GARY BURTLESS

1 *An Economic View of Retirement*

Retirement is an event with profound personal, social, and economic consequences. Economists, not surprisingly, focus on its financial ramifications. When workers withdraw from the labor force in old age, their earnings cease and must be replaced by some other source of income. Nowadays the main alternative sources of income are employer-provided pensions and social security. But these only became important within the past half century. Before that, workers who wished to retire were ordinarily forced to rely on transfers from relatives, their own savings, public aid, or charity.

Economists' interest in retirement has focused on three interrelated phenomena: saving in anticipation of retirement, the timing of retirement, and the impact on the economy of transfers needed to support a large retired population. This chapter focuses on the timing of retirement and its relationship with lifetime consumption.

Trends in Retirement

When workers retire, they withdraw from their normal occupations and reduce their work effort or stop work altogether. At the turn of the twen-

I gratefully acknowledge the research assistance of J. J. Prescott and Stacy Sneeringer.

tieth century, retirement was relatively rare but not unknown; two out of three men past age 65 were employed, but one-third were not.[1] By mid-century retirement was far more common: less than half the men 65 and older held jobs in 1950. By 1990 just 16 percent of men older than 65 were employed or actively seeking a job. The proportion of women past age 65 who were employed also fell during the century, but the reduction was far smaller than among men because the percentage of older women in paid work has always been low.

The decline in labor force participation among older men has not been confined to the United States. It is characteristic of all rich industrialized countries. In some European countries, employment rates among the elderly are now significantly below those in the United States.[2] Along with a shrinking workweek and rising labor force participation among women, earlier retirement among men has been a distinctive feature of economic development in the rich countries.

That older men are increasingly retired is clearly evident in figure 1-1. Each line traces the labor force participation rate of older American men, by age, during a different year of the twentieth century (a person is considered a labor force participant if he or she holds a job or is actively searching for work). The top line shows age-specific participation rates of older men in 1910. There was a clear pattern of labor market withdrawal with advancing age. Even at age 74, however, the participation rate in 1910 was only slightly below 50 percent. Participation rates in 1940, 1970, and 1995–96 also show a characteristic pattern of labor market withdrawal as men grow older, but the fall-off begins at an earlier age and proceeds at a faster pace.

The decline in male participation has been neither smooth nor uniform over the century. Figure 1-2 shows the amount of decline in participation at each year of age, measured as a percentage of the 1910 participation rate at the same age. By far the largest proportionate declines have occurred among men past the age of 66. In 1996, for example, the participation rate among 74-year-olds was nearly 80 percent below the equivalent rate in 1910. The fall-off in participation has been proportionately smaller at younger ages.

The shading scheme in the figure shows how fast participation rates fell in different periods. In general, large declines occurred early in the cen-

1. Bureau of the Census (1975, p. 132). Retirement patterns are much more difficult to measure among women because most worked primarily within the home (and without pay) during most of their adult lives.

2. Quinn and Burkhauser (1994, pp. 56–61).

Figure 1-1. *Male Labor Force Participation, by Age, Selected Years, 1910–96*

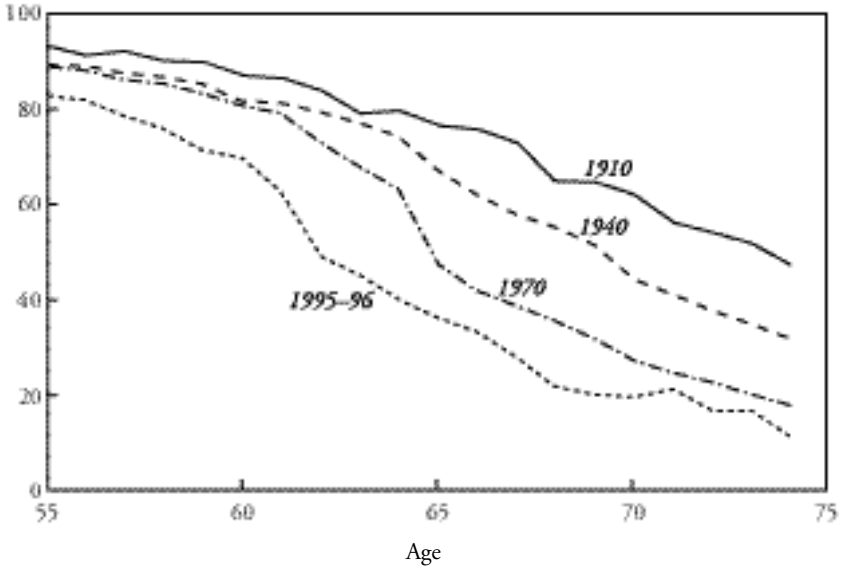

Source: Labor force participation rates for 1910, 1940, and 1970 are based on responses to employment questions in the decennial censuses. See Ransom, Sutch, and Williamson (1991, pp. 45–46); and Munnell (1977, p. 70). Rates for 1995–96 are the arithmetic average participation rates on the March 1995 and March 1996 Current Population Survey (CPS) files.

tury for the oldest age groups; large declines have occurred more recently among younger groups. The largest percentage declines among men older than 70 occurred between 1910 and 1940. The fastest declines among those aged 65 to 69 took place between 1940 and 1970. The biggest declines among men younger than age 65 did not occur until after 1970. As we shall see, this pattern of labor market withdrawal is consistent with the view that the introduction and liberalization of social security was an important factor in pushing down the rates. Social security old-age pensions were first paid in 1940, and they were first made available to men aged 62 to 64 in 1961.

The basic pattern in figures 1-1 and 1-2 is that, although retirement has been present throughout the twentieth century, it is now more prevalent and occurs at a much younger age. Figure 1-3 shows the trend in the "average" retirement age, if that age is defined as the youngest age at which fewer than half the men in the age group remain in the labor force. Under

Figure 1-2. *Reductions in Male Labor Force Participation Rate, by Age, Selected Periods, 1910–96*

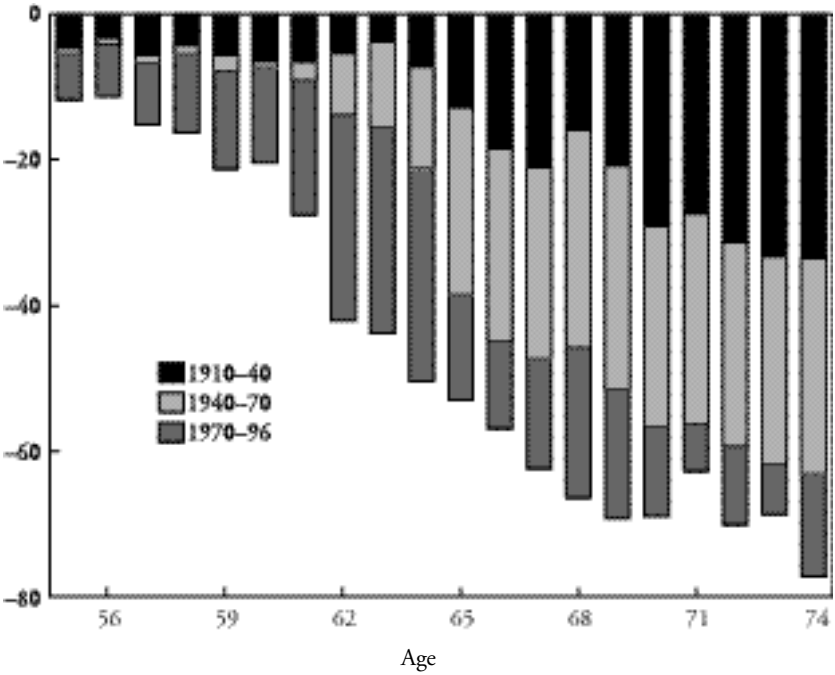

Source: Ransom, Sutch, and Williamson (1991, p. 45); Munnell (1977, p. 70); and author's tabulations of Current Population Survey, March files.

this definition, the average male retirement age fell from 74 to 62 between 1910 and 1996, a drop of about 1.4 years a decade.

The decline in the average retirement age has occurred in an environment of rising life expectancy among older Americans, especially since 1940 (table 1-1). Falling mortality rates among the elderly have added 3 years to the expected life span of a 65-year-old man and 5.5 years to the life expectancy of a 65-year-old woman since 1940. Because expected male life spans increased about 0.6 year a decade during a period in which the retirement age dropped 1.4 years a decade, the amount of the male life span devoted to retirement climbed about 2 years a decade. Retirement now represents a substantial part of a typical worker's life. For many if not most workers, retirement will last longer than the period from birth until full-time entry into the job market.

Figure 1-3. *Average Retirement Age of Men, 1910–96*[a]

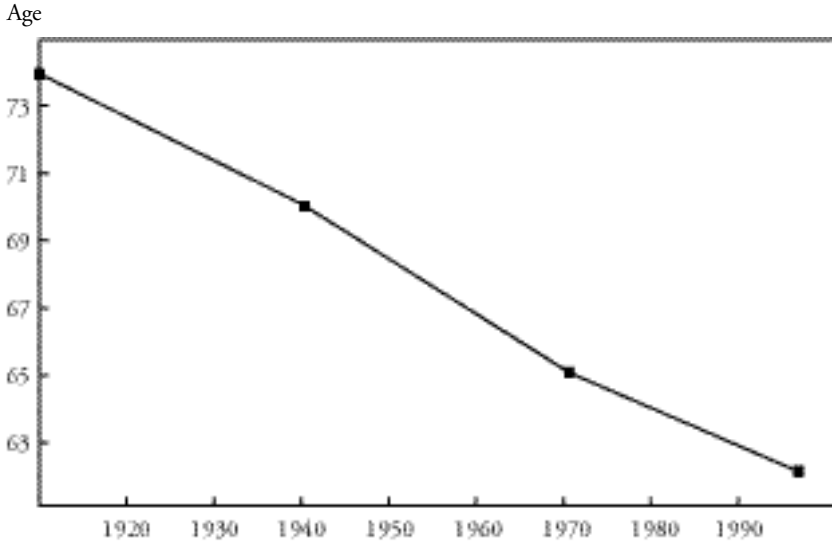

Source: Author's tabulations based on sources for figure 1-2.

a. The average retirement age is the earliest age at which the labor force participation rate of older men drops below 50 percent.

