

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Themes in the Economics of Aging

Volume Author/Editor: David A. Wise, editor

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-90284-6

Volume URL: <http://www.nber.org/books/wise01-1>

Publication Date: January 2001

Chapter Title: Social Security Incentives for Retirement

Chapter Author: Courtney Coile, Jonathan Gruber

Chapter URL: <http://www.nber.org/chapters/c10332>

Chapter pages in book: (p. 311 - 354)

Social Security Incentives for Retirement

Courtney Coile and Jonathan Gruber

One of the most striking labor force phenomena of the second half of the twentieth century has been the rapid decline in the labor force participation rate of older men. In 1950, for example, 81 percent of sixty-two-year-old men were in the labor force; by 1995, this figure had fallen to 51 percent, although it has rebounded slightly in the past few years (Quinn 1999). Declines have been seen for all groups of older men, as illustrated by figure 10.1. For women, these declines with age have been offset by an overall rising trend in labor force participation, as shown in figure 10.2.¹

Much has been written about the proximate causes of this important trend among older men, and in particular about the role of the Social Security program. A large number of articles have documented pronounced spikes in retirement at ages sixty-two and sixty-five, which correspond to the early and normal retirement ages for Social Security, respectively. While there are some other explanations for a spike at age sixty-five, such as entitlement to health insurance under the Medicare program or rounding error in surveys, there is little reason to see a spike at sixty-two other than the Social Security program. Indeed, as Burtless and Moffitt (1984) document, this spike at age sixty-two emerged only after the early retirement eligibility age for men was introduced in 1961.

Courtney Coile is assistant professor of economics at Wellesley College and a faculty research fellow at the National Bureau of Economic Research. Jonathan Gruber is professor of economics at the Massachusetts Institute of Technology and a research associate and director of the program on children at the National Bureau of Economic Research.

This work builds on earlier joint research with Peter Diamond, and we are grateful to him for his continuing guidance throughout this project. We are also grateful to Dean Karlan for excellent research assistance, and to the National Institute on Aging for financial support.

1. Both figures are from Diamond and Gruber (1998).




Fig. 10.1 Historical trends in labor force participation of older men
 Source: Data from Bureau of Labor Statistics website (<http://www.bls.gov>)




Fig. 10.2 Historical trends in labor force participation of older women
Source: See fig. 10.1.

The presence of these strong patterns in retirement data suggests that Social Security is playing a critical role in determining retirement decisions. But in order to model the impact of Social Security reform on retirement behavior, it is critical to understand what this role is. The evidence of spikes at age sixty-two, for example, is consistent with at least three alternative hypotheses. The first is that there is an actuarial unfairness built into the system penalizing work past age sixty-two, so that there is a tax effect that leads workers to leave at that age. The second is that workers are liquidity constrained; they would like to retire before age sixty-two, but cannot because they are unable to borrow against their Social Security benefits and have no other sources of retirement support. In this case, there will be a large exit at age sixty-two as benefits first become available. The third explanation is that workers are information constrained or myopic; they either do not understand or do not appreciate the actuarial incentives for additional work past age sixty-two, so they retire as soon as benefits become available.

The existing evidence would appear to refute the first explanation. Diamond and Gruber (1998) calculate for a typical individual the implicit tax on continued work at each age from the Social Security system and find that there is actually a small subsidy to continued work at age sixty-two. There is some supportive evidence for the second view; Kahn (1988) finds a pronounced spike in the retirement hazard at age sixty-two for those with low wealth, but that the much larger spike is at age sixty-five for those with higher wealth. There is little work on the third view, other than a recent careful exposition of the model by Diamond and Koszegi (1999).

This paper provides a more thorough investigation of the first effect, the tax effect, along four dimensions. First, we assess whether the tax rate Diamond and Gruber compute using a synthetic individual with annual earnings at the median of his cohort is similar to the tax rate of the real median person. We might expect a difference, as the shape of the earnings history is a significant determinant of Social Security incentives through the dropout-years provision, and this is not appropriately reflected with a synthetic earnings history. Second, we assess the *distribution* of retirement incentives across the population. Even if there is no significant disincentive for the typical worker, disincentives for a large subset of workers could still be associated with a spike in the aggregate retirement data. Third, we assess the importance of considering incentives for retirement in the next year versus incentives for retirement over all possible years, drawing on the insights of the option value model of Stock and Wise (1990a,b). Finally, we incorporate the role of private pensions, an important determinant of retirement for a large share of workers.

Our strategy is to apply the model of Diamond and Gruber to a set of real individuals, the older persons surveyed by the Health and Retirement Survey (HRS). This is a very rich survey with information on individual

Social Security earnings histories, private pension plan details, and demographics. These data allow us to compute carefully the incentives for retirement from Social Security and pensions, both for the median individual and across the distribution.

Our paper proceeds as follows. We begin, in section 10.1, with background on the relevant institutional features of the Social Security system and the previous literature in this area. In section 10.2 we describe our data and empirical strategy. Section 10.3 presents our basic results for the accrual of Social Security wealth with additional work and the associated tax/subsidy relative to potential earnings, both on average and across the distribution. Section 10.4 then highlights the fundamental weakness of simple one-year accrual measures of this type: Many Social Security wealth trajectories are nonmonotonic, suggesting that the appropriate measure must look across all years to find the optimal retirement date. We then present calculations for what we label “peak value,” an incentive measure that provides a middle ground between accrual and the utility-based option value metric of Stock and Wise (1990a,b) by comparing retirement wealth at the current retirement date to retirement wealth at its global maximum. In this section, we also extend the results to incorporate private pensions. Section 10.5 concludes by discussing the implications of our findings and the directions for future research.

10.1 Background

10.1.1 Institutional Features of Social Security

The Social Security system is financed by a payroll tax that is levied equally on workers and firms. The total payroll tax paid by each party is 7.65 percentage points; 5.3 percentage points are devoted to the Old-Age and Survivors Insurance (OASI) program, with 0.9 percentage points funding the Disability Insurance (DI) system and 1.45 percentage points funding Medicare’s Hospital Insurance (HI) program.² The payroll tax that funds OASI and DI is levied on earnings up to the taxable maximum, \$72,600 in 1999; the HI tax is uncapped.

Individuals qualify for an OASI pension by working for forty quarters in covered employment, which now encompasses most sectors of the economy. Benefits are determined in several steps. The first step is computation of the worker’s averaged indexed monthly earnings (AIME), which is one-twelfth the average of the worker’s annual earnings in covered employment, indexed by a national wage index. A key feature of this process is

2. The total OASI + DI contribution rate has been 6.20 percent since 1990, although the division between the two parts has varied slightly from year to year; the OASI portion is 5.35 percent in 1999 and will be 5.30 percent starting in 2000.

that additional higher earnings years can replace earlier lower earnings years, since only the highest thirty-five years of earnings are used in the calculation (the “dropout year” provision).³

The next step of the benefits calculation is to convert the AIME into the primary insurance amount (PIA). This is done by applying a three-piece linear progressive schedule to an individual’s average earnings, whereby ninety cents of the first dollar of earnings is converted to benefits, while only fifteen cents of the last dollar of earnings (up to the taxable maximum) is so converted. As a result, the rate at which Social Security replaces past earnings (the replacement rate) falls with the level of lifetime earnings. Although up to 85 percent of Social Security benefits are subject to tax for retirees with sufficiently high incomes (couples with non-Social Security income above \$32,000 in 1999), all of earnings are taxed (including the employee portion of the payroll tax), raising the effective replacement rate of the program.

The final step is to adjust the PIA based on the age at which benefits are first claimed. For workers commencing benefit receipt at the normal retirement age (NRA; currently sixty-five, but legislated to increase slowly to age sixty-seven), the monthly benefit is the PIA. For workers claiming before the NRA, benefits are decreased by an actuarial reduction factor of five-ninths of one percent per month; thus, a worker claiming on his sixty-second birthday receives 80 percent of the PIA.⁴ Individuals can also delay the receipt of benefits beyond the NRA and receive a delayed retirement credit (DRC). For workers reaching age sixty-five in 1999, an additional 5.5 percent is paid for each year of delay; this amount will steadily increase until it reaches 8 percent per year in 2008.

While a worker may claim as early as age sixty-two, receipt of Social Security benefits is conditioned on the earnings test until the worker reaches age seventy. A worker aged sixty-two to sixty-five may earn up to \$9,600 in 1999 without the loss of any benefits; then benefits are reduced \$1 for each \$2 of earnings above this amount. For workers aged sixty-five to sixty-nine, the earnings test floor is \$15,500 and benefits are reduced at a rate of \$1 for each \$3 in earnings. Months of benefits lost through the earnings test are treated as delayed receipt, entitling the worker to a DRC on the lost benefits when he or she does claim benefits. Despite this, the earnings test appears to have a pronounced effect on retirement decisions,

3. In particular, although earnings through age fifty-nine are converted to real dollars for averaging, earnings after age sixty are treated nominally. There is a two-year lag in availability of the wage index, calling for a base in the year in which the worker turns sixty in order to compute benefits for workers retiring at their sixty-second birthdays. Although it would be possible to make adjustments as data become available, this is not done. This gap would become important if we had large and varying inflation rates.

4. The reduction factor will be only five-twelfths of 1 percent for months beyond thirty-six months before the NRA, which will become relevant once the delay in NRA becomes effective.

with evidence of extreme piling-up of the earnings distribution among elderly workers at the earnings test limit (Friedberg 1998).

One of the most important features of Social Security is that it also provides benefits to dependents of covered workers. Spouses of Social Security beneficiaries receive a dependent spouse benefit equal to 50 percent of the worker's PIA, which is available once the worker has claimed benefits and the spouse has reached age sixty-two; however, the spouse receives only the larger of this and his or her own entitlement as a worker.⁵ Dependent children are also each eligible for 50 percent of the PIA, but the total family benefit cannot exceed a maximum, which is roughly 175 percent of the PIA. Surviving spouses receive 100 percent of the PIA, beginning at age sixty, although there is an actuarial reduction for claiming benefits before age sixty-five or if the worker had an actuarial reduction. In practice, estimating a family's total benefits is complicated by the fact that both spouses may qualify for Social Security benefits as retired workers. Finally, benefit payments are adjusted for increases in the consumer price index (CPI) after the worker has reached age sixty-two; thus, Social Security provides a real annuity.

10.1.2 Previous Related Literature

There are two broad strands of the literature on Social Security that are related to this paper. The first strand attempts to document the labor force disincentives inherent in Social Security, or implicit Social Security "tax rates." Feldstein and Samwick (1992) model the tax rates on the marginal earnings decision for simulated workers of different ages, earnings, and marital status. They find that there are significant marginal tax rates on earnings for higher-income workers and secondary earners, and for younger workers as well.⁶

A subsequent paper by Diamond and Gruber (1998) focuses more directly on tax rates around the time of retirement. They build a simulation model similar to that used here and compute Social Security tax rates for simulated workers. As noted above, they find that for the median worker, there is little net incentive or disincentive for continued work at age sixty-two, although there is a sizeable positive tax rate at age sixty-five and beyond due to the unfair DRC still in place. They also find that tax rates are higher for single workers, because they do not benefit from dependent and survivors benefits, and that tax rates are initially lower for low earners (who benefit from the redistributive nature of benefits) but eventually higher (since they are penalized more by actuarial unfairness after age sixty-five).

5. Spousal benefits can begin earlier if there is a dependent child in the household; spousal benefits are also subject to actuarial reduction if receipt commences before the spouse's NRA.

6. Earlier work by Blinder, Gordon, and Wise (1980, 1981) and Burkhauser and Turner (1981) calculates tax rates under the pre-1977 Social Security rules.

While suggestive, both of these studies suffer from a key limitation: They do not consider the incentives facing real individuals. This is important because of the dropout year provision, which implies that the actual pattern of earnings, and not just the level of average or final earnings, matters for benefits determination. As we will show later, even for workers with the same average and final earnings, there is considerable heterogeneity in Social Security tax rates. By considering a real sample of individuals, we will be able to measure appropriately both the incentives for the median worker and the underlying heterogeneity in these incentives.

