The Behavioral and Consumption Effects of Social Security Changes

Wenliang Hou and Geoffrey Sanzenbacher
Boston College

19th Annual Joint Meeting of the Retirement Research Consortium
August 3-4, 2017
Washington, DC

The research reported herein was pursuant to a grant from the U.S. Social Security Administration (SSA), funded as part of the Retirement Research Consortium. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the federal government, or the Center for Retirement Research at Boston College.

Abstract

Social Security's Trust Fund is projected to be exhausted in 2034. A variety of changes to the program have been put forward that would either push this date out into the future or delay it indefinitely. Some of these changes cut benefits – e.g., pushing the Full Retirement Age (FRA) out to 69 – while others increase program revenue – e.g., increasing the payroll tax. While Social Security's Office of the Chief Actuary projects the financial impact on the program of a wide variety of changes, understanding the impact of these changes on recipients' behavior and well-being is also a valuable exercise. After all, any programmatic change can be calibrated to reduce Social Security's financial shortfall by a given amount, so a useful tie-breaker for policymakers may be to consider the effect on program recipients. This paper uses the Gustman and Steinmeier structural model to analyze the effect of four changes to the Social Security program on recipients' retirement timing and household consumption. All four of the changes reduce the financial shortfall by about 1 percent of payroll.

This paper found that:

- The two policies analyzed that reduce benefit size – an increase in the FRA to 69 and a reduction in the COLA of 0.5 percentage points – increase the length of workers' careers by delaying retirement.
- The two policies analyzed that increase revenue – an increase in the payroll tax to 7.75 percent and an increase in the cap to cover 90 percent of earnings – have a negligible impact on retirement timing.
- For the benefit based policies, the reduction in consumption relative to current policy is relatively high post-retirement, with the COLA adjustment having an increasing effect with time.
- Policies that increase revenue have little effect on consumption after retirement but have a consistent effect during the working life.

The policy implications of this paper are:

- Policy makers can expect a larger behavioral response of delayed retirement to accompany changes to Social Security that reduces benefits than for those increasing revenue.

- Policies that cut benefits lead to a larger reduction in consumption during retirement than those increasing revenue, but revenue increasing policies reduce consumption over a longer period of time.

Introduction

According to the 2017 Trustees Report, Social Security's Trust Fund is projected to be exhausted in 2034. Some changes to the Social Security program designed to push this date back cut benefits – e.g., pushing the Full Retirement Age (FRA) out to 69 – while others increase revenue – e.g., increasing the payroll tax. While Social Security's Office of the Chief Actuary projects the financial impact on the program of a wide variety of changes, understanding the impact on recipients' behavior and well-being is also a valuable exercise. After all, programmatic changes can often be calibrated to reduce Social Security's financial shortfall by a given amount, so a useful tie breaker for policymakers may be to consider the effect on beneficiaries. In addition, understanding behavioral effects can inform future financial projections, since the financial impact of a change like increasing the Full Retirement Age (FRA) depends on whether workers stand pat and accept a lower benefit or lengthen their careers and receive the same benefit while paying into the system longer. This paper will examine the effect of four policy changes on the retirement timing and consumption of workers across the income distribution.

The issue with projecting these behavioral effects is that in many cases the policies are completely untried. For example, one policy change that has been suggested (and that this paper will examine) is a reduction in the Cost-of-Living Adjustment (COLA). This change would effectively reduce the real benefit by an increasing amount as people age. To date, this type of benefit reduction – one that gets larger with age post-claiming – has not been attempted, so the behavioral impacts and any effect on consumption in retirement of such a policy are largely unknown.

In the absence of a natural experiment, one tool economists use to predict the effect of counter-factual policies is called structural modeling. In structural models, the preference parameters that govern individuals' choices are estimated by observing their behavior given their constraints. For example, in the context of retirement timing, if individuals with few assets retire early when they become ill, then the model would estimate that the disutility from working while ill is so high people are willing to tolerate low consumption in retirement. Once this kind of preference is estimated, the econometrician can change the environment – e.g., by making workers healthier – and see how workers' behavior is likely to respond.

This paper uses one particular structural model, from Gustman and Steinmeier (2006), to analyze the effect of changes to the Social Security program on recipients' retirement timing and household consumption. All four of the changes reduce the financial shortfall by roughly 1 percent of payroll (depending on how they are phased in), with two of the analyzed changes being benefit reductions and the other two increases in program revenue.[1] The four changes considered in this paper are: 1) an increase in the Full Retirement Age from 67 to 69 with a Delayed Retirement Credit available for delaying until age 70; 2) a decrease in the Cost-of-Living-Adjustment by 0.5 percentage points; 3) an increase in the employee portion of the payroll tax from 6.2 percent to 7.75 percent; and 4) an increase in the taxable maximum to cover 90 percent of earnings (roughly $270,000 in 2016 dollars). The first two policies, which both reduce benefits, would be expected to lead to delayed retirement as workers try to balance the disutility from continued work with the need to make-up for a reduced Social Security benefit. The second two policies, which both increase program revenue by reducing pre-retirement income, would be expected to have two offsetting effects: 1) they could speed retirement as the benefits of continued work are lower; but 2) they could delay retirement to the extent less is saved at younger ages.

The results confirm this general intuition, although the effect of revenue increasing policies on retirement timing seems negligible. With respect to increasing the FRA to 69 or reducing the COLA, the effects are a 3.9 and 2.5 percent reduction respectively in the share completely retired at 69 (the model assumes all individuals claim at 70 at the latest). The increase in the payroll tax increases the share completely retired by just 0.3 percent at age 69 and increasing the cap has an effect of less than 0.1 percent.

The reduction in consumption also differs across policies and, more so than for the reductions in retirement timing, across people with different incomes. For example, at age 70 the average reduction in consumption under the increase in the FRA (relative to current policy) is 5.6 percent for individuals in the lowest third of the income distribution and 2.2 percent for those in the top third. The corresponding numbers for the COLA adjustment are 2.8 and 0.8 percent. However, the effect of the COLA adjustment increases with age, so that for those living to age 90 the reduction is 10.5 and 4.0 percent for the lowest and highest third. Prior to retirement, the

[1] For simplicity, the policies simulated in this paper are assumed to be implemented in one shot. In reality, the financial impact of most changes to the Social Security program assumes a more gradual phase in.

effect of benefit reductions on consumption is estimated to be relatively small. On the other hand, an increase in the payroll tax decreases consumption primarily during the working life – by between 1.3 and 1.5 percent between ages 25 and 55 for the lowest third and 1.2 and 1.4 percent for the highest third. The effect during retirement of revenue increasing policies is smaller and operates through reduced savings prior to retirement.

