

The importance of precautionary motives in explaining individual and aggregate saving*

R. Glenn Hubbard[†]

Columbia University, New York, NY 10027

National Bureau of Economic Research

Jonathan Skinner

University of Virginia, Charlottesville, VA 22901

National Bureau of Economic Research

and

Stephen P. Zeldes

The Wharton School, University of Pennsylvania, Philadelphia, PA 19104

National Bureau of Economic Research

Abstract

This paper examines predictions of a life-cycle simulation model — in which individuals face uncertainty regarding their length of life, earnings, and out-of-pocket medical expenditures, and imperfect insurance and lending markets — for

*We are very grateful to Eric Engen, Robert Reider, and Charlie Ye for assistance with computer programming, and to Keith Dallara, Leslie Jeng, David Kerko, Vasan Kesavan, Richard McGrath, and Dan Zabinski for research assistance. We received helpful comments from Miles Kimball, Allan Meltzer, Masao Ogaki, Edgar Olsen, and attendees at the Carnegie-Rochester Conference on Public Policy, April 23–24, 1993, and the NBER Summer Institute Working Group on Consumption. This research was supported by the Harry and Lynde Bradley Foundation and the National Science Foundation. In addition, Hubbard acknowledges financial support from the Faculty Research Fund of Columbia Business School, Skinner from the National Institute on Aging, and Zeldes from the Alfred P. Sloan Foundation. Major computational work was conducted using the Cornell National Supercomputer Facility, and was funded by the Cornell Theory Center and the National Science Foundation.

[†]*Correspondence to:* R. Glenn Hubbard, Columbia University, Graduate School of Business, Uris Hall 609, New York, NY 10027.

individual and aggregate wealth accumulation. Relative to life-cycle or buffer-stock alternatives, our augmented life-cycle model better matches a variety of features of U.S. data, including: (1) aggregate wealth, (2) cross-sectional differences in wealth-age and consumption-age profiles by education group, and (3) short-run time-series comovements of consumption and income.

I. Introduction

Since the pioneering work of Modigliani and Brumberg (1954) and Ando and Modigliani (1963), the life-cycle model has been the leading theory of aggregate wealth accumulation. The theory has induced volumes of research, and has guided the structure of many models used in economic policy.¹ Yet whether the life-cycle model can actually explain observed patterns in either individual or aggregate wealth accumulation is controversial (see, *e.g.*, Deaton, 1992). Empirical evidence has suggested that standard certainty (or certainty-equivalence) life-cycle models fail along at least two dimensions. First, such models tend to underpredict aggregate wealth in the United States (Kotlikoff and Summers, 1981; White, 1978; Darby, 1979). The underprediction of aggregate wealth accumulation by the life-cycle model could mean that bequest motives play a strong role in saving behavior, but it could also mean that other motives not captured by the model are important as well.

Second, the life-cycle model overpredicts individual wealth for a significant fraction of households. In the life-cycle model, nearly everyone should save for retirement, given the decline in income after retirement for nearly every individual. Yet median financial assets of families nearing retirement are typically only a small fraction of current income (see, *e.g.*, Venti and Wise, 1987; Bernheim and Scholz, 1993; Hubbard, Skinner, and Zeldes, 1993).

A parallel literature has sought to examine how idiosyncratic risk affects consumption and saving, in the absence of frictionless markets for insurance and lending. While some researchers have focused on earnings uncertainty (see *e.g.*, Skinner, 1988; and Zeldes, 1989b), others have instead analyzed uncertainty about length of life (see, *e.g.*, Hubbard and Judd, 1987) or health expenses (see, *e.g.*, Kotlikoff, 1988a).²

This paper combines these two strains of the literature by considering the theoretical and empirical implications of a life-cycle model of consumption, saving, and wealth accumulation subject to the three sources of uninsured idiosyncratic risk that we think are most important: uncertainty about earnings, medical expenses, and lifespan. That is, we examine how the *precautionary* motive for saving affects both the theoretical predictions and the

¹See Auerbach and Kotlikoff (1987) for the preeminent model of this type.

²There is also a growing literature on the effects of idiosyncratic risk on equilibrium asset-pricing; see, *e.g.*, Heaton and Lucas (1992) and Aiyagari and Gertler (1991).

policy prescriptions of the life-cycle model.

We use dynamic programming techniques to solve the model numerically. Previous numerical models have included only one or two sources of uncertainty using stylized models. We create a more realistic model in which families live for many periods, working for part of their lives and retiring later in life. Families face uncertainty about earnings, lifespan, and medical expenses. Thus we incorporate both traditional life-cycle aspects and uncertainty into our model. We also include a realistic modeling of means-tested public-welfare programs such as AFDC and Medicaid.³ To parameterize the uncertainty actually facing households, we estimate the stochastic processes using a variety of cross-section and panel data sets on households. The result is a model with inputs that are significantly more realistic than existing work: more realistic than conventional life-cycle models because we explicitly incorporate uncertainty, and more realistic than existing models of precautionary saving because we include accurate life-cycle earnings profiles and parameterizations of the three major forms of uncertainty facing households.

Three questions naturally arise. First, do the theoretical implications of our model differ from those of earlier models and are the effects, *implied by the model*, of idiosyncratic risk on consumption and capital accumulation economically large and unambiguously of one direction or another? Second, if so, are the implications of the model (in particular the ones that differ from previous work) more consistent with the real world? Third, how do the policy implications of our model differ from those of certainty models?

