Journal of Pension Economics and Finance (2020), **19**, 442–457
doi:10.1017/S1474747218000331

CAMBRIDGE UNIVERSITY PRESS

POLICY PAPER/BRIEF

Why are US men retiring later?

Wenliang Hou*, Alicia Munnell, Geoffrey Todd Sanzenbacher and Yinji Li

Boston College, Center for Retirement Research, 140 Commonwealth Avenue, Hovey House, 256 Hammond Street, Chestnut Hill, Massachusetts, USA
*Corresponding author. E-mail: Houw@bc.edu

(Received 9 April 2018; revised 5 November 2018; accepted 5 November 2018; first published online 26 December 2018)

Abstract

Over the past two decades, the share of individuals claiming Social Security at the Early Eligibility Age has dropped and the average retirement age has increased. At the same time, Social Security rules have changed substantially, employer-sponsored retirement plans have shifted from defined benefit (DB) to defined contribution (DC), health has improved, and mortality has decreased. In theory, all of these changes could lead to a trend toward later claiming. Disentangling the effect of any one change is difficult because they have been occurring simultaneously. This paper uses the Gustman and Steinmeier structural model of retirement timing to investigate which of these changes matter most by simulating their effects on the original cohort (1931–1941) of the Health and Retirement Study (HRS). The predicted behavior is then compared with the actual retirements of the Early Boomer cohort (1948–1953) to see how much of the later cohort's delayed claiming and retirement can be explained by these changes. The Early Boomer cohort was less likely to be fully retired than the HRS cohort at both age 62 (36.7% vs. 44.0%) and age 64 (49.5% vs. 53.9%). The model suggests that the shift from DB toward DC plans was the biggest contributor to these declines, followed by better health. Social Security rules and improvements in mortality played smaller roles.

Key words: Pension plans; retirement decision; retirement timing; Social Security; structural model

JEL Codes: J22; J26

1. Introduction

Between 1996 and 2013, the average retirement age increased by 18 months for men and, relatedly, the share claiming Social Security at age 62 declined from 56.0% to 35.6% (Munnell, 2015; Munnell and Chen, 2015). A variety of explanations have been put forward to explain these changes. Several Social Security parameters have changed in recent decades: the Full Retirement Age (FRA) increased from age 65 to 66, the earnings test was eliminated for those over the FRA, and the delayed retirement credit (DRC) was increased. The pension landscape also changed from one comprised primarily of defined benefit (DB) plans to one comprised of defined contribution (DC) plans. In the USA, most DB plans are structured so that once a certain age is reached (often called the normal retirement age), benefits can be received with little increase for additional work, thus encouraging workers to reach that age and then offering little incentive to work longer. On the other hand, most DC plans vest (i.e., they can keep employer contributions) early in a workers' career, so that each additional year of work near retirement is of roughly equal value. Therefore, the shift from DB to DC plans likely created incentives to work longer. At the same time, health improved, meaning people can work longer, and longevity increased, meaning people need to work longer. How much of the delay in retirement can be explained by these changes, and what portion is left over, unexplained?

The problem in answering this question is that all of these factors have been changing simultaneously, making isolating the impact of any one on the retirement age difficult. Yet, isolating these effects is important, because the rate at which they will change in the future differs. For example, the FRA held at 66 for retirees reaching 62 between 2005 and 2016, but has begun its ascent toward 67 where it will remain for the foreseeable future. On the other hand, health and mortality seem likely to continue their improvement over the next several decades (despite some recent flattening). Knowing, for example, that health has had a large impact in delayed retirement would suggest continued delay should be expected in the future, whereas a large effect of changing Social Security rules would only have a short-term impact. This, of course, would have implications for government programs like Social Security as well as for private sector employers seeking to manage their workforces.

To isolate the effects of these changes, this paper will use the structural model of retirement timing developed by Gustman and Steinmeier (2006). Structural models are useful for this purpose because they can isolate the separate impact of each factor discussed above on retirement ages, holding preferences constant. To do so, this paper will conduct four simulations that explore how the retirement age of the Health and Retirement Study (HRS) cohort would react if some aspects of the Early Boomer cohort's experience were applied to it. Specifically, this paper conducts four experiments that replace the HRS cohort's actual data with data mirroring that of the Early Boomer cohort: (1) their Social Security rules; (2) their pensions; (3) their health; and (4) their mortality rates. This paper then combines all four changes to see how similar the HRS cohort's behavior up to age 64 (the last age the Early Boomer cohort is observed) would be to that of the Early Boomer cohort. Essentially, combining the four experiments explores how much of the difference between the retirement behavior of the HRS cohort and the Early Boomer cohort can be explained by these four factors.

The results suggest that of the four factors examined the shift to DC plans has had the greatest impact on the retirement timing of the Early Boomer cohort, explaining over one third of the drop in retirement that occurred relative to the HRS cohort at age 62 and over three quarters of the drop at age 64. Health also played a significant role, especially at later ages – improving health explained a quarter of the drop between the two cohorts at age 64. Changing Social Security rules had a smaller effect and may have actually served to increase claiming at 62, while decreasing it slightly at 64.

The next section describes the literature on retirement timing and on structural models. The following section reviews the model of Gustman and Steinmeier used here and the fourth describes the data and estimation of that model. The fifth section compares the behavior of the Early Boomer cohort to the HRS cohort and performs simulations to see how much of any change in retirement timing can be explained by the four factors considered. The final section concludes that because much of the change in retirement age from the HRS cohort to the Early Boomer cohort was due to changes in pension structure, the upward trend in retirement ages is likely to continue as DC plans continue to replace DB plans among new retirees.

2. Literature review

The literature on the four factors' impact on retirement timing spans both reduced form and structural techniques and covers each of the four issues being tested here.

2.1 Social Security benefits

Starting with the effect of Social Security rules on retirement timing, Blau and Goodstein (2010) estimate a reduced form approximation of a life-cycle model and find that changes in the FRA and in the DRC could explain about a quarter to a half of the increase in labor force participation seen since the 1980s. On the structural side, Gustman and Steinmeier (2009) predict that the share of the HRS cohort retiring between 62 and 66 would have been reduced by about 1 to 2 percentage points under the retirement rules experienced by the Early Boomer cohort – but they cannot look at how well this compares to what actually occurred. While the share of the increase explained by the Gustman and

Steinmeier study is slightly smaller than that explained by Blau and Goodstein, both suggest that changes in Social Security rules do explain some of the increase in the retirement age discussed above.[1]

