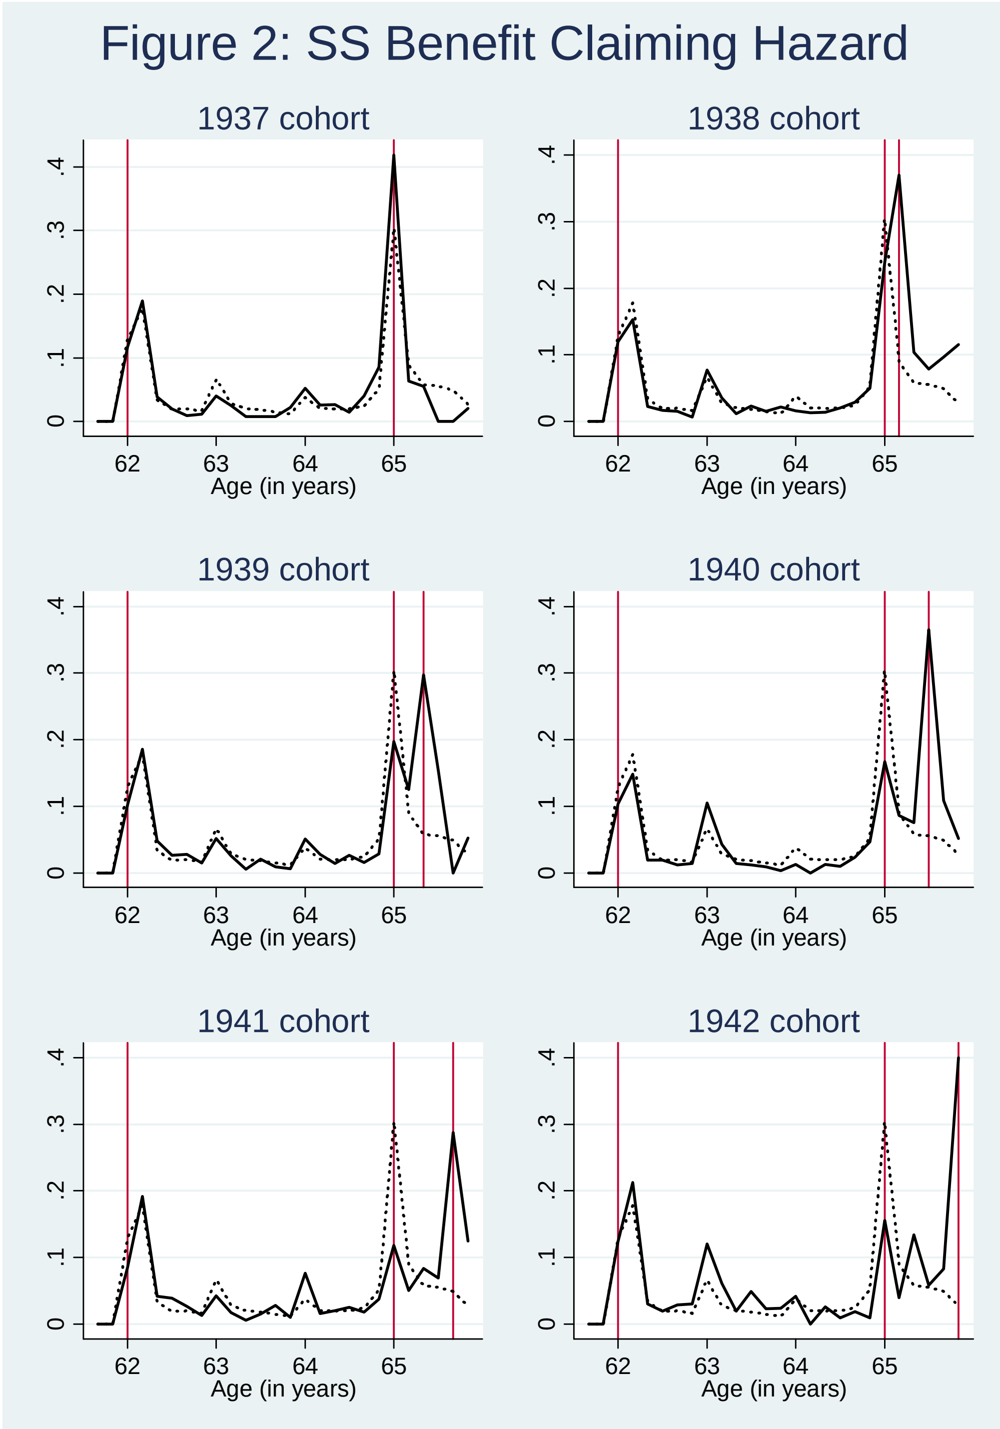

Notes: The graphs show the average monthly claiming hazard rates for pre and post-reform cohorts. The claiming hazard rate is defined as the probability of claiming at a given age, conditional on not having claimed previously. Age is measured at a bimonthly frequency; e.g. age 65 denotes age 65 0/12 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