Most of the early U.S. research on retirement trends was conducted by analysts in the Social Security Administration using survey information from retired workers receiving social security benefits or workers who had recently retired. This research, which dates to the mid-1940s, has been summarized by Joseph Quinn, Richard Burkhauser, and Daniel Myers.[3] In the earliest surveys an overwhelming majority of male respondents reported retiring because they were laid off by their latest employer or were in such poor health that further work was unappealing or impossible. In fact, these explanations for retirement dominated the survey responses from the 1940s through the early 1970s. Only a very small percentage of men reported leaving work because they wanted to retire. Quinn quotes an early analyst as suggesting that "most old people work as long as they can and retire only because they are forced to do so. . . . [O]nly a small proportion of old people leave the labor market for good unless they have to."[4]

In more recent surveys of new social security beneficiaries the proportion of workers who say they have retired to enjoy additional leisure or

3. See Quinn, Burkhauser, and Myers (1990, pp. 43–53); and Quinn (1991, pp. 119–23).

4. Quinn (1991, p. 120).

Table 1-1. *Life Expectancy, by Gender and Decade, 1900–2070*[a]

	Life expectancy at birth		*Life expectancy at age 65*	
Year	*Male*	*Female*	*Male*	*Female*
Actual				
1900	46.4	49.0	11.4	11.7
1910	50.1	53.6	11.4	12.1
1920	54.5	56.3	11.8	12.3
1930	58.0	61.3	11.8	12.9
1940	61.4	65.7	11.9	13.4
1950	65.6	71.1	12.8	15.1
1960	66.7	73.2	12.9	15.9
1970	67.1	74.9	13.1	17.1
1980	69.9	77.5	14.0	18.4
1990[b]	71.1	78.8	14.9	18.9
Projected[c]				
2000	72.6	79.7	15.4	19.4
2010	74.0	80.5	15.8	19.7
2020	74.7	81.2	16.3	20.2
2030	75.3	81.8	16.7	20.6
2040	75.9	82.4	17.1	21.1
2050	76.5	82.9	17.5	21.5
2060	77.0	83.5	17.9	22.0
2070	77.5	84.0	18.3	22.4

Source: Office of the Actuary, Social Security Admnistration.

a. Life expectancy for any year is the average number of years of life remaining for a person if that person were to experience the death rates by age observed in, or assumed for, the selected years.

b. Estimated.

c. Based on the intermediate mortality assumptions of the 1993 *Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds.*

other purely voluntary reasons is plainly on the increase. Figure 1-4 presents Quinn's summary of survey responses by men aged 65 and older to the question "Why did you retire?" He has divided responses into four broad categories: "Lost last job," "Health reasons," "Wished to retire," and "Other." This classification is not precise, because survey questionnaires were not always consistent in the way they framed the question or the possible responses a retiree might offer.

The growing importance of voluntary retirement is clearly apparent in figure 1-4. In the 1940s and early 1950s fewer than 5 percent of new retirees reported leaving work because of a wish to retire or enjoy more

Figure 1-4. *Reasons for Retirement among Retired Men Aged 65 and Older, Selected Years, 1940–80*

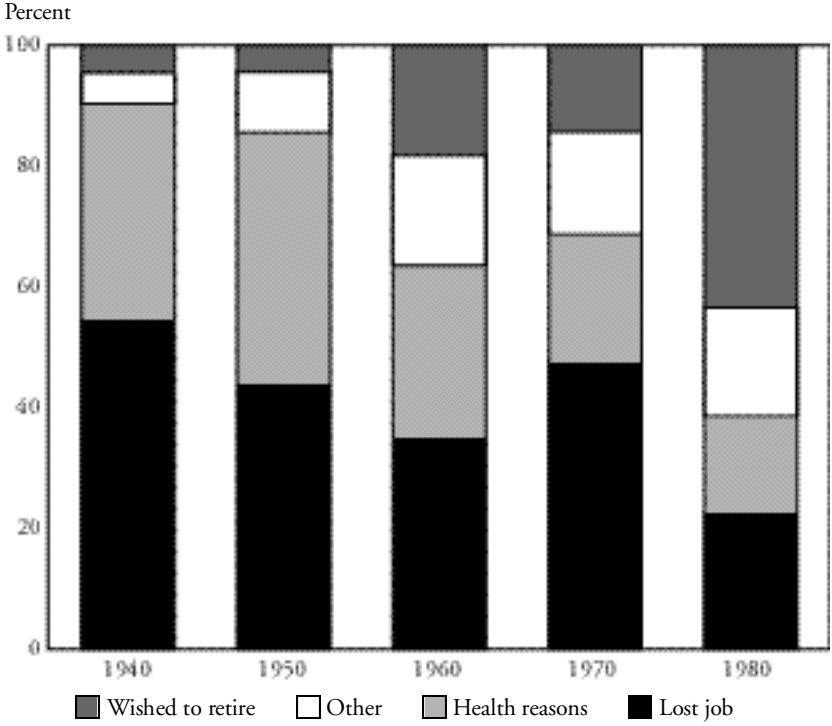

Source: Quinn (1991, p. 123). Years given are approximate year of surveys.

leisure. About 90 percent left because of poor health or a layoff. By the early 1980s the desire to leave work explained nearly half of all retirements among men aged 65 or older, while poor health accounted for only a little over a fifth and involuntary layoff about 15 percent.

Many readers might accept these responses at face value, but economists treat them more skeptically. For example, from 1940 through the early 1970s well over a third of respondents explained their entry into retirement as the result of involuntary job loss. Although this explanation might seem reasonable to noneconomists, hard-headed practitioners of the dismal science realize that millions of workers lose their jobs each year without becoming retired. The overwhelming majority of workers who offer "job loss" as the explanation for their retirement probably lost several jobs during their careers, but on no previous occasion did a layoff cause them to withdraw permanently from the work force. When forced into unem-

ployment at a younger age, they looked for another job and eventually found one. It is natural to ask why job loss pushed them into retirement on this one occasion but not on the others. Even the health explanation arouses skepticism among some economists. Social security beneficiaries may justify their retirement with the excuse that poor health leaves them no alternative, but many economists wonder whether their decision would have been any different if social security were not available.

Economist and noneconomist can both agree, however, that the persistent trend toward earlier retirement has an important voluntary component. The survey responses plainly show it, and the responses are consistent with what we know about increasing wealth, wider eligibility for public and private pensions, and the increasing generosity of pensions.

Retirement and Consumption

The downward drift in the labor force participation of older men has had a significant effect on the growth in aggregate labor supply, both in the United States and in other industrialized countries. The long-term trend in the retirement age did not attract much notice from economists until the 1970s. Curiously, the first aspect of retirement to command their attention was the effect—or hypothesized effect—of retirement on individual and national saving. The classic statement of this relationship is contained in articles written or coauthored by the economist Franco Modigliani.[5] His theory has had a wide influence on economists' thinking about the timing of retirement as well as the determination of saving. For that reason it is useful to describe the theory in some detail.

Theory

Modigliani's basic idea was that farsighted workers will rationally plan their consumption over a full lifetime. In devising their lifetime consumption plans, they will take account of the likely path of their labor earnings as they age and will prudently accumulate savings in anticipation of their retirement. The goal of a good consumption plan is to maximize the worker's lifetime well-being, subject to the constraint that lifetime consumption cannot exceed the worker's lifetime wealth. Lifetime wealth consists of the worker's initial assets and the present discounted value of antic-

5. Modigliani and Brumberg (1954); and Ando and Modigliani (1963).

ipated labor earnings and other kinds of income such as public assistance and inheritances that are not derived from initial assets or labor earnings. Rational and farsighted workers will plan to avoid situations in which all their lifetime wealth has been consumed long before they expect to die. In the absence of transfers from relatives, public aid, or private charity, the consequences of this kind of planning error might be unappealing.

It is possible to sketch out the implications of this theory for the path of consumption and wealth accumulation using a few simple assumptions. The top panel in figure 1-5 shows the path of a worker's expected earnings over a lifetime. When he first enters the labor force at age 20 his earnings are just $10,000 a year, but they will climb rapidly as he gains work experience. His earnings reach a peak around age 50 and then gradually decline. He withdraws from the work force on his sixty-fifty birthday, at which point his earnings cease. By assumption, he is completely certain about the path of his future earnings, his age at death (85), and the interest rate throughout his life (5 percent). The worker is able to borrow freely at this interest rate, and if he accumulates savings he will receive the same interest rate on his investments. It is also assumed that the worker has stable preferences throughout his life.

A worker who successfully solves the consumption planning problem under these assumptions will plot out a desired path of consumption for each future year of life and will then stick with the plan. The best plan will depend on the relationship between the worker's subjective rate of time preference and the interest rate he can obtain on his savings. The rate of time preference is a measure of the worker's impatience in consumption. People who insist on consuming nine-tenths of a box of chocolate truffles today, leaving only one-tenth of it for consumption tomorrow, are said to have a high rate of time preference; they are very impatient in their consumption.