The second literature is that on the retirement effects of Social Security. A number of studies use aggregate information on the labor force behavior of workers at different ages, such as that documented in the introduction, to infer the role that is played by Social Security. Hurd (1990) and Ruhm (1995) emphasize the spike in the age pattern of retirement at age sixty-two; as Hurd (1990) states, "there are no other institutional or economic reasons for the peak" (597). Using precise quarterly data, Blau (1994) finds that almost one-fourth of the men remaining in the labor force at their sixty-fifth birthdays retire within the next three months; this hazard rate is more than 2.5 times as large as the rate in surrounding quarters. However, Lumsdaine and Wise (1994) document that this penalty alone cannot account for this excess retirement at age sixty-five, nor can the incentives embedded in private pension plans or the availability of retirement health insurance through the Medicare program. This does not rule out a role for Social Security; by setting up the focal point of a normal retirement age, the program may be the causal factor in explaining this spike.

The main body of the retirement-incentives literature attempts to model specifically the role that potential Social Security benefits play in determining retirement. The general strategy followed by this literature is to use microdata sets with information on potential Social Security benefit determinants (earnings histories) or ex post benefit levels to measure the incentives to retire across individuals in the data.⁷ Then, retirement models are estimated as a function of these incentive measures. While the exact modeling technique differs substantially across papers,⁸ the conclusions drawn are fairly similar: Social Security has large effects on retirement,

7. The data used are generally the Retirement History Survey (Boskin and Hurd 1978; Burtless 1986; Burtless and Moffitt 1984; Hurd and Boskin 1984; Fields and Mitchell 1984; Blau 1994), although some authors have relied on the National Longitudinal Survey of Older Men (Diamond and Hausman 1984), and recent work uses the Survey of Consumer Finances (Samwick 1998).

8. The earliest studies (Boskin and Hurd 1978; Fields and Mitchell 1984) used standard linear or nonlinear regression techniques. Later research (Burtless 1986; Burtless and Moffitt 1984) used nonlinear budget constraint estimation to capture the richness of Social Security's effects on the opportunity set. The most recent work (Diamond and Hausman 1984; Hausman and Wise 1985; Samwick 1998; Blau 1994) uses dynamic estimation of the retirement transition.

but the effects are small relative to the trends over time documented in figures 10.1 and 10.2. For example, Burtless (1986) found that the 20 percent benefit rise of the 1969–72 period raised the probability of retirement at sixty-two and sixty-five by about 2 percentage points. Over this period, however, the labor force participation of older men fell by more than 6 percent, so that Social Security can explain only about one-third of the change.⁹

This literature suffers from two important limitations. First, the key regressor, Social Security benefits, is a nonlinear function of past earnings, and retirement propensities are clearly correlated with past earnings levels. This problem is common to the social insurance literature in the United States.¹⁰ For other social insurance programs, however, there is often variation along dimensions arguably exogenous to individual tastes, such as different legislative regimes across locations or within locations over time, that can be used to identify behavioral models. There is no comparable variation in Social Security, which is a nationally homogeneous program. Of course, this criticism does not necessarily imply that the estimates of this cross-sectional literature are flawed; as Hurd (1990) emphasizes, the nonlinearities in the Social Security benefits determination process are unlikely to be correlated with retirement propensities. However, there has been little serious effort to decompose the sources of variation in Social Security benefits in an effort to assess whether the determinants that drive retirement behavior are plausibly excluded from a retirement equation.¹¹

This criticism is levied most compellingly by Krueger and Pischke (1992), who note that there is a unique “natural experiment” provided by the end of double-indexing for the “notch generation” that retired in the late 1970s and early 1980s. For this cohort, Social Security benefits were greatly reduced relative to what they would have expected based on the experience of the early to mid-1970s. Yet, the dramatic fall in labor force participation continued unabated in this era. This raises important questions about the identification of this cross-sectional literature.

The second problem with this literature is that it generally focuses on only one of the two key Social Security benefits variables, including Social Security benefits or wealth but ignoring the Social Security tax/subsidy

9. One exception is Hurd and Boskin (1984), who claim that the large benefits increases of the 1968–73 period can explain all of the change in labor force participation in those years.

10. See Meyer (1989) for a careful discussion of this issue in the context of Unemployment Insurance (UI).

11. At a minimum, one would want to include the level of lifetime earnings as a regressor, but most studies include only earnings in a recent year (i.e., Boskin and Hurd 1978; Burtless 1986). In addition, even using a somewhat longer time frame for measuring the earnings control (as do Diamond and Hausman 1984) does not solve the problem; one could imagine that certain features of the lifetime pattern of earnings are correlated with both benefit levels and retirement decisions, such as the ratio of earnings around age sixty-two to earnings at earlier ages (since individuals who have relatively high earnings at older ages may have better labor market opportunities around the age of retirement and may therefore work longer).

rate documented above. In theory, as discussed above, both of these factors play an important role in determining retirement behavior. Studies that include the tax/subsidy rate find it to have a significant role in explaining retirement (Fields and Mitchell 1984; Samwick 1998); indeed, even in Krueger and Pischke's (1992) paper the accrual rate is often right-signed and significant, even as the wealth effect is insignificant. More recently, Stock and Wise (1990a,b) note that the correct regressor for considering both Social Security and pension incentives for retirement is not the year-to-year accrual rate, but the return to working this year relative to retiring at some future optimal date.

Our findings are relevant to addressing both of these shortcomings. To the extent that we find substantial variation in the retirement incentives facing workers under the Social Security system, even after conditioning on correlates of the retirement decision such as earnings, it suggests that there are significant nonlinearities in the determination of Social Security incentives that can help identify retirement impacts. We will also compare the retirement incentives over the subsequent year with those over all future years, following the insights of Stock and Wise.

10.2 Data and Empirical Strategy

10.2.1 Data

Our data for this analysis comes from the Health and Retirement Study (HRS), a survey of individuals aged fifty-one to sixty-one in 1992 with reinterviews every two years. The first two waves of the survey (1992 and 1994) and preliminary data for the third and fourth waves (1996 and 1998) are available at this time. Spouses of respondents are also interviewed, so the total age range covered by the survey is much wider.

A key feature of the HRS is that it includes Social Security earnings histories back to 1951 for most respondents. This provides two advantages for our empirical work. First, it allows us appropriately to calculate benefit entitlements, which depend (through the dropout year provision) on the entire history of earnings.¹² Second, it allows us to construct a large sample of person-year observations by using the earnings histories to compute Social Security retirement incentives and labor force participation at each age. We use all person-year observations on men aged fifty-five to sixty-nine for our analysis, subject to the exclusions detailed below.

Our sample is selected conditional on working, so that we examine the incentives for retirement conditional on being in the labor force. Work is

12. Only earnings since 1950 are required to compute Social Security benefits for our sample's age range; the benefit rules specify that a shorter averaging period is used for persons born prior to 1929.

Table 10.1 Sample for Analysis

Category	Number of Obs		Obs Lost	
	Obs	Person-Year Obs	Obs	Person-Year Obs
Men in HRS aged 55–69, 1980–97	6,173	40,614	—	—
Drop if born before 1922	6,052	39,658	121	956
Drop if missing earnings history	4,305	29,110	1,747	10,548
Drop if not working	3,445	20,059	860	9,051
Drop if missing spouse's earnings history	3,231	18,903	214	1,156
Drop if reentered labor force	3,231	17,547	—	1,356

Note: Obs is the number of persons for which we have Health and Retirement Survey (HRS) data. Person-Year Obs is the number of person-year observations for which we have data. Each row in the first two columns shows the number of observations after the exclusion labeled in that row. Each row in the second two columns shows the number of observations lost through the exclusion labeled in that row.

defined in one of two ways. For those person-years before 1992, when we are using earnings histories, we define work as positive earnings in two consecutive years; if earnings are positive this year but zero the next, we consider the person to have retired this year. For person-years from 1992 onward, when we have the actual survey responses, we cannot use this earnings-based definition, since we have earnings at two year intervals only. For this era, we use information on self-reported retirement status and dates of retirement to construct annual retirement measures. Although these are somewhat different constructs, the retirement rates by age are similar across the two samples, so we combine them for precision purposes. We also consider individuals before their first retirement only; if a person who is categorized as retired reenters the labor force, the later observations are not used.

Our sample selection criteria are documented in table 10.1. There are 6,173 men who participate in one or more waves of the HRS. We exclude 121 men who were born before 1922 and thus are subject to different Social Security benefit rules. We lose an additional 1,747 men due to a lack of Social Security earnings history data.¹³ We lose 860 men who ceased working prior to age fifty-five, and an additional 214 men due to a lack of information on their wives' Social Security earnings histories (necessary due to the family structure of benefits).¹⁴ The 3,231 remaining observa-

13. Individuals were required to sign a permission form in order for their Social Security records to be attached; approximately 75 percent of the sample gave permission. Haider and Solon (1999) find that willingness to give permission varies only weakly with observable characteristics.

14. We keep observations for which the wife's Social Security earnings records are not available, but we can ascertain from the self-reported labor force histories that the wife worked less than half as many years as the husband.

tions are converted into 18,903 person-year observations by creating one observation for each year from 1980 through 1997 in which the individual is between the ages of fifty-five and sixty-nine and working. Finally, we lose 1,356 person-year observations where the individual is working after a previous retirement. The final sample size is 17,547 person-year observations.

10.2.2 Empirical Strategy

Our goal is to measure the retirement incentives inherent in Social Security and private pension systems. For the case of Social Security, we begin with the calculation of an individual's Social Security wealth. The basis for this calculation is a simulation model that we have developed to compute for any individual his or her Social Security entitlement for any age of retirement. This is based on a careful modeling of Social Security benefits rules, and our simulation model has been cross-checked against the Social Security Administration's ANYPIA model for accuracy.

The next step in our simulation is to take these monthly benefit entitlements and compute an expected net present discounted value of Social Security wealth. This requires projecting benefits out until workers reach age 120, and then taking a weighted sum that discounts future benefits by both the individual discount rate and the probability that the worker will live to a given future age. Our methodology for doing so is described in the appendix. For the worker himself, this is fairly straightforward; it is simply a sum of future benefits, discounted backward by time-preference rates and mortality rates. For dependent and survivor benefits it is more complicated, since we must account for the joint likelihood of survival of the worker and the dependent. In our base case, we use a real discount rate of 3 percent. To adjust for mortality prospects, we use the sex/age-specific U.S. lifetables from the 1995 OASDI Trustees Report (intermediate assumption case).¹⁵ All figures are discounted back to age fifty-five by both time preference rates and mortality risk.

For the output of the simulations, we calculate several different concepts. The first is the level of Social Security wealth. The second is the accrual, or the dollar change in Social Security wealth from the previous year. We then compute the "after-tax accrual," which subtracts from this dollar change the payroll taxes paid by the worker and his employer (assuming full tax incidence on wages). Finally, since it is natural to think about these incentives relative to the returns from additional work, we also follow Diamond and Gruber (1998) in calculating the implicit tax/subsidy rate on additional work, which normalizes the negative of the accrual by the

15. In principle, individual-specific mortality prospects should be used to compute SSW and related retirement incentives. In future work, we plan to use the richer information in the HRS on health and even subjective mortality evaluation to do these richer calculations.

potential wage for that year; a positive accrual implies a negative tax rate and vice versa. Thus, if the tax rate is positive, it implies that the Social Security system causes a disincentive to additional work through foregone Social Security wealth. To measure the full tax wedge, we use the gross wage in the denominator; under the assumption that the employer portion of payroll taxes is reflected in wages, we increase reported wages by 6.2 percent.

For assessing the accrual rate and related concepts used later in the paper, we must project the worker's earnings over the next year (or all future years) if he continues to work. We considered a number of different projection methodologies, and found that the best predictive performance was from a model that simply grew earnings from the last observation by 1 percent real growth per year, so this is the assumption we use for our simulations.