The next section describes the literature on how Social Security policy changes or changes in taxation affect retirement timing. The following section reviews the structural model being used in this paper and the next section the data used to estimate the model. The next section begins with a discussion of the internal and external validity of the Gustman and Steinmeier model and presents the results of the policy experiments. The final section concludes that the model estimates benefit reductions tend to have a greater behavioral impact than revenue increases as well as a greater reduction in consumption in retirement, but much of the more muted effect is due to the fact that the reductions in consumption due to tax increases are spread over a longer period of time.

Literature Review

The literature on the effects of benefit and the effect of taxation on retirement timing and consumption are discussed below.

Benefit Reductions

Increases in the FRA and their effect on retirement timing have been explored somewhat extensively in the academic literature. The most relevant work with respect to this paper is done by Gustman and Steinmeier themselves. Gustman and Steinmeier (2006) find that the effect of increasing the FRA from 65 to 67 is to increase retirement around the early eligibility age and decrease retirement at ages 65 and on. They argue that the reason for the counterintuitive increase in retirement at earlier ages is that the incentive to keep working for those that would have retired lower is lower given the benefit reduction (i.e., an additional year of delay yields lower additional income). At ages 65 and later, the effect is intuitive, with the share completely retired declining by roughly 3 percent. Gustman and Steinmeier (2009) subsequently investigate what share of the increase in labor force participation between 1992 and 2004 changes in Social Security rules over the same time period (which include an increase in the FRA) explain and find they account for about a sixth. Blau and Goodstein (2010) explore the period since the 1980s and find a larger effect of Social Security rule changes that included an increase in the FRA but also other programmatic changes. Their finding is that changes in Social Security rules could explain about a quarter to a half of the increase in labor force participation over the time period of their study. Matrobuoni (2009) investigates the retirement behavior of birth cohorts turning 62 after 2000 and investigates the effect of increases in the FRA specifically. That work suggests that a one year increase in the FRA would result in a six month increase in the average age of retirement.

The decrease in the COLA represents a less explored case, since past changes to the Social Security program have not included this kind of policy. And although a COLA reduction is essentially a real benefit cut, it is a benefit cut with odd timing – the effect on the benefit is minimal at first but increases as the distance from the time claimed increases. For example, if the COLA is reduced by 0.5 percentage points, after 10 years an individual's PIA will have been shrunk in real terms by 5 percent relative to having the COLA as it is now, 20 years later by 10 percent, and 30 years later by 17 percent. So while other authors using structural models – notably Gustman and Steinmeier (2006) and Van der Klauuw and Wolpin (2008) – find fairly large effects in terms of delayed retirement from benefit reductions on the order of 25 percent, it is unclear the results will be as strong for a COLA reduction given that individuals tend to discount the future. It is worth noting that earlier research focused on the effect of benefit *increases* between 1969 and 1972, for example Burtless (1986), and found that these increases were a minor contributor to an overall speeding up of retirement for the cohorts affected.

The effect on consumption of policies that decrease benefits has received somewhat less attention. Using their structural model, Van der Klauuw and Wolpin (2008) examine consumption effects for lower income households and estimate that a reduction in benefits of 25 percent reduces household consumption by about 4 percent, a somewhat small reduction relative to the benefit decline. Benefit increases, on the other hand, have been frequently investigated within the context of the life cycle model, but the focus has been on consumption changes within households instead of on consumption across households experiencing different benefit levels. For example, Wilcox (1986) finds that households who experienced a somewhat unexpected increase in benefits (typically informed less than one year but more than one month in advance) saw month-to-month increases in consumption once receiving their benefit, in violation of a simple formulation of the life cycle model.

Revenue Increases

The effect of tax increases generally and the payroll tax specifically on retirement timing have been explored in several prior studies, although the suggested effect on retirement timing is somewhat unclear. Using a reduced form approach, Alpert and Powell (2012) find that the elimination of the employee portion of the FICA tax would reduce retirement for men, although the magnitude of the effect was relatively small at 4 percent.[2] The structural model of Keane and Wolpin (2008) agrees with the direction of the effect but predicts a larger magnitude – it predicts a significant and large decrease in labor force participation following an increase in the employee portion of the tax to 15 percent (a much larger increase than examined here). These findings are consistent with the substitution effect of taxes dominating – if the tax increases (decreases) work will decrease (increase) because its value has been lowered (increased). However, this finding is not uniform. Gurley-Calvez and Hill (2011) find that increases in the state income tax actually decrease retirement, suggesting that at least with respect to state taxes the income effect may dominate.

Model

The structural model used in this paper is described in detail in Gustman and Steinmeier (2006), but a review of it with information on how Social Security benefit reductions and tax increases are ultimately incorporated is useful. To review the basics, the model takes a life-cycle approach and focuses on the retirement behavior of males who begin their time in the *Health and Retirement Study* (HRS) as a part of a married couple. These individuals are assumed to maximize their expected lifetime utility subject to a budget constraint incorporating income from labor, pension and Social Security income, and the returns from other assets. Individuals face uncertainty over the rate of return on assets, their utility from continued work, and over their mortality.[3] The paths of an individual's potential wages (realized if they work), health, and the

[2] Schmidt and Sevak (2009) come to a similar conclusion using a reduced form approach, albeit with a slightly larger effect.
[3] Individuals know their mortality rates, but not when they will die.

formulation of the Social Security benefit are not stochastic and are assumed to be known by the individual.[4] Given their constraints, workers are assumed to make choices about their labor supply and about their consumption. Lifetime utility is given by the following equation:

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

In equation (1), the individual's marital status is indicated by m, which can be married with: 1) both spouses alive; 2) only the married male surviving; and 3) only the spouse surviving.[5] The vector $s_{m,t}$ represents the likelihood of those situations arising by time t and reflects mortality risk. The vector h_t indicates the individual's leisure preference, which is allowed to vary as a function of a person's health and as a function of their age. $L_{m,t}$, represents the labor/leisure decision which can be working ($L_{m,t}$= 0), partially retired ($L_{m,t}$= 0.5), or fully retired ($L_{m,t}$= 1.0). The vector h_t is also where health enters the model. Specifically, the model specifies that the leisure term is:

$$h_t = e^{\beta X_t + \varepsilon_t} \tag{2}$$

In equation (2), X_t is a vector that contains a constant, the individual's age, and whether they are in poor health.[6] The term ε_t reflects the aforementioned uncertainty over the utility from leisure.[7] A positive coefficient on poor health would indicate that it makes leisure more attractive than work. The parameter γ represents the individual's preference for partial retirement.[8] Finally, the parameters ρ and α capture an individual's time preference and the shape of the utility function regarding consumption, respectively.

