

Do Changes in Pension Incentives Affect Retirement?
A Longitudinal Study of Subjective Retirement Expectations

February 2001

Sewin Chan

Robert F. Wagner School of Public Service
New York University
sewin.chan@nyu.edu

Ann Huff Stevens

Department of Economics
Yale University
ann.stevens@yale.edu

The authors thank Stephen Haider, Ron Miller and seminar participants at the University of California, Davis, University of Connecticut, Hunter College, University of Maryland, New York University and University of Toronto for their comments and discussions, and gratefully acknowledge financial support from the National Science Foundation, grants 9905275 and 9907824. Excellent research assistance was provided by Min Qi and Matthew Johnson. A previous version of this paper was circulated under the title “Retirement Incentives and Expectations.”

Do Changes in Pension Incentives Affect Retirement?

A Longitudinal Study of Subjective Retirement Expectations

Abstract

This paper investigates the responsiveness of individuals' retirement decisions to forward-looking measures of pension accumulations. In contrast to previous research, we are able to use within-person variation in retirement incentives and control for unobserved heterogeneity in tastes for retirement that may be correlated with retirement incentives, by studying a panel of subjective retirement expectations. We find that the response to pension incentives estimated from fixed-effects regressions is at most half the magnitude of that estimated from similar OLS regressions. Our results suggest potentially large biases from the failure to control for unobserved heterogeneity in previous estimates of such relationships.

JEL classification: J2

I. Introduction

Understanding the determinants of retirement from the labor force is crucial for designing retirement programs and public policies that affect older individuals. While there has been much research in this area, virtually all of the existing empirical studies that use longitudinal data have relied heavily on cross-sectional variation in pension incentives, and have paid little attention to the possibility of unobserved heterogeneity in tastes for retirement that are correlated with both retirement behavior and pension incentives. This paper makes use of repeated observations of individuals' retirement expectations to investigate the effects of retirement incentives on those expectations, and to evaluate the robustness of the results to controls for unobserved heterogeneity. We find evidence of substantial biases in estimates of the effect of forward-looking retirement incentives when there are no controls for unobserved heterogeneity. Thus, previous research may have substantially overstated the effects of retirement incentives on retirement behavior.

The retirement problem does not lend itself directly to approaches based on within-person variation because actual retirement (defined as a permanent exit from the labor force) will be observed only once. Thus, repeated observations of actual behavior will provide limited information about responses to *changes* in retirement incentives and may be subject to bias from unobserved heterogeneity in tastes for retirement. Factors that determine the nature of pension plans and Social Security benefits, such as employer choice and earnings history, are very likely to be correlated with labor force attachment, leading to an overstatement of the retirement response to incentives.

We are able to estimate the effect of incentives on retirement using within-person variation by examining a different dependent variable: individuals' subjective probabilities of

continuing work beyond age 62, or beyond age 65, from a survey question in the Health and Retirement Study (HRS). The advantage of this measure is that we observe these subjective probabilities at up to three different dates. Thus, we can control for any fixed individual characteristics and learn much more about the updating of retirement plans in response to new information.

Most longitudinal studies of the effect of pension incentives on retirement behavior use cross-sectional variation as the main source of identification and ignore the likely correlation between pension incentives and tastes for retirement. Even studies that are supposedly based on time-varying incentives are often identified from the cross-section since there is little unanticipated within-person variation over time.

One exception is Krueger and Pischke (1992) who investigated the labor supply effect of Social Security using cohort analysis within an aggregate panel data set. Their source of within-person variation comes from a natural experiment: a substantial and unanticipated decline in Social Security benefits for those born between 1917 and 1921 (the so-called notch generation), compared with otherwise identical individuals who were born before 1917. They find an insignificant relationship between Social Security wealth and labor force participation for the notch generation and their results suggest that previous measures of a large negative Social Security wealth effect were overestimated. Although they were not able to investigate the role of private pensions or other sources of wealth that affect the retirement decision, their general conclusion is consistent with findings in this paper.

Using more recent longitudinal data from the HRS, Coile and Gruber (2000) have tried to overcome the reliance on cross-sectional variation in retirement incentives by exploiting nonlinearities in the Social Security program rules as their main source of variation in Social

Security wealth accumulation. This, they argue, isolates a source of variation in retirement income accumulation that is less subject to bias from omitted unobserved characteristics. However, they do not consider the analogous potential correlation arising from private pensions. If earnings (and thus Social Security benefits) are correlated with tastes for leisure or retirement, it also seems likely that private pension structures will be similarly correlated. Individuals with a preference for early retirement are more likely to seek out and remain with employers offering generous plans for early retirement. Thus, their approach is still potentially subject to bias from unobserved heterogeneity in tastes for retirement that is correlated with retirement incentives.

Since we are not using a natural experiment in our analysis, it is important to consider what might be driving within-person variation in the incentives to retire. Among workers who are employed by the same firm throughout their 50s and 60s, the rules of their pension eligibility and benefits should be known to them from the very beginning of their job tenure. This assumption is the basis of the optimizing behavior inherent in forward-looking models of retirement. This also suggests, however, that unless new information is revealed from one year to the next, within-person variation in expectations will not be informative. Below we describe possible sources of within-person variation, including information acquisition, early retirement incentives and job separations, and we show that our main results are robust to a sample restriction to individuals who have experienced an identifiable and arguably exogenous change in retirement incentives.

This paper contributes to our understanding of the relationship between retirement incentives and retirement by providing a crucial cross check of estimates based on studying actual retirement. We demonstrate that ignoring the likely correlation between retirement incentives and tastes for retirement leads to a substantial overstatement of the effects of pension

incentives on retirement decisions. Since we are using a different dependent variable – retirement expectations rather than actual retirement – our findings will not be directly comparable to the previous literature. However, our evidence on the sensitivity of the incentive estimate to unobserved heterogeneity will have important implications for any work that seeks to understand the timing of retirement.

The next section of the paper discusses in more detail the relationship between our approach and that taken in the recent literature on retirement decisions. Section III presents our data and econometric strategy. Section IV discusses the sources of within-person variation that is driving our estimates. Section V presents our results and the final section concludes.

II. Relation to Previous Literature

The literature on the determinants of retirement has emphasized the importance of future retirement income accumulation that comes from Social Security benefit formulas and pension structures. The behavioral assumption that underlies the empirical models is that individuals decide whether to retire based on an evaluation of lifetime utility associated with current and future retirement dates.¹ The approach that we take in this paper also adopts this fundamental idea. Our estimating framework is most directly related to the reduced form implementations of the Stock and Wise (1990) option value model of retirement, first discussed in Lumsdaine, Stock and Wise (1992). We briefly summarize their methodology and subsequent applications here.

¹Estimated structural models of retirement, including Rust (1989), Gustman and Steinmeier (1986), Stock and Wise (1990) and Rust and Phelan (1997), all share this same basic feature.