For the purposes of the simulations below, we assume that workers claim Social Security benefits at the point of retirement, or when they become eligible if they retire before the point of eligibility. In fact, this is not necessarily true; retirement and claiming are two distinct events, and for certain values of mortality prospects and discount rates it is optimal to delay claiming until some time after retirement (due to the actuarial adjustment of benefits). Coile et al. (2001) investigate this issue in some detail, using simulation analysis to document the gains to delaying claiming and showing that a nontrivial share of individuals do delay claiming past age sixty-two. In this case, our calculations will overstate any subsidies to continued work, since part of this subsidy will come from delayed claiming that could be obtained without delaying retirement.

Also, it is important to highlight that our work is focused on the impact of Social Security on the labor force participation decision. A separate and interesting issue is the impact of Social Security on the marginal labor supply decision among those participating in the labor force, which was the focus of the Feldstein and Samwick (1992) analysis. This is more complicated for those around retirement age, since it involves incorporating the role of the earnings test, which we avoid with our analysis of participation. This, in turn, would involve modeling expectations about the earnings test, since individuals appear not to understand that this is only a benefits delay instead of a benefits cut. This is clearly a fruitful avenue for further research.

10.3 Social Security Accruals and Tax/Subsidies

10.3.1 Median Worker

We begin by considering the incentives facing the median worker at each age. These results are presented in table 10.2. Each row represents

Table 10.2 Accrual and Tax Rate, Medians by Age

Age	Obs	SSW	Benefit Accrual	After-Tax Accrual	Tax/Subsidy Rate	Diamond- Gruber Tax Rate
55	2,811	154,928	2,277	-933	0.047	-0.022
56	2,746	157,205	2,136	-1,136	0.052	0.046
57	2,444	159,341	1,958	-1,314	0.057	0.060
58	2,131	161,299	1,791	-1,517	0.061	0.069
59	1,822	163,090	1,687	-1,781	0.067	0.072
60	1,547	164,777	1,563	-1,848	0.073	0.071
61	1,252	166,340	1,643	-1,848	0.073	0.064
62	1,010	167,983	3,855	-48	0.002	-0.028
63	688	171,838	4,019	46	-0.001	-0.005
64	443	175,857	2,849	-843	0.027	0.031
65	313	178,706	-902	-4,831	0.145	0.188
66	159	177,804	-2,074	-5,833	0.176	0.225
67	91	175,730	-2,908	-6,418	0.249	0.269
68	57	172,822	-4,190	-6,989	0.334	0.439
69	33	168,632	-4,043	-7,138	0.252	0.455

Notes: Each row reflects the incentives workers face for continued work that year (e.g., the Age-55 row is the incentive to delay retirement until age 56). SSW is the net present discounted value of Social Security wealth at the beginning of the year. Benefit accrual is the change in SSW that results from working that year. After-tax accrual is the benefit accrual net of Social Security taxes paid during the year. Tax rate is the negative of the after-tax accrual divided by annual earnings. Diamond-Gruber tax rate replicates results from table 1 in Diamond and Gruber (1999).

the incentives facing a worker whose last year of work is labeled in the first column; that is, the “Age 55” row represents the incentives facing a worker who decides to retire on his fifty-sixth birthday.¹⁶ We show for each age the Social Security wealth, the accrual, the after-tax accrual, and tax/subsidy rate. In the final column, we show the tax/subsidy rate from Diamond and Gruber (1999), for comparison; their results are for a married male, which is appropriate since 90 percent of our sample is married.

We find that the median Social Security wealth for workers who retire on their fifty-fifth birthdays is \$154,928. Social Security wealth grows steadily through age sixty-five, then declines. This is shown most clearly

16. The SSW value is calculated from the data for age fifty-five, and is then constrained to follow the pattern of accruals from age fifty-six onward. We do this because the actual median SSW at each age does not correspond to the accrual pattern. If we use the SSW of the person with the median accrual, the pattern of SSW is nonsensical (with large shifts from year to year), since that person is different at every year. If we use the median SSW across the sample in each year (picking the median SSW in our sample, and not the SSW of the median accrual person), the SSW rises substantially over all years, due to sample selection. Another alternative is to project retirement incentives up to age sixty-nine for the sample working at age fifty-five; doing so, we find that the median SSW follows the same pattern as accruals, and that the accrual and tax variables are very similar to what we report here.

in the next column, which presents the benefit accruals at each age. From ages fifty-five to sixty-one, these accruals are positive due to the dropout year provision; the median worker is increasing his Social Security wealth by replacing lower earnings years in his earnings average. These accruals then get much larger between ages sixty-two and sixty-four, due to the actuarial adjustment. That is, the fact that accruals are larger after age sixty-two suggests that the actuarial adjustment is more than fair for the median worker; the gain to delaying receipt outweighs the fact that benefits are received for fewer years. At age sixty-five and thereafter, however, there are negative accruals for working additional years because the delayed retirement credit is not sufficiently large to compensate workers fairly.

The next column amends the benefit accrual by incorporating the fact that the worker and his employer must pay payroll taxes for additional work. This reverses the signs on the accruals at ages fifty-five to sixty-one, which are now negative, as the small benefit of AIME recomputation is outweighed by paying 12.4 percent of wages in tax. However, at ages sixty-two to sixty-four, the larger benefit accruals approximately offset the taxes incurred through additional work, so that the after-tax accrual for the median person is near zero. The after-tax accrual then turns sharply negative from age sixty-five onward.

The next column converts these after-tax accruals into tax/subsidy rates by dividing by the gross wage. There are positive taxes on work from ages fifty-five to sixty-one, but these taxes are significantly lower than the statutory 12.4 percent payroll tax rate, due to the benefit of additional earnings through the dropout year provision. The tax rate is near zero for the median worker at ages sixty-two and sixty-three and is 2.7 percent at age sixty-four. From age sixty-five onward, the tax rate is positive and very large. By age sixty-eight, the tax rate exceeds 30 percent, it drops back down again at age sixty-nine, but the samples at these ages in the HRS are very small.

These results are very similar to those in Diamond and Gruber, in spite of several important differences in methodology. First, Diamond and Gruber use a smooth age-earnings profile, which underestimates the value of the dropout year provision for people with real earnings trajectories with more variance. Second, Diamond and Gruber take an individual aged fifty-five and simulate his incentives to work at each future age, while the current calculations potentially incorporate some selection effects by using only those individuals still working at each age. The most notable differences between the two sets of results are at age fifty-five, where Diamond and Gruber find a subsidy to work (by construction, their individual replaces a zero year of earnings with his fifty-fifth year of work), and at age sixty-two, where Diamond and Gruber find a subsidy of 2.8 percent and we find a zero tax rate. The bottom line is very similar, however: Small

taxes on work up through age sixty-one, tax rates near zero at ages sixty-two to sixty-four, and more sizeable taxes after age sixty-five. Thus, we reaffirm the important conclusion of previous studies that the Social Security system does not place a significant tax on work at age sixty-two for the median worker.

10.3.2 Heterogeneity

As emphasized earlier, the incentives facing the median worker may mask considerable heterogeneity across the population in retirement incentives. Substantial heterogeneity may in turn be associated with an increase in retirement rates at age sixty-two, even if the incentives are small for the median worker. If, for example, there are large tax rates on work for 50 percent of the population, then there may be a zero tax rate for the median worker, but still potentially a large amount of retirement at age sixty-two.

We explore the heterogeneity in incentives in table 10.3, which shows the distribution of after-tax accruals and of tax/subsidy rates by age. As is immediately apparent, there is a substantial amount of heterogeneity in the accruals and tax rates. For example, from age fifty-five to sixty-one, while the median tax rate is positive and nontrivial, roughly one-sixth of the sample actually has a subsidy to additional work. At age sixty-two, while there is a zero tax rate for the median worker, 10 percent of the sample faces a tax rate of 6.8 percent or higher, and the standard deviation of the tax rate is 17.8 percent. After age sixty-five, while virtually all of the sample faces positive tax rates, there remains substantial variation in the magnitude of the tax rate; at age sixty-five, the standard deviation is nearly twice as large as the median tax rate.

What explains this substantial heterogeneity in Social Security incentives? This is an important question both for understanding how Social Security incentives work and for considering the validity of empirical work which relies on Social Security incentives to identify retirement behavior. As highlighted by Krueger and Pischke's (1992) criticism of the previous literature, if the vast majority of the variation in these incentives comes from factors such as wages or marital status, which are themselves likely to be independently correlated with retirement decisions, we might worry that incentive measures are capturing these other aspects of retirement decisions. If, however, as suggested by Hurd's (1990) rebuttal to this line of criticism, there are significant nonlinearities and interactions otherwise (likely) uncorrelated with retirement that primarily identify the impact of these incentive measures, one might feel more confident about retirement estimates.

We next turn to regression modeling of Social Security accruals and tax/subsidy rates to address this question. We consider in turn various potential determinants of the variation in incentives:

Table 10.3 Heterogeneity in Accrual and Tax Rate

Age	After-Tax Accrual			Tax/Subsidy Rate			Percent with Tax < 0
	10th Percentile	90th Percentile	Standard Deviation	10th Percentile	90th Percentile	Standard Deviation	
55	-3,977	92	2,553	-0.003	0.090	0.059	0.175
56	-4,100	154	3,416	-0.006	0.093	0.073	0.176
57	-4,205	178	3,492	-0.008	0.098	0.081	0.152
58	-4,342	212	2,778	-0.010	0.103	0.069	0.144
59	-4,476	212	3,458	-0.012	0.109	0.076	0.135
60	-4,478	272	4,005	-0.014	0.116	0.073	0.140
61	-4,365	340	4,971	-0.017	0.113	0.076	0.150
62	-1,592	2,291	2,945	-0.139	0.068	0.178	0.493
63	-1,657	2,490	2,888	-0.156	0.089	0.172	0.515
64	-2,497	1,628	1,985	-0.110	0.117	0.161	0.336
65	-7,462	-882	4,196	0.044	0.556	0.268	0.058
66	-8,365	-1,591	3,073	0.093	0.771	0.290	0.050
67	-8,627	-1,321	11,069	0.117	1.000	0.386	0.055
68	-9,391	-2,010	3,397	0.117	1.000	0.351	0.053
69	-9,878	-278	3,630	0.117	1.000	0.360	0.061

Notes: This table shows the 10th and 90th percentiles and the standard deviations of the distribution of after-tax accruals and tax rates, as well as (in the last column) the percent of the sample with negative tax rates. For a description of the incentive variables, see notes to table 10.2.

- *Age*. As shown earlier, there is important variation in tax rates with age.
- *Earnings in the last year of work*. This is the denominator of the tax rate, and will also enter through the dropout year provision. This may, as a result, have both linear and nonlinear effects, so we try both a linear earnings term and an earnings quartic.
- *AIME*. Average lifetime earnings is the primary determinant of benefits. Once again, the effects will be nonlinear, through the redistributive function that determines the PIA.
- *Marital status and age difference with spouse*. Marital status will be an important determinant of tax rates through the dependent benefits structure. In addition, the larger the positive age difference between spouses (a larger number of years by which the husband is older), the larger the value of the dependent spouse and survivor benefits.
- *Earnings in lowest year*. In combination with earnings in the last year, earnings in the lowest year will determine the value of the dropout year provision. We also include the number of years in the thirty-five-year earnings history with earnings below current earnings.

The results of this exercise are shown in table 10.4. We find that the explanatory power on the accrual is much more substantial than on the tax/subsidy rate, which is not surprising since the tax/subsidy rate introduces additional variation simply by normalizing by the wage. Thus, we focus on the accrual in our discussion.