We find that precautionary saving is large in these models; the precautionary motive plays an important role in determining aggregate saving. We also show that along a number of dimensions, our model better replicates empirical regularities in consumption, saving, and aggregate wealth data than existing models. We focus on the model's ability to explain three sets of empirical facts: corresponding to (1) aggregate wealth and aggregate saving rate, (2) cross-sectional differences in wealth-age and consumption-age profiles by education group, and (3) short-run time-series properties of consumption, income, and wealth. We find that the aggregate wealth-income ratios implied by this model better approximate the empirical wealth-income ratios than do those generated by alternative models. However, our data do not replicate the sharp differences in average saving rates between high-school dropouts and college-educated families, largely because of the higher relative uncertainty faced by low-income families in our model. In addition,

³This paper is one of two papers that we have written to date using this model. In Hubbard, Skinner, and Zeldes (1993), we focus on the implicit taxation of saving due to the asset-based means testing of many social insurance programs (such as AFDC, food stamps, and Medicaid). We show in that paper that the combination of such programs and uncertainty can explain much of the heterogeneity in accumulated assets and the large fraction of the population with very little wealth.

we show that the model's prediction about the portion of households approximately consuming their current income match the data more closely than alternative models. Finally, the policy implications of our model differ from conventional perfect certainty models. First, the simulation results emphasize the importance of the *variance* in earnings and health expenses in determining aggregate saving and wealth. These results suggest, for example, that tax policy may affect saving as much by shifting the variance of income as by changing after-tax rate of return. Second, the model suggests that the design of government expenditure programs can have an important impact on aggregate saving behavior.

The paper is organized as follows. The second section describes the multiperiod dynamic programming model and our numerical solution technique. The third section describes our parameterization of the model, much of which is based on original empirical work with panel datasets. In the fourth section, we examine how well our model explains aggregate wealth in the United States, and estimate the importance of uncertainty for total wealth accumulation. The fifth section focuses on the question of whether our model can explain wealth-age and consumption-age profiles. In the sixth section, we examine the implications of our model for the short-run time-series properties of consumption, income, and wealth. The seventh section concludes.

II. A multiperiod model of consumption under uncertainty

Virtually all of the research on effects of uncertainty on consumption that has solved *analytically* for the optimal level of consumption and saving has done so by assuming either no uncertainty or a quadratic utility function in consumption. Consider a model with uncertain earnings. With a quadratic utility function (ignoring nonnegativity constraints on consumption), consumption under earnings uncertainty is identical to what it would be if earnings were set equal to its expected value. This basic certainty-equivalence result is that optimal consumption is based on expected lifetime resources, calculated by adding financial wealth to the present value of expected future income. Consumption each period equals then the annuity value (possibly allowing for growth) of this summary statistic (see Zeldes, 1989b, for a further discussion of this point). In such models, the receipt of a dollar today has the identical effect on consumption as a dollar (in present value terms) awarded twenty years from now. The certainty or certainty-equivalence model is standard in almost all of the literature on the life-cycle hypothesis and permanent income hypothesis that examines the level of consumption, and is the basis of much of our intuition about determinants of consumption.

A number of authors have examined the effects of uncertainty with preferences other than the quadratic form. If preferences are such that the third

derivative of the utility function is positive (exhibiting what Kimball, 1990, has termed “prudence”), then, at any level of wealth, saving under uncertainty is greater than saving under certainty (with income set equal to its expected value). Leland (1968) dubbed this difference “precautionary saving.” Leland (1968); Sandmo (1970); Drèze and Modigliani (1972); Sibley (1975); Miller (1976); Bewley (1977); Schechtman and Escudero (1977); Levhari, Mirman, and Zilcha (1980); Cantor (1985); Clarida (1987); Caballero (1990); Kimball (1990, 1991), and others have derived analytically some properties of the solution under uncertainty.

When certainty equivalence does not hold, analytical solutions to the problem do not in general exist.⁴ Therefore, it is necessary to use numerical techniques to solve for optimal consumption. Numerical studies of precautionary saving include those by Skinner (1988), Zeldes (1989b), Hubbard and Judd (1987), Deaton (1991), Carroll (1992), Kotlikoff (1988a), Aiyagari (1992), and others. Recently, Deaton (1991) and Carroll (1992) have argued that existing precautionary saving models failed to predict the large fraction of the population with extremely low wealth accumulation. In their models of precautionary saving, therefore, they assume that the rate of time preference is high relative to the real interest rate.⁵ In response to earnings uncertainty (possibly augmented by borrowing constraints), individuals maintain a “buffer stock” or contingency fund against income downturns. Impatience keeps these buffer stocks small, thereby providing a possible explanation of why many households save little throughout their life. We return to this model later in the paper.

Our model builds on the set of earlier work on precautionary saving. Previous numerical work included only one or two sources of uncertainty in the context of highly stylized models. We create a richer model: families live through many periods; they work for part of their lives, and retire later in life. Families face uncertainty about earnings, lifespan, and health expenses. Thus, we incorporate both life-cycle aspects and multiple sources of uncertainty into our model. To parameterize the uncertainty actually facing households, we estimate the stochastic processes using a variety of cross-section and panel data sets on households. An important element of this model is the realistic modeling of means-tested public-welfare programs, which we describe later in the paper and in Appendix A.

⁴An exception is when the utility function is exponential. See, *e.g.*, Cantor (1985), Kimball and Mankiw (1989), and Caballero (1991).

⁵What is required are assumptions such that in a certainty model, households would prefer to borrow against future income.

The consumer's problem

The consumer acts to maximize expected lifetime utility, given all of the relevant constraints. At each age t , a level of consumption is chosen which maximizes

$$E_t \sum_{s=t}^T D_s \cdot U(C_s) / (1 + \delta)^{s-t}, \quad (1)$$

subject to the transition equation

$$A_s = A_{s-1}(1 + r) + W_s + TR_s - C_s - M_s, \quad (2)$$

where

$$TR_s = TR(W_s, M_s, A_{s-1}(1 + r)). \quad (3)$$

with the borrowing constraints and terminal condition:

$$A_s \geq 0, \forall s. \quad (4)$$

Equation (1) indicates that consumption C_s is chosen to maximize expected lifetime utility (where E_t is the expectations operator conditional on information at time t .) The period utility is discounted at the rate of time preference δ . To account for a random date of death, D_s is a state variable that is equal to unity if the individual is alive and zero otherwise; death ($D_s = 0$) is, of course, an absorbing state. The limit T is the maximum potential lifespan, which in the numerical work will correspond to age 100.

