

Future Social Security Entitlements and the Retirement Decision

Courtney Coile, Wellesley College and NBER
Jonathan Gruber, MIT and NBER

Revised: June 2006

A critical question for Social Security policy is how program incentives affect retirement behavior. We use the Health and Retirement Survey (HRS) to examine the impact of Social Security incentives on male retirement. We implement forward-looking models whereby individuals consider the incentives to work in all future years. We find that forward-looking incentive measures for Social Security are significant determinants of retirement. We also find that private pension incentives have roughly similar effects. Our findings suggest that Social Security policies that increase the incentives to work at older ages can significantly reduce the labor force exit rate of older workers.

JEL: H55, J26

We are grateful to Dean Karlan for research assistance and to Peter Diamond, Alan Gustman, Jim Poterba, Andrew Samwick, Roger Gordon, seminar participants at MIT, Harvard, the NBER, and the Social Security Administration, and several anonymous referees for very helpful comments. Coile gratefully acknowledges support from the National Institute on Aging through the National Bureau of Economic Research. Gruber acknowledges support from the National Institute on Aging and from the Social Security Administration through the Center for Retirement Research at Boston College. Contact information: ccoile@wellesley.edu or (781) 283-2408; gruberj@mit.edu or (617)253-8892.

One of the most striking labor force phenomena of the second half of the twentieth century has been the rapid decline in the labor force participation rate of older men. In 1950, for example, 81% of 62-year-old men were in the labor force; by 1995, this figure had fallen to 51%, though it has rebounded slightly in the past few years (Quinn, 1999).

Much has been written about the proximate causes of this important trend among older men, and in particular about the role of the Social Security program. A large number of articles have documented pronounced “spikes” in retirement at ages 62 and 65, which correspond to the early eligibility age and full benefit age for Social Security. The presence of these strong patterns in retirement data suggest that the underlying structure of Social Security plays a critical role in determining retirement decisions. But the impact of changes in Social Security generosity on retirement decisions is less obvious. A large literature dating from the mid-1970s has investigated this relationship, and the broad conclusion of that literature is that the level of Social Security benefits has a significant, but modest, effect on retirement dates.

The purpose of our paper is to revisit the impact of Social Security on retirement, using the best available recent data on retirement behavior, the Health and Retirement Study (HRS). This is a comprehensive data set from the 1990s that contains information on demographic and job characteristics, labor force attachment, earnings histories, and the features of private pension plans for a large sample of individuals near retirement age.

We use these data to revisit the important observation of Stock and Wise (1990a,b) that it is not simply the level of retirement wealth or the increment with one additional year of work that matters, but the entire evolution of future wealth with further work. Their “option value” model posited retirement decisions as a function of the difference between the utility of retirement at the current date and at the date that maximizes one’s utility. They used data from a

sample of firms with pension plans to show that the option of future benefit increments affected retirement decisions today in this framework. This finding is echoed in a reduced form context by Samwick (1998), who used data across many firms in the Survey of Consumer Finances.

This past work, however, was subject to three potentially important limitations. First, these papers have focused on all retirement income, including pensions, and have not isolated the impact of Social Security per se. There are a number of reasons to believe that the responsiveness of retirement to Social Security and private pension incentives may differ, and it is only the former that is relevant for Social Security policy-making. Second, retirement decisions in the option value framework are a function of both pension incentives and wages, the latter of which provides the vast majority of variation across individuals in option value. If wage differences across individuals capture partly heterogeneity in tastes for work, however, then building wage variation into the retirement incentive measure can lead to misleading estimates of the responsiveness to financial incentives. Finally, the previous data have either been specific to a handful of firms, or have not contained sufficient detail to compute correctly Social Security incentives for retirement.

Our work remedies these deficiencies by building on the strengths of this impressive new data source. We draw on the insights of Stock and Wise in developing forward-looking measures of incentives, but we also consider in more detail the sources of variation in these incentives, by controlling in a rich way for past earnings, and developing an alternative forward-looking measure that is not primarily driven by wage differences. This measure allows us to examine jointly the impact of Social Security and pensions on retirement, as well as to separate and compare the effects of these different incentives.

We have at least three major findings. First, retirement decisions respond in a significant way to both the accrual of retirement wealth with future work (negatively) and the level of retirement wealth (positively). Second, retirement decisions are much more sensitive to incentive variables defined with reference to the entire future stream of retirement incentives than to the accrual in retirement wealth over the next year alone, indicating that forward-looking measures of this type are important variables to include in retirement models. Third, retirement is roughly equally responsive to a comparable change in Social Security and pension incentives.

Our paper proceeds as follows. We begin, in Part I, with background on both the relevant institutional features of the Social Security system and on the previous literature in this area. We then move on, in Part II, to describe our data and empirical strategy. Part III presents our basic results, and Part IV concludes.

Part I: Background

Institutional Features of Social Security

The social security system is financed by a payroll tax of 7.65 percentage points levied both on workers and firms on wages up to a taxable maximum, currently \$90,000. Workers are eligible for retired worker benefits once they have accumulated 40 quarters of work in covered employment. Benefits are calculated in several steps. First, the worker's Average Indexed Monthly Earnings (AIME) is calculated as their average earnings over their best 35 years, with past earnings adjusted by a national wage index. Second, the AIME is converted to the Primary Insurance Amount (PIA) through a three-piece progressive schedule, whereby 90 cents of the first dollar of earnings is converted to benefits, while only 15 cents of the last dollar of earnings (up to the taxable maximum) is so converted. Finally, workers receive a monthly benefit equal

to the PIA if they commence benefit receipt at the Full Benefit Age (FBA), historically 65 but legislated to slowly rise to 67. If they claim before the FBA, benefits are reduced by an actuarial reduction factor (equal to $5/9$ of one percent per month for the first 36 months of delay and $5/12$ per month thereafter); thus, a worker with a FBA of 65 who claims at the Early Eligibility Age (EEA) of age 62 receives 80% of the PIA. If workers claim after the PIA, benefits are increased by a delayed retirement credit (legislated to rise over time from 3% to 8% per year of delay).

A few other details are worthy of note. For workers claiming before the FBA, an earnings test applies; workers may earn up to a ceiling amount, currently \$12,000, but benefits are reduced by \$1 for each \$2 of earnings above this amount (although months of lost benefits are treated as delayed receipt, entitling the worker to a delayed retirement credit when he resumes full benefit receipt). Since 2000, there has been no earnings test for workers above the FBA. Spouses of beneficiaries receive a dependent spouse benefit equal to 50% of the worker's PIA, which is available once the worker has claimed benefits and the spouse has reached age 62; however, the spouse only receives the larger of this and her own entitlement as a worker. Surviving spouses receive 100% of the PIA and may claim as early as age 60, although there is an actuarial reduction for claiming benefits before age 65 or if the worker had an actuarial reduction. Finally, benefit payments are adjusted for increases in the Consumer Price Index (CPI) after the worker has reached age 62; thus, Social Security provides a real annuity.

Previous Related Literature

A number of studies have used aggregate information on the labor force behavior of workers at different ages to infer the role that is played by Social Security. Hurd (1990) and Ruhm (1995) emphasize the spike in the age pattern of retirement at age 62; as Hurd (1990)

states, "there are no other institutional or economic reasons for the peak". Indeed, as Burtless and Moffitt (1984) document, this spike at age 62 only emerged after the early retirement eligibility age for men was introduced in 1961. Using precise quarterly data, Blau (1994) finds that almost one-quarter of the men remaining in the labor force at their 65th birthday retire within the next three months; this hazard rate is over 2.5 times as large as the rate in surrounding quarters. Lumsdaine and Wise (1994) examine this "excess" retirement at 65 and conclude that it cannot be explained by the change in the actuarial adjustment at this age, nor by the incentives embedded in private pension plans or the availability of retirement health insurance through Medicare. However, this does not rule out a role for Social Security; by setting up the "focal point" of a full benefit age, the program may be the causal factor in explaining this spike.