Our overall conclusion is that, although these factors have some ability

Table 10.4 Variance Decomposition, Accrual, and Tax/Subsidy Rate

Variable	After-Tax Accrual		Tax/Subsidy Rate	
	R^2 of Variable	Cumulative R^2	R^2 of Variable	Cumulative R^2
Age dummies	0.062	—	0.044	—
Earnings	0.087	0.155	0.000	0.045
Earnings quartic	0.092	0.158	0.000	0.045
AIME	0.156	0.230	0.012	0.059
AIME quartic	0.171	0.238	0.015	0.064
Earn * AIME quartic	0.198	0.266	0.039	0.081
Married, agediff	0.002	0.270	0.001	0.082
Spouse earn * AIME ⁴	0.169	0.280	0.033	0.110
Low earn year	0.074	0.284	0.010	0.110

Notes: The second and fourth columns of the table show the R^2 from regressions of the after-tax accrual or tax/subsidy rate on the variable in the first column; the third and fifth columns show the cumulative R^2 from including that variable and all previous variables in the regression. Low earn year includes earning in lowest year and number of years with earnings below current years. For a description of the incentive variables, see notes to table 10.2.

to explain accrual patterns, the overall explanatory power is small. Factors clearly (potentially) correlated with tastes for retirement such as age, current earnings, and lifetime earnings, even when the former is entered as a series of dummies and the latter as flexible cubic functions, explains less than 25 percent of the variation in accruals. Even if we include a full set of interactions of these cubic functions of earnings and AIME, we explain only 27 percent of the variation. Adding marital status, age difference with spouse, spouse's earnings, and the low earnings year explains only another 2 percent of variation. Thus there appears to be a substantial amount of variation in the accrual that is not explained by factors that would plausibly otherwise be correlated with retirement.

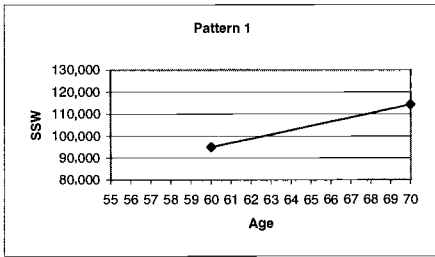
10.4 Peak Value Calculations

10.4.1 Motivation

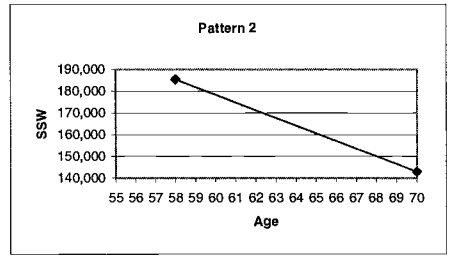
The results thus far have focused on one-year accruals of Social Security wealth and the associated tax/subsidy rates on an additional year of work. As noted above, a key insight of Stock and Wise (1988) in the private pension context is that one-year forward measures of this type may be misleading if there are substantial incentives or disincentives for retirement in future years. This was a natural concern in the context of private pensions, which often have dramatic and explicit retirement incentives at certain ages, such as the plan's early and normal retirement ages. However, is this an important issue in the context of Social Security?

In fact, the critical importance of considering the entire future path of incentives is illustrated in figure 10.3. This figure shows the most common patterns of after-tax Social Security wealth evolution (including payroll taxes paid for additional work) across our sample. In each figure, we graph for a group of workers the pattern of Social Security wealth evolution over all future years; this is done for the full cross-sectional sample, comparable to table 10.2, in which each worker contributes an observation for up to fourteen years. Each observation is then the pattern of Social Security wealth from that year forward, based on that year's characteristics. Under each graph is a figure for the percentage of our full cross-sectional sample that is in each case, and the cumulative share across the cases. For example, as shown in Pattern 1, 1 percent of the sample has an Social Security wealth that is everywhere increasing, while Pattern 3 shows that 14 percent of the sample has an Social Security wealth that first rises, then falls. In each case, the length of each segment is defined by the median starting and ending ages of the segment, and the slope of the segment is determined by the median Social Security wealth at the beginning and end of the segment.


As these graphs illustrate, substantial nonmonotonicities of the type




Percent=1.0, Cumulative Percent=1.0




Percent=47.7, Cumulative Percent=48.7




Percent=13.6, Cumulative Percent=62.3




Percent=1.8, Cumulative Percent=64.1




Percent=0.2, Cumulative Percent=64.3




Percent=29.0, Cumulative Percent=93.3



Percent=4.1, Cumulative Percent=97.4



Percent=0.3, Cumulative Percent=97.8



Percent=2.0, Cumulative Percent=99.7

Fig. 10.3 After-tax Social Security Wealth Patterns

seen for private pensions also exist for Social Security. For 38 percent of our sample, there is a local maximum that is not a global maximum. The most common data pattern (Pattern 2), which applies to 48 percent of the sample, is one in which after-tax Social Security wealth is always declining. However, the second most common pattern (Pattern 6), which applies to 29 percent of the sample, is one in which after-tax Social Security wealth declines from fifty-five to sixty-one, rises from sixty-two to sixty-four, then falls.

This is a striking finding, because it highlights an important weakness of the accrual measure. For any given year from age fifty-five to sixty-one, a typical worker will lose money on net through the Social Security system by working. However, by working, that worker is also *buying an option* on the more than fair actuarial adjustment that exists from age sixty-two to sixty-four. Incorporating this option, as shown in Pattern 6, leads to the conclusion that there may overall be net subsidies to work before age sixty-two for many workers through the Social Security system.

10.4.2 Peak Value

To incorporate this feature into our incentive calculations, we move away from the accrual and tax/subsidy rates to a more forward looking measure of incentives, which we call “peak value.” This is the value of continuing to work until the future year when Social Security wealth is maximized, or the difference between the expected present discounted value (PDV) of Social Security wealth at its highest possible value in the future and the expected PDV of Social Security wealth if one retires this year. This is thus like the typical accrual concept, except that the individual looks forward to the optimal year, rather than only to next year. If the individual is at an age beyond the Social Security wealth optimum, then the peak value is the difference between retirement this year and next year, which is exactly the accrual rate. Once again, it is natural to think about this type of concept relative to potential earnings, but here what is relevant is the entire stream of earnings until the optimal Social Security wealth is reached. That is, if the optimum is \$5,000 higher than Social Security wealth today and is one year away, then this is a larger subsidy to continuing to work than if the optimum were higher by the same amount but is five years away. We therefore normalize this peak value by the expected PDV of wages over the period between this year and the year of maximal Social Security wealth. Thus, this concept captures the benefits of continuing to work toward the peak Social Security wealth year, relative to earnings over that period.

We show our peak value calculations in table 10.5. On a pretax basis, peak value is \$22,426 at age fifty-five and falls steadily, becoming negative at age sixty-five. For the median worker, post-tax peak value is negative at all ages except for ages sixty-two to sixty-three. However, 30–40 percent of workers have positive after-tax peak values at ages fifty-five to sixty-

Table 10.5 Peak Value, Medians by Age

Age	Obs	SSW	Peak Value, Pretax	Peak Value, After-Tax	Percent with After-Tax PV > 0
55	2,811	154,928	22,426	-820	0.307
56	2,746	157,205	20,477	-1,018	0.292
57	2,444	159,341	18,339	-1,213	0.275
58	2,131	161,299	16,395	-1,399	0.282
59	1,822	163,090	15,228	-1,675	0.288
60	1,547	164,777	13,500	-1,701	0.326
61	1,252	166,340	12,245	-1,694	0.380
62	1,010	167,983	10,812	192	0.525
63	688	171,838	7,652	170	0.538
64	443	175,857	3,280	-758	0.359
65	313	178,706	-864	-4,808	0.077
66	159	177,804	-1,984	-5,799	0.069
67	91	175,730	-2,908	-6,418	0.066
68	57	172,822	-4,190	-6,989	0.053
69	33	168,632	-4,043	-7,138	0.061

Notes: PV is peak value, which is the change in SSW (Social Security wealth) that results from working until the age at which SSW is maximized (if peak has passed, PV is the after-tax accrual). Peak Value, Pretax excludes Social Security payroll taxes; Peak Value, After-Tax is net of taxes.

one. For these workers, the option value of a more than fair actuarial adjustment after age sixty-two outweighs payroll tax payments before age sixty-two. As a result, Social Security is actually providing a subsidy to additional work throughout all ages from fifty-five to sixty-four. This subsidy is rather small relative to earnings; the median after-tax peak value for those with positive value is about \$3,000. After age sixty-five, there is a negative return to additional work for the vast majority of workers.

Thus, viewed from a year-to-year perspective, the Social Security system taxes work between ages fifty-five and sixty-one at a modest rate for more than 80 percent of workers; but, viewed from a more forward-looking perspective, there are actually modest subsidies at those ages for 30–40 percent of workers because workers are buying the option of delaying claiming at a more than actuarially fair rate. Of course, this conclusion is somewhat overstated, for two reasons. First, as noted above, exercising the option to delay claiming does not require additional work, but simply delayed claiming. Second, for an individual who was planning to retire and claim at sixty-two for other reasons, there is no option value from delayed claiming. Thus, whether peak value is the relevant concept for actual retirement decisions at this age is an empirical question, and one we plan to explore in further work. The fact remains, however, that this option exists and is not recognized by the accrual concept.

It is worth noting an apparent inconsistency between table 10.2 and

Table 10.6 Accrual and Peak Value for Sample Observation

Age	Benefit Accrual	After-Tax Accrual	Tax/Subsidy Rate	Peak Value, Pretax	Peak Value, After-Tax
55	2,120	-1,388	0.046	19,170	-1,388
56	2,093	-1,218	0.041	17,050	-1,218
57	2,029	-1,159	0.039	14,957	-1,159
58	1,730	-1,360	0.046	12,928	-1,360
59	1,018	-1,996	0.066	11,198	-1,996
60	674	-2,287	0.074	10,180	-2,287
61	781	-2,077	0.066	9,505	-2,077
62	2,863	110	-0.003	8,725	274
63	2,817	164	-0.005	5,862	164
64	1,747	-804	0.025	3,045	-804
65	999	-1,451	0.045	1,297	-1,451
66	298	-2,052	0.063	298	-2,052
67	-368	-2,621	0.080	-368	-2,621
68	-960	-3,116	0.094	-960	-3,116
69	-1,480	-3,540	0.112	-1,480	-3,540

Notes: Table shows the incentives for one sample observation. For a description of the incentive variables, see notes to tables 10.2 and 10.5.

table 10.5. As should be obvious, the peak value at any given age is just the sum of all future accruals to the year when Social Security wealth is maximized. Yet the sum of the benefit accruals from age fifty-five forward in table 10.2 is not equal to the age-fifty-five pretax peak value from table 10.5. The reason for this apparent inconsistency is simply composition effects; the median individuals at each age and across the two tables are different. There is no clear way to address this in the aggregate while still representing the median values for our incentive variables of interest.

Instead, we illustrate that this is not a problem at the individual level in table 10.6, where we show these concepts for a typical individual in our data, who was fifty-five years old in 1992 and who has roughly the median earnings of his age cohort in our data. For this individual, the sum of the benefit accruals from any given age forward to the peak of Social Security wealth (age sixty-six) does equal the pretax peak value at that age. That is, at age fifty-five, the peak value pretax is \$19,170, which is the sum of the pre-tax accruals from age fifty-five through age sixty-six. Thus, for a typical individual in our data, we see that there is no inconsistency across these concepts; it arises only when we try to compare sample medians across the concepts.

10.4.3 Heterogeneity

As with after-tax accruals, there is a substantial amount of heterogeneity in after-tax peak values, as illustrated in table 10.7. At age sixty-two, for example, the median after-tax peak value is \$192. However, at that age,

Table 10.7 Heterogeneity in Peak Value

Age	Peak Value, After-Tax			Percent with PV > 0
	10th Percentile	90th Percentile	Standard Deviation	
55	-3,936	3,441	8,027	0.307
56	-4,081	4,095	8,384	0.292
57	-4,171	4,808	8,019	0.275
58	-4,310	5,221	8,260	0.282
59	-4,443	5,587	8,803	0.288
60	-4,441	6,589	9,435	0.326
61	-4,338	6,690	9,400	0.380
62	-1,515	6,299	7,852	0.525
63	-1,484	4,628	7,527	0.538
64	-2,428	2,023	5,684	0.359
65	-7,462	-215	7,943	0.077
66	-8,365	-994	10,680	0.069
67	-8,627	-1,200	12,274	0.066
68	-9,391	-2,010	3,849	0.053
69	-9,878	-278	3,630	0.061

Notes: PV is peak value. For a description of the incentive variables, see notes to table 10.5.