In Stock and Wise (1990), individuals chose their retirement date, R , to maximize a lifetime utility function V , of the following form:

$$(1) \quad V_t(R) = \sum_{s=t}^{R-1} \beta^{s-t} U_W(Y_s) + \sum_{s=R}^S \beta^{s-t} U_R(B_s(R))$$

where β is a subjective discount factor, U_W is utility during working years, U_R is utility while retired, Y is income while working, B is retirement benefits or income while retired, and S is the final period. Individuals choose their retirement date to maximize the present value of utility from the streams of income while working and after retirement. Stock and Wise implement the optimization undertaken here with a simplification of the full dynamic programming rule: in every period individuals compare the lifetime utility of retiring today versus retirement at a future “optimal” date, the retirement date at which lifetime utility is maximized. This difference between the utility from immediate retirement and delayed optimal retirement is referred to as the option value of work or delayed retirement. Reduced form implementations of the model involve calculating the option value variable using some assumed utility function parameter values and estimating its effect on retirement.

Samwick (1998) applies this methodology using detailed pension plan information from the Survey of Consumer Finances and Pension Provider Surveys. His calculated option value closely follows the framework of the Stock and Wise model, and is thus a non-linear function of earnings, Social Security benefit accruals, and private pension benefit accruals. The option value variable is found to be highly significant, as are measures that focus on the accrual of wealth from working just one additional year (as opposed to the option value measures which consider the gain from working until some future optimal retirement date). Levels of pension and other wealth are not found to be major determinants of retirement. Samwick’s contribution is

important since it uses a dataset from a nationally based sample, while Stock and Wise's results were based on data from workers at a single firm.

Coile and Gruber (2000) raise a concern about using option value and similar measures in a reduced-form framework. They argue that a problem with the reduced-form approaches (and to some extent the structural models as well) is a "lack of careful attention to the sources of identification of the retirement incentive effects they estimate." (Page 9). Coile and Gruber focus their criticism on the fact that the option value combines information on private pensions, Social Security benefits, and wages. Most importantly, they note that option value depends critically on an individual's wage, both because the wage enters directly as the main part of income while working, and because earnings are an important determinant of Social Security benefits. Their primary concern is that wages are likely to be correlated with underlying tastes for retirement, and so using cross-sectional variation in option value may result in biased estimates of the effects of changes in retirement incentives on retirement.

Coile and Gruber's alternative is to develop measures that capture the forward-looking accumulation of retirement wealth or income inherent in the option value, but to focus on individual variables that separately measure accumulation of Social Security benefits and accumulation of pension wealth, after controlling for a host of wage and earnings measures. Rather than using a utility-based measure such as the Stock and Wise option value, they calculate a measure they call "peak value" that is equal to the difference between total discounted Social Security wealth at its maximum expected value and its value if retirement occurs immediately; a similar peak value measure is calculated including both pension and Social Security accumulations. Unsurprisingly, peak value is equivalent to the Stock and Wise option value under a set of utility function parameter restrictions, as explained in Samwick (1999). Because

Social Security and pensions are no longer combined with one another and with wages, as in the utility-based option value measures, Coile and Gruber can control separately for wages, and use nonlinearities in the Social Security program rules as their main sources of variation in Social Security wealth accumulation. This, they argue, isolates a source of variation in retirement income accumulation that is less subject to bias from omitted unobserved characteristics in the reduced-form estimating equations.

We share the concerns raised by Coile and Gruber about the source of variation in reduced-form implementations of the option value model, but it is likely that Coile and Gruber's approach does not go far enough in eliminating variation in retirement incentives that could be correlated with unobservables such as tastes for retirement. As mentioned above, if wages are likely to be correlated with tastes for leisure or retirement, it also seems likely that private pension structures will be similarly correlated. Thus, simply using variation in pension wealth accumulation may result in similar problems from omitted variable bias.

In this paper, we develop an identification strategy that controls for unobserved heterogeneity that is potentially correlated with both private pension incentives and retirement incentives coming through Social Security benefits. Our estimating model is a modification of the option value framework that takes account of the different dependent variable based on retirement expectations.

III. Data and Econometric Methodology

To examine the effects of retirement incentives on retirement expectations, we use data from the first three waves of the Health and Retirement Study (HRS), conducted in 1992, 1994 and 1996. The HRS interviewed individuals aged 51 to 60 as of the first survey wave in 1992, along with their spouses, and collected detailed information on demographics characteristics, employment, employment history, pensions, assets, income, health, and subjective expectations of various future events. Our dependent variable comes from the following question asked of all employed individuals age 62 and younger: “Thinking about work generally and not just your present job, what do you think are the chances that you will be working full-time after you reach age 62?” A similar question is asked with respect to the chances of continuing work after age 65. For ease of exposition, we refer to the age 62 questions, although we have performed many of our analyses using the age 65 question as well.

We begin by presenting some simple descriptive statistics on our dependent variables, the subjective probabilities of working beyond age 62 or beyond age 65, in Table 1. There is substantial variation in individual responses to this question, with a large number of individuals who are “certain” that they either will or will not work beyond the given age, leading to responses that are clustered at 0 and 1. One-third of women and one-quarter of men report that there is no chance they will be working full-time after they reach the age of 62, and around half report a zero probability of working after age 65. Approximately 17 percent of women and 23 percent of men respond that they are certain to be working after age 62, while 6 and 9 percent of women and men are certain that they will be working after age 65. There is also a smaller clustering of responses around 0.5.

Because we are using these retirement expectation responses in a context in which actual retirement behavior is the primary concern, it is important to relate these subjective expectations to some measure of actual behavior. To do this, we take responses to the retirement expectations from the initial wave of the survey in 1992 and then form a sample of those individuals who have reached age 63 by the time of the third wave survey, approximately four years later.² Some simple conditional mean calculations make clear that subjective retirement expectations from the wave 1 survey are related to employment status at wave 3. Among individuals over the age of 62 by wave 3 who are no longer working full time (at least 35 hours per week), the mean of the wave 1 expectation of work beyond age 62 is 0.43, compared to 0.67 for those who are working full time at wave 3. Furthermore, the correlation coefficient between wave 1 expectations and whether individuals (over age 62) are working full time at wave 3 is 0.35 for men and 0.34 for women, while the analogous correlation coefficient based on wave 2 expectations is significantly higher at 0.60 for men and 0.54 for women. The significant increase in predictive power between wave 1 and wave 2 remains evident when we estimate linear probability models for whether an individual (over age 62) is working full time at wave 3, controlling for the wave 1 or wave 2 expectation as well as age, health, marital status and education. This is consistent with new information arriving between waves 1 and 2 that allows individuals to adjust their expectations. As individuals grow closer to the relevant retirement age, their expectations are stronger predictors of their future work or retirement status.

² We have performed the same exercises using individuals aged 66 or over by wave 3, and the expectations of working beyond age 65. The results are similar to those for the sample aged 63 and over, although the sample sizes are much smaller with the higher age requirement.

The extensive information on earnings, employment, pensions and assets collected at each survey wave allows us to calculate retirement incentives among older workers in the HRS. Employment and pension information is collected at the wave 1 survey on earnings and pensions connected with both current jobs and jobs that have ended. Pension details include eligibility, structure, and benefit amounts on up to three different pension plans associated with each employer. At each subsequent survey wave, respondents are asked whether their employment and/or pension information has changed and what those changes are. Even for those who do not report a change in pension plans, information is again collected at each survey wave on the pension benefits and eligibility rules.