If the worker's rate of time preference is equal to the market interest rate, the consumption path will be level throughout the worker's life, as I have drawn it in the figure. If instead the rate of time preference is higher than the interest rate, he will attempt to shift his consumption toward the early part of his life, and his consumption will fall as he grows older. People with a very low rate of time preference, those who are very patient in their consumption, will shift consumption to later stages of their life and will plan to increase consumption as they age. Workers may wish to leave bequests to survivors, in which case they will consume all their lifetime wealth except the amount needed to leave to their heirs. I assume in figure 1-5 that the worker plans to make no bequests and will thus consume all his lifetime wealth by the time he dies.

Figure 1-5. *Hypothetical Life-Cycle Consumption and Wealth Accumulation*

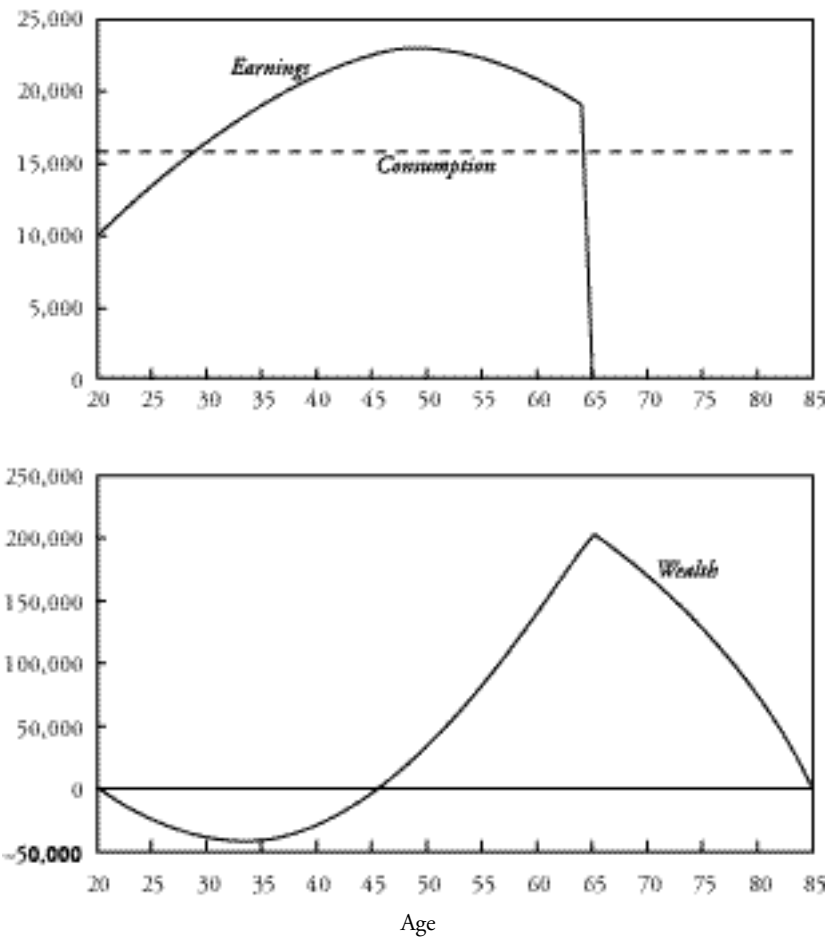

The resulting consumption path is shown as the horizontal line in the top panel. Because the worker's consumption is initially higher than his earnings, he must borrow money when he is young to finance consumption.[6] In the benign world assumed here, he can borrow as much money as he wants as long as he can repay the loan out of his lifetime wealth. The lower panel in figure 1-5 shows the lifetime path of the worker's asset holdings. He accumulates increasing amounts of debt to finance his con-

6. He must borrow more than the difference between his earnings and his consumption because he also needs to borrow enough money to make interest payments on his outstanding debt.

sumption until he reaches age 33, when rising earnings allow him to begin paying off liabilities. By age 45 he is free of debt and begins to accumulate assets. The peak of his asset accumulation occurs on his sixty-fifth birthday, the day he retires. His assets are then used up over the last twenty years of his life, when he has no labor income to help pay for consumption.

The characteristic pattern of increasing and then declining asset holdings over the life cycle is one of the central empirical predictions of Modigliani's consumption theory. The asset buildup would not be needed if workers did not expect to retire. In the absence of retirement, saving would be needed mainly to finance bequests and smooth out consumption in comparison with earnings. Readers may notice that some of the assumptions I have mentioned are highly stylized and are unlikely to be true in life. For example, workers cannot borrow money at the same interest rate they obtain on their investments: usually they must pay a much higher rate than the one they can safely earn. More important, few workers can borrow large sums of money to finance current consumption. These are comparatively minor matters for many workers, however, and addressing them does not fundamentally alter many of the implications of the theory.[7]

A more serious problem arises when a realistic picture of the worker's uncertainty about the future is introduced. In formulating an ideal consumption plan, it obviously helps if the worker is completely confident about his future earnings, his age at retirement and at death, and the future interest rate. But no one can predict these with much confidence. In formulating the consumption plan, one must take account of the possibility that the future may turn out to be more or less congenial than anticipated.

In one respect, the life-cycle consumption theory and the closely related permanent-income model represented a major advance in economists' understanding of how consumers handle unexpected events.[8] The two models make a clear and plausible distinction between (unanticipated) changes in flows of income that can be expected to last and changes that are only temporary. According to both theories, an unexpected income

7. Even if borrowing constraints prevented workers from borrowing money early in their careers, as they are implicitly able to do in figure 1-5, the life-cycle model still predicts that workers who expect to retire will accumulate substantial wealth holdings toward the end of their active careers. Thus retirement decisively affects the lifetime pattern of asset accumulation and deaccumulation.

8. The permanent-income model was introduced by Milton Friedman (1957) around the same time that Modigliani and his collaborators proposed the life-cycle consumption model. The empirical predictions of the two theories are virtually identical, leading many students of consumption to refer to a common life-cycle–permanent-income model. I have emphasized the life-cycle variant here because early versions of it treated the issue of retirement planning as the central motivation for consumer saving.

improvement that is permanent, such as an earnings gain that accompanies a promotion, will have a much larger impact on a worker's consumption than an improvement that is only temporary, such as a one-time bonus for outstanding job performance. By the logic of the life-cycle model, a person who wins a lottery that pays $10,000 a year for thirty years will plan to make a much bigger change in short-term consumption habits than the person who wins a one-time prize of $10,000. By the same reasoning, the lottery winner who obtains a prize paying a modest annual amount (say $700 a year) that has a present discounted value of $10,000 will alter consumption by the same amount as the winner of a one-time prize equal to $10,000.

Consumers still face the problem of deciding whether an income change will be long-lasting or only temporary. And if it is long-lasting, when will it cease? These considerations are crucial in determining how much workers should adjust their flow of consumption once they have obtained new information about future income flows. In theory, alert consumers will formulate a new lifetime consumption plan every time they receive new information about the future. If an employer's quarterly earnings statement shows an unexpected drop in profits, employees in the company should scale back their consumption in anticipation of layoffs or slower future wage growth. If interest rates rise, workers may postpone consumption until later in life to take greater advantage of improved earnings on their investments. If a worker suffers an unexpected heart attack, he may boost his saving in anticipation of an earlier retirement and lower lifetime earnings.

New information about the future state of the world is seldom clear cut. Does a heart attack mean that retirement will last longer because the victim may be forced to leave the work force earlier? Or will it shorten retirement because the worker can expect an earlier death? The two outcomes, if fully anticipated, would have opposite effects on the rate of consumption over the remainder of life, but a farsighted worker will take account of both possibilities in formulating a consumption plan. Will an interest rate hike be temporary or permanent? Even financial market specialists do not have enough information to answer this question confidently.

Evidence

To say that solving the consumption planning problem under uncertainty is difficult does not mean it is impossible. Workers who devote enough intelligence and attention to the problem will usually make more prudent

and satisfying decisions about consumption than those who approach the issue casually or ignore it altogether. The life-cycle–permanent-income theory has produced important insights into consumption planning. Among economists it remains by far the most influential model of consumption. Whether it provides an accurate explanation for observed consumption behavior remains an open question.

Some evidence supports the theory. Most empirical research suggests that the model is correct in emphasizing that households discount short-run fluctuations in their income when determining current consumption and that retirement is one important motive for saving. There is competing evidence, however, that consumption is more volatile and closely related to current income changes than would be the case if there were complete smoothing of consumption over full lifetime resources.[9] As the theory predicts, economists observe a tendency among many workers to steadily but gradually build up their wealth, increasing their rates of saving in peak earning years and as they approach retirement. The life-cycle theory's implication that consumers have a target wealth-income ratio that increases with age up to retirement also seems to be valid for many households.

Nonetheless, some economists are doubtful of the theory because simple versions of it are not very successful in accounting for important aspects of personal saving. For example, many American workers enter retirement without any assets. A large percentage of others who do have assets apparently continue to add to them after retirement. Neither fact is easy to reconcile with simple versions of the life-cycle model. Theorists are thus forced to adopt modifications in the basic theory to account for obvious empirical contradictions.[10] Different theorists have proposed different modifications to rescue the basic model. Whatever their criticisms of the model, however, few have strayed far from it in trying to explain retirement behavior.