48 percent of the sample has peak values that are less than zero, the 90th percentile value is \$6,299, and the standard deviation is \$7,852, nearly thirty times the median.

The variation in peak values is more readily explained by the other factors that might naturally be included in a retirement model, as is shown in table 10.8. For the after-tax peak value, the inclusion of flexible functions of age, earnings, and AIME can explain about half of the variation; adding marital status, spousal characteristics, and lowest earnings year can explain another 3 percent. Still, a substantial share of the variation in peak value remains unexplained, suggesting that there is useful identifying variation available for inclusion in a retirement model.

10.4.4 Incorporating Private Pensions

We can also incorporate private pension incentives into our analysis. The HRS collected detailed pension-determination information from employers for roughly 60 percent of the individuals with pensions in the sample.¹⁷ They then used this information to create a pension benefits calculator that is comparable to the PIA simulation model we developed for Social Security. We use these calculated pension benefits at each retire-

17. Conversations with HRS staff indicate the HRS did not attempt to collect pension information for people in firms with fewer than 100 employees and that the nonresponse rate among employers they did contact was about 30 percent.

Table 10.8 Variance Decomposition, Peak Value (full sample)

Variable	Peak Value, After-Tax	
	R^2 of Variable	Cumulative R^2
Age dummies	0.014	—
Earnings	0.024	0.040
Earnings quartic	0.025	0.041
AIME	0.254	0.318
AIME quartic	0.297	0.350
Earn * AIME quartic	0.473	0.491
Married, agediff	0.006	0.497
Spouse earn * AIME ⁴	0.162	0.518
Low earn year	0.100	0.521

Notes: See notes to table 10.4 for description of table layout. Low earn year includes earnings in lowest year and number of years with earnings below current earnings. For a description of the incentive variables, see notes to table 10.5.

Table 10.9 Peak Value Including Pensions, Medians by Age

Age	Obs	SSW	Accrual After-Tax	Tax Rate	Peak Value, Pretax	Peak Value, After-Tax
55	2,811	183,138	-342	0.019	28,417	326
56	2,746	185,767	-381	0.018	25,843	197
57	2,444	188,380	-540	0.032	23,120	-61
58	2,131	190,675	-649	0.040	20,524	-18
59	1,822	192,781	-857	0.047	18,757	-90
60	1,547	194,637	-956	0.051	16,518	329
61	1,252	196,384	-1,003	0.053	14,341	1,143
62	1,010	198,149	672	-0.030	11,431	1,857
63	688	202,891	662	-0.025	7,949	863
64	443	208,312	-381	0.014	3,382	-172
65	313	212,733	-4,237	0.145	-774	-4,179
66	159	211,980	-5,384	0.189	-1,984	-5,369
67	91	210,024	-6,201	0.246	-3,538	-6,137
68	57	207,164	-6,916	0.344	-4,330	-6,916
69	33	203,236	-7,138	0.278	-4,452	-7,138

Notes: For a description of the incentive variables, see notes to tables 10.2 and 10.5. These results differ through the inclusion of pension incentives as well as Social Security program incentives.

ment age to create an analogous set of retirement incentive variables that include pensions.

The results of doing this for both of our incentive concepts (accrual and peak value) are presented in table 10.9. The patterns of incentives by age for the median worker are very similar to those shown in tables 10.2 and 10.5. This should not be surprising because only 40 percent of the individuals in our sample have pensions. Annual accruals are roughly \$700 larger

at ages fifty-five to sixty-four when pensions are included. The median tax/subsidy rate with pensions is roughly 2–3 percentage points lower at ages fifty-five to sixty-four than the median tax rate without pensions, and is similar at ages sixty-five to sixty-nine. With the inclusion of pension, the after-tax peak value is now often positive for the median person at ages fifty-five to sixty-four.

Table 10.10 shows the impact of including pensions on the distribution of the incentive measures. There is a substantial increase in the variation for all measures, particularly at ages fifty-five to sixty-one. For example, at age sixty, the ratio of the standard deviation to the median is close to 1 for the tax rate in the no-pensions case, but is higher than 3 in the with-pensions case. For the peak value measures, there are similarly large increases in heterogeneity, particularly at younger ages. Thus, adding pensions to the analysis does not dramatically change the incentives at the median, but does add substantial variation to the distribution of incentives.

10.5 Conclusions

The substantial time series decline in older male labor-force participation, as well as the striking correlation between the labor force departure rates of older workers and the early and normal retirement ages for Social Security, has motivated an enormous body of literature on how the Social Security program affects retirement. Yet there has been little recognition of a fundamental mystery in the relationship between Social Security incentives and retirement behavior: There is no evidence of a substantial disincentive to continued work at age sixty-two, despite the enormous increase in labor force exit at that age. This point was highlighted by Diamond and Gruber (1998), but this was based on a typical (simulated) individual.

In this paper, we have expanded on the earlier analysis in four ways. First, we have considered the impact of Social Security retirement incentives in real data, the HRS. We confirm in these data that there is, in fact, no tax on work at ages sixty-two to sixty-four at the median, further heightening the disconnect between observed retirement patterns and the pattern of Social Security retirement incentives. Second, however, we have shown that there is a substantial amount of heterogeneity in these incentives across our sample, and that (for example) there is a net tax on work at age sixty-two for about one-half of our sample. This would be more consistent with a spike in the hazard rate at that age, if it is those individuals being taxed who are responding by retiring. We also show that factors that otherwise might be expected naturally to impact retirement decisions can explain only a small share of the variation in accruals, suggesting that these are fruitful regressors for explaining retirement decisions.

Table 10.10 Heterogeneity in Peak Value and Accrual, Including Pensions

Age	Accrual, After-Tax			Tax Rate			Peak Value, After-Tax			Percent with PV < 0
	10th Percentile	90th Percentile	Standard Deviation	10th Percentile	90th Percentile	Standard Deviation	10th Percentile	90th Percentile	Standard Deviation	
55	-3,700	6,213	12,980	-0.177	0.091	0.195	-3.511	73,702	63,294	0.454
56	-3,701	6,914	10,779	-0.191	0.093	0.181	-3,570	69,620	57,529	0.476
57	-3,954	6,506	11,202	-0.185	0.100	0.186	-3,739	58,252	50,959	0.505
58	-4,158	6,517	9,477	-0.182	0.109	0.194	-3,973	50,340	38,199	0.501
59	-4,436	6,054	11,703	-0.171	0.117	0.189	-4,251	41,994	34,032	0.504
60	-4,410	5,021	9,189	-0.147	0.117	0.177	-4,245	34,273	29,530	0.483
61	-4,373	5,339	9,833	-0.149	0.117	0.166	-4,210	30,574	27,320	0.446
62	-1,677	6,512	7,269	-0.283	0.081	0.216	-1,495	24,078	21,155	0.347
63	-2,085	5,669	7,410	-0.272	0.106	0.214	-1,815	17,647	19,821	0.362
64	-2,637	4,540	7,662	-0.188	0.123	0.211	-2,535	8,479	16,403	0.517
65	-7,762	657	9,306	-0.028	0.556	0.304	-7,702	1,949	14,390	0.853
66	-9,392	11	5,459	-0.003	0.771	0.308	-9,392	987	13,680	0.874
67	-9,039	358	11,588	-0.015	1.000	0.403	-9,039	1,590	14,935	0.857
68	-10,830	-485	11,266	0.067	1.000	0.396	-10,830	-485	11,267	0.947
69	-10,370	-278	4,407	0.109	1.000	0.358	-10,370	-278	4,407	0.939

Notes: PV is peak value. For a description of the incentive variables, see notes to tables 10.2 and 10.5. These results differ through the inclusion of pension incentives as well as Social Security program incentives.

Third, we have suggested that the focus on next-year measures such as accruals and tax/subsidy rates might be misleading, particularly at ages fifty-five to sixty-one, because they ignore the option value of reaching age sixty-two and taking advantage of a (for many workers) more than fair actuarial adjustment. Thus, we also have considered a peak-value concept that compares wealth accruals not between this year and the next, but instead between this year and the year in which Social Security wealth reaches its peak. We find that using peak values instead of accruals leads to very similar results at the median after age sixty-two, but to subsidies to work rather than to taxes at ages fifty-five to sixty-one for a large share of the sample. Finally, we incorporate private pensions into our analysis; we find that the addition of pensions increases the return to additional work modestly at the median and substantially increases heterogeneity in the measures.

Our findings have two important implications for future empirical work on Social Security and retirement. First, our results suggest that if researchers are careful to condition on the determinants of both retirement and Social Security incentive measures, there may be sufficient remaining variation to identify the impact of these measures on retirement decisions. Second, our results suggest that, even in a Social Security-only context, it is important to consider forward-looking measures of the type pioneered by Stock and Wise (1990a,b). In preliminary work on retirement decisions (Coile and Gruber 1999), we have found that these forward-looking measures are indeed an important determinant of retirement behavior, while there is a much weaker relationship between retirement and accruals.

Appendix

In this appendix, we provide the formula for the computation of Social Security wealth.

Notation

t = year of observation

R = year of retirement

T = last year either spouse could be alive (maximum age is 120)

$pr_{h,s|t}$ = probability husband is alive at time s conditional on being alive at time t

$pr_{w,s|t}$ = probability wife is alive at time s conditional on being alive at time t

d = real discount rate (.03 in base case)

$age62_{h,s}$ = indicator variable equal to 1 if husband is age sixty-two or over at time s

$\text{age62}_{w,s}$ = indicator variable equal to 1 if wife is age sixty-two or over at time s

$\text{age60}_{h,s}$ = indicator variable equal to 1 if husband is age sixty or over at time s

$\text{age60}_{w,s}$ = indicator variable equal to 1 if wife is age sixty or over at time s

$\text{rwb}_{h,s}$ = retired worker benefit of husband if husband retires at time s

$\text{rwb}_{w,62}$ = retired worker benefit of wife if wife retires at age sixty-two

$\text{dsb}_{h,62}$ = dependent spouse benefit of husband if wife retires at age sixty-two

$\text{dsb}_{w,s}$ = dependent spouse benefit of wife if husband retires at time s

$\text{svb}_{h,s}$ = survivor benefit of husband if wife dies at time s

$\text{svb}_{w,s}$ = survivor benefit of wife if husband dies at time s

s, k = simple counting variables

Formula

$$\begin{aligned} \text{SSW}_t(R) = & \sum_{s=t}^{R-1} (1 + d)^{-(s-t)} \left[\text{pr}_{h,s|t} * \text{pr}_{w,s|t} * \text{age62}_{w,s} * \text{rwb}_{w,62} \right. \\ & + (1 - \text{pr}_{h,s|t}) * \text{pr}_{w,s|t} * \text{age62}_{w,s} \\ & * \max \left(\text{rwb}_{w,62}, \sum_{k=t}^{s-1} \frac{\text{pr}_{h,k|t} - \text{pr}_{h,(k+1)|t}}{\text{pr}_{h,t|t} - \text{pr}_{h,s|t}} \text{svb}_{w,k} \right) \Big] \\ & + \sum_{s=R}^T (1 + d)^{-(s-t)} \left\langle \text{pr}_{h,s|t} * \text{pr}_{w,s|t} * \{ \text{age62}_{h,s} * [\text{rwb}_{h,R} + \text{age62}_{w,s} \right. \\ & * \max(0, \text{dsb}_{h,62} - \text{rwb}_{h,R})] + \text{age62}_{w,s} * [\text{rwb}_{w,62} + \text{age62}_{h,s} \\ & * \max(0, \text{dsb}_{w,R} - \text{rwb}_{w,62})] \} + \text{pr}_{h,s|t} * (1 - \text{pr}_{w,s|t}) * \text{age62}_{h,s} \\ & * \max \left(\text{rwb}_{h,R}, \sum_{k=t}^{s-1} \frac{\text{pr}_{w,k|t} - \text{pr}_{w,(k+1)|t}}{\text{pr}_{w,t|t} - \text{pr}_{w,s|t}} \text{svb}_{h,k} \right) + (1 - \text{pr}_{h,s|t}) * \text{pr}_{w,s|t} \\ & * \text{age62}_{w,s} * \max \left(\text{rwb}_{w,62}, \sum_{k=t}^{s-1} \frac{\text{pr}_{h,k|t} - \text{pr}_{h,(k+1)|t}}{\text{pr}_{h,t|t} - \text{pr}_{h,s|t}} \text{svb}_{w,k} \right) \Big\rangle \end{aligned}$$

An important assumption built into the calculation is that the spouse retires at age sixty-two.