The transition equation (2) describes asset accumulation. The household begins period s with financial assets from the previous period plus accumulated interest, $A_{s-1}(1 + r)$, where r is the exogenous, nonstochastic real after-tax rate of interest. It then receives exogenous earnings, W_s , pays out necessary medical expenses, M_s , and receives government transfers, TR_s . As a result, the consumer is left each period with

$$X_s = A_{s-1}(1 + r) + W_s - M_s + TR_s, \quad (5)$$

which, following Deaton (1991, 1992), we refer to as “cash on hand.” Consumption is chosen given X_s , and what is remaining equals end-of-period assets in period s , A_s . We assume that no utility is derived from medical expenditures per se; the costs are required to offset the consequences of ill health.

Transfers received depend on earnings, medical expenses, and financial assets. This formulation permits us to incorporate transfer programs with income-based means testing, payments tied to medical expenses, and asset-based means testing (including such payments as Medicaid benefits, which depend upon medical expenses and assets, and food stamps and other cash assistance, which depend on earnings and assets).

For this paper, the transfers function in (3) is set equal to:

$$TR_s = \max\{[(\bar{C} + M_s) - (A_{s-1}(1+r) + W_s)], 0\}. \quad (6)$$

That is, the consumer receives transfers equal to a consumption “floor” \bar{C} — a minimum level of consumption to be guaranteed by transfers — plus medical expenses minus all available resources, if that combined amount is positive, and zero otherwise. We assume that transfers are reduced dollar for dollar of additional assets or current earnings received.⁶ This simplified transfer function captures the high implicit marginal tax rates on saving in means-tested programs such as Medicaid, AFDC, and food stamps.⁷

Borrowing constraints prevent negative assets. We introduce borrowing constraints because of the existence of the guaranteed consumption floor \bar{C} . If borrowing were allowed, an individual could borrow in one period, then default and consume \bar{C} in the following period. Given that individuals have the option of declaring bankruptcy and consuming \bar{C} , we assume that there is no market for unsecured loans.⁸

As outlined, the individual’s problem cannot be solved analytically except in special cases, so we will use numerical stochastic dynamic programming techniques to approximate closely the solution. Using this method, we will calculate explicit decision rules for optimal consumption as well as the consumer’s value function. The solution technique is described in detail below and in Appendix B.

III. Model parameterization and sources of uncertainty

In this section, we describe the parameterization of the model just outlined, highlighting our definition of the utility function; the estimation of uncertainty over lifespan, earnings, and out-of-pocket medical expenses; the estimation of the average age profiles of earnings and out-of-pocket medical expenses; and the specification of the consumption floor. Readers more interested in applications of the model can continue to the fourth section.

⁶This is an approximation to the type of means testing in U.S. transfer programs. Appendix A contains a detailed description of actual means testing.

⁷Hubbard, Skinner, and Zeldes (1993) focus on the effects of this implicit taxation, inherent in programs that use asset-based means testing, on optimal saving under uncertainty.

⁸In the parameterizations of our model under uncertainty, the maximum realization of medical expenses is always greater than the minimum possible earnings realization, i.e., the minimum net earnings draw in any period is negative. In the case when \bar{C} is set to zero, and the utility function is such that $U'(0) = \infty$, individuals choose never to borrow and the liquidity constraint is never binding (see the related discussion in Zeldes, 1989b). Therefore, in the uncertainty model, we are, in effect, preventing borrowing against the future guaranteed consumption floor.

Utility function

For simplicity and to preserve comparison with other studies, we assume that the period utility function in (1) is isoelastic. That is,

$$U(C_s) = (C_s^{1-\gamma} - 1)/(1 - \gamma) \quad (7)$$

The coefficient γ serves multiple roles in this utility function: γ is the coefficient of relative risk aversion, $(1/\gamma)$ is the intertemporal elasticity of substitution in consumption, and $(\gamma + 1)$ is the coefficient of relative prudence (Kimball, 1990). The third derivative of this utility function is positive, which will generate precautionary saving in response to uncertainty regarding earnings and out-of-pocket medical expenses. Our preferred value for γ is 3, which is consistent with many empirical studies.⁹ Our benchmark case assumes that the rate of time preference, δ , is 3 percent per annum. Finally, the real after-tax rate of interest r is assumed to be 3 percent per annum. In the simulation exercises reported in the third through sixth sections, we conduct a sensitivity analysis of the effects of alternative values for γ and δ on our results.

Sources of uncertainty

Uncertain lifetimes. Several researchers have documented the importance of uncertain lifespan on life-cycle consumption, either in theoretical models (Yaari, 1965; Davies, 1981; Abel, 1985; and Hubbard, 1987) or empirical studies (Hubbard and Judd, 1987; Hurd, 1989; and Engen, 1992b). Access to a fair annuity market could remove the influence of lifetime uncertainty on consumption. Individuals could exchange a portion of their wealth when young to smooth consumption in old age. The existence of a competitive equilibrium may, however, be precluded by asymmetries of information between individuals and insurers, since individuals have better information concerning their life expectancy (see *e.g.*, Rothschild and Stiglitz, 1976; and Wilson, 1977). Annuity markets in the United States are indeed small, and “load factors” (*i.e.*, premia above actuarially fair prices) are high (see Friedman and Warshawsky, 1988). In the absence of annuities, substantial precautionary saving is likely to accompany uncertainty over longevity (Hubbard, 1987).