Households make their leisure and consumption decisions under the following budget constraint:

---

[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. Gustman and Steinmeier (2015) incorporated uncertainty over the future of Social Security and found it did improve the models ability to fit retirement and claiming decisions.

[5] The individual being modeled is assumed to not get any utility once they die, regardless of the spouse's survival status.

[6] In the model, once people enter poor health they are assumed to remain in poor health.

[7] This uncertainty initially comes from a mean zero distribution and holds a single value until retirement from the career job. After retirement from this career job, ε_t, is allowed to vary over time, with a year-to-year correlation parameter of ρ_ε.

[8] The preference for partial retirement has an individual-specific component, δ, and a "shifter" that potentially increases that preference with age, δ_a.

$$A_t = (1 + r_t)A_{t-1} + W_t\big(1 - L_{m,t}\big) + E_{m,t} + B_{m,t} - C_{m,t} \qquad (3)$$

Equation (3) shows how the assets available today (and ultimately in the future) depend on: 1) the assets available in the last period (A_{t-1}), and the returns on those assets (r_t);[9] 2) wages (W_t), which are allowed to be lower once an individual has quit the career job they entered the model with; 3) income from the spouse which includes both wage and pension income ($E_{m,t}$);[10] and 4) income from the Social Security and from the individual's own pension ($B_{m,t}$).[11]

Equations (1) and (3) paint an intuitive picture of how individuals make decisions within this model: individual make decisions about retirement and consumption today knowing those decisions will alter their utility, level of income, and amount of assets both today and in the future. For example, continuing work today brings some amount of disutility in the moment, but allows career wage income to continue, potentially increases the future Social Security benefit, and delays claiming one more year so that the benefit increases (if the individual is over the Early Eligibility Age). Individuals who continue working also know that they are able to continue accumulating assets and potentially pension wealth. So the decision affects utility today through the disutility from work and the availability of wage income and affects it tomorrow through the potential increase in Social Security wealth, pension wealth, and savings. When the disutility from work overcomes those current and future benefits, people retire. Before discussing the data used to estimate this model, it is worth delving deeper into the two aspects of the model that will be changed in this paper to simulate policy, Social Security benefits and taxation.

Social Security Benefits

The modeling of Social Security benefits, contained in $B_{m,t}$ in equation (2), is straightforward. Given that the future potential income (if the individual works) from labor is known with certainty, the benefit from claiming Social Security at any age is also known in

[9] The current paper uses returns calculated by Ibbotson Associates (2002) after taking into account the distribution of financial assets in the HRS, as in Gustman and Steinmeier (2006).

[10] The spouse's wage income and pensions are treated as exogenous, as is their retirement decision.

[11] 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" and by the rules of their pension plan. If those rules leave little to gain for working past certain ages or large gains of working to certain ages, they may have a large impact on retirement timing.

advance. Individuals are assumed to claim Social Security as soon as they are eligible and their earnings fall under the earnings test.[12] At this point, they receive their PIA reduced by any actuarial adjustment (or increased by the delayed retirement credit) as determined by their birth cohort and the age at which they retired.[13] These assumptions mean that Social Security benefits in the model tend to keep people in the workforce until age 62, and that delaying retirement has the benefit of increasing post-retirement income. Increasing the FRA or decreasing the COLA can be incorporated by altering the benefit formula.

Payroll Taxation

Taxes are not explicitly included in the Gustman and Steinmeier (2006) model (i.e., the budget constraint is assumed to reflect post-tax wages, assets, and consumption). To simulate changes in the payroll tax, the income from work will be reduced by the appropriate amount given the rate assumed and the assumed cap.

Data and Estimation

This section first describes the data gathered to estimate the model described above and then describes how estimation is conducted.

Data

The data used in this analysis come from the HRS and specifically the original HRS Cohort waves 1 to 6, as was used in Gustman and Steinmeier (2006). As in Gustman and Steinmeier (2006), the sample begins with married males and then eliminates individuals by introducing a number of restrictions, as listed in Table 1. 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. The most important restrictions are that individuals with retirement plans but

[12] This assumption is based on evidence from Coile et al. (2002) that people typically claim Social Security as soon as possible, although later version of the Gustman and Steinmeier allow claiming to be investigated separately.
[13] 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.

without data on those pensions are dropped as are individuals with unstable employment histories.[14]

Perhaps the most important individual variable for estimation of the model is an individual's labor participation status. In the data, individuals working at least 30 hours each week and 1,560 hours a year are counted as full-time workers; between 100 hours and 1,250 hours a year and who do not exceed 25 hours per week are labeled as part-time and partially retired; and individuals working less than 100 hours per year fully retired. The paper will focus on the time at which people completely retire from work, so Figure 1 shows how the share of workers who are completely retired changes with age. The most notable feature of the data is the large increase in the share completely retired increases substantially at the Early Eligibility Age of 62.

Data to fill in the variables from equation (3) come from a variety of sources. For those HRS sample members that are linked to Social Security earnings records, those records are used. For the roughly a quarter of respondents remaining, earnings histories are constructed from self-report on the current job or any previous long-term job lasting more than 5 years. Because behavior is modeled for most individuals beyond their time in the data – not all HRS Cohort members were 69 years old in 2004 when the data used ends – earnings must be projected into the future. This projection is done following Gustman and Steinmeier (2006) using an individual's experience and tenure to conduct a regression and form predictions. The individual's Social Security Primary Insurance Amount needed to calculate their benefit is calculated based on this earnings history and updated each year the individual works. Their benefit is determined by the rules when they claim, which (as mentioned above) is taken as the first time when eligible to claim and their earnings falls below the earnings test.