The benefit variables (rwb, dsb, and svb) are adjusted appropriately for actuarial adjustment or delayed retirement credit. The adjustment depends on R , the birth year of each spouse (since Social Security rules differ by birth cohort), and age difference between the spouses. Where an individual first claims retired worker benefits and later tops them up to the level of the dependent spouse benefit, the appropriate actuarial adjustment is applied to each part of the total benefit.

Claiming is assumed to occur at first eligibility (the age of retirement or age sixty-two, whichever is later). For simplicity, survivor benefits are assumed to be claimed no earlier than age sixty-two, though individuals are allowed to claim them at age sixty, or earlier if there are dependent children.

The calculations including pensions are analogous, except that pension receipt commences as soon as the individual retires (not at age sixty-two).

References

- Blau, David M. 1994. Labor force dynamics of older men. *Econometrica* 62 (1): 117–56.
- Blinder, Alan S., Roger H. Gordon, and Donald E. Wise. 1981. Rhetoric and reality in Social Security—A rejoinder. *National Tax Journal* 34 (4): 473–78.
- . 1980. Reconsidering the work disincentive effects of Social Security. *National Tax Journal* 33 (4): 431–42.
- Boskin, Michael J., and Michael D. Hurd. 1978. The effect of Social Security on early retirement. *Journal of Public Economics* 10 (3): 361–77.
- Burkhauser, Richard V., and John Turner. 1981. Can twenty-five million Americans be wrong?—A response to Blinder, Gordon, and Wise. *National Tax Journal* 34 (4): 467–72.
- Burtless, Gary. 1986. Social Security, unanticipated benefit increases, and the timing of retirement. *Review of Economic Studies* 53 (5): 781–805.
- Burtless, Gary, and Robert Moffitt. 1984. The effect of Social Security benefits on the labor supply of the aged. In *Retirement and economic behavior*, ed. Henry J. Aaron and Gary T. Burtless, 135–75. Washington, D.C.: Brookings Institution.
- Coile, Courtney, Peter Diamond, Jonathan Gruber, and Alain Jouten. (2001). Delays in claiming Social Security benefits. *Journal of Public Economics*, forthcoming.
- Coile, Courtney, and Jonathan Gruber. 1999. Social Security and retirement. NBER Working Paper no. 7830. Cambridge, Mass.: National Bureau of Economic Research.
- Diamond, Peter, and Jonathan Gruber. 1999. Social Security and retirement in the United States. In *Social Security and retirement around the world*, ed. Jonathan Gruber and David A. Wise, 437–73. Chicago: University of Chicago Press.
- Diamond, Peter, and Jerry Hausman. 1984. Retirement and unemployment behavior of older men. In *Retirement and economic behavior*, ed. Henry J. Aaron and Gary T. Burtless, 97–135. Washington, D.C.: Brookings Institution.
- Diamond, Peter, and Botond Koszegi. (1999). Quasi-hyperbolic discounting and retirement. MIT, Mimeograph.
- Feldstein, Martin A., and Andrew A. Samwick. 1992. Social Security rules and marginal tax rates. *National Tax Journal* 45 (1): 1–22.
- Fields, Gary S., and Olivia S. Mitchell. 1984. Economic determinants of the optimal retirement age: An empirical investigation. *Journal of Human Resources* 19 (2): 245–62.
- Friedberg, Leora. (2000). The labor supply effects of the Social Security earnings test. *Review of Economics and Statistics* 82 (1): 48–63.
- Haider, Steven, and Gary Solon. (1999). Nonrandom selection in the HRS Social

- Security earnings sample. Unpublished paper, RAND and University of Michigan, Department of Economics.
- Hausman, Jerry A., and David A. Wise. 1985. Social Security, health status, and retirement. In *Pensions, labor, and individual choice*, 159–91. ed. David A. Wise, Chicago: University of Chicago Press.
- Hurd, Michael D. 1990. Research on the elderly: Economic status, retirement, and consumption and saving. *Journal of Economic Literature* 28 (2): 565–637.
- Hurd, Michael D., and Michael J. Boskin. 1984. The effect of Social Security on retirement in the early 1970s. *Quarterly Journal of Economics* 99 (4): 767–90.
- Kahn, James A. (1988). Social Security, liquidity, and early retirement. *Journal of Public Economics* 35 (1): 97–117.
- Krueger, Alan B., and Jorn-Steffan Pischke. 1992. The effect of Social Security on labor supply: A cohort analysis of the notch generation. *Journal of Labor Economics* 10 (4): 412–37.
- Lumsdaine, Robin, and David Wise. 1994. Aging and labor force participation: A review of trends and explanations. In *Aging in the United States and Japan: Economic Trends*, ed. Y. Noguchi and D. Wise, 7–41. Chicago: University of Chicago Press.
- Meyer, Bruce D. 1989. A quasi-experimental approach to the effects of unemployment insurance. NBER Working Paper no. 3159. Cambridge, Mass.: National Bureau of Economic Research.
- Quinn, Joseph. 1999. Retirement patterns and bridge jobs in the 1990s. Policy Brief for the Employee Benefit Research Institute. Washington, D.C.: EBRI.
- Ruhm, Christopher. 1995. Secular changes in the work and retirement patterns of older men. *Journal of Human Resources* 30 (2): 362–85.
- Samwick, Andrew A. 1998. New evidence on pensions, Social Security, and the timing of retirement. *Journal of Public Economics* 70 (2): 207–36.
- Stock, James H., and David A. Wise. 1990a. Pensions, the option value of work, and retirement. *Econometrica* 58 (5): 1151–80.
- . 1990b. The pension inducement to retire: An option value analysis. In *Issues in the economics of aging*, ed. David A. Wise, 205–29. Chicago: University of Chicago Press.

Comment Andrew A. Samwick

The Social Security program in the United States grew substantially in coverage and generosity during the postwar period. Motivated by the large, contemporaneous reductions in the labor force participation of older workers, a substantial literature emerged during the 1970s and 1980s on the effects of Social Security on the timing of retirement. Surveys of the literature can be found in Atkinson (1987) and Quinn, Burkhauser, and Myers (1990). Most studies estimate statistically significant relationships between the level of Social Security benefits and the likelihood of retirement at various ages. However, these estimated relationships typically im-

ply a small economic impact of altering Social Security benefits on the average age of retirement or probability of retirement.

This finding is rather surprising, and I believe that it is a consequence of examining retirement primarily in a cross-section. I doubt that it has convinced too many people that the differences in Social Security replacement rates between cohorts who turned sixty-five in 1950 and those who turned sixty-five in 1975 have nothing to do with the differences in the cohorts' respective retirement ages. The finding does tell us that, within a given cohort, workers who are observed to have different entitlements under Social Security are not observed to retire at dramatically different ages. Even without time-series implications, this is an interesting result that is worthy of continued study.

There are several possible explanations that can be explored as a response to this finding. The first is that investigators have made an inappropriate simplifying assumption regarding an important aspect of the retirement decision. For example, Gustman and Steinmeier (2000) and Coile (1999) have shown that retirement decisions of spouses should be modeled jointly, rather than independently, as is done in most of the literature. Another simplification is the separation of the household's intertemporal consumption problem from its retirement decisions. Although option value models have made important advances toward incorporating intertemporal tradeoffs in the labor market, they typically assume uniformity in the way households make those choices. Recent work on consumption has shown that household consumption choices are characterized by higher and more variable discount rates than are typically assumed in retirement analyses. I will return to this issue at the end of my comments.

A second explanation is that the primary determinant of cross-sectional retirement patterns is some other phenomenon, such as the incentives from employer-provided pensions. This is the motivation behind the original Stock and Wise (1990a, b) option value model and the many subsequent papers based on it. As documented in Kotlikoff and Wise (1987), defined benefit (DB) pension plans typically provide large financial incentives to delay or hasten retirement at various ages. These incentives vary substantially across plans. Further, DB plans also grew in coverage and generosity during the postwar period and, as shown in Luzadis and Mitchell (1991), the ages at which retirement incentives are strongest has been falling over time. The key ages for pensions are currently younger than the Social Security early retirement age.

A third explanation is the one pursued by the authors in this chapter. The thrust of their argument is that Social Security itself provides a variety of different incentives to workers with different earnings histories. In order to calculate these incentives, rich data on earnings histories are required. The failure to find an effect of Social Security on retirement may in fact be due to measurement problems. It would be tempting for a reader of the

option value literature to date to conclude that Social Security incentives are not an important source of cross-sectional variation in retirement outcomes in the United States. The authors have done the literature on the incentive effects of retirement programs a great service in carefully measuring Social Security incentives. The main lesson the reader takes away from this chapter is that rumors of Social Security's irrelevance may have been greatly exaggerated.

Beginning with the seminal work of Stock and Wise (1990a, b), studies of the option value of retirement have utilized strong data on pensions and weak data on Social Security. These studies have also found that in comparison to incentives from pensions, Social Security incentives have little explanatory power for retirement decisions. Samwick (1998b) acknowledges the link between the poor quality of his Social Security data and the imprecision of the estimates of retirement incentive effects based on them. Note that Social Security calculations place a greater demand on the earnings data. Pension incentives depend only on the earnings history at the current employer and a forecast of future earnings with that same employer.¹ In addition to these earnings, Social Security incentives depend on earnings in years prior to those spent with the current employer, earnings at other jobs worked while with the current employer, and a forecast of earnings that may be received after the employee leaves the current employer but before claiming benefits.

The present chapter provides a basis for quantifying how important Social Security retirement incentives are when measured with as much attention to detail as the pension incentives in past option value studies. It is a tantalizing prelude to econometric analyses of retirement for male workers in Coile and Gruber (1999) and working couples in Coile (1999). My comments will therefore include some issues that are relevant to the authors' broader research agenda. In the present chapter, the authors make two arguments. The first is that careful calculation and rich data yield substantial heterogeneity in financial incentives to retire from Social Security. The second is that the authors' measure of "peak value" is a comprehensive and robust measure of incentives to hasten or delay retirement. My main remarks will assess the validity of these points and make suggestions for further research on the effect of financial incentives on retirement.

Heterogeneity in Social Security Retirement Incentives

The Health and Retirement Study (HRS) is ideally suited to examining the financial incentives inherent in the Social Security system. The initial HRS cross-section in 1992 is a nationally representative sample of households between the ages of fifty-one and sixty-one. Subsequent interview

1. However, as discussed below, pension calculations are far more susceptible to measurement error in the earnings histories.

waves are conducted every two years, and the authors have matched the household data to Social Security earnings histories. This allows them to calculate Social Security wealth at various ages and thereby estimate the financial incentives to retire at each age.

There are two characteristics of the data that must be demonstrated in order to suggest that they will be useful for further analyses of retirement behavior. The first is the heterogeneity in incentives, controlling for age. The results are presented in tables 10.2 and 10.3 for the one-year accrual of wealth and tables 10.5 and 10.7 for peak value. Median accruals are positive before age sixty-two, large and positive from age sixty-two to age sixty-four, and negative and declining at ages beyond sixty-five. Tables 10.3 and 10.7 suggest that there is variation in the magnitude of accruals by age around the medians, although I will argue below that this variation is small compared to the variation in incentives caused by pensions. Tables 10.4 and 10.8 give an indication that much of this variation in accruals is not simply the result of differences in the factors (e.g., age, earnings, and marital status) that help determine benefits but might also affect retirement directly.