We use mortality data from the National Center for Health Statistics and the Social Security Administration (Faber, 1982) to estimate the conditional probability at each age of surviving one more year. We consider that the family “dies” when the last member, assumed to be the wife, dies; we, therefore, use mortality data on women. The model initiates household decisions

⁹See *e.g.*, Ghez and Becker (1975), Skinner (1985), Engen (1992b), and the more detailed discussion in Auerbach and Kotlikoff (1987).

at age 21, and the maximum length of life is set equal to 100. The conditional mortality probabilities are presented in Appendix A.

Earnings uncertainty. Available evidence suggests that individuals' earnings uncertainty is substantial (see, for example, Hall and Mishkin, 1982; MaCurdy, 1982; and Abowd and Card, 1989). It is clear that there is an absence of markets through which human capital returns can be explicitly traded and of securities with which individual-specific income risks can be perfectly hedged. It is likely that insurance against individual-specific earnings fluctuations is severely limited for reasons of adverse selection and moral hazard. In addition, the cost of credit to individuals with uncertain earnings will likely be high (all else equal) because of asymmetries of information between borrowers and lenders; that is, individuals may have better knowledge about the source, persistence, and realizations of shocks to earnings than do lenders. Hence, in the presence of uncertain earnings, there are likely to be borrowing constraints as well.

As noted by Skinner (1988), the impact of earnings uncertainty on precautionary saving depends in part on the time-series properties of shocks to individual earnings. Time-series patterns of earnings and wages have been the subject of many studies. Previous research on earnings uncertainty (notably Lillard and Willis, 1978; Hall and Mishkin, 1982; MaCurdy, 1982; and Abowd and Card, 1989) has focused purely on labor earnings. The results of this research have typically indicated substantial variability in earnings, and, in some cases, evidence of near random-walk behavior in the log of earnings.¹⁰ Our empirical analysis differs in two respects from these previous studies. First, we include unemployment insurance and subtract taxes in our measure of "earnings"; these adjustments should reduce earnings variability. Second, we partition our sample into the three educational categories – no high-school degree, high-school degree, and college degree. For each education group, we consider a (log) earnings equation:

$$y_{it} = Z_{it}\beta + u_{it} + v_{it}, \quad (8)$$

where y_{it} is the log of earnings, Z_{it} is a vector of explanatory variables (see Appendix A), v_{it} is transitory earnings (and measurement error, discussed below), and $u_{it} = \rho u_{it-1} + e_{it}$. To estimate the pattern of earnings uncertainty,

¹⁰A study of earnings and wage structure by MaCurdy (1982), using ten years of data from the PSID, found a high degree of serial correlation in the error terms of earnings, suggesting that shocks to earnings are persistent. MaCurdy's estimates suggest, roughly, that each new shift in earnings or wages is roughly half a transitory shock (or measurement error), while the remaining half is expected to persist in the future, decaying slowly over time. Thus, according to MaCurdy and Hall and Mishkin (1982), the permanent component is highly serially correlated and not far from a random walk. Skinner (1988) and Zeldes (1989b) have shown that precautionary saving is substantial given these risk patterns.

we use panel data from the Panel Study of Income Dynamics (PSID) on families during the years 1983–1987 (reporting on earnings in years 1982–1986) who were neither retired nor reporting less than \$3000 in the sum of earnings and transfer income.¹¹ We define “labor income” as earnings of the head and spouse (if present) plus unemployment insurance. We calculate an average tax rate for each family as the ratio of total tax payments to total family income. We then calculate after-tax labor income assuming that the overall average rate held for labor income.¹² We estimate earnings regressions (in levels) separately for each education group, with estimated coefficients presented in Appendix A.

Residuals from log earnings regressions are used to estimate the education-specific $AR(1)$ process for the residuals. We assume that the transitory component of earnings ν_{it} is entirely measurement error, thereby effectively removing much of the true transitory randomness in earnings.

The estimated parameters for each educational class are presented in Appendix A. There are differences across educational class; families headed by high-school dropouts are subject to considerably more risk of earnings fluctuations over time than those with college educations. For all three education groups, roughly half of the innovation in log earnings is transitory or measurement error, while the other represents a highly persistent shock to earnings. We also estimate receipts from retirement annuities by education group (see Appendix A).

Uncertainty over expenses for medical care. Out-of-pocket medical-care expenditures are a significant fraction of household budgets, especially for elderly families. A recent report from the U.S. House Committee on Aging (1990) found that out-of-pocket medical-care expenditures in elderly households were 18.2 percent of mean income, a substantial rise from only 12.3 percent in 1977. Furthermore, Feenberg and Skinner (1992) present evidence suggesting that these out-of-pocket expenses are very persistent over time.

Our empirical model of (log) medical expenses is represented by:

$$m_{it} = G_{it}\Gamma + \mu_{it} + \omega_{it}, \quad (9)$$

where m is the log of medical expenses, G is a vector of explanatory variables, ω_{it} is measurement error plus transitory medical expenses (attributed

¹¹We also excluded families for which education levels were not apparent. We did not sample on the basis of employment; workers who were unemployed or not in the labor force were included in the sample as long as total family income exceeded \$3000.

