

ERIC M. ENGEN

*Federal Reserve Board*

WILLIAM G. GALE

*Brookings Institution*

CORI E. UCCELLO

*Urban Institute*

## *The Adequacy of Household Saving*

DURING THE PAST half century, retirement income security in the United States has been based on a combination of social security, employer-sponsored pensions, and households' own saving. Social security was intended to provide a retirement income base. Pensions generated additional retirement income. Households' own saving supplemented these sources. In many ways this combination has served retirees well, but recent and impending developments have raised concerns about the adequacy of households' preparations for retirement.<sup>1</sup>

Social security faces a long-term financial imbalance, owing to lengthening life spans, earlier retirement, and, over the next few decades, the demographic bulge created by the retirement of the baby boom generation. Any viable solution must in some way reduce retirement benefits or raise

We thank Benjamin Harris, Jasper Hoek, Joseph Milano, Kevin Perese, Rumke Saha, James Sly, and Joshua Twilley for outstanding research assistance. We have received helpful comments from Henry Aaron, John Ameriks, Orazio Attanasio, Christopher Carroll, Darrel Cohen, Victor Fuchs, David Laibson, Dean Maki, Christopher Mayer, Olivia Mitchell, Michael Palumbo, Wayne Passmore, Andrew Samwick, John Karl Scholz, Todd Sinai, Jonathan Skinner, Kent Smetters, Nicholas Souleles, Mark Warshawsky, Stephen Zeldes, and seminar participants at the American Economic Association meetings and at Brookings, the University of Delaware, George Mason University, Georgetown University, the Wharton School of Business, and the NBER Summer Institute. William Gale gratefully acknowledges financial support from the National Institute on Aging, the Social Security Administration, and the TIAA-CREF Institute. Cori Uccello's research was supported by the Urban Institute's Retirement Project, funded through a generous grant from the Andrew W. Mellon Foundation. The views expressed in the paper are those of the authors and should not be attributed to the staff, officers, or trustees of any of the organizations with which they are associated, or any of the funding organizations.

1. Recent studies showing varying degrees of concern include American Academy of Actuaries (1998), Committee for Economic Development (1995), Employee Benefit Research Institute (1994), Gist, Wu, and Ford (1999), and U.S. Department of Labor (1998).

taxes.<sup>2</sup> Private pensions have shifted away from defined benefit (DB) plans to defined contribution (DC) plans. Although more portable than DB plans, DC plans also give workers increased discretion over participation, contribution, investment, and withdrawal decisions, and thus raise concerns about how effectively workers will use these instruments to finance retirement.<sup>3</sup> Other household saving has fallen dramatically in recent years, according to the National Income and Product Accounts.<sup>4</sup> Concerns generated by these aggregate trends have been increased by numerous microeconomic studies, reviewed below, which conclude from observed wealth accumulation patterns that a significant portion of today's working-age households will be unable to maintain current living standards in retirement.

These changing prospects for retirement income security raise a host of issues. Key research issues revolve around the extent to which households are forward-looking and able to save sufficiently to meet future consumption needs. Policy issues focus on the stability and direction of the nation's retirement system, prospects for the living standards of future retirees, and the extent to which a lack of financial preparation will translate into pressure for increased government assistance.

This paper provides a new analysis of the adequacy of household saving and is organized as follows. In the first section we define "adequate" saving as wealth accumulation that is sufficient for households to smooth the marginal utility of their consumption over time. We contrast our definition with alternative definitions and discuss how our definition affects the scope of the subsequent analysis.

In the second section we develop a stochastic life-cycle simulation model in which households save both for retirement and as a precaution against uncertain future earnings. The model formalizes our definition of adequate saving and provides a set of quantitative benchmarks against which to measure the adequacy of actual saving behavior. We use the sim-

2. Diamond (1997); Engen and Gale (1997).

3. Mitchell and Schieber (1998); Gale, Papke, and Vanderhei (1999).

4. Gale and Sabelhaus (1999) find that official saving rates vastly underreport saving. This is consistent with the fact that the aggregate household wealth-income ratio is at historically high levels (Federal Reserve Board, 1999). But this finding may be misleading for assessing the adequacy of a typical household's financial preparations for retirement, because wealth holdings are skewed in the United States (Poterba and Samwick, 1995; Kenickell, Starr-McCluer, and Sunden, 1997).

ulation model to generate three results regarding optimal wealth and consumption patterns. First, uncertainty about earnings implies that there will be a *distribution* of optimal wealth-earnings ratios, rather than a single benchmark ratio, among households that are otherwise observationally equivalent (that is, have the same age, education, pension status, marital status, and current wage). This finding fundamentally changes the interpretation of observed saving patterns. In particular, it implies that some households should be expected to exhibit low wealth-earnings ratios even if every household is forward-looking and making optimal choices.

Second, because of earnings uncertainty, optimal consumption rises with age during the working years for a wide range of time preference rates, holding interest rates and family size constant. This implies that stated preferences about age-consumption profiles cannot be mapped easily onto an implied time preference rate, unless the full economic situation is specified very carefully.

Third, owing to increases in mortality risk, optimal consumption generally declines as households reach and transit through retirement. As a result, their optimal wealth decumulation involves the eventual exhaustion of nonannuitized assets well before the longest possible life span. These results will prove critical in reconciling our empirical results with those of other studies.

The third section presents the central empirical work. Using the Health and Retirement Survey (HRS) of the University of Michigan's Institute for Social Research and the Survey of Consumer Finances (SCF) conducted by the Federal Reserve, we examine wealth and earnings data for married couples where the husband works full time. Because the simulation generates a distribution of optimal wealth-earnings ratios for a given set of household characteristics, we are unable to determine any household's precise optimal level of wealth. This finding shapes our empirical strategy, which focuses mainly on two issues: determining the proportion of households in the data who exceed the simulated median wealth-earnings ratio for households with their characteristics, and comparing the distributions of observed and simulated wealth-earnings ratios. Using what we regard as the most reasonable specification—a time preference rate of 3 percent and a definition of retirement wealth that includes half or more of housing equity—we find that more than half of households have actual wealth-earnings ratios that exceed the median simulated wealth-earnings ratio for households with the same characteristics. Indeed, in some cases the pro-

portion is well above half. In addition, the simulation model underestimates actual wealth among households with high ratios of wealth to earnings. Both of these results suggest that wealth accumulation is adequate for a majority of households in the sample. However, among households with low wealth-earnings ratios there is mixed evidence of undersaving at the 5th and the 25th percentiles of the wealth-earnings distribution.

Households that exceed the median simulated wealth-earnings ratios differ in predictable and plausible ways from others. After other factors are controlled for, they have more education and more pension coverage, and they are more likely to be self-employed, to plan on retiring early, to have thought about retirement, to have a long financial horizon, and to have received a large inheritance. The adequacy of saving fell somewhat between 1983 and 1995, but even for 1995 the aggregate figures indicate that more than half of all households exceeded the simulated median wealth-earnings ratios.

In the paper's fourth section we present the results of sensitivity analysis and discuss potential biases in, and extensions of, the underlying model. Our results are sensitive to the treatment of housing as retirement wealth and to the time preference rate employed. The combination of excluding all housing wealth and using a time preference rate of zero in the simulation generates significant undersaving. We also show that variations in key preference parameters that are small—in the sense that they cannot be ruled out on the basis of existing empirical work—can nevertheless have significant effects on the benchmark wealth-earnings ratios. This suggests considerable uncertainty regarding any assessment of the adequacy of saving. Perhaps surprisingly, however, we find that a 30 percent reduction in social security benefits, or a stock market decline of 40 percent, would have relatively small effects on the proportion of households whose wealth-earnings ratios exceed the simulated median ratios. Increased health care costs in retirement and increases in life span are estimated to have larger effects.

We then compare our findings with other sources of information about retirement saving. In the fifth section we examine some examples of popular financial advice, which often recommends that households aim to replace a certain portion—usually between 65 and 85 percent—of pre-retirement income in retirement. We show that, under plausible conditions regarding social security, employer-provided pensions, part-time work during retirement, and other factors, reaching these replacement rates often does not require much in the way of discretionary financial saving.

That is, saving “enough” does not necessarily imply that households need to accumulate very much in the form of financial assets. Thus the common observation that households on the eve of retirement have low levels of financial assets is not in itself evidence of systematic undersaving. We also show that the simulation model generates average replacement rates—including social security, employer-provided pensions, and asset income—between 70 and 80 percent of final earnings. This suggests that the popular financial advice is not inconsistent with rational utility maximization, and it supports the view that relatively low accumulations of financial savings can be perfectly consistent with optimizing behavior.

In the sixth section of the paper we examine previous microeconomic studies. We show that most previous studies that have been interpreted as showing inadequate household saving can be largely reconciled with our findings. Some of the differences in interpretation stem from the use of different benchmarks for adequate saving. In particular, no previous study incorporates the notion that there should be a distribution of optimal wealth-earnings ratios among observationally equivalent households, and many do not allow for declining consumption as retired households age. Other differences stem from different measures of wealth—we argue it is appropriate to include at least a significant portion of housing wealth in retirement wealth calculations—and from a variety of other factors.

In the penultimate section we discuss findings from surveys that ask respondents if they feel well prepared for retirement. We suggest that although some of the surveys suggest very little retirement preparation, the survey answers are sometimes difficult to interpret, and many surveys appear to suggest quite significant amounts of preparation. We conclude the paper by placing our results in a broader context.

### **What Is Adequate Saving?**

A number of alternative definitions of adequate saving could be employed. At the aggregate level, adequate saving might be associated with the golden-rule level of capital accumulation.<sup>5</sup> At the household level,

5. See Phelps (1961). The golden-rule capital stock level is obtained when the marginal product of capital, net of depreciation, is equal to the sum of the rate of labor-augmenting technological change plus the rate of growth of the labor force.

adequacy can be defined relative to the proportion of the future elderly population who will be in or near poverty, or in terms of the living standards of the future elderly relative to today's elderly. Each of these definitions is useful and valid for some purposes, but none matches the definition we employ.

We define a household to be saving adequately if it is accumulating enough wealth to be able to smooth its marginal utility of consumption over time in accordance with the optimizing model of consumption described in the next section. Several features of this definition are worth emphasizing. First, our definition is model-based and as such will depend on all of the features of the underlying model. Second, we define adequacy at the household level. As a result, our definition bears no relation to golden-rule levels of aggregate capital accumulation, because households in the model make choices that are conditional on government spending programs and taxes, which are not relevant considerations in determining the golden rule. Thus, even if all households are saving optimally, given government policies, the economy could still be below the golden-rule level of aggregate capital accumulation.

Third, our definition is based on comparisons of the marginal utility of pre- and postretirement consumption (adjusted for family size). As a result, there is no relation between poverty rates among the elderly and our definition of adequacy. A household in poverty during its working years and in retirement may still be considered to be saving adequately by the definition we employ, if the marginal utility of postretirement consumption is not high relative to that of preretirement consumption.

Fourth, our definition is different from "saving enough to maintain pre-retirement living standards in retirement." The latter requires smoothing of consumption levels (adjusted for family size) over time, whereas our definition requires smoothing of the discounted marginal utility of consumption over time. Maintaining living standards is a special case of smoothing the marginal utility of consumption. In theory, our definition could require either more or less saving than maintaining living standards would.

Our definition is motivated by research controversies regarding whether households are forward-looking. It corresponds to the definition of adequacy used in most of the literature we review below, and it is a natural way to examine adequacy from the perspective of economic research. Some analysts are clearly more prepared than others to assume that

observed saving behavior is optimal.<sup>6</sup> As discussed by Douglas Bernheim,<sup>7</sup> lifetime saving choices may be suboptimal. Many types of behavior can be optimized gradually through a trial-and-error approach. In contrast, people proceed through the life cycle only once and do not have the opportunity to replay their saving history. Moreover, there is no market or feedback mechanism to punish suboptimal saving behavior before retirement. Finally, retirement saving choices can be very complex, especially in the presence of realistic types of uncertainty. Bernheim concludes, “It would be astonishing if the average individual, with no practice and little or no training, could on his first try act as if he was a perfectly rational, far-sighted utility maximizer.”<sup>8</sup>

Our analysis takes these concerns seriously. We do not assume that observed saving behavior is optimal. Rather, we simulate optimal behavior with the model and then compare the model results with actual wealth accumulation patterns. Thus, whereas in most other contexts deviations between a model and the data indicate that the model is flawed, we will interpret any shortfall of actual wealth relative to the model’s wealth patterns as evidence that *behavior* is flawed, that is, that actual saving is too low. This assessment, of course, is subject to any qualifications about features of the model that do not accurately capture the full set of incentives and opportunities facing households.

However, our results will speak only to whether the observed levels of wealth are consistent with the patterns of an optimizing model. They cannot in any way prove that people are actually solving the optimization problem defined in the simulation model. Nor do the results speak to several important related issues, such as whether social security or pensions are responsible for observed levels of wealth accumulation, or whether higher saving would raise households’ or national welfare. As a result, the policy implications of our findings may not be direct or obvious. In particular, we make no claim that attaining our definition of adequacy is the most appropriate goal for retirement income policy. For example, if it were determined that workers are saving inadequately by our definition,

6. See, for example, Lazear (1994) and the comments by Bernheim (1994a) and Skinner (1994).

7. Bernheim (1994a).

8. Bernheim (1994a, p. 172).

but that the shortfall is small, society might well decide that there are more pressing uses of limited public resources than raising the living standards of future retirees.

## **Modeling the Adequacy of Saving**

This section describes the model used to produce benchmarks for adequate saving. After discussing the underlying structure and base-case parameter specifications, we examine the model's implications for optimal wealth accumulation and consumption over the life cycle.

### *A Stochastic Life-Cycle Model of Saving*

Appendix A describes the model in detail.<sup>9</sup> Here we summarize the main features.

OVERVIEW. Households enter the model with two adults aged twenty-one.<sup>10</sup> One child is added at age twenty-five and a second at age twenty-eight. Each child leaves the home at age twenty-one.<sup>11</sup> Families are not linked across generations. Each adult faces an age-varying probability of dying, with a maximum life span of 110 years. Each year, the assets of those who die are bequeathed to members of the generation that is then forty-five years old. The bequests are distributed in accordance with the wealth distribution of these forty-five-year-olds, thus capturing the empirically established tendency of wealthier households to receive larger inheritances.<sup>12</sup> The inheritance is assumed to be unanticipated.

In each period, forward-looking households maximize expected lifetime utility by choosing total consumption (consumption per capita times the number of people in the household) and total saving subject to a lifetime budget constraint, nonnegativity constraints on net assets, income and payroll taxes, and uncertainty regarding future earnings, life span, and inheritances. There are no markets for insurance against these uncertain-

9. See also Engen (1993b); Engen, Gale, and Scholz (1994); Engen and Gale (1996).

10. The model developed in this paper examines married couples because our empirical analysis is focused on couples.

11. To smooth the effects of children entering and leaving the households, we allow each transition to occur on a pro rata basis over four years.

12. Gale and Scholz (1994).



ties. Because there is a positive probability of death at each age, borrowing against the uncertain portion of future income and inheritances is not permitted.

Utility is separable over time, and separable within a time period between consumption and leisure. The utility function for consumption exhibits constant relative risk aversion, a constant intertemporal elasticity of substitution, and constant prudence, which implies that risky income and uncertain life spans lead to precautionary saving. Thus households save for retirement and as a precaution against downturns in future income and the possibility of outliving assets once retired.

Before retirement, consumption may be financed by labor earnings, decumulations of previously accumulated assets, or inheritances received. After retirement, consumption is financed by assets accumulated earlier, which are fully taxable, and by annuity income from social security and private DB pensions. Labor supply is exogenous and retirement occurs at a predetermined age. Household earnings are modeled as the sum of a stochastic component and a nonstochastic component. The latter follows a hump-shaped pattern with respect to age and varies by education class.


Because the model does not have a closed-form solution, and the analytical solution would be intractable, we use a numerical solution method to solve households' consumption-saving problem. Earnings shocks over the life cycle are simulated with a random number generator for each of 10,000 households. Because households receive different earnings shocks, they end up with different realized income, consumption, saving, and wealth.

**PARAMETER VALUES.** The model requires specification of numerous parameter values. We highlight the most important specification issues here. Appendix B contains additional details.

**Mortality risk.** The conditional survival probabilities—the probability of living to age  $t + 1$  conditional on being alive at age  $t$ —used in the model are calculated as one minus the estimated conditional mortality probability. Mortality probabilities for males and females are based on estimates from the life tables for 1994 used by the Social Security Administration.<sup>13</sup> Conditional survival probabilities for each sex are shown in figure 1 up to age 110, the maximum life span in the model (an individual in the model

13. U.S. Social Security Administration (1997).

Figure 1. Conditional Survival Probabilities by Sex<sup>a</sup>



Source: Authors' calculations based on U.S. Social Security Administration (1997).

a. Probabilities are calculated as one minus the Social Security Administration's estimated conditional mortality probability.

dies with certainty after age 110). For males that have lived to age fifty, for example, the probability of living to age fifty-one is 99.4 percent; that for females is 99.6 percent. By age sixty-five these one-year survival probabilities drop to 97.6 and 98.7 percent, respectively, and they decline further to 92 and 95 percent at age eighty. By age ninety-five, the one-year survival probabilities are 74 percent for a male and 79 percent for a female. Life expectancy at age twenty-one is seventy-four years for males and eighty years for females.

*Retirement age.* We specify a retirement age of sixty-two years in the base case. Peter Diamond and Jonathan Gruber show that 50 percent or more of men and women are out of the labor force by that age.<sup>14</sup>

*Age-earnings profiles.* Because saving is the difference between income (which before retirement consists largely of labor earnings) and consumption, the specification of the age-earnings profile is an important determinant of optimal saving patterns. To estimate the mean age-earnings profile, we use panel data on earnings of employed heads of households and their spouses from the Panel Survey of Income Dynamics, conducted by the University of Michigan's Institute for Social Research, from 1980 to 1992. We exclude the self-employed and households where the household head is over sixty-five years old. We estimate a fixed-effects model with the logarithm of earnings as a function of age, age squared, and year dummies to control for macroeconomic effects (that is, aggregate wage growth; table 1 reports the regression results). Separate equations were estimated for household heads with sixteen or more years of education and for those with less education. Earnings for the group with more education are always higher, rise and fall more steeply, and peak at later ages than for the group with less education (figure 2). The wages of all age groups are assumed to rise by 1 percent per year to reflect aggregate growth in the economy. These age-earnings profiles are generally similar to those used by other researchers.<sup>15</sup>

*Earnings shocks.* Available empirical evidence suggests that individuals face substantial uncertainty in their labor earnings, and that the largest share of the variance is idiosyncratic to households rather than stemming from common aggregate shocks (that is, the business cycle). In a previ-

14. Diamond and Gruber (1999, figure 11.6).

15. For example, Carroll (1997a); Hubbard, Skinner, and Zeldes (1995); and Laibson, Reppetto, and Tobacman (1998).

**Table 1. Estimated Age-Earnings Profiles**

	<i>Years of education</i>	
	<i>&lt;16</i>	<i>≥16</i>
Regression estimates of earnings from labor <sup>a</sup>		
Constant	7.906	6.850
Coefficient on age	0.105	0.165
Coefficient on age squared	-0.0012	-0.0017
Peak earnings		
As percentage of earnings at age 25	184	326
As percentage of earnings at age 62	128	123
Age at which earnings peak	48	52
Replacement rate as a percentage of final earnings		
Social security only	35	21
Social security and private defined benefit pension	64	57

Source: Authors' estimates based on regressions using data from the Panel Survey of Income Dynamics.


a. The regressions relate the log of real earnings of the household head and spouse to a household fixed effect, age, age squared, and year dummy variables. The sample consists of households where the head is employed (but not self-employed) and is between the ages of twenty-one and sixty-four. The data cover the 1980–92 period.

ous study, because of concern that measurement error in self-reported earnings might overstate the variation in actual earnings, Eric Engen used data from the Internal Revenue Service–Michigan tax panel to estimate the stochastic process for the logarithm of earnings variations.<sup>16</sup> Measurement error is less of a problem with earnings data collected from Internal Revenue Service W-2 forms filed with income tax returns, because wages are directly reported by employers, who are required and have the incentive to keep accurate records of earnings paid. Based on that analysis, we model the stochastic process for labor earnings shocks as a first-order autoregressive process with a persistence parameter of 0.85 and a variance of 0.05. Under this specification, about half of a given shock to earnings remains after five years.<sup>17</sup> Also, this specification for earnings shocks implies a variance of log earnings equal to 0.18 (calculated as the variance of earnings shocks divided by one minus the square of the persistence parameter:  $0.05/[1 - 0.85^2]$ ) and a covariance of log earnings one year

16. Engen (1993b).

17. This specification for earnings shocks is similar to that of Hubbard, Skinner, and Zeldes (1995) and generates less variation in earnings than the random walk process specified in Zeldes (1989a) and Carroll (1997a).

**Figure 2. Estimated Age-Earnings Profiles by Pension Status<sup>a</sup>**



Source: Authors' estimates based on data from the Panel Survey of Income Dynamics, as described in the text.  
a. Specified as in table 1.

apart equal to 0.15 (the persistence parameter multiplied by the variance:  $0.85 \times 0.18$ ).<sup>18</sup>

*Income and payroll taxes.* We impose a progressive income tax structure, similar to the actual U.S. system in 1998, with statutory marginal rates of 15 percent, 28 percent, 31 percent, 36 percent, and 39.6 percent. The taxable income brackets, in dollars, are those effective in 1998 for joint tax filers. Households are allowed a standard deduction of \$7,100 and an exemption of \$2,650 for each person. To capture the effect of preferential capital gains tax rates and tax-preferred saving vehicles, without introducing the substantial complication of explicitly modeling tax-favored saving, tax rates on capital income are capped at 20 percent. The social security payroll tax is modeled by taxing labor earnings up to a limit of \$68,400 at a 6.2 percent rate—the employee share of the payroll tax.<sup>19</sup>

*After-tax real rate of return.* The only asset in the model has a riskless return. Nevertheless, we do not feel it is appropriate to use empirical values of the risk-free rate in the simulation. In the model, the interest rate has two roles: it affects the growth of consumption and the overall return on saving. If the model had a safe asset and risky assets, the Euler equation for optimal consumption growth would be determined by the return on the safe asset,<sup>20</sup> and the overall return on saving would be a weighted average of these assets.<sup>21</sup> The real risk-free rate of return on short-term Treasury bills has averaged about 1 percent historically. Longer-term government and corporate bonds have yielded about 2 percent in real terms, and the equity market about 9 percent, in the postwar period.<sup>22</sup> A market-weighted basket of these returns gives a real return of about 5 percent. To capture the dual roles played by a single asset return in the model, we take a midpoint of the historical real risk-free rate of return and a mix of all

18. These generate less variation in earnings shocks than in MaCurdy (1982) and Abowd and Card (1987, 1989). See Engen (1993b) for further discussion of the estimation of the earnings shock parameters and how they compare with other studies of earnings variation.

19. We assume that employees bear the full burden of the payroll tax, and therefore reported earnings have already been reduced by the employer share of the payroll tax.

20. See, for example, Kocherlakota (1996).

21. See, for example, Zeldes (1989b).

22. See, for example, Cochrane (1998).

returns, and thus use an average after-tax real rate of return of 3 percent (the average tax rate on capital income is used here).<sup>23</sup>

*Coefficient of relative risk aversion.* We use a base-case value of 3 for the coefficient of relative risk aversion, which implies a value of 0.33 for the intertemporal elasticity of substitution and is based on empirical estimates of the model.<sup>24</sup> This value is within the range of estimated values for risk aversion and intertemporal substitution from aggregate and micro-economic studies.<sup>25</sup> It is also similar to the risk aversion and intertemporal substitution values used in other simulations.<sup>26</sup>

*Replacement rates of social security and defined benefit pensions.* The model assumes that each household receives social security and pension benefits that are based on features of the average age-earnings profile of its education class, not on its actual wage profile. For example, among households without pensions, social security is assumed to replace 35 percent of average final earnings for those with less than sixteen years of education, and 21 percent of average final earnings for those with sixteen or more years of education. For households with both pensions and social security, the replacement rates of the two combined are 64 percent and 57 percent of final earnings for the two education groups, respectively. These values are based on pension data from the HRS and social security data from actual beneficiaries, as described in appendix B.<sup>27</sup> Real private DB pension benefits are assumed to decline by 1 percent per year.

23. This rate of return is equivalent to that used in Hubbard, Skinner, and Zeldes (1995), and somewhat lower than the 5 percent (before-tax) rate of return used in Laibson, Reppetto, and Tobacman (1998). Many other simulation studies have used a rate of return in the 3 to 5 percent range. Carroll (1997a) and Bernheim and Scholz (1993), however, used a risk-free rate of return.

24. Engen (1993a).

25. See Barsky and others (1997) and Deaton (1992) for discussions of other estimates. There is little empirical consensus on the value of the risk aversion coefficient, as estimates range from 1 to the very large double-digit values implied by the asset pricing literature. However, the equity premium has generally been considered a puzzle, because economists have typically assumed that a risk aversion coefficient greater than around 5 seems inconsistent with other observed behavior regarding risk (see Cochrane, 1998).

26. See, for example, Auerbach and Kotlikoff (1987); Carroll (1997a); Hubbard, Skinner, and Zeldes (1995); Laibson, Reppetto, and Tobacman (1998).

27. These replacement rates appear to be generally comparable to or lower than those used in Carroll (1997a); Hubbard, Skinner, and Zeldes (1995); and Laibson, Reppetto, and Tobacman (1998).

*Time preference rates.* Specifying the appropriate time preference rate is difficult but crucial. The goal of the model is to describe optimal (and, implicitly, time-consistent) behavior, rather than actual behavior. As a result, choosing the rate so that the model is well calibrated with household wealth data, or using estimates of time preference rates from previous empirical studies, would inappropriately impose the assumption that households' actual behavior was optimal. Basing the choice on time preference rates used in other simulation models would also be misleading, since most of these models aim to explain actual behavior.

An alternative is to examine respondents' responses to survey questions about their most desired consumption profiles, but this creates problems. First, the choices are typically hypothetical or involve only small amounts of money. Under these circumstances there is no reason to believe that respondents' answers are more representative of their true preferences than their actions are.<sup>28</sup> Second, we show below (figure 6) that consumption profiles that look similar can nevertheless be based on very different time preference rates.<sup>29</sup>

Faced with these constraints, we choose two values for the time preference rate: the average after-tax real interest rate (3 percent) and zero. Setting the time preference rate equal to the real average after-tax rate of return is a natural benchmark. It implies that, holding family size constant and setting mortality risk equal to zero, the household facing the average marginal tax rate would equate the marginal utility of consumption across each period of its existence. This is consistent with the basic notion of consumption smoothing (technically, marginal utility smoothing) that is central to dynamic optimization models. A time preference rate of

28. It is unclear whether the surveys are eliciting answers about respondents' true preferences, or what respondents wish their true preferences to be. If the latter, it is also unclear what to do with the information. For example, if high school-educated adults in a household indicate that they would prefer to have a college education, it would be difficult to justify modeling them in the simulation as having a college education. Likewise, it is difficult to decide on a true time preference rate for a household that says its time preference rate is one value but acts as if it is a different value. Becker and Mulligan (1997) provide a model of how individuals can invest in future-oriented capital and thus alter their time preference rates.

29. In particular, with nonstochastic earnings and life span, the slope of the consumption profile will be directly related to the difference between the interest rate and the time preference rate. With uncertainty, however, this need not occur. Upward-sloping consumption profiles during the working years can result from time preference rates larger than, equal to, or less than the interest rate.



3 percent is lower than previous empirical estimates based on consumption and saving behavior.<sup>30</sup> It is also low compared with values often used in previous simulation studies. We emphasize, however, that this value was chosen because it sets the after-tax interest rate equal to the time preference rate, not because it necessarily corresponds with observed behavior.

As a still lower alternative, it is plausible to consider a time preference rate of zero as another benchmark. If the household is thought of as a planner designing its own age-consumption profile, it may make sense not to value consumption in any period over consumption in any other period. There are, however, some problems with using a zero time preference rate. With no borrowing constraints, no uncertainty, and a 3 percent real after-tax rate of return on its assets, the amount of consumption that the household would like to defer is huge. The household would choose to consume so little when young that the marginal utility of consumption at age twenty-five would be almost six times that at age eighty-five.<sup>31</sup> With a constant relative risk aversion utility function and an intertemporal elasticity of substitution of one-third, as in our simulation, the consumption level at age eighty-five would be 81 percent higher than at age twenty-five (in the absence of productivity growth and holding family size constant). Even during middle age, the household would face significantly depressed consumption relative to old age. The marginal utility of consumption at age forty-five would be more than three times as high as the marginal utility at age eighty-five, and consumption would be 48 percent higher at age eighty-five than at age forty-five. Thus we view a zero time preference rate as assuming that households are extremely patient.

It would be possible to consider other rates, of course. However, with higher rates of time preference, undersaving would be less of a problem

30. See Dynan (1993); Engen (1993a); Lawrance (1991); Samwick (1998). Studies of household durables purchases have yielded a wide variety of estimated time preference rates. See Hausman (1979); Gately (1980); Hassett and Metcalf (1995).

31. With discrete time periods and no uncertainty, the Euler equation linking marginal utility over time equates the marginal utility of consumption at age twenty-five with the marginal utility of consumption at age eighty-five multiplied by a factor equal to one plus the interest rate divided by one plus the time preference rate, all raised to the 60th power:

$$U'(C_{25}) = \{([1 + r]/[1 + \delta])^{(85-25)}\} U'(C_{85}),$$

where  $U'(C_i)$  is the marginal utility of consumption at age  $i$ , for  $i = 25, 85$ . With  $r = 0.03$  and  $\delta = 0$ , the term in brackets equals 5.89.

than reported below. We also considered using a negative time preference rate but rejected the idea for reasons discussed in appendix B.

### *Model Results*

We begin by reporting results in terms of the ratio of current wealth to current earnings.<sup>32</sup> This is done to normalize the results with respect to changes in productivity and inflation over time and for ease of comparison with some previous studies.<sup>33</sup> The wealth measure excludes the present value of social security and DB pension benefits, to facilitate comparisons with the empirical results presented later. Because variations in the time preference rate proved to be important determinants of the results, we present findings using both time preference rates specified above.

**MEDIAN WEALTH-EARNINGS RATIOS.** Optimal wealth-earnings ratios will evolve differently for different households for two reasons. The first is that households differ by education status and private pension coverage. Differences in education affect the level and shape of the age-earnings profile. Differences in pension coverage affect retirement income. These factors have been included in other studies.