The second characteristic is the nonmonotonic pattern of accruals as workers approach retirement. These patterns are shown in figure 10.3 for the authors' measure of After-Tax Social Security wealth. The results here are somewhat disappointing, in that nearly half, or 47.7 percent, have monotonically decreasing Social Security wealth as given by Pattern 2. Pattern 6 is the next most frequent pattern, with 29.0 percent of the sample. This pattern differs from Pattern 2 in that the negative accruals are interrupted by an interval of positive accruals. These positive accruals are fairly small and do not last very long, suggesting at least from the figure that they are not too important in dollar terms. Taken together, these two patterns—with little if any departure from a simple profile of declining Social Security wealth—account for more than three-fourths of the observations. As the authors note, in the absence of nonmonotonicities, it becomes more difficult to econometrically identify the effects of Social Security from other factors that may be increasing or decreasing with age.

It may be that the use of medians and piecewise-linear segments in figure 10.3 masks some important differences within the patterns. For example, there may be an important subset of workers who have very large positive accruals between ages sixty-two and sixty-five, despite the low median. It could also be that some of the workers in Pattern 2 look a lot like those in Patterns 3 and 7, and that some look like Patterns 4 and 10 (with flat segments instead of increases). The lack of variation indicated in this descriptive work and the "success" of the initial econometric estimations in Coile and Gruber (1999) and Coile (1999) present a bit of a puzzle. More explanation is warranted if the latter estimates are to be

thought of as identified by Social Security wealth rather than as other factors that change monotonically for a given worker over time.

As a final note about the Social Security wealth calculations, it is worth emphasizing that the idea for calculating Social Security retirement incentives is not original to this chapter. The real novelty in this chapter is the data that are used, particularly the Social Security earnings histories. Because obtaining the earnings histories and working with them are not trivial processes, the authors could do a useful service in this chapter if they also calculated incentives using only the data on prior earnings available in the public use files of the HRS. Comparisons of the pattern of incentives with and without the earnings histories would help determine how useful the extra data are. What will users of the HRS miss if they do not use the detailed earnings histories?

Retirement Incentives from Pensions

The next interesting question to ask about the calculated Social Security incentives concerns their magnitude—are these numbers large or small? Since most of the recent research on the accrual of retirement wealth is based on option value models of pensions, a natural benchmark is the distribution of pension incentives. The authors make use of the Pension Provider Survey (PPS) that accompanied the first wave of the HRS. As implemented, the PPS includes detailed pension information on roughly 60 percent of the workers eligible for pensions.

The authors do not present analogous distributions of the incentives from pensions in the HRS.² However, they report distributions of peak value including pensions in tables 10.9 and 10.10. The effect of adding the pension data is substantial. Comparing median peak (pretax) values in tables 10.5 and 10.9, the inclusion of pensions adds about \$6,000 at age fifty-five and \$3,000 at age sixty. Tables 10.7 and 10.10 show that the most dramatic effects are (for this sample) at the higher retirement income levels. The standard deviation of peak after-tax value at age fifty-five increases from \$8,027 to \$63,294 when pensions are included, and the 90th percentile increases from \$3,441 to \$73,702. The inclusion of the pension data has a smaller effect at older ages, because most pensions have their early retirement ages (and therefore largest incentives) between the ages of fifty and sixty. Very few workers who have pensions are still working past age sixty-five.

To get an idea of the whole distribution of incentives, figure 10c.1 is a graphical representation of incentives from pensions, taken from Samwick (1998b). It shows the mean, 10th percentile, and 90th percentile of the one-year benefit accruals (see table 10.2, column 4) from the pensions in

2. See Gustman et al. (1999, 2000) for preliminary tabulations.

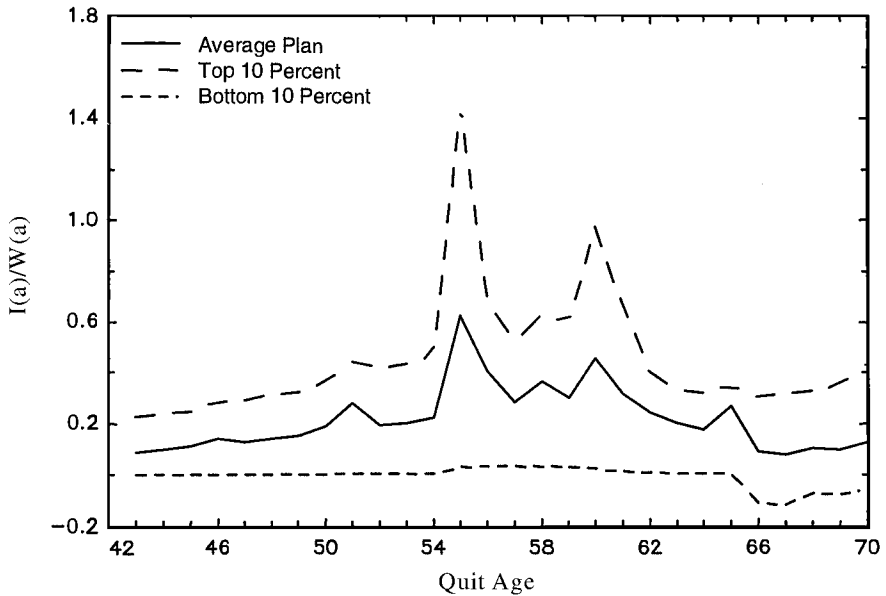


Fig. 10C.1 Pension accruals relative to wages, weighted average of all plans:

Hired at age 31, average wage

Source: Samwick (1998b, fig. 2)

the PPS that accompanied the Survey of Consumer Finances (SCF) 1983. This distribution is for a worker with average earnings who started working for his current employer at age thirty-one. That worker's benefits under all pension plans in the sample are calculated and appropriately weighted to reflect the distribution of workers in those plans. The vertical axis shows the accrual as a share of current earnings. The graph shows that accruals due to pensions are many times larger than those from Social Security. The 90th percentiles at ages fifty-five and sixty are roughly 1.4 and 1.0 times earnings in those years. Even the mean accruals are close to 60 percent and 40 percent of current earnings at those ages.³

The main conclusion to be drawn from the pension results is that incentives from pensions will dwarf incentives from Social Security for workers who are covered by both. This is the general result found in Stock and Wise's (1990a) sample of salesmen covered by the same pension. This result strongly suggests that researchers should examine workers with pen-

3. It is not uncommon for pension plans to provide multiple spikes in a worker's accrual of pension wealth. Workers who start working for their employers at older ages typically have fewer and later (but larger) spikes in their accrual profiles.

sions separately from those without pensions. It would also be interesting to see how much the inclusion of pensions changes the optimal retirement date in the peak value formula (discussed below).

Regarding the specific calculations with pensions in this chapter, I have two primary concerns. The first is that, when the distributions with pensions are presented, only 60 percent of the pension-covered population is represented. And considering the way that observations get lost in moving from the HRS to the PPS, it seems clear that the subsample of pensions that are found will be disproportionately those from the government sector and large corporations. These employers typically are easier to locate, have better record keeping, and account for multiple respondents to the survey. This was certainly the case in the SCF 1983, which had about 70 percent of the observations with pensions in its PPS, and the authors note (see n. 17) that the same is true here. Thus, the subsample of the pension-covered population that is included in this PPS is more homogeneous than the excluded subsample, which includes many fewer observations per plan and hence more plans. Even the very large effects of pensions on the distribution shown in tables 10.9 and 10.10 are likely underestimates of the true contributions of pensions to heterogeneity in retirement incentives.

One possible remedy would be to control for this sample selection directly, perhaps by modeling it as a function of firm size, occupation, and industry. This would essentially re-weight the sample to increase the representation of the workers who are in the HRS and PPS with characteristics, such as working at a small firm, that are associated with being omitted from the PPS. In the absence of such a remedy, the variance decompositions of the sort presented in table 10.8 are likely to miss a lot of the variation in retirement incentives due to pensions.

My second concern is that pension calculations are extremely sensitive to the quality of the inputs, particularly age, tenure, and the details of the pension formulas themselves. The potential for coding errors is large, and although the Survey Research Center staff and other researchers have worked diligently to make the formulas error free, pension calculations must be checked for anomalous results in each application. It is a tedious but necessary process. For example, suppose that, either because the respondent reported a year of hire that was one year later than the actual starting year or the pension formula was miscoded in the PPS to have early retirement eligibility at an age one year too late. Consider a worker who paid a lot of attention to financial incentives and therefore stayed with the firm until the large spike in the accrual profile at early retirement, and then retired. We will observe this worker's retiring just before a large financial incentive to delay retirement, because we have gotten the timing of that accrual wrong by only one year. Note as well that the effect on a maximum likelihood estimation will be asymmetric—retiring one year

after a large spike is not so different from retiring at the spike, but retiring one year before the large spike strongly suggests that financial incentives do not affect retirement.

The point is that fairly small errors in either the household or the pension survey can have potentially large effects on econometric analyses. Because the present chapter seeks only to present a distribution of incentives without estimating a behavioral effect, the problem is less severe. Some of the large accruals are probably off by a year or two in the calculations underlying tables 10.9 and 10.10; however, these sorts of errors may in large measure offset each other, unless tenure or age is systematically over- or underreported.

At this stage, I can only offer suggestions to minimize the exposure to these problems in subsequent work. An example is the way the authors have backcast pension formulas to the years prior to the initial survey (1992). Comparisons of pension formulas in the 1983 and 1989 SCF Pension Provider Surveys in Samwick (1993) and Gustman and Steinmeier (1998) show that pension formulas changed substantially during the 1980s. In particular, early retirement ages continued their downward trend, and the growing inequality in wages was also reflected in growing inequality in pension entitlements. Using the pension formula in 1992 to describe pension incentives during the 1980s is likely to expose the calculations to error. It is quite unlikely that the historical provisions of the plans are accurately represented in the summary plan description as of 1992. For example, do we see evidence of the temporary early retirement windows that were popular in the early 1980s and analyzed in Stock and Wise (1990b) in the current formulas?

Peak Value versus Option Value

The authors begin to make the case in this chapter, and follow it up in Coile and Gruber (1999), that their measure of peak value is a robust measure of retirement incentives. *Peak value* is defined as the maximum increment to the actuarial present value of future retirement benefits for any possible year of retirement. Like the option value on which it is based, it has the advantage of being forward-looking. It evaluates the financial gains to delaying retirement to the most advantageous year, not just the next year as in a typical retirement wealth accrual variable.⁴ Peak value is also easily scaled by the present value of future wages until that date.

In fact, peak value is equivalent to the option value under a set of pa-

4. Peak value also shares with the option value the potential drawback that it is not a full dynamic programming model—it is calculated conditional on a given R^* and considers only the value of retiring in the best year. Lumsdaine, Stock, and Wise (1992) compare the option value model to a dynamic programming model that compares the value of retirement in the current year to the value of reevaluating all retirement possibilities next year. The distinction is shown not to be important in their work.

parameter restrictions. To see this, consider the original derivation of the option value from Stock and Wise (1990a). It begins with the indirect utility function

$$(1) \quad V_t(R) = \sum_{s=t}^{R-1} \beta^{s-t} \pi(s|t) E_t(y_s^\gamma) + \sum_{s=R}^T \beta^{s-t} \pi(s|t) E_t[k \cdot B_s(R)^\gamma].$$

In this expression, R denotes the date of retirement and t denotes the current date. The probability of living to year s conditional on being alive in year t is $\pi(s|t)$. The discount factor is denoted by β , and risk aversion (or intertemporal substitution) is measured by γ . The final parameter is k , which represents the additional value of receiving income during retirement rather than the working year. It is designed to capture the utility value of leisure and is expected to be greater than unity. In evaluating retirement possibilities, a worker would trade one dollar of income while working for k dollars during retirement. Income takes the form of earnings, y_s , during the remaining years of work or retirement benefits, $B_s(R)$, during retirement. Note that benefits are a function of the year in which retirement occurs.