The main body of this retirement incentives literature attempts to specifically model the role that potential Social Security benefits play in determining retirement. Broadly speaking, there are four classes of studies in this literature. The earliest work in this area, from the early 1970s through the mid-1980s, considered reduced form models of the retirement decision as a function of Social Security wealth and pension levels. Much of this literature is reviewed in Mitchell and Fields (1982); more recent cites include Diamond and Hausman (1984), and Blau (1994). While these articles differ in the estimation strategies employed, with the more recent work using richer models such as non-linear 2SLS or hazard modeling, the results consistently suggest a significant role for Social Security, but a role that is small relative to the time trends in retirement behavior documented in the introduction.

A key limitation of this first class of study is that it considers social security effects at a point in time, but not any impacts on the retirement decision arising from the time pattern of SSW accruals. This was remedied in three different ways by subsequent literatures. The first

was to consider structural modeling of retirement decisions by workers facing a lifetime budget constraint; examples here include Burtless (1986), Burtless and Moffitt (1984), Gustman and Steinmeier (1985, 1986), and Rust and Phelan (1997). The second was to continue to estimate reduced form models, but to incorporate the accrual of SSW with a year of additional work; examples here include Fields and Mitchell (1984), Hausman and Wise (1985), and Sueyoshi (1989). Both of these types of studies continued to find an important, but modest, role for Social Security, and some (e.g. Fields and Mitchell) indicated a larger role for private pensions. The final strand of this literature is the option value work of Stock and Wise noted above.

A final article that deserves particular mention here is that of Krueger and Pischke (1992). They note that the key regressor in many of the articles summarized here, Social Security benefits, is a non-linear function of past earnings, and retirement propensities are clearly correlated with past earnings levels. They solve this problem by using a unique "natural experiment" provided by the end of double-indexing for the "notch generation" that retired in the late 1970s and early 1980s. Workers in this cohort experienced a substantial, unexpected reduction in benefits, yet the dramatic fall in labor force participation continued unabated in this era. This raises important questions about the identification of the cross-sectional literature. However, even with this natural experiment, Krueger and Pischke find significant and sizeable impacts of Social Security accruals on retirement, which highlights the value of the dynamic approach, and suggests that the additional non-linearities that govern the evolution of SSW (as opposed to its level) may be a fruitful source of identification for retirement models.

Each of these dynamic literatures has important limitations. The first suffers from the perhaps untenable assumptions that are required to identify these very complicated structural models; see MaCurdy (1981) for a criticism of this type in the context of general labor supply

models. The second suffers from the limited way in which dynamic retirement incentives are specified. Some of these problems are remedied by the option value literature. But the problem with this approach is that the non-linearities embedded in the assumed utility function make it difficult to test whether financial incentives arising from Social Security and from private pensions have similar effects on retirement.¹ If all dollars of retirement wealth are weighted equally by potential retirees, then pension differences provide a legitimate source of identification of retirement income effects. But if they are not, either because individuals understand their firm's pension incentives better than Social Security incentives, or because the real annuity provided by Social Security is valued differently than the nominal annuity provided by most defined benefit pensions, then it is important to separately estimate Social Security and private pension impacts.

In addition, all of these studies suffer from important data deficiencies, either because the estimates are based on data from the 1970s, when the structure of the Social Security system was fairly different, or data from only a handful of firms, or data without complete information on Social Security incentives. Finally, all of these literatures suffer from a lack of careful attention to the sources of identification of the retirement incentive effects that they estimate. This problem is not necessarily surmounted, and is potentially compounded, by the option value literature, which uses a measure of future incentives that is largely determined by wage differences across individuals and only secondarily influenced by the structure of retirement incentives; we discuss this issue at more length below.

Part II: Data and Empirical Strategy

Data

Our data for this analysis comes from the Health and Retirement Study. The HRS began in 1992 as a survey of individuals then aged 51-61, with re-interviews of these individuals every two years; we use the first five waves of the survey, 1992 through 2000. Spouses of respondents are also interviewed, so the total age range covered by the survey is much wider.

A key feature of the HRS is that it includes Social Security earnings histories from 1951-1991 for most respondents. This provides two advantages for our empirical work. First, it allows us to appropriately calculate benefit entitlements, which depend on the entire history of earnings, by providing a complete and accurate earnings history; some previous studies have been seriously hampered by having to construct earnings histories based on limited self-reported job history information.² Second, it allows us to construct a large sample of person-year observations by using the earnings histories to compute Social Security retirement incentives and labor force participation at each age. We use all person-year observations on men age 55-69 for our analysis, subject to the exclusions detailed below.

Our analysis examines how financial incentives from Social Security and private pensions affect retirement decisions.³ Following the previous literature, we focus on males in this analysis and ignore joint retirement issues, as handling such issues properly is beyond the scope of this study; for more evidence on women's retirement decisions and the spillover effects of incentives on spouses' retirement decisions, see Coile (2004). As we discuss in more detail below, we treat Social Security benefit claiming as occurring simultaneously with retirement.

Our sample is selected conditional on working, so that we examine the incentives for retirement conditional on being in the labor force.⁴ Work is defined in one of two ways. For those person-years before 1992, when we are using earnings histories, we define work as positive earnings in two consecutive years; if earnings are positive this year but zero the next (and if the

year of zero earnings occurs at or after age 55), we consider the person to have retired this year.⁵ For person-years from 1992 onwards, when we have the actual survey responses, we cannot use this earnings-based definition, since we only have earnings at two-year intervals. For this era, we use information on self-reported retirement status and dates of retirement to construct retirement measures.⁶

While these are somewhat different constructs, the hazard rates in the two samples by age are similar, as is illustrated in Figure 1.⁷ Although the hazard rate at the oldest ages becomes noisy for the pre-1992 sample, due to small sizes, the key tendencies in the data, most notably the spikes at age 62 and 65, are present and quite similar in both halves of the sample. Thus, we combine them for precision purposes.⁸ It is important to note that we only consider individuals before their first retirement; if a person who is categorized as retired reenters the labor force, the later observations are not used.⁹

Our core sample consists of all men in the HRS who (a) are born (and have spouses born) in 1922 or later, since the earlier cohort is subject to quite different Social Security benefit rules, (b) have non-missing Social Security earnings history data for the worker and his spouse, and (c) have non-missing information on private pension plan characteristics.¹⁰ The final sample consists of 14,825 observations on 2,467 men. The means of our key variables are shown in Table 1. On average, 7 % of our sample retires each year. The typical man in our sample is 59 years old, is married and 4 years older than his wife, and earns \$30,000 per year with average lifetime earnings of \$24,000 per year. We discuss the other control variables below.

Incentive Variable Calculation - Accrual

Our goal is to measure the retirement incentives inherent in Social Security and private pension systems. The first step in this calculation uses a simulation model we have developed to compute the PIA for any individual at all possible future retirement dates. This process is based on a careful modeling of Social Security benefits rules and has been cross-checked against the Social Security Administrations's ANYPIA model for accuracy. The appropriate actuarial adjustment is applied to the PIA to obtain the monthly benefit entitlement.

