This article analyzes the relationship between retirement and wealth. Using data from the first four waves of the longitudinal Health and Retirement Study—a cohort of individuals born from 1931 to 1941—we estimate reduced-form retirement and wealth equations. Our results show that those who retire earlier do not necessarily save more and that even if one's primary interest is in the relationship between Social Security policy and the decision to retire, it is important to incorporate saving behavior and other key decisions into the analysis.

*Alan L. Gustman is Loren Berry Professor of Economics at Dartmouth College, Department of Economics. Thomas L. Steinmeier is Professor of Economics, Texas Tech University, Department of Economics.

Retirement and Wealth

by Alan L. Gustman and Thomas L. Steinmeier*

Summary

The decision to retire is related to the decision to save and to a number of other decisions, including decisions of when to claim Social Security benefits and what share of assets to hold as pensions, Social Security, and in other forms. This article explores the relationships among these various decisions and then explains why it is important to take them into account when attempting to understand the effects of changing Social Security and related policies on retirement outcomes.

To understand how Social Security benefits affect retirement behavior, and the implications of changing such features as the Social Security early retirement age, the Social Security Administration and others have begun to estimate and use single-equation models of retirement. We explain why the kind of simple model they use is likely to provide a misleading guide for policy. Even if one's primary interest is in the relationship between Social Security policy and the decision to retire, it is important to incorporate other key decisions into the analysis.

These simple models relate the probability of retiring to measures of changes in the value of Social Security benefits when retirement is postponed. The basic problem is that because the omitted factors are related systematically both to retirement outcomes and to the measured reward to postponing retirement, a simple retirement equation credits the effects of the omitted factors to the included measures of changes in Social Security benefits. New policies will change the relationship between retirement and the increase in the value of Social Security benefits with postponed retirement, resulting in incorrect predictions of the effects of new policies.

When we fit single-equation retirement models, we find a variety of evidence that important behaviors have been omitted. These models include variables measuring the age of the respondent. These age variables suggest there is a sharp increase in the probability of retirement at age 62. This is a sign that even though the equations include measures of the increase in the value of Social Security with delayed retirement, the cause of the increased retirement behavior at age 62 has not been included in the model. In addition, the estimated effect of a variable measuring the future value of Social Security and pensions on retirement suggests that if the Social Security early retirement age were to be abolished, more people would retire earlier rather than later—a counterintuitive prediction.

There is even more direct evidence of the need for a more comprehensive model of behavior. We show that if individuals' preferences for leisure time were unrelated to their preferences for

saving, then a simple retirement equation would yield an unbiased estimate of the effects of Social Security on retirement. An implication of such a model is that those who retire earlier for particular reasons would also save more for those same reasons. But when we estimate an equation with wealth accumulated through 1992 as a dependent variable, together with the simple retirement equation, we do not observe that the factors associated with earlier retirement are also associated with higher saving. These and related findings suggest that those who wish to retire earlier also have a weaker preference for saving, a relationship that is ignored in the simple model and can only be measured in a more complex model.

Still other evidence also warns of internal inconsistencies in the simple retirement equations that are being estimated. Social Security incentives are often measured by the increment in the value of benefits associated with deferred retirement, but the incremental value depends on when benefits are claimed. Our findings show that those who retire completely are claiming their benefits too early to be maximizing the expected value of the benefits. Yet the measures of Social Security benefit accrual used in these retirement models often include the increase in the value of benefits from deferred claiming in their measure of the gain to deferring retirement. On the one hand, early retirees are seen not to defer benefit acceptance despite the actuarial advantage. On the other hand, later retirees are said to defer their retirement in order to gain the advantage of deferring benefit acceptance.

Our empirical analysis is based on data from the first four waves of the Health and Retirement Study (HRS), a longitudinal survey of 12,652 respondents from 7,607 households with at least one respondent who was born from 1931 to 1941. Our analysis also uses linked pension and Social Security data together with respondents' records from the HRS.

We also evaluate a number of specific features of retirement models and suggest improvements. We develop a measure of the future value of pensions and Social Security—the premium value—that is not subject to a problem plaguing other measures in that it handles the accrual of benefits under defined contribution plans very well. We also introduce a new definition of retirement status that blends information on objective hours worked with subjective self-reports of retirement status. Our findings also explore the effects of Social Security incentives on partial retirement and consider the importance of incorporating partial retirement in any study of the relation of Social Security to retirement behavior.

Introduction

Researchers often analyze the relation between retirement and the incentives created by pensions and Social Security in the context of a single-equation, reduced-form

model. Such models are routinely used for behavioral and policy analysis. For example, the Social Security Administration has contracted to use such a model to predict the effects of an important change in current policy, namely, increasing the age of eligibility for early Social Security retirement benefits.

Under certain conditions, the coefficients estimated in retirement equations for variables indicating the future reward from Social Security and pensions to continued work will allow us to predict the individual's response to a change in the reward. For example, if people behave according to a simple life-cycle model and if capital markets are perfect, the estimated relationship between retirement outcomes and measures of the change in wealth from Social Security or pensions with continued work will indicate how these financial incentives influence retirement outcomes and how changes in these programs will influence retirement behavior. Under other conditions, however, those measures will not be stable indicators. Thus, for example, if capital markets are imperfect, so that some people are liquidity constrained, the coefficient on a variable measuring the change in the future value of pensions and Social Security cannot be used to predict the effect of a change in Social Security policy. The value of future work depends on unobserved preferences. Consequently, the coefficient estimated in the retirement equation will change as policy changes.

This article examines the efficacy of a single-equation approach to understanding the effects of current and proposed Social Security policies and changes in pensions on retirement outcomes. We would like to determine whether one can interpret the coefficients estimated for variables measuring the future reward to continued work as deep structural parameters, or whether the coefficients commonly estimated are composites that can be expected to change as policies are changed and so are unreliable predictors of the effects of changes in policies on retirement outcomes.

To gain further insight into the underlying behavior, we focus on two outcomes that are jointly determined with retirement: accumulated wealth and the timing of benefit claiming. Our analysis first sketches a theoretical structure that generates various relationships between retirement and wealth in accordance with the correlation between leisure and time preference. We then conduct a number of empirical tests to determine whether the observed parameters obtained in reduced-form retirement equations are likely to be useful for behavioral and policy analysis or whether it is necessary to specify and estimate a structural model that specifically incorporates tastes for leisure and time preference, incorporates liquidity constraints for some, and allows the influence of preferences for retirement and saving to be separated from the effects of future pension and Social Security

rewards. The tests include an analysis of the relation between the residuals from reduced-form retirement and wealth equations. They also consider whether exogenous factors symmetrically affect retirement and wealth, as would be expected in simple models with uncorrelated tastes for leisure and time preference, and whether particular age dummy variables continue to have significant effects on retirement outcomes even after measures of the timing of Social Security and pension incentives are specifically included in the retirement equation. Next we consider the effects of delayed benefit claiming on the value of future rewards to Social Security and pensions and discuss an improved measure of the option value of pensions and Social Security, which we call the premium value.

Findings from these tests raise questions about using a single reduced-form retirement equation to analyze Social Security or pension policies. Parameters from a reduced-form retirement model predict counterintuitively, for example, that raising Social Security's early retirement age will increase the number of early retirements. Although reduced-form models of retirement and wealth accumulation can be improved by modifying both the measure of the retirement variable and the pension premium variable and by incorporating measures of liquidity constraints, these improvements are probably not sufficient to allow their use in policy analysis.

The empirical analysis presented in this article is based on data from the first four waves of the Health and Retirement Study (HRS), a longitudinal survey of a nationally representative sample of the population who were 51 to 61 years old in 1992. Incentives created by Social Security and pensions are measured using linked data. Earnings histories for work through 1991 have been obtained from the Social Security Administration for respondents who signed permission forms allowing their earnings records to be used. Detailed descriptions of pension plan provisions have been obtained from the employers of respondents who indicated they were covered by a pension on present or past jobs.

Measures of the accrual in pension and Social Security values with continued employment play a central role in any study of the relation of pensions and Social Security to retirement and saving behavior. In this article, we measure these incentives by the immediate per-period accrual in benefits from postponing retirement by 1 year and by what we call the premium value—the difference between the value of potential future benefits, including spikes in benefit accrual at early and normal retirement ages, and the value from the basic accrual in each period. Thus, the premium value is positive for a person who has a defined benefit plan but has not yet reached early retirement age—the point at which the plan has a sharp spike in the accrual profile at early retirement. But the

premium value is zero for a defined contribution plan with benefits that accrue evenly each period.

Our analysis also shows that when Social Security incentives are computed on the assumption that respondents accept benefits immediately upon retiring, the calculated incentives to retire are much sharper than when the date of benefit acceptance is timed to maximize the present value of benefits. If covered individuals have to claim benefits immediately because of, say, liquidity constraints, then the reward to postponing retirement (that is, continuing to work) includes the value from postponing benefit receipt. In fact, most of those entering retirement claim their benefits immediately upon retiring. That fact raises a question about whether liquidity constraints or other complexities not reflected in a simple retirement model act to enhance the rewards to immediate retirement or whether the decision to claim benefits should be treated as independent of the decision to retire.

Evidence from previous studies suggests that in attempting to interpret estimated coefficients in retirement equations, it may be necessary to modify assumptions about perfectly operating capital markets, covered workers' understanding of the Social Security system and pension plans, equal valuation of own and spouse's benefits, and other key assumptions. Bearing these caveats in mind, we turn first to a discussion of what has been found in the previous literature and then to our analysis.

What Previous Studies Indicate About Underlying Behavior

Studies of retirement and saving typically are conducted independently of each other and at times involve inconsistent assumptions.[1] Most studies of saving take retirement behavior to be fixed. At best, the retirement horizon or expected retirement date is included as a right-hand side variable.[2] Studies of retirement typically assume that capital markets are perfect, so that saving and consumption decisions are made in the background and do not affect the retirement decision.[3] Nevertheless, previous studies of retirement and saving contain a great deal of information that help clarify the relation between retirement and saving behavior.

Studies of retirement recognize that pension and Social Security benefit formulas affect the reward to continued work and therefore incorporate those incentives.[4] The literature on saving, however, is only now beginning to fully incorporate the influence of pensions and Social Security on saving. Although Social Security and pensions represent half the wealth accumulated for retirement (Gustman and others 1999), many studies of saving ignore pension and Social Security wealth. Moreover, it

is not just a question of whether pensions and Social Security are accounted for when analyzing saving. Even when pensions are counted as part of wealth, fundamental questions remain.

Gale (1998) argues that it is important to properly measure pensions, wealth, and lifetime earnings and to include indicators of the stage of the life cycle if one is to correctly estimate the pension offset in a wealth equation. Consistent with an uncomplicated life-cycle model, he finds indications of large offsets when using data from the Survey of Consumer Finances. Yet when Gustman and Steinmeier (1999) follow Gale's prescription and estimate the pension offset in wealth equations using HRS data, contrary to Gale's predictions they find very little pension offset. Major advantages of the HRS data include the fact that pension values are estimated using detailed descriptions of pension plans obtained from respondent employers; lifetime earnings are estimated using both self-reported earnings histories and earnings histories obtained from the Social Security Administration; and since members of the sample are approaching retirement, their lifetime earnings and lifetime wealth can be estimated fairly accurately. Gustman and Steinmeier (1999) find that those with pensions accumulate more total wealth than those without pensions, holding lifetime income and the retirement horizon constant. As a result, a wealth equation cannot treat pensions simply as a tax-favored method of saving that is a substitute for other forms of saving.

Considerable progress has been made in measuring the future value promised by a pension or Social Security and in using those measures to explain retirement or job mobility. The "option value" of the pension is how Lazear and Moore (1988) and Stock and Wise (1990a and 1990b) refer to the potential value of the pension resulting from continued work at the firm for a number of years in the future. A related measure is the difference between the projected liability and the legal liability of the pension—that is, the value of a defined benefit plan that accrues from future expected employment but is not legally owed to the worker on the basis of employment to date. This measure is used by Ippolito (1986) to evaluate the implicit pension contract. Gustman and Steinmeier (1993 and 1995) use a measure of pension backloading to estimate the disincentive to mobility from pensions. Coile and Gruber (2000 and 2001) adopt a measure they call the peak value, which is the maximum found for all future dates of retirement, and use it to evaluate retirement incentives from Social Security.