Information on DB pensions 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. The final source of data needed to complete equation (3) is on any wage income, pension income, or Social Security income coming from spouses. Generally, these quantities are

[14] The justification for these restrictions is discussed at length in Gustman and Steinmeier (2006).

calculated in the same way as they are for males, with the exception that they are treated as exogenous (unlike the male's earnings, which can be affected by the decision to work).

Aside from the resource constraint, information on an individual's health is also required to determine the utility from continued work as modeled by equation (2). A person's health is treated as binary variable equal to 1 if the individual self-reports being in poor health and 0 otherwise.

Estimation

The goal of estimating the model is to obtain the preference parameters in equations (1) and (2). Once these quantities are estimated, the model can be used to simulate individual behavior and ultimately conduct the four policy experiments laid out in the introduction. The parameters include the consumption parameter from equation (1), the distribution of individual preferences for partial retirement also from equation (1), the distribution and year-to-year correlation of uncertainty in a person's leisure preference from equation (2), and the effect of age and health on the preference for leisure, also from equation (2).

As is standard in dynamic stochastic structural models like that used in this paper, the model is solved recursively and then estimated iteratively. 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 any combination of earnings, assets, and career versus non-career job they could arrive at that time with. Because it is the final period of the model, the individual does not need to consider the future and so their decision is relatively easy – they consume everything and work if the extra income more than balances out the disutility of labor.

At the period T-1, the decision becomes harder because a future period must be incorporated into the individual's decision. Again, for a given set of parameters, the current utilities are calculated at each possible model outcome. Although one choice will yield the highest utility at time T-1, the individual must also consider that the choice impacts the assets available tomorrow and thus their expected utility in the future. Fortunately, given the individual is assumed to know the distribution of the rate of return, the individual can calculate the probability they end up at any particular asset level tomorrow given their choice of consumption today. And because they have solved the model already at time T at each asset level, they can determine their expected utility from their work given they are assumed to know the distribution of uncertainty in their preference for leisure. This knowledge allows them to make a labor and consumption choice today that maximizes not just their utility today, but their expected utility from this T-1 forward. Work and consumption decisions are modeled in this way going back to the first period of the model.

Once this iterative process is complete all the way back to age 51, each individual will have a simulated work history. The goal of the estimation routine is to closely match these work histories to reality. To accomplish this goal, the model uses the Generalized Method of Moments (GMM). The GMM methodology is relatively straightforward and compares moments from the actual data to the same moments from the simulated data. One example of such a moment is the share of people in poor health at any given age that is completely retired. If the parameters that the model was solved on are close to the true parameters, then the difference between the simulated and actual moments will be small. The estimation process proceeds iteratively – if the model predicts that people who are sick work far more often than they actually do, then the next time the parameters guessed will reflect a larger penalty for working while in poor health. In this iterative way, the model arrives at parameters for which the model fits the chosen moments.[15]

Results

Before discussing the results of the policy simulations, it is useful to first do two preliminary exercises to ensure those policy simulations are meaningful. The first exercise is to show that the model at least fits the data it was estimated on – i.e., to internally validate the model. Given that policy simulations by definition are "out of sample" exercises, to believe them it is important that the model at least yield good "in sample" predictions. On this point, the Gustman and Steinmeier (2006) model performs fairly well with respect to the main behavioral outcome of interest in this study, complete retirement. Figure 2 shows that up until age 65, the predicted shares completely retired trend very closely to the actual share. After age 65, the model slightly under predicts the share of people completely retired.

[15] Results can be found in the Appendix. 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 to 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.

The second exercise is to attempt to "externally validate" the model. Because the goal of this paper is to predict how people are likely to respond to policy changes, it is useful to know if the model can accurately predict people's responses to changes that have already occurred. An analysis in Hou et al. (2016) using the Gustman and Steinmeier (2006) model yields an encouraging result. In that paper, the authors looked at how much of the change in behavior between the HRS Cohort and the Early Baby Boomer Cohort (born roughly 12 years later) could be explained by four major changes that played out over that time: 1) changes in Social Security rules; 2) changes in pension structure; 3) improvements in health; and 4) improvements in mortality. To accomplish this comparison, the authors re-simulated the model and replaced the actual characteristics of the HRS Cohort along those four dimensions with simulated characteristics based on the Early Baby Boomer Cohort. For example, many workers in the original HRS Cohort with DB pensions were assigned DC pensions with the same pension wealth, but without the retirement rules that correspond to DB plans. Although these four changes do not represent all changes in the economy over those four years, one would expect they explain a good portion of any change in behavior. Figure 3 suggests that these changes do explain much of the decrease in retirement seen between the two cohorts and bodes well for the counterfactual experiments discussed below.

Reductions in Benefits

The simulations begin with today's baseline policy – an FRA of 67, a Delayed Retirement Credit that spans to 70, and the elimination of the earnings test over age 67.[16] To simulate the FRA increasing to 69 the benefit reduction of retiring prior to 67 was increased to reflect two additional years, a benefit reduction consistent with two years of early claiming was incorporated at 67, a benefit reduction consistent with one year of early claiming was incorporated at 68, and only one year of a Delayed Retirement Credit was allowed. To simulate the effect of a COLA reduction by 0.5 percent, the individual's real benefit was divided by $1.005^{(a-c)}$ where *a* is their current age and *c* is the age at which they claimed. This means that in

[16] This "baseline" policy is not the original policy from the Gustman and Steinmeier model, which used the Social Security policy relevant for each individual given their birth cohort. However, since the policy changes being discussed build off today's policy, a slightly different approach is needed.

this simulation delaying retirement has a small side benefit – it delays the time in which the benefit can basically be deflated.