The next step is to compute the expected net present discounted value of Social Security Wealth (SSW) associated with each retirement date. Our methodology for doing so is described in Coile and Gruber (2001). For single workers, this is simply a sum of future benefits, discounted by time preference rates and survival probabilities. For married workers it is more complicated, since we must include dependent spouse and survivor benefits (we discuss these in more detail below) and account for the joint likelihood of survival of the worker and dependent. We use a real discount rate of 6% and survival probabilities from the age and sex specific U.S. life tables used in the 1995 annual report of the Social Security Trustees.¹¹

We next compute the other Social Security incentive variables. We initially follow the literature and focus on the accrual, the change in SSW resulting from an additional year of work. There are two routes through which an additional year of work affects SSW. First, the additional year of earnings will be used in the recomputation of Social Security benefits. For workers who have not yet worked 35 years, this replaces a zero in the benefits computation; for workers who have worked 35 years, it may replace a previous low earnings year. So the recomputation raises SSW (or leaves it unchanged). Second, at ages 62 and beyond, the additional year of work implies a delay in claiming; this raises future benefits through the actuarial adjustment, but reduces the number of years of benefit receipt, so the net effect is uncertain. Both of these

factors will affect workers differently, depending on their potential earnings next year, earnings history, mortality prospects (which will vary over time and cohort in our data), family structure, and spouse's earnings. Thus, the net effect of an additional year of work on SSW is theoretically ambiguous and will vary significantly across people.

Computing the accrual and other incentive variables requires projecting the worker's potential earnings next year (or in all future years). It is difficult to estimate such a projection from observed data, since true wage growth is confounded with reductions in intensity of work. We considered a number of different projection methodologies, and found that the best predictive performance was from a model that simply grew real earnings from the last observation by 1% per year, so we use this assumption in our simulations.¹² Since this assumption may be most realistic for movements out of full-time "career" jobs, we have considered the sensitivity of our results to redefining retirement as decreasing weekly hours of work to less than 30, which will often indicate such a job change; our results are quite similar using this alternative definition.

Our Social Security incentive variables incorporate dependent spouse and survivor benefits, since these are important components of SSW. For men with non-working wives or wives whose benefits entitlement is less than one-half of the husband's, these benefits are based on the husband's earnings record. For men whose wives have a larger benefit entitlement on their own, these benefits are based on her record but are also included in SSW. Since a full modeling of the joint retirement decision is beyond the scope of this paper, we simply assume that the wives in this sample who are working will retire at age 62; this seems reasonable, given that the median retirement age is 62 for married women in the HRS who are working at age 50.

For the simulations below, we assume that workers claim Social Security benefits at retirement, or when they become eligible (age 62) if they retire before then. In fact, this is not necessarily true; retirement and claiming are two distinct events, and for certain values of mortality prospects and discount rates it is optimal to delay claiming until some time after retirement, due to the actuarial adjustment of benefits. Coile, Diamond, Gruber, and Josten (2002) investigate this issue in some detail, and they find that a relatively small share of those retiring before age 62 delay claiming past age 62 (about 10%), and that virtually none of those retiring at age 62 or later delay claiming. Given these findings, we choose not to jointly model delayed claiming here. Our incentive measures will therefore slightly overstate any subsidies to continued work, since part of this subsidy will come from delayed claiming that could be obtained without delaying retirement.

We also incorporate private pension incentives into our analysis. The HRS collected detailed pension determination information from employers for some of those with pensions and used this information to create a pension benefits calculator comparable to the PIA simulation model we developed for Social Security. We use the calculated pension benefits at each retirement age to create an analogous set of retirement incentive variables including pensions.

These pension data, unfortunately, have two key weaknesses. First, they are available for only about 60% of the observations who self-report positive pension wealth, and the response patterns appear to be non-random. Among those who report having a pension, men with missing pension data work at smaller firms and have lower retirement rates, less education, lower earnings, and shorter job tenure. To avoid incorporating significant measurement error in our incentive variables, we drop observations with missing pension data, defined as self-reporting positive pension wealth but having no employer-provided pension data.¹³ Second, pension data

was matched at wave 1 for the current job (or last job for those not working) and for past jobs lasting at least 5 years; therefore, the data may misstate incentives if individuals change jobs after wave 1 or if the provisions of the pension plan change over time. As we show below, however, our basic conclusions about the effects of Social Security on retirement are similar if estimated on the full sample.

Table 2 shows the evolution of retirement wealth by age, both with and without pensions. The sample used here is all age-eligible men in the HRS (those born between 1931 and 1941) who allowed their Social Security earnings records to be attached to the survey and are working at age 55; the incentives shown are those projected for the worker at age 55, regardless of his later labor supply decisions. The median SSW rises from \$93,018 at age 55 to a peak of \$100,586 at age 63, then falls to \$92,656 at age 69. The age pattern of accruals demonstrates how the various effects of working an additional year enter in at different ages. From ages 55 to 61, accruals are falling, because the value of the dropout year provision decreases as more zero and low earnings years have been replaced. At ages 62-64, workers may benefit from the delayed claiming effect, whereby an additional year of work increases the actuarial adjustment and raises future benefits, though they are also giving up a year of benefits by continuing to work; with a 6% real discount rate (combined with a rising mortality rate and falling value of dropout year), the actuarial adjustment is less than fair for the median worker starting at age 62.¹⁴ Starting at age 65, the negative accruals become very large, as the delayed retirement credit is far too small to compensate for the value of lost benefits. When pension benefits are included, median retirement wealth at age 55 is substantially higher, \$111,645. Relative to the accruals from Social Security only, median accruals with pensions are somewhat larger before age 62 and somewhat more negative after age 65.

Most importantly for our analysis, there is enormous heterogeneity in accruals, as is also shown in Table 2. The standard deviation in Social Security accruals is substantial, averaging roughly \$1,400 per year. At 62, for example, while the median accrual is near zero, the 10th percentile person has an accrual of -\$1,132, and the 90th percentile person has an accrual of \$1,275; the standard deviation at that age is \$1,539. Once pensions are included, the variation is far greater. It is this sizeable variation that identifies our models.

Incentive Variable Calculation - Forward Looking Measures

As noted earlier, the more recent work on pension incentives and retirement has focused not on accruals, but rather on more forward-looking incentive models that incorporate the entire future path of retirement incentives. This literature highlights an important weakness of the accrual measure. For any given year, as we show in Table 2, a typical worker sees a small positive accrual from additional work. But, by working, that worker is also buying an option on the additional accruals that result from work in future years. Incorporating this option dramatically changes the nature of retirement incentives. This point is emphasized for Social Security in Coile and Gruber (2001), where we document the important differences in single year versus multi-year accruals. For a sizeable minority of workers, accrual patterns are non-monotonic, so that forward looking measures can deliver very different incentives than one year forward accruals. Once pensions are included, the non-monotonicities can be even greater.¹⁵

As noted above, Stock and Wise (1990a) suggested an approach to account for these option values, by contrasting the utility of retiring today versus at the optimal point in the future. Their option value model is based on the individual's indirect utility function over work and leisure:

$$(1) \quad V_t(R) = \sum_{s=t}^{R-1} p_{30t} d^{s-t} (y_s)^\gamma + \sum_{s=R}^T p_{30t} d^{s-t} [k * B_s(R)]^\gamma$$

where R is the retirement date, d is the discount rate, p is the probability of being alive at some future date conditional on being alive today, y is income while working, B is retirement benefits, gamma is a parameter of risk aversion, k is a parameter to account for disutility of labor ($k \geq 1$), and T is maximum life length.

In this model, additional work has three effects. First, it raises total wage earnings, increasing utility. Second, it reduces the number of years over which benefits are received, lowering utility. Third, it may raise or lower the benefit amount, depending on the shape of the benefit function, B(R). The latter two effects are weighted more heavily because of the disutility of labor, which acts as a devaluation of wage income relative to retirement income. The optimal date of retirement is therefore the date where the utility gained from the increase in earnings resulting from additional work is outweighed by the utility lost from the decrease in retirement income. The “option value” is the difference between the indirect utility from retirement at the optimal date, R^* , and the indirect utility from retirement today.