The top panel of table 2 reports median optimal wealth-earnings ratios for households classified by age, education, and pension status, assuming a time preference rate of 3 percent. Simulated optimal wealth-earnings ratios rise over the life cycle. When education status is controlled for, households with pensions have lower optimal wealth-earnings ratios than those without, because pensions provide retirement income. When pension status is controlled for, college graduates have lower optimal wealth-earnings ratios when young and almost equal or higher ratios when old

32. In previous work we have shown that, parameterized to represent actual behavior, the model generates plausible wealth-income ratios, saving elasticities, and aggregate saving rates (Engen, 1993b; Engen and Gale, 1996). At the household level, consumption tracks income more closely in this model than in certainty models, and simulated consumption is more sensitive to income shocks in this model than in a certainty-equivalent model. These are well-documented features of actual consumption and saving data (Engen, 1993b). Both the model results and the subsequent data analysis focus on before-tax values of wealth and earnings.

33. Despite our reporting the results this way, our model should not be confused with a "buffer stock" or target saving model (see Carroll, 1992). In our model, as already noted, households save both for retirement and as a precaution against uncertain income and life span.

**Table 2. Median Simulated Wealth-Earnings Ratios by Age, Education, and Pension Status**

Age	<i>Education &lt;16 years</i>		<i>Education ≥ 16 years</i>	
	<i>No pension</i>	<i>Pension</i>	<i>No pension</i>	<i>Pension</i>
<i>Time preference rate = 3 percent</i>				
30–34	0.25	0.23	0.06	0.06
35–39	0.56	0.46	0.20	0.14
40–44	1.08	0.83	0.62	0.35
45–49	1.84	1.36	1.39	0.78
50–54	2.70	1.97	2.40	1.39
55–59	3.76	2.66	3.67	2.19
60–62	4.74	3.28	4.91	2.92
<i>Time preference rate = 0</i>				
30–34	0.57	0.46	0.14	0.10
35–39	1.19	0.94	0.52	0.31
40–44	2.00	1.60	1.24	0.76
45–49	3.00	2.42	2.23	1.48
50–54	4.10	3.32	3.41	2.35
55–59	5.24	4.30	4.82	3.39
60–62	6.47	5.16	6.20	4.37

Source: Authors' estimates as described in the text.

than do other households. This reflects the steeper and later-peaking earnings profiles of college graduates than of other households.

The median wealth-earnings ratios in table 2 are significantly higher than similar targets calculated by Bernheim and John Karl Scholz.<sup>34</sup> For households in their sixties, our median wealth-earnings ratio exceeds the Bernheim-Scholz target by about 45 percent for households without pensions, by 37 percent for non-college-educated households with pensions, and by 12 percent for college-educated households with pensions (see appendix table D1). The higher values in our model arise from the existence of precautionary saving, the earlier retirement age (sixty-two as opposed to sixty-five years), and other factors. Thus we believe that the basic model with a 3 percent time preference rate provides a conservative basis on which to judge the adequacy of saving.

When the time preference rate is assumed to be zero (bottom panel of table 2), several changes occur. Naturally, the median wealth-earning ratios are higher in this case, and the change is substantial. By ages sixty

34. Bernheim and Scholz (1993).

to sixty-two, the median wealth-earnings ratios are greater by between 26 percent and 57 percent, depending on the specification of education and pension status. When education is controlled for, the change in time preference raises the ratios for households with pensions by a larger proportionate amount than for households without pensions. When pension status is controlled for, the lower time preference rate raises wealth-earnings ratios more for those with fewer years of education. For households in their sixties, the wealth-earnings ratios with a time preference rate of zero range between 170 and 215 percent of the Bernheim-Scholz targets.

**DISTRIBUTION OF WEALTH-EARNINGS RATIOS.** The second reason why wealth-earnings ratios vary across households is distinct to our methodology: households receive different earnings shocks over time and at a given point in time. As a result, households that are observationally equivalent in the data—that is, that are identical with respect to age, current earnings, family size, life expectancy, education, and pension status—will have different optimal wealth-earnings ratios.

Table 3 shows the importance of heterogeneous earnings shocks in generating a *distribution* of wealth-earnings ratios. The table focuses on college graduates with pensions. (Appendix tables D2 through D4 report similar results for other households.) For these households, with a time preference rate of 3 percent, wealth-earnings ratios among thirty-five- to thirty-nine-year-olds vary by a factor of 100, from 0.01 at the 5th percentile to 1.02 at the 95th percentile. Among sixty- to sixty-two-year-olds, wealth-earnings ratios vary by a factor of almost 20, from 0.37 at the 5th percentile to 7.05 at the 95th percentile.

With a time preference rate of zero, the range is almost as significant, with the optimal ratios varying by a factor of 160 among thirty-five- to thirty-nine-year-olds and a factor of 7 among sixty- to sixty-two-year-olds. Reducing the time preference rate to zero raises the optimal wealth-earnings ratio at ages sixty to sixty-two by between 0.9 and 1.7 at the various percentile points marked in the table. The ratio rises by 243 percent at the 5th percentile, 50 percent at the median, and 26 percent at the 95th percentile. For this age group, the ratio at the 25th percentile of the wealth-earnings distribution with a time preference rate of zero (2.94) roughly equals the ratio at the median of the distribution with a time preference rate of 3 percent (2.92).

Several features of these results merit comment. Most important, these observed ratios represent households' *optimal* responses to the pattern of

**Table 3. Distribution of Simulated Wealth-Earnings Ratios by Age Among Households with Sixteen or More Years of Education and with Private Pensions**

Age	5th percentile	25th percentile	Median	75th percentile	95th percentile
<i>Time preference rate = 3 percent</i>					
30–34	0.00	0.02	0.06	0.17	0.52
35–39	0.01	0.04	0.14	0.39	1.02
40–44	0.01	0.11	0.35	0.81	1.82
45–49	0.04	0.31	0.78	1.48	2.94
50–54	0.12	0.69	1.39	2.35	4.15
55–59	0.29	1.22	2.19	3.41	5.77
60–62	0.37	1.68	2.92	4.35	7.05
<i>Time preference rate = 0</i>					
30–34	0.00	0.03	0.10	0.29	0.82
35–39	0.01	0.08	0.31	0.72	1.63
40–44	0.03	0.29	0.76	1.44	2.75
45–49	0.13	0.76	1.48	2.38	4.19
50–54	0.43	1.44	2.35	3.50	5.62
55–59	0.89	2.25	3.39	4.86	7.47
60–62	1.27	2.94	4.37	6.05	8.88

Source: Authors' estimates as described in the text.

earnings shocks they receive. The low wealth accumulation exhibited by a significant minority of households in the simulation model is consistent with optimizing behavior and in no way implies a retirement saving short-fall owing to myopia, irrationality, or poor information.

Moreover, the wide variation is not owing to differences in current earnings. Table 4 shows the distribution of simulated wealth-earnings ratios by current earnings quintile among college-educated households with pensions for two age groups: those aged fifty to fifty-four and those aged sixty to sixty-two. Appendix tables D5 through D7 show results for other households in these age groups. Even within relatively narrow earnings bands, the variation in optimal wealth-earnings ratios is substantial. In the fourth earnings quintile in table 4, for example, the optimal ratios vary between the 5th and 95th percentiles by a factor of 13 for fifty- to fifty-four-year-olds and by a factor of 8 for sixty- to sixty-two-year-olds.

OPTIMAL CONSUMPTION AND WEALTH PROFILES. Figures 3 and 4 display model results for optimal median consumption per adult equivalent as a function of age, education, and pension status. By focusing on consumption per adult equivalent rather than on total household consumption,

**Table 4. Distribution of Simulated Wealth-Earnings Ratios by Current Earnings Among Households with Sixteen or More Years of Education and with Private Pensions<sup>a</sup>**

<i>Earnings quintile</i>	<i>5th percentile</i>	<i>25th percentile</i>	<i>Median</i>	<i>75th percentile</i>	<i>95th percentile</i>
<i>Ages 50–54</i>					
Lowest	0.02	0.17	0.74	1.80	4.11
2nd	0.10	0.47	1.10	2.14	4.05
3rd	0.19	0.65	1.28	2.21	3.92
4th	0.30	0.83	1.48	2.34	3.87
Highest	0.49	1.02	1.53	2.25	3.39
<i>Ages 60–62</i>					
Lowest	0.10	1.02	2.86	4.95	8.57
2nd	0.35	1.57	3.08	4.80	7.52
3rd	0.50	1.82	3.18	4.66	6.84
4th	0.73	1.84	2.93	4.17	6.20
Highest	0.96	1.89	2.69	3.69	5.28

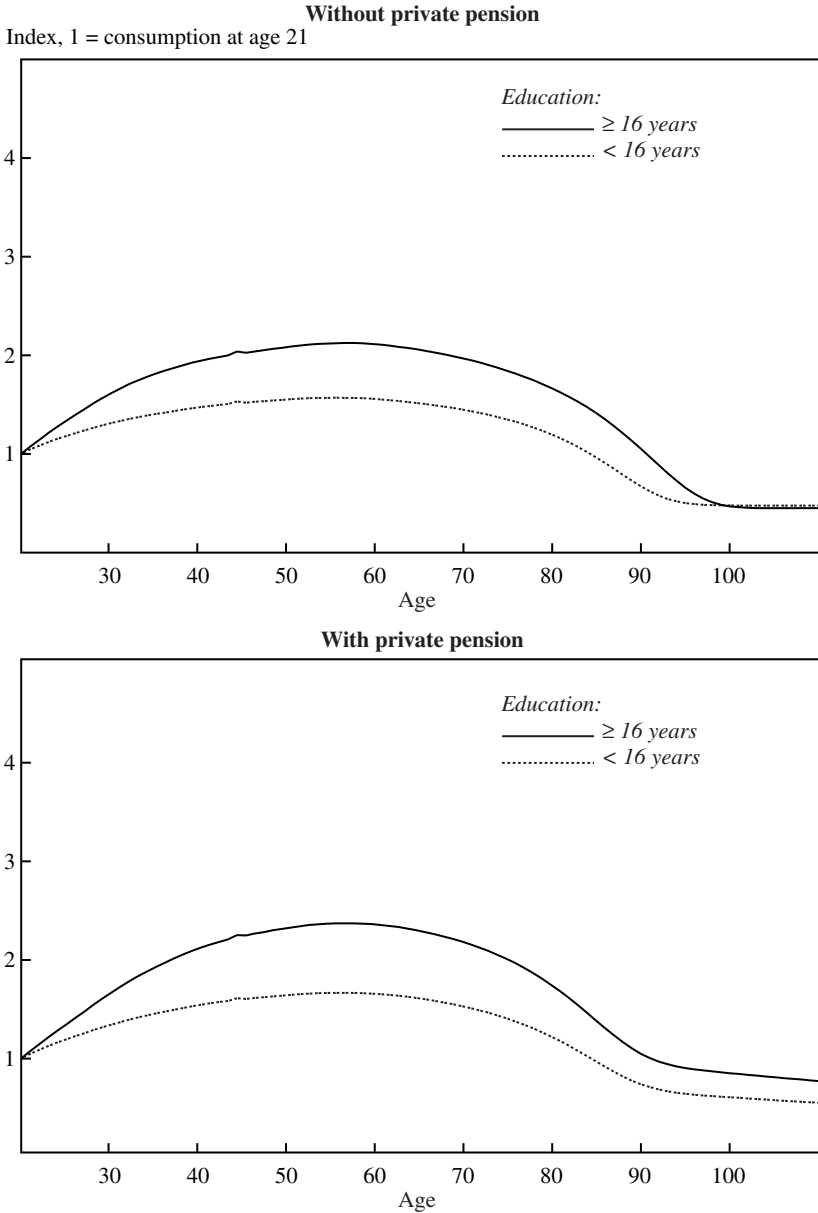
Source: Authors' estimates as described in the text.

a. These simulations employ a time preference rate of 3 percent.

the figures remove any impact of changes in family size on the consumption profile. For each profile the data are normalized so that consumption at age twenty-one is equal to 1. Figure 3 reports the results of using a 3 percent time preference rate; figure 4 uses a time preference rate of zero. For either time preference rate, households with a private pension have higher lifetime consumption, holding education status constant, because the pension gives them greater lifetime wealth. When pension status is held constant, those with greater education have higher lifetime consumption paths because they have greater human capital (as reflected in higher lifetime wages).


The model implies hump-shaped consumption profiles—rising when young, peaking before retirement, and then generally falling throughout old age—regardless of education or pension status. Consumption is low when households are young (even when the real rate of return on assets is equal to the time preference rate), because households desire to build up their precautionary saving. As households age, income and wealth rise, some income uncertainty is resolved, and the precautionary motive for saving edges off; all these factors lead to rising consumption during the working years. Consumption declines in old age as an increasing mortality probability effectively makes households less patient and less willing to

**Figure 3. Optimal Age-Consumption Profiles by Pension Status with a Time Preference Rate of 3 Percent<sup>a</sup>**



Source: Authors' calculations using the stochastic life-cycle simulation model described in the text.  
a. Consumption is per adult equivalent.

**Figure 4. Optimal Age-Consumption Profiles by Pension Status with a Time Preference Rate of Zero<sup>a</sup>**



Source: Authors' calculations using the stochastic life-cycle simulation model described in the text.

a. Consumption is per adult equivalent.



defer consumption to an uncertain future. Hump-shaped consumption profiles like those in figures 3 and 4 are consistent with observed age-consumption patterns of households.<sup>35</sup>

When the time preference rate is zero rather than 3 percent, the consumption profile peaks at a later age, because households are more patient and thus willing to defer more consumption. However, the general pattern of consumption is similar with either time preference rate.

If household members live long enough, optimal behavior in this model, which includes annuity income from social security and in some cases a private pension, suggests that the household should at some time optimally deplete its financial wealth and rely solely on annuity income to finance consumption in late old age.<sup>36</sup> This can be seen in figures 3 and 4 where consumption flattens out some time after age ninety. This depletion of nonannuity wealth is even more evident in figure 5, which shows the median age-wealth profiles corresponding to a time preference rate of 3 percent.<sup>37</sup> In the cases shown, it is optimal to deplete all nonannuity wealth by about age ninety to ninety-five.

Lastly, it is important to note that the presence of earnings uncertainty, as in our model, has important implications for the relation between the time preference rate and the slope of the age-consumption profile. With no uncertainty, the slope of the age-consumption profile is uniquely determined by the difference between the time preference rate and the interest rate. For example, the top panel of figure 6 shows—in a nonstochastic model with an interest rate of 3 percent—that with a time preference rate of zero, consumption rises over the life cycle; with a time preference rate of 3 percent, consumption is flat; and with a time preference rate of 6 percent, consumption falls over the entire life cycle.

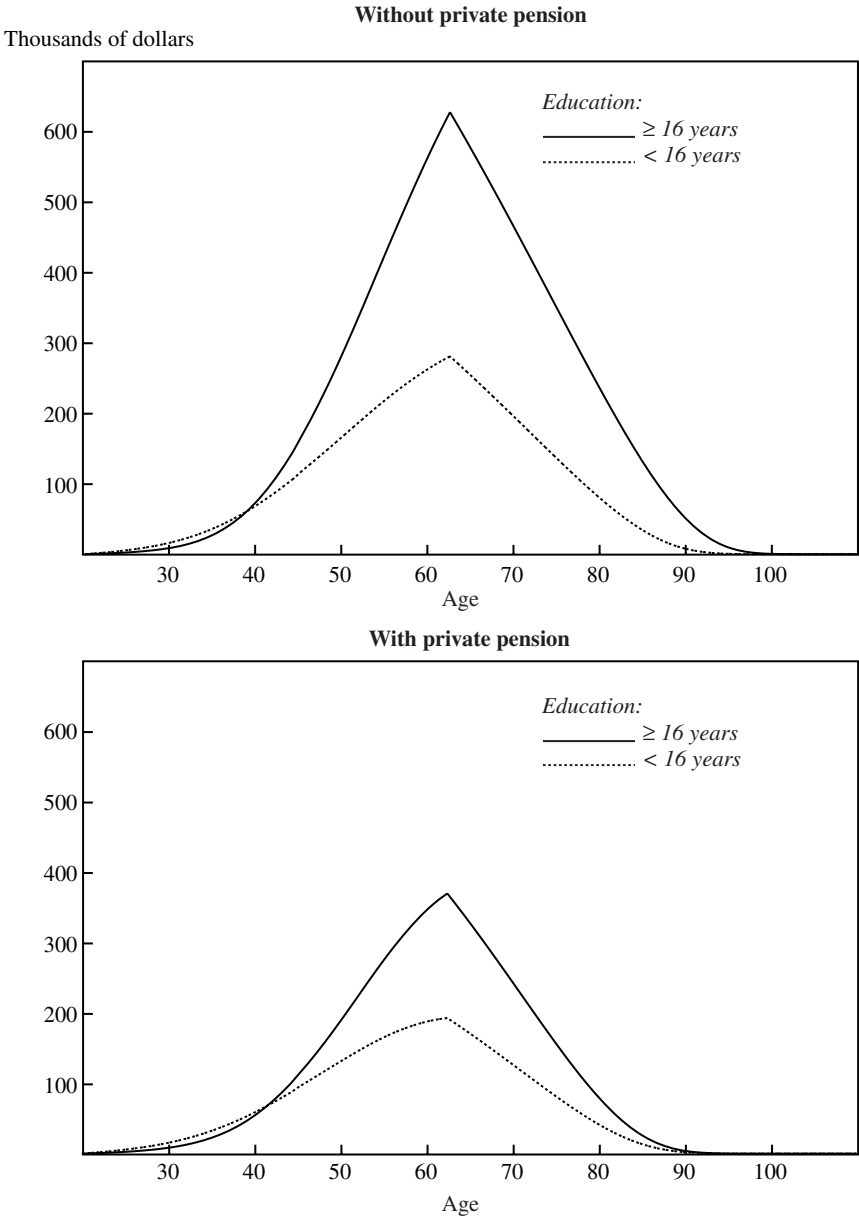
In a stochastic model, however, there is a much looser link between the slope of the age-consumption profile and the difference between the time preference rate and the interest rate. The bottom panel of figure 6 shows optimal age-consumption profiles for the same three time preference rates, but with stochastic earnings and uncertain life span. All three consumption profiles rise through the working years at roughly the same

35. See, for example, Attanasio (1993).

36. Leung (1994).

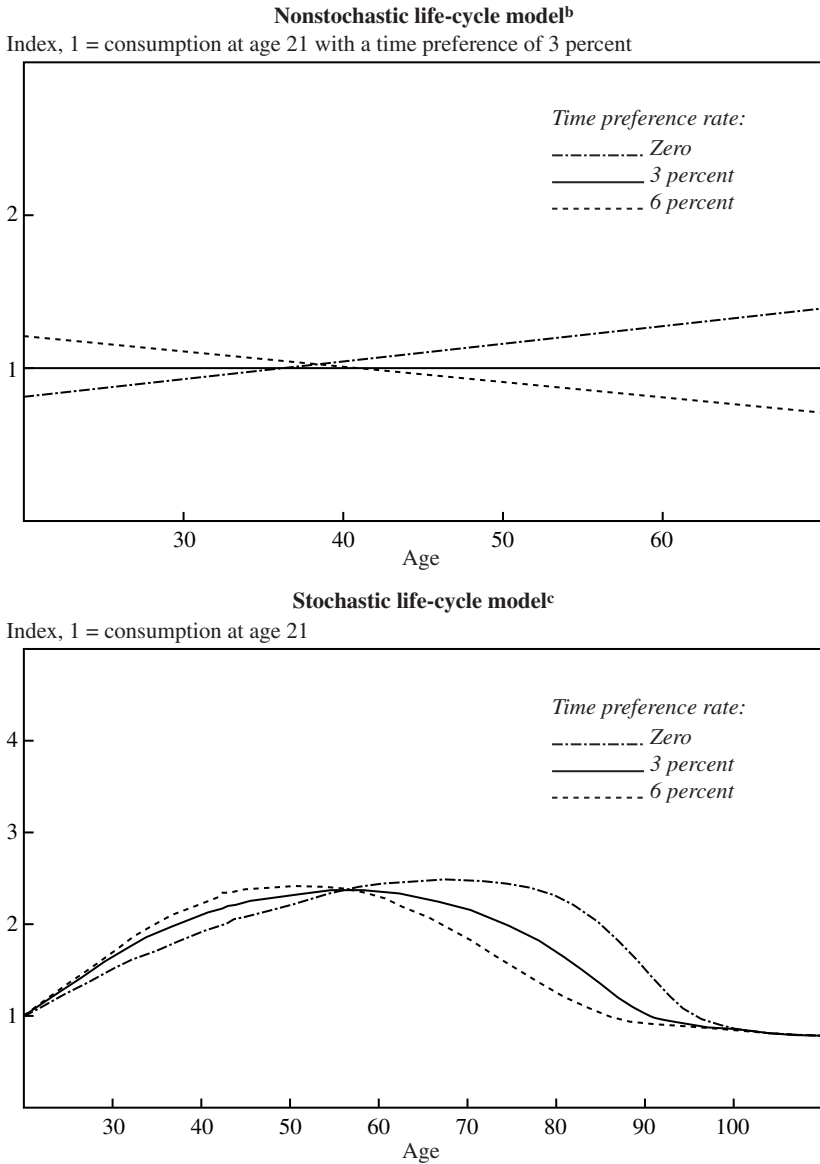
37. To be clear, these are household wealth profiles and thus are based on household consumption profiles.

**Figure 5. Optimal Age-Wealth Profiles by Pension Status with a Time Preference Rate of 3 Percent<sup>a</sup>**



Source: Authors' calculations using the stochastic life-cycle simulation model described in the text.  
 a. Wealth is nonannuity wealth for the median household.

**Figure 6. Nonstochastic and Stochastic Optimal Age-Consumption Profiles for Households with Sixteen or More Years of Education<sup>a</sup>**



Source: Authors' calculations.  
 a. A real interest rate of 3 percent is assumed.  
 b. Assumes an expected life span of seventy-four.  
 c. Assumes a maximum life span of 110 years.

rate, despite the fact that the time preference rates straddle the interest rate and differ by a total of 6 percentage points. This implies that, in the real world, it is not possible to infer a household's time preference rate (or the difference between its time preference rate and the interest rate) from its choice of an upward-sloping consumption profile over a flat or downward-sloping one.

## **Empirical Analysis**

### *Data Issues*

Our analyses use data from the 1992 Health and Retirement Survey and the 1983, 1989, 1992, and 1995 Surveys of Consumer Finances. The HRS is conducted by the Institute for Social Research at the University of Michigan. In 1992 the survey gathered data on a nationally representative sample of persons born in 1931 to 1941 and on their spouses regardless of age. Reinterviews have occurred every two years since then. The survey oversamples blacks, Hispanics, and Florida residents and contains detailed information on wealth, pensions, income, employment, demographics, and health. Our HRS sample consists of the 2,626 married households where the husband was born between 1931 and 1941 and worked at least twenty hours per week in the 1992 survey.<sup>38</sup>

The SCF is a triennial survey undertaken by the Federal Reserve Board with the cooperation of the Department of the Treasury. The survey oversamples high-income households and is designed to provide detailed information on family balance sheets, pension status, income, and demographics. We use data for married households where the husband is between the ages of twenty-five and sixty-two and works at least twenty hours per week. This generates sample sizes between 1,300 and 1,800 in each year.<sup>39</sup>

All of our results using both data sets are weighted in accordance with a nationally representative population. Sample sizes in subgroups of each data set stratified by age, education, and pension status are reported in

38. Background information on the HRS is provided in Juster and Suzman (1995).

39. For discussion of the SCF see Avery, Elliehausen, and Canner (1984a, 1984b) and Kennickell, Starr-McCluer, and Sunden (1997).

appendix table D8. It is worth keeping in mind throughout the discussion of the empirical results that some of the sample sizes are small.

**BASIC WEALTH AND EARNINGS MEASURES.** To measure earnings we use the sum of current earnings by husband and wife. Measuring wealth is more complicated. Because the simulation model accounts for both precautionary saving and saving for retirement, our empirical wealth measure needs to be broad enough to account for both. We define three measures of wealth. What we call broad wealth is essentially all net worth other than equity in vehicles. Specifically, broad wealth is the sum of equity in the primary residence, other real estate equity, equity in businesses, and net financial assets; financial assets include balances in DC plans, 401(k) plans, Individual Retirement Accounts, and Keogh plans as well as non-tax-advantaged financial assets, less consumer debt. Narrow wealth is broad wealth less all equity in the primary residence. Intermediate wealth is broad wealth less half of equity in the primary residence.

**SOCIAL SECURITY AND PENSION WEALTH.** All of the wealth measures above exclude social security. The treatment of DB pension wealth, however, differs in the two data sets. In the SCF, DB pensions are excluded from the empirical wealth calculations. Households in which at least one adult has a DB pension from his or her current job are assumed to receive pension benefits as estimated in table 1, and their wealth, excluding DB pensions and social security, is compared with the simulation benchmarks developed above for households with pension coverage. In effect, this treatment provides each household that has a DB pension from the current job with average DB pension benefits, conditional on education status, as shown in table 1.<sup>40</sup>

In the HRS data, however, estimates of expected DB pension benefits can be generated. We use the additional information provided by the DB pension wealth data in the HRS and therefore include DB pension wealth in the empirical wealth measures (see appendix C).<sup>41</sup> Thus pension wealth varies across households that have pensions, even after controlling for education status. Because we include DB pensions in wealth, we compare

40. We somewhat underestimate DB pension coverage for SCF households because households with DB plans from prior jobs but not on the current job are treated as not having DB plan coverage.

41. This approach follows that of Moore and Mitchell (1997) and Gustman and Steinmeier (1998).

the resulting wealth measures from the HRS data with the simulation's wealth benchmarks above for households that do *not* have pension coverage. We do this to avoid double counting DB pension assets.

HOUSING WEALTH. Whether it is conceptually appropriate to include housing equity in measures of the adequacy of retirement saving has generated significant debate. The controversy is heightened by the fact that housing equity constitutes a large portion of nonpension net worth for most households. Excluding housing wealth is defended on at least three grounds: some surveys suggest that people do not like to move when they are old; others indicate that people do not want to consume housing equity to finance retirement; and some evidence suggests that younger elderly households in the 1970s chose not to reduce housing equity.<sup>42</sup>

Although we do not dispute these findings, we do not believe they support the view that *all* housing wealth should be excluded from considerations of retirement income adequacy. Households can extract housing equity without moving, by means of reverse mortgages. And retirees would surely like not to have to deplete financial assets in retirement either, yet those are counted in retirement income calculations; therefore housing wealth should be counted also.

There are many reasons why housing equity should be included as retirement wealth even though younger elderly households in the 1970s chose to retain housing equity. First, it may make sense to consume housing wealth only after consuming other assets, both because housing wealth is illiquid and somewhat difficult to tap, and because it is tax-preferred. The value of a tax-preferred asset relative to a fully taxable asset typically grows as the asset is held over time, because the advantages of tax deferral cumulate. If so, studies of the younger elderly may not pick up this effect. Indeed, other studies indicate that older elderly households do eventually consume their housing wealth.<sup>43</sup> Second, in the 1970s housing was a highly profitable investment, and thus people would have tended to invest more in housing rather than disinvested. In the 1980s and 1990s, as housing has become a less attractive asset to hold for both demographic and tax reasons, people may be more willing to extract equity from their houses.

42. Venti and Wise (1990) provide evidence on whether households choose to reduce housing equity. Bernheim (1992, 1994b, 1997) summarizes the case against including housing equity in measuring the adequacy of retirement saving.

43. Sheiner and Weil (1992); Hurd (1995).

Third, recent policy changes have eliminated the taxation of the first \$500,000 of capital gains on a house. This may induce more retired people to sell their homes in the future and allow them to consume some of their housing wealth, even if they would not have done so in the past.

Fourth, the elderly in the 1970s lived through World War I, the Great Depression, and World War II, and so may have had different attitudes toward the importance of maintaining a precautionary stock of wealth. Baby boomers, in contrast, have been among the major participants in home equity lending booms in the 1980s and 1990s. Thus, determining whether the current elderly are willing to extract their housing equity does not resolve whether current generations of workers will be willing to extract their housing equity when they are old. Rather the question is whether the baby boomers will behave when they are elderly more like they themselves did when they were young, or more like the elderly do today.

Fifth and most important, housing provides consumption services and thus represents wealth. Certainly, if a household had a negative housing equity position, that fact would be relevant to its retirement income security. Consider two families, identical in every way except that one owns a \$300,000 house that is fully paid off and the other rents. Ignoring housing equity would amount to concluding that these two households are equally well prepared for retirement. But common sense as well as economic reasoning indicates that the homeowner would be in much better shape.

Finally, it makes sense from a policy perspective to consider housing wealth. A retired couple that lives in a \$300,000 house with no mortgage and has little cash or financial assets, but refuses to dip into housing equity, may not be considered to have pressing retirement needs from a social policy perspective.

For all of these reasons, we believe it is appropriate to include housing equity in retirement saving calculations. Nevertheless, it may not be appropriate to include every dollar of equity, since liquidating housing wealth through sale or reverse mortgages imposes some transactions costs. Excluding half of housing wealth—as we do with our intermediate wealth measure—to account for transactions costs certainly overestimates such costs. Therefore we believe that our intermediate and broad wealth measures generate the most reasonable empirical results. Nevertheless, we pre-

sent most of our results for all three wealth measures described above, which together bound all the possible effects of including housing equity.

**SAMPLE SELECTION BIAS.** By including only married couples that are still in the work force, we introduce a potentially important sample selection bias. As a cohort nears retirement age, it is plausible that the more prosperous households in the cohort retire earlier rather than later. Thus, as we examine progressively older households in our data, we may be examining cross sections that are progressively poorer. To the extent that this shows up as a decline in the adequacy of saving for older households, the decline would be spurious.

There is some evidence of such bias. In the HRS the overall wealth—including pensions from previous jobs—of married couples that are *not* in the sample (that is, where the husband does not work full-time) is lower than that of married couples in the sample for fifty-one- to fifty-four-year-olds and for fifty-five- to fifty-nine-year-olds. However, among sixty- and sixty-one-year-olds the relationship is reversed: households that are not in the labor force have more wealth than do those in the labor force.

Similar evidence arises in the 1992 SCF data. Among college-educated households aged fifty-five to fifty-nine and sixty to sixty-two with pensions, married couples that are *not* in the labor force have higher median broad wealth than do married couples that remain in the labor force. At younger ages the relationship is reversed. Additional supporting evidence comes from the fact that, for college-educated households without pensions in our sample, median wealth rises with age through ages fifty-five to fifty-nine and then falls for those ages sixty to sixty-two. The same pattern occurs for non-college-educated households with pensions.

All of these results suggest that a disproportionate number of wealthy households are retiring before age sixty-two. Thus any decline in the measured adequacy of saving among older working households relative to younger households may be partly spurious.