Economic Models of Retirement

Economic theories of retirement naturally focus on financial aspects of workers' decisions. This section describes some of the financial considera-

9. See, for example, Skinner (1988); Zeldes (1989); and, in particular, Carroll and Summers (1991).

10. For example, Hubbard, Skinner, and Zeldes (1994) rescue the model by theorizing that income uncertainty and asset-tested transfer programs erode the incentive to save for a sizable minority of households.

tions affecting workers' choice of retirement age. It then considers the theories economists have advanced to explain retirement choice and the evidence they have used to test them. Although economic studies on choice of retirement age did not begin in earnest until the mid-1970s, their number has grown explosively since then.[11] Rather than provide another survey of them, I will focus on a handful to highlight the growing complexity of economists' models of worker decisionmaking.

Financial Aspects of Retirement

Modigliani's life-cycle consumption model emphasizes the single most important financial aspect of retirement—the sharp reduction or complete cessation of labor earnings. Most worker households rely heavily on labor earnings to pay for consumption. When earnings cease at retirement, workers must find another way to pay for it. Modigliani stressed personal saving as an alternative source of support in old age. Even though other income sources are now more important (and still others may have been more important in the past), it is useful to think about the choice of retirement age in a world in which retired workers rely solely on their own savings to finance consumption.

Consider a worker who can earn exactly $10,000 in each year she is employed. If she begins working at age 20 and has been reliably informed she will expire on her seventieth birthday, she can work for up to 50 years, potentially earning as much as $500,000 over her life. To keep the calculations simple, let us also assume the interest rate and the worker's rate of time preference are the same and are exactly 0 percent. Under that assumption, she will plan to consume her lifetime wealth at a constant rate over her life span. Her lifetime wealth in this case is simply her lifetime earnings, which in turn are equal to $10,000 times the number of years she chooses to work. If she works 40 years, for example, her lifetime wealth will be $400,000, and she will consume this amount at a rate of $8,000 a year ($400,000 / 50 years).

In this highly stylized case, the worker's retirement choice can be described as a simple trade-off between a higher flow of consumption per year and a longer period of time spent in retirement (figure 1-6). If the worker retires at age 20, she will earn no money, accumulate no assets, and consume $0 a year throughout her life (which is likely to be brief). If

11. For good surveys, see Quinn, Burkhauser, and Myers (1990); Leonesio (1993); Quinn and Burkhauser (1994); and Lumsdaine (1996).

Figure 1-6. *Trade-off between Annual Consumption and Age at Retirement*

she works until age 70, she can consume all of her annual wages ($10,000) each year. It seems reasonable to expect that the worker would prefer to consume more each year (holding fixed the amount she works) and would prefer to retire sooner rather than later (holding fixed the amount she consumes). The illustration in figure 1-6 should be familiar to anyone who has studied freshman economics. The worker's problem is to select the best possible combination of consumption and years of retirement in light of her preferences and the trade-off shown in the diagram. In the figure, the most desirable combination occurs if the worker retires at age 50 after working 30 years and consumes her lifetime earnings at a rate of $6,000 a year. Other workers facing the same trade-off might choose to retire at a younger or older age, depending on their preferences for consumption and retirement.

Although the analysis may seem trivial, it can shed light on the retirement trends discussed earlier. Retirement ages have declined in the twentieth century, and the simple economic model suggests three reasons this might have occurred. Workers may now enter their careers with a higher flow of income from inherited assets. This provides a straightforward explanation for earlier retirement because under plausible assumptions about worker preferences greater initial wealth will induce most workers to spend

more of their life in retirement.[12] Workers also earn much higher wages in the 1990s than they did in 1910 or 1940. The higher rate of pay provides greater lifetime wealth at any fixed retirement age, which may induce workers to spend more years in retirement.[13] Of course, a higher wage also increases the financial penalty on workers who retire early. If a worker's earnings doubled from $10,000 to $20,000 a year, she would give up $20,000 in lifetime wealth rather than only $10,000 if she chose to retire one year earlier. The larger penalty for early retirement might offset some or all of the effect of higher lifetime wealth. Finally, worker preferences may have changed. Independent of the change in initial wealth or yearly pay, workers today may simply prefer to spend more of their life in retirement.

Social Security and Pensions

Other changes in the environment have also affected the trade-off between consumption and the age at retirement. Employer-sponsored pensions are now much more prevalent than was the case fifty or a hundred years ago. Workers covered by a pension are provided with a potential source of income in addition to their own savings for financing retirement. Social security, which was introduced in 1935 and greatly liberalized between 1950 and 1975, also offers a reliable source of income in old age.

Social security and pensions affect the lifetime trade-off between consumption and retirement in a complicated way. Here I will focus on social security. Workers who become eligible to receive benefits under social security are entitled to receive a pension starting at age 62 or when they retire, whichever occurs later. Because the system has historically been very generous, all generations retiring up to the present have received larger pensions than their contributions could have paid for if the contributions had been invested in safe assets. In effect, this generosity increased the lifetime wealth of older workers who became vested in the system. If they consumed all of the benefits paid to them, they enjoyed higher lifetime consumption than their labor income alone could have financed.[14] The for-

12. In terms of figure 1-6, higher initial wealth will lift the lifetime budget constraint to a position that is above and parallel to the budget constraint in the figure. Because of inherited wealth, the worker might be able to consume $4,000 a year rather than only $0 if she retired at age 20.

13. Higher wages cause the budget constraint drawn in figure 1-6 to rotate counterclockwise around the origin.

14. This kind of generosity is possible in a pay-as-you-go retirement system because early contributors to the system make much smaller contributions to the program than will be needed when the system is fully mature. Future generations will not receive such a good deal under social security. Young workers today may actually receive lower benefits under the system than their contributions would have earned if they had been invested in safe assets, such as government bonds.

tunate generations that received this windfall may have retired earlier than they would have if social security had not been introduced or if it had offered less generous pensions.

The effect of social security on retirement depends on the social security tax and the benefit formula linking monthly pensions to a worker's past covered earnings. Employers and workers pay into the system a combined tax equal to 12.4 percent of wages. The tax thus reduces workers' wages by about 12 percent in comparison with the wages they would receive if the program were abolished. But, of course, contributions allow a worker to earn credits toward a social security pension. The pension entitlement goes up as the worker's covered lifetime wages increase. Whether the increase in the pension entitlement is large enough to compensate a worker for his extra contributions is an empirical question. Low-wage workers receive favorable treatment under social security, so they usually receive a generous return on their contributions. High-wage workers typically receive lower returns.

Workers who delay their retirement until after age 62 are at least temporarily passing up the opportunity to receive a social security check, which can begin immediately after the worker's sixty-second birthday. If a worker is entitled to a pension of $500 a month, for example, he is sacrificing $500 in retirement income every month he delays retirement past age 62. If his regular monthly pay is $10,000, this represents a small sacrifice. But if his usual pay is $1,000, the sacrifice amounts to half his wage. Between the ages of 62 and 64 the social security formula offers workers a fair compensation for giving up a year's benefits. Monthly benefits are adjusted upward about 8.5 percent for each year's delay in claiming a pension. For workers with average life expectancy and a moderate rate of time preference, this adjustment is just large enough so that the sacrifice of a year's benefits is compensated by eligibility for a higher pension in future years. After age 65, however, the benefit formula is much less generous toward delayed retirement. Postponement of retirement after that age is not fairly compensated by increases in the monthly pension.

The reason that many people must retire in order to collect a social security check is that the program imposes an earnings test in calculating the annual pension. Workers who are between 62 and 64 and who earn more than $8,640 a year lose $1 in annual benefits for every $2 in earnings they receive in excess of $8,640. Workers between 65 and 69 lose $1 in benefits for every $3 in annual earnings in excess of $13,500. (Pensioners age 70 and older do not face an earnings test.) At one time the earnings limitations were much lower and the tax on excess earnings was much higher, dis-

couraging pensioners from work and possibly encouraging them to postpone claiming a pension until they were confident their earnings would remain low.

Social security has, then, boosted the lifetime wealth of older people who have received benefits under the program. Its complicated benefit formula provides an incentive for people to become entitled to benefits but may discourage extra work by high-wage workers, who do not receive good returns on their marginal contributions to the program. At age 62, when pensions can first be claimed, eligible people who continue to work give up a year's benefits every year they postpone their retirement. For workers who have a high rate of time preference or who do not expect to live long, this sacrifice is equivalent to a big cut in annual pay. Between the ages of 62 and 64, however, the pension formula fairly compensates most workers for this sacrifice. Starting at age 65 the compensation formula is much less generous, so workers must sacrifice some lifetime benefits every year they delay retirement past 65. Social security does not provide a simple annuity. Most workers must substantially cut their earnings to collect a full pension.