Using this indirect utility function, the worker chooses the optimal date of retirement, R^* , as the one that maximizes $V_t(R)$. The option value of continued work is the excess of the indirect utility of retiring at R^* rather than the current date t . The option value is $V_t(R^*) - V_t(t)$, or

$$(2) \quad \begin{aligned} OV_t(R^*) = & \sum_{s=t}^{R^*-1} \beta^{s-t} \pi(s|t) E_t(y_s^\gamma) \\ & + \sum_{s=R^*}^T \beta^{s-t} \pi(s|t) E_t[k \cdot B_s(R^*)] \\ & - \sum_{s=t}^T \beta^{s-t} \pi(s|t) E_t[k \cdot B_s(t)^\gamma]. \end{aligned}$$

The first two terms are $V_t(R^*)$, and the last term is $V_t(t)$. When retirement occurs at the current date, there are no terms reflecting the utility from additional years of labor income.

By comparison, the authors' peak value calculation is given by

$$(3) \quad \begin{aligned} PKV_t(R^*) = & \sum_{s=R^*}^T \beta^{s-t} \pi(s|t) E_t[B_s(R^*)] \\ & - \sum_{s=t}^T \beta^{s-t} \pi(s|t) E_t[B_s(t)]. \end{aligned}$$

Peak value imposes three restrictions on the option value. The first restriction is that the first term in equation (2), pertaining to future earnings, is dropped. This restriction should be innocuous. Variance decompositions of the sort done in table 10.8 would show that earnings explain a larger

fraction of option values than peak values, but this would not necessarily suggest that one measure is preferred to the other in a regression. The present value of earnings through the optimal retirement date can simply be included in the regression in addition to the option value or the peak value. For example, Samwick (1998b) finds that, controlling for the option value of retirement, earnings do not have a significant effect on retirement in most specifications. Additionally, this first term (with the subsequent two restrictions imposed) is the quantity that the authors use to scale peak value in the right-most columns of tables 10.5 through 10.8.

The second restriction is that k is equal to unity. In an option value calculation, this restriction implies that there is no disutility of working relative to being retired. In that context, such a restriction would be counterintuitive and is not supported by the estimates of about 1.6 from Stock and Wise (1990a, table 4). However, peak value compares income flows only during retirement, so this assumption is without loss of generality. Additionally, a value of k can be estimated in a simple regression as long as the first term from the option value calculation (the present value of future earnings) is included as a regressor along with peak value. The value of k would be the ratio of the coefficient on the peak value term to the coefficient on the earnings term.

The third restriction is that γ is equal to unity. Setting $\gamma = 1$ implies that workers are indifferent to whether income and retirement benefit payments vary across years. Values less than unity, such as the estimates of approximately 0.75 in Stock and Wise (1990a), are consistent with individuals who derive more utility from smooth income flows than from variable ones. In the option value calculation, the main variation in income across periods pertains to the retirement replacement rate. Since the peak value calculation pertains only to income received in retirement and both Social Security and DB pensions typically pay annuities, the choice of γ is much less critical in the peak value calculation. However, it would still be useful to experiment with different values of γ , applied to both peak value and the present value of earnings in a regression, in order to determine how sensitive empirical estimates are to the smoothness of income over time.

A more subtle difference between the two concepts is that the optimal retirement date, R^* , may change when the present value of future earnings is dropped from the option value to get peak value. Given the central role of the optimal retirement date, the authors should investigate the extent to which the two concepts yield different values of R^* . Doing so would require a more sophisticated procedure for forecasting future wages than simply assuming a 1 percent growth rate, as assumed here.

Parameter Selection and Heterogeneity

As discussed above, the main difference between the formulas for peak value and option value is whether the first term of the option value, repre-

senting the present value of future earnings until the optimal retirement date, is included in the calculation. Much of that difference can be eliminated in practice by simply including that first term as an additional explanatory variable in a reduced-form specification. Since future earnings are an appropriate variable in that regression anyway, the real issue with using peak value rather than the option value is whether the values of the parameters are estimated, as in Stock and Wise (1990a), or simply assumed, as in reduced-form approaches in Samwick (1998b) or Coile and Gruber (1999). This issue is relevant but perhaps not too important for the risk aversion and leisure parameters, γ and k .

The remaining parameter in both the option value and peak value calculations is the discount factor, β , in the formulas above. Mathematically, $\beta = 1/(1 + \delta)$, where δ is the discount rate or rate of time preference. The choice of discount rate is not a trivial matter in retirement analyses. Berkovec and Stern (1991) report that in their dynamic programming model, “the estimation algorithm would not converge” when they attempted to estimate the discount factor. They fixed it at 0.95. My own efforts to estimate option value models in Samwick (1998b) were similarly unsuccessful, as it was not possible to obtain reliable estimates for both the discount factor and the leisure parameter, k . As in this chapter, the value of β was set to correspond to a discount rate of 3 percent per year. Stock and Wise (1990a) estimate values for β of approximately 0.75 across several specifications.

Part of the explanation for Stock and Wise’s estimate of such a high discount rate (nearly 33 percent per year) is that it is discounting utility rather than dollars. Another part of the explanation, however, is that discount rates may in fact be very large. The best evidence for this comes from wealth levels—even among households on the verge of retirement, median financial wealth corresponds to less than a year’s worth of income. If discount rates were as low as 3 percent over the life cycle, then household wealth accumulation would be much larger than we observe for the median household (by perhaps a factor of 4 or 5, based on simulations in Samwick [1998a]). The median peak value of about \$25,000 in table 10.5 would not be very large in comparison.

Other evidence on the magnitude of discount rates comes from the cross-sectional variation in wealth. Rates of time-preference as low as 3 percent are inconsistent with the observed sensitivities of household wealth to income uncertainty (Carroll and Samwick 1997) and to the retirement replacement rate (Samwick 1995). In both cases, the observed sensitivities are of the correct sign but are too low to be consistent with patient consumers. If households discounted the future at only 3 percent, then they would react more strongly to differences in income uncertainty and differences in the retirement replacement rate that imply very different possibilities for future resources.




Fig. 10C.2 Distribution of discount rates, survey of consumer finances 1992
Source: Samwick (1998a, fig. 3)

Another characteristic of the distribution of wealth is its wide dispersion. Although many factors affect wealth accumulation, it seems hard to believe that we could generate the observed distribution of wealth if all households had the same 3 percent rate of time preference. Instead, we should infer from the distribution of wealth that there are also distributions of underlying preference parameters, such as the discount rate. In a first attempt to make such inference, Samwick (1998a) estimates using the 1992 SCF that the median discount rate (δ) is 7.63 percent, compared to an assumed interest rate of 4 percent. A household's estimated discount rate reflects the value that would be needed in a stochastic life-cycle model of consumption to generate a predicted wealth-to-income ratio that matches the household's observed wealth-to-income ratio.

More importantly, although factors such as the retirement replacement rate and income uncertainty are accounted for in the estimation, much of the heterogeneity in wealth holdings appears as heterogeneity in discount rates. The 25th and 75th percentiles of the distribution are 2.93 percent and 14.66 percent, respectively. The difference in behavior for households with these two discount rates is enormous. The full distribution is shown in figure 10C.2.⁵ Samwick (1998a) also compares the estimated discount rates across different responses to a question regarding the household's

5. Values higher than 20 percent and lower than -15 percent have been truncated at those levels.

most important financial planning horizon. Households who responded “the next few months” had average discount rates of 10.43 percent, compared to an average of 5.91 percent for households who reported “ten years or more.”

Evidence from wealth holding allows us to determine whether parameters used in studying the financial incentives to retire are reasonable. Although a uniform discount rate of 3 percent is fine for the descriptive, univariate analyses in this chapter, more attention must be paid to the choice of parameters in subsequent econometric models. The benchmark should be a higher median value that is consistent with the median household’s low financial wealth accumulation over the life cycle. The heterogeneity in discount rates, implied by the heterogeneity of wealth holdings, should also be incorporated into the econometric specification as a topic for further research. This will necessitate a return to the structural estimation of parameters in the option value, as in the original work of Stock and Wise (1990a, b).

Conclusion

Overall, this chapter provides an interesting starting point for the growing literature that uses the rich data in the HRS to estimate the effect of financial incentives on the timing of retirement. The main contribution is to show that, when modeled with careful attention to detail, Social Security itself yields both heterogeneity and, to a lesser extent, nonmonotonicities in retirement incentives. These findings suggest a potentially larger role for Social Security in explaining cross-sectional retirement patterns than is currently believed.

The authors also make a strong case that, for the purpose of providing descriptive statistics on forward-looking measures of incentives, peak value is useful because it does not rely on parameters of an indirect utility function. Most of that advantage will not be present in econometric analyses, for two reasons. First, including the present value of earnings in the regression accounts for most of the difference between calculations of peak value and option value. Second, we actually do care about the parameters that are chosen and about gaining a deeper understanding of the appropriate model for retirement decisions. Further improvements should enable researchers to estimate the role of pensions more precisely, and to provide better links between retirement and saving behavior.

References

- Atkinson, Anthony B. 1987. Income maintenance and social insurance. In *Handbook of public economics, Vol. 2*, ed. Alan J. Auerbach and Martin S. Feldstein. New York: North-Holland.
- Berkovec, James, and Steven Stern. 1991. Job exit behavior of older men. *Econometrica* 59:189–210.

- Carroll, Christopher D., and Andrew A. Samwick. 1997. The nature of precautionary wealth. *Journal of Monetary Economics* 40:41–72.
- Coile, Courtney. 1999. Social Security, pensions, and the retirement decisions of couples. Manuscript, Massachusetts Institute of Technology, April.
- Coile, Courtney, and Jonathan Gruber. 2000. Social Security and retirement. NBER Working Paper no. W7830. Cambridge, Mass.: National Bureau of Economic Research, August.
- 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*, ed. Robert Willis and James Smith, 150–208. Ann Arbor: University of Michigan Press.
- . 2000. Evaluating pension entitlements. In *Forecasting retirement needs and retirement wealth*, ed. P. Brett Hammond, Olivia Mitchell, and Anna Rappaport, 309–26. Philadelphia: University of Pennsylvania Press.
- Gustman, Alan L., and Thomas L. Steinmeier. 1998. Changing pensions in cross-section and panel data: Analysis with employer provided plan descriptions. *National Tax Association Proceedings—1998*, 371–77.
- . 2000. Retirement in dual-career families: A structural model. *Journal of Labor Economics* 18:503–45.
- Kotlikoff, Laurence J., and David A. Wise. 1987. The incentive effects of private pension plans. In *Issues in pension economics*, ed. Zvi Bodie, John B. Shoven, and David A. Wise, 283–339. Chicago: University of Chicago Press.
- Lumsdaine, Robin L., and James H. Stock, and David A. Wise. 1992. Three models of retirement: Computational complexity versus predictive validity. In *Topics in the economics of aging*, ed. David A. Wise, 19–57. Chicago: University of Chicago Press.
- Luzadis, Rebecca A., and Olivia S. Mitchell. 1991. Explaining pension dynamics. *Journal of Human Resources* 26:679–703.
- Quinn, Joseph, Richard Burkhauser, and Daniel Myers. 1990. *Passing the torch: The influence of economic incentives on work and retirement*. Kalamazoo, Mich.: W. E. Upjohn Institute for Employment Research.
- Samwick, Andrew A. 1993. The effect of tax reform act of 1986 on pension coverage: Evidence from pension provider surveys. Manuscript, National Bureau of Economic Research, July.
- . 1995. The limited offset between pension wealth and other private wealth: Implications of buffer-stock saving. Manuscript, Dartmouth College, December.
- . 1998a. Discount rate heterogeneity and Social Security reform. *Journal of Development Economics* 57:177–46.
- . 1998b. New evidence on pensions, Social Security, and the timing of retirement. *Journal of Public Economics* 70:207–36.
- Stock, James H., and David A. Wise. 1990a. Pensions, the option value of work, and retirement. *Econometrica* 58:1151–80.
- . 1990b. The pension inducement to retire: An option value analysis. In *Issues in the economics of aging*, ed. David A. Wise, 205–24. Chicago: University of Chicago Press.