### *Basic Results Using the Health and Retirement Survey*

For a household with a given set of observable characteristics, the simulation model generates a distribution of optimal wealth-earnings ratios, rather than a single optimal level. This implies that we cannot determine precisely the optimal wealth-earnings ratio for any particular household. Instead, we compare the distributions of observed and simulated wealth-



earnings data for married households with a given set of characteristics: age, current earnings, education, and pension status. Thus our strategy for examining the adequacy of saving focuses mainly on two issues: determining the proportion of households whose wealth-earnings ratios exceed the median simulated wealth-earnings ratio for households with the same characteristics; and comparing wealth-earnings ratios at different percentiles of the actual and simulated distributions. Both approaches provide valuable information, but neither permits us to identify which particular households are saving adequately or inadequately.<sup>44</sup>

**MEDIAN WEALTH-EARNINGS RATIOS.** Table 5 reports the results of comparing each HRS household's wealth-earnings ratio with the median of the distribution of wealth-earnings ratios from the simulation for households with the same characteristics. For the full sample, the table shows that, with a time preference rate of 3 percent, 60.5 percent of households have ratios of broad wealth to earnings that exceed the median simulated wealth-earnings ratio for households with the same observable characteristics.

The *interpretation* of this result depends on the fact that the saving benchmark is derived from a stochastic rather than a nonstochastic model. In a nonstochastic model, all households of the same age, current earnings, education, and pension status would be assigned the same optimal wealth-earnings ratio, and the finding above would be interpreted as showing that 60.5 percent of households exceed the optimal ratio. That would mean that almost 40 percent of households fall short of their assigned optimal wealth-earnings ratio. This would (erroneously) suggest that a significant portion of the population is undersaving.

In contrast, once it is recognized that households face uncertainty about their future earnings, it is appropriate to use a stochastic model as the benchmark. This in turn implies that one would expect only 50 percent of households to exceed the median wealth-earnings ratio. Thus the same fact—that 60.5 percent of actual households exceed the simulated median—would instead suggest adequate, indeed somewhat more than adequate, amounts of wealth accumulation relative to the benchmark at the median of the distribution.

For reasons noted above, we believe the most reasonable estimates stem from using intermediate or broad wealth in the data and a time preference

44. The identification of exact optimal rates may be facilitated with household panel data with a history of previous earnings, an analysis we hope to undertake in future work.

**Table 5. Shares of All Households and of Selected Subgroups from the 1992 Health and Retirement Survey with Wealth-Earnings Ratios At or Above the Simulated Median**

Sample	Time preference rate = 3 percent			Time preference rate = 0		
	Narrow wealth <sup>a</sup>	Intermediate wealth <sup>b</sup>	Broad wealth <sup>c</sup>	Narrow wealth	Intermediate wealth	Broad wealth
Full sample	43.4	51.9	60.5	32.9	39.3	45.7
Households with pension coverage						
All	48.9	58.3	66.7	35.9	42.7	49.9
≥16 years of education	64.6	72.6	77.8	52.7	60.1	65.2
<16 years of education	42.7	52.7	62.3	29.3	35.8	44.0
Households without pension coverage						
All	34.8	41.9	50.8	28.2	34.0	39.1
≥16 years of education	52.6	60.4	65.4	45.7	53.1	57.9
<16 years of education	29.5	36.3	46.4	23.0	28.3	33.4
All households with ≥16 years of education	60.5	68.4	73.5	50.3	57.7	62.6
All households with <16 years of education	37.3	46.0	55.8	26.7	32.7	39.7
Age						
51–54	42.9	51.2	60.7	31.3	38.0	45.4
55–59	45.0	52.7	60.2	35.4	41.1	46.6
60–61	39.7	51.5	60.5	29.9	37.2	43.8
Earnings						
0–\$10,000	39.6	54.0	65.4	30.8	47.6	56.0
\$10,000–\$20,000	29.5	41.2	52.5	25.3	31.9	43.2
\$20,000–\$30,000	33.5	43.1	52.1	26.7	33.2	39.6
\$30,000–\$40,000	34.7	44.2	54.7	25.7	31.3	38.6
\$40,000–\$50,000	38.3	47.6	58.8	26.9	34.1	40.4
\$50,000–\$75,000	51.0	57.7	64.4	36.5	43.1	49.5
≥\$75,000	57.6	64.6	70.1	46.7	52.0	55.7

Source: Authors' calculations based on HRS data.

a. Broad wealth less all equity in the primary residence.

b. Broad wealth less half of all equity in the primary residence.

c. The sum of equity in the primary residence, other real estate equity, equity in businesses, and net financial assets.

rate of 3 percent in the simulation. Thus our central finding from table 5 is that estimates using that specification show that between 52 and 61 percent of the sample exceed the median wealth-earnings ratio. These results suggest that households are saving adequately, at least around the median of the distribution.

In addition, however, table 5 shows that the treatment of housing wealth and the choice of time preference rate can have significant effects on the results. When housing equity is excluded (that is, under the narrow wealth specification) but the time preference rate of 3 percent is retained, the proportion of households whose wealth-earnings ratios exceed the median ratio falls to 43 percent. Reducing the time preference rate to zero also produces significantly poorer results than in the base case. Nevertheless, 46 percent of households still have broad wealth-earnings ratios that exceed the median simulated ratios, and 39 percent have intermediate wealth-earnings ratios that exceed the median simulated ratio. For all combinations of the two time preference rates and the three treatments of housing, our results show that between 33 and 60 percent of all households exceed the median wealth-earnings ratio. We emphasize that all of these results should be compared against a benchmark expectation that only 50 percent of households will exceed the median.

The table shows several other interesting results as well. When education is controlled for, households with pensions appear to be saving significantly more adequately than households without pensions. Having a pension is associated with an increase of about 12 percentage points in the proportion of households that exceed the median target wealth-earnings ratio when the time preference rate is 3 percent, and an increase of about 7 percentage points when the time preference rate is zero. Alternatively, when pensions are controlled for, households with more education are saving more adequately than are households with less education. Having more education is associated with a 15- to 25-percentage-point increase in the likelihood of exceeding the simulated median wealth-earnings ratio. These qualitative results are consistent with those of numerous previous studies.<sup>45</sup> As with previous studies of the adequacy of saving,

45. See Banks, Blundell, and Tanner (1998); Bernheim (1992); Bernheim and Scholz (1993); Gale (1997); Kotlikoff, Spivak, and Summers (1982); Mitchell, Moore, and Phillips (1998); Moore and Mitchell (1997); Robb and Burbidge (1989); Warshawsky and Ameriks (1998).

we do not determine whether the results are due to the direct effects of pensions and education or to unobserved characteristics that affect household saving and are correlated with pension coverage and education.

The results do not vary significantly with respect to age. The proportion of households whose wealth-earnings ratios exceed the median simulated ratio is higher for those with current earnings below \$10,000 a year than it is for households with current earnings between \$10,000 and \$50,000 a year. It is highest for households with earnings above \$75,000 a year. This suggests that high-earnings households may have some important difference in tastes or opportunities for saving compared with others.<sup>46</sup>

DISTRIBUTION OF WEALTH-EARNINGS RATIOS. Table 6 provides evidence on the distribution of wealth-earnings ratios. The top three panels report data from the HRS. The bottom two provide simulated wealth-earnings ratios from the model, using the same distribution of households across education groups as is found in the HRS and using each of the two benchmark time preference rates.

In the simulations using a time preference rate of 3 percent, several results stand out. First, replicating the results in table 5, the median wealth-earnings ratios in the data exceeds the median in the simulation when intermediate and broad measures of wealth are used. Second, the model *underestimates* wealth-earnings ratios at the high end of the distribution. That is, there is a significant amount of real-world wealth accumulation that the model does not include. This may not be particularly surprising, because the model does not include bequest motives or the possibility of receiving a very high rate of return, perhaps on an entrepreneurial investment.

Third, at the 25th percentile the broad wealth-earnings ratio is almost exactly equal to the simulated ratio, whereas the intermediate wealth-earnings ratio falls below the simulated ratio by about 0.4 for fifty-one- to fifty-four-year-olds, by 0.7 for fifty-five- to fifty-nine-year-olds, and by 0.8 for sixty- to sixty-one-year-olds. Thus there is some evidence of a shortfall at the 25th percentile. At the 5th percentile, actual wealth-earnings ratios are far below the simulated optimal ratios. For example, among fifty-five-

46. For further evidence on these issues see Carroll (2000); Dynan, Skinner, and Zeldes (1996); Gentry and Hubbard (1998).

**Table 6. Distribution of Actual Wealth-Earnings Ratios of Households in the 1992 Health and Retirement Survey and Simulated Ratios<sup>a</sup>**

<i>Wealth measure<sup>b</sup></i>	<i>Age</i>	<i>5th percentile</i>	<i>25th percentile</i>	<i>Median</i>	<i>75th percentile</i>	<i>95th percentile</i>
<i>Actual HRS</i>						
Narrow	51–54	–0.04	0.83	2.14	4.93	13.71
	55–59	0.00	1.20	3.11	7.17	20.62
	60–61	0.06	1.67	3.63	7.42	20.04
	All <sup>c</sup>	0.00	0.99	2.75	6.29	17.41
Intermediate	51–54	0.08	1.34	2.83	5.95	15.19
	55–59	0.19	1.83	3.95	8.03	21.88
	60–61	0.49	2.51	4.73	8.58	21.77
	All	0.17	1.65	3.59	7.29	19.50
Broad	51–54	0.09	1.79	3.63	6.96	17.07
	55–59	0.20	2.50	4.93	8.97	23.54
	60–61	0.69	3.17	5.80	9.84	23.50
	All	0.23	2.19	4.40	8.37	21.94
<i>Simulation with time preference rate = 3 percent</i>						
	51–54	0.63	1.75	2.73	3.97	6.25
	55–59	1.21	2.55	3.74	5.22	7.89
	60–61	1.77	3.27	4.63	6.25	9.13
	All	0.96	2.28	3.49	5.03	7.78
<i>Simulation with time preference rate = 0</i>						
	51–54	1.50	2.84	4.04	5.47	7.99
	55–59	2.24	3.82	5.21	6.90	9.76
	60–61	2.94	4.68	6.22	8.05	11.08
	All	1.92	3.50	4.92	6.68	9.65

Source: Authors' calculations based on HRS data.

a. Simulated ratios are weighted to reflect the distribution of households across education groups. As noted in the text, because DB pensions are included in the empirical HRS wealth measure, the simulated ratios are based on the assumption that no households have pension coverage.

b. Wealth measures are as defined in table 5.

c. All households aged fifty-one to sixty-one.

to fifty-nine-year-olds, the simulated optimal ratio is 1.2, compared with a ratio of broad or intermediate wealth to earnings of 0.2.

In summary, we view the results as showing that, with a time preference rate of 3 percent in the simulation, actual wealth-earnings ratios unambiguously fall below the simulated ratios somewhere in the bottom 25 percent of the distribution. This result is consistent with systematic undersaving in this portion of the sample. It is also consistent, however, with other explanations that the model does not take into account. In particular, the model omits any sort of government-provided consumption

floor.<sup>47</sup> We return to this issue when we discuss sensitivity analysis and extensions below.

Naturally, the results are significantly less encouraging when the data are compared with the simulation results that use a time preference rate of zero. As table 5 shows, a majority of the full sample have wealth-earning ratios below the simulated median. The difference becomes significantly larger at the 25th and the 5th percentiles (table 6). If one accepts a zero time preference rate as the correct value, the data suggest that a significant portion of the population is undersaving by substantial amounts.

CHARACTERISTICS AND CORRELATES OF HIGH SAVERS. Table 7 shows the characteristics of “high savers” and “low savers.” We define high savers as households whose intermediate wealth-earnings ratios exceed the median ratio from the simulation using the 3 percent time preference rate, and low savers as those below the median.<sup>48</sup> Before examining these characteristics, it is worth emphasizing that these designations may be misleading. Because the optimal wealth-earnings ratio varies among observationally equivalent households, there is no way to determine, with the current data, whether any particular household is actually saving more than it needs for retirement. It could be that, given its earnings history, the household has an optimal wealth-earnings ratio that is higher than its actual ratio, even though its actual ratio exceeds the median ratio for households with its characteristics. Nevertheless, the typical determinants of households above and below the median target are of interest.

Table 7 shows that the average high saver household has more wealth and higher wages than the average low saver. High savers are more likely to have received an inheritance, and among those who have received an inheritance, theirs tend to be larger. High savers also have fewer children living at home; they are more likely to be self-employed, to be college graduates, and to have pension coverage; and they are less likely to be nonwhite or Hispanic. High savers are also less likely to smoke and more likely to say they have thought about retirement “a lot,” and they have slightly longer financial horizons. They are more likely to believe they will live to age seventy-five, and they expect to retire earlier than low savers.

47. Hubbard, Skinner, and Zeldes (1995).

48. Similar qualitative patterns emerge if we use definitions of high savers based on broad or narrow wealth and/or the simulation model with a time preference rate of zero.

**Table 7. Characteristics of High and Low Savers in the 1992 Health and Retirement Survey<sup>a</sup>**

Percent of all respondents except where stated otherwise

<i>Characteristic</i>	<i>Low savers</i>	<i>High savers</i>
Narrow wealth (dollars) <sup>b</sup>	41,933	321,294
Broad wealth (dollars) <sup>b</sup>	88,350	420,598
Age of household head (years) <sup>b</sup>	55	56
Combined wages of head and spouse (dollars) <sup>b</sup>	41,697	51,172
No. of children living at home	0.849	0.774
Household head has $\geq 16$ years of education	17.2	34.6
Household head is self-employed	20.9	28.4
Household head has pension coverage	52.6	68.3
Either spouse is nonwhite	12.9	5.9
Either spouse is Hispanic	8.2	3.2
Husband is in fair or poor health	12.7	8.5
Wife is in fair or poor health	15.4	10.2
Husband smokes	30.9	17.7
Wife smokes	25.2	17.4
Husband's relative mortality optimism (age 75)	-0.020	0.031
Husband's mortality optimism index missing	0.102	0.098
Husband certain he will not attain age 75	6.3	3.0
Wife's relative mortality optimism (age 75)	-0.120	-0.065
Wife's mortality optimism index missing	0.034	0.032
Wife certain she will not attain age 75	4.8	2.2
Expected retirement age (years)	63	62
Expect never to retire	13.9	10.4
Thought about retirement		
Hardly at all	21.0	12.4
Little	12.7	10.8
Some	21.2	26.9
A lot	21.9	30.6
No answer	23.3	19.3
Financial horizon		
<1 year	12.6	7.1
1-5 years	38.3	39.9
5-10 years	32.4	32.8
$\geq 10$ years	6.0	9.9
No answer	10.7	10.3
Risk aversion		
Level 1	12.1	9.3
Level 2	10.9	8.2
Level 3	10.1	11.9
Level 4 (most risk averse)	66.8	70.6
Received inheritance	16.6	25.8
Value of inheritance, given receipt (dollars) <sup>b</sup>	9,000	20,000

Source: Author's calculations based on 1992 HRS data.

a. A high saver is defined as a household whose intermediate wealth-earnings ratio exceeds the median simulated ratio for households with the same characteristics, when the simulation model uses a time preference rate of 3 percent.

b. Values are medians for households with the stated characteristic; values for other characteristics are means. Narrow and broad wealth are as defined in table 5.

c. The mortality optimism index is the difference between the respondent's subjective expectation of life expectancy and an objective measure of that respondent's life expectancy, as a percentage of the latter.

Table 8 presents estimates of two probit models of whether a household is a high saver. Model 1 contains basic household demographic and earnings variables. Model 2 adds indicators for health, inheritances, retirement, and preferences. The demographic variables largely have effects similar to previous estimates. Households with college degrees have a higher likelihood of being a high saver, by 14 percentage points. Households that have pensions or are self-employed have higher likelihoods, by about 20 percentage points. Nonwhites and Hispanics have lower likelihoods of being high savers, by 15 and 9 percentage points, respectively. The only slightly anomalous finding involves income. The likelihood of being a high saver is lower for households with higher incomes, when other factors are controlled for. There is no apparent pattern with respect to age.

The added indicators have plausible signs as well. The likelihood of being a high saver is low for smokers, perhaps because of a higher time preference rate for those households. The likelihood rises with declines in expected retirement age, with the extent to which the household has thought about retirement, and with the household's financial horizon. It is also higher for households who have contacted the Social Security Administration to find out about their benefits, and for households who have received a large inheritance.

### *Basic Results Using the Survey of Consumer Finances*

Empirical analysis with the SCF data allows consideration of a number of additional items, because the surveys span a wider age group and a longer period of time than does the 1992 HRS. For comparability with the HRS results, we focus first on the 1992 SCF and then examine results from several years of data.

**MEDIAN WEALTH-EARNINGS RATIOS.** Table 9 uses the 1992 SCF sample of married households where the husband is between the ages of twenty-five and sixty-two and works full-time. The aggregate results are somewhat more favorable than the HRS results. Using a time preference rate of 3 percent, the proportions of households exceeding the simulated median wealth-earnings ratio are 66 percent, 60 percent, and 47 percent for measures using broad, intermediate, and narrow wealth, respectively. Using a time preference rate of zero, the analogous figures are 54 percent, 46 percent, and 36 percent. Thus between 35 and 66 percent of households exceed the median wealth-earnings ratio, depending on the



**Table 8. Probit Regression Results for Households in the 1992 Health and Retirement Survey<sup>a</sup>**

<i>Variable</i>	<i>Model 1</i>		<i>Model 2</i>	
	<i>Coefficient estimate</i>	<i>Marginal probability</i>	<i>Coefficient estimate</i>	<i>Marginal probability</i>
Age				
55–59	0.044	0.018	–0.030	–0.012
60–61	0.033	0.013	–0.017	–0.007
Household income				
\$20,000–\$30,000	–0.176	–0.070	–0.277**	–0.110
\$30,000–\$40,000	–0.214**	–0.085	–0.466***	–0.182
\$40,000–\$50,000	–0.214**	–0.085	–0.443***	–0.173
\$50,000–\$75,000	–0.035	–0.014	–0.279**	–0.111
≥\$75,000	–0.024	–0.010	–0.324**	–0.128
Children living at home				
1	–0.046	–0.018	–0.053	–0.021
2	–0.085	–0.034	–0.042	–0.017
≥3	–0.254**	–0.100	–0.156	–0.062
Head of household has	0.484***	0.190	0.358***	0.141
≥16 years of education				
Either spouse is self-employed	0.472***	0.186	0.554***	0.216
Either spouse has pension coverage	0.554***	0.218	0.527***	0.208
Household head is nonwhite	–0.488***	–0.189	–0.390***	–0.153
Household head is Hispanic	–0.288***	–0.113	–0.238**	–0.094
Husband in fair or poor health			–0.093	–0.037
Wife in fair or poor health			–0.111	–0.044
Husband smokes			–0.154**	–0.061
Wife smokes			–0.283***	–0.112
Husband's relative mortality optimism (age 75)			0.030	0.012
Husband's mortality optimism index missing			–0.056	–0.022
Husband certain he will not reach 75			–0.302*	–0.119
Wife's relative mortality optimism (age 75)			0.132	0.053
Wife's mortality optimism index missing			–0.056	–0.022
Wife certain she will not reach 75			–0.135	–0.054
Expected retirement age				
56–59			0.397**	0.155
60–61			0.437***	0.170
62			–0.007	–0.003
63–64			–0.061	–0.024
65			–0.429***	–0.168

*(continued next page)*

**Table 8. Probit Regression Results for Households in the 1992 Health and Retirement Survey<sup>a</sup> (continued)**

Variable	Model 2	
	Coefficient estimate	Marginal probability
Expected retirement age ( <i>continued</i> )		
66–69	–0.376***	–0.147
70	–0.338**	–0.133
Will never retire	0.204	0.081
Thought about retirement		
A little	0.124	0.050
Some	0.209**	0.083
A lot	0.290***	0.115
Financial horizon (years)		
1–5	0.226**	0.090
5–10	0.177*	0.071
≥10	0.386***	0.151
Risk aversion		
Level 2	–0.042	–0.017
Level 3	0.204*	0.081
Level 4 (most averse)	0.081	0.032
Ever contacted Social Security	0.169***	0.067
Inheritance		
<\$5,000	–0.284*	–0.112
\$5,000–\$10,000	–0.124	–0.049
\$10,000–\$25,000	0.053	0.021
\$25,000–\$100,000	0.461***	0.179
≥\$100,000	0.840***	0.303
Constant	–0.365***	–0.336*
<i>N</i>	2,626	2,378

Source: Authors' calculations.

\*\*\*  $p < 0.01$ ; \*\*  $p < 0.05$ ; \*  $p < 0.10$ .

a. The dependent variable takes a value of one if the household's actual wealth-earnings ratio exceeds the median simulated wealth-earnings ratio, and zero otherwise.

time preference rate and the treatment of housing. Although the two surveys correspond to the same time period, comparing the aggregate SCF and HRS results is a little misleading, because the two data sets span different age groups.

Table 9 also shows that a higher proportion of younger than of older SCF households exceed the median wealth-earnings ratios. For fifty- to sixty-two-year-olds, the SCF data generate about the same results as do the

**Table 9. Shares of All Households and of Selected Subgroups from the 1992 Survey of Current Finances with Wealth-Earnings Ratios At or Above the Simulated Median**

Sample	Time preference rate = 3 percent			Time preference rate = 0		
	Narrow wealth <sup>a</sup>	Intermediate wealth	Broad wealth	Narrow wealth <sup>a</sup>	Intermediate wealth	Broad wealth
Full sample	47.0	59.7	66.4	35.5	46.0	54.0
Households with pension coverage						
All	49.0	64.2	72.3	35.5	47.0	57.3
≥16 years of education	64.9	78.5	85.2	52.6	63.4	74.2
<16 years of education	37.3	53.6	62.8	22.9	35.0	44.9
Households without pension coverage						
All	45.4	56.2	61.7	35.4	45.1	51.3
≥16 years of education	69.8	77.7	80.6	58.6	67.0	72.0
<16 years of education	32.3	44.7	51.6	23.1	33.4	40.2
All households with ≥16 years of education	67.4	78.1	82.9	55.6	65.2	73.1
All households with <16 years of education	34.4	48.4	56.3	23.0	34.1	42.2
Age						
25-29	50.6	62.2	62.2	41.4	55.3	56.4
30-34	57.3	73.0	73.0	47.1	59.8	66.6
35-39	51.6	68.2	73.6	36.1	50.3	60.8
40-44	43.1	53.6	68.2	31.8	43.0	50.7
45-49	38.1	49.1	57.4	28.1	34.8	44.6
50-54	39.9	50.0	57.5	27.3	33.6	43.3
55-59	43.9	57.5	67.8	34.6	41.3	49.4
60-62	42.6	50.4	57.3	26.7	31.2	46.2
Earnings						
0-\$10,000	38.3	51.0	52.8	34.6	45.7	52.3
\$10,000-\$20,000	27.0	40.0	47.9	23.6	33.3	37.0
\$20,000-\$30,000	40.1	54.7	58.8	31.3	42.3	49.2
\$30,000-\$40,000	46.7	61.9	72.5	31.7	47.8	54.3
\$40,000-\$50,000	45.1	58.9	67.9	33.0	45.4	54.3
\$50,000-\$75,000	53.3	64.2	69.4	38.7	46.2	56.9
≥\$75,000	62.0	72.4	78.3	48.5	56.2	65.6

Source: Authors' calculations based on SCF data.  
a. Measures of wealth are as defined in table 5.

HRS data (table 5). In the SCF between 57 and 68 percent of households in the age groups over fifty had broad wealth–earnings ratios exceeding the median simulated ratio, using a time preference rate of 3 percent, compared with about 60 percent in the HRS data. Between 43 and 49 percent of households in the same age group had broad wealth–earnings ratios exceeding the median simulated ratio with a time preference rate of zero, compared with a range of 43 to 47 percent in the HRS data.

Like the HRS data, the SCF data show higher proportions of households with pension coverage and of households with more education exceeding their simulated median wealth-earnings ratios. The SCF data, like the HRS data, also show that the proportion of households that exceed the simulated median ratio is higher for households with very low earnings (less than \$10,000 a year) than for households with low to moderate earnings (\$10,000 to \$40,000). The proportion exceeding the median wealth-earnings ratio is also higher for the households in the highest earnings categories than in slightly lower earnings categories.

**DISTRIBUTION OF WEALTH-EARNINGS RATIOS.** Table 10 reports the distribution of wealth-earnings ratios among fifty-one- to sixty-one-year-olds in the 1992 SCF and in the simulation model. The model results are generated by creating an artificial sample with the same proportion of households by education and pension status that is found among fifty-one- to sixty-one-year-olds in the 1992 SCF.<sup>49</sup> The SCF data show that, with a time preference rate of 3 percent, the simulation model understates actual saving at the high end of the wealth-earnings distribution and (for the broad wealth definition) at the median, which is consistent with the HRS results.

The SCF data show strong wealth accumulation at the bottom of the wealth-earnings distribution. The broad wealth–earnings ratio in the SCF exceeds the simulated ratio at the 25th percentile and at the 5th percentile of the distribution for fifty-five- to fifty-nine-year-olds and for sixty- and sixty-one-year-olds. The intermediate wealth–earnings ratio in the SCF

49. The weighted simulation results in table 10 are not strictly comparable to those in table 6. Table 10 uses weights based on households' college and DB pension status, because in the SCF data we do not include DB pension wealth in households' wealth calculation. The simulation results in table 6 assume that all households do not have a pension—because in the HRS data, DB pension wealth is included as a component of measured wealth—and simply weight between the proportion of households with a college education and those without.

**Table 10. Distribution of Actual Wealth-Earnings Ratios of Households in the 1992 Survey of Current Finances and Simulated Ratios<sup>a</sup>**

<i>Wealth measure<sup>b</sup></i>	<i>Age</i>	<i>5th percentile</i>	<i>25th percentile</i>	<i>Median</i>	<i>75th percentile</i>	<i>95th percentile</i>
<i>Actual SCF</i>						
Narrow	51–54	–0.12	0.14	1.33	4.12	11.94
	55–59	0.00	0.77	2.62	5.89	20.45
	60–61	0.33	1.19	2.64	6.88	35.79
	All <sup>c</sup>	–0.07	0.55	1.93	5.33	20.46
Intermediate	51–54	–0.07	0.51	1.87	5.05	14.13
	55–59	0.62	1.54	3.30	6.69	24.42
	60–61	0.98	1.96	3.83	9.05	38.78
	All	0.05	1.28	2.78	6.44	21.64
Broad	51–54	–0.04	0.78	2.38	6.00	16.69
	55–59	0.99	2.52	4.51	7.62	27.03
	60–61	1.22	2.73	4.82	10.09	43.10
	All	0.07	1.81	3.53	7.57	23.97
<i>Simulation with time preference rate = 3 percent</i>						
	51–54	0.42	1.37	2.30	3.50	5.75
	55–59	0.69	1.97	3.13	4.58	7.24
	60–61	1.08	2.50	3.85	5.46	8.33
	All	0.62	1.77	2.92	4.39	7.10
<i>Simulation with time preference rate = 0</i>						
	51–54	1.11	2.42	3.59	5.01	7.52
	55–59	1.67	3.25	4.63	6.31	9.20
	60–61	2.18	3.96	5.50	7.33	10.43
	All	1.44	2.98	4.37	6.08	9.06

Source: Authors' calculations based on SCF data.

a. Simulated ratios are weighted to reflect the distribution of households across education groups and pension status.

b. Wealth measures are as defined in table 5.

c. All households aged fifty-one to sixty-one.

falls below the simulated ratio by about 0.5 in these same age groups at the 25th percentile, but is very close to the simulated ratio at the 5th percentile. In summary, table 10 provides some evidence of undersaving in the SCF data, but the evidence is hardly conclusive.

**CHARACTERISTICS OF HIGH AND LOW SAVERS.** In further analysis, the results of which are not shown, we also examine the characteristics of SCF households that were high savers and low savers, as defined above. The data patterns generally parallel the results from the HRS. High savers had more wealth, income, education, and pension coverage than did low savers. In addition, 25 percent of high savers expected to receive an inheritance, compared with only 16 percent of low savers.

CHANGES IN WEALTH-EARNINGS RATIOS OVER TIME. Table 11 reports the proportion of SCF households whose wealth-earnings ratios exceeded the simulated median ratio over time. Using the intermediate wealth measure and a time preference rate of 3 percent, the proportion of households who exceeded the median simulated wealth-earnings ratio for households with their characteristics fell slightly from 1983 to 1995 but remained at 58 percent or more in each sample year. Using the broad wealth measure, the proportion of households that exceeded the median simulated wealth-earnings ratio fell from 71 percent in 1983 to 66 percent in 1992 and remained at that level in 1995. With a time preference rate of zero, fewer households surpassed the median simulated wealth-earnings ratios. Nevertheless, between 44 and 47 percent of households had intermediate wealth above the median simulated ratio, and over 50 percent had broad wealth that exceeded the median ratio. Notably, narrow wealth-earnings ratios rose over the same period, presumably in part because of the large buildup of financial assets during this period.<sup>50</sup> These results indicate that different measures of the adequacy of saving can move in different directions over the same time period.

The table also shows how different cohorts have fared over time relative to their median simulated wealth-earnings ratios. Wealth accumulation for younger baby boomers (those born between 1956 and 1964) improved relative to the simulated medians over the 1989–95 period, as they aged from a range of twenty-five to thirty-three years to thirty-one to thirty-nine years. Depending on whether a time preference rate of 3 percent or zero is used, and whether the wealth measure is intermediate or broad, between 50 and 76 percent of younger boomers exceeded the simulated median wealth ratios in 1995.

For older boomers (those born between 1946 and 1955), wealth accumulation has not increased as much as the optimal simulated ratios have, but overall wealth accumulation is still fairly high. Using a time preference rate of 3 percent, in 1995, when the older boomers were between forty and forty-nine years old, between 55 and 65 percent of them exceeded the median wealth-earnings ratios, depending on whether the intermediate or the broad measure of wealth is used.

Wealth accumulation among the younger HRS cohort (those born between 1937 and 1941) fell dramatically relative to the simulated median

50. See Engen, Gale, and Scholz (1996).

**Table 11. Shares of Households of Different Cohorts with Wealth-Earnings Ratios At or Above the Simulated Median, 1983-95**

Sample	Year	Age	Time preference rate = 3 percent			Time preference rate = 0		
			Narrow wealth <sup>a</sup>	Intermediate wealth	Broad wealth	Narrow wealth <sup>a</sup>	Intermediate wealth	Broad wealth
All households	1983	25-62	42.9	61.7	71.0	32.9	46.9	57.5
	1989	25-62	44.4	62.3	69.3	32.4	47.0	57.6
	1992	25-62	47.0	59.7	66.4	35.5	46.0	54.0
	1995	25-62	46.3	58.4	66.3	34.7	43.9	52.0
Younger boomers (born 1956-64)	1983	19-27	...	...	...	...	...	...
	1989	25-33	48.3	63.6	67.0	38.6	54.3	59.9
	1992	28-36	54.5	68.8	71.3	44.0	56.6	62.9
	1995	31-39	53.8	69.1	75.6	38.7	50.8	61.1
Older boomers (born 1946-55)	1983	28-37	53.8	72.9	75.9	43.2	61.1	69.1
	1989	34-43	50.3	72.5	78.7	36.0	51.9	64.5
	1992	37-46	44.4	58.6	68.9	31.4	43.3	52.9
	1995	40-49	43.3	55.6	65.1	30.4	38.6	47.4
Younger HRS cohort (born 1937-41)	1983	42-46	41.3	61.7	80.0	32.8	44.8	62.0
	1989	48-52	35.9	51.7	63.9	22.7	34.9	51.2
	1992	51-55	41.0	49.1	55.1	27.4	34.3	40.9
	1995	54-58	28.1	36.0	43.7	18.8	26.1	30.0
Older HRS cohort (born 1931-36)	1983	47-52	34.3	49.8	66.1	20.7	30.6	44.3
	1989	53-58	26.0	38.5	50.0	18.7	24.7	37.9
	1992	56-61	45.5	57.7	67.3	35.9	41.8	52.4
	1995	59-64	...	...	...	...	...	...