2.2 Retirement benefits

Like Social Security rules, the rules of an individual's retirement plan can also influence retirement timing, and the topic has received extensive attention in the literature. In general, the research suggests that DB plans, which in the USA often have early and normal retirement ages at which time people can receive a reduced and then normal benefits with little incentive to work longer, have tended to decrease retirement ages.[2] For example, reduced form work by Samwick (1998) using the Survey of Consumer Finances linked to the Pension Provider Survey showed that increases in DB coverage could explain roughly a quarter of the decline in labor force participation at older ages observed during the 1980s. But as DC plans have come onto the scene, those trends have reversed. After all, the value of working an additional year with a DC plan is the value of contributions from the employee and the employer, which decreases with age only because returns have less time to accrue. Perhaps not surprisingly then, Munnell *et al.* (2004) found that workers with DC plans had higher expected retirement ages than those with DB plans. Friedman and Webb (2005) estimated that the shift from DB plans to DC plans would increase the median retirement age by 9 to 12 months for workers over the three decades between 1983 and 2015. More recently, Rutledge *et al.* (2015) noted that the continued decline of DB pensions is likely to push the average retirement age up in the future. To the extent DC plans became more common in between the HRS and Early Boomer cohorts, this literature suggests the effect will augment the effect of changing Social Security rules. Indeed, Munnell *et al.* (2016) provide ample evidence of this trend, showing that between the HRS cohort and the Early Boomer cohort the percent of households covered by a retirement plan with at least some DB wealth dropped from 79.9% to 68.6% and the share with some DC wealth increased from 56.2% to 74.0%.

2.3 Health

Like the impact of retirement plans on retirement, the relationship between health and retirement has been extensively researched. This literature generally places an emphasis on differentiating the causal impact of health on retirement separately from either (1) reverse causality; or (2) justification bias. For example, Blau and Gilleskie (2001) use a semiparametric random effect estimator and a variety of health measures, spanning self-report and diagnosed conditions, to get around these two issues and find a substantial impact of health on the timing of retirement. Bound *et al.* (2010) measure health using an index reflecting both self-reported and true measures and find those in poor health are five to ten times as likely to exit the labor force prior to age 62 as those in average health. Maurer *et al.* (2011) account for justification bias by using self-reported variables on day-to-day functioning (which presumably individuals do not manipulate) to predict an individual's overall health and find that health is again a significant determinant of retirement. Broadly, it is clear from the literature that health has a large and causal effect on retirement timing. To the extent health has improved between the HRS and Early Boomer cohorts, one expects it to have an additive effect with changes in Social Security rules and the nature of pension benefits.

[1]Other studies have examined policies not explored here, typically because those policies were counterfactual and thus would not explain observed behavior. For example, Gustman and Steinmeier (2015) found that increasing the Earliest Eligibility Age (EEA) from 62 to 64 had large effects on the share of workers at ages 62 and 63. Van der Klaauw and Wolpin (2008) explored policies like the elimination of early retirement, focusing only on low-income workers without pensions. French (2005) finds that removing the Social Security earnings test would delay job exit by a year – a larger effect than benefit cuts.

[2]An excellent description of these incentive effects can be found in Kotlikoff and Wise (1987).

2.4 Mortality

The final factor this study examines is the effect that any increase in life expectancy that occurred between the HRS and Early Boomer cohorts might have on retirement timing. While it is clear that a longer life expectancy should yield later retirement ages, especially for those with limited annuitized wealth, research on whether this relationship actually shows up in the data is somewhat limited. This limited research likely follows from the obvious correlation between mortality and other factors correlated with retirement timing, such as health and wealth. Instead, much of the literature has relied on variation in subjective life expectancy to identify any relationship between life expectancy and retirement timing.[3] Hurd *et al.*, (2002) use HRS data and show that those with a low perceived probability of surviving until age 85 retire earlier than those with higher perceived probabilities. In a similar study using data from the Netherlands, Van Solinge and Henkens (2010) show that subjective life expectancy is correlated with longer anticipated careers.

The structural model developed by Gustman and Steinmeier incorporates these four features and allows them to be manipulated to see how, in the context of a life-cycle model, they should influence behavior. This paper turns to a description of that model next.

3. Model

The structural model used in this paper has been described elsewhere in detail, perhaps most completely in Gustman and Steinmeier (2006). Still, a review of the model and then a focus on the four issues discussed above is useful.

To review the basics, the model takes a standard life-cycle approach and estimates the retirement behavior of males that start their time in the HRS as part of a married couple. These individuals are assumed to maximize their expected lifetime utility subject to constraints and uncertainty regarding the returns to assets, the mortality of themselves and their spouses, and their preferences. The individual's potential wages (realized if they work) and health are treated as non-stochastic – their paths are assumed to be known by the individual.[4] The spouse's wage income and pensions are also treated as exogenous, as is their retirement decision (i.e., only the male member of the couple's retirement decision is being modeled). If the spouse is no longer alive, this term is zero. The main choices in the model are the decision to work and the decision of how much to consume. Lifetime utility is given by the following equation:

$$EU_i = E_i\left[\sum_{t=a}^{T}\left\{e^{-\rho t}\sum_{m=1}^{3} s_{m,t}\left(\frac{1}{\alpha}C_{m,t}^{\alpha} + h_t L_{m,t}^{\gamma}\right)\right\}\right] \qquad (1)$$

In equation (1), m indicates the individual's marital status, which can be married with: (1) both spouses alive and present; (2) only the married male remaining; and (3) only the spouse remaining. The vector $s_{m,t}$ represents the likelihood of the male's marital status arising by time t and thus encompasses mortality risk. The model assumes that the decision-maker (in this model the male member of the couple) does not get any utility once they are dead and thus does not account for the consumption possibilities of their spouse once they are deceased. It also implies that individuals in the model do not have a bequest motive; any non-annuitized wealth left over at death is left accidentally and yields no utility. The vector h_t represents the individual's preference for leisure, which is allowed to vary by age in a linear manner and as a function of the person's health. $L_{m,t}$ represents the leisure decision itself which can be working ($L_{m,t} = 0$), partially retired ($L_{m,t} = 0.5$), or fully retired ($L_{m,t} = 1.0$), with γ indicating an individual's preference for partial retirement (see Section 4.1 on data for the study's

[3]Other research has shown that subjective survival probabilities do in fact contain information about actual mortality. For example, see Hurd and McGarry (2002).

[4]A later version of the model, Gustman and Steinmeier (2014), incorporated uncertainty in health and found that this addition did not change the overall conclusions of the model nor its implications for policy.

definition of retirement).[5] Finally, the parameters ρ and α represent an individual's time preference and the concavity of their utility function with respect to consumption, respectively.

This formulation of the utility function makes several assumptions worth further discussion. The first assumption is that the utility being maximized is at the household-level and without regard to the presence of the spouse – a single male with a deceased spouse and a married male would get the same amount of utility from the same amount of consumption. Having his wife die would affect the amount of consumption available to the male through the loss of any income due to her, but would not alter the translation of consumption to utility. The second assumption is that leisure and consumption are assumed to be additive. This assumption means that the decision to retire is basically a comparison of the consumption value of lost future income vs. the disutility of continued work and that the value of leisure cannot be improved by a longer career and thus more available income. The final assumption is that while health influences the leisure decision, it cannot be influenced by the individual's behavior (i.e., it is not a commodity that can be improved by devoting income do it).