We follow Stock and Wise in assuming values of 1.5 and 0.75 for k and g, respectively. But we found that the fit of our model was much better with a more reasonable assumption for d of 0.06, relative to the very large estimate of 0.25 obtained from their model. This assumption also has the advantage that it is the interest rate which Coile et al. (2002) argue would be needed to justify the high level of early claiming we see in reality, so that it is consistent with our assumption of claiming upon retirement. Our results are not very sensitive to reasonable variations in this assumed discount rate.

This approach to modeling retirement incentives has the important advantage, particularly when considering private pensions, of allowing the individual to be forward looking,

and consider incentives beyond the coming year. While theoretically attractive, however, implementation of the option value model runs into two important difficulties in a retirement regression context. First, the vast majority of the variation in the option value derives from the variation in wages. Indeed, in our HRS sample, a set of age dummies plus a quartic in earnings alone explains 66% of the variation in option value. This potentially poses problems for the option value measure if the goal of the empirical exercise is to measure the impact of Social Security policy changes on retirement behavior. If, for example, wages are correlated in some way with underlying tastes for retirement (e.g. high wage individuals are those motivated individuals with tastes for continued work, even conditional on wage), then variation in wages does not provide a legitimate source of identification for learning about retirement income effects.

In principle, this problem can be surmounted by structural estimation of the option value model, which will identify the difference in the impacts of wages and retirement income on retirement decisions, through the value of leisure parameter. But, in practice, this is only true if the particular utility structure is correct, for example if the additional leisure of utility enters the model only as a multiplier on post-retirement income and not in some other way.

The second limitation of the option value approach is that because of its non-linear structure, there is no straightforward decomposition of the measure that could be used to assess whether individuals respond similarly to financial incentives from different retirement income sources. Testing for a differential response is necessary in order to evaluate the separate contributions of Social Security and private pensions to retirement decisions.

We take two approaches to addressing these potential shortcomings with the option value model. One is to include a set of flexible controls for earnings directly in the model, in order to

capture the heterogeneity that may bias these estimates. This is imperfect, however, since wages enter highly non-linearly in the option value model, and the form of heterogeneity is unknown, so that even rich wage controls may not fully capture the underlying correspondence between option value and tastes for work.

The second is to construct a measure which incorporates the insights of the option value measure, but focuses only on the variation in retirement income and not wages, and which can be decomposed easily into Social Security and pension components. We create a forward-looking measure of incentives which we call “peak value”.¹⁶ This is comparable to the accrual, but looks forward more than just one year: it calculates the difference between SSW at its *maximum expected value* and SSW at today’s value, to measure the incentive to continued work. In this way, the peak value appropriately considers the tradeoff between retiring today and working to a period with much higher SSW, thereby capturing the option value of continued work even before Social Security entitlement ages are reached.¹⁷¹⁸ If the individual is at an age that is beyond the SSW optimum, then the peak value is the difference between retirement this year and next year, which is exactly the accrual rate. Since PV does not include wages directly (only indirectly, insofar as wages affect Social Security and pension accruals), there is much more variation from the structure of the Social Security entitlement; an earnings quartic and age dummies explain only 16% of the variation in peak value.

The tradeoffs between using the option value and peak value measures parallel the tradeoffs between more reduced form and structural modeling in general. On the one hand, if the option value model is “right,” it is a richer specification that more should accurately reflect the incentives inherent in retirement programs. On the other hand, if the specification is wrong (e.g. if the parameterization chosen is wrong, or if the findings are driven by wage heterogeneity and

not retirement incentive effects), then peak value may give a better approximation to incentive effects. We therefore use both measures below and discuss their relative performance. But, in either case, one clear advantage of the peak value approach is that its linear structure allows for a straightforward decomposition into Social Security and private pension incentive effects.

Regression Framework

In a standard retirement model, Social Security and private pensions will play two roles in the decision whether to retire this year or to continue working. The first is through wealth effects: higher retirement wealth (RW) will induce individuals to consume more of all goods, including leisure, and to retire earlier. The second is through accrual effects: the individual's decision to continue to work is a function of the increase in retirement consumption resulting from an additional year of work, relative to the value of an additional year of leisure.

Following this discussion, we use the incentive variables described above to run regressions of the form:

$$(2) \quad R_{it} = \beta_0 + \beta_1 RW_{it} + \beta_2 INCENT_{it} + \beta_3 X_{it} + \beta_4 AGE_{it} + \beta_5 EARN_{it} + \beta_6 AIME_{it} \\ + \beta_7 MAR_{it} + \beta_8 AGEDIFF_{it} + \beta_9 SPEARN_{it} + \beta_{10} SPAIME_{it} + \beta_{11} Y_t + \varepsilon_{it}$$

where R is a dummy variable equal to one if the worker retires during the year, RW is the expected PDV of retirement benefits that is available to the person if he retires that year (t); INCENT is one of the incentive measures noted above (accrual, option value, peak value); X is a vector of control variables that may importantly influence the retirement decision but do not enter directly into the calculation of RW (education, race, veteran status, born in the U.S., region of residence, experience in the labor market and its square, tenure at the firm and its square, 13 major industry dummies, 17 major occupation dummies); AGE is a set of dummies for each age

55-69; EARN is a control for potential earnings in the next year; AIME is a control for average monthly lifetime earnings as of period t ; MAR is a dummy for marital status; AGEDIFF controls for the age difference with the spouse; SPEARN and SPAIME are the spouse's next year and average lifetime earnings; and Y is a series of year dummies. Since our dependent variable is dichotomous, we estimate the model as a probit. We have also estimated these models as Cox proportional hazard models and the results were very similar; this is not surprising, given that the models all include a full set of age dummies, which pick up the same factors captured by the baseline in the hazard model.

This model parallels the types of models used in the first round of research on Social Security and retirement, with two important exceptions. The first is the use of age dummies, rather than linear age. The disadvantage of using age dummies is that they will capture some of the important variation in incentive measures at key ages such as 62 and 65. But the advantage is that we do not impose the assumption that the taste for leisure changes linearly over time. Our results clearly indicate that age dummies are preferred to linear age, so we include them.

The second difference is earnings controls. Most articles in this literature did not control for earnings, and no articles controlled for both earnings around retirement and average lifetime earnings. Yet both of these variables are clearly important determinants of both Social Security incentives and retirement decisions, so excluding them from the model imparts a potential omitted variables bias. Moreover, there is no reason to suspect that heterogeneity is a purely linear function of earnings. Thus, for each of the earnings controls above, we include squared, cubed, and quartic terms as well. Moreover, it is possible that heterogeneity in retirement is also related to the relationship between current and average lifetime earnings; we therefore include as well a full set of interactions between the EARN and AIME quartics to reflect this.

Part III: Results

Retirement Incentives and Retirement Decisions

In Table 3 we report marginal effects from probit estimation of equation (2) for each of our three incentive measures, incorporating both Social Security and private pensions.¹⁹ Peak value, accrual, and retirement wealth (RW) are expressed in units of \$100,000; option value is expressed in units of 10,000. We also report the implied percentage point impact of a one standard deviation increase in the variable in square brackets.

For the accrual model, we estimate a positive and significant impact of retirement wealth levels, as expected. The coefficient implies that a one standard deviation increase in RW increases the probability of retirement by 0.96 percentage points, or 14% of the sample average retirement rate of 6.9%; evaluated at the mean, this corresponds to an elasticity of non-participation with respect to benefits of 0.16. The coefficient on the accrual is right signed, but insignificant, and implies that a one standard deviation increase in ACC decreases the probability of retirement by 0.12 percentage points or 2% of the sample average retirement rate; the elasticity is -0.003. This suggests that there is only a modest impact of one-year forward incentives on retirement decisions. This could reflect the fact that individuals are not at all forward-looking in their decisions. Alternatively, given non-linearities in future accruals, it could represent the fact that individuals are not considering solely the accrual to the next year but the entire future path of incentives.