In a reduced-form setting, the challenge is to properly value current and future benefits, especially the spikes in the pension accrual profile seen at the early and normal retirement dates. Yet one will downplay the relative importance of the spikes in the benefit accrual profile at early and normal retirement ages by simply adding up the expected future benefit for each year of future employment. For example, when benefits are simply summed, a defined contribution plan will have a misleadingly large future value. In the discussion below, we will blend available measures for valuing future benefits, basing our evaluation of the expected future value of the pension or Social Security on the premium value. The "premium value" differs from "peak value" used by Coile and Gruber (2000 and 2001) in that the peak value counts all increases in benefits with continued work—and thus continues to increase in time as benefits are accumulated in defined contribution plans—but the premium value does not.

Many findings from the literature raise questions about the behavior governing retirement and saving decisions. People are not very well informed about the details of their pensions. Many cannot identify what type of pension they have (Mitchell 1988; Gustman and Steinmeier 1989 and forthcoming). Respondents are especially poorly informed about the location and size of the spikes in pension benefit accruals created by their defined benefit plans, which are key determinants of the incentives that pensions create for retirement behavior (Gustman and Steinmeier forthcoming). Imperfect information about pensions leads to two kinds of problems. One is that descriptions of pensions (or Social Security) obtained from respondents may be misleading. For example, when respondents misidentify their type of plan, they typically are asked follow-up questions about the wrong type of plan. This problem can be remedied by using linked pension and Social Security data obtained from employers and the Social Security Administration. A second problem is that the respondents may be guided in their saving or retirement decisions by a misunderstanding of their pensions. This problem cannot be fixed by using better data; rather, it must be modeled.

There also are questions about the behavior that determines when people claim their Social Security benefits. There is a literature analyzing when it is optimal to claim benefits.[5] From an expected value perspective, it is often optimal to delay claiming benefits when first eligible so as to disproportionately increase the value of benefits, especially of spouse and survivor benefits.[6] However, there are reasons for some people to claim benefits before the present value is maximized. For example, those with private information who believe that they are likely to die at a younger age, or who are (mortality) risk averse, will claim their benefits earlier. Models of retirement and saving should be reconciled with observed behavior of benefit claiming.

There are many reasons why Social Security beneficiaries may not delay their acceptance of benefits to the optimal time. One possibility is that the primary benefi-

ciary places less weight on spouse and survivor benefits than on his or her own benefits, which would lead to earlier claiming.[7] Another possibility is a high discount rate. Perhaps because they have high discount rates, some consider themselves to be overannuitized and liquidity constrained. A household with little liquid wealth will not be able to support consumption between retirement and the time of first receipt of delayed benefits. Positively correlated leisure and time preferences may also make early claiming more likely among retirees. Or perhaps some people believe the government will pay lower Social Security benefits than they have been promised; such persons attach a great deal of risk to the future payments promised by Social Security and therefore believe it is in their interest to collect their benefits as early as possible.

It is important to understand claiming behavior in order to properly measure how Social Security affects the incentive to retire. We show below that when people claim their Social Security benefits so as to maximize expected value, the reward to postponing retirement is lower than if they claim benefits as soon as they retire. Even if benefits are claimed immediately upon retirement, as the evidence suggests in most cases it is, retirement and claiming behavior may not be tied in the respondent's mind. Accordingly, retirement decisions may not be influenced by the actuarial increase in the value of Social Security benefits from delayed claiming. Still another possibility is that individuals may be sophisticated enough to understand the actuarial return to postponing benefits but not sophisticated enough to divorce the decision to retire from the decision to accept benefits. Thus the extent to which Social Security creates incentives that influence retirement outcomes depends on claiming behavior, and the valuation of deferred Social Security benefits in turn depends on the reason why most retirees do not defer their benefit claims.

Among persons who are working part time and are earning enough to be subject to the earnings test, more are willing to postpone accepting benefits.[8] A person who is working part time and making more than the earnings test disregard is in roughly the same actuarial position with regard to the lost benefits as a person who postpones benefit receipt. Both will have their future benefits increased by a similar amount to cover their lost benefits.

We are aware of a number of other issues affecting the specification of retirement and saving equations. Findings are sensitive to how retirement is measured—based on self-reported status, hours of work, or some combination (Gustman, Mitchell, and Steinmeier 1995; Gustman and Steinmeier 2001). Findings will also be influenced by whether the partially retired are counted as retired or not retired (Gustman and Steinmeier 1984). We address these issues below.

Joint Determination of Retirement and Wealth in a Simple Model

To facilitate the discussion of the relationship between retirement and wealth, let us examine a simple yet instructive model. In this model, the consumer maximizes a lifetime utility function:

$$U \ = \ \int_0^T e^{-\rho t} u[C(t)] \, dt$$

subject to a lifetime budget constraint

$$\int_0^T C(t) \, dt \ = \ WR$$

where C(t) is consumption at time t, W is the (constant) wage rate, R is the retirement age, and T is the lifetime.

This model solves for consumption and wealth, given the optimal retirement date. The effect of variation in the taste for retirement on saving is then simulated by varying the date of retirement. A more complete analysis would include leisure in the utility function and allow for heterogeneity in the leisure parameter. The results demonstrated here also hold in a more general model in which leisure is included in the utility function and retirement is endogenously determined. We have done the required calculations, and they are quite extensive. This simple model, however, illustrates the major points without undue complications.[9]

The Euler-Lagrange condition for this problem is

$$U'[C(t)] \ = \ \lambda e^{\rho t}$$

where λ is a Lagrangian multiplier that, in this problem, is constant over time. Differentiating this condition with respect to the retirement date R yields

$$U''[C(t)] \ \frac{\partial C}{\partial R} \ = \ \frac{\partial \lambda}{\partial R} e^{\rho t}$$

Since $U'' < 0$, this condition implies that $\partial C / \partial R$ and $\partial \lambda / \partial R$ are of opposite signs, and furthermore, since λ is constant over time, that the sign of $\partial C / \partial R$ is uniform over time.

Differentiating the budget constraint with respect to R gives

$$\int_0^T \frac{\partial C}{\partial R} \, dt \ = \ W \ > \ 0$$

Since $\partial C/\partial R$ has a uniform sign over time, that sign must be positive. Assets at any point in time before retirement are simply the difference between the cumulative wages and the cumulative consumption:

$$A(t) \;=\; Wt \,-\, \int_0^t C(t')\,dt'$$

Since an increase in the retirement age uniformly increases consumption over time, it must reduce the level of assets at any point in time: $\partial A/\partial R < 0$.

Implications of Heterogeneous Leisure Preferences

Suppose that different individuals have characteristics (either observed or unobserved) that make them either more or less inclined to retire early. Let X_i be one such characteristic, one such that high values of X_i are associated with earlier retirement: $\partial R/\partial X_i < 0$. We can also ask what the effect of X_i is on asset holdings at some time prior to retirement. Since X_i operates indirectly through the retirement age in the model above and not directly on either assets or consumption, $\partial A(t)/\partial X_i = \partial A(t)/\partial R \; \partial R/\partial X_i > 0$. Holding all other things equal, a characteristic that makes an individual more inclined to retire early also induces that individual to hold more assets than otherwise.

A simple interpretation is that if the individual plans to retire early, he or she will hold more preretirement assets in order to finance the longer period of retirement without a sharp cutback in consumption. This finding is noted in the top panel of Table 1. There, an earlier retirement is associated with an increased level of assets at any preretirement age.

Implications of Heterogeneous Time Preferences

Next, we investigate the effects of heterogeneous time preferences, holding leisure preferences (and hence the retirement date) constant. Without going through the details of the derivation in the model above, it can be shown that $\partial A(t)/\partial\rho < 0$. Heuristically, an increase in time preference is associated in the consumption formula with a more rapid decline in consumption over the lifetime and, hence, with a tendency to consume more in the early years. Increased consumption in the early years will lower the amount of accumulated savings with a given level of wages.

As shown in the middle panel of Table 1, a higher level of time preference will have no effect on the retirement age, given the assumption that leisure preferences are constant. However, the higher level of time preference will result in lower rates of asset accumulation and lower levels of assets at any given age.

Correlated Leisure Preferences and Time Preferences

The previous sections have examined either heterogeneous leisure preferences (holding time preference constant) or heterogeneous time preferences (holding leisure preferences constant). If the two sets of preferences were independent, then the correlation between early retirement and higher wealth levels that are implied from the top panel in Table 1 would prevail overall. That is, an individual with high leisure preferences would be more likely to retire early and hold more wealth. Because there is no systematic correlation with leisure preferences, heterogeneous time preference does not change this relationship, although it does spread out the wealth distribution for a given leisure preference. The net result is that allowing for both preferences but requiring that they be independent implies that there is still a positive association between early retirement and wealth holdings but that they are not as tightly correlated as when we considered heterogeneous leisure with a given time preference.

However, there is no particular reason to assume that leisure preferences and time preferences are uncorrelated, and arguments for a correlation are relatively easy to make. A high time preference is symptomatic of an increased desire for short-term gratification, the "I want it **now**" attitude. The same desire for short-term gratification is likely to carry over into the leisure/work decision, where it manifests itself as an increased desire for leisure. Thus, it is plausible to argue for a positive correlation between time preference and leisure preference.

The bottom panel of Table 1 gives the results of combining heterogeneous leisure preferences with positively associated heterogeneous time preferences. An individual with high leisure preferences is more likely

---

Table 1.

Effects of leisure preferences and time preferences on retirement and wealth

Preference	Effects on retirement decision	Effects on level of wealth
Leisure preference		
Low	Late	Low
High	Early	High
Time preference		
Low	No effect	High
High	No effect	Low
Positively correlated leisure and time preferences		
Leisure preference is low	Late	Ambiguous
Leisure preference is high	Early	Ambiguous

to retire early. Because of the longer retirement period, there is an incentive to have higher levels of wealth in the years leading up to retirement. However, offsetting this finding is the fact that such an individual is likely to have high levels of time preference as well. High levels of time preference work in the opposite direction in terms of wealth accumulation and tend to lower the level of wealth. Which effect is dominant is a priori unclear; hence, the wealth of individuals with high leisure preferences is labeled as "ambiguous." The net result is that in this situation early retirement may be associated with either high or low levels of wealth, and the direction of the correlation between retirement and wealth is not determined.

Implications

One of the purposes of this study is to find out what kinds of models are generally consistent with the data. A model that allows for individual heterogeneity in preferences for leisure but assumes that all individuals have the same time preferences implies a negative relationship between retirement ages and wealth levels. A slightly more general version of this model, which includes both heterogeneous leisure preferences and heterogeneous time preferences and allows for these preferences to be correlated in plausible ways, can accommodate cases in which retirement ages and wealth levels are not correlated or are positively correlated.[10]

A structural model that explicitly incorporates the distributions of leisure and time preference will allow the data to tell the story. Evidence on the relation between wealth and retirement will provide the first piece of evidence as to whether the story is consistent with the simple model that must underlie a reduced-form approach if the coefficients estimated for pension and Social Security wealth are to reflect the behavioral response to the incentives created by those plans, or whether the estimated coefficients are composites that will change in value if pension and Social Security rules are changed. Other evidence on whether a simple reduced-form approach is adequate for understanding the effects of pension and Social Security policies on retirement outcomes is also developed.

Data and Variables

The data used to investigate the relationship between retirement and wealth come from the first four waves of the original cohorts of the Health and Retirement Study. The HRS began in 1992 with about 9,800 respondents who were born between 1931 and 1941. Spouses were also interviewed, but they are not included in the retirement portion of this study unless they were born in that time period; otherwise they would not be representative

of their respective cohorts. The study continued to interview the respondents at 2-year intervals, and the current study uses these interviews through 1998, which is the last interview available as of this writing.