The results of the experiments are shown in Table 2. The effect of increasing the FRA from 67 to 69 is similar in magnitude to effect of increasing it to 67 found in Gustman and Steinmeier (2006). Between the ages of 62 and 69 the reduction can be as small as 0.8 percent or as large as 4.2 percent. The effect of the COLA decrease is generally smaller but also more consistent over the age range, with decreases of between 2.0 and 2.5 percent in the share completely retired. In general, it seems either policy will delay retirement. It is worth noting these behavioral effects are slightly larger for low income individuals, as Table 3 shows. For example, for the FRA increase, the reduction in the share completely retired at ages 62 to 69 for individuals in the bottom third of the income distribution ranges from 4.2 to 7.5 percent and for individuals in the top third from an *increase* of 1.7 percent to a decrease of 2.4.[17]

Figure 4 shows the percent reduction in consumption during the working years and then in retirement by income tercile for an increase in the FRA and Figure 5 for a decrease in the COLA. The figure shows that both policies reduce consumption slightly as workers approach the Early Eligibility Age and the need to save more to offset benefit reductions becomes less distant and so less discounted. However, the size of the consumption reduction differs considerably across the two policies. Increasing the FRA to 69 results in a larger decline in consumption early in the retirement period while the COLA results in a larger reduction later in retirement. For example, Figure 4 shows that for low income individuals, the reduction in consumption is about 2.0 percent under an FRA of 69 and 0.5 percent under the COLA reduction. At age 90, those relative magnitudes are reversed to 7.2 percent and 10.5 percent since the COLA reduction results in a larger decrease in benefits with age.[18] Both policies decrease the consumption of lower income individuals by a higher amount than higher income individuals, who are less reliant on Social Security.

[17] See the literature review for an explanation of how increases in the share retired could occur under an FRA increase, as Gustman and Steinmeier (2006) reports for some ages.

[18] The effect of both policies increases with age as individuals spend down their other wealth and rely more and more on Social Security income.

Increase in Taxes

Table 4 shows the share completely retired under the two policies and indicates essentially no effect on retirement timing. This result stands in contrast to that Van der Klauuw and Wolpin (2008) who found a larger effect, although they also looked at a much larger tax increase. In any case, neither the payroll tax increase to 7.75 percent nor the raising of the cap is expected to have a large effect on retirement timing according to the Gustman and Steinmeier model. With respect to changes in the timing of retirement, little difference exists between the income groups.

The effect on consumption of a payroll tax increase is shown in Figure 6. For simplicity, only the average decline is shown since the reduction is similar across income groups (if slightly smaller for high income individuals who rely less on labor income during their working lives). The main takeaway is that an increase in the payroll tax from 6.2 to 7.75 percent has the predictable effect of decreasing consumption during the working life of about 1.5 percent with the effect dissipating as people retire and cease relying on wage earnings. Still, a small residual effect remains even once all workers are assumed retired, since less was saved during the working life. Finally, Figure 7 shows the reduction in consumption under a payroll cap increase for individuals in the highest income tercile (the only ones affected by such a policy in the model). The result shows a similar, but muted, pattern to the payroll tax increase.

Conclusion

As policymakers consider ways to lessen Social Security's financial shortfall, a variety of policy changes are possible that could have similar effects on the program's finances. This paper shows the differing effects on recipients of some of those policies. The results suggest that benefit reductions tend to have a larger behavioral effect than tax increases that yield the same effect on program finances, likely because the decrease in retirement income is larger albeit concentrated over a smaller time period. The reduction in consumption under benefit reductions is substantially larger during retirement than under a payroll tax or cap increase, but those policies decrease consumption more during the working lives of individuals which are considerably longer than retirement. Therefore, this research suggests that one tradeoff policymakers face is between a long period of lower consumption pre-retirement versus a shorter but more pronounced period of lower consumption post-retirement.

References

Alpert, Abby and David Powell. 2012. "Tax Elasticity of Labor Earnings for Older Individuals." Working Paper 2012-272. Ann Arbor, MI: University of Michigan Retirement Research Center.

Blau, David M. and Ryan M. Goodstein. 2010. "Can Social Security Explain Trends in Labor Force Participation of Older Men in the United States." *Journal of Human Resources* 45(2): 328-363.

Burtless, Gary. 1986. "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement." 1986. *Review of Economic Studies* 53(5): 781-805.

Coile, Courtney, Peter Diamond, Jonathan Gruver, and Alain Jousten. 2002. "Delays in Claiming Social Security Benefits." *Journal of Public Economics* 84(3):357-385.

Gurley-Calvez, Tami and Brian Hill. 2011. "Time to Retire? The Effect of State Fiscal Policies on Retirement Decisions." *American Economic Review: Papers & Proceedings* 101(3): 35-39.

Gustman, Alan L. and Thomas L. Steinmeier. 2006. "Social Security and Retirement Dynamics." Working Paper 2006-121. Ann Arbor, MI: University of Michigan Retirement Research Center.

Gustman, Alan L. and Thomas Steinmeier. 2009. "How Changes in Social Security Affect Recent Retirement Trends." *Research on Aging* 31(2): 261-290.

Gustman, Alan L. and Thomas Steinmeier. 2015. "Effects of Social Security Policies on Benefit Claiming, Retirement, and Saving." *Journal of Public Economics* 129(9): 51-62.

Hou, Wenliang, Alicia H. Munnell, Geoffrey T. Sanzenbacher, and Yinji Li. "Why are U.S. Households Claiming Social Security Later?" Working Paper 2017-3. Chestnut Hill, MA: Center for Retirement Research at Boston College.

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

Schmidt, Lucie and Purvi Sevak. 2009. "Taxes, Wages, and the Labor Supply of Older Americans." *Research on Aging* 31(2): 207-32.

Van der Klaauw, Wilbert and Kenneth I. Wolpin. 2008. "Social Security and the Retirement and Saving Behavior of Low-Income Households." *Journal of Econometrics* 145(1): 21-42.

Table 1. *Sample Size Following Restrictions Imposed by Gustman and Steinmeier (2006)*

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

Source: Authors' calculation from *Health and Retirement Study* (HRS) and Gustman and Steinmeier (2006).