This possibility is addressed in the next two columns, which show the estimates from the peak value and option value models. For both models, we once again estimate a positive and significant coefficient on RW; the coefficient is quite similar in the PV model and falls in the OV model. But, in both cases, we now estimate significant negative impacts of the forward looking

incentive measures for retirement decisions. We find that a one standard deviation increase in the peak value lowers retirement by 1.1 percentage points, or 16% of the sample average retirement rate; this corresponds to an elasticity of non-participation with respect to benefits of -0.07. For option value, we find the largest effects of all, with a one standard deviation increase in option value lowering retirement by 2.6 percentage points, or 38%.

The log likelihood values for all three models are also shown Table 3. We find that there is a substantial improvement in model fit moving from the accrual to PV models. There is then a further improvement moving to the OV model. These findings suggest that the forward-looking models of the type advocated by Stock and Wise are very important for explaining retirement behavior. Individuals do appear to recognize the future path of SSW accumulation, and take this into account in making their retirement decisions. They also highlight the tradeoff raised earlier: the OV model provides better explanatory power, but it potentially does so through exploiting individual heterogeneity in wages, and not true incentive effects.

The other variables in the regression have their expected impacts. Retirement propensity rises with age, with particularly large effects at ages 62 and 65. The set of higher order earnings controls is statistically significant at the 10% level (slightly less so in the OV model, since that measure more directly incorporates variation in earnings), although re-estimating the models with only linear age controls does not yield statistically different results. Being married and having a larger age difference with one's wife decrease the probability of retirement, though only the former is significant. More experience lowers the odds of retirement, conditional on age, but this relationship is decreasing in absolute value. There is no distinct relationship with tenure, although there is a very significant positive impact of being in the 6% of the sample with missing tenure data; this is consistent with lower labor force attachment among those in jobs of short

duration. The industry and occupation dummies, not shown, do not show a particularly strong pattern, with the exception of higher retirement rates in the armed forces and the cleaning and building services occupation. There is no significant time pattern to retirement behavior, which is consistent with Quinn (1999), who shows that the strong time series trend towards earlier retirement was arrested beginning in the mid-1980s.²⁰ There is no strong regional pattern, other than a significantly higher retirement rate in the western pacific region and a lower rate in New England.²¹

Social Security vs. Pension Incentives

As highlighted earlier, a limitation of earlier work has been its inability to separate the responsiveness of retirement decisions to Social Security and private pension incentives. The one successful attempt to do so, Diamond and Hausman (1984), suggested that retirement was actually more responsive to Social Security than to private pensions. This is a somewhat surprising result, since the accrual patterns in private pensions are generally sharper than for Social Security (Samwick, 1998). Moreover, it might seem that workers are more well-informed about the details of their private retirement plans than about Social Security. On the other hand, however, Social Security provides a real annuity that may be valued more than the nominal annuities provided by many defined benefit pension plans. And the uniformity of the system across everyone, not just coworkers, may actually make information more readily accessible.

We assess this issue by splitting our accrual and PV measures, along with the associated PDVs, into their Social Security and pension components. For accrual, and for PDV, this is straightforward. For PV, it requires separately assessing the peak for each retirement income stream, and computing the difference between the PDV of each stream from the perspective of

today and from the perspective of that peak. As highlighted earlier, this computation is not really feasible for our reduced-form version of OV, since by its non-linear nature it cannot be easily decomposed.

The results of this exercise are shown in columns 1 and 5 of Table 4. For the accrual regression, there are larger coefficients on both Social Security accrual and wealth than for pensions, although neither accrual coefficient is significant. For the PV regression, there is a highly significant coefficient on PV for both Social Security and private pensions, and the SSPV coefficient is roughly four times as large as that for pensions. The Social Security wealth (SSW) coefficient is also about five times as large as the pension wealth (PW) coefficient, but the SSW coefficient itself is not significant; this reflects the fact that SSW levels are hard to identify in models that control for lifetime earnings.²²

At the same time, it is also true that there is much more variation in pension PV than in Social Security PV. Thus, when we evaluate the impact of a one standard deviation change in incentives in square brackets, the implied effects are almost identical for Social Security and for pensions: in each case, a one standard deviation increase in PV leads to a reduction in retirement probabilities of about 1 percentage point, or 14% of baseline retirement probability. Thus, it appears that individuals are roughly equally responsive to comparable changes in Social Security and pension incentives.

One further issue with this approach is the treatment of defined benefit (DB) vs. defined contribution (DC) plans. While DB pensions feature strong incentives to work to particular entitlement ages, and often disincentives to working beyond those ages, DC plans simply accumulate based on contribution rates and employer matches. In the construction of our incentive measures, we follow the HRS pension calculator in assuming that voluntary employee

contributions occur at the same rate as in the past and that real DC balances grow at 2% annually. Appendix Table 1 shows the peak age distribution for Social Security, DB pensions, and DC pensions. Both DB and DC pensions attain their peak for about 10% of the sample at each age from 62 to 66; DB plans have relatively more mass before 62 and DC plans have relatively more mass after 66, including a large mass at 70 (the last possible retirement age in the analysis). The pattern is of course quite different for Social Security, with virtually all of the mass at ages 62 through 65.

The similarity of DB and DC pension effects is reflected in columns 2 and 6 of Table 4, which show a model that breaks out accrual and peak value into Social Security, DB, and DC components. Once again, the Social Security component is much larger than either of the two pension components. And, the two pension components themselves have quite similar coefficients (although there is a much larger effect of DB wealth than of DC wealth). In this case, the variation in pension wealth is not as large as when the two types of pensions are combined, so that the implied effect of a one standard deviation change in PV is only about 0.6 percentage points, or 9% of baseline retirement rates, which is smaller than the implied Social Security effect.

A final issue is sample selection. As noted earlier, the set of persons with non-missing pension information from HRS is not a random sample. Thus, in columns 4 and 8 of Table 4, we show the accrual and PV results using Social Security incentives only, excluding pension information, for the full sample. The cost of this approach is the potential omitted variables from excluding pension incentives, but the benefit is that the sample is now representative. For this full sample, our estimated coefficients are smaller than the Social Security coefficients from the other specifications and not statistically significant at traditional levels, but the pattern of results

is very similar. Comparing these results to those in columns 3 and 7, which show the effect of using Social Security incentives only for the smaller sample of 14,825 observations, we can see that the smaller coefficients arise primarily from the use of a different sample rather than from the omission of pension incentives.

Part IV: Conclusion

The Social Security program is the most important source of retirement income support for older Americans. As such, it is possible that the incentives embodied in this system for continued work or retirement at various ages are a critical determinant of retirement decisions. Understanding the influence that Social Security has on retirement decisions is particularly important now, as any of the proposed reforms to the Social Security system will change the structure of the program in a manner that has important impacts on retirement incentives.

Our paper has used the richest available current data, the Health and Retirement Study, to provide new evidence on the impact of Social Security on retirement. We have drawn on the insights of the option value approach of Stock and Wise, but we have augmented this with a new measure, peak value, which both is more strictly identified by retirement incentive changes rather than variation in wages, and which allows for a straightforward decomposition into Social Security and private pension incentive effects. We also include a much richer set of earnings controls than the previous literature, to ensure that our estimates capture legitimate effects of retirement incentives on retirement and are not biased by possible non-linear effects of earnings on retirement.

We find that future entitlements do matter; both the peak value and option value models perform much better than an accrual only model. We find that a one standard deviation change

in incentives has similar impacts for Social Security and private pensions. These results suggest that changes to the structure of Social Security can have important effects on retirement decisions. Indeed, in Coile and Gruber (2004), we extend this analysis to simulate the effect of several policy changes, such as raising the early entitlement and full benefit ages, and conclude that these changes could have significant impacts on retirement decisions.