If any of these assumptions were far off the mark, the likely consequence would be the tendency of the model to predict retirement that is too early relative to the actual data. Maximizing household-size adjusted consumption would emphasize early consumption and encourage a longer period of work that the model would have difficulty fitting. If consumption and leisure were modeled as complements the likely result would also be a longer career, as the individual would look to maintain higher consumption longer. Finally, the ability to manipulate health with income would also push toward later retirement. However, as will be discussed below, the model does a good job of predicting actual behavior, suggesting that these assumptions do not significantly impair the model's ability to fit the data.

The amount of assets a household has available for consumption is governed by the asset constraint:

$$A_t = (1 + r_t)A_{t-1} + W_t(1 - L_{m,t}) + E_{m,t} + B_{m,t} - C_{m,t} \tag{2}$$

Equation (2) is a standard budget constraint where the difference between assets today and tomorrow depends on: (1) yesterday's assets (A_{t-1}) and any returns those assets accumulated (r_t); (2) wages (W_t), which are allowed to be lower once an individual has quit their first job; (3) pension and wage income from the spouse ($E_{m,t}$); and (4) income from the individual's pensions and from Social Security ($B_{m,t}$). The rate of return is stochastic from the point of view of the individual and is distributed normally based on the mean and standard deviation of a portfolio containing the average amounts of stocks and bonds held by HRS sample members.

The intuition behind equations (1) and (2) is simple: individuals make their retirement and consumption decisions today knowing that their decisions alter their utility and their level of income and assets available to them both today and in the future. For example, retiring today brings some amount of utility in the moment, eliminates wage income, and introduces Social Security and retirement income into today's budget. It also affects future consumption and utility by locking in the actuarial reduction on Social Security benefits and by lowering the wages available if the individual decides to work again. With these basics in hand, the modeling of Social Security benefits, retirement benefits, health, and mortality can be discussed.

3.1 Social Security benefits

The modeling of Social Security benefits, a part of $B_{m,t}$ in equation (2), is straightforward. Because potential income from labor is known to the individual, the benefit from claiming Social Security at each given age is known in advance given the Social Security rules the person's birth cohort faces. Individuals are assumed to claim Social Security as soon as they are eligible and their earnings

[5]In practice, the preference for partial retirement is allowed to be a function of some individual-specific preference, δ, and a 'shifter' that potentially increases that preference with age, δ_a.

fall under the earnings test. Although this assumption is based on somewhat dated evidence from Coile *et al*. (2002) that people typically claim Social Security as soon as possible, MacInnis (2009) suggests that while retirement and claiming are slightly less correlated than in the past, they are still highly correlated. At this point, they receive their primary insurance amount reduced by any actuarial adjustment (or increased by the DRC) as determined by their birth cohort and the age at which they retired.[6] These assumptions mean that Social Security benefits in the model tend to keep people in the workforce until age 62, and that delay has the benefit of increasing post-retirement income.

3.2 Retirement benefits

The Gustman and Steinmeier model includes both DB and DC plans. For DB pensions, the evolution of the benefit is determined by the individual's wages on their 'career job' – the job the person has until they initially retire or partially retire – and by the rules of their pension plan. When the person retires, the annuity benefit is fixed (up to cost of living adjustments) and is received until death. The benefit an individual receives at each retirement age is based on the actual rules of their retirement plan as contained in the HRS data (see Section 4.1 for a discussion). To the extent most US DB plans have early eligibility ages after which point individuals can receive a benefit and normal retirement ages after which point there is little to gain from continued work, then these features will be reflected in the data and will significantly affect retirement timing.

DC plans on the other hand do not have such features. In the model, DC plans behave much as they do in the real world – contributions are placed in the account, they grow at an uncertain rate (whose distribution is discussed along with equation (2)), and when the individual retires the account is available for consumption.[7] Because the DC plan is not assumed to be annuitized – as Johnson *et al*., (2004) find only 10% do – and because DC plans accrue assets smoothly they are likely to be associated with later retirement relative to DB plans.

3.3 Health

Health enters the model through h_t in equation (1), which indicates the individual's utility from leisure. Specifically, the model specifies that the leisure term is:

$$h_t = e^{\beta X_t + \varepsilon_t} \tag{3}$$

In equation (3), X_t is a vector that contains a constant, the individual's age, and whether they self-report to be in poor health.[8] In the model, people are not assumed to go back and forth between poor health – once they are in poor health they are assumed to stay in poor health. The term ε_t reflects an uncertain component in the individual's preference for leisure that could draw them into or out of retirement. In the model, this uncertainty initially comes from a distribution with mean zero and holds a single value until the individual retires from their career job. After retirement from this job, ε_t, is allowed to vary from year to year, with a correlation parameter of ρ_ε. A positive coefficient on poor health would indicate that it speeds up retirement by increasing the value of leisure relative to work.

[6]Because the spouse's income is treated as exogenous, their claiming decision is not modeled, but the size of the household benefit does depend on whether both members are alive and on the lifetime earnings and subsequent benefit size of the spouse.

[7]The distribution of rates of return was calculated by Gustman and Steinmeier (2006) using returns calculated by Ibbotson Associates (2002) and taking into account the distribution of financial assets in the HRS. The exception to the real world modeling of DC plans is that, in the model, DC wealth is treated like other savings and therefore does not have any penalty placed on withdrawals before retirement or any required minimum distributions after age 70 and half.

[8]Although an imperfect measure, Dwyer and Mitchell (1999) find that it performs similarly to more complicated indices of health conditions.

3.4 Mortality

Mortality enters the model through the term $s_{m,t}$ which indicates whether both the male respondent and their spouse, the male respondent only, or the spouse only is still alive. Individuals know these probabilities in advance, but they do not know how long they will actually live. Mortality probabilities are based on Social Security Cohort mortality tables for the relevant birth cohort. In the model, a decrease in mortality would lead to delayed retirement to the extent that longer lifespans require higher levels of wealth or higher Social Security benefits at retirement. The effect of mortality interacts with the effect of pensions, since DB plans provide income until death but DC plans would not.

4. Data and estimation

This section begins by describing the data gathered from the HRS to estimate the parameters of the model described above and then describes how estimation is carried out.