Figure 1-7 summarizes the potential effects of social security on the consumption–retirement age trade-off. The trade-off is similar to the one shown in figure 1-6, except that the diagram focuses on the retirement age choice late in life. The lower line represents the consumption–retirement trade-off the worker would face in the absence of social security. The trade-off is shown to be linear because of my assumption that the worker can earn exactly $10,000 a year from continued work, an amount that does not change as the worker grows older. The higher dashed line represents the trade-off for the same worker when social security is introduced. The trade-off under the system permits the worker to enjoy higher annual consumption at any fixed retirement age because it increases the worker's pension income by more than it reduces net earnings. The benefit formula also affects the slope of the trade-off, particularly at ages 62 and 65, for reasons mentioned earlier. At age 65 the slope certainly declines because the worker must give up pension benefits for continued work without receiving a fair compensation for the sacrifice. The change in slope at age 62 is less clear cut. From 62 to 64 workers receive better compensation for postponing retirement than they do at age 65, but the compensation may not be regarded as fair by a worker with a high rate of time preference or a short life expectancy. In that case, the change in the slope would look the same as the change in slope at age 65. Workers with a low rate of time preference and long life expectancy might regard the compensation as

Figure 1-7. *Trade-off between Consumption and Age at Retirement under Social Security*

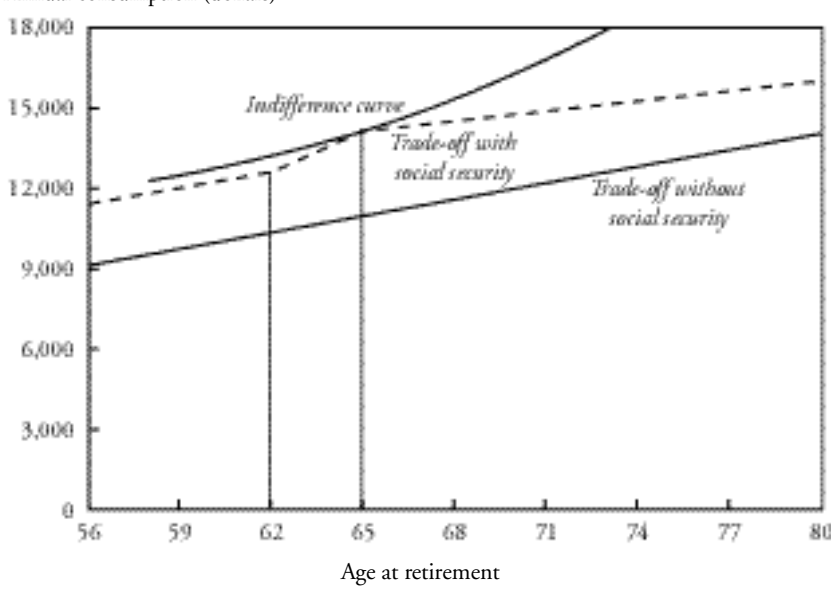

more than fair. In other words, the compensation they receive is more than enough to make up for the loss in pensions connected with delay in their retirement. The slope in that case would increase at age 62, which is the situation shown in figure 1-7.

Employer-sponsored pensions introduce further complications in the trade-off between consumption and retirement. These are similar to the complications introduced by social security. There is one important distinction between social security and employer-sponsored pensions, however. Social security imposes an earnings test on income received from *all* employment, including self-employment. Employer-sponsored pensions may impose an even more stringent earnings test, but the test applies only to earnings received from the sponsoring employer or perhaps a group of related employers. Workers who wish to claim a pension may be forced to leave the job on which they earned the pension, but they are not forced to leave work altogether. Nevertheless, the effects of employer-sponsored pensions on retirement may be similar to those of social security, because many older workers find it hard to obtain good job offers after they have retired from a career job.

Health and Job Opportunities

Changes over time in a worker's health and job opportunities also influence the timing of retirement. Many workers experience deteriorating health as they age, which makes it physically more difficult or less pleasant to hold a job and in some cases can make work practically impossible. If deteriorating health is fully anticipated when a person begins to work, it will affect the perceived attractiveness of a consumption plan that postpones retirement until very late in life. This will be reflected in the worker's preferences across various combinations of average consumption and age at retirement. Anticipating that she will be in poor health when she reaches age 70, the worker will avoid lifetime consumption plans that require her to work past that age unless the plans also provide exceptional amounts of lifetime income. Unanticipated changes in health at older ages force workers to reconsider their lifetime consumption plans in light of their new knowledge about their work capacity. In many cases these health changes will force them to retire earlier than they originally planned.

Many workers face worse job prospects when they are old than they did when they were young. This fact is reflected in figure 1-5, which shows declining wages after age 50. It is not reflected in figures 1-6 or 1-7, which are based on the assumption that workers can earn $10,000 a year no matter how long they work. If potential annual earnings decline as workers grow older, the trade-off between consumption and retirement age will look less favorable than it does in figures 1-6 and 1-7. Each added year of work past some critical age (say, age 50) will produce a smaller gain in annual consumption. The decline in potential earnings is an important reason that older workers withdraw from the labor force rather than search for work when they are laid off. At a younger age, when their job opportunities were better, they would have looked for and accepted a job after being laid off.

Evidence

Before turning to the specific models economists have tried to estimate, it is worth reconsidering the historical evidence in light of this discussion of retirement incentives. Do simple economic models shed any light on the historical trends? Some rough indication of the possible influence of social security on retirement is provided by examining the relationship between social security incentives and the observed distribution of retirement ages.

Social security is now the principal source of cash income of households headed by someone aged 65 or older. Tabulations of the Current Population Survey show that social security benefits account for slightly more than 40 percent of the total cash income of the aged. Among aged households in the bottom 60 percent of the elderly income distribution, social security accounts for more than two-thirds of cash income (figure 1-8). Until 1941 social security provided no income to the aged. Today the program replaces about 42 percent of the final wage earned by a full-career single worker who earns the average wage and claims a pension at age 65. If the worker has a nonworking dependent spouse, the benefit replaces 63 percent of the final wage. Benefits are clearly high enough so they can be economically significant in influencing the choice of retirement age.

The distributions of male retirement ages in 1940, 1970, and 1995–96 are plotted in figure 1-9. The figure shows the percentage of men leaving the labor force at each age from 56 to 70, computed as a fraction of the men in the labor force at age 55.[15] The calculations are based on the data reflected in figure 1-1. Not surprisingly, the retirement distributions for 1970 and especially for 1995–96 show that labor force withdrawal occurred at earlier ages than in 1940. Both the 1970 and 1995–96 distributions also show evidence of clustering in retirement at particular ages. In 1970 the peak rate of retirement occurred at age 65; by 1995–96 the peak occurred at age 62. There are peaks in the distribution of retirements in 1940 at ages 65 and 70, but these are far below the peaks in 1970 and 1995–96.

The discussion of social security suggests an explanation for the clustering of retirements at ages 62 and 65, at least in years after 1940. Workers who continue to work beyond age 65 give up social security benefits for which they are not fairly compensated. We can anticipate that this feature of the benefit formula will encourage retirement at age 65. The clustering of retirements at age 62 is not much more difficult to explain. Starting in 1961 age 62 became the earliest age at which men could claim a social security pension. Before 1961 there was no evidence of retirements clustering at age 62, but by 1970 retirement was more common at 62 than at any other age except 65. By the mid-1990s, age 62 was by a wide margin

15. If the labor force participation rate at age 63 is designated $LFPR_{63}$, the retirement rate at age 63 is calculated as $(LFPR_{63} - LFPR_{62}) / LFPR_{55}$. This calculation ignores the complications involved in computing true cohort distributions and the effects of mortality rates, immigration, and temporary withdrawal from the labor force. It offers a picture of the timing of labor market withdrawal based on the participation choices of men aged 55 through 70 in a particular year.

Figure 1-8. *Social Security as a Share of Total Cash Income of Elderly Households, by Quintile, 1996*

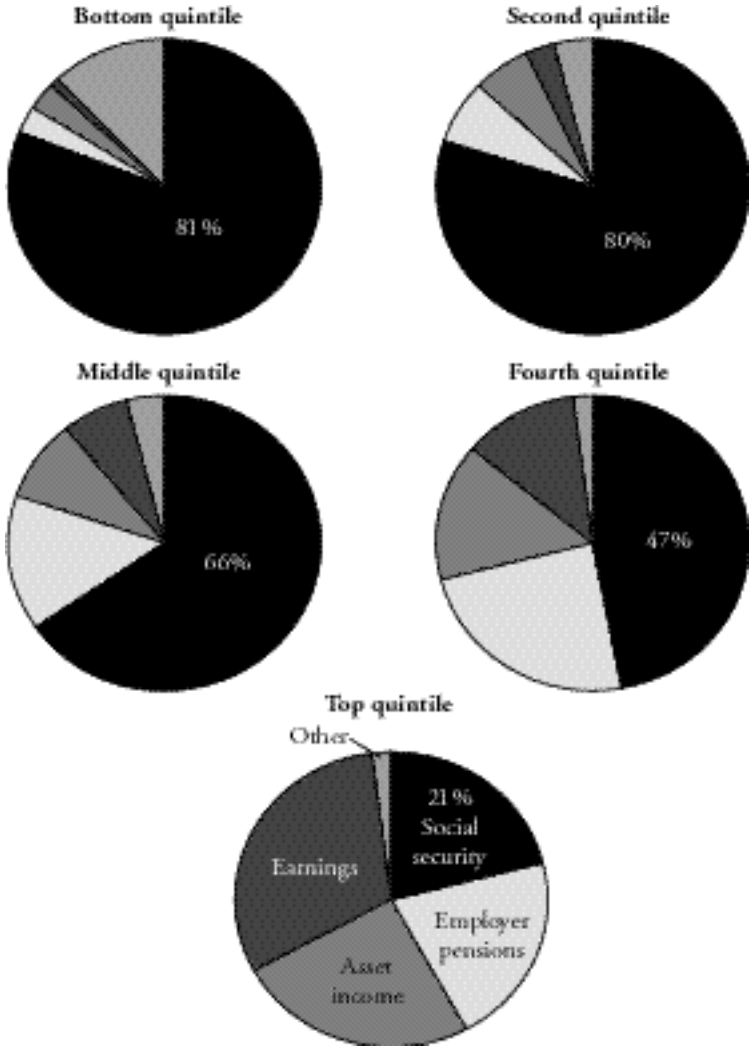

Source: Social Security Administration (1998, pp. 15–16).

the most popular age of retirement. In principle, the social security formula fairly compensates workers if they delay claiming a pension past age 62. But a worker with a high rate of time preference or short life expectancy might not regard the compensation as fair. In those instances one should expect many workers to prefer retiring at age 62 rather than later.