¹²The marginal tax rate is relevant for calculating optimal labor given the net-of-tax wage rate. However, our model assumes that labor supply is exogenous. We use the average tax rate instead because it is the relevant tax measure in the budget constraint. For example, saving is equal to net-of-tax income (*i.e.*, income reduced by the *average* tax rate) less consumption.

entirely to measurement error), and $\mu_{it} = \rho_m \mu_{it-1} + \epsilon_{it}$, where ϵ is the white-noise innovation in health expenses.¹³ We estimate the time-series properties of medical costs using a panel of tax returns. During the period 1968–73, households could deduct medical expenses in excess of 3 percent of adjusted gross income. Based on declared medical expenses by households over age 54 during this period, we estimated the dynamic properties of medical costs using the smoothed simulated maximum likelihood tobit methodology described in Feenberg and Skinner (1992). We then merge these estimates with cross-sectional data from the 1977 National Health Care Expenditure Survey and the 1977 National Nursing Home Survey to characterize the risk properties of out-of-pocket medical expenses. (Our estimation results are presented in Appendix A.) These estimates suggest high persistence in medical-care expenses, with $\rho_m = 0.901$ and a sizable standard deviation of the white-noise error term ϵ_{it} . These medical-care uncertainty parameters are specific to data on the elderly, so there may be a bias in attributing these risk parameters to the nonelderly. However, the mean square error in the cross-sectional log medical-care regressions are nearly identical for the nonelderly (1.27) and elderly (1.31).

The consumption floor

The consumption floor, the minimum government-guaranteed level of consumption for a family, depends on many factors, such as age, number of children, and state of residence. We attempt to calculate the value of the floor for a representative eligible family, and assume that recipients treat the transfers as if they were cash.

It is important to distinguish between entitlement and nonentitlement programs in calculating the floor. Under entitlement programs, anyone who is eligible may sign up; accordingly, we assume that eligible individuals participate. Hence, entitlement benefits are included in the floor with no adjustment. A nonentitlement program, such as housing subsidies, is limited by the annual program budget, so that there can be queues for assistance. In this case, we adjust benefits by the percentage of the population actually receiving benefits. We also adjust state Supplemental Security Income benefits because some states do not provide state supplements. We describe the exact construction of consumption floors for the nonelderly and elderly populations in Appendix A. In our model, we assume a consumption floor of \$7000.

¹³We assume that the outcomes of earnings and medical expenses are uncorrelated. In practice, it is likely that a health shock will be associated with a change in earnings. Hence assuming independence between the two outcomes will tend to understate the true financial risk of poor health outcomes.

Numerical solution of the dynamic model

The household's problem. We believe that the household's problem as stated cannot be solved analytically, so we use numerical stochastic dynamic programming techniques to approximate closely the solution. Using these methods, we calculate explicit decision rules for optimal consumption as well as the value function. We summarize the technique below. It is described more fully in Appendix B.

As noted above, earnings and medical expenditures are assumed to follow first-order autoregressive processes around a deterministic trend. The deviation from the trend is discretized into 9 nodes. Hence earnings and health deviations from trend are first-order Markov processes, with the probability of realizing a given discrete outcome in period $t + 1$ a function of the outcome in period t .

We divide the maximum feasible range for cash on hand, X , in each period into 61 nodes. The nodes are evenly spaced on the basis of the log of cash on hand, in order to get finer intervals at lower absolute levels of cash on hand, where nonlinearities in the consumption function are most likely.

The dynamic program therefore has three state variables in addition to age: cash on hand, earnings, and medical expenses.¹⁴ The problem is solved by starting in the last possible period of life (T) and solving backwards. At period T , C_T equals X_T . In periods prior to T , we calculate optimal consumption for each possible combination of nodes, using stored information about the subsequent period's optimal consumption and value function. We do not discretize consumption, but allow it to be a continuous variable. Because of possible multiple local maxima, we use information about both the value function and expected marginal utility in our search for optimal consumption. Optimal consumption is calculated by searching for levels of consumption that maximize the value function and that (with the exception of corner solutions) equate the marginal utility of consumption at t to the (appropriately discounted) expected marginal utility of consumption in period $t + 1$. Solving the household's problem numerically involves extensive computation.¹⁵

Once we determine the optimal consumption function for all possible nodes, it is straightforward to simulate a history for each of a large number (4000 in most cases) of families. For each family, we use the following procedure. In any period, we begin with the level of assets from the previous period and multiply by $(1 + r)$. We draw random realizations for earnings and

¹⁴In years after retirement, the earnings state variable is a trivial one, leaving us with two state variables.

¹⁵All computer work was performed using the vectorizing capabilities of the Cornell National Supercomputer Facility, a resource of the Cornell Theory Center, funded by the National Science Foundation, the IBM Corporation, the state of New York, and members of the Corporate Research Institute.

medical expenses from the appropriate distributions. We then add the realized earnings and subtract the realized medical expenses, resulting in a value for cash on hand. Since realized cash on hand will not generally be equal to one of the nodes for cash on hand, we interpolate the optimal consumption function, using the two nearest nodes for cash on hand, for the given levels of earnings and medical expenses. This gives us the realized value for consumption. Subtracting this consumption from cash on hand gives us end-of-period assets. We then follow each family over time, recording the realized levels of earnings, consumption, and assets for each period.

Aggregation of individual consumption, earnings, and assets. Having solved for individual consumption functions and simulated time paths of consumption, earnings, and assets, we next need to calculate aggregate consumption, earnings, and assets. We use educational attainment as a proxy for lifetime income, and focus on three education groups composed of individuals with less than 12 years of education and thus without a high-school degree (the “no-high-school” group), high-school graduates (the “high-school” group), and college graduates (the “college” group). Recent papers by Bernheim and Scholz (1993) and Attanasio (1993) also examine differences in saving behavior of different education groups. We calculate aggregate values for each education group. The annual growth rate of the population is assumed to be one percent, and the growth rate of productivity is assumed to be zero. The aggregation procedure is outlined in Appendix B.

The next step is to examine the implications of the model and assess whether those implications are consistent with patterns in individual and aggregate data.

IV. Can the augmented life-cycle model explain aggregate saving?

The ability of the standard perfect-certainty version of the life-cycle model to explain aggregate saving in the United States has been debated for some time. Tobin (1967), Boskin (1978), and Modigliani (1988a,b) have argued that “life-cycle wealth” constitutes the majority of observed household wealth, a view challenged by White (1978), Darby (1979), and Kotlikoff and Summers (1981), and Kotlikoff (1988b); more intermediate views are presented in Tobin and Dolde (1971) and Hubbard and Judd (1987).

