

Social Norms, Rules of Thumb, and Retirement: Evidence for Rationality in Retirement Planning

by

Gary Burtless
THE BROOKINGS INSTITUTION

October 2004

This paper was prepared for a conference on “Social Structures, Aging, and Self-Regulation in the Elderly,” at the Penn State Gerontology Center, State College, Pennsylvania, October 4-5, 2004. I am grateful to Lisa Bell for excellent research assistance and to Ellen Peters of Decision Research and Joseph Quinn of Boston College for providing unusually helpful and illuminating comments on an earlier draft of the paper. The views are solely those of the author and should not be attributed to the Gerontology Center or the Brookings Institution.

Social Norms, Rules of Thumb, and Retirement: Evidence for Rationality in Retirement Planning

WORKERS WHO ANTICIPATE reaching old age must make three choices about their retirement: The age when they will retire; the percentage of their wages to set aside so they can live comfortably when earnings cease; and the allocation of their retirement savings across different kinds of investments, such as stocks, bonds, bank and insurance accounts, and real estate. The three decisions are interrelated. Workers who do not expect to retire until shortly before they die do not need to save much for retirement. Those who anticipate retiring in their early 50s should plan to save a sizeable fraction of their pay. People who invest in risky assets, like biotechnology stocks and swamp real estate, may obtain terrific rewards for accepting great risk. If they are lucky they can use some of the rewards to retire young and live sumptuously. Less fortunate investors may be left with too little savings to retire.

Over the past couple of decades economists have devoted increasing effort to understanding retirement decisions. Labor economists focus on the timing of retirement and on retirees' work patterns after they leave career jobs. Microeconomists have examined workers' pre-retirement saving and post-retirement consumption behavior. Finance economists have theorized about workers' choices regarding the allocation of retirement saving across risky and less risky investments.

In thinking about decision making, economists' usually assume agents are well informed, far sighted, and rational. By "rational" I mean that agents use sound logic when deciding on a course of action. They combine all the information at their disposal, including knowledge of their own preferences and long-term interests, to make logical choices that maximize their well-being given the constraints they face. This does not mean people never make choices they later regret. Stock prices might fall 50 percent and house prices rise 50 percent after a worker has placed all his retirement savings in the stock market. When these outcomes become known, most people would regret investing in stocks rather than real estate. At the time the worker decides to invest in stocks, however, his choice is assumed to be based on good information and a prudent evaluation of the potential risks and rewards of alternative investments. Economists' presumption that agents are "rational" simply means they think people make decisions calculated

to produce the greatest long-term satisfaction given the trade-offs and uncertainties which an actor faces when the decision is made.

Many observers, including some economists, are skeptical that workers make retirement decisions in the far sighted and logical manner just described. Unlike other economic choices, which are repeated many times over the course of adulthood, the decision of when to retire is made only once. Workers are not given the opportunity to improve on their decision making through constant repetition, as is the case when consumers learn how to budget and shop for groceries, clothing, apartments, and even marriage partners. It is therefore hard to argue that workers can eventually learn from experience about choosing an advantageous retirement age or an optimal rate of saving. When people decide when to retire or how much to save for retirement, their choices may be poorly informed, short sighted, and less than rational.

What alternative decision making rules might workers follow when making these choices? One possibility is that workers economize on independent decision making by imitating the decisions of other workers. If a worker's colleagues or friends retire around age 60, age 60 may seem like a desirable age to leave a job. Another possibility is that workers use simple rules of thumb, which may not be far sighted or fully rational. Instead of considering the financial implications of retiring at every future age, some workers may retire as soon as available pension income replaces a target percentage of earnings, say, 75 percent of pre-retirement pay. Simple rules of thumb may produce decisions that are correlated with optimal decisions. But rules of thumb can also produce a retirement that occurs much too early or much too late compared with the age selected under a far-sighted plan. For example, a 75-percent income replacement rate may support a comfortable retirement for someone who is in excellent health, which is likely to be the case in the first years after a worker retires. But the same replacement rate may be grossly inadequate to pay for necessities and medical bills if the retiree's health deteriorates. A far-sighted planner would take account of this risk by delaying retirement or building up a bigger nest egg before leaving a career job.

Does it matter which of these decision-making models is more accurate in describing workers' retirement choices? It matters if we are concerned about the well-being of the aged. This is particularly true if we are considering policies that place greater responsibility on individual workers to save and invest for their own retirement. At the moment, all industrial countries have state sponsored pension and health insurance programs to support consumption in

old age. Workers are automatically enrolled in these programs when they take a job. Industrial countries also face the prospect of rapidly aging populations, a development which is likely to lead to the insolvency of state pension schemes unless pension contributions are raised and benefits cut. One way to reduce the long-term burden of state pensions is gradually to replace them with individual investment accounts. Under this kind of system, workers might have to decide how much savings to place in their accounts, how to allocate their savings across different investment options, and when and how fast to make withdrawals from their accounts. After they retire, workers would depend on withdrawals from their accounts to pay for consumption in old age. Some countries, including Japan and Great Britain, have moved in this direction, and policymakers in other countries, including the United States, urge reforms along the same lines.

Policies that rely on workers to make their own decisions about retirement saving and investment seem reasonable if most workers make these choices rationally and competently. The same policies look less appealing when a large fraction of workers base their retirement and saving choices on herd behavior, faulty logic, or defective information. If workers assume more responsibility to save for their own retirement, allocate their pension contributions, and withdraw funds in retirement, policymakers should be confident that few workers will make big planning mistakes. A big mistake can lead to serious hardship if state pensions are small. By the time an aged worker discovers he has retired too early or saved too little, he may have no opportunity to undo his mistake by saving more or returning to work.

In the remainder of this essay I consider what economists have learned about the rationality and far-sightedness of workers' retirement decisions. Do workers retire at a reasonable age? Do they save enough to afford the retirement ages they choose? Do they invest their retirement savings in a rational and prudent way?

The literature on these topics is lengthy and growing. Not surprisingly, it is open to competing interpretations. When polled about their preparations for retirement, large minorities of Americans acknowledge they have given no thought to the subject, have saved little or nothing in pension and other retirement accounts, and lack confidence they will be able to afford retirement. A number of economists and financial planners have published alarming studies suggesting that middle-aged and older workers face large saving shortfalls compared with the nest eggs needed to retire at the typical retirement age. On the other hand, careful analysis of the best survey evidence on workers' earnings, pension accumulations, Social Security entitlements,

and non-pension saving reveals that very few workers have nest eggs that obviously and substantially fall short of what is “optimal” under some conceivably rational plan. To be sure, many workers have little or no retirement saving, as already noted. However, most low-saving workers can fall back on Social Security or public assistance to support minimal consumption in old age. While the annual consumption reported by newly retired workers is lower than workers’ consumption before they retired, the decline is relatively small, is anticipated by many people approaching retirement, and is not clearly associated with a drop in retirees’ well-being.

Recent research findings on worker savings and retiree consumption do not *prove* that retirement and saving decisions are made in a fully rational and far-sighted way. The evidence only shows it is hard to rule out rationality and far-sightedness using available information on households’ consumption and savings. One problem is that we do not directly observe the underlying preferences of individual workers. This makes it nearly impossible to rule out rational decision making, even when we observe very odd patterns of work, saving, and consumption. When a worker suffers a 50-percent drop in consumption upon retirement we might interpret this as evidence of poor (irrational) planning. Alternatively, it may reflect the unfortunate effect of an unexpected early exit from the labor force, possibly because of a factory shutdown or onset of serious disease. Even if the possibility of a plant closing or poor health were fully reflected in the worker’s saving plan, she may have rationally intended to accept a big cut in consumption if her career came to a premature end. This means that we cannot rule out rational foresight when workers enter retirement with little savings or when they experience big drops in consumption over the course of their retirement. Bad outcomes may have been fully anticipated – and accepted – in a far-sighted and plausible retirement plan.

In the next section of the paper I consider what is known about the determinants of retirement and the role of far-sighted planning in the choice of retirement age. What evidence have economists found that the age of retirement is the result of a deliberate and rational plan involving lifetime optimization? The following section contains a review of saving behavior. Do workers save enough to retire comfortably? Or is there substantial under-saving compared with what is needed in light of the distribution of retirement ages? I will not treat the rationality of workers’ observed investment choices. Since there is deep division among economists and financial planners on what constitutes a sensible investment strategy, the soundness of most

workers' portfolios is not easy to evaluate. The paper will conclude with a brief summary of implications.