Figure 1-9. *Male Retirement Rate, by Age, 1940, 1970, 1995–96*[a]

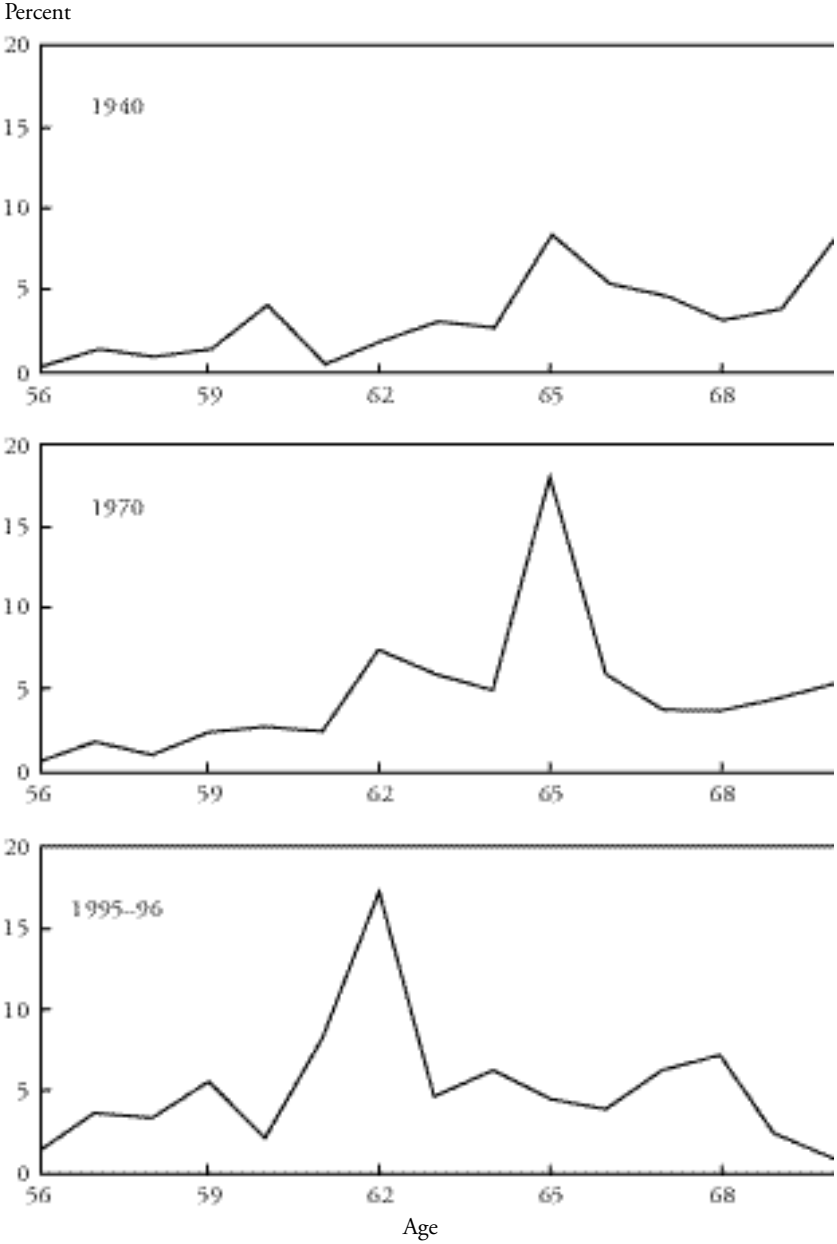

Source: Munnell (1977, p. 70) and author's tabulations of Current Population Survey, March files.

a. Percent reflects the constructed number of men who retire at the designated age, measured as a percentage of men who are labor force participants at age 55.

Of course, the clustering of retirements at ages 62 and 65 may be due to factors other than the availability of social security benefits. It is hard to believe, however, that health or work opportunities decline abruptly at a particular age. Another explanation is that some workers are affected by mandatory retirement rules. This explanation may have been valid in 1940 and 1970, when employers were permitted to dismiss older workers under mandatory retirement rules, but it is not convincing today. Amendments to the Age Discrimination in Employment Act passed in 1986 prohibit employers from dismissing workers solely on account of their age. The simplest alternative explanation for the clustering is that workers are affected by employer-sponsored pension plans. For many older workers private pensions are unlikely to be an important determinant of retirement age, however, because they are not covered by an employer plan. Figure 1-8 shows that employer-sponsored pensions do not provide a large percentage of income to older Americans, except in more affluent households. Among workers who are covered by a private pension plan, some features of the plan may encourage workers to retire at a particular age.

Early Economic Models

Two of the earliest economic studies of retirement were based on very straightforward models of the retirement decision. Both focused on the choice between work and retirement faced by workers in a particular year. Michael Boskin analyzed the choices made by a small number of older men who were interviewed in the Panel Study of Income Dynamics, and Joseph Quinn examined the 1969 retirement choices of a much larger sample of older men interviewed in the Retirement History Survey.[16] In essence, both researchers treated the retirement decision as if it were made independently in each year of a worker's later life based on the financial and health circumstances facing the worker in that year. Under this view, workers eligible for a large social security check or employer pension in a particular year might be tempted to leave the work force in that year. Workers capable of earning a high wage might be induced to remain at work.

This model may actually reflect how workers decide when to retire, but it is based on the assumption that workers are not very forward looking. For example, neither Boskin nor Quinn included variables in his model to reflect the future pensions workers could obtain if they delayed their retirement. A 62-year-old worker who delays retirement one year can increase his

16. Boskin (1977); and Quinn (1977).

monthly social security check by as much as 10 percent. Boskin and Quinn disregarded this aspect of the pension formula under the assumption that workers are motivated by the pension they can collect in the current period rather than the increase in their lifetime wealth (including higher future pensions) that they can obtain by delaying retirement.

Most research on older workers suggests that for a majority retirement is a once-in-a-lifetime event. Recognizing this, well-informed workers may not make an independent decision about whether to be retired each year, but may instead treat retirement as a choice they are likely to make only once. They may want to select the most advantageous retirement age possible in light of the pension for which they qualify at each potential retirement age. Researchers do not know, however, whether workers are actually this farsighted. Under the assumption that they are not very farsighted, both Boskin and Quinn found that social security had a very large impact on their behavior.

Life-Cycle Models

Beginning with a series of articles by Richard Burkhauser, economists started to adopt the life-cycle perspective in their analyses of retirement-age choice.[17] In some cases they examined workers' retirement behavior under specific company pension plans, but most research focused on retirement choice under social security. The information requirements to perform a plausible study are formidable. To analyze the retirement choices of a single cohort of workers requires detailed information about the workers' health, wages, family situation, employment status, and assets over a period that might stretch eight or ten years. To accurately calculate a worker's pension entitlement and determine how the entitlement changes with added work experience, analysts also need extensive information about their past wages and pension service credits. (The social security pension is based on earnings over a forty-year period, for example.)

As economists assembled detailed financial data and put them in a framework that realistically represented the evolution of workers' lifetime pension assets and wealth, their analytical models became increasingly complicated. Part of the complexity reflects the genuine complexity of human behavior. Although many people work in a career job for many years and then leave the job to withdraw permanently from work, others move into transitional jobs with lower pay, less responsibility, and shorter

17. See, for example, Burkhauser (1979).

hours. To reflect the wide range of observed or potential paths into full retirement, analysts sometimes specified in great detail the alternative possibilities from which workers could choose.[18] A wider range of choice usually means that the worker's retirement decision problem becomes more difficult to solve, both for the worker and for the researcher who is attempting to analyze it.

The richest source of information about Americans' retirement behavior is provided by the Retirement History Survey, a ten-year panel survey covering about 11,000 families headed by people who were between 58 and 63 years old when the survey began in 1969. Retirement behavior in these 11,000 families has been analyzed by a number of researchers who applied the life-cycle framework in their studies. In many respects the pattern of retirement and postretirement work effort reflected in the survey offered strong confirmation that social security financial incentives matter in determining labor supply late in life.

Figure 1-10 displays a pair of labor supply distributions based on behavior recorded in the survey.[19] The top panel shows the distribution of retirement ages among nondisabled men who were observed to retire by the end of the survey in 1979, when respondents were between 68 and 73 years old. To determine the retirement age, the analysts examined the lifetime pattern of respondents' work effort and selected the point in each worker's life when he made a discontinuous and apparently permanent reduction in labor supply. This definition excludes spells of unemployment or nonemployment that end with the worker's return to a full-time job. However, the definition would include movements from steady full-time into part-time employment. The picture misses the retirements of some men who did not retire before their last completed interviews, and this omission will lead to some underrepresentation of retirements that occur after age 67. Taking account of the different populations included in the tabulations and the differing definitions of retirement, the pattern of retirement in figure 1-10 is broadly similar to that shown for 1970 in figure 1-2.