4.1 Data

The data used in this analysis come from the HRS. This paper uses data from the 1992 HRS cohort (born 1931–1941), as was used in Gustman and Steinmeier (2006, 2009), and also uses the 2004 Early Baby Boom cohort (born 1948–1954). For both cohorts, the sample begins with married males and, as in Gustman and Steinmeier (2006), reduces the sample by introducing a number of restrictions in it. Table 1 lists these restrictions and shows how the sample is reduced from the initial sample of married males is to reduce the final sample used in the estimation of the model and ultimately in the simulations. Table 1 indicates that these restrictions limit the HRS cohort from an initial size of just over 12,000 individuals down to about 2,200 observations and the Early Boomer cohort from just over 3,000 observations to about 400 observations. In both cohorts, the most important restrictions are that individuals with retirement plans but without data on those pensions from their employers are dropped, as are individuals with unstable employment histories.[9] For each cohort, six waves of data are used; for the HRS cohort Waves 1–6 are used (1992–2002) and for the Early Boomer cohort Waves 7–12 (2004–14).

Perhaps the most important individual variable collected is an individual's status as working, partially retired, or fully retired. In the data, individuals working at least 30 h each week and 1,560 h a year are counted as full-time workers. Workers who work between 100 h and 1,250 h a year and who do not exceed 25 h per week are labeled as part-time and partially retired. Workers who work between 1,250 h and 1,560 h per year or between 25 h and 30 h a week are classified based on self-report. Workers who work fewer than 100 h per week are said to be fully retired. The focus of this paper is on complete retirement, so Figure 1 shows how the share of workers who are completely retired changes with age by cohort. The figure illustrates that workers in the Early Boomer cohort retire at similar rates as the HRS cohort up until age 60, but are less likely to retire from that point on. For example, at age 64 (the last age a significant number of Early Boomers are observed) about 53.9% of the HRS cohort is fully retired and 49.5% of the Early Boomer cohort.

Data to determine the resource constraint, equation (2), come from a variety of sources. Where possible, earnings come from actual Social Security records. For the remaining respondents (roughly a quarter) who are not linked to such records, earnings histories are constructed from self-report on any current job or any previous job that lasted more than 5 years. Because behavior is modeled for most individuals beyond their time in the data – the oldest Early Boomer member is only 66 by 2014 – and because earnings are assumed to be known by the individual, earnings must be projected into the future. This projection is accomplished using an individual's experience and tenure and the same regression used in Gustman and Steinmeier (2006). An individual's Social Security Primary Insurance Amount is calculated based on this earnings history and updated as they continue working.

[9]The appropriateness of these restrictions is discussed at length in Gustman and Steinmeier (2006).

Table 1. Sample size following restrictions imposed by Gustman and Steinmeier (2006)

Restriction criteria	HRS cohort Observations dropped	HRS cohort Observations left	Early Boomer cohort Observations dropped	Early Boomer cohort Observations left
Total sample of individuals		12,652		3,330
Males married at first sampled wave	6,785	5,867	1,801	1,529
Divorced or widowed after age 35	1,578	4,289	489	1,040
Spouse not interviewed in first wave	133	4,156	39	1,001
Respondent does not have career job	497	3,659	182	819
Ambiguity about Social Security coverage	49	3,610	6	813
No full-time earnings	36	3,574	69	744
No self-reported earnings, Social Security earnings over limit	31	3,543	11	733
Relatively large business assets	291	3,252	165	568
No pension provider record in last job	865	2,387	92	476
Fulltime years of work unavailable for spouse	156	2,231	80	396

Source: Authors' calculation from HRS and Gustman and Steinmeier (2006).

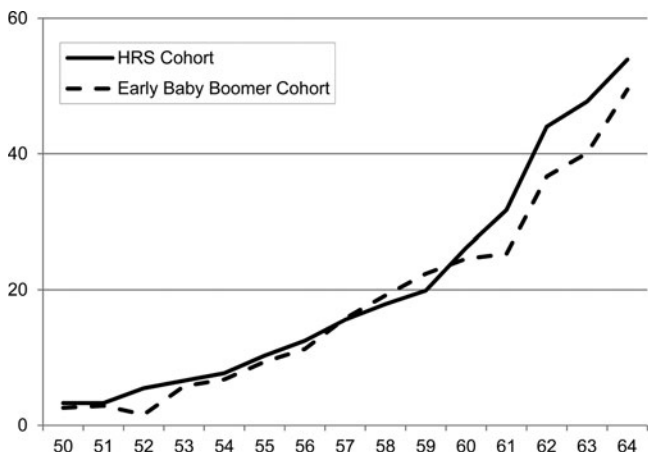

Figure 1. Share of sample completely retired age 50–64, by cohort.
Source: Authors' calculations from HRS.

Their benefit is determined by the rules when they claim, which again is taken as the first instance when eligible to claim and their earnings falls below the earnings test.

Information on the pension annuity benefit comes from summary plan descriptions provided by the individuals' employers. For DC plans, the amount of money in an account is updated year by year based on the individual's earnings and contribution rates from them and their employer. The assets in the DC account are assumed to grow at the same stochastic rate of return as any other assets. Figure 2 shows that between the HRS and Early Boomer cohort a shift away from DB and toward DC plans occurred.

The final source of data needed to complete equation (2) is on resources coming from spouses. For spouses, earnings, retirement benefits, and DC wealth are calculated much as they are for the male being modeled. But these streams of income are treated as exogenous to the male being modeled, i.e., they are taken as given by the male being modeled (unlike his own earnings, which can be affected by the decision to work).

Equation (3) makes clear that another piece of information needed to estimate the model is information on health, since health impacts the disutility of work. In the model, a person's health is treated as binary variable equal to 1 if the individual self-reports being in poor health and 0 otherwise. Figure 3 shows how the share of sample members in poor health evolves over time for the HRS

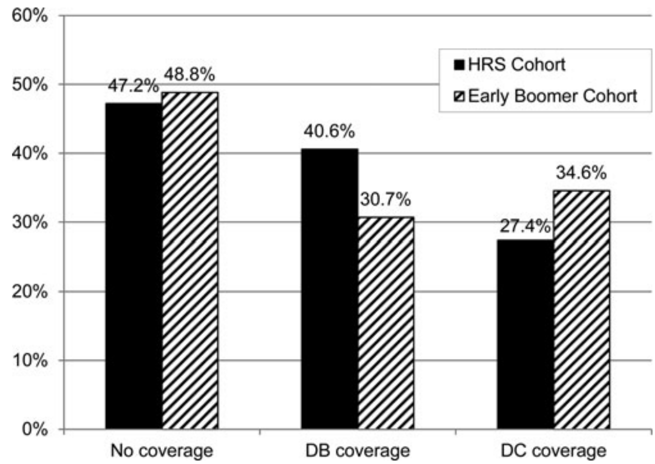

Figure 2. Share of sample with retirement plan coverage, by cohort.
Source: Authors' calculations from HRS.
Note: Numbers within a cohort may sum to over 100% due to those with multiple plans.