Source: Authors' calculations based on data from the SCF.

a. Wealth measures are defined as in table 5.

wealth-earnings ratios between 1983 and 1995. In contrast, for the older HRS cohort (those born between 1931 and 1936), wealth accumulation in 1992 was quite strong relative to earlier years and relative to the simulated ratios. In 1992 between 58 and 67 percent exceeded the median ratios using the intermediate or the broad wealth definition and a 3 percent time preference rate. Using a zero percent time preference rate and the same wealth data, between 42 and 52 percent exceeded the median.

These results show that trends in wealth accumulation can vary significantly across cohorts. In two of the four cohorts, the proportion of households that exceeded the median wealth-earnings ratios declined. But in the other two that proportion rose, and for the one cohort that actually reached retirement age in the sample years—the older HRS cohort—observed wealth was higher in 1992 relative to the benchmarks than in either of the earlier years. Finally, it is worth noting that at least some of the decline as cohorts near retirement age is probably owing to the sample selection biases noted earlier.

### **Sensitivity Analysis and Extensions**

In this section we examine the sensitivity of our findings to variations in model parameter values, retirement wealth estimates, and consumption needs (table 12). Appendix table D9 shows how selected median simulated wealth-earnings ratios change across the various sensitivity analyses. We also explore a number of possible extensions of the underlying model that might influence the findings.

#### *Sensitivity Analysis*

We focus on our preferred specification: simulation results using a 3 percent time preference rate and intermediate and broad definitions of wealth. For comparison purposes, the first line of table 12 reports base-case results derived earlier using these specifications, showing the percentage of households whose wealth-earnings ratios exceed the median simulated wealth-earnings ratio for households with the same characteristics.

**BASIC MODEL PARAMETERS.** Variations in the intertemporal elasticity of substitution (IES) between 0.25 and 0.50 change the proportion of households exceeding the median wealth-earnings target by between 4 and 8



**Table 12. Sensitivity Analysis: Shares of Households with Wealth-Earnings Ratios At or Above the Simulated Median Wealth-Earnings Ratio Under Alternative Scenarios**  
Percent

<i>Case</i>	<i>1992 HRS</i>		<i>1992 SCF</i>	
	<i>Intermediate wealth</i>	<i>Broad wealth</i>	<i>Intermediate wealth</i>	<i>Broad wealth</i>
Base case <sup>a</sup>	51.9	60.5	59.7	66.4
Changes to base-case parameters <sup>b</sup>				
Time preference rate = 0	39.3	45.7	46.0	54.0
IES = 0.5	59.2	68.5	65.4	72.0
IES = 0.25	47.9	56.2	54.5	62.1
Persistence parameter = 0.99	40.9	47.6	28.8	35.2
After-tax return = 5 percent	50.0	58.2	55.0	62.1
After-tax return = 1 percent	60.3	69.0	65.6	72.5
Changes to wealth measures <sup>b</sup>				
Exclude business wealth	48.3	57.3	56.3	63.9
40 percent decline in stock market	49.6	58.8	58.2	65.6
30 percent cut in social security benefits	46.9	55.0	57.3	64.3
Add expected inheritances	—	—	62.5	68.8
Retire at age 65	57.0	65.9	61.5	68.5
Social security and pensions as function of household's final earnings	49.4	57.4	47.9	56.4
Changes to consumption needs <sup>b</sup>				
20 percent increase in all simulated wealth-earnings ratios	45.1	53.0	55.7	63.4
10 percent increase in survival rates	42.3	50.2	53.6	60.6
Substitution of Bernheim-Scholz targets <sup>b</sup>	66.2	73.6	62.5	70.6

Source: Authors' calculations based on HRS and SCF data.

a. The parameters of the base case are as follows: time preference rate = 0.03, intertemporal elasticity of substitution (IES) = 0.33, autoregressive persistence parameter = 0.85, retirement at age sixty-two, a real after-tax rate of return of 3 percent, and social security and pension income derived from the average final earnings of one's own education class.

b. Specifications of the sensitivity analysis are described in the text. Wealth measures are as defined in table 5.

percentage points (table 12). Raising the real after-tax rate of return on assets to 5 percent reduces the proportion slightly, but reducing it to 1 percent raises the proportion by 9 percentage points in the HRS data and 6 percentage points in the SCF data. To increase the persistence of a given earnings shock, we increase the first-order autoregression coefficient to 0.99 (a coefficient of one would correspond to a permanent shock). More-persistent shocks greatly increase the need for precautionary saving, especially early in the life cycle. Thus the adequacy of saving declines significantly in the HRS data, which cover households between the ages of

fifty-one and sixty-one, but declines by even larger amounts in the SCF data, which focus more heavily on younger households. These results suggest that the basic findings are sensitive to the appropriate specification of intertemporal substitution, the rate of return, and earnings shocks. However, for most of these specifications, half or close to half of all households are still above the simulated median wealth-earnings ratios.

WEALTH MEASURES. The next panel of table 12 explores the impact of changing the definition or the amount of wealth. For example, equity in a business may reflect human capital that is specific to the owner. Households may be unable to cash in such wealth to finance retirement. Excluding all business wealth from the estimates, however, does not change the results for households at the median benchmark very much. Only about one-quarter of the HRS households are self-employed, and they tend to have significantly higher wealth-earnings ratios than average, as shown in table 8.

To simulate the effects of a substantial decline in the stock market, we reduce each household's actual wealth by 40 percent of its stock and mutual fund holdings and, on the assumption that retirement funds are divided equally between stocks and other assets, by 20 percent of balances in DC pensions, Individual Retirement Accounts, Keogh plans, and 401(k) plans. This has a very small impact on the results for the median household, presumably because stock holdings are concentrated among the wealthiest families.<sup>51</sup>

To simulate possible changes to social security, we reduce benefits by 30 percent, an amount sufficient to restore long-term balance to the social security system.<sup>52</sup> In this scenario the proportion of households whose broad wealth exceeds the median target falls by 5.5 percentage points in the HRS sample and by 2.1 percentage points in the SCF sample. Given the centrality of social security in the retirement income of many elderly households,<sup>53</sup> these effects seem small. However, it is likely that the effects are larger at the lower end of the wealth-earnings distribution.

51. Since 1992, stocks have risen significantly in value, which would make the impact of a similar decline today even smaller. Stock ownership is also more widespread now than it was in 1992, and this could increase the impact of a crash on the results reported here, but stock holdings for the median household are generally small (Kennickell, Starr-McCluer, and Sunden, 1997).

52. U.S. Board of Trustees of the Federal OASDI and Disability Insurance Trust Funds (1999).

53. Aaron and Reischauer (1998).

Using the SCF data, we assign to each household that expects an inheritance a current wealth increment, such that if that increment grew at a real rate of 3 percent per year, the resulting balance would be \$25,000 at age sixty-five. This addition to wealth has a small effect on the overall results, raising the proportion of households who exceed the median wealth-earnings target by 2 to 3 percentage points.

Raising the predetermined retirement age to sixty-five increases lifetime earnings in the simulations and raises the proportion of households who exceed the median wealth-earnings ratio by about 5 percentage points in the HRS data. In the SCF data, most of the households are younger than fifty, so the impact is smaller.<sup>54</sup>

Lastly, the base-case model specifies that each household receives social security and pension income based on the *average* final earnings of its education class—taken from the mean age-earnings profile—rather than based on the household’s actual wage profile. This implies that some households with very low earnings relative to their education class would have a very high actual replacement rate and therefore need to save very little. Likewise, households with very high earnings would have very low replacement rates relative to their actual earnings and thus must save more. This generates (inappropriate) variance in wealth and possibly in wealth-earnings ratios as well. To test the importance of this effect, we redesigned the model to allow social security and pension income to depend on each household’s actual final earnings. In effect, this *overstates* the uncertainty of pension and social security benefits, because real-world social security benefits are a function of years of coverage and lifetime average wages, and real-world benefits from DB plans are often a function of several years of earnings and years of coverage.

This change increases optimal precautionary saving by substantial amounts for younger households but only by small amounts for older households. As a result, the effect of this change was relatively small on HRS households in their fifties. However, among SCF households aged twenty-five to sixty-two, the change reduced the proportion exceeding the median wealth-earnings target by 10 to 12 percentage points.

54. To analyze retirement at age sixty-five, we raise the pension and social security replacement rates by 10 percent and allow earnings to continue between ages sixty-two and sixty-five according to the same age-earnings profile and the same stochastic process used in the rest of the analysis.

CONSUMPTION NEEDS. As a further sensitivity test, we raised all simulated wealth-earnings ratios by 20 percent. This scenario could cover a number of possibilities. For example, if health care accounts for 10 percent of household expenditure before retirement, this amounts to tripling health expenditure in retirement; if preretirement health expenditure is 20 percent, it represents a doubling. Likewise, raising the simulated wealth-earnings ratios could be a rough way to proxy for uncertainty regarding health expenses or income in retirement. Although this may not be a worst-case scenario, a 20 percent increase does reduce the proportion of households who exceed the simulated median wealth-earnings ratios by 7 percentage points in the HRS data and by less in the SCF data.<sup>55</sup>

Some observers believe that the assumptions used in the social security trustees' forecasts and in this paper systematically understate the typical future life span.<sup>56</sup> To account for this possibility, we raise survival rates by 10 percent. This increases life expectancy at birth by about 7.5 years for men and 8 years for women. The resulting survival rates are higher than those in the Social Security Administration's high-cost scenario. This change has a significant impact on the results, reducing the proportion of households that exceed the median saving benchmark by 10 percentage points in the HRS data, and by 6 percentage points in the SCF data.

BERNHEIM-SCHOLZ TARGETS. Lastly, we examine the impact of using the simulation benchmarks of Bernheim and Scholz.<sup>57</sup> Both the HRS and the SCF data show very high levels of saving using these targets. Two-thirds of households in the HRS data and over 60 percent of those in the SCF data have intermediate wealth-earnings ratios that exceed the wealth-earnings benchmark. These findings lend support to the view that,

55. Fuchs (1998a) cites data showing that health expenditure per capita for persons over age sixty-five is more than three times greater than that before age sixty-five, but this includes government-provided care as well as out-of-pocket expense. Fuchs (1998b) notes that if health expenditures continue to grow at the same rate as they have in the past, health care for the elderly will absorb 10 percent of GDP in 2020, compared with 4.3 percent in 1995. He estimates that this will require either a sizable increase in public health expenditure or a reduction in the amount of nonhealth private goods and services the elderly can purchase compared with earlier years, or both. See Hubbard, Skinner, and Zeldes (1994) for information on the age profile of health expenditure, and Dick, Garber, and MaCurdy (1994) for an analysis of nursing home stays.

56. Lee and Skinner (1999).

57. Bernheim and Scholz (1993).

even with a time preference rate of 3 percent, our model generates high wealth benchmarks relative to previous work.

### *Extensions*

Several features of the model should be kept in mind in interpreting the results and serve as points of comparison with other studies and as a source of possible future research.

**ADDITIONAL UNCERTAINTY.** By ignoring uncertainty regarding asset returns, income during retirement, and health care expenditure, the model understates the demand for precautionary saving and overstates the adequacy of actual saving. However, Jonathan Skinner shows that uncertainty in the rate of return has a negligible impact on saving, especially relative to the importance of uncertain earnings.<sup>58</sup> This is so because, for most households for most of their lives, the vast majority of lifetime wealth is in the form of human capital. Our simulation of a stock market decline above is also intended to capture some of the possible effects of changes in asset prices.

The impact of income uncertainty in retirement may also be small. In the model, retirement income consists of social security, pensions, and the return on existing assets. Social security is fixed in real terms in the model and in the real world (subject to the legislative risk noted above), and Skinner's finding suggests that plausible variation in the overall return to existing assets does not generate much extra precautionary saving.

Adding uncertainty regarding health care expenditure could have a significant effect. Daniel Feenberg and Skinner find substantial time persistence in large medical expenditures.<sup>59</sup> Glenn Hubbard, Skinner, and Stephen Zeldes show that uncertainty about health expenses increases precautionary saving.<sup>60</sup> However, during the working years, earnings uncertainty has a much larger impact. We chose not to model health care uncertainty explicitly. However, we believe that raising the wealth-earnings ratios by 20 percent, as in table 12, would account for a large portion, if not all, of realistic uncertainty about health care expense.

58. Skinner (1988).

59. Feenberg and Skinner (1992).

60. Hubbard, Skinner, and Zeldes (1995).

PRIVATE AND SOCIAL INSURANCE. The effects of introducing uncertainty can be overstated, however, unless such changes are coupled with the introduction of plausible insurance schemes to protect against that uncertainty. The simulation model has no government-provided consumption floor. Hubbard, Skinner, and Zeldes show that such a floor reduces precautionary saving dramatically, especially among lower-income households.<sup>61</sup> Because social security is progressive, low-income households have optimal wealth–earnings ratios that are lower than those of other households even in the absence of a consumption floor. Thus, incorporating a consumption floor would reduce the optimal wealth–earnings ratio for those who already have low optimal ratios. This could explain a significant portion of the results in tables 6 and 10 that show that actual wealth–earnings ratios sometimes fall below simulated wealth–earnings ratios at the bottom of the distribution.

Likewise, the introduction of private annuity markets could insure against the risk of outliving one's assets and reduce precautionary saving in the model. Although private annuity markets used to be quite small, they have grown dramatically in recent years.<sup>62</sup> A significant minority of households in their fifties now hold annuities.<sup>63</sup> Thus, omitting an annuity market, like the absence of social insurance, raises the simulated wealth–earnings ratios relative to those in a model that contains private annuity markets. This biases our analysis toward concluding that actual saving is inadequate.

OTHER ARGUMENTS OF THE UTILITY FUNCTION. In the model, utility depends only on consumption. Extending the model to incorporate leisure as a substitute for consumption would allow for an optimal drop in consumption upon retirement, as households increase their leisure by around 1,000 to 2,000 hours per year and effectively substitute time for money. Alternatively, leisure and certain forms of consumption expenditure, such as travel, may be complements, which would increase some spending on those items after retirement.

Another important extension would consider health in the utility function. A household's standard of living, and the marginal utility of a given basket of consumption goods, undoubtedly depend on the household's

61. Hubbard, Skinner, and Zeldes (1995).

62. Gentry and Milano (1998); Poterba (1997).

63. Brown (1999).

health status, which could be expected to shift markedly between pre- and postretirement periods. The model, however, ignores health status, implicitly assuming it is held constant. The bias created by omitting health from the utility function is ambiguous. Deteriorating health may increase out-of-pocket health expenditures, but it may also simultaneously reduce other expenditures, which households would have preferred to have made when healthy.

**INTERGENERATIONAL TRANSFERS.** Adding a bequest motive would raise the level of required saving, but this may not be very important for studying the adequacy of saving, because the desire to leave a bequest is likely to be less important than the adequacy of one's own saving and is likely to be an issue mainly for wealthy households.<sup>64</sup> Moreover, clear evidence of intentional bequest motives has proven difficult to generate.<sup>65</sup> Allowing for transfers from children to parents would, on the other hand, reduce the required level of saving in the model, but such transfers are rare in practice.<sup>66</sup>

**RETIREMENT FLEXIBILITY AND PARTIAL RETIREMENT.** By setting retirement at a predetermined age, the model overstates required saving. In practice, workers who reach a given age and find that they have insufficient wealth for retirement usually have the option of continuing to work. Thus the ability to vary the date of retirement is to some extent a substitute for saving. In addition, the model overstates required saving by omitting partial retirement, which is growing in importance (and is discussed further below), and by omitting the decline in work-related expenses for those who do fully retire.

**HOUSING.** The model creates different biases with respect to housing. The model does not require people to build up a down payment in order to buy a house. This leads to an understatement of required wealth for very young households. But these are not the households on the verge of retirement that our analysis and policymakers are most concerned about. On the other hand, the model does not account for the fact that, when mortgages are paid off, the household can consume the same amount with lower expenditure. This leads to an overstatement of required wealth for older households.

64. See Carroll (2000); Gale and Scholz (1994).

65. Hurd (1987); Brown (1999).

66. Gale and Scholz (1994).

TAXES. Both the model and the data focus on ratios of before-tax wealth to before-tax earnings. Ideally, the analysis would be carried out exclusively in after-tax terms. Although the distinction between taxable and tax-preferred assets is important in some contexts,<sup>67</sup> it is unclear how our focus on before-tax values biases our results. Our best judgment is that the net effect is small.

### **Comparisons with Popular Financial Advice**

Popular financial advice often suggests that households should aim to replace between 65 and 85 percent of preretirement income in retirement.<sup>68</sup> This section compares such advice with the model and results above by developing three sets of results. First, we discuss why 65 to 85 percent of preretirement income might be considered a sufficient target. Second, we show that our simulation model generates replacement rates in this range. Third, using these targets, we show that many households can cobble together sufficient retirement income without large amounts of saving in financial assets. This finding is important because it suggests that optimization and adequate saving are not inconsistent with the widespread empirical finding that many households accumulate little in the way of financial assets.

#### *Developing Replacement Rate Targets*

A household can maintain the same consumption during retirement with less income than before retirement for several reasons. The need to save for retirement ceases, or at least diminishes substantially. Taxes decline because payroll taxes are no longer due, because income is generally lower, because social security benefits receive more favorable income tax treatment than wages, and because of the extra exemption for those over sixty-five. Work-related expenses such as for commuting and clothing decline. Family size declines as the grown children leave the household.

67. See Engen, Gale, and Scholz (1996).

68. For example, see American Savings Education Council (undated); U.S. Department of Labor (1997); Palmer (1994); Tacchino and Saltzman (1999); Tyson (1997). Warshawsky and Ameriks (1998) and Gokhale, Kotlikoff, and Warshawsky (1999) provide interesting analyses of the economic features of popular financial planning models.



Family size declines further, of course, if either spouse dies. Households eventually pay off their mortgages, which allows for continued consumption of housing services at less expense than before. Finally, households can consume some of their asset principal—not just income—in retirement.

Consider a household that contributed 5 percent of wages to a 401(k) plan while its members were working, paid payroll taxes of 7.65 percent of earnings, had commuting expenses of 3 percent of income, paid off a mortgage whose annual payment was 6 percent of preretirement income, and saw its average federal and state income tax payments fall by 4 percent of income after retirement. A postretirement income of 75 percent of preretirement income would be sufficient to maintain that household's living standard.<sup>69</sup> The household could maintain its consumption per capita on even less income if family size fell or if the household slowly liquidated its assets. If it values the increase in leisure at retirement, the household could maintain its living standard with even less income.

#### *Replacement Rates in the Simulation*

The simulation model developed above can be used to calculate replacement rates, which we define as the sum of social security benefits, pension benefits, and the return (but not any principal) on other assets, as a percentage of preretirement income. These rates vary, of course, by education and pension status. Using a time preference rate of 3 percent, the average of the median replacement rate in each of the four education and pension groups is 72 percent. Using a time preference rate of zero, the average of the median replacement rates is 80 percent. Households in the model, however, will optimally consume some of their asset principal in retirement, as well as draw on the income sources noted above. These findings provide some support for the reasonableness of common financial advice to replace between 65 and 85 percent of preretirement income.

#### *Saving Enough Versus Saving a Lot*

Even without saving a large share of income in terms of financial assets, households can easily achieve replacement rates that are within the range recommended by financial planners and by the simulation model. The

69. These expense figures are taken from various tables in McGill and others (1996).

starting point is social security. As discussed in appendix B, social security replaced about 49 to 62 percent of final earnings for the typical new beneficiary couple in 1982. Adjusting for a 20 percent decline in social security benefits since then, as described in appendix B, suggests a replacement rate of 40 to 50 percent. Alternatively, the Social Security Administration shows that replacement rates for low, average, and high earners who first receive benefits at age sixty-five in 1990 or 2000 would be about 58 percent, 43 percent, and 25 percent, respectively.<sup>70</sup> Consideration of a spouse would raise these replacement rates at the household level by up to one-half. Addition of a typical DB pension plan raises these figures still further (see appendix B). Thus the combination of social security and pensions can provide all or most of the income needed to finance an adequate retirement, by financial planning standards, for some people even in the absence of any additional saving.

If social security and pensions are not sufficient, households have the option of working part-time in retirement. Growing evidence suggests that many workers prefer to reduce their hours gradually rather than abruptly. Joseph Quinn finds that between one-third and one-half of older Americans will work on a bridge job before retiring completely.<sup>71</sup> Leora Friedberg shows that the frequency of part-time work among older men almost doubled between 1980 and 1995 and rose slightly for older women.<sup>72</sup> A recent survey of well-off retirees and near-retirees suggests that half of recent retirees and over 60 percent of preretirees expect to work in retirement, but that few expect to do so full-time.<sup>73</sup> According to a recent survey by the American Association of Retired Persons, 80 percent of baby boomers say they plan to work at least part-time during retirement, whereas just 16 percent say they will not work at all.<sup>74</sup> About 35 percent of baby boomers say they will be working part-time mainly for the sake of interest and enjoyment; 23 percent say they will work part-time mainly for the income it provides. In 1996, earnings from work represented about 34 percent of the income of households with heads aged sixty-five to sixty-

70. Definitions of each level of earner are provided in appendix B.

71. Quinn (1999).

72. Friedberg (1999).

73. Forum for Investor Advice (1999).

74. American Association of Retired Persons (1999).

nine, and about 39 percent of such households had positive earnings.<sup>75</sup> The increasing tendency for people to hold bridge jobs implies that, in the future, retired households may acquire an even more significant portion of their income from working.

In addition, many households could obtain significant income from the equity in their homes. In our HRS sample, 30 percent of households have completely paid off their mortgages, and 44 percent have homes worth at least twice their annual earnings. Thus, for many households, a reverse mortgage or sale of a house could provide a nontrivial income source.

Finally, a significant fraction of households will receive inheritances that will help provide retirement income. For all of these reasons, many households can piece together sufficient retirement income without necessarily saving much—or, as in the examples above, without saving anything—in the way of (nonannuitized) financial assets. This may not be a recommended strategy, but it does help reconcile our results with the common finding that many households have few if any financial assets, even on the eve of retirement.<sup>76</sup>

### **Comparisons with Previous Findings**

At first glance, our results stand in contrast to much—but not all—of the previous literature, which has largely concluded that many, if not most, households are saving less than is adequate for retirement. In this section we present evidence suggesting that previous results and methodologies are, in fact, largely consistent with the general tenor of our findings.

75. U.S. Social Security Administration (1998).

76. Venti and Wise (1992). Sabelhaus (1997) reaches a similar conclusion. He calculates that, given the current structure of U.S. taxes and public pensions, low saving rates in other forms of saving are sufficient to maintain lifetime consumption levels in retirement. For example, if the interest rate is 4 percent, and households save (including pensions and housing) just 7.5 percent of their gross income, estimated consumption replacement rates in retirement exceed 100 percent for income levels below \$50,000, 90 percent or higher for incomes between \$50,000 and \$70,000, and between 83 and 74 percent for incomes between \$80,000 and \$100,000. If the rate of return is 6 percent, a saving rate of 7.5 percent ensures a retirement consumption replacement rate of 100 percent or higher for all income groups up to \$90,000 and 98 percent for those with income of \$100,000.

*Simulation Models of Optimal Wealth Accumulation*

This paper and several earlier studies analyze the adequacy of saving by comparing simulated optimal saving behavior with actual household data. This approach raises several general concerns. Most important, other simulation studies produce a single optimal wealth-earnings ratio (or wealth level) for a group of households with the same current characteristics. Recognition of earnings uncertainty, however, requires that the optimal ratio be interpreted as a mean or median, not a minimum. We will assume that previous studies can be interpreted as reporting median ratios. Since the mean wealth ratio will typically be higher than the median, if we assumed that previous nonstochastic studies have generated a mean ratio, it would be easier to reconcile previous results with our own.

A second concern is that, as noted above, there almost certainly will be biases in the simulation model, the data, or both, that over- or understate the severity of the saving problem. A third is that the manner in which any saving shortfalls are reported can have important effects on the interpretation of the results.

Bernheim models households' optimal saving and consumption choices as a function of family size, education, earnings, age, social security, pensions, and other factors.<sup>77</sup> He then compares households' actual saving with the simulation results. His primary finding, summarized in a "baby boomer retirement index," is that boomers' retirement saving averages only about one-third of that needed to maintain preretirement living standards in retirement.<sup>78</sup>

The main issue in interpreting these results is understanding what the baby boomer index measures. It does not measure the adequacy of saving by the ratio of total retirement resources (social security, pensions, and other assets) to total retirement needs (the wealth necessary on the eve of

77. Bernheim (1992, 1995).

78. Bernheim (1997, p. 43) characterizes his work as "consistently find[ing] that baby boomers are saving at 33 to 38 percent of the rate required to cover their expected costs of retirement." However, the consistency of his results is open to question. The first survey found that saving in all nonhousing assets was 34 percent of the required level (Bernheim, 1992). The second study found that saving in all nonhousing assets had increased substantially—to 56 percent of required saving—but it also noted that saving in retirement assets was only 16 percent of required saving. By taking an average of the 56 and 16 percent figures, the second survey produced a boomer retirement index of 36 percent (Bernheim, 1994c).

retirement to maintain preretirement living standards). Instead, it examines the ratio of actual saving in financial assets to the total required amount of saving less social security and pensions.

Table 13 helps explain how the index is constructed. In case A, a hypothetical household (or group of households) needs to accumulate 100 units of wealth. It is on course to generate 61 units in social security, 30 in pensions, and 3 in other assets.<sup>79</sup> Total retirement resources are therefore projected to be 94 percent of what is needed to maintain living standards. But according to the baby boomer index, the household is saving only 33 percent ( $= 3/[100 - 61 - 30]$ ) of what it needs.

Thus one problem is that the *level* of the baby boomer index understates the overall adequacy of retirement preparations, and that understatement can be vast. In particular, having the baby boomer index stand at one-third does not at all imply that, unless they change their saving behavior, boomers will have living standards in retirement equal to one-third of their current living standards. It would only have that meaning if both social security and pensions were equal to zero (as in case B). It could mean that retirement living standards will be 94 percent of current living standards (case A), or 60 percent (case C), or any figure from 33 percent to more than 99 percent.<sup>80</sup>

A second problem is that *changes* in the baby boomer index over time, or differences across groups, do not correspond to changes or differences in the adequacy of overall retirement saving. If, as in case D, the household in case A rolls over its pension into an Individual Retirement Account (IRA), the baby boomer index rises dramatically, even though total retirement resources are unchanged. If, as in case E, the household in case A rolls over half of its pension into an IRA and spends the rest on a vacation, the household clearly is less well prepared for retirement than before—its index of total retirement resources falls. Yet it obtains a higher baby boomer index than in case A.

A third problem is that the baby boomer index can be extremely sensitive to estimates of retirement needs. In case F, retirement needs are only

79. The discussion of social security and private pension replacement rates in earlier sections suggests that the figures used in case A are not atypical.

80. A boomer index of one-third would imply retirement living standards over 99 percent of preretirement living standards if retirement needs were 100, social security and pension wealth were 99, and other saving were 0.33.

**Table 13. Performance of Alternative Measures of Retirement Saving Adequacy Under Selected Scenarios**

<i>Case</i>	<i>Social security</i>	<i>Pension</i>	<i>Other assets</i>	<i>Total retirement resources<sup>a</sup></i>	<i>Needs<sup>b</sup></i>	<i>Boomer index (percent)<sup>c</sup></i>	<i>Total resources index (percent)<sup>d</sup></i>
A	61	30	3	94	100	33	94
B	0	0	33	33	100	33	33
C	20	20	20	60	100	33	60
D	61	0	33	94	100	85	94
E	61	0	18	78	100	45	78
F	61	30	3	94	95	75	99
G	61	30	3	94	93	150	101

Source: Authors' calculations.

a. Social security plus pensions plus other assets.

b. Needs are defined such that 100 equals accumulated wealth on the eve of retirement sufficient to keep a constant living standard before and after retirement.

c. The baby boomer retirement index (Bernheim, 1992, 1995) is defined as (other assets)/(needs – social security – pensions).

d. The total resources index is defined as (total retirement resources)/(needs).

5 percent lower than in case A, but the boomer index rises from 33 percent to 75 percent. In case G, retirement needs are only 7 percent lower than in A, yet the boomer index rises from 33 percent to 150 percent. For all of these reasons, we conclude that the boomer index is not useful as a guide to understanding the adequacy of retirement saving.

Bernheim and Scholz use Bernheim's model,<sup>81</sup> but rather than report a baby boomer index, they compare wealth accumulation targets from the simulation with actual household data. Their sample, taken from the 1983–86 SCF, focuses on married households where the husband works full time and is between the ages of twenty-five and sixty-four. Looking at five-year age groups within the sample, they find that the median college-educated household in each age group is accumulating about what the simulation indicates is optimal. For non-college-educated households up to ages forty-five to forty-nine, the median household is also roughly on target. However, median older non-college-educated households have less than optimal wealth. The authors conclude that “many Americans, particularly those without a college education, save too little.”

We believe that their results are not necessarily evidence of under-saving. Because earnings are in fact stochastic, the model targets should be interpreted as medians of the distribution of optimal wealth-earnings ratios. Thus the Bernheim-Scholz findings for college-educated house-

81. Bernheim and Scholz (1993).

holds in all age groups and non-college-educated households up to age forty-nine should not be interpreted as showing that half of those households are saving too little. Rather, the results show no evidence of undersaving for these groups. Only the group of non-college-educated households over age fifty shows any signs of undersaving. However, as Bernheim and Scholz note, they use a narrow definition of wealth that excludes housing equity. Using our simulation model with a time preference rate of 3 percent, we find that, using this narrow definition, only 28 percent of non-college-educated households aged fifty to sixty-two in the 1992 SCF exceed the median simulated wealth targets. However, when we use the broad wealth definition (which, again, includes all of housing equity), the figure rises to 48 percent. Using the Bernheim-Scholz targets with the 1992 SCF, 36 percent of non-college-educated households have narrow wealth-earnings ratios exceeding the target, and 65 percent have broad wealth-earnings ratios exceeding the target. Thus inclusion of housing wealth eliminates most or all of the estimated shortfall between median actual and median simulated wealth-earnings ratios for these households, depending on the benchmark used.