The lower panel displays the pattern of earnings among retired but working men who are 62 years old or older in the first interview after they retire. Approximately one-fifth of retiring men were still working within the first two years after their retirements, and on average they worked a little more than sixteeen hours a week. The panel shows the distribution of their earnings in relation to the earnings-exempt amount in the social

18. Gustman and Steinmeier (1986).

19. The calculations are described and presented in Burtless and Moffitt (1985).

Figure 1-10. *Retirement Age and Postretirement Earnings Distributions in the 1969–79 Retirement History Survey*

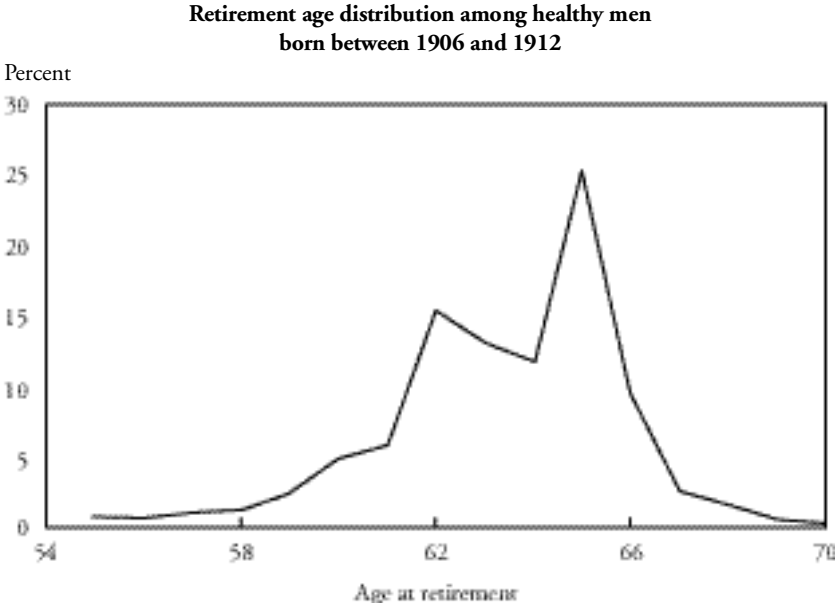

Source: Burtless and Moffitt (1985, p. 225).

security benefit formula. Earnings below the exempt amount had no effect on a worker's pension; earnings above the exempt amount caused benefits to be reduced by 50 percent of the amount of excess wages over the exempt amount. (The tax rate on excess earnings has subsequently been reduced for older retirees.)

Casual observation of the top and bottom panels suggests that social security had a powerful effect on both retirement ages and postretirement work effort. The age distribution of retirements has two peaks, a lower one at age 62, when benefits can first be claimed, and a much higher one at age 65, when the social security formula stops making generous adjustments for further delays in claiming a pension. The distribution of postretirement work effort shows an even larger effect of social security. Workers appear acutely sensitive to the high implicit tax on their earnings when annual wages exceed the exempt amount. More than a quarter of working retirees earn within 10 percent of the exempt amount, and more than half earn within 30 percent of it. Although retirees may underreport their true earnings to social security to avoid paying the high implicit tax, the earnings estimates in figure 1-10 are based on workers' responses to a census interviewer, not their earnings reports to the Social Security Administration. Thus, the distribution is likely to reflect a genuine effect of social security on postretirement hours of work.

In view of the apparently dramatic effect of social security on the exact timing of retirement and on postretirement hours worked, readers may be surprised to learn that most of the life-cycle studies in the 1980s found that even big changes in social security produced only modest effects on lifetime work effort. The diagrams may be a bit deceptive, however. Because fewer than a fifth of retirees choose to work after claiming a social security check and only about two-thirds of these earn enough wages to be remotely affected by the earnings test, the potential effect of a change in the earnings test is limited. Even if the postretirement hours of affected workers rose one-third, the average work effort of all retired men would only climb by about one hour a week (from 3.25 hours to 4.25). The effect is even smaller than this because workers do not remain in their postretirement jobs very long. Gustman and Steinmeier reported that the average duration of partial retirement is only about three years.[20] By implication, the apparently dramatic effect of the social security earnings test has only a tiny effect on workers' lifetime labor supply.

20. Gustman and Steinmeier (1986).

Even the effect on the timing of retirement may not be particularly large. It is true that many more men retire at ages 62 and 65 than at other ages, exactly as predicted by economists' life-cycle model. It does not follow, however, that sharp reductions in social security would cause workers retiring at ages 62 or 65 to retire very much later or that big increases in pensions would cause them to retire much earlier. After a reduction in benefits, for example, the distribution of retirement ages might show less clustering of retirement at particular ages, but the overall effect of the reduction on the average retirement age might be small. This was the conclusion reached in the great majority of the life-cycle studies conducted in the 1980s and early 1990s. Most analysts reached this conclusion by comparing the cross-sectional distribution of retirements within a group of men born around the same years. The behavior of workers with a large amount of social security or pension wealth was compared with that of workers with less wealth. Differences in the rates at which workers could accumulate social security or pension wealth and earn additional wages were also taken into account. Most studies found that even big changes in social security benefits would cause only small changes in the average retirement age. Gary Burtless and Robert Moffitt predicted, for example, that raising the normal retirement age in social security from 65 to 68 would add only a little more than four months to the full-time working careers of nondisabled men.[21]

Another way to analyze the impact of social security incentives is to examine the behavioral differences among people who face different incentives because the program has been altered in an unanticipated way. In 1969 and again in 1972 social security benefits were increased much faster relative to wages than at any time in the recent past. By 1973 benefits were 20 percent higher in inflation-adjusted terms than would have been the case if pensions had grown with wages as they did during the 1950s and 1960s. In 1977 Congress passed amendments to the Social Security Act that sharply reduced benefits to workers born in 1917 and later years (the "notch" babies) in comparison with benefits payable to workers born before 1917. I examined the first episode, and Alan Krueger and Jörn-Steffan Pishke examined the second using the life-cycle framework.[22] In the period I analyzed, workers born in earlier years planned their retirements when social security was comparatively less generous; workers born in later

21. Burtless and Moffitt (1985, p. 230). For a discussion of other economists' predictions, see Quinn, Burkhauser, and Myers (1990, pp. 108–11).

22. Burtless (1986); and Krueger and Pischke (1991).

years planned their retirement when social security was significantly more liberal. Krueger and Pischke analyzed a period in which younger workers received significantly less generous pensions than those available to older workers. Both studies reached an identical conclusion: major changes in social security generosity produced small effects on the retirement behavior and labor force participation of older men. Burtless estimated, for example, that the 20 percent benefit hike between 1969 and 1993 caused only a two-month reduction in the work career of men who were fully covered by the more generous formula. This was equivalent to a reduction in the labor force participation rates of 62-year-old and 65-year-old men of less than 2 percentage points. The effects of the 1977 amendments found by Krueger and Pischke were even smaller. Even if their largest estimated effects are accepted at face value, less than one-sixth of the drop in labor force participation during the 1970s could be explained by changes in social security.[23]

John Rust has proposed by far the most ambitious framework for understanding retirement decisions and their relationship to workers' choices about consumption, wealth accumulation, and full- or part-time work over an entire career.[24] Many students of life-cycle consumption behavior take the retirement age as fixed and try to model the implications for wealth accumulation during the career (see figure 1-5). Labor economists typically take asset holdings at a particular age (say, 54) as fixed and try to model retirement behavior. In contrast, Rust proposes a model that explains the rate of saving and consumption, the path of annual work hours, and the timing of retirement within a single comprehensive framework. In his setup, workers must use dynamic programming techniques to solve their decision problem. At each point during their lives, they seek out the optimal path for future work, consumption, and wealth accumulation, taking account of their preferences as well as their best current information about future wages, interest rates, rules regarding pension benefits, and expected health. Their behavior in a particular period is guided by the solution they have reached to the decision problem. However, their choice in a future period may turn out to be different from the choice they anticipate in the current period because new information can affect the optimal path of work and retirement.

Noneconomists may find this framework unappealing or unrealistic, but some empirical evidence suggests that the choice of retirement age is often

23. Burtless (1986, p. 800); and Krueger and Pischke (1991, p. 24).
24. Rust (1989, 1990).

guided by the application of a sophisticated decisionmaking rule. Robin Lumsdaine, James Stock, and David A. Wise examined the retirement choices of workers under company pension plans using three decisionmaking rules, one of which was based on application of a simple rule while the other two were based on more sophisticated decisionmaking approaches.[25] (One used an "option value" technique for evaluating the value of pension offers, and the other used a dynamic reprogramming rule.) They estimated their models using information from one period and then tried to predict retirement patterns in a later period under the three models. Perhaps surprisingly, they found that the models based on more complex decisionmaking rules were much more successful in predicting future retirement patterns. This evidence suggests that at least some workers use information in a sophisticated way to decide when to retire. Of course, within a well-established company plan that covers many workers in the same workplace, information helpful in choosing an optimal retirement age discovered by one worker can easily be shared with co-workers. Where information sharing is more difficult, workers might rely on simpler decisionmaking rules, and some workers may end up retiring at an age that is less than optimal. Also, the evidence collected by the authors does not show that workers use dynamic programming methods that are as complicated as those assumed by John Rust. Lumsdaine and colleagues analyzed the retirement decision in isolation. They did not assume workers were making farsighted and fully consistent plans for both work and consumption over a time horizon of many years. Even the most sophisticated decision rule they consider is much simpler than the planning methods assumed by Rust.

What Have We Learned?

Research by economists and others has shed light on retirement in the United States. It has shown that the average retirement age of men has declined throughout the twentieth century. Economists have assembled powerful evidence that the exact distribution of male retirement ages is influenced by the financial incentives in social security and private pension plans. However, they have not offered convincing evidence that changes in these incentives can explain a major share of the trend in male labor force

25. Lumsdaine, Stock, and Wise (1992).

participation. Figures 1-1 and 1-3 suggest that the trend toward earlier retirement was already under way during the first decades of the century.