I. Evidence on Age at Retirement

When a worker retires, he withdraws from his normal occupation and reduces his work effort or stops work altogether. At the beginning of the last century retirement was relatively rare but not unknown. Two out of three American men past age 65 were employed, but one-third were not.¹ By the middle of the century retirement was much more common. Fewer than half of men 65 and older held a job in 1950. In 2000 the proportion at work had fallen still further. Just 18 percent of men over 65 were employed or actively seeking a job. Eighty-two percent were outside the active labor force. The proportion of women past 65 who were employed also fell during the twentieth 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 most European countries employment rates among the aged are now significantly below those in the United States (Quinn and Burkhauser, 1994; Burtless, 2004). Along with a shrinking work week and rising labor force participation among women, earlier retirement has been a distinctive feature of economic development in all the rich countries.

Theory. When economists think about retirement they naturally focus on the financial aspects of the decision. In this section I consider the basic theory economists have used to explain retirement choice and the evidence they have analyzed to test the theory.² Although the economic literature on retirement-age choice did not begin in earnest until the mid-1970s, it has grown rapidly since that time.³ The first aspect of retirement to attract economists' interest was the effect – or hypothesized effect – of retirement on individual and national saving. The classic statement of this relationship is contained in a series of articles written or co-authored by Nobel-prize-winning economist Franco Modigliani.⁴ His theory has had a wide influence on economists' thinking about the timing of retirement as well as the determination of saving.

Modigliani's basic hypothesis was that far-sighted workers will rationally plan their consumption over a full lifetime. In devising their lifetime consumption plans, they take account of the likely path of their labor earnings as they age and then 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 anticipated labor earnings and other kinds of income that are not derived from initial assets or labor earnings. Rational and far-sighted workers will plan to avoid situations in which all of their lifetime wealth has been consumed long before they expect to die. In the absence of public aid or private charity, the effects of such a planning error are not very appealing.

When a worker first enters the labor force his earnings are low, but wages typically rise as the worker gains extra skill and experience. In the United States earnings often reach a peak around age 50 and then begin gradually to decline. Since 1985 the average age at retirement for U.S. men and women has been about 62. One half of the people who work in their early 50s have withdrawn from the labor force by the time they reach age 63. When most workers retire their earnings cease altogether, although about one-fifth of retirees continue to earn reduced wages, at least temporarily. In simple versions of the life-cycle consumption model, a worker is well informed about the path of his future earnings, his age at death, and the interest rate he is able to earn on his savings. If the worker has stable preferences throughout his life, his planning problem is formidable but tractable.

A worker who successfully solves the consumption planning problem will plot out a desired path of consumption for each future year of life, and he will stick with the plan unless there is an unanticipated change in his financial outlook. The most advantageous 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 immediately consume nine-tenths of a box of cookies, 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 (the situation assumed in the top panel of Figure 1). If instead his 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 – who are very patient in their consumption – will shift consumption to

later stages of their life, and will plan to increase their 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 they plan to leave to heirs.

Figure 1 shows one solution to the life-cycle consumption problem for a worker who expects to live to age 85, faces a 5 percent interest rate, chooses to retire at age 65, and has no intention of leaving anything to an heir. The worker's potential wage is \$10,000 a year when she enters the workforce at age 20, gradually rises to a little more than \$22,000 a year when she reaches age 50, and then declines at later ages. The value of her expected lifetime earnings is a little more than \$850,000, but her total lifetime consumption can be greater than this if she saves part of her earnings and invests her savings in a 5-percent-a-year interest-bearing account. If she consumes her entire lifetime income (including anticipated interest earnings) steadily throughout her life, her annual consumption will be greater than her annual earnings early in her career. Her consumption will be less than her earnings in the middle of her career when her wages reach their peak. After she retires her earnings will be zero, and she will live off accumulated wealth. The lower panel in Figure 1 shows the worker's wealth holdings at successive stages of her career. Note that her nest egg is negative when she is young, because she must borrow money in order to consume more than she earns. (Young workers who cannot borrow are limited to consuming what they earn, but most consumption surveys suggest that Americans in their twenties and early thirties consume more than they earn and accumulate substantial debt.) By age 30 the worker's wages exceed her annual consumption, and she can begin to repay some of her debts. At age 45 her debts are paid off and she begins to accumulate wealth. Her wealth holdings reach a lifetime peak at the point of her retirement, and then decline as her nest egg is depleted to pay for old-age consumption.

Modigliani's life-cycle consumption model emphasizes the single most important financial aspect of retirement, namely, the sharp drop or complete cessation of labor earnings when work hours decline. 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 their consumption. Modigliani stressed personal saving as an alternative source of support in old age. Even though other income sources are nowadays quite important, 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. A crucial implication of the life-cycle theory is that far-sighted workers

will simultaneously select both a retirement age and a pattern of lifetime consumption. Their choice will be decisively affected by the expected pattern of their wage income, the interest rate they pay on money they borrow, and the interest earnings they obtain on money they save. Another implication of the theory is that year-to-year changes in consumption should be much smaller than year-to-year changes in earnings, especially around the planned age of retirement. Workers with far-sighted plans will smooth consumption using saving and dissaving over the course of their careers.

The characteristic pattern of increasing and then declining asset holdings over the life-cycle, displayed in the bottom panel of Figure 1, is a central empirical prediction of life-cycle consumption theory. The asset build-up 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 to smooth out consumption in comparison with earnings. Several of the assumptions I have mentioned are highly stylized and are not likely to be true in real life. For example, workers cannot borrow money at the same interest rate they obtain on their investments. Usually they must pay a higher interest rate for loans than the one they can safely earn on their investments. More important, few workers can borrow huge sums of money to finance current consumption. They are constrained in the amount they can borrow. These are comparatively minor issues for many workers, however, and addressing them does not really change most of the implications of the theory.

A more serious problem arises when we introduce a realistic picture of the worker's uncertainty about the future. In formulating an ideal retirement and consumption plan, it obviously helps if the worker is completely confident about her future earnings, her age at death, and the future interest rate. Unfortunately, few people can predict these things very accurately. In formulating their consumption plans, workers must take account of the possibility that the future may turn out to be more or less congenial than they anticipate. When they experience an unexpected change in earnings, workers face the problem of deciding whether the change will be long-lasting or only temporary. 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 report 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 higher 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 a job earlier? Or will it shorten retirement because the worker can expect an earlier reunion with his deceased relatives? The two outcomes, if fully anticipated, would have opposite effects on the rate of consumption over the remainder of life, but a far-sighted 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 with much confidence.

Setting aside for a moment the effects of uncertainty, the life-cycle model can be used to analyze workers' choice of a retirement age. To simplify the analysis, assume that workers will stop working completely when they retire. If a worker's potential wages at each future age are known with reasonable certainty, her planning problem is to select the most satisfying combination of years at work and lifetime consumption that is available to her. Economists usually assume that, other things equal, workers would prefer to work fewer years (holding constant their lifetime consumption) and to consume more goods and services (holding constant their years at work). In other words, additional consumption of goods and services is a "good" and an additional year at work is a "bad." If workers postpone their retirement (accepting more of a "bad"), they can also consume more over their lifetime (a "good").

Retirement incentives. One reason workers retire is that their potential earnings decline in old age, so the payoff from accepting a longer work life grows smaller with advancing age. When the payoff falls below the perceived value of the extra goods and services a worker can consume as a result of working longer, she will retire. Another reason work may appear less attractive at older ages is that employer and government pension plans can reduce the financial payoff from extra work. Social Security and company pensions affect the lifetime trade-off between consumption and retirement in a complicated way. Consider the effects of Social Security. The impact of Social Security on retirement depends on the contributions workers must make for the pension and on the benefit formula that links monthly pensions to a worker's

past covered earnings. U.S. employers and workers currently pay a combined tax equal to 12.4 percent of wages into the system. The tax thus reduces workers' wages by about 12 percent in comparison with the wages they would receive if the program were abolished. On the other hand, contributions allow a worker to earn credits toward a bigger Social Security pension. The monthly pension 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 her 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 and long-service workers typically receive much lower returns.

Workers who become eligible to receive Social Security benefits are entitled to receive a pension starting at age 62 or when they retire, whichever occurs later.⁵ Because the Social Security system has historically been quite 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 raised the lifetime wealth of older workers who became entitled to pensions under the system. If they consumed all of the Social Security benefits paid to them, they enjoyed higher lifetime consumption than their labor income alone could have financed. The fortunate generations that received this windfall may have retired earlier than they would have if Social Security had not been introduced or if the program had offered less generous pensions.

On the margin, Social Security can have another effect on the payoff from extra work. 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 when the worker turns 62. If a worker is entitled to a \$800-per-month pension, for example, she gives up \$800 in retirement income every month she delays retirement past age 62. If her regular monthly pay is \$10,000, this represents a comparatively small sacrifice. If her usual pay is \$1,000 a month, the sacrifice amounts to 80 percent of her wage. If she can only earn \$800 a month, the sacrifice is equal to her *entire* monthly pay. The sacrifice is so large the worker would be an idiot to continue working past age 62, unless some adjustment is made in her future pension to compensate her for the sacrifice.