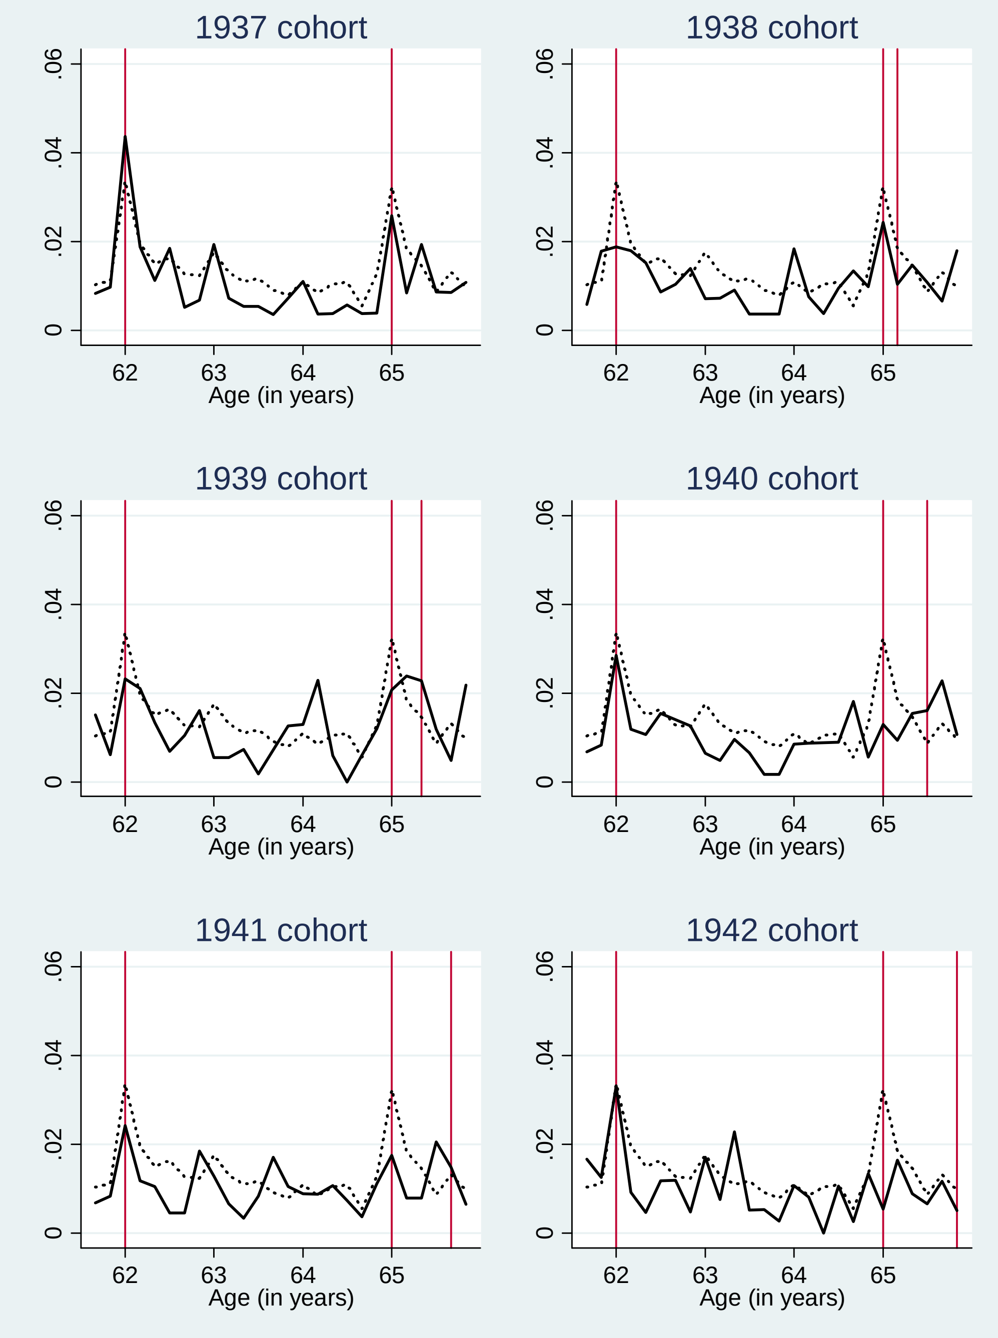

Notes: see notes to figure 2. The hazard of exit from employment is the probability of not being employed for any part of a given month, conditional on having been employed in at least part of the previous month.