Bernheim and Scholz point out that the underlying model understates the retirement saving problem because it assumes no reduction in social security benefits, no increase in life span, and no increase in health care costs at retirement, and it ignores motives for saving other than retirement.<sup>82</sup> Other factors, however, bias the model the other way. Besides housing wealth, the model also omits all earnings after retirement, any decline in work-related expenses for those who do retire fully, and all inheritances. Table 12 provides estimates of the relative importance of these biases using our model, which allows for both precautionary and retirement saving.<sup>83</sup>

82. Bernheim (1992, 1995); Bernheim and Scholz (1993).

83. Gale (1997) estimates the proportion of households that exceed the Bernheim and Scholz (1993) wealth targets. Using a sample similar to that of Bernheim and Scholz and ignoring housing equity, he finds that 47 percent of households had wealth above the target levels in 1992. When housing is included, the figure rises to 71 percent. These estimates provide very high measures of the adequacy of saving, especially when the wealth targets are interpreted as medians of the distribution of targets rather than minima. Gale interpreted his results as showing that many households were saving adequately but that a significant minority were not. This interpretation, however, is flawed because it, too, is based on the assumption that the target wealth ratios are minima rather than medians.

Warshawsky and Ameriks apply data from the 1992 SCF to a popular financial planning program.<sup>84</sup> Their sample includes households where either the respondent or the respondent's spouse or partner was employed full time, the respondent or spouse was between twenty-five and seventy years old, nonfinancial assets were below \$1 million, and annual salary was below \$125,000. They find that 52 percent of these "working, middle-class American households will not have fully funded retirements." Among households with wealth shortfalls, the median undiscounted shortfall is \$273,000, and the mean is \$297,000. Although these results appear, at first glance, to represent significant amounts of undersaving, we believe the findings are consistent with little or no undersaving, for several reasons.

First, the paper assumes that earnings are nonstochastic. Allowing for stochastic earnings, the Warshawsky and Ameriks results indicate that almost half (48 percent) of households exceed the *median* optimal wealth-earnings ratio. This finding is consistent with the simulation model we have developed above.

Second, wealth available to finance retirement consumption may be understated. Households are forbidden in the model to use 75 percent of their primary housing equity to finance consumption. They may not sell any businesses, second homes, or other nonfinancial assets, and these assets are assumed to generate no cash income. All inheritances and earnings from part-time work in retirement are ignored.<sup>85</sup>

Third, current consumption is set equal to wage income less taxes, pension saving contributions, housing expenses, and debt payments. This is likely to overstate consumption because it assumes there are no active contributions (as opposed to interest accruals) to discretionary saving, including fully taxable accounts or Individual Retirement Accounts.<sup>86</sup> Because

84. Warshawsky and Ameriks (1998).

85. The model may also understate lifetime labor earnings before age sixty-five, because the data on retirement age may be problematic. Some households apparently plan to retire as young as twenty-seven, and a significant number indicate they will retire before turning fifty.

86. In addition, the assumption of zero current active contributions to discretionary saving, combined with the authors' assumption that earnings grow faster than living expenses up until age fifty, implies that active contributions to discretionary saving in previous years would have to have been *negative* for those younger than fifty. But this is inconsistent with observed positive holdings of discretionary assets. One solution to this inconsistency is to reduce current consumption levels. This would imply lower consumption in all future years and reduce the wealth targets.



current consumption determines the entire consumption stream, any overstatement could have significant effects on lifetime consumption.<sup>87</sup>

Fourth, the authors add ten years to each person's life span, which raises lifetime consumption needs significantly. As shown in table 12, raising life span can have a significantly deleterious impact on the measured adequacy of saving. Warshawsky and Ameriks explain this adjustment as a way to account for uninsured health care and long-term care, uncertain life span, imperfect annuity markets, and expected increases in life expectancy over time.<sup>88</sup> It is difficult to determine the appropriateness of this adjustment.<sup>89</sup>

Even if one ignores these issues, it would be interesting to determine the size of any implied shortfall in retirement consumption. As a rough calculation, with an average retirement period of forty-one years, as in the Warshawsky and Ameriks paper, a real return of 5 percent, and an average undiscounted shortfall in retirement wealth among those with inadequate wealth of \$297,000, the equivalent shortfall in annual retirement consumption comes to only \$2,213 per year. It would be interesting to know what percentage of retirement consumption that figure represents and to perform such calculations on a household-by-household basis.

#### *Estimates of the Value of Annuitized Wealth*

A second approach is to compare households' preretirement consumption or income with the consumption or income that could be generated by converting the households' wealth into a hypothetical annuity. Alan Gustman and Thomas Steinmeier, using data from the 1992 HRS, find that wealth accumulated through 1992—not through retirement—would

87. Consumption is held constant at its current level during the working years, falls 20 percent at retirement, and then falls another 20 percent when the first spouse dies. By not allowing consumption to rise through the working years (as our model does, as shown in figures 3 and 4), the authors understate consumption and therefore understate required wealth accumulation. However, by not reducing consumption needs when the children leave the home, the model overstates required wealth accumulation. The net effect is that required wealth accumulation will be overstated for households near retirement, but possibly understated for younger households.

88. They also note that only 3 percent of the financial plans in the sample fail in the last ten years of the lifetime.

89. Life expectancy is expected to rise by only three years between now and 2080 (Lee and Skinner, 1999). Although health expenditure rises in retirement, consumption of other goods may optimally decline as households age and health deteriorates.

finance a nominal annuity replacing 86 percent of projected final earnings, on average.<sup>90</sup> The nominal replacement rate for households in the median 10 percent of the lifetime earnings distribution is 97 percent. The corresponding real replacement rates are 60 percent overall and 66 percent for the median 10 percent. The average replacement rates for nominal and real annuities for all groups of lifetime earners except the top 5 percent are at least 83 percent and 59 percent, respectively, of final earnings, using current assets alone. Gustman and Steinmeier conclude that “it is hard to find evidence of a massive crisis in retirement undersaving of the type that has been promoted in the media.”<sup>91</sup>

James Moore and Olivia Mitchell use the same data set to estimate how much respondents need to save between 1992 and their time of retirement if they wish to preserve preretirement consumption levels after retirement.<sup>92</sup> They find that the median household will need to save 16 percent of annual earnings, in addition to saving that occurs through mortgage repayment, accruing interest on net financial assets, and accruals in pension value, between 1992 and retirement at age sixty-two in order to equate pre- and postretirement consumption. Note that this does not imply that households have to save 16 percent of their income *more* than they currently do. If retirement occurs at age sixty-five, the median household needs to save 7 percent of annual earnings for the remainder of its adult members' careers. Moore and Mitchell also find substantial diversity in required saving rates. More than 30 percent of households require no additional saving for retirement at age sixty-two, but at least 40 percent of households have a prescribed saving rate of 20 percent or higher. They interpret their findings as showing that “despite seemingly large accumulations of total retirement wealth, the majority of older households will not be able to maintain current levels of consumption into retirement” without saving positive discretionary amounts between 1992 and their age of retirement. This conclusion is consistent with the results, but it is also consistent with little or no undersaving in the population, for five reasons.

First, most of the households in the HRS are typically still working. Thus it is not surprising that they have not amassed sufficient retirement

90. Gustman and Steinmeier (1998). They assume that couples purchase a two-thirds joint and survivor annuity.

91. Gustman and Steinmeier (1998, p. 23).

92. Moore and Mitchell (1997).

wealth before they have actually retired. In the typical optimizing model, households reach sufficiency of retirement wealth only in the last period before retirement.

Second, the model assumes that earnings are nonstochastic. Suppose a household is saving adequately for retirement but in 1992 experiences a large positive wage shock. Moore and Mitchell, using the household's current earnings, would conclude that the household would have to save very large amounts to maintain current living standards in retirement. Thus a significant portion of the households that are thought to need to save large amounts to maintain preretirement living standards may be households with positive earnings shocks in that year. Indeed, in subsequent work, Mitchell, Moore, and John Phillips find that households with high current earnings are much more likely to have saving shortfalls, precisely the pattern that would be expected if earnings are stochastic.<sup>93</sup>

Third, having households maintain current levels of consumption into retirement may be too strict a test for the adequacy of saving. Allowing for mortality risk, setting the time preference rate equal to 3 percent, and using a 3 percent real rate of return—but ignoring any reductions in work-related expenses, mortgage costs, or family size—the present value of optimal consumption generated by our simulation model between retirement at age sixty-two and the end of life is 28 percent lower than it would be if consumption were flat over time.<sup>94</sup> It is 16 percent lower if the rate of time preference is zero.

Fourth, even if the median household saved nothing, rather than 16 percent of its income, between 1992 and retirement, its retirement consumption would not fall very much. The median household has current wealth of \$325,000, which is projected to rise to \$382,000 at age sixty-two, even if the household does no additional discretionary saving. The household needs to save 16 percent of its current income to maintain consumption in retirement. If the adult household members are fifty-six years old (the HRS covers households aged fifty-one to sixty-one) and earn \$35,000 a

93. Mitchell, Moore, and Phillips (1998). More generally, households with high current earnings could be households that had two earners in 1992 but did not have two earners for most of their careers. Similarly, a large portion of the 31 percent of households that appear to do no further discretionary saving before retirement could be households with temporarily low earnings.

94. These calculations are based on the age-consumption profiles shown in figures 3 and 4. The percentage declines do not differ significantly across education and pension classes.

year (the average of 1992 earnings in the fifth and sixth earnings deciles),<sup>95</sup> the household needs to save \$5,600 per year for six years. Accumulating these funds at a real rate of 5 percent would generate \$40,000 in additional wealth by age sixty-two. This would raise the household's wealth at that age to \$422,000. Thus, even if the median household had no discretionary saving between 1992 and retirement, it would be able to finance retirement consumption equal to over 90 percent (\$382,000 divided by \$422,000) of the level that Moore and Mitchell describe as optimal.<sup>96</sup>

Lastly, the study ignores other possible sources of retirement income, including part-time work and inheritances. The results above suggest that earning a total of just \$40,000 (after taxes and work expenses) from part-time work during the remaining lifetime after retirement at age sixty-two would be sufficient to generate the optimal retirement consumption level calculated by Moore and Mitchell.

Laurence Kotlikoff, Avia Spivak, and Lawrence Summers use the Retirement History Survey (RHS), conducted by the Social Security Administration, to compare the constant level of consumption that could be financed by the elderly, based on their resources in old age, and the constant level of lifetime consumption that the same households could have financed based on lifetime resources available at the start of their life.<sup>97</sup> They find that over 90 percent of married couples could afford consumption in old age exceeding 80 percent of their affordable lifetime consumption level. In addition, depending on assumptions about annuitization, between 72 and 85 percent of them could afford higher consumption in old age than over the rest of their lifetime. The authors conclude that the results suggest little undersaving in the RHS cohort.

#### *Analysis of Changes in Consumption at Retirement*

If households are saving inadequately, their consumption has to fall in retirement. However, interpreting tests of how consumption changes upon retirement raises several difficult issues. First, how much *should* consumption fall at retirement? As the household head ages and reaches retire-

95. See Moore and Mitchell (1997), table 3B.

96. Similar calculations suggest that if the household had no discretionary saving between 1992 and retirement at age sixty-five, its consumption in retirement would only be 6 percent less than if it saved optimally.

97. Kotlikoff, Spivak, and Summers (1982).

ment, consumption will optimally fall as family size shrinks, work-related expenses decline, and mortality risk rises.<sup>98</sup> Consumption will fall further if time preference rates are positive and if consumption and leisure are substitutes.

A second concern is the distinction between consumption services and consumption expenditure. In optimizing models, households smooth the former, not the latter. As households age, many pay off their mortgages. When they do, the household can maintain the same level of consumption services with less expenditure than before.

A third issue is whether retirement occurs as part of the household's optimal plan or as an involuntary and unanticipated event—such as the onset of disability—that conveys new information to the household about its lifetime income. In the latter case, any reduction in consumption beyond a benchmark level may be the result of the new information conveyed by the event that prompted retirement, rather than by inadequate saving in the original plan.

James Banks, Richard Blundell, and Sarah Tanner carefully examine cohort data on British households and find that, even after controlling for labor force participation, age, mortality risk, and other factors, there remains an unexplained drop in consumption expenditure on nondurables around the time of retirement.<sup>99</sup> They suggest that the findings are consistent with households' realization of a negative surprise at the time of retirement. For example, households may have just found out how little they had actually saved for retirement.

We interpret their findings as suggesting that the implied retirement income shortfall is fairly small. Between the ages of sixty-one and sixty-six (during which period almost half of recent British cohorts retire), actual cohort consumption expenditure falls by about 12 percent. The authors' model, which includes age, demographics, mortality risk, and leisure, can explain about 10 percentage points of the decline.<sup>100</sup> Thus,

98. An increase in mortality risk has the same effect on the age-consumption profile as a rise in the time preference rate, tilting consumption toward the present and away from the future.

99. Banks, Blundell, and Tanner (1998).

100. Data were kindly provided by Sarah Tanner. For cohort retiree patterns and the actual and predicted consumption declines, see figures 1 and 6 of Banks, Blundell, and Tanner (1998).

even if other factors cannot explain the residual, only a 2 percent decline in cohort consumption can be attributed to a shortfall in saving. This seems small in economic terms.

Bernheim, Skinner, and Steven Weinberg provide an intriguing analysis of consumption, income, and wealth around retirement using panel data from the Panel Survey of Income Dynamics between 1978 and 1990.<sup>101</sup> They find that households with lower wealth-earnings ratios before retirement and households with lower income replacement ratios in retirement have larger reductions in consumption at retirement. They note that these patterns hold even after controlling for unanticipated shocks that could lead to early retirement. They interpret the results as posing “a significant challenge to the validity of standard life-cycle models” and, like Banks, Blundell, and Tanner, suggest that many households may receive an unpleasant surprise upon retirement.

Surprisingly, however, the results in the Bernheim, Skinner, and Weinberg paper actually show that most households experience an *increase* in consumption after retirement, again after controlling for unanticipated events that may cause retirement and affect lifetime income simultaneously, and for other factors. To examine the manner in which consumption responds to predictable events that affect the probability of retirement, Bernheim, Skinner, and Weinberg proceed in two steps. First, they estimate probit models that explain retirement status as a function of household education, family size, sex of family head, and marital status. A separate model is estimated for each age between fifty-four and seventy. Second, they estimate the log of household consumption as a function of the household-specific predicted probability of retirement, that same probability interacted with dummies representing the various wealth and income quartiles, and other household characteristics, including those used in the first-stage probits. They estimate the second-stage equation with and without age as an independent variable. However, because the retirement probability is naturally a function of age, it makes sense to control for age separately in the regression to isolate the impact of predicted retirement on consumption. Controlling for age, they estimate the following equation:

101. Bernheim, Skinner, and Weinberg (1997).

$$\begin{aligned}
 \log C = & -0.274 * P \\
 & (2.8) \\
 & + 0.073 * WQ_2 * P + 0.212 * WQ_3 * P + 0.195 * WQ_4 * P \\
 & (0.9) \qquad (2.2) \qquad (1.9) \\
 & + 0.157 * IQ_2 * P + 0.331 * IQ_3 * P + 0.298 * IQ_4 * P \\
 & (1.5) \qquad (3.4) \qquad (2.9) \\
 & + \text{other factors,}^{102}
 \end{aligned}$$

where  $C$  is consumption,  $P$  is the household-specific fitted probability of retirement, and  $WQ_i$  and  $IQ_i$  are wealth and income quartiles, respectively. The coefficient on  $P$  represents the effect of changes in the predicted probability of retirement on log consumption for households that are in both the first income quartile and the first wealth quartile.

Table 14 uses these results to display the implied change in log consumption in each combination of wealth and income quartiles. For example, for households in the second income quartile and the third wealth quartile, table 14 shows that an increase in the probability of retirement raises log consumption by 0.10. This is the sum of three coefficients:  $-0.274$ , associated with the probability of retirement  $P$ ;  $0.157$ , associated with  $IQ_2 * P$ ; and  $0.212$ , associated with  $WQ_3 * P$ .

Table 14 shows that 60 percent of households are in subgroups where average consumption *rises* at retirement, once the other factors noted above have been controlled for.<sup>103</sup> More than 25 percent of all households (those in the top two quartiles of both income and wealth) are in cells where consumption rises at retirement by about 20 percent or more.

About another 25 percent of households are in wealth and income cells where consumption falls by between 4 and 8 percent at retirement. For reasons noted above, these changes do not strike us as worthy of concern. Even in optimizing models, consumption should be expected to fall by some amount as households age and retire.<sup>104</sup> Approximately 16 percent of

102. Numbers in parentheses are z-statistics. This equation is based on Bernheim, Skinner, and Weinberg (1997), table 5, column c, which uses financial assets to generate the wealth quartiles. Using total wealth, in their table 5, column d, yields very similar results.

103. More precisely, the finding is that consumption rises for these households in response to predictable events that positively affect the probability of retirement.

104. It is also worth noting that it can take more than one year to “retire” in their model (Bernheim, Skinner, and Weinberg, 1997, p. 27), so the consumption changes estimated in the table could be occurring over several years.

**Table 14. Estimated Change in Consumption at Retirement by Income and Wealth Quartile**

<i>Income quartile</i>	<i>Wealth quartile</i>	<i>Percent of total sample in cell</i>	<i>Estimated change in log consumption at retirement</i>
1	1	6.2	-0.27
	2	4.5	-0.20
	3	6.6	-0.06
	4	7.8	-0.08
2	1	5.4	-0.12
	2	9.2	-0.04
	3	5.2	0.10
	4	5.2	0.08
3	1	6.9	0.06
	2	5.0	0.13
	3	7.8	0.27
	4	5.0	0.25
4	1	6.6	0.02
	2	6.1	0.10
	3	5.4	0.23
	4	6.9	0.22

Source: Bernheim, Skinner, and Weinberg (1997, table 2a and table 5, column C) and authors' calculations.

households in the sample are in cells where significant declines in consumption are expected to occur. Thus we view Bernheim, Skinner, and Weinberg's results as showing support for the view that the vast majority of households are saving adequately, and that significantly inadequate retirement saving is concentrated in only a small minority of households.

Daniel Hamermesh finds that current consumption spending among newly retired couples in the RHS was 14 percent higher than the annuitized value of their wealth (including physical, financial, social security, and pension wealth) and interprets this as evidence of inadequate saving.<sup>105</sup> However, consumption expenditure by a retired couple *should* fall over time. As noted above, the simulation results shown in figures 3 and 4 imply that assuming consumption is constant from age sixty-two until death overstates optimal consumption by 16 percent, even with a time preference rate of zero, and by 28 percent if the time preference is 3 percent. Moreover, these figures do not adjust for reductions in work expenses, mortgage

105. Hamermesh (1984).



costs, or reductions in family size. Thus there is nothing implausible or irrational about having consumption at age sixty-two exceed the annuitized value of wealth by 14 percent.<sup>106</sup>

Hausman and Paquette, using RHS data, find that food consumption per person falls by 10 percent in households headed by men who retire voluntarily and by 30 percent in households headed by men who retire involuntarily.<sup>107</sup> These results are consistent with inadequate retirement saving, but are also consistent with the view that planned saving was adequate, but that involuntary, unanticipated retirement contains news about lifetime income prospects. Robb and Burbidge find that consumption falls by about 15 percent upon retirement for a typical Canadian blue-collar household, but does not fall for white-collar workers.<sup>108</sup>

### *Intergenerational Comparisons*

A fourth approach compares wealth accumulation patterns across generations. The Congressional Budget Office (CBO) compared households aged twenty-five to forty-four in 1989 (roughly the baby boomer cohort) with households the same age (roughly the parents of the boomers) in 1962.<sup>109</sup> Baby boomer households were shown to have more real income and a *higher* ratio of wealth to income than the previous generation. Richard Easterlin, Christine Schaeffer, and Diane Macunovich compare income, wealth, and demographic status across numerous generations.<sup>110</sup> They conclude that boomers are doing “considerably better” than the previous generation and project that boomers’ status in retirement is likely to be better than that of previous generations, but not by as much as the

106. Hamermesh also shows that, among households that were retired in 1973 and stayed retired in 1975, real consumption fell by an average of 5 percent per year for two years. He concludes that these households cut back on their consumption in response to their observed saving shortfall, but there are other possibilities. One is that a significant portion of households paid off their mortgages, or took long-awaited trips, just after retiring. Another is that the severe recession that occurred at the end of this period caused retirees to cut back on their expenditure.

107. Hausman and Paquette (1987).

108. Robb and Burbidge (1989). Mariger (1987) finds that adults over age sixty-four consume 47 percent less than a younger adult, and he concludes that consumption drops at retirement. However, his result is based on a cross-sectional comparison of different cohorts and so does not speak to the impact of retirement on consumption.

109. U.S. Congressional Budget Office (1993).

110. Easterlin, Schaeffer, and Macunovich (1993).

boomers' increase in living standards during their working years relative to previous generations.

The last two studies establish the simple and important fact that baby boomers, to date, are accumulating more wealth than previous generations. Nevertheless, there is considerable controversy over how to interpret such findings. For example, the CBO study only implies that baby boomers are going to do well in retirement if the current generation of elderly is thought to be doing well; if the retirement needs of the two generations are the same; if baby boomers will be satisfied to do as well in retirement as today's retirees; and if the experience from middle age to retirement is the same for both generations. None of these conditions is guaranteed.

First, although the current generation of elderly is generally thought to be doing well in retirement, some 18 percent were living below 125 percent of the poverty line in 1995. Second, retirement needs may be higher for boomers because they will live longer than the previous generation, may retire earlier, and will likely face higher health costs. Third, baby boomers may view doing as well as the previous generation as a less than satisfactory accomplishment from a personal or a policy perspective.

The most difficult issue is whether the boomers' experience from middle age to retirement will resemble that of previous generations. Bernheim argues that earlier generations benefited from growth of real social security benefits in the 1970s and from the general inflation that occurred between 1965 and 1985, which dramatically raised housing values and reduced the real value of mortgage debt.<sup>111</sup> Joyce Manchester, however, notes that the CBO analysis is based on data from 1962—before the prior generation received windfall gains in housing and social security.<sup>112</sup> Thus the boomers are currently accumulating more than their parents did before their parents received windfall gains.<sup>113</sup>

111. Bernheim (1994b, 1997). However, if housing is omitted from measures of the adequacy of saving, as Bernheim advocates, it is unclear why an increase in housing wealth should be thought of as having helped the previous generation.

112. Manchester (1994).

113. Both Attanasio (1993) and Bosworth, Burtless, and Sabelhaus (1991) found a steep decline in age-specific saving rates in the 1980s for households born between 1925 and 1940. This suggests that these households responded rationally to the windfall gains they received in housing and social security. If the boomers do not receive such large windfalls, they are less likely to reduce their future saving as the previous generation did. It is also worth noting that the cohort identified by these two studies is essentially the HRS cohort, which was born between 1931 and 1941. But our evidence above, the results from Gust-

A number of other factors could evolve differently for boomers and their parents. Although boomers have benefited from the stock market boom of the last fifteen years, they may face a weak market if all of them try to cash in their funds at the same ages. Boomers have had fewer children than their parents, which implies lower living expenses, but also smaller reductions in living expenses when their children leave home. Boomers are having children later, and so they have less time to accumulate wealth after their children finish college and before retirement. And even with fewer children, boomers will likely pay more college tuition for their children, because more of their children will attend college and college will cost more in real terms.<sup>114</sup>

Several factors, however, will clearly help the boomers' retirement prospects relative to those of previous generations. More female boomers are working, and those that work are earning more than women in previous generations did. This will raise pension coverage and benefits. It also implies that boomers will have a bigger drop in work-related expenses when they do fully retire. Pension coverage and benefits will also rise for other reasons. Gustman and Steinmeier show that, between 1969 and 1992, if one ignores the effects of changes in wages and years of work covered, changes in pension coverage and plan provisions would have raised the total wealth of each household in the HRS by \$67,000 in 1992.<sup>115</sup> This is so because older plans were less generous and covered fewer people, and older generations were covered for fewer years.

Moreover, lifetime earnings may peak later for boomers than for the previous generation, because boomers are more likely to be in white-collar jobs than in jobs that emphasize physical effort. This means that, at any given age, relative to the previous generation, boomers have a greater proportion of their lifetime income (from which to save) ahead of them. Finally, as noted above, boomers appear to be much more likely to be willing to engage in partial retirement. Earnings for the baby boomers from part-time work after age sixty-two could effectively supplement traditional sources of retirement income to an extent unseen in the past.

---

man and Steinmeier (1998), and, we would argue, those of Moore and Mitchell (1997) suggest that this cohort is doing quite well. This supports the view that the decline in saving in the 1980s was not irrational for this cohort.

114. Bernheim (1994b, 1997).

115. Gustman and Steinmeier (1999).

### Comparisons with Surveys

The sophisticated methods developed in the previous literature are one way to glean information about households' preparedness for retirement. Another is simply to ask people how they think are doing in terms of accumulating retirement wealth. Unfortunately, deciphering the answers is not as straightforward as it might seem.

A 1997 survey by Public Agenda found that 76 percent of those responding thought that they should be saving more for retirement.<sup>116</sup> Of those who felt they were at a point in life where they should be saving seriously, only 6 percent felt they were ahead of where they should be, whereas 55 percent felt they were behind. A 1993 Luntz Webber–Merrill Lynch survey asked baby boomers the proportion of income they thought they should save and the proportion they actually saved. The median difference was 10 percentage points. More than three-quarters of respondents said they were saving less than they thought they should. Similar examples abound.

The first problem is how to interpret these answers. They could mean that households are vastly undersaving. But we would not jump to that conclusion so quickly. After all, in conventional economic theory, consumers have *unlimited* wants and limited resources. In equilibrium, the consumer sets the ratio of marginal utilities between any two goods equal to their price ratios. Thus consumers derive positive marginal utility from consuming more of all goods. So if asked whether they would like to save more or “should” save more, consumers may well say yes, because the utility they would receive from more saving is positive, holding other forms of consumption constant. Likewise, a worker who is making optimal labor-leisure trade-offs may nonetheless say that he or she would like to have more leisure, other things equal.<sup>117</sup>

A second problem is the astonishing range of answers one can obtain from such surveys. In contrast to the answers given above, over 70 per-

116. These examples are taken from Laibson, Repetto, and Tobacman (1998).

117. Perhaps a professional example will be useful. Suppose we surveyed economists to see how many articles they read and how many they thought they should read. Most would undoubtedly answer that they should read many more than they do; after all, how many economists feel that they read every article that they should? But such an answer would not necessarily indicate any sort of irrationality. It may simply indicate the impossibility of learning everything one would like to know.

cent of workers surveyed in the 1999 Retirement Confidence Survey, conducted by the Employee Benefit Research Institute, reported being either very or somewhat confident that they will have enough money to live comfortably throughout their retirement years.<sup>118</sup> Similarly, two-thirds of workers surveyed in the 1997 Workplace Pulse Survey said they expect to live as well or better in retirement than they did while working.<sup>119</sup> A 1999 survey by Roper Starch Worldwide Inc. and the American Association of Retired Persons found that two-thirds of boomers were satisfied with the amount of money they are putting away today for retirement.<sup>120</sup>

A third issue is the apparently inconsistent answers within some surveys. The same relatively optimistic surveys noted above suggest that the high levels of confidence regarding retirement well-being may not be well-founded. Although a majority of workers in the Retirement Confidence Survey claimed to have set aside personal savings for retirement, only about half had determined how much they need to save before they retire in order to achieve their planned consumption levels. Of the workers surveyed in the 1997 Workplace Pulse Survey, only 40 percent have put a retirement plan in place.

## **Conclusion**

This paper has presented new theory and evidence, and reexamined existing evidence, regarding the adequacy of household saving. Our study differs from previous work in that it uses a stochastic life-cycle model to generate optimal wealth accumulation benchmarks. Because of the uncertainty of earnings, the model generates a distribution of optimal wealth-earnings ratios among households that are observationally equivalent. This distribution implies that some households that have very low wealth-earnings ratios are nonetheless saving optimally for retirement.

Applying the model to data from the HRS and the SCF suggests, in the base specification, that more than half of households exceed the median wealth-earnings ratios from the simulation. In addition, households at the

118. Results of the Retirement Confidence Survey may be found on the Employee Benefit Research Institute's World Wide Web site: [www.ebri.org/rcs/index.htm](http://www.ebri.org/rcs/index.htm).

119. Workplace Pulse Surveys (1997).

120. American Association of Retired Persons (1999).

75th and 95th percentiles of the wealth-earnings distribution exceed the models' wealth benchmarks. There is some mixed evidence of inadequate saving among households with low wealth-earnings ratios. Our results appear, at least at first glance, to be significantly more optimistic than the interpretations provided in previous research. However, a careful interpretation of previous work indicates that earlier results are in fact largely consistent with ours.

Nevertheless, several caveats and limitations should be kept in mind. Our results should not be interpreted as indicating that all households are necessarily acting rationally, or that there is no saving problem at all. Rather, we show that there is significant uncertainty regarding how the adequacy of saving is affected by assumptions regarding housing wealth, the time preference rate, other model parameters, and real-world contingencies relating to health care expenses, life span, and other factors. Nor do our results have immediate implications for the welfare effects of raising private or national saving.

Clearly, there are significant warning signs that portend potential saving problems in the future. Lusardi notes that one-third of households in their fifties appear not to have thought much about retirement.<sup>121</sup> Approximately 20 percent of households, including 45 percent of black households and 37 percent of fifty-one- to sixty-one-year-olds who have less than a high school education, do not have a checking or a savings account.<sup>122</sup> Our results do not include single workers or unemployed couples, although both of these groups have been shown in previous work to be more at risk than married couples who work full-time.<sup>123</sup>

Thus perhaps the best way to interpret our results is that they show that how a saving benchmark is established has important implications for measuring the adequacy of saving. This insight, coupled with our best judgments regarding the model and the data, suggest that households are largely saving adequately, but other interpretations are possible.

We close by addressing how it could be possible that household saving may be adequate given that aggregate saving rates have fallen so far in recent years and remained so low. There is, for better or worse, no neces-

121. Lusardi (1999).

122. Carney and Gale (1998); Lusardi (1999).

123. Employee Benefit Research Institute (1994); U.S. Congressional Budget Office (1993); Mitchell, Moore, and Phillips (1998).

sary connection between aggregate saving figures and microeconomic findings about the adequacy of saving. The official personal saving measures do not measure wealth accumulation in the form of capital gains, which have been quite substantial in recent years. They provide inconsistent treatment of durable goods, payments from corporations, inflation, and taxes.<sup>124</sup> They are affected by demographic factors, and they provide no information on the distribution of saving across households. In fact, our estimates in table 11 show that some measures of retirement preparedness have improved or remained roughly constant between 1983 and 1995, even though official saving figures plummeted during this period.