Between the ages of 62 and 64 the Social Security formula offers workers a fair compensation for giving up a year's pension. Monthly benefits are adjusted upwards about 8

percent for each year's delay in claiming a pension. For workers who have average life expectancy and a moderate rate of time preference, this adjustment is large enough so that the sacrifice of a year's benefits is compensated by eligibility for a bigger pension in the future. After age 65, however, the benefit formula was historically much less generous toward delayed retirement. Postponement of retirement after that age was not fairly compensated by increases in the monthly pension.⁶ In effect, the historical benefit formula required workers to give up part of the accumulated value of their lifetime pensions if they delayed retirement after age 65. It should not be surprising under these circumstances if a sizeable fraction of workers stopped working at 65 and began collecting Social Security pensions.

It is worth noting that very few workers are exactly "average." A benefit calculation rule that is age-neutral (or "actuarially fair") on average can still provide strong financial incentives to retire for a worker who has below-average life expectancy. The worker may not expect to live long enough for the future benefit increase to make up for the benefits he gives up by delaying retirement for one more year. Similarly, a worker who applies a high discount factor when evaluating future benefits may not be impressed that the pension adjustment is "fair" for an average worker. For workers who are impatient to consume, an 8-percent hike in benefits starting one year from today may not be enough to compensate for the loss of twelve monthly pension checks over the next year. Even an actuarially fair pension adjustment might be too small to persuade workers who are tired of their jobs to delay retirement.

The reason 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 \$11,640 a year lose \$1 in annual benefits for every \$2 in earnings they receive in excess of \$11,640.⁷ At one time the earnings limitation was much lower, the tax on excess earnings was much higher, and the age range in which the earnings test applied was much wider. Social Security pensioners were discouraged from working, and some workers may have been induced to postpone claiming a pension until they were confident their earnings would remain low.

Many company pension plans are structured similarly to Social Security pensions. Workers who are covered under an old-fashioned defined-benefit plan earn pension credits for as long as they work for the employer that sponsors the plan (sometimes up to a maximum number of years). The longer they work, the higher their monthly pension. Most defined-benefit plans

are structured to encourage workers to remain with the employer for a minimal period – say, 10 years – or until a critical age – say, age 55. Workers who stay for shorter periods may receive very little under the plan. On the other hand, workers who stay in the job too long may see the value of their pension accumulation shrink. This would happen if the plan offered benefits to workers starting at age 55 but then failed to significantly increase the monthly benefit for workers who delayed retirement after age 55. If a 55-year-old worker can collect a monthly pension of \$2,000 when she retires immediately and a monthly check of \$2,001 if she delays her retirement one year, she will clearly lose a substantial amount of lifetime benefits – nearly \$24,000 – for each year she postpones retirement. The worker essentially suffers a pay cut when she reaches age 55, and the cut is equal to the loss in lifetime benefits she suffers by postponing retirement. Many employers find this kind of pension formula to be an effective inducement to push workers into early retirement.

Evidence. 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 American households headed by someone 65 or older. Tabulations of the Current Population Survey show that Social Security accounts for slightly more than 40 percent of the total cash income of the aged. For almost three-fifths of aged families, Social Security represents half or more of the family's total cash income (Employee Benefit Research Institute, 2002, Figure 1; and Grad, 2002, Table 6.A2). Until 1941, Social Security provided no income at all to the aged. Today the program replaces about 40 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 non-working dependent spouse, the benefit replaces approximately 60 percent of the worker's final wage. Benefits are clearly high enough so they can be economically significant in influencing the choice of retirement age.

As noted above, labor force participation rates at older ages fell substantially over the twentieth century. What role did Social Security incentives play in this trend? The distributions of male retirement ages in selected years between 1940 and 2000 are plotted in Figure 2. The figure shows the percentage of men leaving the labor force at each age from 55 to 72, computed as a fraction of men in the labor force at age 54.⁸ The calculations are based on male labor force participation rates for successive years of age in each of the indicated years. Not surprisingly,

the retirement distributions for more recent years are skewed toward the left, reflecting the fact that men have withdrawn from the workforce at younger and younger ages. The tabulations in all of the years show evidence of clustering in retirement at particular ages. There are peaks in the 1940 distribution at ages 60, 65, and 70, indicating that retirement at those ages was more common than at other ages. By 1960, however, there is only one main peak in the retirement distribution, at age 65. In 1970 there is evidence of a secondary peak in the distribution, at age 62. By 1980, the percentage of retirements occurring at age 62 was almost as high as the percentage at age 65. In 1990 and 2000, retirement at age 62 was much more common than retirement at 65.

The earlier discussion of Social Security incentives suggests an explanation for the clustering of retirements at ages 62 and 65. Between 1941 and 2000, workers eligible for Social Security who continued to work beyond age 65 gave up pensions for which they were not fairly compensated. The earnings penalty in the benefit formula encouraged workers to retire at age 65. The clustering of retirements at age 62, which began after 1960, is also easy to explain. Starting in 1961, age 62 became the earliest age at which men could claim a reduced Social Security pension. Before 1961 men could not claim a pension until 65, and there was no evidence of clustering in retirements at age 62. By 1970 retirement was more common at 62 than at any other age except 65. By 1990, age 62 was by a wide margin the most popular age of retirement. In principle, the Social Security formula fairly compensates workers if they delay claiming a pension past age 62. As we have seen, however, a worker with a high rate of time preference or short life expectancy might not regard the compensation as fair. In that case, some workers will prefer retiring at age 62 rather than a later age.

The retirement age distributions displayed in Figure 2 are based on a crude approximation of workers' behavior in each of the indicated years. If the labor force participation rate of 60-year-old men is 5 percentage points lower than the participation rate at age 59, and if 90 percent of 54-year-old men are in the labor force in the same year, the calculation assumes that the retirement rate at age 60 is 5.5 percent $[(5 \div 90) \times 100]$. A more refined estimate of workers' retirement ages can be obtained by interviewing the same people several times as they approach the end of their careers. The U.S. government has mounted two such panel surveys, the Longitudinal Retirement History Survey (LRHS), conducted between 1969 and 1979, and the Health and Retirement Survey (HRS), which began in 1992.

The LRHS was a 10-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. Figure 3 displays information on the retirement behavior of men in the LRHS sample who had no disabilities. The top panel shows the distribution of retirement ages among men who were observed to retire by the end of the survey in 1979, when respondents were between 68 and 73 years old.⁹ To determine the exact retirement age, I 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 that end with the worker's return to a full-time job. However, the definition includes movements from steady full-time work into part-time jobs. The picture misses the retirements of some men who did not retire before their last completed interviews, and this omission will lead to some under-representation 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 the top panel of Figure 3 is broadly similar to that shown for 1970 in Figure 2. As in Figure 2 there is a clustering of retirements at ages 62 and 65, with a much higher peak at the latter age.

The lower panel displays the pattern of earnings among retired but working men who are 62 years old or older in the first LRHS interview after they retired. 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 16 hours a week. The lower panel in Figure 3 shows the distribution of their earnings in relation to the earnings exempt amount in the Social 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.¹⁰

Casual observation of the top and bottom panels of Figure 3 suggests Social Security had a powerful effect on both retirement ages and post-retirement earnings. Note that the age distribution of retirements had two peaks, a lower one at age 62, when Social Security 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 post-retirement 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. Over a quarter of working retirees earn within 10 percent of the exempt amount, and over half earn within 30 percent of it. While retirees may under-report their true earnings to Social Security to avoid paying the high implicit tax, the earnings estimates displayed in Figure 3 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 post-retirement hours of work.

The evidence in Figures 2 and 3 strongly suggests that some fraction of men are quite sensitive to Social Security incentives when they retire. It is less obvious whether this shows most of them are choosing their retirement age on the basis of a far-sighted and rational plan. As noted earlier, workers following a simple rule of thumb may retire as soon as available retirement income replaces a target percentage of their monthly pay. It does not require long-term planning to recognize this target is more likely to be met when the worker can first claim a Social Security pension. To be sure, the shifts in the peak of male retirement ages shown in Figure 2 conform broadly with the shifting incentives provided by Social Security. Still, it seems surprising that men were so slow to respond to the availability of early pensions, which began in 1961. The percentage of men retiring at age 62 approximately doubled between 1970 and 1980, yet it is hard to see how the incentives for retirement at that particular age changed appreciably over the decade. The innovation in pension rules occurred in 1961 when early retirement benefits were first made available.

Robert Axtell and Joshua Epstein (1999) argue that the slow evolution of retirement ages after the 1961 rule change actually provides powerful evidence *against* the view that workers are fully rational in their choice of retirement age. They suggest instead that "... imitative behavior and social interactions – factors absent from traditional economic models – may be fundamental in explaining the sluggish response to policy" (Axtell and Epstein, 1999, p. 162). They argue that only a small percentage of workers may have the capacity or willingness to understand program rules and interpret their meaning for the choice of an optimal retirement age. Most workers imitate the behavior of their "neighbors," that is, older relatives, colleagues at work, or actual neighbors whose retirement behavior can be directly observed. If an imitator's neighborhood happens to include one or more far-sighted planners, it is more likely the imitator will respond to new incentives in a far-sighted way, because the behavior imitated is more likely to be optimal.¹¹

Axtell and Epstein show how rational behavior can cascade through a social network, even though very few members of the network may be far-sighted or fully rational in their decision making. Eventually, retirement patterns attain a new equilibrium in which the rational behavior predominates.