Table 2. *Share Completely Retired Under Various Policies*

Age	Share completely retired			Percent change from baseline	
	Baseline	FRA 69	COLA reduction	FRA 69	COLA reduction
62	42.6%	40.8%	41.7%	-4.2%	-2.0%
63	46.8%	45.0%	45.9%	-3.9%	-2.1%
64	49.8%	49.7%	48.8%	-0.3%	-2.0%
65	55.0%	54.5%	53.8%	-0.8%	-2.0%
66	59.5%	56.9%	58.3%	-4.4%	-2.1%
67	63.0%	61.5%	61.4%	-2.3%	-2.5%
68	66.6%	65.3%	65.0%	-2.0%	-2.4%
69	69.2%	66.5%	67.5%	-3.9%	-2.4%

Source: Authors' calculation from *Health and Retirement Study* (HRS) and Gustman and Steinmeier (2006).

Table 3. *Percent Change from Baseline in Share Completely Retired under Various Policies, by Income Tercile*

Age	FRA69			COLA Reduction		
	Low	Middle	High	Low	Middle	High
62	-7.1%	-4.1%	-2.4%	-3.5%	-2.2%	-1.0%
63	-6.4%	-3.7%	-2.3%	-3.6%	-2.3%	-1.1%
64	-4.2%	0.7%	1.7%	-3.6%	-2.0%	-0.9%
65	-4.5%	0.1%	1.2%	-3.6%	-2.1%	-0.9%
66	-7.0%	-4.3%	-2.4%	-3.6%	-2.2%	-1.0%
67	-5.0%	-1.7%	-0.9%	-3.7%	-2.7%	-1.5%
68	-4.8%	-1.4%	-0.3%	-3.6%	-2.6%	-1.5%
69	-5.5%	-4.1%	-2.4%	-3.5%	-2.5%	-1.5%

Source: Authors' calculation from *Health and Retirement Study* (HRS) and Gustman and Steinmeier (2006).

Table 4. *Percent Change from Baseline in Share Completely Retired under Various Policies, by Income Tercile*

Age	Share completely retired			Percent change from baseline	
	Baseline	Payroll tax increase	Payroll cap increase	Payroll tax increase	Payroll cap increase
62	42.6%	42.6%	42.6%	0.1%	<0.1%
63	46.8%	46.9%	46.8%	0.1%	<0.1
64	49.8%	49.9%	49.8%	0.2%	<0.1
65	55.0%	55.1%	55.0%	0.2%	<0.1
66	59.5%	59.7%	59.5%	0.2%	<0.1
67	63.0%	63.2%	63.0%	0.3%	<0.1
68	66.6%	66.8%	66.6%	0.3%	<0.1
69	69.2%	69.3%	69.2%	0.3%	<0.1

Source: Authors' calculation from *Health and Retirement Study* (HRS) and Gustman and Steinmeier (2006).