But our modeling comes quite far from “explaining” retirement decisions. There are still enormous spikes in the retirement hazard at ages 62 and 65 which are not explained by the incentive measures we include in this regression. Clearly, much more work is needed to more fully understand retirement behavior, incorporating both individual (supply-side) and employer (demand-side) decision-making.

References

- Blau, David M.. "Labor Force Dynamics of Older Men," *Econometrica* 62 (January 1994), 117-156.
- Brown, Jeffrey R., Jeffrey B. Liebman, and Joshua Pollet. "Appendix: Estimating Life Tables That Reflect Socioeconomic Differences in Mortality," in *The Distributional Aspects of Social Security and Social Security Reform*, Martin Feldstein and Jeffrey B. Liebman, eds. Chicago: University of Chicago Press (2002).
- Burtless, Gary. "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement," *Review of Economic Studies* 53 (October 1986), 781-805.
- Burtless, Gary and Robert Moffitt. "The Effect of Social Security Benefits on the Labor Supply of the Aged," in *Retirement and Economic Behavior*, H. Aaron and G. Burtless, eds. Washington: Brookings Institution (1984).
- Coile, Courtney. "Retirement Incentives and Couples' Retirement Decisions," Topics in Economic Analysis & Policy 4 (2004), article 17.
- Coile, Courtney, Peter Diamond, Jonathan Gruber, Alain Jouten. "Delays in Claiming Social Security Benefits," *Journal of Public Economics* 84 (June 2002), 357-385.
- Coile, Courtney and Jonathan Gruber. "The Effect of Social Security on Retirement in the United States," in *Social Security Programs and Retirement Around the World: Micro-Estimation*, Jonathan Gruber and David A. Wise, eds. Chicago: University of Chicago Press (2004).
- Coile, Courtney and Jonathan Gruber. "Social Security Incentives for Retirement," in *Themes in the Economics of Aging*, David A. Wise, ed. Chicago: University of Chicago Press (2001).

- Coile, Courtney and Jonathan Gruber (2000). "Social Security and Retirement," National Bureau of Economic Research Working Paper 7830.
- Diamond, Peter and Jerry Hausman. "Retirement and Unemployment Behavior of Older Men," in *Retirement and Economic Behavior*, H. Aaron and G. Burtless, eds. Washington: Brookings Institution (1984).
- Fields, Gary S. and Olivia S. Mitchell. "Economic Determinants of the Optimal Retirement Age: An Empirical Investigation," *Journal of Human Resources* 19 (Spring 1984), 245-262.
- Friedberg, Leora, and Anthony Webb. "Retirement and the Evolution of the Pension Structure," *Journal of Human Resources* 40 (Spring 2005), 281-308.
- Gruber, Jonathan and David Wise. *Social Security Programs and Retirement Around the World: Micro-Estimation*. Chicago: University of Chicago Press (2004).
- Gustman, Alan L. and Thomas L. Steinmeier. "Retirement and Wealth," *Social Security Bulletin* (September 2002), 66-91.
- Gustman, Alan L. and Thomas L. Steinmeier. "A Structural Retirement Model," *Econometrica* 54 (May 1986), 555-584.
- Gustman, Alan L. and Thomas L. Steinmeier. "The 1983 Social Security Reforms and Labor Supply Adjustments of Older Individuals in the Long Run," *Journal of Labor Economics* 3 (April 1985), 237-253.
- Haider, Steven and Gary Solon. "Nonrandom Selection in the HRS Social Security Earnings Sample." RAND Labor and Population Program Paper 00-01 (2000).
- Hausman, Jerry A. and David A. Wise. "Social Security, Health Status, and Retirement," in David A. Wise, ed., *Pensions, Labor, and Individual Choice*. Chicago: University of Chicago Press (1985).

- Hurd, Michael D. "Research on the Elderly: Economic Status, Retirement, and Consumption and Saving," *Journal of Economic Literature* 28 (June 1990), 565-637.
- Krueger, Alan B. and Jorn-Steffan Pischke. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation," *Journal of Labor Economics* 10 (October 1992), 412-437.
- Lumsdaine, Robin and David Wise. "Aging and Labor Force Participation: A Review of Trends and Explanations," in *Aging in the United States and Japan: Economic Trends*, Y. Noguchi and D. Wise, eds. Chicago: University of Chicago Press (1994).
- MaCurdy, Thomas E. "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy* 89 (December 1981), 1059-1085.
- Mitchell, Olivia S. and Gary S. Fields. "The Effects of Pensions and Earnings on Retirement," in *Research in Labor Economics, Volume 5*. Greenwich, Connecticut: JAI Press (1982).
- Panis, Constantijn, Michael Hurd, David Loughran, Julie Zissimopoulos, Steven Haider, and Patricia StClair. "The Effects of Changing Social Security Administration's Early Entitlement Age and the Normal Retirement Age," report prepared for the Social Security Administration (2002).
- Quinn, Joseph. "Retirement Patterns and Bridge Jobs in the 1990s," *Policy Brief for the Employee Benefit Research Institute*. Washington, DC: EBRI (1999).
- Ruhm, Christopher. "Secular Changes in the Work and Retirement Patterns of Older Men," *Journal of Human Resources* 30 (Spring 1995), 362-385.
- Rust, John and Christopher Phelan. "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica* 65 (July 1997), 781-831.

- Samwick, Andrew A. “New Evidence on Pensions, Social Security, and the Timing of Retirement,” *Journal of Public Economics* 70 (November 1998), 207-236.
- Stock, James H. and David A. Wise. “Pensions, the Option Value of Work, and Retirement,” *Econometrica* 58 (September 1990a), 1151-1180.
- Stock, James H. and David A. Wise. “The Pension Inducement to Retire: An Option Value Analysis,” in David A. Wise, ed., *Issues in the Economics of Aging*. Chicago: University of Chicago Press (1990b).
- Sueyoshi, Glenn T. “Social Security and the Determinants of Full and Partial Retirement: A Competing Risks Analysis,” NBER Working Paper #3113 (1989).

Figure 1: Empirical Retirement Hazard Function




Table 1: Summary Statistics for HRS Person-Year Sample

Variable	Mean	Standard Deviation
Retired	0.069	0.253
Age	58.8	3.1
Married	0.902	0.298
Age Difference with Wife	4.2	5.0
Black	0.109	0.312
Other Nonwhite	0.095	0.293
Educ: <12 yrs	0.243	0.429
Educ: 12 yrs	0.353	0.478
Educ: 13-15 yrs	0.141	0.348
Earnings	29,384	17,725
AIME	2,061	850
Spouse's earnings	9,325	12,590
Spouse's AIME	531	580
Experience	40.5	4.6
Job tenure	17.1	12.7
Missing job tenure	0.074	0.262
Veteran	0.616	0.486
Born in US	0.913	0.282
Number of Obs	14,825	

Table 2: Retirement Incentives by Age, Age-Eligible Men in HRS

Age	SSW Median	SSW Accrual Median	SSW Accrual Std. Dev.	SS+Pens Wealth Median	SS+Pens Accrual Median	SS+Pens Accrual Std. Dev.
55	93,018	1,164	1,084	111,645	1,283	9,568
56	94,529	1,054	2,251	113,457	1,225	6,541
57	95,740	938	1,954	114,706	1,067	7,405
58	96,917	828	1,174	116,927	948	7,565
59	98,207	741	1,684	118,023	832	9,655
60	99,170	659	899	119,190	742	7,664
61	99,790	647	1,414	120,546	702	6,178
62	100,500	-114	1,539	121,435	-41	5,091
63	100,586	-31	845	121,470	-55	4,114
64	100,423	-514	804	121,672	-568	3,895
65	100,072	-1,223	1,036	121,421	-1,423	4,333
66	98,604	-1,474	1,212	119,589	-1,776	4,238
67	97,093	-1,748	1,361	118,133	-2,066	4,476
68	95,067	-1,947	1,842	115,983	-2,309	4,547
69	92,656	-2,109	1,573	113,538	-2,469	4,480

Notes:

1. The sample is all age-eligible men in the HRS (born between 1931-1941) who allowed their Social Security records to be attached to the survey and who were working at age 55; N=2272 observations.
2. The table shows the incentives that are projected at age 55 for all individuals, regardless of their later labor supply decisions.