We are relying primarily on self-reported pension information from the three survey waves of the HRS to construct our measures of retirement incentives. Although Pension-Provider and Social Security Administration data matched to the HRS files have been used by researchers to forecast pension accumulation and Social Security wealth³, the matched employer pension plan file provides details on pensions at wave 1 only. Our interest in changes in retirement incentives means that we must use the self-reported longitudinal data. While concerns have been raised regarding the accuracy and completeness of self-reported pensions in the HRS, for our purposes, the self-reports match well with the subjective retirement expectations since it is individuals' understanding of their own pensions that should affect their expectations.

³ See, for example, Gustman and Steinmeier (1998), Moore and Mitchell (1997), and Uccello and Perese (1999).

An important consequence of our use of pension plan information is that we must delete from our samples those individuals who claim to have a pension, but who do not provide any information on the benefit or account amounts, or on age of eligibility. While this is a matter for some concern, we have few alternatives. Note that a similar problem arises with the matched employer pension data since a substantial fraction of individuals refused to authorize the collection of pension information from their employers and some “authorized” employers failed to respond. Moreover, if an individual is unaware of their pension eligibility, no information was collected regarding their employer’s pension plan.

To calculate the Social Security benefits to which individuals in our sample are entitled, we make use of the complete Social Security covered earnings histories of individuals in the HRS sample. We calculate each individual’s benefit entitlement based on years of covered earnings and the level of earnings in each year. This involves calculating each individual’s Annual Indexed Monthly Earnings (AIME) based on their highest 35 years of covered earnings and using the Social Security program rules to convert the AIME into annual benefit amounts for alternative retirement dates. Since the earnings histories are only available up to 1991, we use self-reported earnings for subsequent years in which respondents are still working.⁴ We also calculate eligibility for spousal benefits, and assume that individuals will receive the maximum of their own benefit entitlement, or half of their spouse’s, following Social Security program rules.

⁴ We assume that earnings remain constant between one earnings report and the subsequent report in the next wave of the HRS.

We next provide some details on how the pension and other data are used to form our main independent variables of interest. This requires some discussion of the models to be estimated. The basic estimation approach that we use follows closely the reduced-form implementations of Lumsdaine, Stock and Wise (1992), Samwick (1998) and Coile and Gruber (2000) discussed above. In these models, an indicator for retirement is the dependent variable, and the main independent variable of interest is the option value of delayed retirement. Since non-utility-based measures of option value (such as Coile and Gruber’s peak value) are nested in the utility-based measures (as in the original Stock and Wise specification), we will henceforth refer to option value broadly defined as encompassing either case, and note that lifetime utility refers to the present value of pension wealth when option value is defined without utility parameters. To distinguish between the two measures, we refer to the measure using only pension wealth (no utility parameters or other income information) as “pension gain.”

In the standard reduced-form approach to Stock and Wise, the lifetime utility $V_t(R)$ associated with each possible retirement age R is calculated for each individual. Option value OV_t is calculated by taking the highest possible lifetime utility (associated with the “optimal” retirement date R^* between time t and the maximum time S), and subtracting from it the lifetime utility available if the individual retires immediately at time t :

$$(2) \quad OV_t = V_t(R^*) - V_t(t)$$

$$\text{where } R^* = \operatorname{argmax}_{\{t \leq R \leq S\}} V_t(R)$$

The likelihood of retirement is then estimated with OV_t as an explanatory variable.

Our key explanatory variables differ from the previous studies in two ways. First, our use of retirement expectations as the dependent variable necessitates a slight modification of this option value approach. In order to derive a measure of retirement incentives that focuses on the pre-/post-age 62 distinction, we calculate the maximum lifetime pension wealth (or utility) from retiring after age 62, and subtract from it the maximum lifetime pension wealth available from retiring prior to or at age 62:

$$(3) \quad OV62_t = V_t(R_{>62}^*) - V_t(R_{\leq 62}^*)$$

$$\text{where } R_{\leq 62}^* = \operatorname{argmax}_{\{t \leq R \leq 62\}} V_t(R)$$

$$R_{>62}^* = \operatorname{argmax}_{\{62 < R \leq S\}} V_t(R)$$

This retains the concept of maximizing behavior inherent in traditional option value approaches, but modifies it to fit the expectation based dependent variable: *OV62* is the option value from delaying retirement until after age 62.

We employ three different measures of the option value of delaying retirement: (i) pension gain, (ii) pension and Social Security benefit gain (iii) option value based on assumed utility function parameters. The first two measures are analogous to the peak value of Coile and Gruber that does not require explicit assumptions about the utility parameters. In the second measure, we look at the present value of the sum of pension and Social Security benefits, and calculate the “gain” variable as described above.

In calculating the value of pension wealth at each potential retirement age, we discount future income and weight this income by age-specific survival probabilities. In the results presented below, we use a discount rate of 3 percent, following both Samwick (1998) and Coile and Gruber (2000). Survival probabilities are taken from the Social Security Administration’s *Annual Statistical Supplement*, 1997.

A second deviation from previous work arises in our measure of utility-based option value. This measure is based on the Stock and Wise framework, but we use an alternative form for the utility function that allows saving and borrowing across time periods. The framework is discussed in Chan and Stevens (1999). We assume that at time t , individuals maximize a lifetime utility function of the form:

$$(4) \quad V_t = \sum_{s=t}^S \beta^{s-t} u(c_s, l_s)$$

subject to a lifetime budget constraint at time t of:

$$(5) \quad \sum_{s=t}^S \left(\frac{1}{1+i} \right)^{s-t} c_s = A_t + \sum_{s=t}^S \left(\frac{1}{1+i} \right)^{s-t} y_s(R)$$

where c is consumption, l is leisure, S is the last possible date, β is the subjective discount factor, i is the real interest rate, A_t is the stock of assets at time t and y_s is income in period s , which depends on the retirement date R . We assume a constant relative risk aversion utility function with risk aversion parameter $1/\alpha$. Further, we assume that the utility of leisure is 0 while working and equal to a constant U_L when retired. Individuals chose the retirement age from the maximization of:

$$(6) \quad \max V_t(R) = \sum_{s=t}^S \beta^{s-t} \frac{\left(\frac{A_t + Y_t(R)}{\sum_{s=t}^S \beta^{s-t}} \right)^{1-\alpha}}{1-\alpha} + \sum_{s=R}^S \beta^{s-t} U_L$$

The retirement age maximizes utility, which is a function of the present value of lifetime wealth,

including earnings, assets, and retirement income, as well as the utility of leisure.⁵ In previous work (Chan and Stevens, 1999) we have experimented with alternative values for the risk aversion and utility of leisure parameters, and found that they make little difference to the performance of our option value measure in reduced-form regressions. Thus, in the work below, we follow Samwick (1998) and Coile and Gruber (2000), in assuming values for the risk aversion and utility of leisure parameters, but do not find results to be sensitive to the particular values assumed. The utility of leisure parameter is set so that on average, the value of an

⁵ This setup is similar in spirit to that of Stock and Wise. Their functional form is power utility with parameter γ and the value of leisure expressed as a multiplicative factor k on retirement benefits during retirement.

$$V_t(R) = \sum_{s=t}^{R-1} \beta^{s-t} Y^\gamma + \sum_{s=R}^S \beta^{s-t} (kB)^\gamma$$

If their income Y and benefits B are replaced with consumption C_w before retirement, and consumption C_r after retirement, and log utility is used (a limiting case of their power utility, and $\alpha=1$ in our setup), we get

$$V_t(R) = \sum_{s=t}^{R-1} \beta^{s-t} \ln(C_{ws}) + \sum_{s=R}^S \beta^{s-t} \ln(kC_{rs})$$

Breaking out the second term:

$$V_t(R) = \sum_{s=t}^{R-1} \beta^{s-t} \ln(C_{ws}) + \sum_{s=R}^S \beta^{s-t} \ln(C_{rs}) + \sum_{s=R}^S \beta^{s-t} \ln(k)$$

First order conditions would have that consumption is equal in all periods, i.e., $C_w=C_r$, and thus we have our equation [6] where our U_L corresponds to their $\ln(k)$.

additional year of leisure is worth 55 to 65 percent of an additional year of income from working, depending on the age of retirement. We assume a risk aversion parameter, α , of $2/3$.