Defining Retirement

One of the focuses of the study is retirement, which in the empirical analysis we will take to be the transition from working in one survey year to being retired in the next. Measures of retirement as of the survey date are probably more precise and do not require us to infer exactly when between two surveys an individual actually retired. To implement this definition of retirement, however, we must define exactly what it means to be working and what it means to be retired.

There are several potential ways to measure retirement in the HRS, but they group into objective measures, such as whether you have a job in the survey week, and subjective measures, such as whether you consider yourself to be retired. These measures are not always consistent. Table 2 gives cross-tabulations of two measures: usual hours per week and self-reported retirement status.[11] The percentages along the diagonal are instances where the two measures agree, and they total to about 83.4 percent of the observations. For the remaining observations, which are about one-sixth of the total, there is disagreement between the objective measure and the self-reported retirement status.

Cases in which the respondent is working more than would be expected with the self-reported retirement status appear above the diagonal. Since the respondent is working, it is probably not appropriate to classify him or her as completely retired. On the other hand, an examination of numerous individual records suggests that if the respondent indicates that he or she is partially or fully retired, there is usually a reason for the response even if the current hours are in the full-time range. Perhaps the respondent has worked for 60 hours per week in previous jobs and is now working only 40 hours a week, or in some

Table 2.

Objective vs. self-reported retirement status (as a percentage of all observations)

Objective measure (usual hours per week)	Self-reported retirement status			
	Not retired	Partially retired	Completely retired	All observations
More than 35	47.6	2.9	0.4	50.9
1 to 35	3.9	3.4	0.8	8.0
0	5.5	3.2	32.4	41.1
Total	57.0	9.5	33.6	100.0

SOURCE: Authors' calculations.

cases there is a noticeable drop in earnings, suggesting an easier job. Frequently the work history contains a change of employer around the date the respondent says he or she partially or fully retired. In any case, it appears to be sensible to treat respondents who are working but say they are partially or completely retired as though they are partially retired, since in most cases there is at least some evidence they are not working as hard as they did at one time.

Below the diagonal are cases in which the respondent is working less than would be expected with the self-reported retirement status. One cell contains respondents who claim to be not retired at all even though their usual hours per week at their present job are below 35. To decide whether such individuals are not retired or partially retired, we looked at previous jobs in the job history. If there were previous jobs with 35 hours of work or more, then there is evidence of a reduction of work effort, and the individuals are classified as partially retired. If there is no evidence of previous jobs with 35 or more hours per week, then there is no evidence of lower work effort, and the respondents' claims that they are not retired at all are accepted. For the respondents who claim to be not retired or partially retired but who do not have current jobs, we look to see whether they also claim to be unemployed and how long ago their last job was. If they say they are unemployed and had a job within the previous 12 months, their self-reported status is accepted. But for the remainder of the respondents, who are the large majority of this group, the claim of not being retired is not accepted, and they are classified as being completely retired.

In short, we are making a new definition of retirement status based on both objective hours and subjective self-reports. By themselves, both self-reports and objective hours have problems. Objective measures have problems with individuals who reduce work effort while still being above 35 hours and with individuals who have always worked less than 35 hours. Self-reports appear to be unreliable both for individuals who have jobs yet say they are completely retired and for individuals who do not have jobs yet claim to be not retired. The hybrid measure of retirement that we are using should ameliorate these deficiencies.

Measuring Wealth

The second focus of the study is on wealth. The dependent variable in wealth regressions is defined as non-Social Security, nonpension wealth. The HRS went to a lot of trouble to gather good data on wealth, including trying to bracket amounts for which the respondents were unable to provide exact numbers. The quality of the data both reduces the need for imputation and probably increases the accuracy of the imputations that

are made, increasing the accuracy of the wealth measures. We use values imputed by the HRS where required information on wealth is missing.

Pensions and Social Security together account for more than half the total wealth of respondents to the HRS (Gustman and Steinmeier 1999). Incentives for retirement are calculated by considering the changes in Social Security and pension wealth associated with additional work. Pension incentives are estimated from the matched pension plan formulas obtained from the employers for covered HRS respondents. The pension plan descriptions were coded by HRS staff, and the plan values are calculated from those descriptions using the reported wage and projecting it backward using the general wage growth rates. Social Security incentives are estimated from the earnings in the Social Security record, with earnings after 1991 projected using the Social Security assumptions about real wage growth rates.

For respondents whose Social Security records could not be obtained, we impute the record before 1991 using information in the HRS main survey. Respondents were asked about the starting date on their current job, starting and ending dates for their last job (that is, the job last held by those not working in 1992), starting and ending dates for the previous 5-year job held before the current or last job, and the starting and ending dates for up to two other pension-covered jobs. Respondents were also asked about earnings at these dates. In addition, the survey asked respondents in wave 3 about the date of entry into the labor force, how many years they worked before the date the previous job was secured, and the dates that the individual was in jobs not covered by Social Security. Wage profiles are forced through all years when the individual implied he or she was working in jobs covered by Social Security, with values for missing years projected backward off the profiles on the basis of experience and education.[12] From the Social Security earnings record (either actual or imputed if missing), we calculate the respondent's average indexed monthly earnings (AIME) amount and from that the Social Security benefit to which the respondent is entitled (the primary insurance amount, or PIA). The benefit amounts, in turn, are used to calculate the value of Social Security and the incentives for retirement arising from Social Security.

The main problem in wealth regressions is one of scale. If wealth is entered in a linear format as a dependent variable, the wealth regressions are likely to be dominated by respondents with high levels of wealth. If instead wealth is entered in a logarithmic format, there is the problem of what to do with respondents who have zero or negative wealth. These problems can be avoided by using as the dependent variable the level of wealth as a percentage of potential wealth, which can be measured
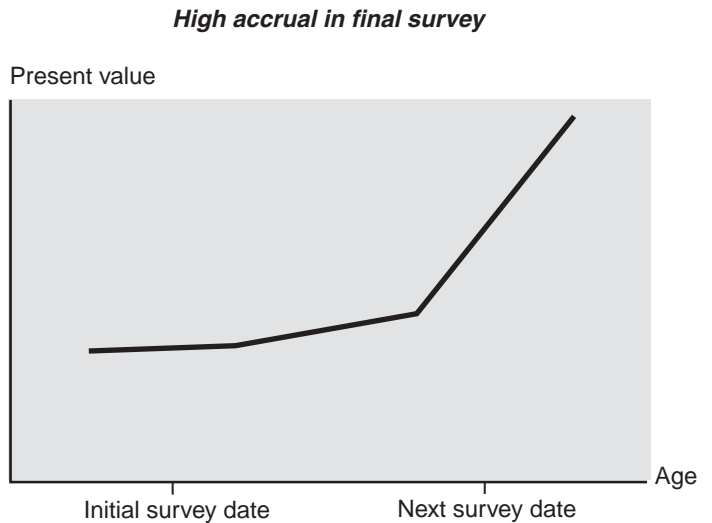
as the real value of lifetime household earnings. Lifetime earnings, in turn, can be measured fairly accurately from the Social Security earnings records that were collected as part of the survey. For instances in which earnings are masked by the Social Security earnings maximum or were not recorded because the respondent was in a noncovered job, actual earnings can be inferred, albeit approximately, from the respondent's reported earnings.[13] The resulting dependent variable, which should lie between zero and one, should not be severely affected by scale. Roughly speaking, this approach treats a household that has $100,000 in assets out of $2,000,000 in lifetime earnings as being in approximately the same situation as a household that has $25,000 in assets out of $500,000 in lifetime earnings.[14]

Most of the explanatory variables in this study are fairly straightforward, and Box 1 includes a short description of selected variables. A few variables, however, merit additional discussion, the most important being those that relate to the incentives that pensions and Social Security provide either to keep on working or to retire.

The first two of these variables measure the *increases* in the present values of future pension and Social Security benefits that come with continued work. They are usually called the pension and Social Security accruals. If we plot the present value of pensions and Social Security as a function of retirement, as in Chart 1, the slope of the present-value line is a measure of the accrual at any point in time.

Since we are looking at the probability of retiring in the period between one survey and the next, two accruals are relevant. In the top panel of Chart 1, the respondent has a large accrual in the initial survey year (initial year after the initial survey date) but a small accrual in the second survey year (second year after the initial survey date). Such an individual would have a high incentive to delay retirement until after the initial survey year but no

Chart 1.
Patterns of accruals

High accrual in initial survey

Present value

Age

Initial survey date   Next survey date

High accrual in final survey

Present value

Age

Initial survey date   Next survey date

Box 1. Partial list of variables	
Wealth	Nonpension, non-Social Security wealth as of 1992
Earnings	1992 earnings (amounts < $100 disregarded)
Social Security value	Household Social Security wealth, assuming spouse works to expected retirement age
Pension value	Pension value as of 1992
Married	Binary variable for being married in the initial year
Health	Binary variable for fair or poor health in initial year
Children	Binary variable for at least one child
Word recall	Number of words recalled in second attempt
Share of lifetime household earnings	Respondent's share of sum of lifetime earnings of respondent and spouse (as of 1992)
Reduced hours	Binary variable if respondent can reduce hours in the current job
Laid off from initial job	Binary variable if respondent was laid off from initial job during the period

additional incentives to further delay. Thus high accruals in the initial survey year should increase retirement during the period.

This result contrasts with the bottom panel of Chart 1, which illustrates a large accrual during the second survey year. In this case, the respondent will have a large incentive to delay retirement until after the second survey year, and a large accrual in that year should be associated with lower retirement. If the accruals were similar in both years, the respondent would have no particular incentive or disincentive to retire during the period, suggesting that the positive effects of an accrual during the first survey year should be of roughly the same magnitude as the negative effects of an accrual during the second.

Social Security and pensions may also provide additional incentives to continue employment into future years that are not necessarily captured by the accruals at either the start or the end of the period over which we are measuring retirement. An example would be a pension that increases sharply in value a couple of years after the end of the second survey year. In this case, a respondent might delay retirement not because the current accruals are high but because of the prospect of a higher pension if he or she waits until the sharp increase in value. This idea is called "option value" by Lazear and Moore (1988) and Stock and Wise (1990a and 1990b) and "peak value" by Coile and Gruber (2000 and 2001). However, neither measure quite embraces the idea that we are trying to capture, which is the potential of a future extra bonus on top of any current accruals. For instance, both the option value and the peak value would increase more or less indefinitely for defined contribution plans, and yet these plans in general are not perceived to provide a strong incentive to retire at any particular time.

For this reason, we are introducing a new measure of future incentives that we call the "premium value." To calculate this measure, for each future year we calculate the value of the pension and compare it with the value the pension would have if the current accruals continued until the future year. The premium value, which is analogous to the measure used in Gustman and Steinmeier (1993 and 1995), is simply the maximum of the present value of these differences. The premium value is illustrated in Chart 2. The solid white line gives the amount that the pension would be worth if it kept accruing value indefinitely at its current rate, and the solid line gives the actual value of the pension. The premium is the maximum vertical difference between the solid line and the dashed line. In this case, the pension jumps notably several years after the current age, perhaps when the respondent becomes eligible for early retirement and as a result can obtain benefits under a more favorable formula than before. As illustrated in the chart, the premium is a measure of the extra value of the pension beyond the value implied in the current accrual.

Note that a defined contribution plan that increases steadily in value will have a zero premium value, since there are no future benefits in this type of plan that are not evident in the current accrual rate. Social Security benefits can also have these premium values if the benefit increases for delaying benefits are more than actuarially fair. Such is frequently the case for married respondents whose spouses will be collecting benefits based on the respondents' earnings.