Figure 4: Hazard of Retirement

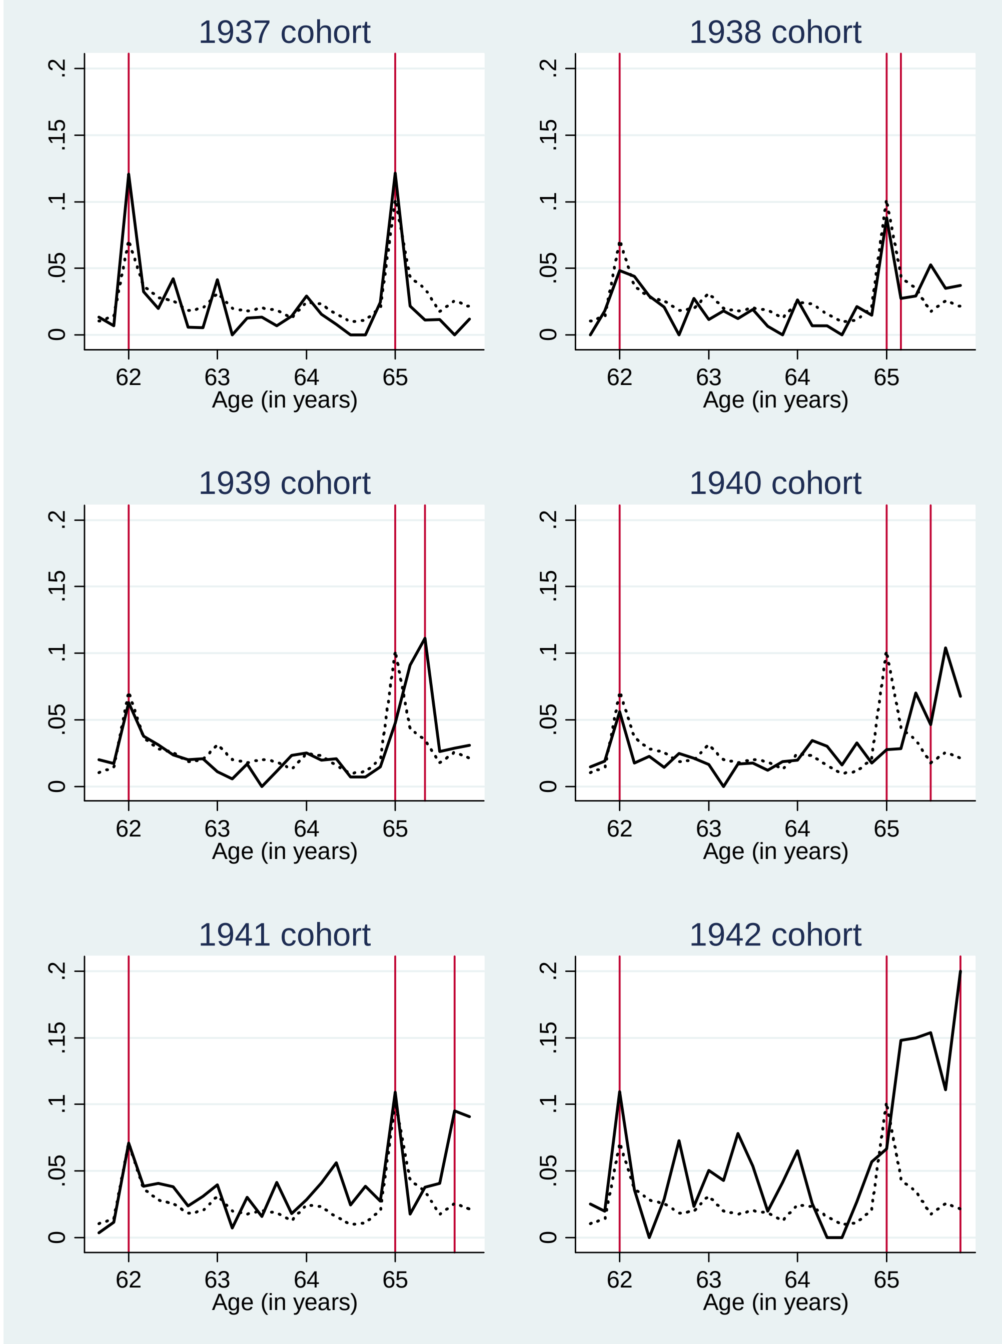

Notes: see notes to figure 2. To keep the same scale on the vertical axis for 1942 as for other birth cohorts, the hazard rate at age 65 10/12 has been arbitrarily set at .2. The observed value is .6 (over only 5 individuals). The hazard of retirement is defined as the probability of reporting being retired in a given month, conditional on not having reported being retired previously.