There is also some question whether the pattern of retirement ages and post-retirement work effort reflects a sensible response to Social Security incentives. Many people retire shortly after their 62nd birthdays, apparently because they can immediately claim a Social Security pension. As noted earlier, however, the Social Security benefit formula compensates workers who delay claiming a pension after 62 by increasing their monthly pension in later years. For workers who have average or above-average life expectancy and who have a savings account that earns less than 5 percent a year, it should make sense to delay claiming Social Security for two or three years after age 62. The rate of return that these workers can obtain through delaying a benefit claim compares favorably to the return they obtain on their savings. For workers in these circumstances there is no more reason to stop working at age 62 than there is at age 61. Only workers who have no liquid savings, who have a short life expectancy, or who apply a high rate of discount when evaluating future income gains have any special reason to retire at 62. Of course, men in those circumstances could account for all the extra retirements observed at 62.

The post-retirement work pattern reflected in the lower panel of Figure 3 seems to reflect a powerful influence of the Social Security retirement earnings test on behavior. Pensioners avoid earning more than the Social Security exempt amount in order to avoid facing a 50-percent marginal tax rate on their earnings. On the other hand, the distribution of post-retirement pay may also reflect a deep misunderstanding of the earnings test. Under the rules of Social Security, workers whose benefits are penalized because of application of the earnings test eventually have their monthly pensions recalculated to reflect this benefit reduction. Suppose a worker earns enough wages above the earnings exempt amount to lose three months of benefit payments when he is 64 years old. According to the rules of the program, his basic monthly pension at age 65 and later years is supposed to be increased to reflect the fact that he failed to receive benefits for three months between ages 62 and 64. The adjustment is exactly the same as the one he would have received if he had postponed claiming benefits for three months. Since this delay in claiming pensions is compensated with an actuarially fair adjustment in monthly pensions, in effect the worker does not lose any lifetime benefits at all when his benefits are reduced because

of application of the retirement earnings test. For pensioners who are between 62 years old and the normal retirement age, the earnings test results in the rearrangement of benefits over time. The worker receives smaller monthly benefits at the time he earns more than the annual exempt amount, but he receives permanently higher monthly benefits starting at a later age. In theory, the later benefit adjustment fairly compensates workers for the temporary reduction in benefits. If workers fully understood these rules, it is a little hard to understand why their post-retirement earnings are so sensitive to the “tax” on earnings above the exempt amount and to changes in the annual exempt amount (Vroman, 1985).

One plausible theory is that workers misunderstand the program rules. Many interpret the rules to mean they face a simple benefit cut whenever their earnings exceed the exempt amount. The clustering of annual earnings around the exempt amount certainly seems consistent with this interpretation. It shows pensioners are responsive to Social Security incentives, but it does not show whether workers are knowledgeable about the true financial implications of the program rules. If they are knowledgeable about the program rules, the clustering of post-retirement earnings at the exempt amount may ironically provide evidence that workers are short-sighted in their response to the earnings test.

Another way to analyze the impact of Social Security incentives is to examine the retirement age 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. Burtless (1986) examined the first episode, and Krueger and Pishke (1992) examined the second using the life-cycle framework. In the period analyzed by Burtless, workers born in earlier years planned their retirements when Social Security was comparatively less generous; workers born in later 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, and the effects went in the direction predicted by the far-sighted planning model. Burtless estimated, for example, that the 20-percent benefit hike between 1969 and 1993 caused a 2-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 about 2 percentage points. The effects of the 1977 amendments found by Krueger and Pischke were smaller, but probably went in the direction predicted by theory. Of course, even if workers are using short-sighted rules of thumb rather than far-sighted optimal planning to choose their retirement age, many will respond to the changed incentives in the same way these economists found. If monthly pensions are hiked 20 percent, as occurred between 1969 and 1972, even short-sighted workers may claim benefits a couple of months earlier. Moreover, the behavioral responses uncovered by Burtless and Krueger and Pischke may have been confined to a relatively small number of retirees. The great majority of workers may have failed to respond at all to the changed incentives.

Some economists have directly posed the question of whether the retirement behavior they observe is guided by simple-minded or far-sighted planning. Robin Lumsdaine, James Stock, and David A. Wise (1992) believe they found evidence suggesting that retirement-age choice is often the result of a sophisticated decision-making rule. They examined the retirement choices of workers in a handful of company pension plans. They assessed these choices using three different decision-making rules, one of which was based on the application of a simple rule while the other two were based on more sophisticated decision-making approaches. (One of these used an “option value” technique for evaluating the value of pension offers, and the other used a “dynamic reprogramming” rule.) The economists estimated their models using information from one period, and then they tried to predict retirement patterns in a later period under each of the three models. Perhaps surprisingly, they found that the models based on more complex and far-sighted decision rules were more successful in predicting future retirement patterns. This evidence suggests at least some workers use information in a sophisticated way to decide when to retire. Of course, within a long-established company plan which covers many workers in the same workplace, information helpful in choosing an optimal retirement age that is discovered by one worker can easily be shared with co-workers. Where information sharing is more difficult, workers might rely on simpler decision-making rules, and some workers may end up retiring at an age that is less than optimal. Lumsdaine *et al.* (1992) analyzed the retirement

decision in isolation. They did not assume workers were making far-sighted and fully consistent plans for both work and consumption over a multi-year time horizon. Even the most sophisticated decision rule they consider is simpler than the planning methods needed to simultaneously select a retirement age and an optimal path for saving and consumption.

Other data are less supportive of the idea that workers use good information and far-sighted plans to select their retirement age. The availability of longitudinal surveys of older workers allows researchers to ask people whether they have made retirement plans and selected an expected age at retirement. Information from later interviews can be used to determine whether respondents follow through on their plans. Katherine Abraham and Susan Hausman (2004) analyzed information from the 1992-2000 HRS to determine how frequently older workers reported a retirement plan and how often they stuck to those plans. Workers in the HRS were in their 50s and early 60s when the question on retirement plans was first posed. This seems like a point in life when long-term planners would have formulated a retirement strategy. Abraham and Hausman report that "...the most common answer (38 percent of responses) was that the respondent had not given much thought to future work and retirement plans, or had no plans" (2004, p. 9). Among workers who reported an expected retirement age or retirement strategy, a large percentage failed to follow through on their plans, even when the planned retirement was within than two years of the time they described their plans. Among respondents reporting they would stop working altogether within two years of an interview, slightly more than one-third were still at work in the next biannual interview. Among workers who claimed they would *never* stop working, about one in seven had actually ceased working within two years. Of course, unexpected events may have intervened between the two surveys, disrupting the best-laid plans of rational workers.

If workers wish to formulate a rational retirement strategy, a minimum requirement is to become familiar with the rules and benefit formulas governing the pension plans in which they are enrolled. The HRS provides a good source of information about older workers' knowledge of their pension plans. Workers were asked to describe some important features of their company plans, and their descriptions were compared to the descriptions of the same plan supplied by employers. Because of the method used to collect and verify the data, we should expect that employers provided more accurate plan descriptions than their employees. Alan Gustman and Thomas Steinmeier (forthcoming) offer a sobering comparison of the pension

descriptions supplied by workers and employers in the HRS. Only about one-half of workers covered by a defined-contribution plan correctly identified the type of pension plan in which they were enrolled. Approximately the same percentage of workers enrolled in a defined-benefit plan correctly reported that their employer offered that form of pension.