Challenges to the importance of life-cycle saving have taken two forms: (1) comparisons of calculated life-cycle assets (based on observed flows of earnings, transfers, and consumption) with total wealth,¹⁶ and (2) compar-

¹⁶One research program has focused on separating actual private net worth into life-cycle wealth and transfer wealth, where the former equals the accumulated value of earnings over consumption and the latter equals accumulated net transfers. Using survey questionnaire data, Modigliani (1988b) and Hurd and Mundaca (1989) estimate that transfers account

isons of total household saving with aggregate saving generated by a life-cycle simulation model. We focus our attention on the latter approach. To approximate life-cycle saving in the second approach, a number of authors have constructed and simulated “certainty” versions of a life-cycle model similar to the one described in the second section, using the isoelastic utility function in (7) (see, *e.g.*, Summers, 1981; Auerbach, Kotlikoff, and Skinner, 1983; Hubbard and Judd, 1986; and Auerbach and Kotlikoff, 1987). White (1978) found that a perfect-certainty version of the life-cycle model with parameters based on available empirical evidence cannot explain a significant fraction of aggregate saving. Like White (1978), Hubbard and Judd (1987) found that with plausible assumptions, the standard certainty life-cycle model generates levels of saving that are too small relative to aggregate saving.

Hubbard and Judd (1987), Skinner (1988), Caballero (1991), and Carroll and Samwick (1992b) have used stylized models of precautionary saving and shown that uncertainty can raise aggregate saving and wealth substantially.¹⁷ Our goal in this section is to examine whether our augmented life-cycle model with the combination of social insurance and precautionary saving against earnings, medical expense, and lifespan uncertainty can explain asset accumulation, aggregated by education group and for the entire economy, more satisfactorily than the standard life-cycle model.

Our first set of experiments examines differences in predicted wealth accumulation of different lifetime income groups (using educational attainment

for less than 20 percent of private net worth. Such analyses are, however, subject to problems of underreporting and to emphasis on intended bequests (as opposed to including unintended bequests or *inter vivos* transfers).

In an influential paper, Kotlikoff and Summers (1981) examined the ability of the life-cycle model to explain the aggregate stock of wealth in the economy. In particular, they compared the relative contributions of life-cycle saving and intergenerational transfers in explaining wealth, and concluded that life-cycle wealth could not explain a significant share of the U.S. capital stock.

Concerns over the Kotlikoff-Summers approach have been voiced in subsequent research. First, as noted by Kessler and Masson (1988), the Kotlikoff-Summers estimates of transfer wealth require a number of controversial assumptions about, *e.g.*, the shape and stability over time of age-consumption and age-earnings profiles, and ages of family formation, retirement, and death. Using different assumptions than Kotlikoff and Summers, Ando and Kennickell (1987) and Modigliani (1988a) estimate that the overwhelming majority of net worth can be explained by life-cycle saving. Second, Blinder (1988) and Modigliani (1988a) have pointed out that the Kotlikoff-Summers conclusions are very sensitive to the treatment of consumer durables. The methodological debate over estimating life-cycle wealth in this research program has not been resolved (see, *e.g.*, the reprise in Kotlikoff and Summers, 1988, or in Kessler and Masson, 1989).

¹⁷Aiyagari (1992), however, argues that these effects are significantly smaller in a general equilibrium model. Direct empirical support for precautionary saving models is mixed; Dardanoni (1991), and Carroll and Samwick (1992a) estimate a large role for precautionary saving in explaining overall saving patterns, while Guiso, Jappelli, and Terlizzese (1992) suggest a much smaller role.

as a proxy). Table 1 reports the simulated aggregate asset-income ratios (total assets divided by total income) and saving rates for the no-high-school, high-school, and college education groups, using five sets of parameter-value assumptions (for γ and δ). The first, in which we assume $\gamma = 3$ and $\delta = 0.03$, is our benchmark case. To examine the consequences of different values of γ (setting $\delta = 0.03$), we consider two additional cases: $\gamma = 1$ and $\gamma = 5$.¹⁸ To examine the sensitivity of these results to different values of δ (setting $\gamma = 3$), we consider: $\delta = 0.015$ and $\delta = 0.10$. (We will later use the case in which $\gamma = 3$ and $\delta = 0.10$ to characterize results of “buffer stock” models.) For comparison, we also present actual asset-income ratios for different groups. These figures, calculated using data for 1984 from the PSID as total net worth for each group divided by total family income for the group, equal 3.69 for the no-high-school group, 3.83 for the high-school group, and 4.80 for the college group.

The first panel of Table 1 corresponds to the “certainty” version of our life-cycle model, with no consumption floor. Here, as well as throughout the paper, when we examine a model with certain earnings, medical expenses, and/or lifespan, we set the variables at every age equal to their expected value for that age as of age 21. For the benchmark case (reported in the first row), simulated asset-income ratios are in the range of 2 to 2.5, and saving rates range from about 2 percent to about 2.5 percent. The predicted asset accumulation is low relative to the average asset-income ratios calculated for the three education groups using the PSID.

The remaining rows of the first panel present simulated asset-income ratios and saving rates corresponding to different choices of γ and δ . Variation in γ does not affect the results in the certainty case, since $r = \delta$. An increase in the rate of time preference to 10 percent (while $\gamma = 3$) leads to still more implausibly low predicted levels of asset accumulation. A decrease in the rate of time preference to 1.5 percent increases asset accumulation relative to the benchmark level, though predicted asset-income ratios are lower than those suggested by empirical evidence.