## APPENDIX A

### *A Stochastic Life-Cycle Simulation Model of Saving*

THIS APPENDIX DESCRIBES the model used to generate the saving benchmarks described in the paper.

#### **Household Preferences**

A household's expected lifetime utility at age  $t$  is defined as:

$$(1) \quad ELU = E_t \left[ \sum_{k=t}^D \pi_t^k (1 + \delta)^{t-k} U(C_k; \theta) \right] \quad \forall t = 1, \dots, D,$$

where  $E_t$  is the expectations operator (conditional on information available at age  $t$ ).

In the stochastic version of the model, the household's age at death is uncertain. If  $D$  is assumed to be the maximum length of life, then  $\pi_t^k$  is the probability of living to age  $k$  (given being alive at age  $t$ ), where:

$$\pi_t^t = 1, \pi_t^{D+1} = 0, \text{ and } 0 < \pi_t^k < 1, \forall k = t, \dots, D.$$

In the certainty version of the model,  $\pi_t^k = 1$ .

Expected lifetime utility as shown in equation (1) is intertemporally separable, with a time preference rate of  $\delta$ . The intraperiod utility function,  $U(\cdot)$ , is a function of household consumption,  $C_k$  (where  $C_k > 0 \forall k$ ), and assumed to have the form:

124. Gale and Sabelhaus (1999).

$$(2) \quad \begin{aligned} U(C_t; \theta) &= \phi_t (C_t / \phi_t)^{1-\gamma} / (1-\gamma) && \text{if } \gamma \neq 1, \\ &= \phi_t \ln(C_t / \phi_t) && \text{if } \gamma = 1, \end{aligned}$$

where the vector of preference parameters  $\theta$  is composed of the elements  $[\gamma, \phi_t]$ . The utility function implies a constant intertemporal elasticity of substitution equal to  $1/\gamma$ , constant relative risk aversion equal to  $\gamma$ , and constant relative prudence equal to  $\gamma + 1$ , where  $\gamma > 0$ . Constant relative risk aversion and prudence imply that the third derivative of utility is positive, so that uncertain income generates precautionary saving. Consumers with a greater degree of prudence will devote a greater fraction of their resources to precautionary saving.

The parameter  $\phi_t$  is a time-varying family structure parameter that captures changes in family size, where  $\phi_t \geq 1$ . Thus the household is effectively maximizing utility with respect to consumption per capita.

Labor is supplied exogenously, and leisure is assumed to be separable in the utility function. Tax revenue collected by the government is spent on social security benefits, which are transferred to households in retirement, and government-provided goods and services that are assumed to be separable from private consumption in the household's utility function.

### Household Budget Constraint

Household wealth,  $W_t$ , at the beginning of period  $t$  (after receiving labor and capital income and paying taxes for period  $t$ , but before consuming in period  $t$ ) is defined according to the wealth transition equation:

$$(3) \quad W_t = [W_{t-1} - C_{t-1}](1 + r_t^a) + Y_t + B_t - T_t + I_t, \quad \forall t = 1, \dots, D.$$

where labor earnings are denoted by  $Y_t$ . Assuming that  $R$  is the exogenous retirement age,  $Y_t \geq 0$  for  $t = 1, \dots, R$ , and  $Y_t = 0$  for  $t = R + 1, \dots, D$ . The average after-tax rate of return on assets held from period  $t - 1$  to  $t$  is  $r_t^a$ . Benefits from private pensions and social security are given by  $B_t$ , where  $B_t = 0$  for  $t = 1, \dots, R$ , and  $B_t \geq 0$  for  $t = R + 1, \dots, D$ . Inheritances are equal to  $I_t$ , where  $i$  is the (exogenous) age at which any inheritance is received, so that  $I_t = 0 \quad \forall t \neq i$ , and  $I_i \geq 0$ .

Total taxes paid (excluding capital income taxes, which are already removed from the budget constraint by using the after-tax rate of return),  $T_t$ , are a function of labor income, social security, pension benefits, and inheritances:  $T_t = \tau\{Y_t, B_t, I_t\}$ . Taxes are imposed as a progressive income tax with a standard deduction and rising marginal rates on all income. To



capture simply the effect of preferential capital gains tax rates and tax-preferred saving vehicles, the tax rate on accruing capital income is capped at 20 percent. A flat payroll tax, similar to social security, is imposed on labor earnings up to an annual threshold. Inheritances are subject to a flat 20 percent tax, which greatly simplifies the actual system of estate taxes.

Households are not allowed to die in debt. If life span is certain, consumers can borrow in periods 1 through  $D - 1$ , as long as net wealth at the end of period  $D$  is positive, where  $D$  is the actual and known time of death. This implies the boundary constraints that  $W_D - C_D \geq 0$  and  $C_D \geq 0$ .

In contrast, if life span is uncertain, the household could die in any period and thus is never allowed to be in a negative net wealth position, so that consumption is restricted by a set of boundary conditions:

$$(4) \quad W_t - C_t \geq 0, \quad \forall t = 1, \dots, D,$$

However, with the utility function in equation (2), households will not want to borrow against the uncertain portion of labor income even if there is no life span uncertainty.<sup>125</sup> If the probability of receiving zero earnings in each period is positive, and there is no government-provided or informal income support, consumers will behave as if they were constrained in their borrowing even though they are not. They will never borrow and risk having nothing to consume in the future.

The wealth transition equation (3) and the net asset constraint in equation (4) imply the following set of budget constraints:

$$(5) \quad \sum_{j=1}^t \frac{Y_j + B_j + I_j - T_j}{\prod_{k=1}^j (1 + r_k^a)} \geq \sum_{j=1}^t \frac{C_j}{\prod_{k=1}^j (1 + r_k^a)}, \quad \forall t = 1, \dots, D.$$

where  $r^a$  is the average after-tax rate of return on assets used to discount future income and consumption. If life span is certain (and thus there are no accidental bequests and inheritances), equation (5) collapses to a single lifetime budget constraint:

$$(5') \quad \sum_{t=1}^D \frac{Y_t + B_t - T_t}{\prod_{k=1}^t (1 + r_k^a)} \geq \sum_{t=1}^D \frac{C_t}{\prod_{k=1}^t (1 + r_k^a)}.$$

125. This occurs because marginal utility approaches infinity as consumption approaches zero. This also implies that the nonnegativity constraints in equation (4) are non-binding.

### Uncertain Income

For most households, labor earnings comprise by far the largest portion of lifetime income, and diversifying into financial assets with returns that are negatively correlated with their labor income is difficult. In the model, larger and more persistent innovations in labor income will generate a larger impact on precautionary saving.

The log of labor earnings,  $y_t = \ln(Y_t)$ , is assumed to be normally distributed, with a stochastic process given by:

$$(6) \quad y_t = \bar{y}_t + u_t,$$

where realized log labor earnings,  $y_t$ , is the sum of the mean of log earnings,  $\bar{y}_t$ , and a random term,  $u_t$ .

To capture the observed persistence of log earnings shocks over time, the random component of log labor earnings is specified as a first-order autoregressive, or AR(1), process:

$$(7) \quad u_t = \rho u_{t-1} + \varepsilon_t,$$

where  $\varepsilon_t$  is Gaussian white noise:

$$(8) \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2).$$

The AR(1) process for the earnings shock implies that future log earnings realizations depend on past realizations:

$$y_{t+1} = \bar{y}_{t+1} + \rho(y_t - \bar{y}_t) + \varepsilon_{t+1},$$

and that the conditional expectation of future log earnings is a function of current log earnings:

$$E(y_{t+1}|y_t) = \bar{y}_{t+1} + \rho(y_t - \bar{y}_t).$$

As a result, current labor income,  $Y_t$ , helps forecast future labor income  $\{Y_{t+1}, \dots, Y_D\}$  and is a state variable, along with wealth,  $W_t$ , and age,  $t$ , in the household's optimization problem.

### Optimal Consumption and Saving

The value function,  $V(W_t, Y_t, t)$ , for the household's dynamic programming problem at any age (or time)  $t$  is defined as the maximized value of expected lifetime utility:

$$(9) \quad V(W_t, Y_t, t) \equiv \max_{\{C_{t+j}\}_{j=0}^{D-t}} E_t \sum_{j=0}^{D-t} \pi_t^{t+j} (1 + \delta)^{-j} U(C_{t+j}; \theta),$$

subject to the wealth transition equation (3) and the net asset constraint equation (4), which are summarized by the budget constraint equation (5). The value function can be rewritten as the following recursive equation (the Bellman equation):

$$(10) \quad V(W_t, Y_t, t) = \max_{\{C_t\}} [U(C_t; \theta) + \frac{\pi_t^{t+1}}{1 + \delta} E_t V(W_{t+1}, Y_{t+1}, t + 1)].$$

Using the net asset constraint in period  $t$ , equation (4), and substituting the wealth transition equation for  $W_{t+1}$ , equation (3), into the value function, the Lagrangian associated with the maximization problem on the right-hand side of equation (10) can be written as:

$$(11) \quad \mathcal{L} = U(C_t; \theta) + \frac{\pi_t^{t+1}}{1 + \delta} E_t V[(W_t - C_t)(1 + r_{t+1}^a) + Y_{t+1} + B_{t+1} + I_{t+1} - T_{t+1}, Y_{t+1}, t + 1] + \lambda_t (W_t - C_t)$$

where  $\lambda_t$  is the shadow price of the net asset constraint (4) in period  $t$ . The Kuhn-Tucker conditions for this maximization problem are:

$$(12a) \quad \frac{\partial \mathcal{L}}{\partial C_t} = U_c(C_t; \theta) - \frac{\pi_t^{t+1} (1 + r_t^m)}{1 + \delta} E_t V_w(W_{t+1}, Y_{t+1}, t + 1) - \lambda_t = 0$$

$$(12b) \quad W_t - C_t \geq 0, \lambda_t \geq 0$$

$$(12c) \quad \lambda_t (W_t - C_t) = 0,$$

where the subscripts on the functions  $U(\cdot)$  and  $V(\cdot)$  denote the partial derivative of the function with respect to the subscripted variable, and  $r^m$  is the marginal after-tax rate of return. Equation (12a) can be rewritten as:

$$(13) \quad U_c(C_t; \theta) = \frac{\pi_t^{t+1} (1 + r_t^m)}{1 + \delta} E_t V_w(W_{t+1}, Y_{t+1}, t + 1) + \lambda_t.$$

Define  $C_t^*$  as the optimal choice of consumption at time  $t$ , and note that it is a function of the state variables  $W_t, Y_t$ , and  $t$ :  $C_t^*(W_t, Y_t, t)$ . Substituting  $C_t^*(W_t, Y_t, t)$  into the Bellman equation (10) and using the wealth transition equation (3) to define  $W_{t+1}$ , gives:

$$V(W_t, Y_t, t) = U(C_t^*; \theta) + \frac{\pi_t^{t+1}}{1 + \delta} E_t V[(W_t - C_t^*)(1 + r_t^a) + Y_{t+1} + B_t + I_t - T_t, Y_{t+1}, t + 1].$$

Differentiating this expression for the Bellman equation with respect to  $W_t$  yields:

$$V_w(W_t, Y_t, t) = U_c(C_t^*; \theta) \frac{\partial C_t^*}{\partial W_t} + \frac{\pi_t^{t+1}}{1 + \delta} E_t V_w[(W_{t+1}, Y_{t+1}, t + 1)(1 + r_t^a)(1 - \frac{\partial C_t^*}{\partial W_t})].$$

Rewriting the right-hand side of this expression gives:

$$V_w(W_t, Y_t, t) = \frac{\pi_t^{t+1}(1 + r_t^m)}{1 + \delta} E_t V_w(W_{t+1}, Y_{t+1}, t + 1) + [U_c(C_t^*; \theta) - \frac{\pi_t^{t+1}(1 + r_t^m)}{1 + \delta} E_t V_w(W_{t+1}, Y_{t+1}, t + 1)] \frac{\partial C_t^*}{\partial W_t}.$$

From the Kuhn-Tucker condition, equation (12a), the coefficient in brackets on  $\partial C_t^*/\partial W_t$  is equal to  $\lambda_t$ , so after substitution:

$$V_w(W_t, Y_t, t) = \frac{\pi_t^{t+1}(1 + r_t^m)}{1 + \delta} E_t V_w(W_{t+1}, Y_{t+1}, t + 1) + \lambda_t \frac{\partial C_t^*}{\partial W_t}.$$

If the net asset constraint in equation (4) is binding, the Kuhn-Tucker conditions in equations (12b) and (12c) show that:

$$\lambda_t > 0 \text{ and } C_t^* = W_t.$$

Therefore:

$$\frac{\partial C_t^*}{\partial W_t} = 1 \text{ and } \lambda_t \frac{\partial C_t^*}{\partial W_t} = \lambda_t.$$

If the constraint is not binding, then:

$$\lambda_t = 0 \text{ and } \lambda_t \frac{\partial C_t^*}{\partial W_t} = 0 = \lambda_t.$$

Substituting  $\lambda_t$  for  $\lambda_t(\partial C_t^*/\partial W_t)$  into the above expression for  $V_w(W_t, Y_t, t)$  yields:

$$(14) \quad V_w(W_t, Y_t, t) = \frac{\pi_t^{t+1}(1+r_t^m)}{1+\delta} E_t V_w(W_{t+1}, Y_{t+1}, t+1) + \lambda_t.$$

Using the expression in equation (14), substituting into the Kuhn-Tucker condition in equation (13), and pushing the time subscript up one period:

$$(15) \quad U_c(C_{t+1}; \theta) = V_w(W_{t+1}, Y_{t+1}, t+1).$$

Substituting this expression in equation (15) into equation (14) yields the following Euler equation:

$$(16) \quad U_c(C_t; \theta) = \frac{\pi_t^{t+1}(1+r_t^m)}{1+\delta} E_t U_c(C_{t+1}; \theta) + \lambda_t,$$

which holds for all time periods  $t = 1, \dots, D$ . This Euler equation (16) shows that when a household is unconstrained by the asset restriction ( $\lambda_t = 0$ ), the household is indifferent between consuming one unit in time  $t$ , thus increasing utility by  $U_c(C_t; \theta)$ , and saving one unit in time  $t$  and having  $(1+r_t^m)$  units to consume in time  $t+1$ , which would increase expected utility by  $\pi_t^{t+1}(1+\delta)^{-1} E_t [U_c(C_{t+1}; \theta)]$ .<sup>126</sup> Marginal utility in period  $t+1$  is an expected value, since  $Y_{t+1}$  is stochastic. When borrowing is constrained ( $\lambda_t > 0$ ), the consumer would like to shift resources from the future (borrow) where the expected marginal utility of one unit of consumption is valued less than one unit of current consumption. The net asset constraint prohibits this resource transfer, and  $\lambda_t$  represents the marginal utility that would be gained if the borrowing constraint could be relaxed. In the terminal period  $D$ , because the functional form of the utility function implies nonsatiation, there is no explicit bequest motive, and  $\pi^{D+1} = 0$ , so that the value function goes to zero in  $D+1$ , and the individual consumes all remaining wealth,  $C_D = W_D$ . Therefore:

$$U_c(C_D; \theta) = U_c(W_D; \theta) = \lambda_D.$$

The net asset constraint is always binding in the terminal time period, since the marginal utility of any wealth in  $D+1$  is zero.

126. If the individual's life span is certain, optimal consumption satisfies the unconstrained Euler equation with  $\pi^{t+1} = 1$ .

### Model Solutions

If income is nonstochastic, a closed-form solution for optimal consumption at time  $t$  can be derived using the Euler equation (16) and the budget constraint equation (5'), and is equal to:

$$(17) \quad C_t = \frac{W_t + \sum_{j=t}^D [(Y_j + B_j + I_j - T_j) \prod_{k=t}^j (1 + r_k^a)]}{H_t},$$

where:

$$H_t = \sum_{j=t}^D \frac{\left[ \frac{\pi_t^j}{(1 + \delta)^{j-t}} \prod_{k=t}^j (1 + r_k^a) \right]^{1/\gamma}}{\prod_{k=t}^j (1 + r_k^a)}.$$

If life span is also nonstochastic, then the solution for optimal consumption has the same functional form, but  $\pi_t^j = 1$  and  $I_j = 0$ .

With stochastic income and constant relative risk aversion utility, the consumer's dynamic programming problem cannot be solved analytically, but only described by the Euler conditions in equation (16). A numerical algorithm that takes advantage of the recursive nature of the dynamic programming problem and the finite lifetime of the consumer is used to solve for optimal consumption and saving. The general idea is to discretize the state space and then, using the terminal boundary condition for net assets and the Euler equation, recursively solve the consumption/saving problem backward from the terminal time period. The first step in the solution algorithm is to discretize the state space that includes wealth, current earnings, and age. It is necessary to limit the state space to a finite grid in order to make numerical computation feasible, but it is also important to make sure that the state space is broad enough to encompass virtually all possible realizations of wealth and earnings.

Given the mean earnings path for an individual, along with the stochastic process for earnings described in equations (7) and (8), an upper bound on the variance for (log) labor earnings is expressed as:

$$\text{var}(y) = \sigma_\varepsilon^2 [1/(1 - \rho^2)],$$

and therefore the maximum and minimum values for the net earnings state space in each time period are approximated by:

$$(18) \quad \begin{aligned} Y_t^{max} &= \exp\{(\bar{y}_t) + \chi\sigma_\varepsilon[1/(1 - \rho^2)]^{1/2}\} \\ Y_t^{min} &= \exp\{(\bar{y}_t) - \chi\sigma_\varepsilon[1/(1 - \rho^2)]^{1/2}\} \end{aligned}$$

where  $\chi$  is defined as the number of standard deviations used to extend the state space above and below the mean; in practice,  $\chi$  is usually set to 2.5.<sup>127</sup>

The limits for the wealth state space are determined by using the stochastic process for labor earnings, and assuming a minimum level of consumption in each period, so that the mean and the variance for wealth in each period are approximated using the wealth transition equation. The mean level of wealth,  $\bar{W}_t$ , in any period  $t$  is:

$$\bar{W}_t = (\bar{W}_{t-1} - C_{t-1}^{min})(1 + r_{t-1}^a) + \bar{Y}_t + \bar{B}_t + \bar{I}_t - \bar{T}_t.$$

Minimum consumption in period  $t$ ,  $C_t^{min}$ , is defined as the solution for optimal consumption if the individual knows that labor earnings will be zero from  $t + 1$  to  $D$ , so that existing mean wealth must be spread out over the rest of the individual's lifetime:

$$C_t^{min} = \bar{W}_t / H_t,$$

where  $H_t$  is defined as above. The variance of wealth in period  $t$  is a function of the variance of net labor earnings in period  $t$  and all previous periods,  $t - 1$  to 1. Substituting the expression for  $C_t^{min}$  into the wealth transition equation, the variance of wealth in period  $t$  can be derived, noting that  $\text{var}(B) = \text{var}(I) = 0$ , and is expressed as:

$$\text{var}(W_t) = \sum_{j=1}^t \{ [1 - (1/H_t)]^{2(t-j)} [\prod_{k=j}^{t-1} (1 + r_k^a)^2] \text{var}(Y_j - T_j) \}.$$

The minimum and maximum values for wealth in period  $t$  are then:

$$(19) \quad \begin{aligned} W_t^{max} &= \bar{W}_t + \chi[\text{var}(W_t)]^{1/2} \\ W_t^{min} &= \max(0, \{ \bar{W}_t - \chi[\text{var}(W_t)]^{1/2} \}), \end{aligned}$$

where the lower bound of wealth must be nonnegative because of the net asset constraint, and once again  $\chi$  represents the number of standard deviations above or below the mean; in practice, this parameter usually is set at

127. The earnings state space encompasses about 99 percent of all possible earnings realizations when  $\chi = 2.5$ . Expanding the size of the earnings space increases computational time but does not noticeably change the numerical calculation.

3.0.<sup>128</sup> To allow numerical computation, the wealth and earnings state spaces are discretized between the maximum and minimum values. The wealth state space is discretized on a log scale so that the steps are smaller for lower levels of wealth, since the consumption function will exhibit more curvature at low wealth levels.<sup>129</sup> The time, or age, state space is already discrete by definition.

Optimal consumption is calculated as follows. Using the terminal constraint, consumption in the final period will be equal to realized wealth;  $C_D = W_D$ . Since labor is not supplied by the individual during the retirement phase of the life cycle, and pension benefits and inheritances are nonstochastic during retirement, optimal consumption can be calculated analytically for any wealth level,  $W_p$ , at any age  $t = R, \dots, D$ , by:

$$C_t = W_t/H_t.$$

If retirement income is stochastic, the numerical procedure described below can be used to calculate optimal consumption during retirement.

During the working phase of the life cycle,  $t = 1, \dots, R - 1$ , optimal consumption is numerically calculated for discrete values over the two-dimensional grid of wealth and labor income in each time period  $t$  by iterating on the Euler equation and working backward through the consumer's lifetime. Values of consumption for wealth and income levels between the discrete steps are linearly interpolated. For instance, the Euler equation (16) is solved at time  $R - 1$  by searching numerically for a value of  $C_{R-1}$  that solves the Euler equation conditional on the state variables  $W_{R-1}$  and  $Y_{R-1}^*$ , and the boundary constraint on assets in equation (4). The expected marginal utility of optimal next-period consumption,  $C_R$ , is integrated numerically as a function of expected wealth,  $W_R$ , and income,  $Y_R$ , in period  $R$ , using the wealth transition equation (3) and the stochastic process for earnings in equations (6) through (8). If the net asset constraint is nonbinding, the marginal utility of current consumption equals the discounted value of the expected marginal utility of optimal consumption in the next period. The optimal choice of consumption is equal to wealth if the net asset constraint is binding. Optimal values of  $C_{R-1}$  are found for

128. Increasing  $\chi$  so that the size of the wealth space is larger does not noticeably alter the solution for optimal consumption but does increase computation time.

129. Optimal consumption increases more rapidly at low levels of wealth when income is uncertain; see Zeldes (1989a), for example.



the different discrete levels of  $W_{R-1}$  and  $Y_{R-1}$ , and optimal consumption is approximated for levels of wealth and income between these levels. The procedure is then repeated for period  $R - 2$ , and the algorithm continues back inductively to the initial period. The result is a numerically approximated optimal consumption function that is a function of a three-dimensional grid of state variables: wealth,  $W_t$ ; current labor income,  $Y_t$ ; and time (age),  $t$ .

A random number generator is used to simulate (log) labor income shocks from a normal distribution, so that a realized labor income path is projected for an individual. For a given realization of earnings, an individual's corresponding lifetime consumption and saving decisions are calculated. This procedure is repeated for 10,000 households in a given cohort, who each receive different earnings shocks drawn from the same distribution, and thus end up with different realized lifetime earnings, consumption, and wealth profiles over the life cycle.

## APPENDIX B

### *Selected Parameter Specifications*

#### **Social Security and Pension Replacement Rates**

Determining the appropriate replacement rates to use is quite difficult, because the rates vary considerably depending on economic circumstances, and because the basic unit of observation in the model and in the data is the household, whereas most available replacement rate information is based on individual workers.

#### *Defined Benefit Pensions*

We calculate DB replacement rates by final earnings and education level using data from the 1992 Health and Retirement Survey. The HRS core survey collects data on job histories, pension coverage, and pension plan specifics. The HRS also collects information directly from the employers of HRS respondents with pension coverage. The Institute for Social Research (ISR) at the University of Michigan developed a software program to evaluate this information. The program combines information on job histories from the core survey, detailed pension formulas from the employers, and user-defined macroeconomic assumptions. The

program also takes into account any integration provisions with social security. One of the variables created by the ISR program is the annual pension benefit as a percentage of final wages—the replacement rate.

Our analysis includes 3,324 workers in the HRS who report having DB pension coverage on their current job. Of these, however, about one-third do not have an employer-provided pension plan match. We use a hot deck procedure to impute a plan for these respondents, based on industry and occupation. Our assumptions regarding interest, inflation, and wage growth correspond to the intermediate assumptions of the 1995 Social Security Administration Trustees report.<sup>130</sup>

To be consistent with our base simulation model, we assume retirement at age sixty-two. The replacement rates calculated here reflect DB pensions on the current job as well as DB pensions from previous jobs, for the sample of workers with current DB coverage.

For workers with sixteen or more years of education, we find that the average replacement rate (projected annual pension benefits divided by projected final earnings) is 37.6 percent when weighted by population weights and by final earnings, and 37.1 percent when weighted only by population weights. For workers with less than sixteen years of education, the averages are 30.9 percent and 29.7 percent, respectively. Hence, we use values of 37 percent for college-educated households and 30 percent for other households.

### *Social Security Replacement Rates*

Unfortunately, we were unable to access the Social Security Earnings Benefit Files associated with the HRS for purposes of this project. We hope to explore this data source in future work. Instead, we gathered information on social security replacement rates from several sources. The Social Security Administration (SSA) defines low earners as those earning 45 percent of the average covered wage, medium earners as those earning 100 percent of the average covered wage, high earners as those earning 160 percent of the average covered wage, and maximum earners as those earning the payroll tax maximum. As an example, in 1997 these wages corresponded to earnings of \$12,341, \$27,426, \$43,881, and \$65,400, respectively. The SSA reports replacement rates for workers

130. U.S. Board of Trustees of the Federal OASDI Trust Funds (1995).

who are either low, medium, high, or maximum earners throughout their career. The replacement rates for retirees who first receive benefits at age sixty-five in 2000 are projected to be about 53 percent, 40 percent, 32 percent, and 24 percent of final earnings, respectively.<sup>131</sup>

These figures suggest, at first glance, that it would be appropriate to use very high replacement rates. However, several issues arise in attempting to convert these results into replacement rates in our model. First, the wage profiles used by the Social Security Administration are hypothetical and highly unlikely. Second, they are not delineated by education class. Third, the figures are for individuals who retire at age sixty-five, rather than for households that retire at age sixty-two. The earlier retirement age in the model suggests reducing the replacement rates by one-fifth. Consideration of a spouse suggests raising the replacement rates by some fraction between zero and one-half. It is useful to note, however, that the SSA replacement rates have fallen by about 15 to 20 percent (not percentage points) since 1980, when they were 68 percent, 51 percent, and 32 percent, for low, medium, and high earners, respectively. By 2040 they are projected to fall to 49 percent, 37 percent, and 24 percent, respectively, for those retiring at age sixty-five. This represents a 20 percent decline for the first two groups, and virtually no change for the high earners.

Grad provides evidence of replacement rates earned by newly retired workers in 1982.<sup>132</sup> As she notes, one of the advantages of using actual replacement rates is that they reflect “the complexities of real life workers.” Grad shows that for retiring men and their wives, the median social security replacement rate was 49 percent of final earnings.<sup>133</sup> For retiring women and their husbands, the analogous rate was 62 percent. An average of these rates is 55 percent. Reducing this figure by 20 percent to account for changes in social security benefits between 1980 and 1990–2000, as determined by the SSA figures above, and by another 20 percent to allow for retirement at age sixty-two in the model, suggests median replacement rates for current retirees of 35 percent.

Using this figure as our base, we adopt replacement rates in the model that are, if anything, lower than warranted. We use a 35 percent replacement rate for the less educated group, even though, on average, they are

131. U.S. Board of Trustees of the Federal OASDI Trust Funds (1999).

132. Grad (1990).

133. Grad (1990, table 17).

likely to have lifetime incomes below the median and thus replacement rates higher than the median, because social security is progressive. We use a 21 percent replacement rate for highly educated households. This is approximately what high-earner households who retire at age sixty-two would receive, based on the SSA data above.

Other sources of data confirm the general validity of our estimates. Wiatrowski uses information on employee benefits in medium and large private establishments in 1989.<sup>134</sup> He finds that total pension income (from the worker's DB plan on the current job, and from the worker's and his or her spouse's social security payment) at age sixty-two ranges between 36 and 60 percent of final earnings by earnings group for full-time workers with twenty years of service. It ranges between 48 and 81 percent for those with thirty years of service.

Second, using data from the Panel Survey of Income Dynamics, Laibson, Reppetto, and Tobacman (1998) estimate replacement rates of 41 to 45 percent for households where the head has less than sixteen years of education, and 55 percent for college graduates. These estimates can be compared with ours in table 1 by weighting the replacement rates in that table by the proportions of households with and without pensions. This generates replacement rates from our model of just under 50 percent for households with less than sixteen years of education and 43 percent for households with sixteen years of education or more. Thus our replacement rates are slightly higher than those in Laibson, Reppetto, and Tobacman for households with less than sixteen years of education and slightly lower for those with more education.

Third, Hubbard, Skinner, and Zeldes (1994) use data from the same survey to estimate labor income and retirement income. Using data from their appendix tables A-2 and A-3, the implied initial replacement rates are 61 to 80 percent of earnings at age sixty-five for those with less than sixteen years of education and 81 percent for those with sixteen years or more. Using earnings at age sixty-two as the base, the replacement rates are 49 to 72 percent and 59 percent, respectively. These figures are, if anything, higher than the replacement rates we use.

134. Wiatrowski (1991).

### **Time Preference Rates**

As noted in the text, we employ values of 3 percent and zero for the time preference rate, but we also examined whether negative values would be appropriate. There is some evidence that people answer survey questions in a manner consistent with having negative discount rates.<sup>135</sup> However, there are several important caveats to adopting negative time preference rates.

First, for every survey that elicits answers that appear to imply negative time preference rates, there are dozens or scores that generate positive time preference rates. There is a large literature that suggests that the time preference rates people employ can vary depending on the magnitude, sign, time delay, and framing issues associated with the discounted event.<sup>136</sup> But the literature almost universally suggests that time preference rates are positive.

Second, there is little evidence that people actually behave as if they had negative time preferences. George Loewenstein and Richard Thaler point to income tax withholding, the preference of teachers to take their nine-month salary over twelve months, and the existence of upward-sloping age-consumption profiles.<sup>137</sup> But even with a positive time preference, overwithholding can be explained by such factors as biased withholding schedules, income uncertainty, itemization of deductions, and fear of penalties and audits. The preference for taking salary over twelve months may have to do more with administrative simplicity than with discount rates. Upward-sloping consumption profiles are standard fare in simulation models with stochastic earnings and precautionary saving, even with positive time preference rates (for example, see figures 3 and 6).

Andrew Samwick estimates the distribution of time preference rates that would reconcile wealth-earnings ratio data in the 1992 Survey of Consumer Finances with the wealth patterns obtained in a stochastic simula-

135. Loewenstein and Sicherman (1991); Barsky and others (1997). In their analysis using 198 households from the Health and Retirement Survey, Barsky and others actually find that the median and modal time preference rates are zero, but that the mean discount rate is negative.