A defined-benefit plan provides vested workers with a pension that is determined by the worker's years of service and final salary while enrolled in the plan. These plans ordinarily have an early entitlement age (when workers can first receive reduced benefits) and a full entitlement age (when workers can claim an unreduced benefit). The monthly pension is guaranteed by the employer. In contrast, a defined-contribution pension is essentially an individual investment account maintained in behalf of individual workers. The employer deposits annual contributions (usually a fixed percentage of a worker's pay), and the ultimate value of the investment account depends on the success of the worker's or employer's investment strategy. The worker bears the risk of poor investment outcomes, but the accumulation in the pension account is the property of the worker even if he leaves the employer long before the standard retirement age. If workers do not know whether they are enrolled in a defined-benefit or a defined-contribution pension plan, it is unlikely they are familiar with the retirement incentives in their plan. Indeed, Gustman and Steinmeier (forthcoming) show startling discrepancies between workers' understanding and employers' descriptions of retirement incentives. Even among workers who correctly stated they were covered by a defined-benefit plan, only a minority accurately reported the youngest age at which they could claim pensions. For example, among workers in plans where the early eligibility age was 55, only 40 percent of workers correctly reported this age. Slightly more than 20 percent believed the early entitlement age was 62, and 7 percent reported it was 65 or higher. Gustman and Steinmeier show that workers are more accurate in describing their Social Security entitlements, although workers with very low entitlements often have an exaggerated estimate of their potential monthly benefits

How does misinformation affect retirement decisions? Sewin Chan and Ann Huff Stevens (2003) offer some fascinating evidence. They focused their analysis on HRS respondents who worked in 1992 and were covered by a company pension plan according to the reports of their employers. Using the employer's description of the worker's pension entitlement, Chan and Stevens could reliably calculate the value of a pension if the worker retired immediately and compare that to the pension value if the worker retired at a future age. The average worker

underestimated the value of the pension by about 55 percent of the amount reported by employers. Moreover, many workers offered wildly inaccurate estimates of the improvement in their pension if they delayed retirement to a later age. Using information from follow-up HRS interviews, Chan and Stevens found that workers' retirement choices were based on their (possibly inaccurate) interpretation of pension rules. For people with accurate information, retirement choices were closely aligned with the financial incentives in their plan. If a worker's understanding of the plan rules was in error, the retirement decision was often based on serious misunderstanding.

One of the most important financial determinants of an optimal retirement age is the increase or decline in the value of a pension if a worker postpones retirement for one or more years. Under some retirement plans, workers can actually lose lifetime pension wealth if they delay their retirement after attaining the plan's early or normal retirement age. Chan and Stevens (2003) found that workers who accurately reported the amount of pension gain from delaying retirement were several times more responsive than average workers to the true financial incentives in their pension plan. Most economists, including Gustman and Steinmeier (1986), Burtless and Moffitt (1985), and Burtless (1986), estimate life-cycle retirement models under the presumption that workers are responding to the true financial incentives in the pension plans in which they are enrolled. On the whole, they find aggregate responses to pension incentives that seem consistent with the basic life-cycle model. Chan and Stevens' findings suggest the pattern of aggregate response may reflect far-sighted responses on the part of some well-informed workers and poorly informed or irrational choices for a sizeable minority or even a majority of workers.

II. Life-Cycle Saving

Workers' saving behavior can also be compared with the predictions of the life-cycle model to help us understand whether saving decisions are guided by rational and far-sighted planning. The model provides some clear predictions about how wealth is accumulated over a career and how workers should respond to unexpected events. The lower panel of Figure 1 implies that workers should build up significant savings in anticipation of retirement and then draw down their wealth during retirement. The model makes a clear and plausible distinction between (unanticipated) changes in flows of income that can be expected to last and changes that

are only temporary. An unexpected income improvement that is permanent, such as an earnings gain that accompanies a promotion, will have a much bigger 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 \$30,000 a year for thirty years will plan to make a much bigger change in his consumption than the person who wins a one-time prize of \$30,000. By the same reasoning, the lottery winner who obtains a prize paying a modest annual amount (say, \$2,000 a year) which has a present discounted value of \$30,000 will alter his consumption by roughly the same amount as the winner of a one-time prize equal to \$30,000. A fully anticipated drop in income, such as the one that accompanies a planned retirement, should have almost no effect on consumption.

Some evidence supports the theory.¹² Most empirical research suggests that the life-cycle 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. As the theory predicts, we 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 skeptical 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 workers who do have assets apparently continue to add to them after they retire. Neither fact is easy to reconcile with simple versions of the life-cycle model. Theorists have made modifications in the basic theory to account for obvious empirical contradictions. Different theorists have proposed different modifications to rescue the basic model. Whatever their criticisms of the model, however, few economists have strayed far from it in trying to explain the connection between saving and retirement behavior.

Pre-retirement saving. Recent empirical research has focused on two questions about saving behavior in relation to retirement. First, do workers typically accumulate enough savings

so that they can live comfortably during retirement? And second, is there evidence to support the prediction of the life-cycle model that consumption changes little when retirement occurs? The first question has aroused considerable controversy because of disagreement over what constitutes adequate saving for retirement. Almost from the beginning of systematic analysis of the wealth distribution, economists have had to confront the fact that many workers reach old age with very little savings (see Diamond and Hausman, 1984). This finding has been confirmed in many studies over the years. Annamaria Lusardi (2001) recently tabulated the wealth holdings of HRS respondents who were not retired at the time of their first interview in 1992. Since these workers were between 50 and 61 years old, it is reasonable to assume most of them were within a decade of retirement. Workers in the bottom one-tenth of the wealth distribution had no wealth at all except their Social Security wealth. Even workers at the 25th percentile had essentially no liquid wealth. That is, the sum of their bank deposits, stocks, and bonds and subtracting their short-term debt left these workers with no liquid savings. These workers' total wealth holdings, including equity in a home or business, IRAs and Keogh plan assets, and vehicles, amounted to less than \$28,000 (1992 dollars). Lusardi points out that if all of this wealth were sold and converted into a lifetime annuity, it would provide workers with an income of less than \$200 a month. One-quarter of 50-61 year-old workers in the HRS had even less wealth than this. One reason for low savings may be that workers have given little or no thought to retirement. Lusardi reports that the median wealth holdings of workers who have thought "hardly at all" about their retirement is less than one-half the median wealth of workers who have thought "some" or "a lot" about retirement.

Although the fraction of older workers who lack wealth may seem shockingly high, is it high enough to cause us to reject the hypothesis that workers save rationally for retirement? In some cases we can show that the optimal rate of pre-retirement saving is zero or very near zero. Many workers who earn low or erratic wages throughout their careers will qualify for a Social Security pension or an old-age public assistance check. The amount of monthly benefits may equal or exceed the average net pay they received in their career. Since some of these workers may not be eligible for public assistance unless their liquid savings are very low, it may make no sense to accumulate pre-retirement wealth. The availability of Social Security, public assistance, and company provided pensions means that the optimal amount of savings depends critically on individual circumstances. Workers who can expect pensions or assistance payments that replace

a large percentage of their net earnings have much less need for savings than workers who do not anticipate pensions or assistance payments.

Douglas Bernheim (1992 and 1995) published two widely cited studies showing that many baby boom workers, including high-wage workers, face large shortfalls in retirement saving. He calculated workers' optimal saving levels taking into account the number of their current and anticipated dependents, earnings, expected Social Security and pension benefits, and other factors. He then compared workers' actual saving with the optimal saving amount and determined whether workers faced a surplus or deficit in their saving. His calculations implied that workers in the baby boom generation were saving at just one-third the rate needed to cover the costs of their retirement.

Other researchers have reached similar conclusions, often using much better survey data. James Moore and Olivia Mitchell (2000) examined the 1992 wealth holdings of HRS respondents and calculated the additional saving they would need to retire without any loss of consumption at retirement. This calculation takes account of the Social Security and pension benefits that workers could obtain if they continued working. The calculations are repeated for two potential retirement ages, 62 and 65, the two ages that are most common (see Figure 2). Moore and Mitchell show that the median HRS household would have to increase its saving rate by 16 percent of earnings to maintain constant consumption after retirement at age 62. If retirement were delayed for three additional years to 65, the required extra savings for the median worker would represent 7 percent of earnings. When Moore and Mitchell compared required savings rates to actual savings rates among households approaching retirement, they found that actual savings rates typically fall far short of the required rate.¹³

Some recent analyses of wealth surveys have produced a very different picture of wealth adequacy. A number show that comparatively few workers have clearly under-saved, and the typical amount of under-saving is quite small. One reason for the difference is that the newer studies explicitly account for the income uncertainty workers face in the years before they retire. Earnings uncertainty is very important for an obvious reason. If workers cannot borrow much money, they must save a very large percentage of their earnings in high earnings years to ensure that their families do not have to reduce their consumption in low earnings years. This effect of earnings uncertainty is not reflected in Figure 1 because it is drawn under the assumption that annual earnings will rise and fall over a worker's career in a completely predictable way. In the

real world, earnings are much less predictable than this. Every year many workers lose their jobs, and some must accept big pay cuts in order to get re-employed. Other workers receive unexpected promotions or take new jobs with higher salaries. If workers want to accumulate enough savings to fully smooth consumption, they must save a very large percentage of their pay to accumulate a buffer stock of savings. Rational planners will save less than the full amount needed to completely smooth consumption, and this will mean that large, unexpected wage reductions will sometimes cause workers to deplete their savings before they retire. As noted by Eric Engen, William Gale, and Cori Uccello (1999), it is wrong to argue that there is a single optimal path of saving for all workers who expect to earn the same lifetime wages. Instead, they find even among workers who share the same preferences, there is a range of optimal saving paths, where each path depends on the exact sequence of earnings “surprises” received by the worker. Workers and retirees also face uncertainty about when they will die. If workers die at an unexpectedly early age, their savings will go unused and will not contribute much to their lifetime happiness. If they die in advanced old age, they may deplete all their savings and face many years of very low consumption. Rational workers will make a savings choice that balances these risks, but for many far-sighted workers the balance will mean their consumption falls as they live longer and longer beyond their retirement age.