Table 3: Retirement Probits (Marginal Effects)

Variable	Incentive Measure					
	Accrual		Peak Value		Option Value	
Incentive Measure	-0.0151	(0.0217)	-0.0394	(0.0093)	-0.0303	(0.0050)
Effect of 1 Std Dev Inc.		[-0.00118]		[-0.01092]		[-0.02642]
Retirement Wealth	0.0052	(0.0010)	0.0055	(0.0010)	0.0021	(0.0011)
Effect of 1 Std Dev Inc.		[0.00962]		[0.01025]		[0.00391]
Age56	-0.0178	(0.0061)	-0.0175	(0.0061)	-0.0185	(0.0059)
Age57	-0.0100	(0.0072)	-0.0101	(0.0072)	-0.0131	(0.0068)
Age58	-0.0027	(0.0087)	-0.0030	(0.0086)	-0.0082	(0.0080)
Age59	-0.0097	(0.0090)	-0.0106	(0.0088)	-0.0162	(0.0080)
Age60	0.0245	(0.0144)	0.0228	(0.0141)	0.0111	(0.0125)
Age61	0.0343	(0.0175)	0.0320	(0.0171)	0.0158	(0.0147)
Age62	0.1930	(0.0368)	0.1866	(0.0363)	0.1457	(0.0332)
Age63	0.0927	(0.0309)	0.0870	(0.0301)	0.0551	(0.0254)
Age64	0.0850	(0.0333)	0.0786	(0.0322)	0.0450	(0.0264)
Age65	0.2310	(0.0551)	0.2188	(0.0540)	0.1583	(0.0480)
Age66	0.0674	(0.0381)	0.0609	(0.0365)	0.0258	(0.0279)
Age67	0.0596	(0.0409)	0.0540	(0.0393)	0.0214	(0.0297)
Age68	0.0356	(0.0405)	0.0300	(0.0383)	0.0029	(0.0276)
Age69	0.0724	(0.0557)	0.0639	(0.0528)	0.0307	(0.0410)
Educ: Less than HS	0.0132	(0.0130)	0.0143	(0.0130)	0.0132	(0.0129)
Educ: HS grad only	0.0156	(0.0089)	0.0163	(0.0089)	0.0158	(0.0088)
Educ: Some college	0.0109	(0.0078)	0.0115	(0.0078)	0.0111	(0.0077)
Marital Status	-0.0298	(0.0103)	-0.0288	(0.0102)	-0.0270	(0.0100)
Age Diff w/ Spouse	-0.0004	(0.0004)	-0.0003	(0.0004)	-0.0003	(0.0004)
Black	0.0034	(0.0063)	0.0037	(0.0063)	0.0032	(0.0062)
Other Non-White	-0.0075	(0.0073)	-0.0079	(0.0072)	-0.0077	(0.0072)
Veteran	0.0035	(0.0039)	0.0038	(0.0039)	0.0035	(0.0039)
Born in US	0.0001	(0.0081)	-0.0014	(0.0082)	-0.0014	(0.0081)
Experience	-0.0184	(0.0064)	-0.0210	(0.0063)	-0.0198	(0.0063)
Experience Squared	0.0002	(0.0001)	0.0002	(0.0001)	0.0002	(0.0001)
Job Tenure	0.0017	(0.0006)	0.0018	(0.0006)	0.0015	(0.0006)
Job Tenure Squared	0.0000	(0.0000)	0.0000	(0.0000)	0.0000	(0.0000)
Job Tenure Missing	0.0785	(0.0152)	0.0750	(0.0149)	0.0729	(0.0147)
Number of Obs	14,825		14,825		14,825	
Log Likelihood Value	-3,221.4		-3,211.5		-3,203.4	

Notes:

- Standard errors are in parentheses, marginal effects of a one standard deviation increase are in brackets.
- Probits also include interacted quartics of current and lifetime earnings, industry and occupation dummies, and region and year dummies (not shown; available from authors on request).

Table 4: Retirement Probits with Multiple Peaks (Marginal Effects)

Variable	Accrual Specifications				Peak Value Specifications			
	SS+Pens (1)	SS, DB, DC (2)	SS (3)	SS, All Obs (4)	SS+Pens (5)	SS, DB, DC (6)	SS (7)	SS, All Obs (8)
Incentive Measures								
Social Security	-0.0880 (0.1341) [-0.00196]	-0.0855 (0.1334) [-0.00191]	-0.0891 (0.1328) [-0.00222]	-0.0654 (0.0813) [-0.00163]	-0.1469 (0.0583) [-0.00966]	-0.1416 (0.0582) [-0.00931]	-0.1575 (0.0328) [-0.01116]	-0.0974 (0.0568) [-0.00690]
Pensions (DB+DC)	-0.0120 (0.0220) [-0.00088]				-0.0362 (0.0092) [-0.00954]			
DB Pensions		0.0135 (0.0229) [0.00090]				-0.0312 (0.0104) [-0.00650]		
DC Pensions		-0.1827 (0.1056) [-0.00534]				-0.0367 (0.0198) [-0.00503]		
Retirement Wealth								
Social Security	0.0305 (0.0156) [0.01374]	0.0316 (0.0156) [0.01422]	0.0302 (0.0185) [0.01342]	0.0205 (0.0144) [0.00910]	0.0245 (0.0155) [0.01103]	0.0257 (0.0155) [0.01156]	0.0242 (0.0159) [0.01075]	0.0141 (0.0139) [0.00625]
Pensions (all)	0.0051 (0.0010) [0.00880]				0.0054 (0.0010) [0.00940]			
DB Pensions		0.0083 (0.0015) [0.01100]				0.0079 (0.0015) [0.01049]		
DC Pensions		0.0026 (0.0018) [0.00240]				0.0026 (0.0018) [0.00235]		
# of Obs	14,825	14,825	14,825	19,871	14,825	14,825	14,825	19,871
Log Likelihood Value	-3,219.7	-3,212.9	-3,232.1	-4,287.5	-3,207.6	-3,204.9	-3,228.5	-4,284.9

Notes:

1. Standard errors are in parentheses; marginal effects of a one standard deviation change are in brackets.
2. Probits also include interacted quartics of current and lifetime earnings, industry and occupation dummies region and year dummies, and same demographic variables as models in Table 3.

Appendix Table 1: Peak Age Distribution

Age	% of Sample with Peak by Age			
	SS	All Pens	DB	DC
55	0.0	2.3	3.2	0.1
56	0.0	2.6	3.7	0.3
57	0.0	3.6	4.6	0.8
58	0.0	3.9	5.0	0.8
59	0.0	4.4	5.6	1.2
60	0.0	5.4	7.4	1.9
61	0.0	6.9	7.6	6.9
62	35.5	9.6	10.0	9.3
63	11.1	10.7	12.0	7.2
64	26.0	7.6	8.1	7.6
65	17.1	12.1	12.7	11.0
66	3.5	12.0	13.9	7.0
67	2.7	3.0	2.9	3.7
68	1.7	2.7	1.1	6.5
69	1.4	2.2	1.1	5.0
70	1.1	11.1	1.2	30.6
#of Obs	13,690	6,107	4,592	2,631

Source: Authors' calculations; each category includes observations with non-zero peak value only.