The distributions of accruals and premium values for both Social Security and pensions are shown in Table 3. The observations are for individuals in pairs of successive surveys. Since there are four surveys, each respondent can have up to three observations; other restrictions are noted later in this section. We refer to the first survey of any pair as the "first survey year" and the second survey of the pair as the "next survey year." The accruals are measured at both survey dates, as suggested by Chart 1. A high accrual in the first survey date and a low accrual in the next survey date would signify that effective compensation dropped over the 2-year period, and that should encourage retirement. The opposite would be true

Chart 2.
Premium values

Present value

Age

Current age
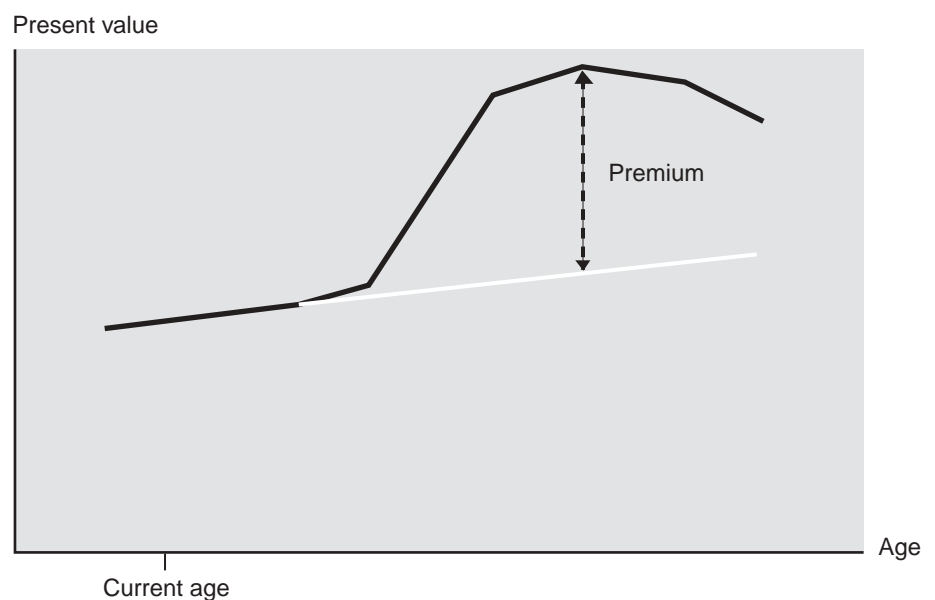
if the accrual on the next survey date was higher than on the first survey date. The premium values are measured at the later of the two survey dates, because it is presumably the premium at that time that would induce respondents to delay retirement. Both the accruals and premium values are expressed as a percentage of current earnings. Presumably the incentives from pensions and Social Security to continue working are more related to the percentage by which they increase regular earnings than they are to the absolute values of the amounts.

Pension and Social Security accruals each average around 6 percent to 8 percent of current earnings, but the variation in pension accruals is almost twice as much as for Social Security accruals. The variation is important because if the estimated effects are the same, the differential impact of the accruals on retirement behavior for the respondents is related to the variance of the accruals and not necessarily to the mean. With regard to the premium values, when averaged across the whole population, the premium is actually higher for Social Security than for pensions, at 18 percent compared with 11 percent, but again the variation in premium values for pensions is somewhat greater than for Social Security. Part of the difference in means comes from the fact that over four times as many respondents have Social Security premium values as have pension premium values. If we look only at respondents with positive premium values (see Table 3), both the mean and variation of the pension premium values are much higher than for the Social Security premium values.

The final data issue is the derivation of the sample to be analyzed from the observations in the data set (see Table 4). The HRS interviewed 12,652 respondents in the initial wave in 1992, and by 1998 the survey had conducted almost 45,000 interviews with those individuals. However, only the respondents born between 1931 and 1941 are a representative sample, and imposing that restriction eliminates about a quarter of the interviews. We require that the individual be initially not retired, that is, working full time, which leaves about 18,000 observations. We require usable age and earnings figures and—if the respondent is married—that the spouse also be interviewed so we can compute household earnings variables. Finally, if

the individual reports a pension on the current job, we require that the pension be included in the employer-provided pension file. We make this last requirement because the respondent interview provides a very poor

Table 3.
Accruals and premium values for pensions and Social Security (as a percentage of current earnings)

Source of accrual	Mean	Standard deviation	Percentage with nonzero values
Accruals at the start of the period			
Pension	8.5	27.6	42.7
Social Security	6.1	11.4	78.0
Combined	14.6	29.8	85.2
Accruals at the end of the period			
Pension	6.6	23.1	43.9
Social Security	5.6	10.8	80.0
Combined	12.2	25.4	86.6
Premium values for all respondents			
Pension	10.6	46.1	14.2
Social Security	17.9	38.4	61.3
Combined	22.2	57.1	50.9
Premium values for respondents with nonzero values			
Pension	74.8	100.9	
Social Security	29.2	45.6	
Combined	43.7	74.0	

SOURCE: Authors' calculations.

Table 4.
Derivation of the sample

Interview	Wave 1	Wave 2	Wave 3	Wave 4	All waves
All interviews	12,652	11,316	10,653	10,119	44,740
Age-eligible interviews[a]	9,824	8,804	8,312	7,886	34,826
In initial year					
Working full time	6,310	4,927	3,845	3,088	18,170
With nonmissing age	6,310	4,742	3,845	3,088	17,985
With nonmissing earnings	5,343	3,962	3,211	2,527	15,043
With nonmissing spouse	5,194	3,847	3,075	2,381	14,497
With nonmissing pension	4,072	3,069	2,523	2,008	11,672
In next survey year					
With interview	3,739	2,844	2,332	0	8,915
With nonmissing work status	3,735	2,842	2,331	0	8,908
With nonmissing age	3,474	2,825	2,331	0	8,630

SOURCE: Authors' calculations from Health and Retirement Study.

a. Interviews with respondents born between 1931 and 1941.

basis for imputing pension accruals and premium values (Gustman and Steinmeier forthcoming). Imposing these restrictions leaves us with about 11,700 observations.

Since retirement is defined as a change in status between one survey and the next, we must consider periods in which the respondents were interviewed in two adjacent waves. Dropping interviews for which there was no subsequent interview leaves about 8,900 observations. There are a couple of minor additional deletions because either the work status or age is not available in the final wave, leaving us with about 8,600 observations used in the retirement part of the analysis.

For the wealth regressions, there are some additional deletions. First, it would seem inappropriate to use the same regression for both married respondents and single respondents. Among the single respondents, there are problems with divorced and widowed respondents because the survey does not interview the former spouses, and hence we cannot tell the earnings potential of the household. The sample of the remaining single respondents, who are the never-married group, is small enough that the results are questionable. Therefore, we only look at married respondents in wealth regressions. This brings the sample down to about 6,300. Second, we further delete anyone in a household that reports any substantial inheritance (more than $10,000) or whose total wealth, including pensions and Social Security, exceeds the real value of the earnings for that household. This leaves around 5,600 observations for the wealth regressions.

Results of the Retirement and Wealth Regressions

The principal results of the retirement and wealth regressions are shown in Table 5. The retirement regression is actually a probit equation, and the figures reported in the table are the marginal effects, that is, the change in probability of retirement that results from a one-unit change in the independent variable.

First consider the retirement probit. The dependent variable in this probit is whether or not a respondent who was fully working in the one survey had completely retired by the next survey, where retirement is as defined in the previous section. The overall probability of retirement between one survey year and the next is about 13.6 percent, so that numbers such as 0.06 or 0.07, though they may appear small, actually represent an increase in retirement rates of about 50 percent. In the retirement equation, the combined pension and Social Security incentive variables are all significant and have the correct sign. We would have expected the two accrual effects to be approximately equal and of opposite sign, whereas the

effect of the final accrual is almost twice as large. However, the difference is not significant using a likelihood ratio test.[15] These coefficients suggest that moving from an accrual value that is one standard deviation below the mean to one that is one standard deviation above the mean (see Table 3) changes retirement by around 3 percentage points, or by roughly one-quarter. A similar variation in the premium value would also change retirement by 2 to 3 percentage points.[16]

The age variables follow the expected path in that the retirement probability steadily increases at older ages. There is almost no evidence of a pure age effect at age 65, although there is a considerable effect at age 62. Recall that age is measured at the beginning of the period and that the period is roughly 2 years, so respondents aged 60 or 61 at the beginning of the period will have passed 62 by the end of the period. Thus, the increases in the coefficients at ages 60 and 61 probably reflect a spike when individuals turn 62. The cause of this spike is still under debate. It could reflect liquidity constraints that are relaxed when the individual is able to collect Social Security benefits, or it could be that individuals do not value the actuarial adjustments to future Social Security benefits very much (or are not aware of them) so that at age 62 it appears that they are giving up benefits by continuing to work. Some analysts argue that it reflects some type of social norm, although this norm is certainly not reflected in the dates of eligibility for early retirement under pension plans, which have a modal value of 55 for those with defined benefit plans in the HRS (Gustman and Steinmeier 2000). In any case, most observers would probably agree that a major part of the cause of the retirement increase at age 62 has something to do with Social Security, even if the exact process remains unclear.

The other variables in the retirement probit behave more or less as expected, to the degree they are significant. The two most important variables are poor health and having been laid off from the initial job, both of which substantially increase the probability of retirement. Tenure (years of service) in the initial job is also highly significant, with higher tenure levels appearing to increase the probability of retirement. Another significant variable is the share that the respondent has contributed to lifetime household income (as of 1992); primary earners retire later. The self-employed are also likely to retire later. There is no systematic difference in retirement associated with sex (although primary earners in the family are significantly more likely to retire later), or with black or Hispanic status once the other independent variables are included in the retirement equation. Also note that the planning horizon, which is the closest direct measure we have to time preference, is not significant in the retirement equation. Finally, those who are able to

Table 5.
Retirement and wealth equations

Independent variable	Full retirement probit		Wealth regression	
	Marginal effect	t-statistic	Coefficient	t-statistic
Constant			0.2272	3.42
Measures of earnings				
Log of annual earnings	-0.0143	-2.48		
Log of lifetime family earnings			-0.0110	-2.48
Pension and Social Security values				
Pension value/lifetime earnings			-0.0548	-2.65
Social Security/lifetime earnings			0.1000	2.57
Pension and Social Security incentives				
Initial accrual/annual earnings	0.0348	3.29	0.0142	2.46
Final accrual/annual earnings	-0.0656	-3.63	0.0111	1.64
Premium value/annual earnings	-0.0216	-2.85	0.0038	1.23
Age binary variables				
50			-0.0407	-1.07
51	0.0114	0.47	-0.0129	-1.33
52	0.0025	0.12	-0.0076	-0.90
53	0.0531	2.66	0.0007	0.09
54	0.0449	2.42	0.0023	0.32
56	0.0290	1.58	0.0121	1.68
57	0.0474	2.57	0.0044	0.62
58	0.0584	3.06	0.0149	2.01
59	0.1100	5.41	0.0074	0.95
60	0.1849	8.35	0.0128	1.65
61	0.2559	10.54	0.0148	1.78
62	0.2514	8.55	0.0009	0.09
63	0.3232	8.92	0.0231	1.69
64	0.3252	6.85	-0.0081	-0.44
65	0.3110	3.69	0.0699	2.05
Personal characteristics				
Female	-0.0026	-0.23	0.0101	1.52
Married	-0.0318	-2.19		
Age difference if married	-0.0003	-0.39	-0.0017	-4.83
Race				
Black	0.0008	0.08	-0.0511	-9.16
Hispanic	-0.0218	-1.60	-0.0171	-2.49
Fair or poor health	0.1035	8.69	-0.0145	-2.62
Not available			-0.0575	-0.46
Education				
Less than high school	0.0149	1.49	-0.0203	-4.13
Some college	-0.0125	-1.27	0.0211	4.41
Undergraduate degree	-0.0032	-0.24	0.0223	3.41
Graduate work	-0.0058	-0.43	0.0366	5.61
Children	0.0102	0.69	0.0029	0.28
Planning horizon				
Next year	0.0048	0.56	-0.0204	-4.63
More than 10 years	0.0077	0.56	0.0084	1.27
Not available	0.0245	1.09	0.0096	0.93
Word recall (number of words)	0	-0.01	0.0006	0.84
Not available	-0.0172	-0.80	-0.0227	-2.15
Share of lifetime household earnings	-0.0773	-3.82	0.0100	0.98

Continued

Table 5.
Continued

Independent variable	Full retirement probit		Wealth regression	
	Marginal effect	t-statistic	Coefficient	t-statistic
Job characteristics				
Self-employed	-0.0298	-2.26	0.1072	16.65
Not available	0.8905	11.82	0.1409	1.02
Years of service	0.0016	4.54	0.0012	7.07
Not available	0.2678	2.43	0.0141	0.30
Industry				
Manufacturing	0.0124	1.08	-0.0038	-0.72
Public administration	0.0170	1.06	-0.0061	-0.81
Occupation				
Management or professional	0.0016	0.14	0.0128	2.52
White collar	-0.0025	-0.27	0.0064	1.39
Covered by union	0.0133	1.45	-0.0106	-2.36
Not available	0.1045	0.85	-0.0197	-0.31
Covered by pension	0.0224	2.23	-0.0075	-1.50
Not available	-0.1186	-31.42	-0.0926	-0.62
Firm with more than 100 employees	0.0174	1.57	-0.0010	-0.18
Not available	0.0309	1.80	-0.0127	-1.56
Availability of reduced hours	-0.0231	-2.16	0.0071	1.27
Laid off from initial job	0.1497	8.00	-0.0080	-0.96
Pseudo R ^{2} or adjusted R ^{2}	0.10		0.15	
Number of observations	8,612		5,608	

SOURCE: Authors' calculations from Health and Retirement Study.