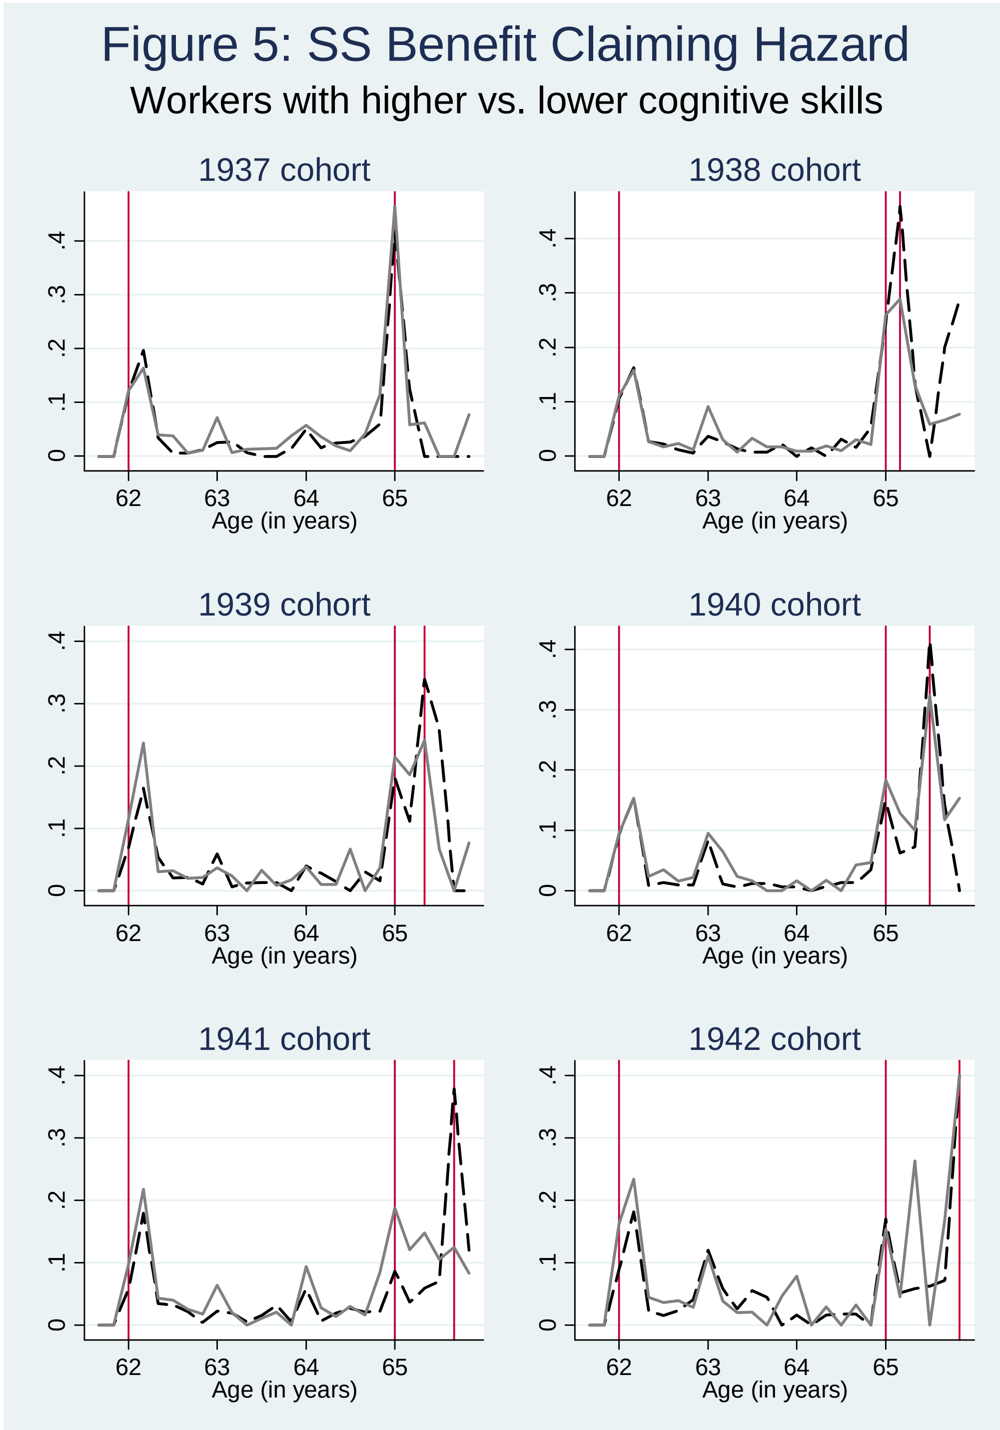

Notes: see notes to figure 2. People with a cognitive index above the median are in black dashed lines, people with a cognitive index below the median are in gray full lines.