The second panel of Table 1 corresponds to the “uncertainty” version of our model (*i.e.*, with uncertain lifespan, earnings, and out-of-pocket medical expenses) with no consumption floor.^{19,20} For the benchmark case ($\delta =$

¹⁸For technical reasons, γ is actually set equal to 1.0001 in the case for which the table is labeled “ $\gamma = 1$.”

¹⁹The steady-state saving rate equals the sum of the steady-state bequest-income ratio and the product of the population growth rate and the steady-state asset-income ratio. See Appendix B for details.

²⁰When the consumption floor is set equal to zero (actually, \$1), the government would still pay medical expenses, but would leave the family with essentially zero nonmedical consumption. Since $U'(0) = \infty$ under the preferences we have assumed, households save sufficiently so that they never rely on this program.

Table 1
Simulated Asset-Income Ratios and Saving Rates

	Parameter Assumptions	----- Asset-Income Ratio -----				----- Saving Rate -----			
		No High School	High School	College	Aggregate	No High School	High School	College	Aggregate
Certain lifespan, earnings, and out of pocket medical expenses	$\delta=.03, \gamma=3$	2.13	2.53	2.27	2.37	.021	.025	.023	.024
	$\delta=.03, \gamma=1$	2.13	2.53	2.27	2.37	.021	.025	.023	.024
	$\delta=.03, \gamma=5$	2.13	2.53	2.27	2.37	.021	.025	.023	.024
	$\delta=.015, \gamma=3$	2.79	3.27	2.78	3.03	.028	.032	.028	.030
	$\delta=.10, \gamma=3$	0.68	0.62	0.86	0.70	.007	.006	.009	.007
Uncertain life-span, earnings, and medical expenses $\bar{C} = \$1$	$\delta=.03, \gamma=3$	7.40	6.06	4.71	5.99	.167	.133	.119	.136
Uncertain life-span, earnings, and medical expenses $\bar{C} = \$7000$	$\delta=.03, \gamma=3$	4.72	4.74	4.20	4.59	.108	.108	.108	.108
	$\delta=.03, \gamma=1$	2.72	2.81	2.31	2.65	.053	.054	.051	.053
	$\delta=.03, \gamma=5$	5.75	5.74	5.23	5.60	.140	.137	.139	.138
	$\delta=.015, \gamma=3$	5.58	5.66	5.09	5.48	.131	.131	.132	.131
	$\delta=.10, \gamma=3$	2.11	1.92	1.59	1.87	.044	.039	.040	.040

Source: Authors' calculations.

Note: The actual ratios (using data from the 1984 PSID) of total net worth to total family income are 3.69, 3.80, and 4.80 for the three groups, respectively. The ratio of private net worth to aggregate disposable income, using 1984 data from the Federal Reserve's *Flow of Funds Accounts*, is 4.64.

0.03, $\gamma = 3$), the simulated average asset-income ratios substantially exceed actual asset-income ratios, with the largest discrepancy for the no-high-school and high-school education groups.

To conclude the first set of experiments, the third panel of Table 1 presents average asset-income ratios and saving rates by education assuming both uncertainty and the presence of a \$7000 consumption floor. For the benchmark case in which $\delta = 0.03$ and $\gamma = 3$, the availability of the floor reduces assets and saving rates in all education groups, though the reductions for the college group are small relative to the declines for the high-school group and, especially, the no-high-school group. This pattern holds qualitatively for the other four sets of parameter assumptions. In Hubbard, Skinner, and Zeldes (1993), we showed that our model can explain why the no-high-school group has a larger fraction of households with very low levels of wealth at any age than do other education groups. That result and the ones presented here together imply that while our model predicts that more households with low educational attainment have low levels of wealth, the *average* level of wealth relative to income is not predicted to differ significantly across education groups.²¹

Our second set of experiments focuses on the predictions of the augmented life-cycle model for *aggregate* wealth accumulation. These simulation results are reported in the fourth and eighth columns of results in each panel of Table 1. (The actual ratio for 1984 of private net worth, calculated from the Federal Reserve's Flow of Funds data, to aggregate disposable income is 4.64). For the benchmark case, the certainty results presented in the first row of the first panel indicate an aggregate asset-income rate of 2.37, which is comparable to simulated asset-income ratios calculated by Auerbach and Kotlikoff (1987) and Hubbard and Judd (1986) (though those authors used somewhat different parameter assumptions).²² This asset-income ratio is low relative to the actual ratio for 1984. The simulated saving rate of 2.4 percent is also much lower than observed saving rates out of disposable income, which were in the 5–8 percent range during the 1980s and early 1990s. While the aggregate asset-income ratio is higher when a lower rate of time preference (1.5 percent) is assumed, it still falls significantly short of the observed value in the data. A higher assumed rate of time preference (10 percent), of course, further reduces predicted aggregate asset accumulation. Hence, consistent with the findings of previous studies, the simulated aggregate asset-income ratio and the saving rate implied by the certainty version of the model are implausibly low.

²¹Saving rates in our base model also do not increase with educational attainment, whereas Attanasio (1993) provides evidence of substantial increases.

²²For example, these studies assume a rate of time preference equal to 1.5 percent per annum. They also use a general equilibrium model with a neoclassical production function.

As shown in the second panel of Table 1, adding uncertainty over lifespan, earnings, and out-of-pocket medical expenses raises the simulated aggregate asset-income ratio and saving rate in the benchmark case substantially to 5.99 and 13.6 percent, respectively. Though not shown, the simulated values increase with lower assumed values of the rate of time preference, δ , or higher assumed values of the coefficient of relative risk aversion, γ .