In light of earnings and life-span uncertainty, Engen *et al.* (1999) ask a somewhat different question about wealth holdings from the one posed by earlier analysts. They ask whether the observed distribution of wealth holdings seems consistent with the distribution that would be observed if each household responded to unexpected earnings changes and life-span uncertainty in an optimal way. Earlier analysts implicitly posed a different question: If the profile of lifetime earnings and date of death were known in advance, how much wealth would an optimizing worker have set aside by the time he reached the age when his wealth holdings were reported to the interviewer? If a worker’s wealth falls short of this threshold, the worker is judged to have inadequate savings. Engen *et al.* (1999) do not actually observe the past sequence of earnings for any member of their sample, but they can use information from other sources to derive reasonable estimates of typical year-to-year variability in earnings. Combining this information with data about the worker’s current earnings and survival probabilities in future years, they simulate the range of wealth holdings that would be observed if workers responded optimally to a simulated sequence earnings fluctuations and the known probabilities of future

death. Their simulations reveal not surprisingly that many prudent and rational savers will have little or no savings if they experience a big, unpleasant earnings surprise. Although Engen *et al.* (1999) conclude that there is probably some under-saving in a few population groups, the shortfall in saving seems quite modest compared with earlier estimates. This conclusion was confirmed in a recent study by John Karl Scholz, Ananth Seshadri, and Surachai Khitatrakun (2004), which used HRS data to calculate optimal saving accumulations based on workers' *actual* lifetime sequence of Social-Security-covered earnings.¹⁴ The optimal accumulations were then compared with wealth holdings reported by the same workers. Scholz *et al.* conclude that fewer than one-fifth of HRS households have less saving than their optimal targets, and the saving shortfall of those households is typically quite small.

Post-retirement consumption. The findings by Engen *et al.* (1999) and Scholz *et al.* (2004) do not prove that the saving behavior of American workers is far-sighted and rational. They demonstrate instead that it is difficult to rule out the hypothesis that saving choices are far-sighted and rational for the overwhelming majority of workers. Some readers might find this conclusion more reassuring if analysts could offer clear evidence that retirees enjoy adequate income or consumption in old age. It is hard to define a reliable benchmark for assessing "adequacy," however. One gauge is the official U.S. poverty line. Census Bureau statistics on income poverty suggest destitution is nowadays no more common among the aged than it is among nonaged adults (see Figure 4). This reassuring fact may suggest that poor planning and short-sightedness do not lead to any more hardship in old age than at younger ages. Of course, if the income needs of the aged are greater than those of the nonelderly, perhaps because of costly health problems, it still might be the case that retirees have inadequate resources to maintain post-retirement consumption.

Economists have found reasonably reliable and consistent evidence suggesting that consumption falls after workers retire, although the implications of this decline are not always clear. Daniel Hamermesh (1985) found that couples' consumption early in retirement is 14 percent higher than their retirement income can support, forcing them to reduce their consumption in later old age. Jerry Hausman and Lynn Paquette (1987) uncovered more compelling evidence of a drop in consumption following retirement. Looking solely at food consumption among families represented in the LRHS, Hausman and Paquette found that retirement led to a decline in expenditures on food of about 14 percent of pre-retirement

consumption. For the workers who were forced to leave their jobs because of a layoff or deterioration in health, the drop in consumption was even bigger – an additional 9 percent of pre-retirement food consumption. For workers who had accumulated below average wealth, the drop in food expenditures was larger still. James Banks, Richard Blundell, and Sarah Tanner (1995) used many years of detailed household consumption data for British families to document the fall in consumption that occurs immediately after workers retire. Part of this decline can be explained by lower spending requirements for people who no longer need to go to work, but much of the falloff in consumption cannot be rationalized. Banks, Blundell, and Tanner conclude that for many households, retirement must have been accompanied by an unwelcome surprise that reduced the families' ability to consume.

Recent studies have shown that the drop in consumption following retirement is at least partly anticipated. Michael Hurd and Susann Rohwedder (2004) used interview responses in the HRS to compare respondents' pre-retirement expectations of consumption after retirement with the actual experiences of workers who have already retired. HRS respondents who were not yet retired were asked whether they expected consumption to fall after retirement and by how much. The responses of people who have not yet retired could be compared with the reported consumption changes of HRS respondents who have retired and already experienced the fall in income that accompanies retirement. Hurd and Rohwedder confirm that consumption falls at retirement. The average decline is about 15 percent to 20 percent of pre-retirement consumption. They also show, however, that the fall in consumption is largely anticipated by workers. In fact, the reported decline in consumption among workers who have already retired is a bit smaller than the average decline predicted by workers who have not yet retired. These findings confirm earlier findings by John Ameriks, Andrew Caplin, and John Leahy (2002). Using survey responses obtained from participants enrolled in the TIAA-CREF pension program, Ameriks, Caplin, and Leahy discovered that a majority of active workers expect their consumption spending to fall after retirement. Among TIAA-CREF participants who have already retired, however, only about one-third report that their consumption actually declined. Forty-four percent report their spending has remained unchanged, and 20 percent say consumption has risen. Thus, many older workers anticipate their spending will decline after retirement, but the actual experience of recent retirees suggests the drop in saving may be smaller than anticipated.

Even granting that consumption falls after retirement, the drop in spending may not reduce retirees' welfare. For example, retirees may spend less on food because they do not need to eat as many meals away from home or because they have more time to shop for bargains. Necessary spending on taxes, clothing, and transportation may also fall. Retirees have the time to produce some goods and services in the home that full-time workers typically purchase in the market. Thus, even if it is true that consumption expenditures fall in retirement, it is not clear the decline is associated with a drop in well-being. The evidence on the reported happiness (or "subjective well-being") of retirees versus older active workers provides little evidence that the retired are systematically less happy than active workers. Retirees on average have more health problems than active workers, but among people who have the same marital status and similar health problems, the retired are about as happy as active workers.¹⁵

In sum, there is overwhelming evidence that many households have very little savings as they approach retirement, but this does not help us decide whether workers' saving decisions were based on short-sighted or irrational decision making. Given the uncertainty of pre-retirement earnings and the availability of means-tested retirement benefits, many forward-looking, rational workers will enter retirement with little savings. There is also pervasive evidence that workers experience significant reductions in consumption after they retire, possibly indicating that they were short-sighted in their saving or unpleasantly surprised by the drop in income that followed retirement. Since many workers anticipate a fall in consumption after retirement, however, another explanation for the fall in consumption is that workers have lower spending needs when they leave work. The drop in consumption spending may not be connected with a decline in welfare.

III. Implications

Squinting their eyes and looking only at aggregate trends, economists find plentiful evidence of rational, far-sighted responses to financial incentives for retirement and saving. American retirement patterns at different ages have gradually evolved to reflect the retirement incentives embodied in the nation's most important pension program. Evidence from other industrialized countries and cross-national comparisons of retirement behavior reinforce the findings based on U.S. evidence. Jonathan Gruber and David Wise (1999) and economists from a number of countries recently examined pension systems and retirement behavior in eleven rich

countries. Some countries allow workers to begin drawing public pensions at age 60 or earlier, while others do not make old-age benefits available until much later. There is also wide variation in the treatment of wage earnings once workers reach the pensionable age. Some countries, like the United States, no longer penalize workers for delaying their retirement beyond the early or normal pensionable ages. Other countries, like France and Belgium, offer very generous pensions and impose heavy financial penalties on workers who remain employed after the pensionable age. Gruber and Wise and their collaborators find a strong correlation between national retirement patterns and the labor supply incentives that are embodied in national pension systems. Countries with modest pensions and generous treatment of earned income after the pensionable age have high employment rates among people between 55 and 70. In contrast, nations that offer generous pensions and impose heavy penalties on earnings after the pensionable age have much lower employment rates. The economists who worked with Gruber and Wise often found striking age patterns in workers' withdrawal from paid employment. These age patterns often corresponded with incentives embedded in the national retirement system. Moreover, as in the United States, the age pattern of actual retirements often moved in parallel with changes in the age pattern of retirement incentives.

The evidence on aggregate retirement patterns seems to dovetail with other recent evidence showing that only a very small proportion of workers enter retirement with wealth that is too low to be consistent with a rational, far-sighted plan for lifetime consumption and retirement. Even the evidence that workers reduce their consumption after retirement, which violates a basic implication of the life-cycle model, can be explained by the lower consumption commitments of retirees and by rational saving and consumption responses to unexpected events that can trigger retirement.