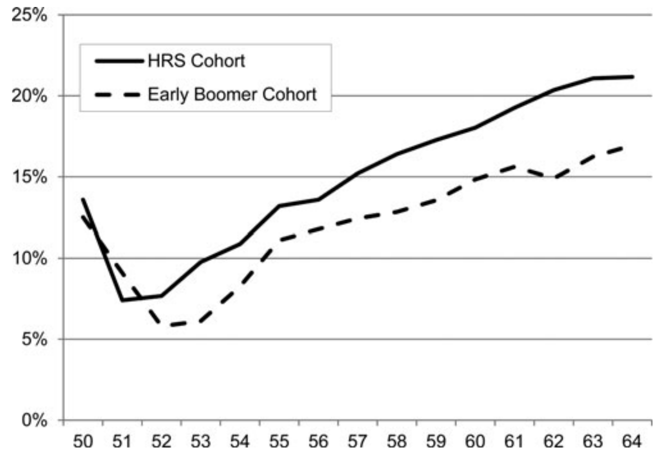

Figure 3. Share of sample in poor health, age 50–64, by cohort.
Source: Authors' calculations from HRS.

and Early Boomer cohorts. The figure shows that, in general, the Early Boomer cohort has a lower share in poor health and that this gap widens with age.

4.2 Estimation

The goal of estimating the model is to obtain the preference parameters in equations (1) and (3) so that behavior can be simulated. These include the consumption parameter in equation (1), the distribution of the preference for partial retirement in equation (1), the distribution and autocorrelation of uncertainty in a person's preference for leisure in equation (3), and the effect of age and health on the preference for leisure, also from equation (3). Each of these parameters has a unique effect on behavior in the model – for example, if the parameter on health in the preference for leisure is positive it will mean people tend to retire when unhealthy. In this paper, the preference parameters estimated for the HRS cohort are held constant as Social Security rules, pension benefits, health, and mortality are changed around them to match up with the Early Boomer cohort.

As is standard in dynamic stochastic structural models like that used in this paper, the model is estimated iteratively and recursively. For a given set of parameters, the model is solved by starting in the final modeled year, T, and determining how the individual would behave at this time for each possible combination of preferences, earnings, assets, Social Security income, and pension income the person could have achieved. Because it is the final period of the model, the individual's decision is relatively easy at each of these points – they consume all their available assets and income and work if it yields higher utility.

Going back a period, the decision is slightly harder. Again, given the parameters guessed, the utilities are calculated at each possible model outcome, but now for each amount consumed a certain amount of assets will be left over next period and the decision to work or not will alter next period's possible income. Still, given that the distribution of the rate of return is known and given the assumed distribution of uncertainty in the preference for leisure, the individual can calculate the probability they end up at any particular point in the next period. And given they know exactly what utility they will have if they get to that point, they can calculate the best course of action this period, both in terms of its utility and tomorrow's expected utility. Work and consumption decisions are modeled in this way going back to the first period of the model.

Given the assumed parameters and the calculations described above, and given draws from the distribution of leisure preferences and the rate of return, each individual in the model will have a simulated track of behavior from the time they are first observed until death based on the data collected. To see how well these simulated tracks fit the data, the model uses the generalized method of moments (GMM), which is less sensitive than the other major approach used, Simulated Maximum Likelihood, to cases where the observed result is not very likely in the model. Essentially, GMM sees how well the 'moments' predicted by the model – for example, the share of people in poor health at a given age that are retired – match up with the actual data. If the guessed parameters are close to the true parameters, these moments should match up well, if not, the estimation procedure will guess again. Returning to the example of the parameter on health in the preference for leisure, if in the data people in poor health retire more frequently than others, the model would only fit with a positive coefficient on this parameter. If a negative coefficient was guessed, the procedure would recognize a poor fit and eventually move to a better one. In this iterative way, the model arrives at parameters for which the model fits the chosen moments.[10]

4.3 Simulation

To simulate the changes discussed above, various aspects of the models are adjusted and new draws are obtained from the preference and returns distribution to allow behavior to be re-simulated. For example, the stream of Social Security benefits available to an individual retiring at the various ages can be changed to reflect an increased FRA. Although these changes will be discussed in more detail below, the approach implicitly assumes that preferences estimated are the same for the HRS Cohort and the Early Boomer Cohort. This assumption is both an advantage and limitation of the approach. The advantage of the assumption is that it allows the effect of the various changes to be isolated from any changes in preferences that may have occurred. The disadvantage is that to the extent preferences did change, the effect on actual behavior would be misstated. However, as will be discussed below, the results suggest that the four changes above can themselves explain much of the change in retirement timing between these two cohorts, suggesting that changes in preferences may be somewhat limited.

5. Results

While the main purpose of this paper is to examine how much of the change in retirement timing between the HRS and the Early Boomer cohort the four explanations presented above can explain, it is useful to quickly revisit how well the Gustman and Steinmeier data fit the HRS cohort it was

[10]More details are provided in Gustman and Steinmeier (2006), including on how iterative process is carried out and how the fit of the model to the moments is evaluated. The moments used include the share fully retired at all ages 54–66, the share partially retired at ages 55, 58, 60, 62, and 65, the share fully retired in the upper and lower third of lifetime income at ages 55, 58, 60, 62, and 65, the share fully or partially retired who are in poor health at ages 55, 58, 60, 62, and 65, and the frequency with which individuals return to full-time work given that they were fully or partially retired. All in all, there are 43 moments.

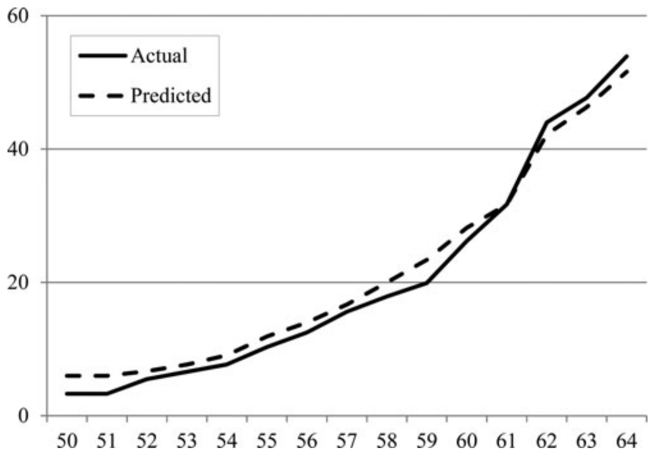

Figure 4. Share of sample completely retired age 50–64, HRS cohort.
Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

estimated on.[11] This exercise is important for two reasons: (1) the believability of any projected changes in behavior due to those factors depends on the model fitting the original data well; and (2) although in this paper a model similar to Gustman and Steinmeier (2006) was exploited and their original work showed good model fit, the data used here were reconstructed and thus could deviate from their results.