Figure 1. *Share of Sample Completely Retired Ages 50-69*

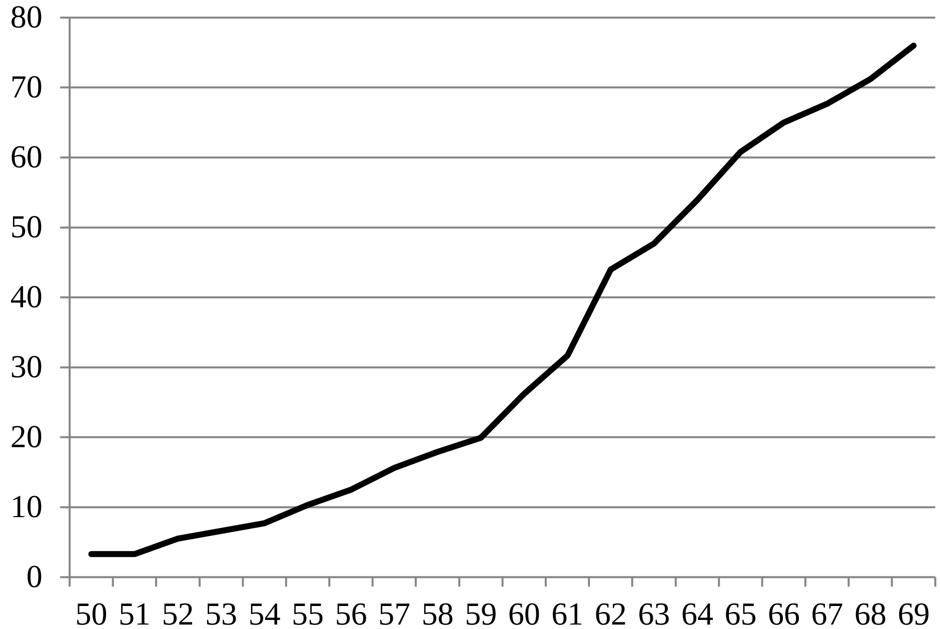

Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

Figure 2. *Actual and Predicted Share Completely Retired Ages 50-69*

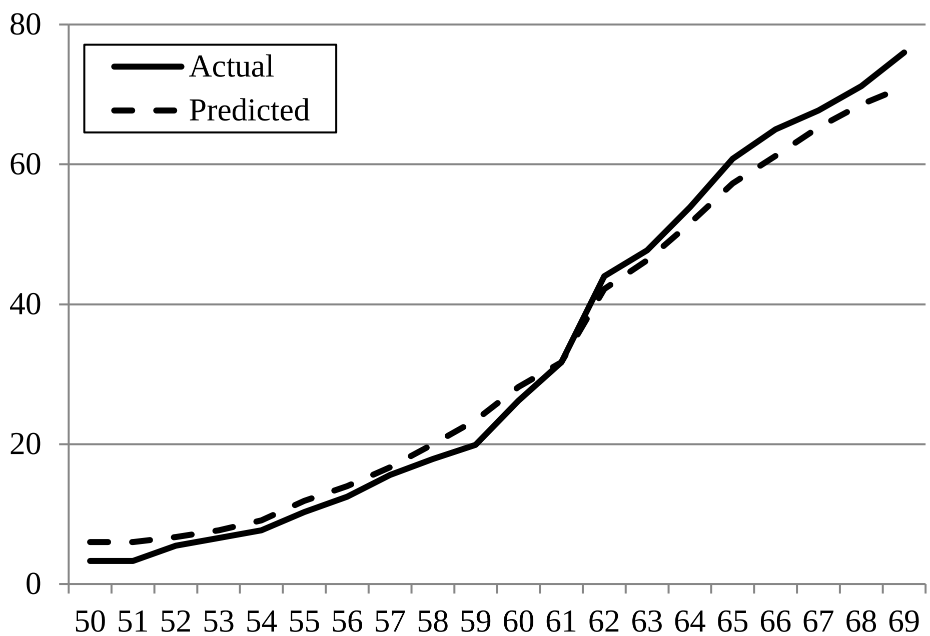

Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

Figure 3. *Share of Sample Completely Retired Age 50-64, by Cohort*

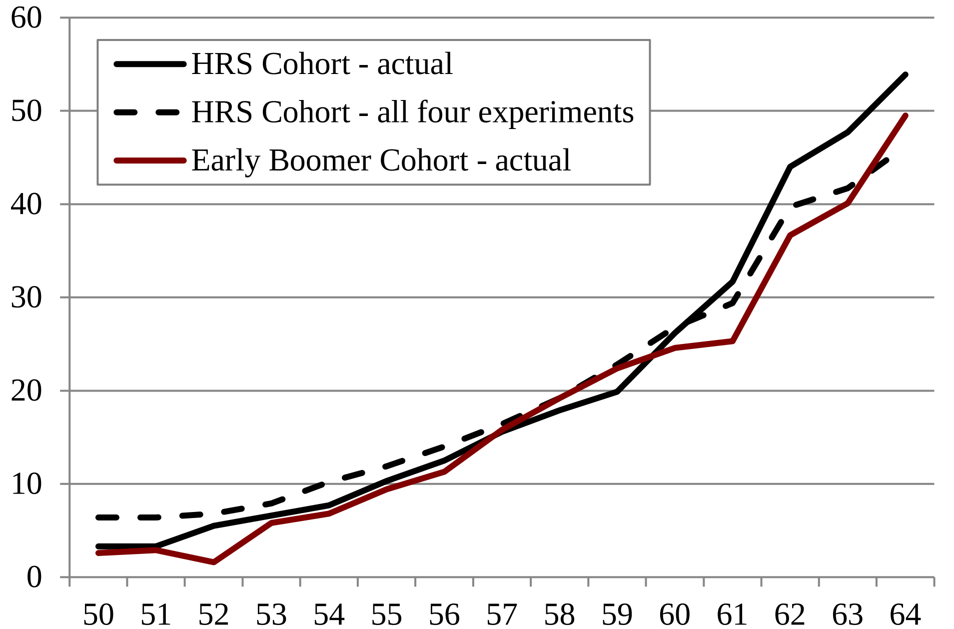

Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

Figure 4. *Percent Reduction in Consumption with FRA Increase Age 25-90, by Income Tercile*

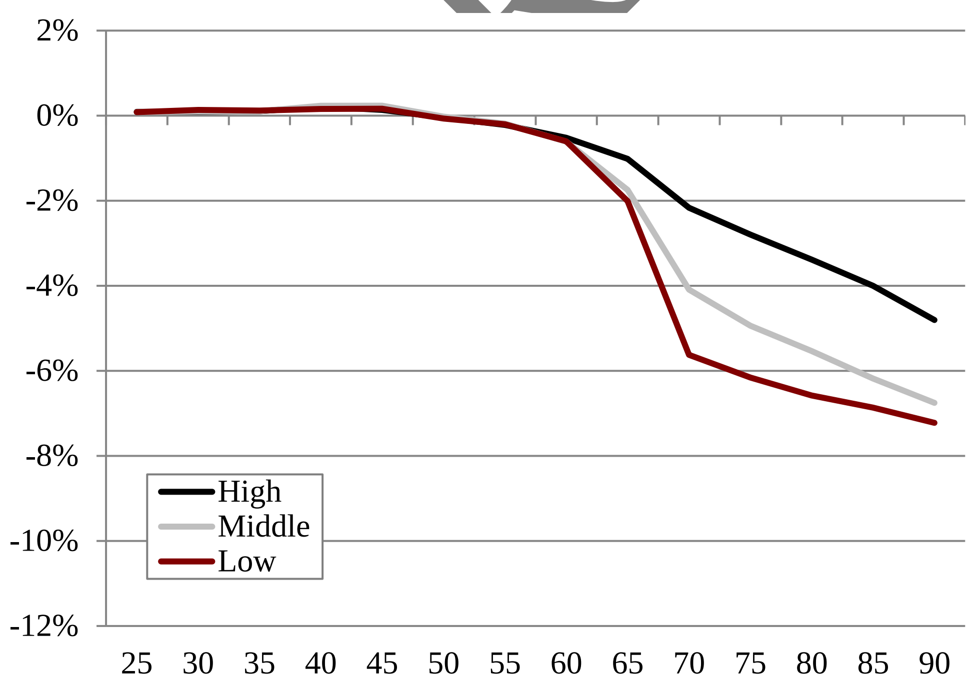

Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

Figure 5. *Percent Reduction in Consumption with COLA Decrease Age 25-90, by Income Tercile*

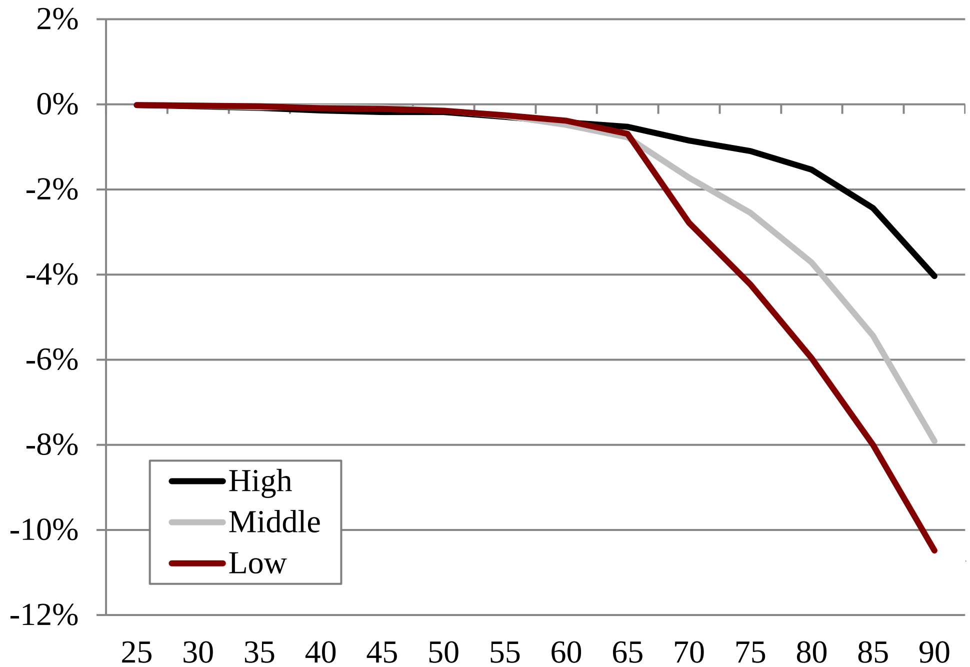

Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

Figure 6. Average *Percent Reduction in Consumption under Payroll Tax Increase Age 25-90*