NOTE: The probit estimates are the marginal effects on the probability of retirement of a one-unit change in the explanatory variable.

reduce their hours of work without leaving their jobs are less likely to proceed directly from full-time work to full retirement, instead either prolonging the length of time spent on a job in which the workload can be modified or partially retiring on such a job.[17]

The wealth regression uses the same observations as the retirement equations, minus single respondents, respondents with substantial inheritances, and respondents whose total wealth exceeds lifetime household earnings. The dependent variable for this regression is the ratio of nonpension, non-Social Security wealth in 1992 to lifetime household earnings. This variable may loosely be interpreted as the fraction of lifetime household resources that have been saved in addition to pensions and Social Security. Since many types of wealth, such as household wealth or financial wealth, cannot really be separated into parts due to each partner, this variable is necessarily a household variable, although the observations are still individuals. As with the retirement variable, the magnitude of the coefficients may be a little deceiving. A value of 0.01 is associated with an increase in household wealth of 1 percent of the lifetime earnings of both spouses, and this can translate into a sizable sum.

The first coefficient is that of the log of total lifetime family earnings.[18] The sign and magnitude of this coefficient suggests that, all other things being equal, a doubling in earnings causes the wealth ratio to drop by about 1 percentage point.[19] The next two variables are the ratio of pension wealth and Social Security wealth to lifetime household earnings.[20] If there were perfect substitution between pension wealth, Social Security wealth, and other types of wealth, these coefficients would be -1, and reductions of other types of wealth would fully offset any pension or Social Security wealth. If there were no offset, the coefficients would be zero. In contrast to the predictions of a simple life-cycle model and consistent with our earlier results with a slightly different specification, these coefficients suggest that the respondents do not reduce the amounts of other types of wealth very much to offset higher levels of pension and Social Security wealth.[21]

For reasons that are not completely clear, the coefficients on the accrual and premium value variables are all positive, although only one of them is significant.[22] Significant coefficients on other variables have effects in plausible directions. These variables include the race

variables, with the ratio of wealth to lifetime household earnings 5 percent lower for blacks; the education variables, with better-educated respondents having considerably more wealth (holding lifetime earnings constant); and the planning horizon variables, with those with short horizons having less wealth. The tenure variable is also highly significant in increasing wealth. There is some tendency of older respondents to have higher wealth ratios, but the tendency is fairly noisy. Households with a larger age difference between spouses, those in poor health, and union workers have lower wealth. Self-employment is associated with much higher wealth, suggesting a unique motivation for wealth accrual by the self-employed.

As indicated above, one of the main interests of these regressions is to see whether retirement and wealth are correlated, as a model with heterogeneous but uncorrelated retirement or time preferences would suggest. Such correlation of retirement and wealth should be evident in Table 5, which lists the results for both the retirement and wealth equations. To facilitate the comparison, the two equations in this table have corresponding observations, except that the wealth equation is limited to married respondents. However, the retirement equation is not much changed when it too is limited to married respondents, and a test of the proposition that married and single respondents have the same coefficients in the retirement equation is not rejected.

Negative correlation of retirement age and wealth, to the extent it exists, should have two implications. First, the independent variables should work in the same direction in the retirement and wealth regressions.[23] Table 6 contains a summary of the significance of the coefficients (other than age) in the two equations. If the independent variables work in the same direction in both equations, there should be a pronounced concentration of entries along the northwest to southeast diagonal. However, the actual pattern does not yield the impression that there is much of any correlation at all between factors affecting early retirement and wealth.

Because the fit in both the retirement probit and the wealth regression is rather poor, most of the action is in the unobserved error terms. This means that perhaps a more important way in which retirement and wealth could be correlated is through a correlation in the error terms. When this correlation is calculated for individuals who are in both equations, however, the correlation is a mere -0.008.[24] This correlation is in rough agreement with the lack of correlation

we observe with regard to the effects of the observed explanatory variables. Both the explanatory variables and the error terms seem to be saying that there is not much relation between retirement and wealth. This means that a model with heterogeneous retirement preferences, even when coupled with heterogenous time preferences, is inconsistent with the observed pattern of retirement and wealth as long as the preferences are not assumed to be correlated, and that any model that is used for structural estimation should probably include correlated retirement preferences and time preferences, or something similar, to break the implication of correlated retirement and wealth.[25]

Separate Pension and Social Security Effects

The equations presented in Table 5 assume that the effects of accruals and premium values are the same whether they operate through pensions or Social Security. Table 7 presents partial results of an additional probit estimation for the retirement regression equation that splits up the effects of accruals and premium values into separate components for pensions and Social Security. The results for the probit in which these variables are

Table 6.
Patterns of coefficients in the retirement and wealth equations

	Retirement equation		
Wealth equation	Significantly negative	Not significant	Significantly positive
Significantly negative	1	7	1
Not significant	5	11	4
Significantly positive	1	4	1

SOURCE: Authors' calculations.

Table 7.
Comparison of pension and Social Security effects in the retirement probit

Selected independent variable	Pension	Social Security	Combined
Initial accrual/annual earnings	0.0402	-0.0077	0.0348
	(3.59)	(-0.22)	(3.29)
Final accrual/annual earnings	-0.0679	-0.0378	-0.0656
	(-3.42)	(-0.88)	(-3.63)
Premium value/annual earnings	-0.0202	-0.0242	-0.0216
	(-2.17)	(-2.11)	(-2.85)
Log likelihood	-3,073.38		-3,074.94

SOURCE: Authors' calculations.

NOTE: t-statistics are given in parentheses.

combined are repeated in the last column for convenience.

At first glance, the effects of the pension and Social Security variables seem to be different. To be sure, all three pension effects are approximately equal to the effects for the combined variables, both in magnitude and significance. For the Social Security variable, the premium value effect is about the same in both magnitude and significance as for the pension variable, but both Social Security accrual variables are smaller in magnitude and are not significant. The effect for the initial level of the Social Security accrual measure is of an unexpected sign, but the magnitude is very small.

However, the confidence intervals of the final accrual variable for Social Security clearly include the point estimate of the pension variable, and the same thing is nearly true for the initial Social Security accrual variable. This raises the possibility that the two sets of estimates for the pension and Social Security variables are not significantly different and invites a test of the differences. Twice the difference in the log likelihoods is 3.12, which is clearly not significant when compared with a chi squared distribution with three degrees of freedom. Recall from Table 3 that the variability of the Social Security accrual variables is less than half as much as it is for the pension accruals. Evidently the lower variation in the Social Security accruals has led to less accurate estimates of these effects, so that we cannot reject the hypothesis that the effects of the Social Security and pension accruals and premium values are the same.[26]

One final note pertains to the finding that the point estimates of the effects of the premium values are approximately the same for pensions and Social Security. Since the variation in premium values for pensions is wider than the variation in Social Security premium values, especially among the group for whom the premium values are positive, the overall effect of pension premium values on retirement appears to be somewhat larger than the effect of Social Security premium values.[27]

Social Security Acceptance Behavior

We have so far assumed that in calculating Social Security accruals and premium values, those who retire accept their Social Security benefits upon retiring, or will accept them at age 62 if they retire before then. In addition, we have assumed that those who do not retire at all do not accept their benefits until they retire. The top panel of Table 8 indicates that the vast majority of those who are retired do claim their benefits, with the share of claimants increasing with age between 62 and 65. At age 62, 69.5 percent of retirees have accepted benefits. By age 65, the acceptance rate is up to 92 percent. The numbers accepting benefits among the

partially retired are just slightly lower. Among those who are not retired, 11.5 percent claim benefits at age 62, rising to 42.1 percent by age 65.

The second panel of Table 8 shows that, using the Social Security interest rate assumptions, benefit acceptance was optimal only for a modest fraction of those aged 62 to 65 who actually accepted them. Among the completely retired who are 62 years old, more than four-fifths of those who accepted benefits would have increased the present value of their benefits if they had delayed collecting them. Among 64-year-old retirees, only a little more than a third of those accepting benefits should have. The main reason is that for a 62-year-old beneficiary, delaying benefits for a year increases future benefits from 80 percent of the full amount to 86.66 percent. This is an increase of about 8.66 percent in

Table 8.
Social Security benefit acceptance, by age of respondent

Characteristic	62	63	64	65
	By current retirement status			
	Actual benefit acceptance rates (percent)			
Not retired	11.5	22.5	20.4	42.1
Partially retired	65.3	77.4	79.8	89.8
Completely retired	69.5	83.8	88.7	91.8
	Percentage of actual acceptors for whom acceptance was optimal			
Not retired	3.9	9.1	9.1	36.3
Partially retired	13.2	23.9	23.1	63.8
Completely retired	16.6	29.0	37.7	65.7
	By sex and marital status			
	Percentage of all potential recipients who should delay benefit acceptance			
Married men	91.9	9.1	87.2	38.1
Single men	93.6	0	0	0
Married women	55.9	51.1	44.4	40.5
Single women	83.8	83.3	60.2	0
	Average present value of delay among those who would gain from delay (dollars)			
Married men	7,991	5,496	2,684	1,806
Single men	293	c	c	c
Married women	7,786	7,260	7,161	7,220
Single women	1,778	654	92	c

SOURCE: Authors' calculations.

NOTE: These numbers understate the fraction of eligible beneficiaries who accept benefits at age 62. See Olson (1999).

a. Social Security receipt refers to the previous year in 1992, the previous month in 1994, and current receipt in 1996 and 1998.

b. Social Security receipt excludes respondents who currently or previously received Social Security Disability Insurance or Supplemental Security Income before age 65.

c. Negative gain from delay.

benefits, and for a married man, the increase applies not only to his own benefits as long as he is alive but most likely to his wife's widow's benefits for as long as she outlives him. Given their joint life expectancy, an increase of 8.66 percent is more than actuarially fair, and it increases the expected present value of the benefits to delay them. The same thing applies to 64-year-old beneficiaries, though with somewhat less force.

The third panel examines who should postpone acceptance, and it includes all individuals in the sample even if they are not currently eligible because of the earnings test. About 90 percent of married men and half of married women should postpone benefit receipt, and no single men over 63 or single women over 65 should delay. For single individuals, the change in the early retirement penalty is barely actuarially advantageous in the 62-64 age range and certainly not after 65. For married couples, the calculations consider the total lifetime household Social Security value. They assume that it would be advantageous for one spouse to delay accepting benefits, but only if the other spouse does also. In the final panel, we see that among those for whom it would be optimal to postpone benefit acceptance, the average gains from doing so are close to $7,000 in total present value for both men and women at age 62. The value declines with age for married men because of the actuarial factors discussed in the previous paragraph. It remains over $7,000 for married women, even at age 64 or 65, because the calculations assume that when the wife postpones, it is optimal for the husband also to postpone so as to increase not only his own benefits but also the widow's benefits and thus to increase the total present value of benefits.[28]

Evidence that benefit claiming is being driven by liquidity constraints, not by the reward to postponing benefit receipt, can be seen in the first two rows of Table 9. The dependent variable is whether the individual had already claimed benefits on the survey date, and the age variable is the age on that date.[29] Among persons who are retired, those with a higher ratio of nonpension, non-Social Security wealth to Social Security wealth are significantly less likely to have accepted benefits. Moreover, among persons who have fully retired, those with the strongest incentive to postpone benefit receipt, as measured by a higher Social Security premium, are most likely to accept benefits. With the overwhelming majority of those who have retired claiming benefits, these regressions appear to distinguish behavior only among a minority of retirees who are on the margin of claiming benefits, and not to tell a clear and consistent story about what is motivating the overwhelming majority of retirees to claim their benefits earlier than optimal.