Our key independent variables of interest, the pension gain, or option value from continuing work beyond ages 62 or 65 are summarized in Appendix Table 1. Note that only 48 percent of women and 66 percent of men in the sample report having any employer-provided pension. This means that a substantial portion of the sample has zero values for the pension gain variables. Below we investigate the sensitivity of the results to focusing only on those individuals with private pension plans. In addition, some workers report having a pension from a previous employer rather than their current employer, and so will have a zero gain to remaining employed with their present employer. Among those with pensions, there are large gains and losses associated with working beyond a particular age. For individuals with a pension eligibility age that is after age 62, the pension gain from working beyond age 62 can be substantial. The average gain among women with positive pension gain values is \$71,179, and for men this average of positive values is \$164,435. For individuals eligible to receive pension benefits at or before age 62, there can also be substantial financial penalties to remaining employed beyond the wealth-maximizing age. Among women with negative pension gain values, the average is –\$58,577, while for men the corresponding average is –\$95,543.

Our basic estimating equation can be expressed as:

$$(7) \quad Y_{it} = \mathbf{X}_{it}\beta + \delta OV62_{it} + \alpha_i + \varepsilon_{it}$$

where Y is the subjective expectation of working beyond age 62, and $OV62$ is the measure of retirement incentives, option value or pension gain. \mathbf{X} includes some standard demographic controls, such as age, education, health status and marital status, and in most specifications, we include earnings and asset measures. Our concern is that the individual fixed-effect, α_i , may be

correlated with the key independent variable $OV62$ and so OLS estimation of equation (7) will lead to biased estimates of δ . To address this concern we estimate fixed-effects models that account for this permanent individual component of the error term.

IV. Within-Person Variation in Retirement Incentives

The fixed-effects specification directly focuses attention on how the within-person variation in option value is generated. We believe that one important source of variation in reported pension details across years for a given individual is the acquisition of information about pension plans. Gustman and Steinmeier (1999) have found that pension beneficiaries in the HRS do not know the details of their pension plans very well, but that this knowledge seems to improve somewhat with age. Because these changes in self-reported pension eligibility and benefits do not represent *actual* changes in pension plans, it might be tempting to dismiss this source of variation as noise. However, for our purposes, this is actually an important source of information since it is the pension incentive an individual believes to be available that should influence his or her retirement plans in a given year, not the actual benefits that may be known only by an employer or the Social Security Administration. The purpose of making this distinction is simply that changes across waves in self-reported pension information need not reflect actual changes in pension plans to provide meaningful variation in our estimations.

Within-person variation in retirement incentives may also come from actual changes in pension plans. In our estimations below, we report results that focus exclusively on those individuals experiencing some observable “events” that are likely to generate changes in individuals’ pension incentives, including involuntary job changes, early retirement incentives,

and other discrete changes in employer-sponsored pension plans. This is intended to focus on individuals for whom there is a relatively clear source of within-person variation in the pension gain or option value variables. Additionally, it seems likely that acquisition of information about a pension plan may be endogenous to the retirement decision. Individuals who have decided to retire in the near future may be most likely to update their knowledge of their pension details. Focusing on a sample of individuals with specific changes in their employment or pension status will reduce the reliance on (potentially endogenous) information acquisition as a source of variation.

Specifically, in some of the results below, we limit the sample to individuals who meet one of three criteria: (i) report having involuntarily lost a job (due to a plant closing or layoff) between waves 1 and 3, (ii) report having received a special “early retirement incentive” at one of the survey waves, or (iii) report that something about their pension plan has changed from one survey wave to the next. Note that we do not include individuals who have experienced *voluntary* job changes in this group. These individuals are likely to have significant changes in their pension incentives, but this may be a result of their planned retirement behavior. As noted by Ruhm (1990), many older workers leave their career job for “bridge jobs” before completely retiring. For such people, the pension gain variable may be determined by a prior decision to retire in the near future.

Job loss will produce a change in pension incentives by, at a minimum, changing a pension from being associated with a current job to one that is associated with a previous job. This means that a decision to retire from a post-displacement job does not have any impact on pension income associated with a pre-displacement job. In previous work (Chan and Stevens, 2001) we show that job loss among older workers typically results in workers maintaining some

pension benefits, and that they often begin receiving pension benefits immediately as a form of severance pay, regardless of their re-employment status. Early retirement incentives may temporarily alter an individual's expected pension benefits in a number of ways such as reducing the minimum age for pension eligibility, or altering the annual benefit amount. Because we have information on when these early incentive plans were offered, we can distinguish between periods when the increased pension incentives were and were not available. Finally, individuals who remain with the same employer across survey waves are asked directly whether some aspect of their pension plan rules has changed since the previous wave. For individuals who answer yes to this question, we again can have more confidence that there is some true within-person variation in the pension incentives that they face.⁶ The within-person variance of the pension gain measure in this "change" sample is approximately 2.5 times larger than the within-person variance for the sample of individuals who do not experience an easily identifiable shock to their pension incentives. Estimation based on this restricted sample produces results very similar to estimation using the full sample.

⁶ Note that, even if individuals do not report a change in their pension rules, they are again asked for information regarding benefit amounts, eligibility ages, account balances, etc.. Thus, there are sometimes significant changes in reported pension details even among individuals who claim that nothing has changed. Our focus on individuals who report an explicit change is meant to focus on a sample that we view as more likely to have experienced a well-defined change in their pension plans across survey waves.

V. Results

We next begin to examine the relationship between workers' subjective retirement expectations and the pension gain or option value. In Table 2, we report OLS and fixed-effects estimates of the expectation of work beyond age 62 on the pension gain measure, estimated separately for men and women. Standard errors are corrected for clustering of observations from the same individuals, and for possible heteroscedasticity arising from information acquisition over time.⁷ The variable of interest here is the pension gain variable. The coefficients on pension gain in the OLS regressions are positive and significant for both men and women. This is consistent with much of the previous literature, showing that pension incentives significantly affect retirement decisions. Larger gains to continuing work beyond age 62 significantly increase individuals' stated expectations of working beyond age 62.