Figure 4 shows the observed share of the HRS cohort completely retired in the actual data and then predicted from the model from ages 50 to 64 – the relevant age range when comparing the HRS cohort to the Early Boomer cohort. The model tends to fit the HRS cohort's behavior fairly well throughout ages 50 to 64, which slight over prediction of retirement in the late 50s and a slight under prediction in the 60s. Importantly, the predicted results show the characteristic jump in retirement at the Early Eligibility Age despite the fact no preference parameter is included in the model to make such a jump occur. The behavior is entirely driven by the presence in the resource constraint of the Social Security benefit. Now, the project turns to the four simulations which will parse out how much of the change in retirement timing between the HRS and Early Boomer cohorts seen in Figure 1 can be explained by the four explanations being considered. The results of all four simulations are shown in Table 2 and the following discusses the findings in turn.

5.1 Social Security benefits

To simulate the effect of changing Social Security rules, each person in the HRS cohort was assigned the benefit rules of someone 12 years younger than themselves and then re-run through the model (importantly, preference parameters were not changed). For most HRS cohort members, this change means their FRA was increased from 65 to 66, the DRC was increased, and the earnings test above the FRA was removed. The effect is shown in the second and third columns of Table 2 and illustrates that the changes to Social Security rules experienced by the Early Boomer cohort relative to the HRS cohort would not be expected to have any effect on retirement ages prior to the early eligibility age. Once the early eligibility age is reached at 62, the initial effect is to actually increase the share completely retired by about 1.1 percentage points. Gustman and Steinmeier (2009) note this somewhat counterintuitive increase as well. However the predicted increase is short-lived – the change in rules decreases the share expected to be fully retired at age 64 by about 0.7 percentage points. This is roughly 16% of the observed gap in complete retirement between the HRS and Early Boomer cohort observed at age 64 and shown in Figure 1.

[11]The estimated coefficients of the model used to generate these predictions are shown in the Appendix.

Table 2. Results from simulations of HRS cohort with various characteristics of Early Boomer cohort

		Social Security rules + 12 years		Retirement benefits of Early Boomers		Health of Early Boomers		Mortality of Early Boomers	
Age	Predicted using raw HRS data (%)	Predicted (%)	Percentage point change (%)	Predicted (%)	Percentage point change (%)	Predicted (%)	Percentage point change (%)	Predicted (%)	Percentage point change (%)
50	6.0	6.1	0.1	6.3	0.3	6.0	0.0	6.0	0.0
51	6.0	6.0	0.0	6.3	0.3	6.1	0.1	6.0	0.0
52	6.7	6.7	0.0	6.7	0.0	6.8	0.1	6.7	0.0
53	7.7	7.7	0.0	7.7	0.0	7.8	0.1	7.6	−0.1
54	9.1	9.1	0.0	10.1	1.0	9.2	0.1	9.0	−0.1
55	11.9	11.9	0.0	11.7	−0.2	11.8	−0.1	11.8	−0.1
56	14.0	14.0	0.0	13.8	−0.2	14.0	0.0	13.9	−0.1
57	16.7	16.7	0.0	16.4	−0.3	16.4	−0.3	16.6	−0.1
58	20.0	20.0	0.0	19.4	−0.6	19.7	−0.3	19.8	−0.2
59	23.4	23.3	−0.1	23.0	−0.4	23.0	−0.4	23.2	−0.2
60	28.2	28.2	0.0	27.0	−1.2	27.9	−0.3	28.0	−0.2
61	31.7	31.7	0.0	29.8	−1.9	31.3	−0.4	31.5	−0.2
62	42.2	43.3	1.1	39.6	−2.6	41.2	−1.0	41.9	−0.3
63	46.3	45.8	−0.5	43.2	−3.1	45.3	−1.0	46.0	−0.3
64	51.6	50.9	−0.7	48.2	−3.4	50.4	−1.2	51.2	−0.4

Source: Authors' calculations from HRS and Gustman and Steinmeier (2006).
Note: Although this table does not show standard errors, in general two sources of uncertainty exist with respect to these simulations. The first is that the size of any difference in behavior could depend on the chosen parameters, which are estimated with uncertainty. However, under reasonable parameter draws, the sign of these effects is unchanged. The other is uncertainty due to the imputation used in the retirement benefit and health simulations. However, in simulations using different imputation draws, the level of uncertainty introduced is very small. Results for both types of uncertainty are available upon request.

5.2 Retirement benefits

To simulate the effect of changing retirement benefits, individuals in the HRS cohort had their retirement coverage and benefits imputed using data from the Early Boomer cohort. This paper accounts for possible correlations between observed variables and pension coverage by using multiple hot-deck imputation (Rubin, 1987). Using occupation, industry of employment, and earnings to link individuals with similar characteristics in the HRS cohort and Early Boomer cohort, the hot-deck process reflects both any changes in the pension rules – i.e., whether the pension was DB or DC, if DB the minimum age at which benefits could be collected, the rate of growth of the benefit with tenure, etc. – and the actual size of the benefit or wealth accrued.[12] As Figure 3 implies, the net effect was to increase the share of all pensions in the HRS cohort that were DC plans. This change is potentially important, since DC plans do not have discrete points when benefits can be claimed and which would tend to trigger retirement and they do not provide benefits for life, potentially leading to the need for delay.

The fourth and fifth columns of Table 2 show that, in fact, retirement plans explain a large amount of the change in retirement behavior between the two cohorts. Early on, the shift from DB pensions to DC plans is predicted to actually increase the share retiring in the early 50s. This effect likely follows from the fact that some people who have DBs in reality but were assigned DCs in the simulation had been remaining in work to hit the age at which they collected DB benefits. With DC plans, no such incentive exists. But once the age crosses the mid-50s, the change in pension structure is uniformly expected to decrease the likelihood of retirement. At age 62, the pension shift is expected to decrease the share completely retired by 2.6 percentage points, or about one-third of the total gap between the HRS and Early Boomer cohort at this age. At age 64 the pension simulation predicts a drop of 3.4 percentage points relative to a 4.4% drop in reality, or about three-fourths of the total drop.

[12]Besides using the information from occupation, industry and earnings as characteristics, the imputation process also controls for the relative level of pension benefit or pension wealth to the individual's earning history to keep the correlation between the pension and earning's profile.

5.3 Health

To simulate the effect of improving health, each person in the HRS Sample had the age at which they experienced their first bout of poor health imputed based on the health distribution of individuals with similar occupations, industries and earnings from the Early Boomer cohort.[13] The effect of this change is shown in the sixth and seventh columns of Table 2 which show that had the HRS cohort experienced the health of the Early Boomer cohort, it would have been less likely to retire at the various ages shown. For example, at age 62 the share expected to be fully retired under this experiment is 1.0 percentage point lower than in the original simulation, or about 14% of the gap between the HRS and Early Boomer cohorts. By age 64, the drop is even larger, at 1.2 percentage points, or 27% of the gap between the two Cohorts shown in Figure 1. Health seems to play a role in the later retirement of the Early Boomer cohort, although the effect is only about a third the effect of retirement plans.