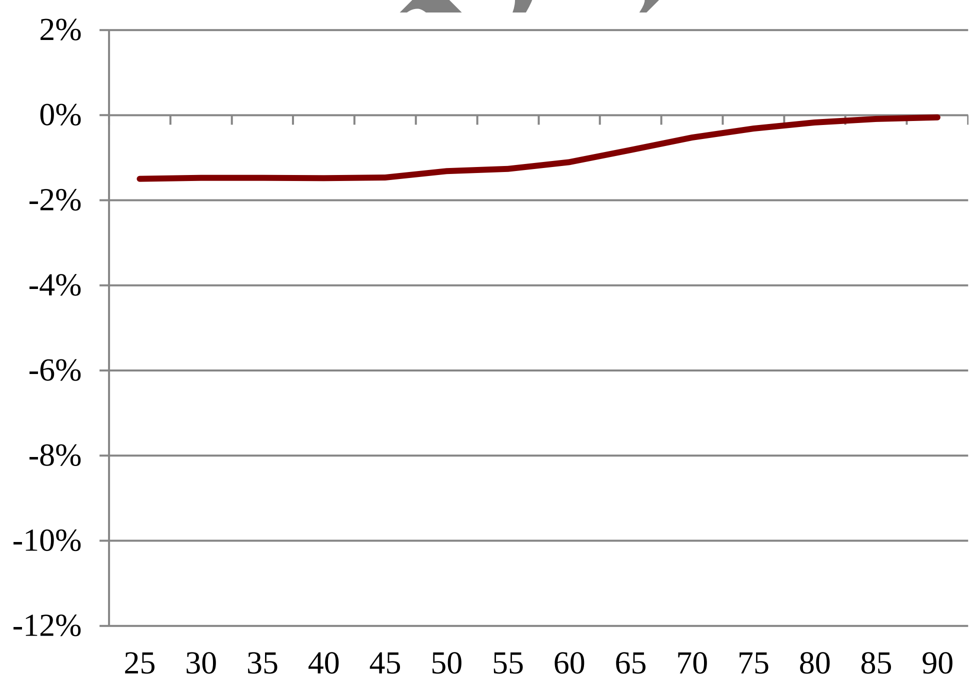

Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

Figure 7. Average *Percent Reduction in Consumption under Payroll Cap Increase Age 25-90, Highest Income Tercile Individuals*

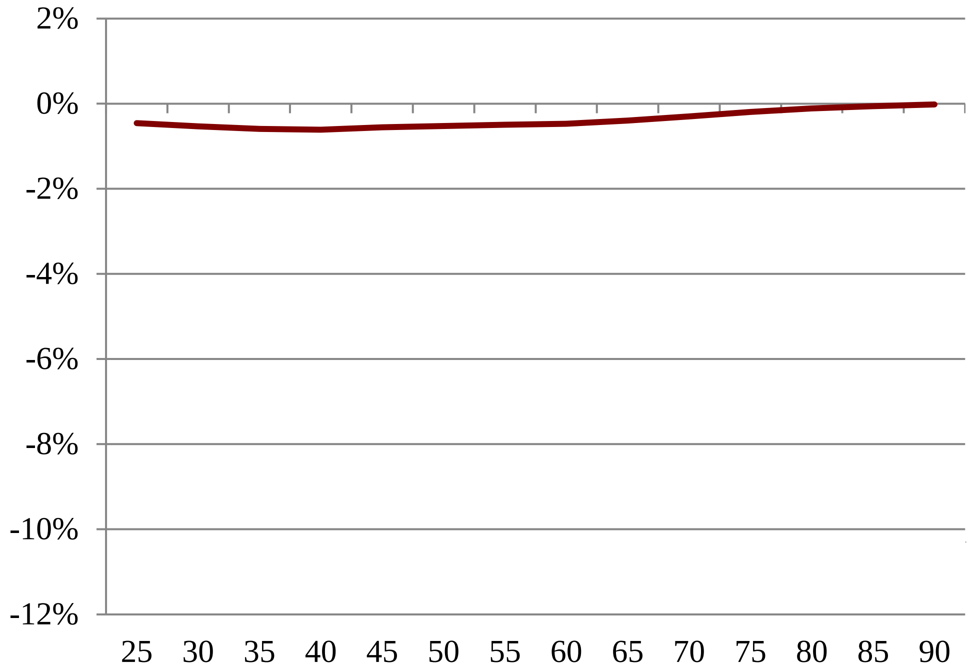

Source: Authors' calculations HRS and Gustman and Steinmeier (2006).

Appendix

Table A1. *Parameter Estimates*

Symbol	Description	Coefficient Value	t-statistic
α	Consumption parameter	-0.24	2.36
	Parameters in β		
$\beta 0$	Constant	-9.82	51.46
βa	Coefficient of Age	0.659	3.68
βh	Coefficient of Health	6.78	7.82
$\rho\varepsilon$	Correlation of ε after retirement	0.88	38.54
	Parameters in δ		
$\delta 0$	Constant	-5.34	2.12
δa	Coefficient of Ageb	-0.48	1.14
$\sigma\varepsilon$	Standard Deviation of ε	5.40	25.58
	q value:	53.32	
	Number of observations:	2,231	

Source: Authors' calculations from *Health and Retirement Study* and Gustman and Steinmeier (2006).