To interpret the magnitude of the pension gain coefficient, consider an individual with the average private pension plan that encourages delaying retirement beyond age 62, and replace that pension with another that has a zero pension gain. This scenario reflects one possible effect of the switch in popularity from defined benefit to defined contribution pension plans that has occurred in recent years. The coefficients from the first and third columns of Table 2 imply that a woman with the average positive pension gain (\$71,179) would have a 0.06 lower expectation of working after age 62 once her pension is changed to having zero pension gain. A man with the average positive pension gain (\$164,435) is predicted to have an expectation that is

⁷ There is also a potential problem arising from the clustering of the expectation responses at zero, 0.5 and one. However, when we estimate the models using a sub-sample of individuals with responses other than zero, 0.5 and one, we obtain similar results.

approximately 0.10 lower once his pension is changed to having zero pension gain. Note that the larger effect for men is driven entirely by the larger average value of the pension gain, while the estimated coefficient on pension gain is smaller for men than for women.

These regressions also include controls for age, health, education, marital status, race, ethnicity, as well as annual earnings and total wealth. Most of the coefficients from the OLS regressions have the expected signs. The coefficient on age is negative and significant, but the squared term is positive. The expectation of working after age 62 is initially declining with age, but begins to increase with age after age 57 or 58. Poor health reduces the expectation of working beyond age 62 for men and women. There are some small differences in retirement expectations by education level. Black women are significantly less likely to expect work beyond age 62. Being married decreases a woman's expected probability of continuing work beyond age 62 by 0.16, but has no effect for men. This is not surprising since married women typically have significant sources of labor and retirement income through their husband's earnings, pensions, and Social Security benefits.

We have also included annual earnings and total non-pension assets to account for differences among workers that might be correlated with their pension wealth and with retirement plans. Gruber and Coile (2000) argue that many option value measures are highly correlated with earnings, and the same may apply to wealth. In the OLS results, these variables have the expected signs, although the asset measures are generally insignificant. Higher earnings increase individuals' expectation of continuing work into their 60s, while increased wealth reduces these expectations.

To address the problem of omitted variable bias, the second and fourth columns of Table 2 repeat the analysis using a fixed-effects estimator. If there are important omitted variables such

as tastes for leisure that are negatively correlated with both the probability of working beyond age 62 and with the pension gain variable, we would expect the fixed-effects estimates to result in a smaller coefficient on the pension gain variable. Indeed, there are sharp reductions in the magnitude of the pension gain coefficient in the fixed-effects specifications. For both men and women, the pension gain effect estimated from within-person changes is less than half the magnitude of the OLS coefficient.

The fixed-effects coefficients imply that a woman with the average positive pension gain value would have an expectation of working after age 62 that is 0.02 lower if her pension changed to having zero pension gain, and that a man under the same circumstances would have an expectation 0.04 lower. Recall that under the OLS estimates, these pension changes would have induced a much larger decrease in expectation (0.06 for women and 0.10 for men). The difference in expectation change is consistent with a strong correlation between unobserved fixed-effects and the pension gain measure and suggests that cross-sectional estimates of the relationship between retirement and retirement wealth or accumulation measures may be subject to significant bias. A Hausman test allows us to reject the hypothesis that the person-specific effects are independent of the regressors.

The second half of the table repeats the OLS and fixed-effects regressions using the variable for the expectation of working beyond age 65, and the equivalently defined pension gain variables. Looking first at the OLS results, the estimated magnitudes of the coefficients on the pension gain from working past age 65 are somewhat smaller than those for age 62. The fixed-effects estimates using the age 65 variables produce dramatically smaller coefficients on the pension gain variable compared with the age 62 results. One possible explanation for this pattern is that relatively few private pensions have peaks after age 65. Thus, there is less variation in the

gain to working beyond age 65 than in the gain to working beyond age 62. In most of the results that follow, we focus on the variables for working beyond age 62, but note that there may be important differences in pension incentive effects depending on the specific retirement ages being considered.⁸

Because the questions about working beyond age 62 or 65 do not differentiate between an individual's desire to continue working, and their ability to continue working, we have also estimated the basic models including some additional controls. If individuals anticipate losing a job or suffering diminished health over the next few years, they may have lower expectations of working in the future for reasons that have little to do with financial incentives for retirement. To test for such a possibility, we have used additional expectations data from the HRS surveys. At each wave, individuals were also asked to provide their expectation of losing their current job, their expectations of being able to find a similar job (in the event of a job loss), and their expectation that health would limit their ability to work in the next several years. While the expectation of losing the current job was never significant, the expectation of easily finding a similar job was consistently positive and significant. This latter variable may be reflecting a high valuation of one's human capital and attachment to the labor force. Surprisingly, the coefficient on the expectation of facing a health limitation in the future was positive and significant. Including responses to these questions as additional controls, however, had little effect on the estimated coefficients on the pension gain variables.

⁸ A similar point is made by Coile and Gruber (1999), who note that the effects of Social Security and pension incentives are much stronger between ages 62 and 69 than between ages 55 and 61.

Some additional specifications of the basic OLS and fixed-effects regressions are summarized in Table 3. We repeat the analysis using both the sum of private pension and Social Security wealth and using the utility-based option value variable (described above) that combines pension, Social Security, earnings, and asset information, along with assumed utility function variables. We believe it is important to incorporate Social Security measures into the incentives measures, although there are some caveats necessary for these results. First, the pension and Social Security measures that are combined in Table 3 are not completely comparable, since the Social Security measures represent actual benefit amounts and eligibility, and the pension measures represent individuals' understanding of their pension benefits and eligibility.⁹ Second, the sources of within-person variation in the Social Security measures are not necessarily the same as for the pension data. The main reason why individuals' Social Security entitlement changes across survey waves is that earnings change across the waves. We control separately for earnings, so the change in Social Security benefits reflects the non-linear way in which earnings translate into benefits. However, if these earnings changes reflect either voluntary job changes or measurement error, the fixed-effects approach may not produce an unbiased estimate.¹⁰ In short, the identification strategy used here is likely to be better suited to the self-reported pension incentives than to the Social Security measures.

⁹ Although the HRS also contains subjective estimates of the magnitude of future Social Security benefits at the subjective expected date of first receipt, it was not clear how these should be adjusted for different retirement dates.

¹⁰ Recall that we have administrative earnings reports with which to compute Social Security benefits only through 1991; for subsequent years we use self-reported annual earnings.

The coefficients on the pension plus Social Security measures are similar to those using pension information only. We have also estimated these models using only the Social Security gain measure (and ignoring private pensions entirely). The coefficients on the Social Security gain variable are wrong-signed and sometimes significant in the OLS regressions, and are not statistically significant in the fixed-effects regressions. This is consistent with there being too little within-person variation in the Social Security gain variable for the fixed-effects estimates to be informative. It is worth noting that Coile and Gruber (1999) found no statistically significant relationship between retirement and the “peak value” defined using Social Security benefits when they restricted their sample to those under 61, as is consistent with our sample here. This may reflect the fact that Social Security becomes more important to retirement decisions once individuals reach the age of eligibility in their 60s.