5.4 Mortality

To simulate the effect of decreasing mortality, each individual in the HRS cohort was assigned the mortality probabilities used to construct $s_{m,t}$ from equation (1) using mortality rates from a birth cohort 12 years younger. The effect of this change is shown in the last two columns of Table 2 and, like health, decreases the predicted likelihood of retirement at each age albeit by a smaller amount. For example, thle effect at age 62 is a 0.3 percentage point reduction, or about 4% of the actual difference between the two cohorts. At age 64, the reduction is similar, at 0.4%, although that represents a larger share, 9%, of the actual gap.

5.5 Combining the experiments

To see what share of the change in retirement timing is due to these four factors, all four experiments were applied to the HRS cohort simultaneously. The result of this exercise is shown in Figure 5. The result is largely consistent with reality – the changes modeled here were not predicted to affect retirement timing much before age 60 and indeed little change occurred between the two cohorts. But the figure also shows that the combination of the four experiments largely explains the decrease in retirement from age 60 on. This suggests that to the extent preferences changed between these two cohorts, the changes were not large; changes in the environment being modeled can explain must of the drop in retirement at the ages examined.

From the prior sections it is clear that the change in the structure of retirement plans is the largest contributor to this decline, with changes in Social Security rules and health playing a smaller role. Interestingly, if anything, at age 64 the experiment over predicts the reduction in retirement that would be experienced by these changes. Figure 5 shows that at age 64 the simulations would predict the share retiring given these four changes would be 44%, 3.3 percentage points lower than was actually observed. This over prediction could reflect the fact that some other factor not tested here, like preferences, changed between the two cohorts. However, it is worth keeping in mind that only about 25% of the Early Boomer sample had made it to age 64 by 2014, so it is worth revisiting this experiment once this cohort ages further to see if the share retired comes more in line with the share predicted, especially given the good fit at younger ages.

Given these four changes to the HRS cohort's characteristics result in behavior that closely mirrors that of the Early Boomer cohort, an interesting final exercise is to see how one might expect the Early Boomer's retirement to play out over time. Figure 6 extends the simulation shown in Figure 5 out to age 69 and shows that the expected gap between the behavior of the actual HRS cohort and its behavior given the simulation would only be expected to grow with time. By age 69, just 64.8% of the HRS cohort would be expected to retire if they had the Social Security rules, pensions, health, and mortality

[13]The imputation process is similar to the one used in pension experiment; one difference is that the imputation is only performed for individuals in the HRS cohort whose age in 1992 is less than or equal to 57, which is the maximum age observed in 2004 for Early Boomer cohort.

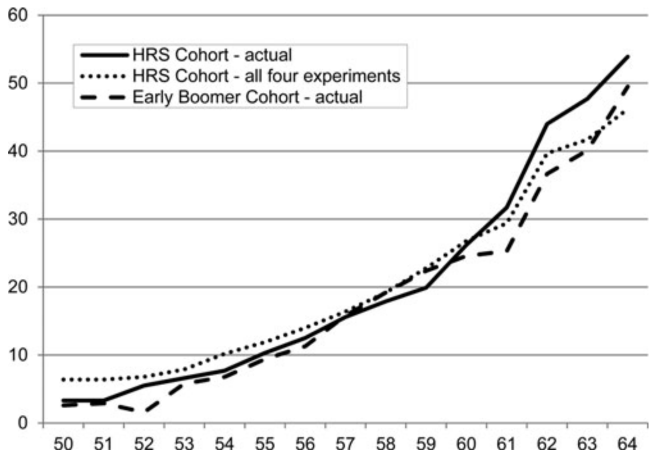

Figure 5. Share of sample completely retired age 50–64, by cohort.
Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

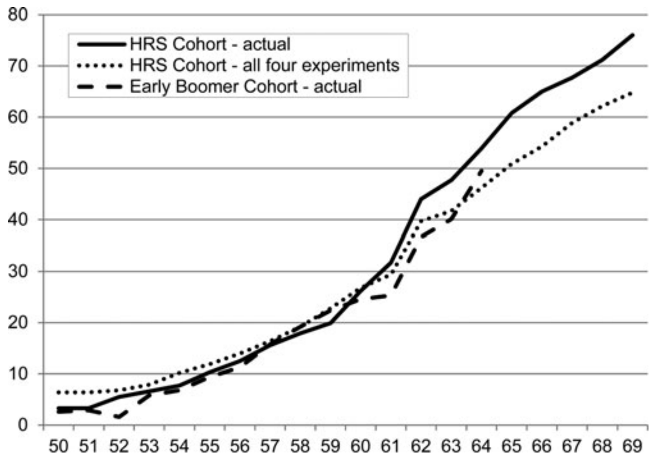

Figure 6. Share of sample completely retired age 50–69, by cohort.
Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

of the Early Boomer cohort, about 11 percentage points lower than what they actually experienced. This makes sense – the Social Security rules faced by the Early Boomer cohort include a push back of the FRA and the elimination of the earnings test, rules that have a much larger impact after age 64. Although the fact that at age 64 the simulations tended to over predict the drop in retirement observed should give some pause in extending the predictions out too far, the experiment suggests that the Early Boomer cohort is likely to continue to have lower retirement rates than the HRS cohort.

5.6 Limitations

Before concluding, a discussion of the limitations of this study is useful. Like many papers involving structural models, a variety of assumptions were made to simplify the computational burden of estimating the model. Perhaps the most limiting in the case of this paper is that to avoid modeling the joint behavior of household members (and to not focus on single individuals, who are relatively uncommon near retirement) the model focuses on married men, taking the behavior of wives as exogenous. Thus, the results of the simulations should be viewed as the expected effect of these policies on married men specifically and assuming any change in their wives behavior does not further affect their decision making. This assumption is restrictive. However, to the extent that the assumption damages model performance, it does not show itself in either the model's ability to predict the HRS Cohort's actual behavior. The fact that when changing the economic environment around the

HRS Cohort to be similar to that of the Early Boomer Cohort with respect to Social Security, pensions, health, and mortality the results look similar to the Early Boomer is also somewhat encouraging. However, given women's increasing presence in the labor force, future research should consider modeling the joint behavior of couples.

The other limitations of this paper revolve around the assumptions that put structure on men's preferences, most of which were discussed as the model was introduced. The largest of these are that men do not internalize the utility of their spouse once he dies, that the utility from consumption and leisure is separable, and that income from earnings and the worker's health are known in advance. These assumptions likely all serve to limit the value of working longer or reduce the insurance value of saving. Still, the model is able to fit observed behavior well, which is at least an indicator that they do not adversely affect the model's ability to fit the data. Indeed, in recent work Gustman and Steinmeier (2014) indicating incorporating health uncertainty into this model did not significantly affect the results.