136. Loewenstein and Thaler (1989) and Loewenstein and Elster (1992) review the literature. Laibson, Reppetto, and Tobacman (1998) study saving behavior in a model that incorporates some of these considerations.

137. Loewenstein and Thaler (1989).

tion model.<sup>138</sup> In his preferred specification, he finds a wide range of time preference rates, with a median of 7.63 percent and an interquartile range of 2.93 percent to 14.66 percent. He finds that fewer than 10 percent of the sample have estimated negative discount rates. He suggests that these findings probably represent households that either received inheritances or plan to leave them—factors not included in Samwick’s simulation model—rather than households with truly negative discount rates. More generally, any factor that raises wealth, but that is not included in the simulation, could contribute to a negative estimated time preference rate. For example, households that received unusually high returns on previous investments would, other things equal, have an estimated time preference rate that was lower than their true rate.

A third problem is the same theoretical concern that affects using a zero discount rate. In particular, as Loewenstein and Drazen Prelec note, “Applied uniformly to all choices, negative time preference would require harsh reductions in present consumption in favor of the future. The fact that one does not observe such sacrifices, even given the additional inducements of a positive interest rate, is normally taken as evidence for positive time preferences.”<sup>139</sup> Loewenstein and Prelec propose a different interpretation of the evidence: “negative time preference is applied selectively, to those events that are seen as part of a meaningful sequence, having a well-defined starting and ending point.” They argue that “the salience of particular intervals . . . depends on . . . perceptual framing.” But if it applied to retirement saving, negative time preference *would* require the harsh reductions in current consumption that Loewenstein and Prelec appear to reject when “applied uniformly to all choices.”

The implied consumption paths seem extreme. With no borrowing constraints, no uncertainty, a 3 percent interest rate, and a time preference rate of –3 percent, the consumer would set the marginal utility of consumption at age twenty-five to be thirty-three times as high as marginal utility at age eighty-five. If the intertemporal elasticity of substitution is one-third, the consumption level at age eighty-five would be 3.3 times as high as at age twenty-five (in the absence of productivity growth or changes in family size), and 2.2 times as high as at age forty-five.

138. Samwick (1998).

139. Loewenstein and Prelec (1991, p. 349). See also Koopmans (1960) and Olson and Bailey (1981).

Loewenstein and Nachum Sicherman claim only that negative time preference cannot be dismissed out of hand in certain circumstances.<sup>140</sup> We do not disagree, but neither do we believe that negative time preference in the overall saving decision has been sufficiently established to merit being used as a benchmark in the simulations.

## APPENDIX C

### *Calculation of Pension Wealth in the Health and Retirement Survey*

THE HRS COLLECTED detailed pension plan information for about two-thirds of respondents who reported pension coverage on a current or previous job. This information was gathered from the respondents' employers or from Summary Plan Description data from the U.S. Department of Labor. The Institute for Social Research (ISR) at the University of Michigan developed a software program that uses this information in conjunction with user-defined macroeconomic assumptions to estimate the present value of future pension benefit payments.

We estimate pension wealth from current jobs, and separately we estimate wealth from DB and DC pension plans, including both 401(k) and non-401(k) plans. To estimate DB pension wealth, we use the restricted pension plan data, the ISR software program, and the long-term intermediate assumptions in the 1995 Social Security Trustees Report.<sup>141</sup> We impute plans to the one-third of HRS respondents with DB plans who lack a pension plan match using a hot deck match based on industry and occupation. DB wealth from the current job reflects work to 1992. This understates DB wealth, since no credit is given for expected future accruals.

We estimate DC pension wealth on the current job using self-reported account balances. Previous research suggests that using the restricted employer-provided pension plan data does not improve upon the self-

140. Loewenstein and Sicherman (1991).

141. U.S. Board of Trustees of the Federal OASDI and Disability Insurance Trust Funds.

reported account balance data.<sup>142</sup> DC wealth reflects self-reported DC balances if these are given, or imputed DC balances if they are not. When imputing missing account balances, we take advantage of the longitudinal nature of the HRS by incorporating wave 2 self-reported account balances when available. For workers with missing wave 1 DC balances who report balances in wave 2, we estimate their wave 1 balance as the wave 2 balance less any contributions and interest earned between the two waves. Contributions are based on self-reported employee and employer contribution rates, if available. If unavailable, we use the sample's median contribution rates of 4.0 percent for employee contributions and 2.0 percent for employer contributions. We also account for increasing wages by assuming nominal wage growth rate of 4.9 percent in 1992. We assume a nominal rate of return of 7.1 percent in 1992 and 6.1 percent in 1993.

We use a regression-based imputation procedure to estimate missing DC account balance information for those missing such information in both wave 1 and wave 2. We estimate a log-linear model of account balances based on wages, employer and employee contribution rates, tenure, occupation, full-time status, sex, and marital status.

142. Johnson, Sambamoorthi, and Crystal (forthcoming).



## APPENDIX D

*Supplemental Tables***Table D1. Median Wealth-Earnings Benchmarks**

Age	<i>Education &lt; 16 years</i>		<i>Education ≥ 16 years</i>	
	<i>No pension</i>	<i>Pension</i>	<i>No pension</i>	<i>Pension</i>
Bernheim and Scholz				
50–54	1.9	1.3	2.2	1.5
55–59	2.4	1.8	2.8	2.2
60–64	3.2	2.4	3.4	2.6
This paper				
	<i>Time preference rate = 3 percent</i>			
50–54	2.70	1.97	2.40	1.39
55–59	3.76	2.66	3.67	2.19
60–62	4.74	3.28	4.91	2.92
	<i>Time preference rate = 0</i>			
50–54	4.10	3.32	3.41	2.35
55–59	5.24	4.30	4.82	3.39
60–62	6.47	5.16	6.20	4.37

Source: Bernheim and Scholz (1993); author's calculations.

**Table D2. Distribution of Simulated Wealth-Earnings Ratios by Age Among Households with Less Than Sixteen Years of Education and Without Private Pensions**

Age	<i>5th percentile</i>	<i>25th percentile</i>	<i>Median</i>	<i>75th percentile</i>	<i>95th percentile</i>
	<i>Time preference rate = 3 percent</i>				
30–34	0.01	0.07	0.25	0.58	1.35
35–39	0.02	0.18	0.56	1.14	2.30
40–44	0.05	0.48	1.08	1.90	3.46
45–49	0.19	0.99	1.84	2.88	4.89
50–54	0.57	1.70	2.70	3.97	6.26
55–59	1.13	2.55	3.76	5.31	8.04
60–62	1.69	3.28	4.74	6.41	9.31
	<i>Time preference rate = 0</i>				
30–34	0.02	0.20	0.57	1.08	2.14
35–39	0.05	0.55	1.19	1.98	3.46
40–44	0.28	1.16	2.00	3.02	4.86
45–49	0.80	1.97	3.00	4.27	6.52
50–54	1.47	2.86	4.10	5.55	8.10
55–59	2.28	3.92	5.34	7.12	10.01
60–62	2.99	4.82	6.47	8.35	11.47

Source: Authors' estimates as described in the text.

**Table D3. Distribution of Simulated Wealth-Earnings Ratios by Age Among Households with Less Than Sixteen Years of Education and with Private Pensions**

Age	5th percentile	25th percentile	Median	75th percentile	95th percentile
<i>Time preference rate = 3 percent</i>					
30–34	0.01	0.07	0.23	0.52	1.23
35–39	0.02	0.14	0.46	0.96	2.04
40–44	0.04	0.33	0.83	1.56	3.01
45–49	0.10	0.65	1.36	2.32	4.25
50–54	0.24	1.06	1.97	3.17	5.39
55–59	0.42	1.53	2.66	4.14	6.92
60–62	0.41	1.85	3.28	4.97	7.89
<i>Time preference rate = 0</i>					
30–34	0.02	0.15	0.46	0.92	1.91
35–39	0.04	0.38	0.94	1.68	3.11
40–44	0.13	0.82	1.60	2.59	4.41
45–49	0.38	1.43	2.42	3.66	5.89
50–54	0.79	2.12	3.32	4.76	7.34
55–59	1.22	2.90	4.30	6.06	9.01
60–62	1.56	3.48	5.16	7.04	10.24

Source: Authors' estimates as described in the text.

**Table D4. Distribution of Simulated Wealth-Earnings Ratios by Age Among Households with Sixteen or More Years of Education and Without Private Pensions**

Age	5th percentile	25th percentile	Median	75th percentile	95th percentile
<i>Time preference rate = 3 percent</i>					
30–34	0.00	0.02	0.06	0.19	0.60
35–39	0.01	0.05	0.20	0.53	1.30
40–44	0.02	0.22	0.62	1.22	2.38
45–49	0.14	0.76	1.39	2.20	3.83
50–54	0.67	1.59	2.40	3.42	5.35
55–59	1.47	2.65	3.67	4.99	7.36
60–62	2.27	3.66	4.91	6.44	9.01
<i>Time preference rate = 0</i>					
30–34	0.00	0.03	0.14	0.39	1.03
35–39	0.01	0.16	0.52	1.04	2.09
40–44	0.09	0.65	1.24	2.00	3.40
45–49	0.55	1.44	2.23	3.21	5.06
50–54	1.35	2.42	3.41	4.59	6.76
55–59	2.31	3.63	4.82	6.33	8.91
60–62	3.18	4.73	6.20	7.88	10.66

Source: Authors' estimates as described in the text.

**Table D5. Distribution of Simulated Wealth-Earnings Ratios by Current Earnings Among Households with Less Than Sixteen Years of Education and Without Private Pensions<sup>a</sup>**

<i>Earnings quintile</i>	<i>5th percentile</i>	<i>25th percentile</i>	<i>Median</i>	<i>75th percentile</i>	<i>95th percentile</i>
<i>Ages 50–54</i>					
Lowest	0.15	1.26	2.77	4.56	7.68
2nd	0.55	1.62	2.83	4.22	6.62
3rd	0.69	1.78	2.82	4.12	6.12
4th	0.82	1.80	2.72	3.79	5.55
Highest	0.99	1.81	2.48	3.38	4.88
<i>Ages 60–62</i>					
Lowest	1.36	3.81	5.87	8.05	11.26
2nd	1.81	3.62	5.23	7.09	9.65
3rd	1.65	3.41	5.06	6.54	8.68
4th	1.77	3.24	4.43	5.79	7.91
Highest	1.78	2.91	3.79	4.89	6.58

Source: Authors' estimates as described in the text.

a. These simulations employ a time preference rate of 3 percent.

**Table D6. Distribution of Simulated Wealth-Earnings Ratios by Current Earnings Among Households with Less Than Sixteen Years of Education with Private Pensions<sup>a</sup>**

<i>Earnings quintile</i>	<i>5th percentile</i>	<i>25th percentile</i>	<i>Median</i>	<i>75th percentile</i>	<i>95th percentile</i>
<i>Ages 50–54</i>					
Lowest	0.04	0.55	1.57	3.20	6.48
2nd	0.74	0.89	1.90	3.29	5.75
3rd	0.36	1.15	2.07	3.33	5.43
4th	0.46	1.25	2.10	3.15	4.95
Highest	0.68	1.40	2.03	2.91	4.45
<i>Ages 60–62</i>					
Lowest	0.15	1.25	3.41	5.92	9.61
2nd	0.36	1.74	3.48	5.50	8.43
3rd	0.49	1.96	3.59	5.22	7.57
4th	0.76	2.00	3.24	4.66	6.96
Highest	0.97	2.03	2.90	4.03	5.77

Source: Authors' estimates as described in the text.

a. These simulations employ a time preference rate of 3 percent.

**Table D7. Distribution of Simulated Wealth-Earnings Ratios by Current Earnings Among Households with Sixteen or More Years of Education Without Private Pensions<sup>a</sup>**

<i>Earnings quintile</i>	<i>5th percentile</i>	<i>25th percentile</i>	<i>Median</i>	<i>75th percentile</i>	<i>95th percentile</i>
<i>Ages 50–54</i>					
Lowest	0.22	1.27	2.42	3.83	6.46
2nd	0.67	1.53	2.49	3.64	5.64
3rd	0.79	1.69	2.51	3.54	5.30
4th	0.86	1.68	2.44	3.32	4.75
Highest	0.98	1.66	2.23	2.99	4.26
<i>Ages 60–62</i>					
Lowest	2.41	4.53	6.21	8.06	10.92
2nd	2.57	4.10	5.48	7.01	9.28
3rd	2.27	3.85	5.13	6.49	8.45
4th	2.26	3.51	4.58	5.73	7.53
Highest	2.13	3.07	3.88	4.84	6.39

Source: Authors' estimates as described in the text.

a. These simulations employ a time preference rate of 3 percent.

**Table D8. Sample Sizes by Age, Education, and Pension Status**

<i>Age</i>	<i>&lt; 16 years of education</i>		<i>≥ 16 years of education</i>		<i>All</i>
	<i>Pension</i>	<i>No pension</i>	<i>Pension</i>	<i>No pension</i>	
<i>1992 HRS</i>					
51–54	545	281	213	54	1,093
55–59	604	281	197	72	1,154
60–61	196	94	65	24	379
All households	1,345	656	475	150	2,626
<i>1992 SCF</i>					
25–29	28	62	7	31	128
30–34	39	80	36	49	204
35–39	50	76	35	74	235
40–44	49	60	45	77	231
45–49	34	62	61	97	254
50–54	38	60	49	72	219
55–59	27	38	39	54	158
60–62	13	20	15	34	82
All households	278	458	287	488	1,511

Source: Authors' calculations based on HRS and SCF data.

**Table D9. Sensitivity Analysis: Simulated Median Wealth-Earnings Ratios for Households Aged 60–62 Under Alternative Scenarios**

<i>Case</i>	<i>Education &lt; 16 years</i>		<i>Education ≥ 16 years</i>	
	<i>No pension</i>	<i>Pension</i>	<i>No pension</i>	<i>Pension</i>
Base case <sup>a</sup>	4.74	3.28	4.91	2.92
Changes to base-case parameters <sup>b</sup>				
Time preference rate = 0	6.47	5.16	6.20	4.37
IES = 0.5	3.95	2.38	4.29	2.23
IES = 0.25	5.25	3.83	5.32	3.38
Persistence parameter = 0.99	5.20	5.54	5.23	2.68
After-tax return = 5 percent	4.95	3.99	4.75	3.44
After-tax return = 1 percent	3.85	1.99	4.73	1.94
Changes to wealth measures <sup>b</sup>				
Exclude business wealth	4.74	3.28	4.91	2.92
40 percent decline in stock market	4.74	3.28	4.91	2.92
30 percent cut in social security benefits	5.52	3.91	5.36	3.28
Add expected inheritances	4.74	3.28	4.91	2.92
Retirement at age 65	4.73	3.14	5.23	2.97
Social security and pensions as function of household's final earnings	4.80	3.28	4.94	2.74
Changes to consumption needs <sup>b</sup>				
20 percent increase in all simulated wealth-earnings ratios	5.69	3.93	5.89	3.50
10 percent increase in survival rates	6.14	4.82	6.01	4.22
Substitution of Bernheim-Scholz targets <sup>b</sup>	2.40	3.20	3.40	2.60

Source: Authors' calculations based on the simulation model described in the text.

a. The parameters of the base case are as follows: time preference rate = 3 percent, intertemporal elasticity of substitution (IES) = 0.33, autoregressive persistence parameter = 0.85, retirement at age sixty-two, a real after-tax rate of return of 3 percent, and social security and pension income derived from the average final earnings of one's own education class.

b. Specifications of the sensitivity analysis are described in the text.

## *Comments and Discussion*

**Christopher D. Carroll:** This paper represents a major accomplishment: it provides the first really credible measure of the extent to which Americans are saving adequately for retirement. Although previous work by Douglas Bernheim and his coauthors has compared observed behavior with the predictions of a certainty-equivalent model, a large literature has recently shown that the optimal solution to the consumption-saving problem changes dramatically when a realistic treatment of uncertainty is incorporated. This paper is the first to compare observed behavior with the prescriptions of a model that takes uncertainty seriously and that is otherwise calibrated to capture a large proportion of the really important features of the lifetime saving problem.

My comments will proceed in three parts. First, I will evaluate how well the model under the authors' baseline parameters is able to match actual data. Second, I will discuss the significance that should (or should not) be attached to the model's ability to match the behavior of the median household under the authors' baseline set of parameter values. Finally, I will discuss whether we should believe that households are truly behaving optimally even if the model under the baseline parameters can match the data.

The fundamental conclusion of the paper is that observed profiles of median household wealth by age look reasonably similar to median profiles that emerge from the dynamic optimization model, if the time preference rate is set at 3 percent. I have argued essentially this point myself in an earlier paper on the basis of a much simpler stochastic optimization model (with a baseline time preference rate of 4 percent).<sup>1</sup> Therefore I am

1. Carroll (1997a).

relieved to see that, broadly speaking, the conclusion holds up in the face of the many important improvements the authors' model contains (such as realistic treatment of taxes, mortality risk, and pensions). The main way in which the authors assess whether their model fits the data is by comparing the profile of median wealth by age predicted by the model with the median profiles observed in the data. They find a reasonably good match. For example, in table 5, using their intermediate definition of wealth, they find that very close to 50 percent of households in the HRS data have wealth above the median predicted by the model. (If the model were exactly true, the figures in this column would all be identically 50 percent.) For the narrow measure of wealth, the proportions tend to be near 40 percent, and for the broad measure they tend to be near 60 percent, but the intermediate measure seems to me the most plausible measure to use for the purposes of these comparisons.

The fit of the model to the data is not perfect, however, and there are some systematic misses. In particular, tables 9 and 11 suggest that the model under its baseline parameterization underpredicts saving at younger ages. For example, in table 9, from ages twenty-five to forty, typically about 65 to 70 percent of households have more wealth than the model predicts, and the figures come down to the 50 percent range when consumers are in their fifties.

A plausible explanation can be found for this underprediction, however: there is good reason to believe that the model's specification understates the amount of income risk that young households face. In particular, the model's income process is a first-order autoregressive, or AR(1), process with a serial correlation coefficient of 0.85. But a substantial body of research beginning with Thomas MaCurdy finds that household income processes have both completely transitory and completely permanent elements.<sup>2</sup> Table 12 presents a sensitivity analysis of the authors' model that shows that when the serial correlation coefficient is increased to 0.99, thus approaching a random walk, the optimal amount of precautionary saving rises substantially. My own interpretation of this pattern of facts is that the assumption of an AR(1) income process understates the amount of precautionary saving that young households need to do, and thus overstates the adequacy of the assets that households actually have. If the

2. MaCurdy (1982).

income process were tweaked to make it somewhat more realistic, I suspect that even the relatively modest divergences between the model's projections and the data could be reconciled.

Unfortunately, the fact that there exists a set of parameter values (in this case, the authors' baseline parameters, perhaps modified slightly to generate more precautionary saving early in life) that can make the model match a particular fact like the median age-wealth profile may not be as impressive as it seems. The reason is what might be called the "dirty little secret" of the modern dynamic stochastic optimization model: with multiple realistic kinds of uncertainty, and with plausible assumptions about the other parameters, the model can predict a very wide range of behavior depending on the precise configuration of parameter values. The authors tend to downplay this point by focusing on the model's implications under their baseline parameter values, but it is important to realize that those baseline values are really no more compelling than many other parametric configurations that would yield quite different predictions for saving behavior.

This point echoes other recent discoveries about the flexibility of the modern version of the dynamic optimization model. For example, Andrew Samwick has shown that under some parameter values the model implies exactly the kind of limited offset between pension wealth and other forms of wealth that has been repeatedly found empirically.<sup>3</sup> Some of my own work has shown that the model can imply high marginal propensities to consume out of transitory income and low marginal propensities to consume out of human wealth. It can also explain the whole range of apparently contradictory results that have been found in the literature on estimating Euler equations.<sup>4</sup> In all of these cases, the results reconcile the dynamic optimization model with empirical findings that had previously been taken as rejections of optimality.

One can take either of two possible views of the newfound flexibility of the dynamic optimization model. The optimistic view is that such results revive dynamic optimization as a plausible description of behavior for the vast majority of households, after a period when regular empirical rejections of the certainty-equivalent versions of the model had put the descriptive accuracy of dynamic optimization in doubt. The less charitable

3. Samwick (1995).

4. Carroll (1997a, 1997b).



view springs from the fundamental scientific principle that, to be useful, a model must be falsifiable. If the dynamic optimization model can explain any conceivable observed behavior by varying the assumptions about essentially unobservable parameters like the time preference rate, it is not a very meaningful scientific model.

The latter view is certainly too extreme; certain patterns of behavior that we could in principle observe (such as saving rates of 90 percent in every year of life) would undoubtedly lead us to reject any plausible version of the model. But the critique contains a large element of truth as well, at least when the model is judged by its ability to match a single fact or a small set of facts. As the authors themselves showed in an earlier draft of the paper, the median asset-income ratio at age fifty-two for a set of consumers behaving exactly according to the model ranges from 0.04 to 2.59, depending on the configuration of parameter values. Given this flexibility, the fact that some set of parameter values exists that matches the actual median wealth-income ratio at age fifty-two (and selected other ages) is not as impressive as it might seem.

Thus, when examining the ability of the model to explain any particular fact, it is essential to examine how the model's predictions with regard to that fact change under alternative assumptions about parameter values. If we find that the model matches the fact under a broad range of parametric assumptions, the assertion that the model is truly capturing actual consumer behavior is more credible than if it can match that fact only under very specific assumptions.

The single most important parametric assumption in the paper is for the value of the time preference rate. Under the authors' baseline assumption that this rate equals 3 percent, the median household's saving is indeed roughly consistent with the predictions of the model. However, this conclusion is not particularly robust with respect to the assumption about the time preference rate. Using the zero-discounting benchmark, which I argue below reflects optimal behavior, the median household aged fifty-one to sixty-one in the HRS data has a wealth shortfall of about 1.3 years' worth of income.<sup>5</sup>

5. This figure is obtained by subtracting the actual median HRS wealth-earnings ratio using the intermediate definition in table 6 for people aged fifty-one to sixty-one, 3.59, from the model's median wealth for consumers of that age, 4.92.

The other appropriate response to the dirty little secret is to up the ante by requiring the model to be able to match a wide variety of facts all at once. Specifically, consider now the question of whether the model, even under a time preference rate of 3 percent, can match features of the data other than the median. The authors rightly emphasize that previous work comparing wealth distributions that result from a stochastic income process with predictions from a nonstochastic model are highly problematic, and that the only conceptually proper thing to do is to compare the distribution that arises from a stochastic model with the distribution that arises in the stochastic real world. However, one pervasive result in this paper is that the model cannot produce a wealth distribution that is nearly as wide as that observed in the data. For example, the comparison where the model has perhaps the closest match to the median data is for the HRS data under the intermediate measure of wealth (table 6). But despite the good fit between model median and HRS median wealth, the model predicts that the poorest 5 percent of households will hold almost six times as much wealth as they actually do (the model and actual wealth-earnings ratios are 0.96 and 0.17, respectively). The absolute gap is even greater at the 95th percentile, where the model predicts a wealth-earnings ratio of 7.78 whereas the actual ratio is 19.5. Of course, as the authors note, excessive wealth does not raise the alarms for public policy that inadequate wealth does. But the inability of even a thoroughly stochastic model to match the dispersion in wealth holdings suggests that some other feature of the model needs to be modified.

Recent work by Hubbard, Skinner, and Zeldes has shown that the lack of wealth at the bottom part of the distribution can be explained if the model is modified to take account of real-world features of the social safety net that discourage low-income people from saving.<sup>6</sup> And in a forthcoming paper I argue that, to match the extreme wealth accumulation at the top of the distribution, it is necessary to modify the model to incorporate a bequest motive in which bequests are a luxury good.<sup>7</sup> However, it must be confessed that such tweaking of the model to force it to fit more and more facts is very much subject to the criticism outlined above. A model loses its scientific status if every rejection by the data can be reconciled by postulating some new feature of the model that had heretofore

6. Hubbard, Skinner, and Zeldes (1995).

7. Carroll (2000).

been omitted. The only recourse is then to derive and test yet more implications of the newly modified model; such a tweaking-and-testing enterprise promises to be able to keep new generations of economists busy for quite a while.

Even if we accept that the authors' version of the model under baseline parameter values is a reasonably good match to the data for the median household, it is not at all clear where that leaves us in judging the optimality of that median household's behavior. The authors' position is that if actual behavior matches the baseline model's behavior under the baseline parameters, people are saving optimally for retirement. But I would certainly not want to take the baseline model's advice myself, because the authors' baseline value for the annual pure time discount factor is an appallingly high 0.97 (a number which is, appropriately, further reduced by mortality risk). Although this value is in the range typically assumed in the literature, consider the implications for a twenty-five-year-old trying to calculate optimal retirement saving. Since  $0.97^{40}$  is about 0.3, the implication is that the rational thing for the consumer to do, if the real long-term interest rate is zero, is to behave in such a way that when he retires at age sixty-five he will be so poor that the additional happiness yielded by spending an extra dollar is more than three times as great as when he was twenty-five.

To sharpen the intuitive case that this is not an optimal allocation, suppose that, on this hypothetical consumer's sixty-fifth birthday, scientists invent a time-travel device, and we bring together the twenty-five-year-old and the sixty-five-year-old versions of the same consumer for what we will politely term a dialogue. Imagine the difficulty the twenty-five-year-old would have in trying to persuade the sixty-five-year-old that the elder self's relative deprivation is really perfectly optimal.

Thinking of the two selves as being able to communicate helps make the point that the essence of time discounting is that the utility of the young "counts more" than the utility of the old. To see this in another light, consider a set of people standing in alphabetical order, and suppose we designate an outsider to allocate some good among them in a way that maximizes their total welfare. We would surely view as suboptimal any rule that allocated the good to people in a way that depended simply on their place in the line. Yet time discounting implies that when the line of people is considered to be a temporal succession of selves, the order matters critically. People near the front of the line receive up to four or five

times as much of the good (assuming linear marginal utility) as those at the end of the line. As a matter of logic and morality, there is little more reason to believe that it is optimal for the utility of the young selves to outweigh that of the old selves than to believe it is optimal for the utility of the alphabetically first to outweigh the utility of the alphabetically last.<sup>8</sup>

This point harkens back to the early debates surrounding the introduction of the dynamic optimization framework. Frank Ramsey, in his famous paper which asked how much a nation should save, argued that a positive rate of time preference was not only irrational but immoral as well.<sup>9</sup> He then proceeded to use a positive rate anyway, because his infinite-horizon model could not be solved with a time preference rate of zero.<sup>10</sup> But Ramsey's problem does not exist in this context: a life-cycle saving problem has a perfectly well defined solution with a time preference factor of one or even greater than one. My impression is that later economists argued for a positive time preference rate not so much on the grounds that it was rational or optimal as on the grounds that people were irrational and behaved suboptimally in a way that could be conveniently captured with a positive time preference rate. Economists now employ the assumption so regularly that it has come to seem rational, but my own instinct remains that a positive pure time preference rate is irrational for an individual optimization problem. Economists' gradual acceptance of a positive rate of time preference as rational may be a good example of what Senator Daniel Patrick Moynihan (D-NY) has called "defining deviancy down."

Another problem with interpreting the model's results as prescriptive of optimal behavior is that, in principle, when one solves this model, the assumption is that one has perfect certainty about what the relevant parameter values are—even the parameter values that describe the degree of uncertainty. That is, the assumption is that there is perfect certainty about the value of the underlying drift term in aggregate productivity growth, and about the future average age-income profile, and about the mean and variance and serial correlation properties of the stochastic shocks to income. Because the precautionary saving motive can be quite powerful in models

8. My thoughts on this subject have been stimulated partly by a compelling recent paper on the social discount rate by Caplin and Leahy (1999).

9. Ramsey (1928).

10. See also Irving Fisher's (1930) discussion arguing that a nonzero time preference rate is irrational, and see Barro (1999) for a recent discussion.

like this, it seems very likely that, under reasonable assumptions about the degree of uncertainty surrounding these parameter values, the model would prescribe substantially more saving than it does.

Another consideration further undermines our degree of confidence about the optimality or nonoptimality of the model's prescribed saving behavior. That is the fact that good reasons exist for questioning certain other important features of the model, in particular the assumption that there is a single, stable utility function over the life cycle. The authors discuss several reasons why this might not be so, including the possibility that the marginal utility of consumption is lower after retirement because work-related expenses such as commuting are diminished. Of course, it is also logically possible that the marginal utility of consumption rises after retirement, because retirees have more free time to spend, for example, on travel. Since we cannot measure marginal utility directly, it is hard to know precisely how to deal with this problem in the context of the model. The most valuable kind of evidence would probably come simply from asking retired consumers in the HRS or some other panel survey whether their standard of living has risen, fallen, or remained about the same during retirement.

Another way in which the standard model's structure may differ importantly from reality is in the assumption that utility is time-separable. There is a growing body of evidence that the utility an agent derives from consumption depends partly on the level of consumption that the agent has experienced in the past—that is, utility exhibits habit formation effects. G. M. Constantinides, A. B. Abel, and John Campbell and J. H. Cochrane have argued that habit formation can explain various asset pricing anomalies.<sup>11</sup> Recent work by Jeffrey Fuhrer shows that habits may help explain the “excess smoothness” documented by Campbell and A. S. Deaton in U.S. time series consumption data.<sup>12</sup> David Weil, Jody Overland, and I have recently argued that habits may be able to explain why many East Asian countries had big runups in their saving rates after their periods of rapid growth had commenced.<sup>13</sup> If utility is really subject to important habit formation effects, there is little reason for confidence that the results from the authors' model are even close to optimal.

11. Constantinides (1990); Abel (1990); Campbell and Cochrane (1999).

12. Campbell and Deaton (1989).

13. Carroll, Overland, and Weil (2000).

These points about how the utility function may change over a person's lifetime are very difficult to deal with in the context of the standard model, because there is no compelling way (at least at present) to calibrate exactly how utility functions change over one's lifetime. This is one of the reasons why I admire the authors' enterprise in obtaining and analyzing data from personal finance guides and other popular sources as an alternative benchmark for comparison with the data. Presumably, if utility exhibits strong time-varying elements, personal financial planners will know this, since financial planners get plenty of feedback from retirees about their marginal utility of consumption (although perhaps expressed in more colorful terms than those in the economic literature). The fact that the advice of financial planners matches up at least roughly with the prescriptions of the model does reassure me somewhat on this point.

In conclusion, I think this paper is a major contribution to both the positive and the normative debates about Americans' saving behavior. Most of the difficulties I have outlined are conceptual problems with the entire enterprise of attempting to judge the optimality of household saving behavior. If we were to wait until all these issues had been resolved before attempting an exercise like this, we would be waiting until doomsday. Despite the conceptual difficulties, it is very useful to see what is the best we can do with the tools we have available, and this paper meets that high standard.