The option value measures based on utility comparisons are similar to those based on pension gain. The fixed-effects estimates are positive and statistically significant, and are substantially below the corresponding OLS estimates. To provide a meaningful interpretation, consider the previous example of replacing an individual’s pension with the average incentive to continuing work past age 62 with a pension that has zero incentive. The fixed-effects coefficient on option value for men suggests that this individual would have an expectation that is 0.01 lower when he no longer has the positive pension incentive while keeping other components of the option value measure the same. Decreasing an individual’s option value by one standard deviation would lower the expectation of working beyond age 62 by 0.03.

In Table 4, we report results from two alternative samples of men and women, focusing again on the pension gain. First, we use a sample that is restricted to individuals who report having a private, employer-provided pension. (The full sample results are repeated in the first set

of columns for comparison.) This follows Samwick (1998) and tests whether there is a distinction between the effect of having no pension versus any pension, and having a more generous or more steeply sloped pension wealth-retirement age profile. Samwick found that a sample restricted to individuals with pensions produced slightly stronger effects of option value on retirement. He interpreted this finding as evidence that the pension effect was not being driven entirely by the contrast between those with pensions and those without and our results confirm this finding. The OLS and fixed-effects estimation using the sample of those with pensions (most comparable to Samwick's specification) produce estimates of pension gain effects that are very close to those based on the full sample. The similarity of the within-person estimates across the two samples is not surprising since very few individuals go from reporting a pension to reporting none in a subsequent wave.

The second new set of results in Table 4 is based on a sample limited to individuals with a well-defined change in their pension structures or incentives between survey waves. As mentioned above, this is intended to focus on individuals for whom there is a relatively clear source of within-person variation in the pension gain or option value variables. The regression results for this sample are consistent with results from the broader samples. For men, the pension gain coefficients are similar in magnitude to the full sample and show significant reductions in moving from the OLS to the fixed-effects estimates. The pension gain coefficient from the fixed-effects regression is approximately half the size of the OLS coefficient. For women the OLS and fixed-effects coefficients are somewhat smaller than in the full sample and the fixed-effects coefficient on pension gain is no longer statistically significant, perhaps because the sample size is so dramatically reduced. Overall our results, particularly for men, are not very

sensitive to these sample changes which may focus on more exogenous sources of within-person variation.

Several previous authors have investigated whether *levels* of retirement wealth (both pension and non-pension) significantly affect retirement behavior, after conditioning on variables capturing the future accumulation of pension wealth such as option value or pension gain. If individuals are liquidity constrained or have relatively low wealth holdings, they may be constrained to continue working. In Table 5, we test whether variables for low pension wealth and low non-pension wealth affect retirement plans. In particular, we include a variable equal to one if total asset are less than \$10,000, and a variable indicating that the value of pension wealth (assuming retirement at age 62) is \$50,000 or less.¹¹

The results in Table 5 suggest that low levels of wealth may be important determinants of retirement planning. While the magnitude of coefficients on both low assets and low levels of retirement wealth are reduced by including fixed effects in the regression, they remain positive and are usually statistically significant. Having very low asset holdings increases the likelihood of continuing work beyond age 62 by 0.05. This is slightly larger than the effect of a change from a zero pension incentive to one the average positive pension incentive. For women, low pension wealth also increases the expectation of working beyond age 62. Inclusion of these variables slightly decreases the magnitudes of the pension gain coefficients, but the differences are not statistically significant. The finding that levels of pension and non-pension wealth are related to retirement expectations differs from recent findings by Samwick (1998). Samwick

¹¹ These results are not sensitive to changes in the definition of “low assets”. Similar results are obtained if we use alternative definitions of low asset holdings up to \$30,000.

notes in his conclusion that it is only the change in pension wealth (through the option value) and not the level that affects the decision to retire. While we do not focus on actual retirement, but on retirement expectations, our findings suggest that wealth levels may be important. Individuals with low levels of wealth are more likely to plan to continue working beyond age 62. Further, these effects remain after controlling for individual fixed effects. While option value or pension gain variables are important, having adequate levels of wealth also appears to be crucial to retirement timing. Coile and Gruber (2000) also provide evidence that wealth levels may be important, even after controlling for retirement wealth accumulation.

We have estimated several additional specifications as robustness checks on our main results. First, because many individuals in their early 50s may not have begun to think in great detail about retirement, we have restricted the sample to include only those over the age of 55. Results for this group are very similar to the results in Table 2. Second, our calculation of the pension gain and option value variables use survival probabilities that differ only by age and gender. Heterogeneity across individuals in health or life expectancy may lead some individuals to more heavily discount future pension returns than others. To account for this, we have included an additional variable in the regressions, indicating an individual's subjective expectation of living beyond age 75, as well as an interaction between this variable and the pension gain or option value measure. None of the interaction terms are statistically significant, and inclusion of the life expectancy variables does not change the coefficients on our main variables of interest. Finally, because single and married women may have very different plans and resources available for retirement, we have estimated the results for women separately by marital status. The effects of measured individual retirement incentives are similar for married and unmarried women.

A final concern with our fixed-effects results is whether they are subject to bias due to measurement error in reported pension incentives. While within-person estimates are often subject to increased bias from measurement error in the independent variable, we do not believe this is a serious problem for the pension incentive measures used here. First, the pension gain (and option value) measures are intended to capture what individuals believe their pension wealth profile to look like, not necessarily the “true” pension plan known to the employer. Thus, deviations between the perceived and the true pension plan should not be viewed as measurement error in this context. Secondly, there is evidence that deviations between self-reported and employer-reported pension plans across waves of the HRS do reflect a lack of knowledge about true pensions, rather than random noise in the reports. Gustman and Steinmeier (1999) report that, as individuals move closer to normal retirement ages, they are more likely to accurately report certain aspects of their pension plans (based on comparison with employer reports).

Although we do not view deviations between the true and reported pension provisions as measurement error in the context of this study of retirement expectations, there do remain potential sources of measurement error in the self-reported pensions. In particular, while the HRS asks respondents an extensive set of questions to describe their pension incentives, it is possible that the responses to these questions do not allow respondents to fully characterize the future profile of their potential retirement wealth. This would mean that the constructed pension gain and option value variables do not fully reflect what individuals *believe* to be their actual profile of pension wealth accumulation. Such a source of measurement error is likely to be highly serially correlated across years since the structure of the HRS questions across years does not change a great deal. In the extreme case of perfectly serially correlated errors, the fixed-

effects estimator would eliminate such measurement error and so it would not be a concern in interpreting our main results.¹²

Finally, it is also possible that individuals do randomly misreport their actual knowledge of their pensions plans, leading to classical measurement error. We have attempted a more direct investigation of this possibility. Because different methods of eliminating the fixed-effects will lead to different degrees of bias from classical measurement error (Griliches and Hausman, 1986), we have repeated our analysis using both first-differences and long-differences – differencing between waves 1 and 3. All of these alternative estimators (fixed-effects, first-differences, long-differences) produce very similar estimates of the coefficients on pension gain and option value. The first-differenced estimator and the fixed-effects estimator can be combined to produce an estimator that is consistent in the face of serially uncorrelated measurement error. This “Griliches-Hausman” estimate is very close to the main estimates presented earlier, but with extremely large standard errors. For men, the pension gain coefficient from the “Griliches-Hausman” estimator is 0.019 with a standard error of 0.04 (comparable to the fixed-effects estimate of 0.0216 from Table 2). The lack of precision likely reflects the fact that we have only three waves of data and so are looking at differences of only 2 versus 4 years.