More fundamentally, economists cannot claim to have offered a persuasive explanation for the trend toward earlier retirement within the terms of their basic model. In a trivial sense, of course, the economic model can "explain" earlier retirement. Some combination of changes in wage rates, inherited wealth, pension plan incentives, the population distribution of health, physical and mental requirements for standard occupations, and individual preferences almost certainly accounts for lower participation rates of older men. However, this is a little like explaining the operation of a television set by saying that some combination of metal, plastic, electricity, and electromagnetic signal produces a moving picture on a piece of glass.

In terms of the basic life-cycle model, it would be useful to know what share of the trend can be explained by each of the factors: higher wages, greater initial wealth, private pension and social security incentives, changing health, evolving physical and mental requirements of typical jobs, and changing preferences for consumption and retirement leisure.

It is particularly important to understand the influence of public and private pensions. The cost of the public retirement programs is mounting rapidly, and many observers believe it will be necessary to make them less generous to keep them affordable and politically acceptable. For two reasons it is important to know whether scaling back benefits will cause retirement patterns to change. Policymakers have a practical reason for wanting to know whether retirement rates are likely to change and, if so, by how much. If they cut benefit levels or raise the early retirement age, they need to know how much overall benefit payments are likely to fall. This will be determined, in part, by the work responses of people affected by the cut. If the response is large, the budgetary impact and wider economic consequences will also be large. Both policymakers and the general public also need to understand how much income and consumption levels among the elderly will fall if benefits are scaled back. Social security accounts for more than three-quarters of the cash incomes of elderly families with incomes in the bottom 60 percent of the income distribution. If benefits were scaled back and the average retirement age remained unchanged, low-income retired workers would face large reductions in their old-age income and consumption.

Of the explanations advanced for earlier retirement, two of the least persuasive are declining health and the changing physical requirements of work. While nearly all good retirement studies find that health plays an

important role in the timing of retirement, I can find no convincing evidence that the health of 62-year-olds or 65-year-olds has been declining over the period in which their labor force participation rates have fallen. Their declining mortality rates as well as recent evidence about the trend in their physical disabilities suggest instead that male health is improving, at least in early old age. Moreover, analyses of the growth of different kinds of occupations and their physical requirements imply that the physical demands of work are now easier to meet than they were in the past. A much smaller proportion of jobs require strenuous physical effort; a larger percentage require only moderate or light physical exertion.[26]

Evidence about the relative influence of the other factors on retirement is less conclusive. Although most studies since the early 1980s suggest that most of the decline in the retirement age after 1970 was due to factors other than social security, we do not have an accounting of the contributions of the other factors. The most important of these may be a change in workers' preferences. Workers may now prefer to spend a larger part of their lives in retirement. If offered the same trade-off between retirement and lifetime consumption that was available to earlier generations, today's elderly might choose to spend less of their life at work and more of it in retirement. In making that choice, workers would willingly accept lower lifetime consumption of the goods and services their wages can buy. Unfortunately, there is as yet no reliable evidence to tell us whether this conjecture is true.

Even if the hypothesis were true, neither economists nor noneconomists have proposed a convincing explanation of why preferences should have changed. In principle, the introduction and liberalization of social security and private pensions may themselves have contributed to the change in preferences. By providing workers with a secure source of income apart from their wages and savings, pensions may have spurred a handful of alert, farsighted workers to consider the possibility of retirement. Once these pioneers decided to retire, less alert workers were offered a pattern of behavior to emulate. The economically "rational" behavior of a handful of workers may have provoked a slow-motion revolution in the preferences of the wider population. Retirement became an acceptable and even prized part of the typical worker's life. This impact of social security and pensions, if it exists, cannot be measured in the type of statistical study described in the previous section. Those studies attempt to find a correlation between workers' choice of retirement age and the financial incentives that might

26. See Manton and Stollard (1994); and Baily (1987).

motivate them to retire at one age rather than another. If the workers' preferences for retirement versus consumption have gradually changed as a result of a social-security-induced shift, this change would be missed in the statistical study.

Some evidence suggests that behavior is slow to change in the face of changing financial incentives. It took a number of years after the introduction of early social security pensions before retirement at age 62 became common. The financial incentives that eventually induced men to retire at age 62 were presumably present from the first day early pensions became available in 1961. It was only in the late 1960s, however, that a sharp fall-off in labor force participation at age 62 became noticeable. By the late 1970s, the fall-off in participation was greater at age 62 than at age 65. This evidence may suggest either that workers are slow to recognize the financial implications of a complicated innovation or that their preferences are formed, in part, by important changes in the financial incentives facing them.

Without a full accounting of the effect of changing worker preferences or a convincing explanation for the shift in preferences, economists' understanding of the retirement process will remain incomplete. If they wish to propose far-reaching policy reforms, their understanding remains dangerously incomplete.

References

Ando, Albert, and Franco Modigliani. 1963. "The 'Life Cycle' Hypothesis of Saving: Aggregate Implications and Tests." *American Economic Review* 53 (March): 55–84.

Baily, Martin N. 1987. "Aging and the Ability to Work: Policy Issues and Recent Trends." In *Work, Health, and Income among the Elderly*, edited by Gary Burtless, 59–96. Brookings.

Boskin, Michael J. 1977. "Social Security and Retirement Decisions." *Economic Inquiry* 15 (January): 1–25.

Bureau of the Census. 1975. *Historical Statistics of the United States: Colonial Times to 1970*. Department of Commerce.

Burkhauser, Richard V. 1979. "The Pension Acceptance Decision of Older Workers." *Journal of Human Resources* 14 (Winter): 63–75.

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

Burtless, Gary, and Robert A. Moffitt. 1985. "The Joint Choice of Retirement Age and Postretirement Hours of Work." *Journal of Labor Economics* 3 (April): 209–36.

Caroll, Christopher R., and Lawrence H. Summers. 1991. "Consumption and Growth: Some New Evidence." In *National Saving and Economic Performance,*

edited by Douglas B. Bernheim and John B. Shoven, 305–43. University of Chicago Press.

Friedman, Milton. 1957. *A Theory of the Consumption Function*. Princeton University Press.

Gustman, Alan A., and Thomas L. Steinmeier. 1986. "A Structural Retirement Model." *Econometrica* 54 (May): 555–84.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1994. "Expanding the Life-Cycle Model: Precautionary Saving and Public Policy." *American Economic Review, Papers and Proceedings* 84 (2): 174–79.

Krueger, Alan B., and Jörn-Steffen Pischke. 1991. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation." Working Paper 3699. Cambridge, Mass.: National Bureau of Economic Research.

Leonesio, Michael V. 1993. "Social Security and Older Workers." In *As the Workforce Ages*, edited by Olivia Mitchell, 183–204. Ithaca, N.Y.: ILR Press.

Lumsdaine, Robin L. 1996. "Factors Affecting Labor Supply Decisions and Retirement Income." In *Assessing Knowledge of Retirement Behavior*, edited by Eric A. Hanushek and Nancy L. Maritato, 61–122. Washington: National Academy Press.

Lumsdaine, Robin L., James Stock, and David A. Wise. 1992. "Three Models of Retirement: Computational Complexity versus Predictive Validity." In *Topics in the Economics of Aging*, edited by David A. Wise, 19–57. University of Chicago Press.

Manton, Kenneth G., and Eric Stollard. 1994. "Medical Demography: Interaction of Disability Dynamics and Mortality." In *Demography of Aging*, edited by Linda G. Martin and Samuel H. Preston, 217–79. Washington: National Academy Press.

Modigliani, Franco, and Richard Brumberg. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data." In *Post Keynesian Economics*, edited by Kenneth K. Kurihara, 388–436. Rutgers University Press.

Munnell, Alicia H. 1977. *The Future of Social Security*. Brookings.

Quinn, Joseph F. 1977. "Microeconomic Determinants of Early Retirement: A Cross-sectional View of White Married Men." *Journal of Human Resources* 12 (Summer): 329–47.

———. 1991. "The Nature of Retirement: Survey and Econometric Evidence." In *Retirement and Public Policy*, edited by Alicia H. Munnell, 115–38. Dubuque, Iowa: Kendall-Hunt.

Quinn, Joseph F., and Richard V. Burkhauser. 1994. "Retirement and Labor Force Behavior of the Elderly." In *Demography of Aging*, edited by Linda G. Martin and Samuel H. Preston. Washington: National Academy Press.

Quinn, Joseph F., Richard V. Burkhauser, and Daniel A. Myers. 1990. *Passing the Torch: The Influence of Economic Incentives on Work and Retirement*. Kalamazoo: Upjohn.

Ransom, Roger L., Richard Sutch, and Samuel H. Williamson. 1991. "Retirement: Past and Present." In *Retirement and Public Policy*, edited by Alicia H. Munnell, 23–50. Dubuque, Iowa: Kendall-Hunt.

Rust, John. 1989. "A Dynamic Programming Model of Retirement Behavior." In *The Economics of Aging*, edited by David A. Wise, 359–98. University of Chicago Press.

———. 1990. "Behavior of Male Workers at the End of the Life Cycle: An Empirical Analysis of States and Controls." In *Issues in the Economics of Aging*, edited by David A. Wise, 317–79. University of Chicago Press.

Skinner, Jonathan. 1988. "Risky Income, Life Cycle Consumption, and Precautionary Savings." *Journal of Monetary Economics* 22 (2): 237–55
Social Security Administration. 1998. *Income of the Aged Chartbook, 1996.*
Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Policical Economy* 97 (2): 305–46.