Economists' rationalizations of observed behavior are a little harder to square with workers' responses when they are asked about their retirement plans and saving habits. A large percentage of Americans say they have given no thought to retirement, have saved too little for their anticipated old-age consumption, and do not know whether they will be able to afford to retire. Moreover, when asked to describe future benefits under their company or Social Security pensions, a majority often shows astonishing ignorance of the most basic provisions determining future retirement incomes. If workers do not take the trouble to learn how their pensions are

determined, it is a little hard to believe they use information sensibly to choose an optimal retirement and saving strategy.

One explanation for workers' ignorance and for the reported absence of planning is that many people anticipate using simple rules of thumb to choose their retirement age. This kind of decision rule may not be far-sighted, but it could be rational if workers do not expect to derive much benefit from a big investment in information gathering and retirement planning. The range of uncertainty about workers' future health, wages, and employment prospects is so wide that many people may believe well-informed, deliberative planning is a waste of effort. The process will need to be repeated every time a worker receives fresh news about health or potential earnings. This prospect is clearly unattractive for people who do not enjoy planning or are unskilled at performing it.

What difference does it make if workers select their saving behavior or choose a retirement age using imitation, simple rules of thumb, or other less-than-rational decision rules? How much is a worker's welfare harmed by use of a poor decision rule? In individual cases, the consequences of a poor choice of retirement age or pre-retirement saving rate can lead to very poor outcomes in old age. On the whole, however, only a small fraction of Americans faces severe material hardship after they retire. Serious deprivation is nowadays no more common among the elderly than it is among younger adults (Figure 4). Workers who use short-sighted decision rules probably enjoy less comfortable material circumstances in old age than they would enjoy if they based their decisions on rational and far-sighted rules. But the shortfall in their retirement income or consumption may be relatively small. Extensive survey evidence on subjective well-being suggests that most people make accommodations to modest changes over time in their income and consumption. People who experience income gains often report temporary increases in subjective well-being, and people who experience economic losses report declines in well-being, but in neither case does the change in well-being appear to persist for very long (Easterlin, 2003). Thus, compared with the happiness that retired workers could obtain under a rational and far-sighted retirement and saving plan, the happiness they actually achieve using less a far-sighted decision process may not be very different.

Economists' models of retirement and saving behavior should probably be considered worthwhile guides to sensible decision-making rather than realistic descriptions of worker thinking or behavior. Some workers follow an approximation of the life-cycle model in

choosing a saving strategy and retirement age. It seems doubtful, however, whether the model describes the actual reasoning of a large percentage of real-world workers. Enough workers follow the model so that it provides a useful explanation of some aggregate patterns in retirement behavior. It even yields helpful predictions of average, long-term responses to changes in retirement incentives. But we may not want to rely on the model to describe actual decision-making among most workers we encounter in real life – unless we rub shoulders with a lot of workers trained in college economics.

NOTES

¹ For more specific data, refer to the U.S. Department of Commerce, Bureau of the Census (1975). Retirement patterns were much more difficult to measure among women because most worked primarily within the home (and without pay) during most of their adult lives.

² A longer summary of the basic model can be found in Burtless (1999).

³ For good surveys, see Quinn, Burkhauser, and Myers (1990); Hurd (1990); Quinn and Burkhauser (1994); and Lumsdaine and Mitchell (1999).

⁴ In 1985 Modigliani was awarded the Nobel Memorial Prize in economics in large measure for his contributions to the theory of consumption. For his work on this subject, see Modigliani Brumberg (1954) and Ando and Modigliani (1963).

⁵ Social Security rules have changed in recent years to allow workers who have attained the “normal retirement age” – or age of entitlement for a full pension – to claim a pension, regardless of whether they have retired. The description of the benefit formula given in the text is meant to cover the rules as they applied during most of the history of the program. Those rules made it difficult or impossible for many workers to collect a full Social Security pension while continuing to work in their career jobs (see below).

⁶ Because of changes in the benefit formula, that is no longer true. Regardless of when workers retire between age 62 and 69, most will receive fair compensation in the form of higher monthly pensions if they delay their retirement by one extra year.

⁷ A more detailed and accurate description of the retirement earnings test is given by the Social Security Administration at <<http://www.ssa.gov/OACT/COLA/rtea.html>>.

⁸ If the labor force participation rate at age 63 is designated $LFPR_{63}$, the retirement rate at age 63 is calculated as $(LFPR_{63} - LFPR_{62}) / LFPR_{54}$. 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 54 through 72 in a particular year.

⁹ The calculations are described in full in Burtless and Moffitt (1985).

¹⁰ The tax rate on excess earnings was subsequently reduced for retirees age 65 and older. The retirement earnings test for workers older than the normal retirement age was eliminated altogether in 2000.

¹¹ It is not obvious, however, whether the optimal, far-sighted behavior of a neighbor offers a good guide to one’s own behavior. A neighbor who has accumulated greater wealth or who expects a shorter life span can comfortably retire at a younger age. Axtell and Epstein’s (1999) model works best in explaining imitative behavior when agents face a common change in incentives. The change in availability of Social Security pensions at age 62 is one example of such a change.

¹² For a good survey of the evidence up through the mid-1990s, see Lusardi and Browning (1996).

¹³ A similar conclusion about the adequacy of household saving was reached by Warshawsky and Ameriks (2001). Following a somewhat different procedure than Moore and Mitchell (2000), Warshawsky and Ameriks used widely available financial planning software to analyze the wealth and financial situation of working families earning less than \$125,000 in 1992. Under the assumption that each family's net income requirements would fall 20 percent when the breadwinner retired and an additional 20 percent when one spouse died, the analysts used Quicken's financial planner to calculate how long each family's resources would last after a worker retired. These calculations imply that about 35 percent of families will exhaust their financial wealth within 10 years of retirement. Since a large majority of worker households contain at least one member who will live longer than 10 years after retirement, this calculation suggests a large percentage of retirement and saving plans will fail sometime in retirement.

¹⁴ The HRS interview data have been matched to Social Security Administration records of workers' lifetime earnings, and researchers have access to the combined interview and earnings record data. See Scholz, Seshadri, and Khitatrakun (2004).

¹⁵ Some of the evidence on subjective well-being among the retired is summarized in Loewenstein, Prelec, and Weber (1999).

REFERENCES

- Aaron, H.J. (1966). "The Social Insurance Paradox," *Canadian Journal of Economics*, 32 (August), pp. 371-74.
- Abraham, K.G., and Hausman, S.N. (2004). "Work and Retirement Plans among Older Americans," Pension Research Council Working Paper. Philadelphia, PA: Pension Research Council, Wharton School.
- Ameriks, J., Caplin, A., and Leahy, J. (2002). "Retirement Consumption: Insights from a Survey," NBER Working Paper 8735. Cambridge, MA: National Bureau of Economic Research.
- Ando, A., and Modigliani, F. (1963). "The Life Cycle Hypothesis of Saving: Aggregate Implications and Tests," *American Economic Review*, vol. 53 (March), pp. 55-84.
- Axtell, R.L., and Epstein, J.M. (1999). "Coordination in Transient Social Networks: An Agent-Based Computational Model of the Timing of Retirement," in H.J. Aaron (Ed.), *Behavioral Dimensions of Retirement Economics*. Washington, D.C.: The Brookings Institution.
- Banks, J., Blundell, R., and Tanner, S. (1998). "Is There a Retirement-Savings Puzzle?" *American Economic Review*, 88(4), pp. 769-788.
- Bernheim, B.D. (1992). *Is the Baby Boom Generation Preparing Adequately for Retirement?* New York: Merrill Lynch and Company.
- Bernheim, B.D. (1995). *The Merrill Lynch Baby Boom Retirement Index: Update '95*, Technical Report. New York: Merrill Lynch and Company.
- Burtless, G. (1986). "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement," *Review of Economic Studies*, vol. 53 (October), pp. 781-805.
- Burtless, G. (1999). "An Economic View of Retirement," in Henry Aaron, ed., *Behavioral Dimensions of Retirement Economics*. Washington, D.C.: The Brookings Institution, pp. 7-42.
- Burtless, G. (2004). "Pension Reform and Labor Force Exit: Cross-National Evidence," Brookings working paper prepared for the "International Forum of the Collaboration Projects" meetings in Tokyo, Japan, February.
- Burtless, G., and Moffitt, R.A. (1985). "The Joint Choice of Retirement Age and Postretirement Hours of Work," *Journal of Labor Economics*, vol. 3 (April), pp. 209-36.
- Chan, S., and Stevens, A.H. (2003). "What You Don't Know Can't Help You: Pension Knowledge and Retirement Decision Making," NBER Working Paper 10185. Cambridge, MA: National Bureau of Economic Research.
- Diamond, P.A., and Hausman, J.A. (1984). "Individual Retirement and Savings Behavior," *Journal of Public Economics*, vol. 23(1-2), pp. 81-114.
- Easterlin, R.A. (2003). "Building a Better Theory of Well-Being," IZA Discussion Paper No. 742. Bonn: Institute for the Study of Labor (IZA).