¹² Similar reasoning would apply to measurement error introduced by our assumptions of discount rates and utility function parameters.

As a result, the combined estimator has very low power.¹³ This admittedly weak evidence, along with consideration of the nature of the self-reported pension data provide some assurance that the difference between our OLS and fixed-effects estimates is not the result of classical measurement error.

VII. Conclusion

This paper investigates the relationship between retirement expectations and incentives for retirement provided by employer-sponsored pension plans and Social Security. Because we focus on repeated observations of individuals' subjective expectations of continuing work, we are able to control for unobserved individual fixed-effects that may be correlated with pension incentives. Our results provide evidence that this unobserved heterogeneity is important, and that existing research which largely ignores such heterogeneity may be substantially overstating the responsiveness of individuals to pension-related incentives.

Our main conclusions can be summarized as follows: first, forward-looking measures of pension wealth only, and broader measures that include earnings, Social Security and assets, are

¹³ We have also attempted an instrumental variables approach, using employer-reports of wave 1 pension incentives, combined with information on job loss, early retirement incentives and discrete pension changes (and interaction terms) as instruments for the change in pension gain across waves. Unfortunately, this also resulted in very imprecisely estimated coefficients on the pension gain variables that were not statistically distinguishable from either the OLS or the fixed-effects estimates.

significantly related to individuals' expectations of continuing work into their 60s. This confirms the main finding of the literature based on actual retirement behavior. Second, this relationship remains significant even when we control for unobserved heterogeneity using fixed-effects estimators. Thus, there can be little doubt that individuals consider these forward-looking incentives when making their retirement plans. Third and most importantly, however, the magnitude of these incentive effects varies dramatically between OLS and differenced estimation strategies. Coefficients on pension incentive measures from fixed-effects regressions are less than half the magnitude of similar OLS regressions.

While most of the literature in this area has focused on actual retirement behavior, our results pertain to individuals' expectations about continuing work beyond age 62. We show that there are strong connections between these subjective expectations and actual future work behavior, and so our results are also relevant to longitudinal studies of actual retirement. These findings pose a challenge to future research on retirement to better control for heterogeneity in tastes for retirement that may be correlated with observed pension and other retirement wealth.

References

- Chan, Sewin and Ann Huff Stevens, "Does Job Loss Alter the Timing of Retirement? Exploring the Role of Earnings, Assets and Pensions," manuscript, New York University (1999).
- Chan, Sewin and Ann Huff Stevens, "The Effects of Job Loss on Older Workers: Employment, Earnings and Wealth," in Peter P. Budetti, Richard V. Burkhauser, Janice M. Gregory and H. Allan Hunt (eds.), *Ensuring Health and Income Security for an Aging Workforce* (Kalamazoo: W.E. Upjohn Institute for Employment Research, 2001).
- Coile, Courtney and Jonathan Gruber, "Social Security and Retirement," manuscript, MIT (1999).
- Coile, Courtney and Jonathan Gruber, "Social Security and Retirement," NBER Working Paper 7830 (2000).
- Griliches, Zvi and Jerry A. Hausman, "Errors in Variables in Panel Data," *Journal of Econometrics* 31:1 (1986), 93-118.
- Gustman, Alan L. and Thomas L. Steinmeier, "A Structural Retirement Model," *Econometrica* 54 (1986), 555-584.
- Gustman, Alan L., and Thomas L. Steinmeier, "Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study," NBER Working Paper 6681 (1998).
- Gustman, Alan L., and Thomas L. Steinmeier, "What People Don't Know about their Pensions and Social Security: An Analysis using Linked Data from the Health and Retirement Study," NBER Working Paper 7368 (1999).
- Krueger, Alan B., and Jorn-Steffen Pischke, "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation," *Journal of Labor Economics* 10:4 (1992), 412-37.

- Lumsdaine, Robin, James Stock and David Wise, “Three Models of Retirement: Computational Complexity Versus Predictive Validity,” in David Wise (ed.) *Topics in the Economics of Aging* (Chicago: University of Chicago Press, 1992).
- Moore, James F., and Olivia S. Mitchell, “Projected Retirement Wealth and Savings Adequacy in the Health and Retirement Study,” NBER Working Paper 6240 (1997).
- Ruhm, Christopher, “Bridge Jobs and Partial Retirement,” *Journal of Labor Economics* 8 (1990), 482-501.
- Rust, John, “A Dynamic Programming Model of Retirement Behavior,” in David Wise (ed.) *The Economics of Aging* (Chicago: University of Chicago Press, 1989).
- Rust, John and Christopher Phelan, “How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets,” *Econometrica* 65:4 (1997), 781-821.
- Samwick, Andrew, “New Evidence on Pensions, Social Security, and the Timing of Retirement,” *Journal of Public Economics* 70:2 (1998), 207-36.
- Samwick, Andrew, “Comments on Coile and Gruber, ‘Social Security Incentives for Retirement,’” manuscript, Dartmouth College (1999).
- Stock, James and David Wise, “Pensions, the Option Value of Work, and Retirement,” *Econometrica* 58 (1990), 1151-1180.
- Uccello, Cori, and Kevin Perese, “Wealth Accumulation in the Health and Retirement Study: The Importance of Including Pension Wealth,” Urban Institute Report (1999).
- Venti, Stephen, and David Wise, “Choice, Chance, and Wealth Dispersion at Retirement,” NBER Working Paper 7521 (2000).

Table 1
Distribution of Subjective Retirement Expectations

	Expectation of working after age 62		Expectation of working after age 65	
	Women	Men	Women	Men
0	33.8%	25.7%	55.1%	44.7%
0.01-0.09	1.1%	1.1%	1.7%	2.2%
0.10-0.19	5.2%	4.7%	6.3%	6.9%
0.20-0.29	6.5%	6.0%	7.3%	7.5%
0.30-0.39	3.5%	3.0%	3.2%	3.1%
0.40-0.49	2.6%	2.3%	2.3%	2.5%
0.50-0.59	15.1%	13.8%	10.3%	12.4%
0.60-0.69	2.1%	1.9%	1.4%	1.9%
0.70-0.79	4.6%	5.5%	2.8%	3.5%
0.80-0.89	4.7%	6.4%	2.3%	3.8%
0.90-0.99	3.7%	6.0%	1.3%	2.4%
1	17.2%	23.4%	5.9%	9.1%
Number of observations	5,703	5,735	5,783	5,948

Notes: Pooled data from the first three waves of the HRS.