To this point we have assumed that even though the older population is failing to postpone benefit receipt so as

to maximize the present value of expected future benefits, they still include the value of delaying a claim to benefits as part of the reward to delayed retirement. To remove the value of that reward, we have reestimated the retirement equation, measuring the Social Security incentives to retire on the assumption that whenever the individual retires, benefits will be claimed at the optimal age. When the regression combines the incentive from Social Security and pensions, the effect of computing Social Security incentives at the optimal retirement age is to drop the coefficient on the premium value by one-third, from -0.022 to -0.014.[30] When pension and Social Security incentives are measured separately, the coefficient on the Social Security premium value is reduced from -0.024 to -0.010, and the t-statistic on the Social Security premium becomes insignificant at -1.08. Accordingly, if increases in Social Security benefits from delaying benefit receipt are not taken into account when deciding on the retirement date, then the size of the reward to delaying retirement is reduced. The effect of each dollar of reward (that is, increased Social Security benefits) on retirement is also reduced to insignificance.

Sensitivity Analysis

Several additional questions might be raised about these results, particularly the retirement equations. In this section, we look at some of these issues.

The first question that might be raised is whether the self-employed respondents are driving the results. Recall that the self-employed have large coefficients, especially in the wealth equation, and that the conditions under which they work may make the retirement decision for these respondents much different from that of the other respondents. The real question is whether the retirement equation will look very different if the self-employed are excluded.

This question is examined in Table 10. The column in Table 10 labeled "Excluding self-employed" estimates the probit only for those who are not self-employed. Compare this with the "Base estimates" column, which estimates the equation for the entire sample. The effects of the two accrual variables are virtually identical, for both the pension and the Social Security versions of the variables. The magnitudes of the premium values for both pensions and Social Security are 20 percent to 25 percent lower with the restricted sample and are no longer significant. However, when the premium-value variables are combined, the resulting variable is significant (this result is not shown in the table). There is no evidence of any difference between the premium-value effect of pensions and the premium-value effect of Social Security.

Another question relates to using some observations where the Social Security values are imputed. In our analysis we exclude observations in which the respondent indicated there was a pension but the pension plan description was not collected from the employer. We do so on the grounds that the imputations of pension incentives (accruals and premium values) from the respondent information only is mostly noise. The same is less true of Social Security, since Social Security operates with a uniform set of rules that are known even if the respondent did not give permission to obtain the Social Security record. Whether including these imputed records has affected the results is an open question, however, since we must still impute the wage history if the Social Security record is missing.

Table 10 also gives the results of the retirement probit when the sample includes only those for whom Social Security records were actually obtained (see the column labeled "Excluding imputed Social Security"). There are some differences here, and the standard errors are generally larger, as one would expect given the reduction in the sample size. The two Social Security accrual variables are still insignificant, but the effect of the Social Security premium-value variable is almost twice as large.

Table 9.
Probits for Social Security acceptance, by retirement status

Independent variable	Not retired		Partially retired		Completely retired	
	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
Liquidity constraint measure [a]	0.0136	0.40	-0.0708	-1.74	-0.0492	-2.91
Social security premium (thousands of dollars)	-0.0030	-0.88	0.0050	1.55	0.0110	6.26
Age						
63	0.1461	4.67	0.0760	2.67	0.1052	8.33
64	0.1628	4.21	0.1297	4.29	0.1330	10.26
65	0.3972	7.69	0.1658	5.33	0.1699	12.78
66	0.6266	9.27	0.1777	5.42	0.1649	11.27
67	0.6207	5.42	0.1760	3.32	0.1468	7.14
Female	0.0395	1.16	0.0133	0.33	-0.0048	-0.32
Married	-0.1184	-2.56	-0.0725	-1.50	-0.0179	-1.12
Age difference if married	-0.0002	-0.07	-0.0036	-1.19	-0.0014	-0.98
Race						
Black	-0.0623	-1.94	-0.0821	-1.95	-0.0940	-5.32
Hispanic	-0.1303	-3.22	-0.1296	-2.01	-0.1101	-4.86
Fair or poor health	0.1380	3.93	-0.0090	-0.24	-0.0374	-2.89
Not available					-0.1979	-0.96
Education						
Less than high school	0.0625	1.94	-0.0256	-0.69	0.0059	0.40
Some college	-0.0437	-1.42	-0.0996	-2.63	0.0005	0.03
Undergraduate degree	-0.0508	-1.26	-0.1519	-3.15	-0.0690	-2.72
Graduate work	-0.1310	-3.76	-0.1779	-3.71	-0.1367	-4.94
Children	0.0837	1.83	0.2268	3.16	-0.0154	-0.71
Not available					0.0256	0.23
Planning horizon						
Next year	0.0906	3.20	0.0291	1.00	-0.0126	-0.97
More than 10 years	0.0796	1.63	-0.0275	-0.59	-0.0166	-0.75
Not available	0.0718	1.16	-0.1266	-1.58	-0.0119	-0.46
Word recall (number of words)	-0.0047	-0.99	0.0074	1.52	0.0061	2.62
Not available	-0.0382	-0.67	0.0646	1.07	-0.0049	-0.17
Share of lifetime household earnings	-0.2573	-4.51	0.0157	0.23	0.0671	4.15
Pseudo R^2		0.19		0.12		0.11
Number of observations		1,446		1,031		4,236

SOURCE Authors' calculations.

NOTE: The liquidity constraint measure is the ratio of nonpension, non-Social Security wealth to Social Security wealth.

a. The probit estimates are the marginal effects on the probability of retirement of a one-unit change in the explanatory variable.

With regard to the pension variables, the premium-value effect and the effect of the final accrual are very close, but the effect of the initial accrual is just less than half as great as in the base estimates.

However, for all six variables, the confidence intervals constructed around the estimates with the restricted sample include the value estimated from the full sample. This leads to the possibility that the effects between the two estimates are not significantly different. One can do the test by estimating over the two subsamples (with and without matched Social Security records) and comparing the log likelihoods with the full sample. When this is done, the test statistic is 69.06, which compares with a 5 percent significance level statistic of 76.88 for 55 degrees of freedom. Thus we would conclude that although the point estimates are different, particularly for a couple of variables, the differences are not statistically significant.

The next question relates to the definition of retirement. In the Data and Variables section, we argued that both objective retirement definitions, such as the one based on hours, and self-reported retirement definitions have problems. We developed a hybrid definition of retirement that combines the information in the objective measures with the self-reports to give what we feel is a more sensible result when the objective measure differs from the self-reported measure. However, we would like to know how sensitive our results are to this approach.

The last two columns of Table 10—labeled "Using self-reported retirement" and "Using objective retire-

ment"—suggest that the point estimates are not too sensitive to the specification of the dependent variable. The significant Social Security accrual variables hold up under the changed specification and the premium variables are of the same sign, but the coefficient estimates slip below significance. There is only one surprise. The coefficient on the final accrual measure for Social Security is positive and almost significant, changing sign in an unexpected direction when the dependent variable is defined using only self-reported status.

Probability of Partial Retirement

To this point we have examined only the flow from full-time work into full retirement. In this section we examine other flows away from full-time work, both the flow from full-time work to any kind of retirement and the flow specifically to partial retirement. Probit estimates for these flows are reported in Table 11, using the same set of explanatory variables as was used in Table 5. The left equation is a probit for leaving full-time work for any retirement, either partial or full, and the right equation is a probit for leaving full-time work for partial retirement only. For some variables, the effect in the partial retirement equation is to amplify the effect in the full retirement equation in Table 5; for others, the effect in the partial retirement equation offsets the effect in the full retirement equation. Although the pattern of significant coefficients is somewhat different for the equation for

Table 10.
Sensitivity tests for retirement probit

Selected independent variable	Base estimates	Excluding self-employed	Excluding imputed Social Security	Using self-reported retirement	Using objective retirement
Pensions					
Initial accrual/annual earnings	0.0402	0.0430	0.0198	0.0355	0.0441
	(3.59)	(3.67)	(1.24)	(3.47)	(3.71)
Final accrual/annual earnings	-0.0679	-0.0626	-0.0512	-0.0634	-0.0704
	(-3.42)	(-3.06)	(-2.38)	(-3.40)	(-3.33)
Premium value/annual earnings	-0.0202	-0.0163	-0.0215	-0.0147	-0.0134
	(-2.17)	(-1.67)	(-1.89)	(-1.71)	(-1.42)
Social Security					
Initial accrual/annual earnings	-0.0077	0.0056	-0.0029	-0.0197	-0.0171
	(-0.22)	(-0.15)	(-0.07)	(-0.63)	(-0.46)
Final accrual/annual earnings	-0.0378	-0.0393	-0.0081	0.0451	-0.0303
	(-0.88)	(-0.80)	(-0.18)	-1.92	(-0.71)
Premium value/annual earnings	-0.0242	-0.0180	-0.0460	-0.0147	-0.0078
	(-2.11)	(-1.48)	(-2.81)	(-1.37)	(-0.72)
Number of observations	8,612	7,377	6,585	8,469	8,513

SOURCE: Authors' calculations.

NOTE: t-statistics are given in parentheses.

Table 11.
Probits for leaving full-time work and for partial retirement

Independent variable	Leaving full-time work		Partial retirement	
	Marginal effect	t-statistic	Marginal effect	t-statistic
Measures of earnings				
Log of annual earnings	-0.0349	-4.90	-0.0160	-4.09
Pension and Social Security incentives				
Initial accrual/annual earnings	0.0271	1.82	-0.0272	-1.98
Final accrual/annual earnings	-0.0671	-3.20	0.0029	0.25
Premium value/annual earnings	-0.0283	-3.05	-0.0068	-1.18
Age binary variables				
51	0.0316	1.05	0.0228	1.16
52	0.0061	0.23	0.0040	0.24
53	0.0557	2.31	0.0029	0.19
54	0.0750	3.30	0.0325	2.15
56	0.0765	3.35	0.0473	3.03
57	0.1137	4.96	0.0678	4.23
58	0.1220	5.18	0.0692	4.19
59	0.1679	6.91	0.0605	3.66
60	0.2987	11.58	0.1276	6.79
61	0.3780	13.79	0.1449	7.20
62	0.3508	10.55	0.1273	5.35
63	0.4574	11.36	0.1830	6.05
64	0.5336	10.12	0.2880	6.88
65	0.4224	4.30	0.1385	2.10
Personal characteristics				
Female	-0.0170	-1.21	-0.0135	-1.58
Married	-0.0196	-1.09	0.0087	0.83
Age difference if married	-0.0015	-1.39	-0.0010	-1.58
Race				
Black	0	0	0	0
Hispanic	-0.0241	1.36	0.0028	0.25
Fair or poor health	0.1061	7.27	-0.0035	-0.42
Education				
Less than high school	0.0001	0	-0.0141	-1.92
Some college	-0.0066	-0.52	0.0040	0.54
Undergraduate degree	-0.0060	-0.35	-0.0041	-0.41
Graduate work	-0.0023	-0.14	0.0031	0.30
Children	0.0053	0.28	-0.0056	-0.47
Planning horizon				
Next year	-0.0009	-0.09	-0.0054	-0.83
More than 10 years	0.0220	1.27	0.0116	1.13
Not available	-0.0390	-1.48	-0.0473	-3.60
Word recall (number of words)	0.0011	0.64	0.0011	1.10
Not available	0.0039	0.13	0.0277	1.39
Share of lifetime household earnings	-0.0926	-3.59	-0.0128	-0.84

Table 11.
Continued

Independent variable	Leaving full-time work		Partial retirement	
	Marginal effect	t-statistic	Marginal effect	t-statistic
Job characteristics				
Self-employed	0.0184	1.08	0.0390	3.68
Not available	0.8118		0.9371	22.77
Years of service	0.0019	4.30	0.0002	0.57
Not available	0.3881	3.00	0.1012	1.37
Industry				
Manufacturing	-0.0071	-0.49	-0.0201	-2.31
Public administration	0.0408	1.96	0.0220	1.65
Occupation				
Management or professional	0.0171	1.23	0.0152	1.81
White collar	0.0094	0.79	0.0123	1.71
Covered by union	0.0299	2.47	0.0141	1.79
Not available	0.0347	0.24	-0.0688	23.09
Covered by pension	-0.0049	-0.38	-0.0266	-3.39
Not available	-0.2030	-12.94	-0.0636	23.49
Firm with more than 100 employees	0.0055	0.39	-0.0097	1.19
Not available	0.0318	1.55	0.0016	0.15
Availability of reduced hours	0.0134	0.95	0.0354	3.99
Laid off from initial job	0.1980	8.67	0.0413	3.07
Pseudo R^2	0.10		0.078	
Number of observations	8,612		8,612	

SOURCE: Authors' calculations.