Figure 6: Heterogeneous preference for leisure and optimal retirement age

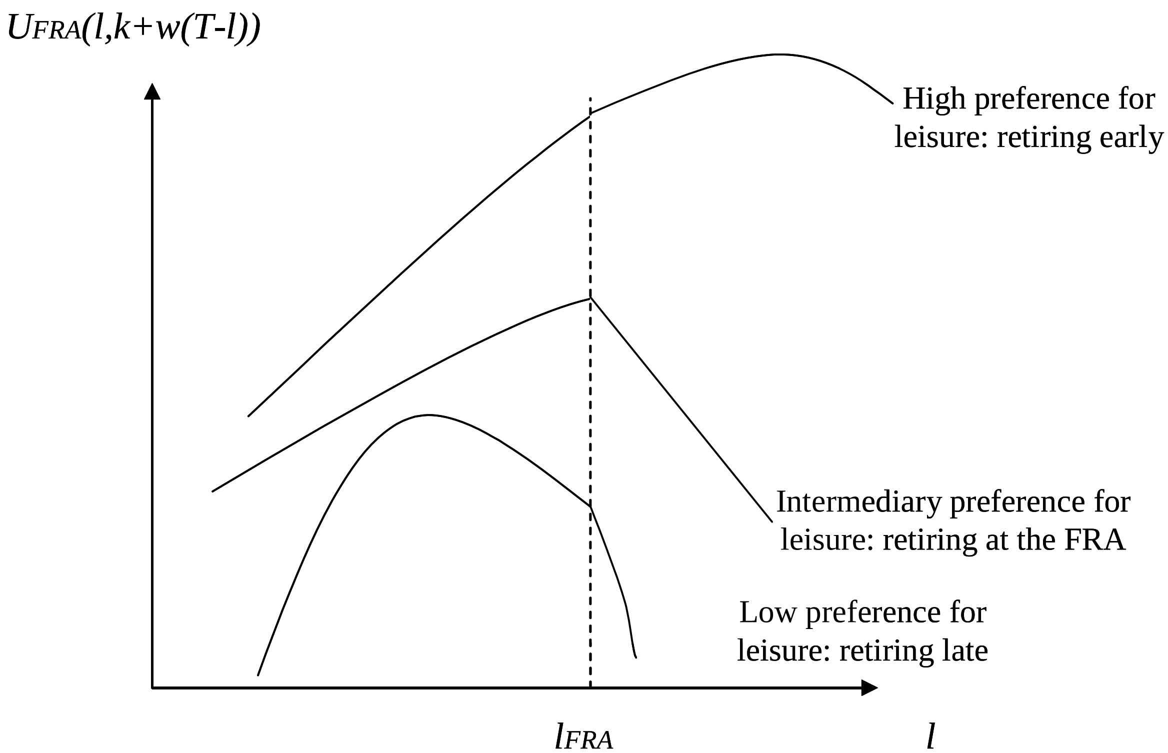

Notes: see discussion in the text.

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

	Claiming social security (OASI) benefits			
	(1)	(2)	(3)	(4)
FRA	13.8***	13.6***	13.3***	13.6***
	(1.9)	(1.9)	(1.9)	(1.9)
SS earnings test removal				3.2***
				(1.0)
Controls	No	Yes	Yes	Yes
Age range	64-66	64-66	62-66	64-66
N	25801	25801	89348	25801
R^2	0.146	0.154	0.162	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 (1) for the specification. The dependent variable is a dummy for claiming social benefits in the current month. The sample includes all person-months from age 62 to 70 in which an individual had not yet claimed. Cases that claimed before age 62 are excluded. 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): race, sex, marital status, education, health, health insurance coverage, retiree health insurance coverage, pension coverage, pension type, household wealth, average hourly earnings, and measures of cognitive capability, planning horizon, and risk aversion. Sample: HRS waves 1992-2008, cohorts born in 1932-41. Age range included in the regression differs 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)
FRA	1.0***	0.9**	0.9**	1.0***
	(0.4)	(0.4)	(0.4)	(0.4)
SS earnings test removal				0.5*
				(0.3)
Controls	No	Yes	Yes	Yes
Age range	64-66	64-66	62-66	64-66
N	68952	68952	152753	68952
R^2	0.016	0.047	0.054	0.047

Notes: see notes to Table 1. The dependent variable is a dummy for non-employment in the current month, conditional on employment in the previous month. The sample includes all person-months in which an individual was employed for any part of the month.