- Employee Benefit Research Institute (2002). "Income of the Elderly, 2000" Washington: EBRI.
- Engen, E.M., Gale, W.G., and Uccello, C.E. (1999). "The Adequacy of Household Saving," *Brookings Papers on Economic Activity*, (1999:2), pp. 65-165.
- Grad, S. (2002). *Income of the Population 55 and Over, 2000*, Washington, DC: U.S. Social Security Administration.
- Gruber, J., and Wise, D.A. (1999). *Social Security and Retirement Around the World*. Chicago: University of Chicago Press.
- Gustman, A.A., and Steinmeier, T.L. (1986). "A Structural Retirement Model," *Econometrica*, vol. 54 (May), pp. 555-84.
- Gustman, A.A., and Steinmeier, T.L. (forthcoming). "What People Don't Know About Their Pensions and Social Security: An Analysis Using Linked Data From the Health and Retirement Study" in W.G. Gale, J.B. Shoven and M.J. Warshawsky (Eds.), *Private Pensions and Public Policies*. Washington, D.C.: Brookings Institution.
- Hamermesh, D.S. (1985). "Consumption during Retirement: The Missing Link in the Life Cycle," *Review of Economics and Statistics* 66(1), pp. 1-7.
- Hausman, J.A., and Paquette, L. (1987). "Involuntary Early Retirement and Consumption," in G. Burtless (Ed.), *Work, Health and Income among the Elderly*. Washington, DC: The Brookings Institution, pp. 151-82.
- Hurd, M.D. (1990). "Research on the Elderly: Economic Status, Retirement, and Consumption and Saving," *Journal of Economic Literature* 28(2), pp. 565-637.
- Hurd, M.D., and Rohwedder, S. (2004). "The Retirement-Consumption Puzzle: Anticipated and Actual Declines in Spending at Retirement," Michigan Retirement Research Center Working Paper 2004-069. Ann Arbor, MI: Michigan Retirement Research Center.
- Krueger, A.B., and Pischke, J.S. (1992). "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation," *Journal of Labor Economics* vol. 10(4) (October), pp. 412-37.
- Loewenstein, G., Prelec, D., and Weber, R., (1999). "What, Me Worry? A Psychological Perspective on Economic Aspects of Retirement," in H.J. Aaron (Ed.), *Behavioral Dimensions of Retirement Economics*. Washington, D.C.: The Brookings Institution.
- Lumsdaine, R., and Mitchell, O.S. (1999). "New Developments in the Economics of Retirement," in O. Ashenfelter and D. Card (Eds.), *Handbook of Labor Economics, Volume 3*. Amsterdam: North Holland, pp. 3261-3308.
- Lumsdaine, R., Stock, J., and Wise, D.A. (1992). "Three Models of Retirement: Computational Complexity versus Predictive Validity" in D.A. Wise (Ed.), *Topics in the Economics of Aging*. Chicago: University of Chicago Press, pp. 19-57.
- Lusardi, A. (2001). "Explaining Why So Many People Do Not Save," Center for Retirement Research Working Paper 2001-05. Chestnut Hill, MA: Center for Retirement Research at Boston College.
- Lusardi, A., and Browning, M. (1996). "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, vol. 34 (December), pp. 1797-1855.

- Modigliani, F., and Brumberg, R. (1954). "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in K.K. Kurihara (Ed.), *Post Keynesian Economics*. New Brunswick, N.J.: Rutgers University Press, pp. 388-436.
- Moore, J., and Mitchell, O.S. (2000). "Projected Retirement Wealth and Saving Adequacy" in O.S. Mitchell, B. Hammond, and A. Rappaport (Eds.), *Forecasting Retirement Needs and Retirement Wealth*. Philadelphia, PA: Univ. of Pennsylvania Press, pp. 68-94.
- Quinn, J.F., and Burkhauser, R.V. (1994). "Retirement and Labor Force Behavior" in L.G. Martin and S.H. Preston (Eds.), *Demography of Aging*. Washington, D.C.: National Academy Press, pp. 50-101.
- Quinn, J.F., and Burkhauser, R.V., and Myers, D.A. (1990). *Passing the Torch: the Influence of Economic Incentives on Work and Retirement*. Kalamazoo, Michigan: W. E. Upjohn Institute for Employment Research.
- Scholz, J.K., Seshadri, A., and Khitatrakun, S. (2004). "Are Americans Saving 'Optimally' for Retirement?" NBER Working Paper 10260. Cambridge, MA: National Bureau of Economic Research.
- U.S. Department of Commerce, Bureau of the Census (1975). *Historical Statistics of the United States: Colonial Times to 1970*. Washington, DC: U.S. Government Printing Office.
- Vroman, W. (1985). "Some Economic Effects of the Social Security Retirement Test," in R. Ehrenberg (Ed.), *Research in Labor Economics*. Greenwich, CT: JAI Press, pp. 31-89.
- Warshawsky, M., and Ameriks, J. (2001). "What Does Financial Planning Software Say about Americans' Preparedness for Retirement?" *Journal of Retirement Planning* (May-June), pp. 27-37.

Figure 1. Life-Cycle Earnings, Consumption, and Wealth Accumulation of a Far-Sighted Saver


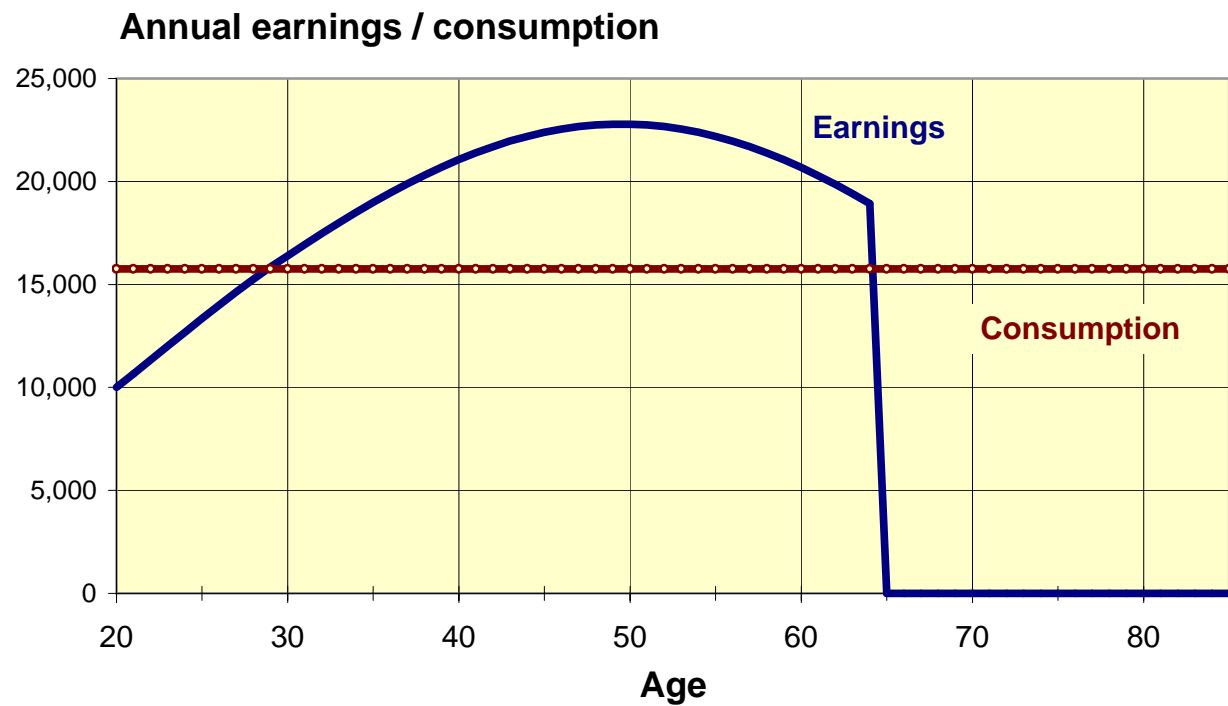





Figure 2. Male Retirement Rate by Age, 1940 - 2000



Note: Percent retiring each year is a constructed number reflecting the fraction of men leaving the workforce at the designated age, measured as a percent of men in the labor force at age 54.


Source: Author's tabulations of participation rates reported by U.S. Census Bureau for 1940, 1960, and 1970 decennial censuses and tabulations of 12 monthly public-use Current Population Survey files for 1980, 1990, and 2000 calendar years.

Figure 3. Retirement Age and Post-Retirement Earnings Distributions in the Longitudinal Retirement History Survey, 1969-1979



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

Figure 4. U.S. Poverty Rate among Aged and Non-Aged Adults, 1959-2003



Source: U.S. Bureau of Census.