Endnotes

¹ Theoretically, it would be possible to estimate a structural option value model that allowed a dollar of retirement income from Social Security and from pensions to have a different effect on utility; we are not aware of any previous studies that have done this.

² Only earnings since 1950 are required to compute SS benefits for persons in our sample; the benefit rules specify that a shorter averaging period (less than the standard 35 years) is used for persons born prior to 1929, who might have started working before 1951. Earnings histories for HRS respondents are not available after 1991 due to the terms of the agreement between the HRS and the Social Security Administration. Below we detail how we project earnings for 1992 and beyond in calculating retirement incentives.

³ Although our work is focused on the labor force participation decision, a separate and interesting issue is the impact of retirement income on the marginal labor supply decision among those participating in the labor force. This is more complicated for those around retirement age, since it involves incorporating the role of the earnings test, which we avoid with our analysis of participation. This, in turn, would involve modeling expectations about the earnings test, since individuals appear not to understand that this is just a benefits delay instead of a benefits cut. This is clearly a fruitful avenue for further research.

⁴ We acknowledge that this introduces the possibility of sample selection bias. For example, it is possible that people who retire early are less responsive to financial incentives than the population as a whole; if true, this would lead to an overestimate of the effect of incentives on

retirement. However, if most early retirees leave the work force at an early age because of other factors, such as work-limiting disability, and not because they are insensitive to financial incentives, then the sample selection bias will be quite small.

⁵ One possible problem with using earnings histories to define retirement is that an individual may move from the private sector to the state and local government sector, in which case he would be classified as retired when in fact he is still working. We check for this by dropping all individuals who list their industry as public administration and find that the results are similar.

⁶ If an individual simultaneously reports his labor supply status as working and retired, we treat him as working.

⁷ Another difference between the samples is that work and retirement definitions are based on behavior during the calendar year for the 1980-1991 observations and on behavior between birthdays (e.g., May 1992 to May 1993) for the 1992-1999 observations. We prefer the latter method, but cannot apply it to the earlier period due to lack of data on month of retirement.

⁸ We have estimated our regressions separately using pre-1992 and post-1992 observations and find no clear pattern of different results across the two periods. We have also experimented with an alternative retirement definition for the post-1992 data that is based on self-reports of earnings every other year supplemented with self-reported information on the date of retirement; the results using this alternative are very similar to our basic results.

⁹ We think it is appropriate to focus on the first labor force exit after age 55, as this is likely to reflect retirement from a “career job” for workers who have such a job. Although examining the effect of retirement incentives on the decision to re-enter the labor force after age 55 or to exit after having re-entered is beyond the scope of our study, we suspect that retirement incentives play a much smaller role in these decisions, particularly in the case of pensions, since the “bridge jobs” that older workers commonly take upon re-entering the workforce are unlikely to have pension benefits. We acknowledge that some of what we classify as retirement may reflect workers who are changing jobs after age 55 and spend some period of time out of the labor force, but unfortunately it is very difficult to distinguish this from actual retirement.

¹⁰ Starting with 5,886 men, 127 are lost for pre-1922 birth years, 1,981 for missing SS data (Haider and Solon, 2000, show that instances of missing SS data are essentially random), 500 for retirement prior to age 55, and 811 for missing pension data. We consider a person to have missing pension data if they self-report pension wealth greater than zero (as calculated by Steve Venti, to whom we express our appreciation for sharing his calculations) but have no firm-provided pension data.

¹¹ We have also tried incorporating differential mortality in our analysis. Specifically, we used relative mortality factors calculated in Brown et. al. (2002) to obtain survival probabilities that vary by race/ethnicity and education as well as age, cohort, and sex. We recalculate the incentive measures using these new survival probabilities and re-estimate our regressions with these incentive measures. The results of this exercise were extremely similar to the results

reported in Tables 3 and 4. As we discuss below, our results also were not very sensitive to the choice of discount rate.

¹² Some of the alternative projection methodologies we used include assuming zero real wage growth, estimating cohort average and individual age-earnings profiles, and varying the earnings base for the projections that assume constant wage growth (e.g., use average earnings over the past three years instead of last year's earnings). Following the previous literature, we do not include uncertainty about future earnings in our analysis; however, this is clearly a fruitful area for future research.

¹³ Pension data is more likely to be missing from previous jobs than from current jobs. We acknowledge that some individuals with missing data from previous jobs may be included in our sample (if they have non-missing data for a pension on a current job). This will affect our estimates of the effect of the PDV of retirement wealth on retirement, but not our estimates of the incentive measures that capture the change in retirement wealth with additional work.

¹⁴ Coile and Gruber (2001), using a 3% real discount rate, find that there is a large subsidy to work at age 62.

¹⁵ The age pattern of accruals in defined benefit pension plans in our data is similar to what has been found in previous work. We find that the most common pattern for DB plan accruals is an "M" shape with two sharp peaks. The median ages for the two peaks are 58 and 63.5. The increase in median DB wealth is \$13,200 at the 1st peak and \$20,200 at the 2nd peak, which

corresponds to 35% and 55% of median annual earnings, respectively. Samwick (1998) finds that for a typical DB plan, the early and normal retirement ages are 55 and 60 and the accruals are 1.8 and 0.5 times annual earnings. Differences between our calculations and previous results could reflect our use of a higher discount rate and our use of the median of many plans rather than a single plan.

¹⁶ Since the release of an earlier version of this paper (Coile and Gruber, 2000), several other papers have used our peak value construct, including: Friedberg and Webb (2005), who delve more deeply into the role of private pensions; Gruber and Wise (2004), who expand this analysis to other developed countries; Panis et. al. (2002), who use the methodology to simulate changes in Social Security rules; and Gustman and Steinmeier (2002), who develop a related concept they call “premium value.”

¹⁷ With Social Security and DB pensions, retirement wealth typically peaks at some point in the worker’s 60s and falls thereafter. With DC pensions, wealth may continue to accumulate indefinitely with continued work if the combined contribution and accumulation rate exceeds the discount rate. We treat 70 as the last possible retirement age, so a DC plan with continuously rising wealth will reach its peak at age 70 by construction; we note that the worker’s total retirement wealth will often peak before 70 even if the DC wealth is still rising. The choice of a discount rate will be particularly important with regards to DC plans; however, as we have previously noted, the results of our analysis are not particularly sensitive to our choice of a discount rate. For an alternative way to treat DC accruals, see Gustman and Steinmeier (2002), whose “premium value” excludes accruals in DC plans that occur at a steady rate.

¹⁸ One potential problem with the PV measure is that it treats a given gain in retirement wealth the same, regardless of how many years of work are required to obtain this gain. We have experimented with dividing PV by the number of years to the peak and find that this measure is also a statistically significant determinant of retirement.

¹⁹ Standard errors are clustered at the person-level to correct for repeated observations on the same individual.

²⁰ We also find no significant time trend in the age dummies: interactions of age dummies with year dummies or a linear time trend are insignificant. Thus there is no evidence that the “social norms” that underlie age dummy effects are changing over time.

²¹ One limitation of these models is that they exclude controls for potentially important alternative determinants of retirement, such as health and wealth. The problem with these controls is that they are also likely endogenous with respect to the work decision, which is why they are excluded in our baseline specification; they are also unavailable prior to 1992. However, we have estimated models for the post-1992 period with and without these controls and our basic results are unchanged by their inclusion; these results are available upon request.

²² One possible explanation for this finding is outliers. There is a much more skewed distribution of incentives under private pensions than under Social Security, and a few outlying very large accruals or PV could skew the estimation. We therefore re-estimated the models

censoring the accrual and PV at the 99th percentiles of the distribution, and also the 95th percentiles; this raised the estimated pension effects somewhat, but also raised the Social Security effects as well, so that the basic pattern of findings was unchanged.