NOTE: The probit estimates are the marginal effects on the probability of retirement of a one-unit change in the explanatory variable.

leaving full-time work in Table 11 as compared with the full retirement equation in Table 5, there appears to be no more correspondence between these coefficients and the wealth equation than there was for the full retirement equation in Table 5. This implies that the conclusions reached in the section on retirement and wealth regressions are not substantially altered by considering retirement as a move from full-time work to either partial retirement or full retirement.

The coefficient on earnings in the partial retirement equation is negative and is about the same size as the negative coefficient in the equation for full retirement, so that higher earnings are twice as effective in slowing the flow from full-time work, as is suggested by the coefficient in the full retirement equation. Similarly, a higher pension premium reduces the flow into partial retirement, in addition to reducing the flow into full retirement. In contrast, the negative coefficient on the measure of initial benefit accrual in the partial retirement equation offsets to some degree the positive coefficient in the equation for full retirement. The result is that while a high benefit

increment in the initial period increases the flow to full retirement in the following period, the effect on the flow out of full-time work is only three-quarters as much.

Notice next that the age effects are significant and in the same direction in the equations for partial and full retirement, but they are substantially smaller in the partial retirement equation. Among the other independent variables, note that while self-employment reduces the flow from full-time work to complete retirement, it increases the flow into partial retirement by even more. Interestingly, those who are free to reduce hours of work on their jobs are 2 percentage points less likely to move from full-time work into full retirement and are 3.5 percentage points more likely to flow into partial retirement. Having experienced a layoff raises the likelihood of moving into full retirement by 15 percent; it also increases the likelihood of moving from full-time work into partial retirement by another 4 percentage points, altogether increasing the likelihood of leaving full-time work by almost a fifth.

Pitfalls in Using Reduced-Form Retirement Equations for Analyzing Social Security Policies

Our findings suggest that caution is required when using reduced-form equations to evaluate new policies. The basic problem is that because the omitted factors are related systematically both to retirement outcomes and to the measured reward to postponing retirement, a simple retirement equation credits the effects of the omitted factors to the included measures of changes in Social Security benefits. New policies will change the relationship between retirement and the increase in the value of Social Security benefits with postponed retirement, resulting in incorrect predictions of the effects of new policies.

When we fit single-equation retirement models, we found a variety of evidence that important behaviors had been omitted. Consider, for example, variables measuring the age of the respondent in the form of a series of dummy variables, one for each year of age. These age variables suggest a sharp increase in the probability of retirement at age 62. This is a sign that even though the equations include measures of the increase in the value of Social Security with delayed retirement, the cause of the increased retirement at age 62 has not been identified by the model. That is, if the measures of Social Security and pension increases that are incorporated in the reduced-form models captured the full effects of the monetary value of Social Security and pension incentives on retirement, there would be no significant, differential effect of an age 62 dummy variable on retirement. Without knowing the cause of the increased retirement at age 62, we cannot determine how much of the increase is due to the Social Security early retirement age. Since the early retirement age has not changed in decades, it cannot be directly included in an analysis. It is conceivable to introduce other variables that would measure the effect, but one would still have to be wary as long as there was any residual increase in retirement around age 62 in the equation.

The omission of important factors from reduced-form retirement equations creates a major problem for policy simulations. An example is estimating the impact of raising the early retirement age. The estimated relation of the probability of retiring to measures of changes in the value of Social Security benefits when retirement is postponed will be biased. This is readily seen from the coefficients on the pension and Social Security premium variables in the retirement equations we have estimated. The negative coefficients on the premium variables suggest that if the Social Security early retirement age were to be abolished, more people would retire earlier rather than later—a counterintuitive prediction. The reason is that raising the early retirement age to 65 would reduce the reward to continued work, since some of that reward results from deferred claiming of benefits. In other words, with an early retirement age raised to 65, for example, there would be no benefit from deferring claiming associated with the decision to postpone retirement from 62 to 63, 64, or 65. Consequently, given available parameter estimates, a reduced-form model with the expected negative sign on the pension premium or a related variable will predict that raising the early retirement age will, if anything, reduce the age of actual retirement.

Yet intuitively we expect that raising the Social Security early retirement age would have an effect in the opposite direction from that predicted by the reduced-form retirement equation. Many members of the population will defer retirement if Social Security benefits are not available until age 65—some because they are liquidity constrained, others for other reasons.

These findings notwithstanding, the Social Security Administration and others have begun to use single-equation retirement models in an effort to understand the effect of Social Security benefits on retirement behavior and the implications of changing such features as the Social Security early retirement age.

Conclusion

This article began with a simple theoretical model of the relationship between retirement and wealth accumulation. If the only heterogeneity were in retirement preferences, those who retire early would be found to accumulate more wealth, enabling them more closely to maintain consumption in retirement. Heterogeneity in time preferences has much the same effect in inducing a positive correlation between wealth and early retirement. A model with both types of heterogeneity maintains this result as long as the two types of heterogeneity are uncorrelated. Only when heterogeneity in retirement preference is positively correlated with heterogeneity in time preference may the positive relation between early retirement and wealth be broken.

We estimated a reduced-form model of retirement and wealth accumulation and asked whether the variables have corresponding effects in both equations, as would be predicted by the simpler versions of the model. We do find some variables that induce early retirement and that also induce higher wealth. However, in many cases the coefficients do not have comparable effects in retirement and wealth equations. Moreover, the unobservables from the retirement and wealth equations are only weakly correlated. This suggests that more than heterogeneous leisure preferences is required to explain the observed patterns.

There is even more direct evidence of the need for a more comprehensive model of behavior, and of internal inconsistencies in the simple retirement equations that are being estimated. Social Security incentives are often measured by the increment in the value of benefits associated with deferred retirement. The increment in the value of benefits depends, however, on when benefits are claimed. Using SSA interest rate projections, we find that those who retire completely are claiming their benefits too early to be maximizing the expected value of these benefits. Yet measures of Social Security benefit accrual used in these retirement models often include the increased value of benefits from deferred claiming in their measure of the gain to deferring retirement. On the one hand, early retirees are seen not to defer benefit acceptance despite the actuarial advantage. On the other hand, later retirees are said to defer their retirement in order to gain the advantage of deferring benefit acceptance.

In addition to the inconsistencies in assumed behavior, there is the direct evidence of omitted variable bias. Age effects remain unexplained by the measures of incentives from pensions and Social Security that are included in the reduced-form retirement models. Moreover, predictions of the effects of major policies, such as raising the eligibility age for Social Security early retirement, appear to be in the wrong direction.

The evidence gathered in this article will be of help in specifying a proper structural model of retirement and wealth determination. The advantage of estimating structural models is that it is possible to investigate the effects of policy changes, such as increasing the early retirement age, even if those changes have not been observed in the data sets used to estimate the model. The evidence suggests that there are more complexities in behavior than those created by either heterogeneous retirement preferences or heterogeneous time preferences alone, and that correlated heterogeneity in retirement preferences and time preferences is also probably required to generate the observed relations between retirement and wealth.

Heterogeneity in time preferences, combined with liquidity constraints that bind for some individuals, implies that individuals with high time preference and an imperfect ability to borrow may value future income from pensions and Social Security much less than the amounts calculated using the market interest rate. The failure of most retirees to delay claiming Social Security benefits suggests that many individuals value future benefits less than using the interest rate would suggest. This raises questions about the way Social Security and pension benefits are calculated as explanatory variables in reduced-form retirement equations. In a world with heterogeneous time preferences and liquidity constraints,

it may not be appropriate to evaluate payment streams using an interest rate that is constant across individuals. Structural models that allow for the possibility of heterogeneous time preferences and liquidity constraints and apply an internal rate for discounting by those who are liquidity constrained may allow for a more natural treatment of this problem. Analyses in which at least some respondents poorly understand the benefit schedule, or do not value spouse and survivor benefits in accordance with their expected value, may also be appropriate.

These findings are unsettling for public policy analysis. Reduced-form equations, such as the retirement equation or the wealth equation, must be used with great caution in analyzing new policy initiatives. Unobserved heterogeneity interacts with the observable variables to produce the estimated coefficients in these equations, but the comparable interactions are not necessarily the same if the policy changes in new ways. Structural models, which depend on the underlying utility parameters, are less subject to this criticism. But such models are almost certainly more difficult to estimate, and the researcher must incorporate the heterogeneity into the model in sensible (and testable) ways.

Notes

[1] Two important exceptions that do consider the joint relation between retirement and saving behavior are Diamond and Hausman (1984) and Kahn (1988).

[2] See Gustman and Juster (1996) for a discussion of the inconsistencies between the saving, retirement, and pension literatures. In the present article, we control for some factors correlated with precautionary and bequest motives but do not explore them systematically. We also do not consider behavioral reactions to uncertain lifetimes, annuities, the demand for life insurance, and related issues.

[3] Rust and Phelan (1997) is an exception. They assume that the capital market is not operative so that the retirement decision affects the path of consumption.

[4] See Lumsdaine (1996) and Lumsdaine and Mitchell (1999) for recent summaries of retirement research.

5 Blinder, Gordon, and Wise (1980) and Clark and Gohman (1983) have discussed the actuarial advantage of delayed claiming of Social Security benefits. See also Feldstein and Samwick (1992). Coile and others (2002) also suggest that it is optimal for many to delay claiming Social Security benefits after obtaining eligibility. All of these studies suggest that it is advantageous to delay claiming benefits, although the behavioral models differ. Coile and others (2002) emphasize the roles of mortality expectations, wealth, and differences in age and relative earnings between spouses to explain delays in benefit claiming. Despite finding coefficients that are consistent with their theoretical discussion—that those who expect to live longer and who have younger wives delay claiming benefits longer—the authors find that relatively few retirees delay claiming their benefits. Our examination of data in the HRS also reveals very little delayed claiming. For example, using data from the New Beneficiary Data System, Coile and others (2002) find that only 10 percent of those retiring before age 62 delay claiming benefits by 1 year or more. The bottom line is that despite all the arguments given in Coile and others (2002) and in earlier papers that would lead one to expect that early retirees postpone benefit claiming, most early retirees do not.

6 Actuarial returns to Social Security vary with family status and age and may be quite generous at younger ages. Using the Social Security benefit reduction rate on the assumption of normal retirement at age 65, at age 62 a 6.67 percent increase in benefits from delaying retirement for 1 year raises the benefit by 6.67/0.8, or 8.33 percent. Given the life tables and assumed interest rate, that adjustment is better than actuarially fair, at least if one's spouse is not over the age of 65.

7 The purpose of the Pension Equity Act is to protect spouses from circumstances in which the primary earner takes a single life annuity and leaves the spouse with no pension income once he or she dies. For an analogous reason, Congress abolished the earnings test for those over 65 but refused to do so for those between the ages of 62 and 65. See Gustman and Steinmeier (1998) for an analysis of how the weight given to spouse and survivor benefits relative to own benefits might affect the decision to accept a private account under a Social Security privatization scheme with voluntary participation.