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

	Retiring			
	(1)	(2)	(3)	(4)
FRA	1.5	1.1	1.1	1.1
	(1.6)	(1.6)	(1.6)	(1.6)
SS earnings test removal				2.2**
				(1.1)
Controls	No	Yes	Yes	Yes
Age range	64-66	64-66	62-66	64-66
N	16387	16387	46737	16387
R^2	0.056	0.082	0.085	0.083

Notes: see notes to Table 1. The dependent variable is a dummy equal to one if the current month corresponds to the respondent's self-reported retirement age.The sample includes all person-months prior to the date at which the individual reported retiring. Cases that gave longitudinally inconsistent reports are excluded.

Table 4: Differential impact of the FRA on SS Benefit Claiming Hazard, by type of worker

	Claiming social security (OASI) benefits					
	(1)	(2)	(3)	(4)	(5)	(6)
FRA*high wealth	3.8	4.4	6.7	6.2	6.0	
	(7.0)	(7.1)	(5.6)	(6.5)	(6.5)	
FRA*(2nd quartile wealth)						3.9
						(7.9)
FRA*(3rd quartile wealth)						0.9
						(8.3)
FRA*(4th quartile wealth)						6.8
						(8.4)
FRA*defined benefits	-12.5**	-12.5**				
	(6.1)	(6.2)				
FRA*high numeracy	3.7					
	(6.7)					
FRA*high memory	14.3**					
	(6.6)					
FRA*high TICS	23.0**					
	(9.7)					
FRA*cognition index		1.7	2.2**	2.3**	2.3**	
		(1.1)	(0.9)	(0.9)	(0.9)	
FRA*(2nd quartile cognition index)						16.0**
						(8.1)
FRA*(3rd quartile cognition index)						24.2***
						(8.0)
FRA*(4th quartile cognition index)						26.4***
						(8.5)
FRA*calculated SS benefits	4.4	5.2	3.7	7.1	7.0	5.6
	(6.2)	(6.2)	(4.8)	(5.2)	(5.2)	(5.1)
FRA*years of education	1.2	1.3	0.0	0.3	0.2	0.1
	(1.1)	(1.1)	(0.9)	(1.1)	(1.1)	(1.1)
FRA*white	-1.1	3.6	-0.6	-8.0	-7.7	-8.3
	(9.1)	(9.4)	(7.3)	(8.2)	(8.1)	(7.7)
FRA*woman	7.5	10.4	-1.5	2.4	2.3	2.2
	(7.6)	(7.2)	(5.0)	(5.6)	(5.6)	(5.7)
FRA*married	6.5	5.5	-4.5	0.6	0.8	1.4
	(7.8)	(8.0)	(5.6)	(6.3)	(6.3)	(6.2)
FRA*stressful job				1.4	1.7	2.8
				(5.2)	(5.2)	(5.2)
FRA*physically demanding job				-1.0	-0.8	-1.8
				(5.9)	(5.9)	(5.9)
FRA*bad health				4.4	3.9	4.5
				(8.9)	(8.9)	(8.7)
FRA*long subjective life expectancy				-1.5	-1.6	-1.3
				(5.3)	(5.3)	(5.3)
FRA*covered by health insurance from job				-8.1	-8.0	-7.4
				(5.5)	(5.5)	(5.5)
FRA*long financial planning horizon				4.7	4.7	5.9
				(5.6)	(5.5)	(5.6)
FRA*risk averse				-1.6	-1.4	-1.8
				(5.1)	(5.1)	(5.2)
SS earnings test removal*cognition index					-0.5	-0.4
					(0.4)	(0.4)
N	28022	28022	57444	45509	45509	45509
R^2	0.12	0.11	0.15	0.14	0.14	0.14

Notes: see notes to Table 1. See equation (2) for the model specification. 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.