Table 2
OLS and Fixed-Effects Estimates of the Effect of Pension Incentives on Retirement Expectations

Dependent variable: Sample: Estimation procedure:	Expectation of work after age 62				Expectation of work after age 65			
	Women		Men		Women		Men	
	OLS (1)	FE (2)	OLS (3)	FE (4)	OLS (5)	FE (6)	OLS (7)	FE (8)
pension gain (\$100,000)	0.0830 (0.0223)	0.0282 (0.0106)	0.0624 (0.0135)	0.0216 (0.0071)	0.0648 (0.0131)	0.0162 (0.0116)	0.0343 (0.0127)	0.0011 (0.0060)
age	-0.2059 (0.0651)	-0.2096 (0.0720)	-0.1365 (0.0713)	-0.0390 (0.0696)	-0.0656 (0.0448)	-0.0297 (0.0522)	0.0289 (0.0503)	0.0495 (0.0528)
age ²	0.0019 (0.0006)	0.0015 (0.0006)	0.0013 (0.0006)	0.0002 (0.0006)	0.0006 (0.0004)	0.0004 (0.0004)	-0.0002 (0.0004)	-0.0005 (0.0004)
self-reported health (1=excellent, 5=poor)	-0.0176 (0.0060)	-0.0081 (0.0089)	-0.0468 (0.0061)	-0.0113 (0.0079)	-0.0142 (0.0048)	0.0015 (0.0069)	-0.0325 (0.0053)	0.0017 (0.0066)
couple	-0.1632 (0.0142)	-0.0052 (0.0505)	0.0151 (0.0198)	0.0950 (0.0489)	-0.1424 (0.0121)	-0.0556 (0.0398)	-0.0249 (0.0186)	0.0349 (0.0411)
some college	0.0381 (0.0164)		0.0253 (0.0181)		0.0270 (0.0133)		0.0292 (0.0154)	
college graduate	0.0260 (0.0179)		0.0370 (0.0188)		0.0444 (0.0153)		0.0588 (0.0162)	
black	-0.0754 (0.0174)		-0.0167 (0.0207)		-0.0491 (0.0137)		-0.0003 (0.0183)	
hispanic	0.0139 (0.0240)		0.0811 (0.0240)		0.0449 (0.0193)		0.0864 (0.0217)	
annual earnings	4.60E-07 (2.06E-07)	2.76E-07 (2.26E-07)	2.68E-07 (3.03E-07)	3.32E-07 (3.59E-07)	8.59E-09 (2.30E-07)	-4.67E-09 (1.86E-07)	2.83E-07 (2.20E-07)	3.75E-07 (3.62E-07)
non-pension assets	-2.10E-08 (1.54E-08)	1.26E-09 (8.63E-09)	-5.54E-09 (9.13E-09)	-1.44E-09 (7.14E-09)	-8.79E-09 (7.83E-09)	-1.16E-10 (5.86E-09)	5.06E-09 (8.03E-09)	-1.50E-09 (5.36E-09)
wave 2	-0.0494 (0.0099)	0.0649 (0.0357)	-0.0500 (0.0104)	0.0016 (0.0341)	-0.0297 (0.0078)	-0.0504 (0.0288)	-0.0471 (0.0084)	-0.0218 (0.0288)
wave 3	0.0072 (0.0139)	0.1798 (0.0729)	-0.0298 (0.0137)	0.0371 (0.0689)	0.0106 (0.0112)	-0.0485 (0.0586)	-0.0188 (0.0117)	0.0007 (0.0581)
intercept	6.2346 (1.8492)	7.5215 (2.1983)	4.2279 (2.0378)	1.9084 (2.1153)	2.2080 (1.2821)	0.7322 (1.6342)	-0.5940 (1.4523)	-1.0483 (1.6497)
No. of observations	5,461		5,445		5,774		5,933	

Notes: OLS and fixed-effect regressions using data from the first three waves of the HRS.

Pension gain is the pension incentive to delay retirement beyond age 62. Standard errors are in parenthesis.

Table 3
OLS and Fixed-Effects Estimates of the Effect of
Alternative Pension Incentive Measures on Retirement Expectations

Dependent variable:	Expectation of work after age 62			
	Women		Men	
Sample:	OLS	FE	OLS	FE
Estimation procedure:	(1)	(2)	(3)	(4)
pension gain (\$100,000)	0.083 (0.022)	0.028 (0.011)	0.062 (0.014)	0.022 (0.007)
pension and Social Security gain (\$100,000)	0.072 (0.017)	0.025 (0.010)	0.059 (0.013)	0.020 (0.007)
option value (000s)	0.836 (0.055)	0.249 (0.080)	0.903 (0.055)	0.237 (0.063)
Number of observations	5,461		5,445	

Notes: OLS and fixed-effect regressions using data from the first three waves of the HRS.

Pension gain, Social Security gain and option value are measures of the pension incentive to delaying retirement beyond age 62.

Each cell in the table represents a separate regression. Other controls are as in table 2. Standard errors are in parenthesis.

Table 4
OLS and Fixed-Effects Estimates of the Effect of Pension Incentives on Retirement Expectations
using Alternative Samples

Dependent variable:	Expectation of work after age 62											
	Full sample				Workers with a private pension				Workers with a "change" in pension			
Sample:	Women		Men		Women		Men		Women		Men	
Estimation procedure:	OLS	FE	OLS	FE	OLS	FE	OLS	FE	OLS	FE	OLS	FE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
pension gain (\$100,000)	0.083 (0.022)	0.028 (0.011)	0.062 (0.014)	0.022 (0.007)	0.081 (0.022)	0.030 (0.010)	0.061 (0.012)	0.022 (0.007)	0.063 (0.023)	0.019 (0.018)	0.057 (0.001)	0.024 (0.009)
No. of observations	5,461		5,445		2,601		3,593		865		1,395	

Notes: OLS and fixed-effect regressions using data from the first three waves of the HRS.

Pension gain is the pension incentive to delay retirement beyond age 62. Other controls are as in table 2. Standard errors are in parenthesis.

Table 5
**OLS and Fixed-Effects Estimates of the Effect of Pension Incentives,
 Low Assets and Low Pension Wealth on Retirement Expectations**

Dependent variable:	Expectation of work after age 62			
	Women		Men	
Sample:	OLS	FE	OLS	FE
Estimation procedure:	(1)	(2)	(3)	(4)
pension gain	0.072 (0.019)	0.025 (0.011)	0.047 (0.010)	0.020 (0.007)
low assets	0.061 (0.017)	0.045 (0.024)	0.059 (0.019)	0.052 (0.026)
low pension wealth upon retirement at age 62	0.119 (0.015)	0.076 (0.023)	0.204 (0.013)	0.019 (0.019)
Number of observations	5,461		5,445	

Notes: OLS and fixed-effect regressions using data from the first three waves of the HRS.

Pension gain is the pension incentive to delay retirement beyond age 62. Other controls are as in table 2.

Standard errors are in parenthesis.

Appendix Table 1
Pension Gain and Option Value Means and Medians

	Women	Men
% with pension	47.6	65.8
% with DB pension	28.0	47.6
% with DC pension	21.4	25.2
Pension gain 62		
mean	\$8,571	\$22,017
median	\$0	\$0
mean(positive values)	\$71,179	\$164,435
mean(negative values)	-\$58,577	-\$95,543
Pension gain 65		
mean	-\$3,781	-\$10,410
median	\$0	\$0
mean(positive values)	\$44,915	\$114,984
mean(negative values)	-\$58,350	-\$106,414
Option value 62		
mean	32.5	58.3
median	6.3	27.8
Option value 65		
mean	-7.9	6.6
median	0.6	2.5
Number of observations	5,809	5,972

Notes: Calculated using data from the first three waves of the HRS.