**David I. Laibson:** Eric Engen, William Gale, and Cori Uccello have written an outstanding paper that advances the consumption literature. Their analysis demonstrates that U.S. households may be saving adequately for retirement after all. Specifically, the authors show that a well-calibrated model assuming rational behavior and stochastically determined income can explain the observed distribution of retirement wealth.

Although the authors' analysis provides an important new benchmark for the literature on saving, it is not definitive for two critical reasons. First, the simulation results depend on numerous institutional and demographic assumptions that reduce the realism of the simulations. Second, the simulation results depend on time preference assumptions that bias the results in favor of finding saving to be adequate.

To make their model computationally operational, the authors adopt a host of simplifying assumptions. In general, I agree with their choices, but I still wonder how their results would change had they made different and more realistic assumptions. For example, the authors assume that

households have no bequest motive, that income is not variable after retirement, and that health expenditure late in life does not rise. These assumptions bias down the normative saving rate and bias up the reported measures of saving adequacy. On the other hand, the authors also make many assumptions that generate biases in the opposite direction. For example, they do not include insurance markets (annuities, for example) or credit markets, and they set an exogenous retirement date. The authors discuss many of these omissions in a section on sensitivity analysis, and I applaud their thorough efforts to evaluate these potential biases. Indeed, I am convinced that on balance all of these modeling omissions roughly cancel out. But the large number of necessary omissions naturally raises questions about the reliability of the benchmark simulations.

The authors' assumptions about time preferences are more problematic, since these assumptions generate a clear bias in the results. In their benchmark simulation, the authors adopt a normative discount rate of 3 percent, with mortality risk providing an additional discount effect above and beyond this base level. In their appendix B, the authors note that a 3 percent discount rate lies within the range of discount rates estimated in the empirical consumption literature. In theory, one can always find a discount rate that "explains" a given level of wealth accumulation. Using this approach, *any* level of retirement wealth accumulation could be rationalized by picking the "right" discount rate.

The authors need an independent measure of time preference, one that is normatively grounded. Unfortunately, we have no universally accepted mechanism for identifying such a normative measure. However, most evidence suggests that a normative discount rate should be either zero or negative. For example, some studies have asked subjects to pick among hypothetical consumption paths. Using this method, Barsky and others find a median discount rate of zero percent and a mean discount rate of -1 percent.<sup>1</sup> These studies make the budget constraint explicit: higher consumption later in life must come at the expense of lower consumption earlier in life.

To further explore our intuitions about normative discounting, at the Brookings Panel conference I asked each of those in attendance to pick a consumption path for themselves out of a set of seven possible deterministic consumption paths (for the life-cycle period from age thirty-six to age sixty-five). The participants were told:

1. Barsky and others (1997).

—You are deciding at age thirty and face no uncertainty (about health or demographics, for example).

—Consumption represents consumption flows (for example, consumption of housing is calculated on a flow basis).

—The path that you pick will be your actual consumption path (in other words, you will not have access to asset markets to make intertemporal reallocations).

—Your household needs will not change over the life cycle (for example, you will have no children to send to college).

—You are guaranteed to survive until at least age sixty-five.

The choices of the participants generated a range of implied discount rates.<sup>2</sup> Assuming a coefficient of relative risk aversion of 3 (consistent with the calibration in the paper), the distribution of discount rates was as follows: –5 percent (15 percent of responses), –3 percent (26 percent), –1 percent (21 percent), +1 percent (18 percent), +3 percent (15 percent), +5 percent (6 percent), and +7 percent (0 percent). The mean and the median discount rate was –1 percent, which echoes the finding of Barsky and others. By contrast, recall that the authors' preferred calibration features a discount rate of +3 percent.

A preference for a rising consumption profile can also be interpreted as evidence of habit formation. Specifically, the instantaneous utility function may take the form  $u = u(c - X)$ , where  $X$  represents the stock of past consumption. But the specific mechanism that generates a preference for a rising consumption profile does not really matter. The important point is that rising consumption profiles are not implied by the benchmark calibration that the authors adopt. Hence their benchmark normative saving profiles are biased downward, and this biases their measures of saving adequacy upward.

The authors have dramatically advanced the debate on the adequacy of saving by identifying the right methodological framework for normatively evaluating saving outcomes. My concerns boil down to an issue of emphasis, since the authors undertake their analysis for both their benchmark case—a 3 percent discount rate—and the zero percent case that I advocate. The authors have provided the right framework for interpreting

2. The consumption paths were generated by assuming a real rate of return of 1 percent, the rate of return assumed in the original draft of the paper. The Euler equation implies that the discount rate is given by the formula:  $0.01 - (\text{coefficient of relative risk aversion}) \times (\text{growth rate of chosen consumption path})$ .



the implications of the normative discount rate. In this very important sense, the authors have clarified and richly expanded the literature on saving adequacy.

**General discussion:** William Gale responded to a couple of David Laibson's comments. He noted that Laibson's survey produced mean and median time preference rate of  $-1$  percent, based on an assumed real interest rate of 1 percent. However, the slope of the age-consumption profile in the survey depends only on the difference between the interest rate and the time preference rate. Thus, the survey might be better interpreted as showing that the mean and median difference between the interest rate and the time preference rate is 2 percent. If so, then using an interest rate of 3 percent and time preference rates of zero and 3 percent, as in the paper, provides differences between the interest and time preference rates of 3 percent and zero, respectively. That is, the range of assumptions for the interest rate minus the time preference rate used in the paper is perfectly consistent with the findings from Laibson's survey. He also disagreed with Laibson's statement that the simulation model does not generate upward-sloping age-consumption profiles. Even if the model is made nonstochastic, with a 3 percent interest rate, age-consumption profiles are upward sloping with a time preference rate of zero and flat with a time preference rate of 3 percent. The stochastic simulation model generates upward-sloping age-consumption profiles for a wide variety of time preference rates.

William Dickens argued that Laibson's experiment did not address the question of what people's true normative discount rate is. People's answers to questionnaires like Laibson's may show that they derive utility not just from income but also from the beliefs that they hold about the future. People may choose upward-sloping consumption profiles because they want to believe the future will be better than the past, but that does not mean that, looking back over their lives at some future time, they would not conclude that they would have been better off choosing the path that maximized total income. Dickens said he found the arguments for a zero normative discount rate convincing.

William Nordhaus thought it worth stressing that the paper, and the discussion thus far, had focused on the personal rate of time preference, whereas Frank Ramsey in his research into this issue had been interested in the social rate. Nordhaus suggested that the discussion of optimality

from the point of view of policy needed to address the possibility that social and private discount rates might well differ.

Giancarlo Corsetti observed that calibration of a model intended to provide a benchmark for determining the optimality of actual saving is a delicate matter. A benchmark is, by definition, potentially counterfactual. Parameter values need to be picked with some confidence in them being accurate measures of fundamentals. Instead, it appeared the authors had picked some parameter values so as to generate simulation results that match the data in important respects—for example, so as to generate plausible wealth-income ratios. Insofar as the model was calibrated to match variables that could themselves reflect nonoptimal behavior, interpretation of the model as a benchmark is problematic.

Christopher Sims argued that it was important to distinguish between privately optimal behavior, given institutional arrangements and market imperfections, and socially optimal outcomes. The model assumes strongly incomplete asset markets, which may result in outcomes far from socially optimal. He suggested that perhaps the behavior simulated in the paper should not be labeled optimal. From the paper's point of view, one should be quite comfortable passing bag ladies in the street if one knows that they knew all along that they might have a stream of bad luck that put them in this position with a small probability. Individuals would like to insure against this kind of outcome, but cannot do so because of the incompleteness of markets. Policy aimed at achieving socially optimal outcomes should seek to create the appropriate insurance. Thus, when assessing whether savings for retirement are adequate from a social point of view, it is important to ask whether social planning for retirement should recognize the need to compensate for asset market incompleteness.

Like Sims, James Duesenberry argued that the issue is not whether people save adequately from the point of view of individual rationality, but whether saving outcomes are satisfactory from a social perspective. In fact, the motivation behind the social security system was exactly such a social concern with people who might end up with very low consumption possibilities under certain contingencies. The concern was less about whether a low consumption outcome resulted from individual myopia or bad luck than about social problems arising from a given behavior. Austan Goolsbee observed that if one accepts the results of the paper, they seem to imply that individuals in countries that save much more than the United States oversave for retirement.

Dickens suggested improving the simulations in the paper by decomposing changes in income into a moving average and a unit root process, as John Abowd and David Card had done.<sup>1</sup> Such decompositions are notoriously difficult to identify. However, they are consistent with the idea that people dissave when they experience temporary unemployment, whereas they adjust their consumption in response to a permanent loss in income, for example, when they lose a job in the steel industry and take a job at McDonald's. Sims observed that the model assumes that differences in luck in income realizations are the only source of the wide range of asset-income ratios. He thought that it should be possible to get some evidence on whether this is the important source of wealth variation by looking at how strong the correlation is between the asset position in retirement and the history of income shocks in the data. If people are irrational, one should observe a lot of variation in saving among people with similar income histories.

Nordhaus suggested placing the paper in the context of the literature that asks whether the life-cycle model can explain national wealth patterns. This literature has failed to come to a consensus, with Franco Modigliani, among others, claiming that the model gives a good explanation and Lawrence Summers, among others, concluding the opposite. In his own work on this topic, Nordhaus was struck by how easy it was for the life-cycle model to generate widely different wealth-income ratios with relatively small changes in the discount rate and in the elasticity of marginal utility with respect to consumption. He noted that just using the simple consumption Euler equation rule and the parameter values in the paper, the optimal growth rate of consumption with certainty equivalence is  $-1.3$  percent, which may rationalize low saving. Hence he was skeptical whether exercises like the one in the paper could provide meaningful tests of the model's predictive ability. However, Nordhaus did think the paper made an interesting contribution to this debate: when one compares the wealth-income distributions from either of the two surveys reported in the paper with the simulation results, the greatest discrepancies arise for the wealthiest groups. At the very top of the wealth distribution, inheritances, differences in rates of return on investment, and differences in rates of investment are probably the key factors that account for the discrepancy between the data and the simulations. Nordhaus thought that this finding was worth relating to the earlier life-cycle literature.

1. Abowd and Card (1989).

### References

- Aaron, Henry J., and Robert D. Reischauer. 1998. *Countdown to Reform: The Great Social Security Debate*. New York: Century Foundation Press.
- Abel, A. B. 1990. "Asset Formation Under Habit Formation and Catching up with the Joneses." *American Economic Review* 40(2): 38–42.
- Abowd, John M., and David Card. 1987. "Intertemporal Labor Supply and Long-Term Employment Contracts." 77(1): 50–68.
- . 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57(2): 411–45 (March).
- American Academy of Actuaries. 1998. "Financing the Retirement of Future Generations." Public Policy Monograph 1. Washington.
- American Association of Retired Persons. 1999. "Baby Boomers Envision Their Retirement: An AARP Segmentation Analysis." (February.)
- American Savings Education Council. Undated. "The Power to Choose." Pamphlet. Washington.
- Attanasio, Orazio P. 1993. "A Cohort Analysis of Saving Behavior by U.S. Households." NBER Working Paper 4454. Cambridge, Mass.: National Bureau of Economic Research.
- Auerbach, Alan J., and Laurence J. Kotlikoff. 1987. *Dynamic Fiscal Policy*. Cambridge, England: Cambridge University Press.
- Avery, Robert B., Gregory E. Elliehausen, and Glenn B. Canner. 1984a. "Survey of Consumer Finances, 1983." *Federal Reserve Bulletin* 70(9): 679–92 (September).
- . 1984b. "Survey of Consumer Finances, 1983: A Second Report." *Federal Reserve Bulletin* 70(12): 857–68 (December).
- Banks, James, Richard Blundell, and Sarah Tanner. 1998. "Is There a Retirement-Savings Puzzle?" *American Economic Review* 88(4): 769–88.
- Barro, R. J. 1999. "Ramsey Meets Laibson in the Neoclassical Growth Model." *Quarterly Journal of Economics*. Forthcoming.
- Barsky, Robert B., and others. 1997. "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study." *Quarterly Journal of Economics* 112(2): 537–79.
- Becker, Gary S., and Casey B. Mulligan. 1997. "The Endogenous Determination of Time Preference." *Quarterly Journal of Economics* 112(3): 729–58 (August).
- Bernheim, B. Douglas. 1992. "Is the Baby Boom Generation Preparing Adequately for Retirement? Technical Report." Unpublished paper. New York: Merrill Lynch and Company (September).
- . 1994a. "Retirement Saving: Comment." In *Studies in the Economics of Aging*, edited by David Wise. University of Chicago Press.
- . 1994b. "Adequacy of Saving for Retirement and the Role of Economic Literacy." In *Retirement in the 21st Century: Ready or Not*, edited by Dallas L.

- Salisbury and Nora Super Jones. Washington: Employee Benefit Research Institute.
- . 1994c. “The Merrill Lynch Baby Boom Retirement Index.” New York: Merrill Lynch and Company (July 14).
- . 1995. “The Merrill Lynch Baby Boom Retirement Index: Update 95.” Unpublished paper. New York: Merrill Lynch and Company (February).
- . 1997. “The Adequacy of Personal Retirement Saving: Issues and Options.” In *Facing the Age Wave*, edited by David A. Wise. Stanford, Calif.: Hoover Press.
- Bernheim, B. Douglas, and John Karl Scholz. 1993. “Private Saving and Public Policy.” In *Tax Policy and the Economy*, vol. 7, edited by James Poterba. MIT Press.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 1997. “What Accounts for the Variation in Retirement Wealth Among U.S. Households?” NBER Working Paper 6227. Cambridge, Mass.: National Bureau of Economic Research (October).
- Bosworth, Barry, Gary Burtless, and John Sabelhaus. 1991. “The Decline in Saving: Evidence from Household Surveys.” *BPEA*, 1:1991, 183–241.
- Brown, Jeffrey R. 1999. “Are the Elderly Really Over-Annuitized? New Evidence on Life Insurance and Bequests.” NBER Working Paper 7193. Cambridge, Mass.: National Bureau of Economic Research (June).
- Campbell, John Y., and J. H. Cochrane. 1999. “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior.” *Journal of Political Economy* 107(2): 205–51.
- Campbell, John Y., and Angus S. Deaton. 1989. “Why Is Consumption So Smooth?” *Review of Economic Studies* 56(3): 357–74.
- Caplin, A., and J. Leahy. 1999. “The Social Discount Rate.” Unpublished paper. Boston University.
- Carney, Stacie, and William G. Gale. 1998. “Asset Accumulation Among Low-Income Households.” Paper prepared for a Ford Foundation conference on “Benefits and Mechanisms for Spreading Asset Ownership in the United States,” New York, December 10–12.
- Carroll, Christopher D. 1992. “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence.” *BPEA*, 2: 1992, 61–135.
- . 1997a. “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis.” *Quarterly Journal of Economics* 112(1): 1–55 (February).
- . 1997b. “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation).” Working Paper 6298. Cambridge, Mass.: National Bureau of Economic Research.
- . 2000. “Why Do the Rich Save So Much?” In *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, edited by Joel B. Slemrod. Harvard University Press (forthcoming).

- Carroll, Christopher D., Jody R. Overland, and David N. Weil. Forthcoming. "Saving and Growth with Habit Formation." *American Economic Review*.
- Cochrane, John H. 1998. "Where Is the Market Going? Uncertain Facts and Novel Theories" NBER Working Paper 6207. Cambridge, Mass.: National Bureau of Economic Research (February).
- Constantinides, G. M. 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy* 98(3): 519–43.
- Committee for Economic Development. 1995. *Who Will Pay for Your Retirement?* New York: Committee for Economic Development.
- Deaton, Angus. 1992. *Understanding Consumption*. New York: Oxford University Press.
- Diamond, Peter A. 1997. "Macroeconomic Aspects of Social Security Reform." *BPEA 1:1997*, 1–66.
- Diamond, Peter A., and Jonathan Gruber. 1999. "Social Security and Retirement in the United States." In *Social Security and Retirement Around the World*, edited by Jonathan Gruber and David A. Wise. University of Chicago Press.
- Dick, Andrew, Alan M. Garber, and Thomas E. MaCurdy. 1994. "Forecasting Nursing Home Utilization of Elderly Americans." In *Studies in the Economics of Aging*, edited by David A. Wise. University of Chicago Press.
- Dynan, Karen E. 1993. "How Prudent Are Consumers?" *Journal of Political Economy* 101(6): 1104–13.
- Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 1996. "Do the Rich Save More?" Unpublished paper. Dartmouth College.
- Easterlin, Richard A., Christine M. Schaeffer, and Diane J. Macunovich. 1993. "Will the Baby Boomers Be Less Well Off Than Their Parents? Income, Wealth, and Family Circumstances Over the Life Cycle in the United States." *Population and Development Review* 19(3): 497–522.
- Employee Benefit Research Institute. 1994. "Baby-Boomers in Retirement: What Are Their Prospects?" EBRI Issue Brief and Special Report 151 (July).
- Engen, Eric M. 1993a. "Consumption, Uncertain Lifespan, and Euler Equations: Estimation with Panel Data." Unpublished paper. University of California, Los Angeles.
- . 1993b. "Consumption and Saving in a Life-Cycle Model with Stochastic Earnings and Uncertain Lifespan." Unpublished paper. Washington: Federal Reserve Board (October).
- Engen, Eric M., and William G. Gale. 1996. "The Effects of Fundamental Tax Reform on Saving." In *Economic Effects of Fundamental Tax Reform*, edited by Henry J. Aaron and William G. Gale. Brookings.
- . 1997. "Effects of Social Security Reform on Private and National Saving." In *Social Security Reform: Conference Proceedings: Links to Saving, Invest-*

- ment, and Growth*, edited by Steven A. Sass and Robert K. Triest. Conference Series no. 41. Federal Reserve Bank of Boston.
- Engen, Eric M., William G. Gale, and John Karl Scholz. 1994. "Do Saving Incentives Work?" *BPEA 1:1994*, 85–151.
- . 1996. "The Illusory Effects of Saving Incentives on Saving." *Journal of Economic Perspectives* 10(4): 113–38.
- Federal Reserve Board. 1999. "Flow of Funds Accounts of the United States." Washington.
- Feenberg, Daniel, and Jonathan Skinner. 1992. "The Risk and Duration of Catastrophic Health Care." NBER Working Paper 4147. Cambridge, Mass.: National Bureau of Economic Research (August).
- Fisher, Irving. 1930. *The Theory of Interest*. New York: Macmillan.
- Forum for Investor Advice. 1999. "Transitioning to Retirement: The Financial Needs of U.S. Retirees." Available on the World Wide Web at [www.investoradvice.org/Articles/keyfindings1999-05.htm](http://www.investoradvice.org/Articles/keyfindings1999-05.htm) (October 6).
- Friedberg, Leora. 1999. "The Trend Towards Part-Time Work Among Older Americans." Unpublished paper. National Bureau of Economic Research.
- Fuchs, Victor R. 1998a. "Provide, Provide: The Economics of Aging." NBER Working Paper 6642. Cambridge, Mass.: National Bureau of Economic Research (July).
- . 1998b. "Health Care for the Elderly: How Much? Who Will Pay for It?" NBER Working Paper 6755. Cambridge, Mass.: National Bureau of Economic Research (October).
- Gale, William. 1997. "Will the Baby Boom Be Ready for Retirement?" *Brookings Review* 15(3): 4–9.
- Gale, William G., Leslie E. Papke, and Jack L. Vanderhei. 1999. "Understanding the Shift From Defined Benefit to Defined Contribution Plans." Unpublished paper. Brookings (September).
- Gale, William G., and John Sabelhaus. 1999. "Perspectives on the Household Saving Rate." *BPEA 1:1999*, 181–224.
- Gale, William G., and John Karl Scholz. 1994. "Intergenerational Transfers and the Accumulation of Wealth." *Journal of Economic Perspectives* 8(4): 145–60.
- Gately, Dermot. 1980. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables: Comment." *Bell Journal of Economics* 11(1): 373–74 (Spring).
- Gentry, William M., and R. Glenn Hubbard. 1998. "Why Do the Wealthy Save So Much? Saving and Investment Decisions of Entrepreneurs." Working paper. Columbia University (July).
- Gentry, William M., and Joseph Milano. 1998. "Taxes and Investment in Annuities." NBER Working Paper 6525. Cambridge, Mass.: National Bureau of Economic Research (April).

- Gist, John R., Ke Bin Wu, and Charles Ford. 1999. "Do Baby-Boomers Save and, If So, What For?" Paper 9906. Washington: AARP Public Policy Institute.
- Gokhale, Jagadeesh, Laurence J. Kotlikoff, and Mark J. Warshawsky. 1999. "Comparing the Economic and Conventional Approaches to Financial Planning." NBER Working Paper 7321. Cambridge, Mass.: National Bureau of Economic Research (August).
- Grad, Susan. 1990. "Earnings Replacement Rates of New Retired Workers." *Social Security Bulletin* 53(10): 2–19.
- Gustman, Alan L., and Thomas L. Steinmeier. 1998. "Effects of Pensions on Saving: Analysis with Data from the Health and Retirement Survey." NBER Working Paper 6681. Cambridge, Mass.: National Bureau of Economic Research (August).
- . 1999. "Pensions and Retiree Health Benefits in Household Wealth: Changes from 1969 to 1992." NBER Working Paper 7320. Cambridge, Mass.: National Bureau of Economic Research (August).
- Hamermesh, Daniel S. 1984. "Consumption During Retirement: The Missing Link in the Life Cycle." *Review of Economics and Statistics* 66(1): 1–7.
- Hassett, Kevin A., and Gilbert E. Metcalf. 1995. "Energy Tax Credits and Residential Conservation Investment: Evidence from Panel Data." *Journal of Public Economics* 57(2): 201–17.
- Hausman, Jerry A. 1979. "Individual Discount Rates and the Purchase and Utilization of Energy-Using Durables." *Bell Journal of Economics* 10(1): 33–54.
- Hausman, Jerry A., and Lynn Paquette. 1987. "Involuntary Early Retirement and Consumption." In *Work, Health, and Income Among the Elderly*, edited by Gary Burtless. Brookings.
- Hubbard, Glen R., Jonathan Skinner, and Stephen P. Zeldes. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy* 40(June): 59–125.
- . 1995. "Precautionary Saving and Social Insurance." *Journal of Public Economy* 103(2): 360–99.
- Hurd, Michael D. 1987. "Savings of the Elderly and Desired Bequests." *American Economic Review* 77(3): 298–312.
- . 1995. "Mortality Risk and Consumption by Couples." Paper presented for the Conference on the Microeconomics of Saving and Consumption Growth, Institute for Fiscal Studies and Bank of Portugal, Lisbon, Portugal (November).
- Johnson, Richard W., Usha Sambamoorthi, and Stephan Crystal. Forthcoming. "Pension Wealth at Midlife: Comparing Self-Reports with Provider Data." *Review of Income and Wealth*.
- Juster, F. Thomas, and Richard Suzman. 1995. "An Overview of the Health and Retirement Study." *Journal of Human Resources* 30(Supplement): S7–S56.



- Kennickell, Arthur B., Martha Starr-McCluer, and Annika E. Sunden. 1997. "Family Finances in the U.S.: Recent Evidence from the Survey of Consumer Finances." *Federal Reserve Bulletin* 83(1): 1–24 (January).
- Kocherlakota, Narayana R. 1996. "The Equity Premium: It's Still a Puzzle." *Journal of Economic Literature* 34(1): 42–71.
- Koopmans, Tjalling C. 1960. "Stationary Ordinal Utility and Impatience." *Econometrica* 28(2): 287–309.
- Kotlikoff, Laurence J., Avia Spivak, and Lawrence H. Summers. 1982. "The Adequacy of Savings." *American Economic Review* 72(5): 1056–69.
- Laibson, David I., Andrea Reppetto, and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement." *BPEA*, 1:1998, 91–172.
- Lawrance, Emily C. 1991. "Poverty and the Rate of Time Preference: Evidence from Panel Data." *Journal of Political Economy* 99(1): 54–77.
- Lazear, Edward P. 1994. "Some Thoughts on Savings." In *Studies in the Economics of Aging*, edited by David A. Wise. University of Chicago Press.
- Lee, Ronald D., and Jonathan Skinner. 1999. "Will Aging Baby Boomers Bust the Federal Budget?" *Journal of Economic Perspectives* 13(1): 117–40.
- Leung, Sui-Fai. 1994. "Uncertain Lifetime, the Theory of the Consumer, and the Life Cycle Hypothesis." *Econometrica* 62(5): 1233–39.
- Loewenstein, George, and Jon Elster, eds. 1992. *Choice Over Time*. New York: Russell Sage Foundation.
- Loewenstein, George, and Drazen Prelec. 1991. "Negative Time Preference." *American Economic Review* 81(2): 347–52.
- Loewenstein, George F., and Nachum Sicherman. 1991. "Do Workers Prefer Increasing Wage Profiles?" *Journal of Labor Economics* 9(1): 67–84.
- Loewenstein, George, and Richard H. Thaler. 1989. "Intertemporal Choice." *Journal of Economic Perspectives* 3(4): 181–93.
- Lusardi, Annamaria. 1999. "Information, Expectations, and Savings for Retirement." In *Behavioral Dimensions of Retirement Economics*, edited by Henry J. Aaron. Brookings.
- MaCurdy, Thomas E. 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis." *Journal of Econometrics* 18(1): 83–114.
- Manchester, Joyce. 1994. "Baby Boomers in Retirement: An Early Perspective," In *Retirement in the 21st Century: Ready or Not?* edited by Dallas L. Salisbury and Nora Super Jones. Washington: Employee Benefit Research Institute.
- Mariger, Randall P. 1987. "A Life-Cycle Consumption with Liquidity Constraints: Theory and Empirical Results." *Econometrica* 55(3): 533–57.
- McGill, Dan M., and others. 1996. *Fundamentals of Private Pensions*. University of Pennsylvania Press.

- Mitchell, Olivia, James Moore, and John Phillips. 1998. "Explaining Retirement Saving Shortfalls." Philadelphia: Pension Research Council.
- Mitchell, Olivia S., and Sylvester J. Schieber, eds. 1998. *Living with Defined Contribution Pensions: Remaking Responsibility for Retirement*. Philadelphia: University of Pennsylvania Press.
- Moore, James F., and Olivia S. Mitchell. 1997. "Projected Retirement Wealth and Savings Adequacy in the Health and Retirement Survey." NBER Working Paper 6240. Cambridge, Mass.: National Bureau of Economic Research (October).
- Olson, Mancur, and Martin J. Bailey. 1981. "Positive Time Preference." *Journal of Political Economy* 89(1): 1–25.
- Palmer, Bruce. 1994. "Retirement Income Ratios: An Update." *Benefits Quarterly* 10(2): 59–75.
- Pelphs, Edmund S. 1961. "The Golden Rule of Accumulation: A Fable for Growth Man." *American Economic Review* 51: 638–43.
- Poterba, James M. 1997. "The History of Annuities in the United States." NBER Working Paper 6001. Cambridge, Mass.: National Bureau of Economic Research (April).
- Poterba, James M., and Andrew A. Samwick. 1995. "Stock Ownership Patterns, Stock Market Fluctuations, and Consumption." *BPEA* 2:1995, 295–357.
- Quinn, Joseph F. 1999. "Retirement Patterns and Bridge Jobs in the 1990's." EBRI Issue Brief 206. Washington: Employee Benefit Research Institute (February).
- Ramsey, Frank. 1928. "A Mathematical Theory of Saving." *Economic Journal* 38(152): 543–59.
- Robb, A. L., and J. B. Burbidge. 1989. "Consumption, Income, and Retirement." *Canadian Journal of Economics* 22(3): 522–42.
- Sabelhaus, John. 1997. "Public Policy and Saving in the United States and Canada." *Canadian Journal of Economics* 30(2): 253–75.
- Samwick, Andrew A. 1995. "The Limited Offset Between Pension Wealth and Other Private Wealth: Implications of Buffer-Stock Saving." Unpublished paper. Dartmouth College.
- . 1998. "Discount Rate Heterogeneity and Social Security Reform." *Journal of Development Economics* 57(1): 117–46.
- Sheiner, Louise, and David N. Weil. 1992. "The Housing Wealth of the Aged." NBER Working Paper 4115. Cambridge, Mass.: National Bureau of Economic Research (July).
- Skinner, Jonathan. 1988. "Risky Income, Life Cycle Consumption, and Precautionary Savings." *Journal of Monetary Economics* 22(2): 237–55.
- . 1994. "Some Thoughts on Savings: Comment." In *Studies in the Economics of Aging*, edited by David A. Wise. University of Chicago Press.
- Tacchino, Kenn, and Cynthia Saltzman. 1999. "Do Accumulation Models Overstate What's Needed to Retire?" *Journal of Financial Planning* 12(2): 62–73.

- Tyson, Eric. 1997. *Personal Finance for Dummies*. Foster City, Calif.: IDG Books Worldwide.
- U.S. Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds. 1995. *1995 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds*. Washington: Government Printing Office.
- . 1999. *1999 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Disability Insurance Trust Funds*. Washington: Government Printing Office.
- U.S. Congressional Budget Office. 1993. "Baby Boomers in Retirement: An Early Perspective." Washington (September).
- U.S. Department of Labor. 1997. "The Top 10 Ways to Beat the Clock and Prepare for Retirement." Consumer pamphlet. Washington.
- . 1998. "Final Report on the National Summit on Retirement Savings. 1998 National Summit on Retirement Savings, June 4–5." Washington.
- U.S. Social Security Administration. 1997. "Social Security Area Population Projections." Actuarial Study 112. Baltimore: Social Security Administration, Office of the Chief Actuary.
- . 1998. *Income of the Population 55 and Older, 1996*. Washington: Social Security Administration.
- Venti, Steven F., and David A. Wise. 1990. "But They Don't Want to Reduce Housing Equity." In *Issues in the Economics of Aging*, edited by David A. Wise. University of Chicago Press.
- . 1992. "Government Policy and Personal Retirement Saving." In *Tax Policy and the Economy*, vol. 6, edited by James Poterba. MIT Press.
- Warshawsky, Mark J., and John Ameriks. 1998. "What Does Financial Planning Software Say About Americans' Preparedness for Retirement?" Unpublished paper. TIAA-CREF Institute, New York (July).
- Wiatrowski, William J. 1991. "New Survey Data on Pension Benefits." *Monthly Labor Review* 114(8): 8–22.
- Workplace Pulse Surveys. 1997. "Savers Interest Piqued by Tax-Advantaged IRAs, Retirement Savings Slowed in 1997, Though Many Workers Expect to Live Better in Retirement Years." Press release. Washington: Workplace Pulse Surveys and Colonial Life and Accident Insurance Company (December 9).
- Zeldes, Stephen P. 1989a. "Optimal Consumption with Stochastic Income: Derivations from Certainty Equivalence." *Quarterly Journal of Economics* 104(2): 275–98.
- . 1989b. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97(2): 305–46.