6. Conclusion

Between the original HRS cohort and the Early Boomer cohort of the HRS, the share completely retired at ages 60 and above decreased substantially, while remaining relatively unchanged during the 50s. The simulations conducted here using the Gustman and Stienmeier (2006) structural model suggest that the reason for this pattern is mostly the change in pension structure. The simulations make clear that while the shift from DB to DC retirement plans would not have large changes on retirement timing of people in their 50s, workers in their late 50s and early 60s become less likely to retire given this shift. Improvements in health also contributed moderately to the decline in retirement observed in the early 60s between the two cohorts. The other changes – Social Security rules changes and decreasing mortality rates – generally reinforce this pattern but have slightly smaller effects.

These simulations shed light on what to expect in the future. For the Early Boomer cohort itself, we should expect continued lower rates of retirement relative to older cohorts with, if anything, a widening gap. And for future cohorts – namely the Mid- and Late-boomers – we should expect a continuation of the trend toward delay. The reason for this trend is mostly the expectation that fewer and fewer retirees will be reaching retirement with DB pensions, although after the entire Boomer Cohort reaches retirement this trend will likely have played out. And any continuation to improvements in health and mortality will cause the trend toward delay to continue, albeit at a slower pace, even once the vast majority of retirees have DC plans.

Acknowledgments. The authors would like to thank Alan Gustman and Thomas Steinmeier for providing us with their model code. The research reported herein was performed pursuant to a grant from the Social Security Administration (SSA) funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the policy or opinions of SSA, any agency of the federal government, or Boston College.

References

Blau DM and Gilleskie D (2001) The effect of health on employment transitions of older men. In Polacheck S (ed.), *Worker Wellbeing in a Changing Labor Market*. Somerville, MA: Emerald Publishing Ltd. pp. 35–65.

Blau DM and Goodstein RM (2010) Can social security explain trends in labor force participation of older men in the United States. *Journal of Human Resources* **45**, 328–363.

Bound J, Stinebrickner T and Waidmann T (2010) Health, economic resources and the work decisions of older men. *Journal of Econometrics* **156**, 106–129.

Coile C, Diamond P, Gruber J and Jousten A (2002) Delays in claiming social security benefits. *Journal of Public Economics* **84**, 357–385.

Dwyer DS and Mitchell OS (1999) Health problems as determinants of retirement: are self-rated measures endogenous? *Journal of Health Economics* **18**, 173–193.

French E (2005) The effects of health, wealth, and wages on labor supply and retirement behaviour. *Review of Economic Studies* **72**, 395–427.

Friedman L and Webb A (2005) Retirement and the evolution of pension structure. *The Journal of Human Resources* **40**, 281–308.

Gustman AL and Steinmeier TL (2006) Social Security and Retirement Dynamics. University of Michigan Retirement Research Center Working Paper 2006-121.

Gustman AL and Steinmeier TL (2009) How changes in social security affect recent retirement trends. *Research on Aging* **31**, 261–290.

Gustman AL and Steinmeier TL (2014) The Role of Health in Retirement. National Bureau of Economic Research Working Paper 19902.

Gustman AL and Steinmeier TL (2015) Effects of social security policies on benefit claiming, retirement, and saving. *Journal of Public Economics* **129**, 51–62.

Hurd MD and McGarry K (2002) The predictive validity of subjective probabilities of survival. *The Economic Journal* **112**, 966–985.

Hurd MD, Smith JP and Zissimopoulos JM (2002) The Effects of Subjective Survival on Retirement and Social Security Claiming. National Bureau of Economic Research Working Paper 9140.

Ibbotson Associates (2002) *Valuation Edition 2002 Yearbook*. Chicago, IL: Ibbotson Associates.

Johnson RW, Burman LE and Kobes DI (2004) *Annuitized Wealth at Older Ages: Evidence From the Health and Retirement Study*. Washington, DC: The Urban Institute.

Kotlikoff LJ and Wise DA (1987) The incentive effects of private pension plans. In Bodie Z, Shoven JB and Wise DA (eds), *Issues in Pension Economics*. Chicago, IL: University of Chicago Press, pp. 283–340.

MacInnis, Bo. (2009) Social Security and the Joint Trends in Labor Supply and Benefits Receipt Among Older Men. Working Paper 2009-22. Center for Retirement Research at Boston College Working Paper 2009-22.

Maurer J, Klein R and Vella F (2011) Subjective health assessments and active labor market participation of older men: evidence from a semiparametric binary choice model with nonadditive correlated individual-specific effects. *Review of Economics and Statistics* **93**, 764–774.

Munnell AH (2015) The Average Retirement Age – An Update. Center for Retirement Research at Boston College Issue in Brief 15-4.

Munnell AH and Chen A (2015) Trends in Social Security Claiming. Center for Retirement Research at Boston College Issue in Brief 15-8.

Munnell AH, Triest RK and Jivan NA (2004) How do Pensions Affect Expected and Actual Retirement Ages. Center for Retirement Research at Boston College Working Paper 2004-27.

Munnell AH, Hou W, Webb A and Li Y (2016) Pension Participation, Wealth, and Income: 1992–2010. Center for Retirement Research at Boston College Working Paper 2016-3.

Rubin DB (1987) *Multiple Imputation for Nonresponse in Surveys*. New York, NY: John Wiley and Sons.

Rutledge MS, Gillis CM and Webb A (2015) Will the Average Retirement Age Continue to Increase? Center for Retirement Research at Boston College Working Paper 2015-16.

Samwick AA (1998) New evidence on pensions, social security, and the timing of retirement. *Journal of Public Economics* **70** (1), 207–236.

Van Solinge H and Henkens K (2010) Living longer, working longer? The impact of subjective life expectancy on retirement intentions and behavior. *The European Journal of Public Health* **20**, 47–51.

Van der Klaauw W and Wolpin KI (2008) Social security and the retirement and saving behavior of low-income households. *Journal of Econometrics* **145**, 21–42.

Appendix

Table A1. Parameter estimates

Symbol	Description	Coefficient value	*t*-statistic
α	Consumption parameter	−0.24	2.36
	Parameters in β		
β_0	Constant	−9.82	51.46
β_a	Coefficient of age	0.659	3.68
β_h	Coefficient of health	6.78	7.82
ρ_ε	Correlation of ε after retirement	0.88	38.54
	Parameters in δ		
δ_0	Constant	−5.34	2.12
δ_a	Coefficient of age *b*	−0.48	1.14
σ_ε	Standard deviation of ε	5.40	25.58
	q value	53.32	
	Number of observations	2,231	

Source: Authors' calculations from HRS and Gustman and Steinmeier (2006).

Cite this article: Hou W, Munnell A, Sanzenbacher GT, Li Y (2020). Why are US men retiring later? *Journal of Pension Economics and Finance* **19**, 442–457. https://doi.org/10.1017/S1474747218000331