8 Many of those who continue to work have the option of immediately claiming some of their benefits, with the remainder postponed because of the earnings test. Gustman and Steinmeier (1991, 742), using 1984 data from the Continuous Work History Survey, found that only 30 percent to 40 percent of working individuals who were eligible for partial benefits at age 62 registered for them.

9 Some other potential extensions to the model—such as inserting a real interest rate into the budget constraint, allowing wages to grow over time, or both—would make the algebra more cluttered but would not affect any of the major conclusions.

10 A negative relation between retirement and wealth may also result if leisure and time preferences are positively correlated, but the effect of leisure on wealth dominates.

11 For self-reported retirement, there is a "not applicable" category, which applies to homemakers and respondents who have not worked for years. Such responses were included in the "completely retired" category.

12 The wage profile coefficients are taken from Anderson, Gustman, and Steinmeier (1999) and are based on data from the Survey of Consumer Finances. Coefficients are experience .0138221, experience squared -.0002827, and experience * education .000996. Note that the wage profiles are not smooth, as they would be if they were based only on the coefficients in the wage equation. Rather, they have sharp discontinuous breaks at points where actual wage observations anchor the profile. Moreover, we do not use wage observations from all years, but only for the number of years worked as reported in the retrospective work history.

13 Reported earnings are used if the Social Security earnings are at the limit in a given year and reported earnings are higher. Reported earnings are also used if the respondent says that the job was a state or federal job or a job not covered by Social Security and the Social Security earnings over the lifetime of the job are less than 60 percent of the reported earnings. If the Social Security earnings are over 60 percent of the reported earnings, we assume that the job is in fact a covered job, and the actual Social Security earnings are used. In point of fact, in most cases in which the Social Security earnings are less than 60 percent of the reported earnings, the Social Security earnings are sporadic and do not amount to more than a few hundred dollars, strongly suggesting that they are coming from some kind of secondary activity and not from a primary job.

14 One can expect a nonlinear relationship between wealth and lifetime earnings on both the low and high ends of the income and wealth distributions. Those with low earnings and wealth are insured against adverse events by a variety of government income- and wealth-tested programs that are not available to those with higher wealth or income (Hubbard, Skinner, and Zeldes 1995). Moreover, the bequest motive, and tax treatment of bequests, may be very different between those at the upper end of the income and wealth distributions and those who have less income and wealth.

15 The fact that the hypothesis that the two accrual coefficients are approximately equal and opposite in sign means that a uniform upward movement in accrual rates, as would occur with a defined contribution pension, would leave retirement relatively unaffected.

16 At the average earnings for the sample, the coefficient on premium value indicates that retirement would decrease by 0.072 percentage points, or about 0.036 percentage points per year since the average period in this study is 2 years, for each $1,000 increase in premium value. This compares with a figure of 0.025 percentage points per year reported by Coile and Gruber (2000, Table 6) for their measure of peak value.

17 See Gustman and Steinmeier (1984) for an analysis of partial retirement both on the main job and on other jobs.

18 When a quartic in family lifetime earnings percentiles is added to the wealth equation, the coefficients are not individu-

ally or jointly significant over and above the log of family lifetime earnings, and the remaining coefficients appear to be hardly affected.

[19] Because the log of lifetime family earnings appears in the denominator of the dependent variable, there will be some downward bias in the coefficient estimated for lifetime earnings because of measurement error. When we fit the wealth regression including only respondents for whom we had Social Security records, and therefore for which any biases arising from errors in measuring lifetime household earnings should be smaller, the coefficient on total lifetime earnings for the family falls from -0.0110 to -0.0161, with a t-statistic of 2.90. This is in the opposite direction from the change that would be caused strictly by measurement error in the lifetime earnings variable, so there must be some systematic difference between the 4,150 observations with an attached Social Security record and the 1,458 observations for families with at least one Social Security record missing. The coefficients on the other covariates are very similar between the two regressions.

[20] These ratios are calculated as of 1992, since the Social Security records provide earnings information up through 1991. For more detailed analysis of the substitution of pensions and wealth, see Gustman and Steinmeier (1999).

[21] We reestimated the wealth equation using median and robust regressions. Among the differences in the significant coefficients, the coefficient on the log of lifetime family earnings turned from small and negative (-0.0110) in the ordinary least squares equation to small and positive in the robust regression (0.0053); and the coefficient on pension value over lifetime earnings turned from small and negative (-0.0548) to small and positive (0.0296).

[22] When we run the regression with a cutoff point of 0.75 for the ratio of wealth to income rather than 1.0, the only one of the key coefficients (measuring the effects of lifetime family earnings, ratios of Social Security and pension wealth to total wealth, and measures of the increments in wealth from additional work) to undergo a significant change is the coefficient on log of lifetime family earnings. Otherwise, the coefficients are not significantly different from those reported for the wealth regression in Table 5.

[23] The omission of the pension and Social Security wealth variables from the retirement probit does not affect these results; when we run the retirement probit adding these variables, there is hardly any change in the coefficients.

[24] For the retirement probit, the error term used in the correlation is either 1 or 0, depending on whether the respondent actually retired, minus the fitted probability of retirement from the estimated probit.

[25] Such a model should also include the other major motivations for saving, as outlined in Gustman and Juster (1996).

[26] On the other hand, when we run a probit without the Social Security accruals, a likelihood ratio test indicates that these two coefficients are not jointly significant.

[27] We also estimated wealth equations with separate measures of pension and Social Security accruals. The

findings were similar to the wealth equation reported in Table 5. The only significant coefficient is for the initial accrual created by pensions.

[28] The observations are respondents as of the four survey dates, and consequently the sample changes from one age to the next. In the last panel, the sample also changes because of differential percentages with positive gains. These changes account for the fact that the amounts in the last panel, which should show a monotonic decrease for a particular individual, do not always do so.

[29] To minimize the recall bias as to exactly when the individual started to receive benefits, the dependent variable is defined as whether the individual had already claimed benefits on the survey date instead of a measure of the age of claiming.

[30] The t-statistic on the measure of the combined premium declines from -2.85 to -2.08. The coefficients on the two delta measures change only very slightly.

References

Anderson, Patricia; Alan Gustman; and Thomas Steinmeier. 1999. "Trends in Male Labor Force Participation and Retirement: Some Evidence on the Role of Pensions and Social Security in the 1970s and 1980s." *Journal of Labor Economics* 17(4), Part 1: 757-783. October.

Blinder, Alan S.; Roger H. Gordon; and Donald E. Wise. 1980. "Reconsidering the Work Disincentive Effects of Social Security." *National Tax Journal* 33(4): 431-442.

Clark, Robert L., and Stephan F. Gohman. 1983. "Retirement and the Acceptance of Social Security Benefits." *National Tax Journal* 36(4): 529-534.

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

Coile, Courtney, and Jonathan Gruber. 2000. "Social Security and Retirement." NBER Working Paper No. 7830. Cambridge, Mass.: National Bureau of Economic Research. August.

_____. 2001. "Social Security Incentives for Retirement." In *Themes in the Economics of Aging*, edited by David A. Wise. Chicago: University of Chicago Press, 311-341.

Diamond, Peter A., and Jerry A. Hausman. 1984. "Individual Retirement and Savings Behavior." *Journal of Public Economics* 23: 81-114.

Feldstein, Martin, and Andrew Samwick. 1992. "Social Security and Marginal Tax Rates." *National Tax Journal* 45: 1-22.

Gale, William G. 1998. "The Effects of Pensions on Wealth: A Re-evaluation of Theory and Evidence." *Journal of Political Economy* 106: 706-723.

Gustman, Alan L., and F. Thomas Juster. 1996. "Income and Wealth of Older American Households: Modeling Issues for Public Policy Analysis." In *Assessing Knowledge of Retirement Behavior,* edited by Eric A. Hanushek and Nancy Maritato. Washington, D.C.: National Academy Press, 11-60.

Gustman, Alan L.; Olivia S. Mitchell; and Thomas L. Steinmeier. 1995. "Retirement Measures in the Health and Retirement Study." *Journal of Human Resources* 30 (Supplement): S57 to S83.

Gustman, Alan L.; Olivia S. Mitchell; Andrew A. Samwick; and Thomas L. Steinmeier. 1999. "Pension and Social Security Wealth in the Health and Retirement Study." In *Wealth, Work and Health, Innovations in Measurement in the Social Sciences*, edited by J. Smith and R. Willis. Ann Arbor: University of Michigan Press, 150-207.

Gustman, Alan L., and Thomas L. Steinmeier. 1984. "Partial Retirement and the Analysis of Retirement Behavior." *Industrial and Labor Relations Review* 37(3): 403-415.

_____. 1989. "An Analysis of Pension Benefit Formulas, Pension Wealth and Incentives from Pensions." *Research in Labor Economics* 10: 53-106.

_____. 1991. "Changing Social Security Rules for Workers Over 65: Proposed Policies and Their Effects." *Industrial and Labor Relations Review* 44(4): 733-745.

_____. 1993. "Pension Portability and Labor Mobility: Evidence from the Survey of Income and Program Participation." *Journal of Public Economics* 50: 299-323.

_____. 1995. *Pension Incentives and Job Mobility.* Kalamazoo, Mich.: W.E. Upjohn Institute for Employment Research.

_____. 1998. "Privatizing Social Security: First-Round Effects of a Generic, Voluntary, Privatized U.S. Social Security System." In *Privatizing Social Security*, edited by Martin Feldstein. Chicago: University of Chicago Press, for the National Bureau of Economic Research, 313-357.

_____. 1999. "Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study." *Carnegie-Rochester Conference Series on Public Policy* 50: 271-326.

_____. 2000. "Employer-Provided Pension Data in the NLS Mature Women's Survey and in the Health and Retirement Study." *Research in Labor Economics* 19: 215-252.

_____. 2001. "Retirement Outcomes in the Health and Retirement Study." *Social Security Bulletin* 64(4): 57-71.

_____. Forthcoming. "What People Don't Know About Their Pensions and Social Security: An Analysis Using Linked Data from the Health and Retirement Study." In *Public Policies and Private Pensions*, edited by William G. Gale, John B. Shoven, and Mark J. Warshawsky. Washington, D.C.: Brookings Institution.

Hubbard, R. Glenn; Jonathan Skinner; and Stephen P. Zeldes. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103: 360-399.

Ippolito, Richard A. 1986. *Pensions, Economics and Public Policy*. Homewood, Ill.: Dow Jones-Irwin, for the Pension Research Council.

Kahn, James A. 1988. "Social Security, Liquidity, and Early Retirement." *Journal of Public Economics* 35: 97-117.

Lazear, Edward P., and Robert L. Moore. 1988. "Pensions and Mobility." In *Pensions in the U.S. Economy,* edited by Zvi Bodie, John Shoven, and David Wise. Chicago: University of Chicago Press, 163-188.

Lumsdaine, Robin L. 1996. "Factors Affecting Labor Supply Decisions and Retirement Income." In *Assessing Knowledge of Retirement Behavior,* edited by Eric A. Hanushek and Nancy Maritato. Washington, D.C.: National Academy Press, 61-115.

Lumsdaine, Robin L., and Olivia S. Mitchell. 1999. "New Developments in the Economic Analysis of Retirement." In *Handbook of Labor Economics,* edited by Orley Ashenfelter and David Card. Amsterdam: North Holland, 3261-3308.

Mitchell, Olivia S. 1988. "Worker Knowledge of Pension Provisions." *Journal of Labor Economics* 6(1): 28-39.

Olson, Janice A. 1999. "Who is "62 Enough"? Identifying Respondents Eligible for Social Security Early Retirement Benefits in the Health and Retirement Study." *Social Security Bulletin* 62(3): 51-56.

Rust, John, and Christopher Phelan. 1997. "How Social Security and Medicare Affect Retirement Behavior." *Econometrica* 65(4): 781-831.

Stock, James, and David Wise. 1990a. "Pensions, the Option Value of Work, and Retirement." *Econometrica* 58(5): 1151-1180.

_____. 1990b. "The Pension Inducement to Retire: An Option Value Analysis." In *Issues in the Economics of Aging,* edited by David A. Wise. Chicago: University of Chicago Press, 205-224.