Finally, as shown in the third panel of Table 1, adding the consumption floor (equal to \$7000) to the uncertainty model reduces the asset-income ratio and saving rate in the benchmark case (relative to the values reported in the second panel) to 4.59 and 10.8 percent, respectively. Particularly regarding the asset-income ratio, the augmented life-cycle model for the benchmark case, with uncertain lifespan, earnings, and out-of-pocket medical expenses, and the consumption floor, matches aggregate data more closely than the standard certainty model.²³ The simulated values of the asset-income ratio for the cases in which the assumed rate of time preference is high ($\delta = 0.10$) or the assumed coefficient of relative risk aversion is low ($\gamma = 1$) are lower than the values for the benchmark case. The aggregate asset-income ratios implied in these cases are significantly smaller than the observed values. That is, “buffer stock” assumptions applied to all individuals in precautionary saving models (see, *e.g.*, Deaton, 1991; and Carroll, 1992) generate too little aggregate wealth. The cases for which (1) $\delta = 0.03, \gamma = 5$ and (2) $\delta = 0.015, \gamma = 3$ produced simulated values of the asset-income ratio higher than, but in the range of, observed values.²⁴

V. Can the augmented life-cycle model explain average wealth-age and consumption-age profiles?

In the certainty life-cycle models described earlier, the shape of the individual’s wealth-age profile depends upon the shape of the earnings-age profile and on the assumed values of the real interest rate (in partial equilibrium models), the rate of time preference, and the intertemporal elasticity of substitution in consumption. For example, using earnings profiles esti-

²³In this version of our model, bequests are not passed to younger generations, but are instead effectively confiscated by the government. An alternative approach would be to award the average bequests to the youngest individuals. Preliminary simulations found that this procedure increased the asset-income ratios by roughly 25 percent. However, this approach is not appropriate because it provides the younger generation with an unrealistically large supply of liquid assets. A more realistic approach would provide individuals with random bequests later in life. Unfortunately, this approach would be difficult to implement in our model.

²⁴For the parameter assumptions in the benchmark case, most of the additional saving generated in the “uncertainty” model can be traced to precautionary saving against earnings uncertainty. The relative importance of uncertainty over lifespan rises the lower is the rate of time preference or the lower is the coefficient of relative risk aversion.

mated in empirical studies, Auerbach, Kotlikoff, and Skinner (1983), Hubbard and Judd (1986), and Auerbach and Kotlikoff (1987) generally found hump-shaped wealth-age profiles, in which individuals dissave when young, accumulate assets in middle age, then dissave in retirement.

The wealth-age profiles predicted by certainty models, however, stand at odds with findings from the empirical studies (to which we referred in the introduction) that: a significant group of households appears to have no life-cycle saving, and the elderly appear to dissave at a rate slower than that predicted by the model. Carroll and Summers (1991), using Consumer Expenditure Survey data for the United States (1960–1961 and 1972–1973 surveys), showed that the consumption-age profile matches the income-age profile more closely than would be predicted by the standard life-cycle model. Carroll and Summers also find that the average consumption-age profile of college graduates is more humped-shaped than those for high-school graduates or high-school dropouts. These gaps in certainty models suggest a number of potential explanations, including heterogeneity in earnings-age profiles, imperfect insurance or capital markets, and the obvious candidate of omitted motives for saving.²⁵


In this section, we use the model developed in earlier sections to illustrate the effects of uncertainty and the “consumption floor” on average wealth-age and consumption-age profiles for the three groups.²⁶ We begin by plotting in Figure 1 wealth-age profiles (average assets at each age) for the three education groups, assuming certain lifespan (80 years), certain earnings, certain out-of-pocket medical expenses, and the parameter assumptions in the benchmark case analyzed earlier (*i.e.*, $\delta = 0.03$ and $\gamma = 3$). While the levels of asset accumulation differ across the education groups owing to differences in lifetime income, the familiar humped-shaped pattern predicted by the standard life-cycle model emerges. On account of binding borrowing constraints, households accumulate no wealth during the first 11 to 16 years of their working life (with the longest period of binding borrowing constraints for college graduates).

Though not illustrated here, simulated wealth-age profiles using alternative assumptions regarding the rate of time preference differ from the wealth-age profile in the benchmark case in a predictable way. With a lower assumed rate of time preference ($\delta = 0.015$), more assets are accumulated prior to retirement, and the borrowing constraint binds for fewer periods than in the benchmark case (ranging from 7 years for high-school dropouts to 13 years for college graduates). When a higher rate of time preference is assumed

²⁵Hubbard and Judd (1986), for example, showed that the introduction of borrowing constraints could explain a flatter wealth-age profile in the certainty model (on account of less dissaving by the constrained young, and, then, less saving in middle-age years).

²⁶These average profiles are equal to the expected values for a household as of age 21.

Figure 1
Average Assets by Age
All Certain, \$1 floor



($\delta = 0.10$), peak individual wealth accumulation prior to retirement falls by about two-thirds relative to the benchmark case, and borrowing constraints bind for about 25 years for all education groups.

The corresponding wealth-age profiles for the benchmark case ($\delta = 0.03$, $\gamma = 3$), assuming uncertain lifespan, earnings, and out-of-pocket medical expenses, are presented for the three education groups in Figures 2(a)-(c). Three features of the profiles are of interest. First, adding uncertain lifespan, earnings, and out-of-pocket medical expenses increases wealth accumulation at all ages; the borrowing constraints do not bind at early ages. Second, because more wealth is accumulated in the early and middle years of life, there is more dissaving than that found in models with only lifespan uncertainty (see, *e.g.*, Hubbard and Judd, 1987).²⁷ Third, the introduction of the consumption floor significantly reduces wealth accumulation at all ages for the no-high-school and high-school education groups, but has only a small negative effect on the wealth accumulation of college graduates.


Though not shown in Figures 2(a)-(c), the positive effect of uncertainty on preretirement asset accumulation is greater: the lower is the assumed rate of time preference, or the higher is the coefficient of relative risk aversion. For each of these alternative sets of parameter assumptions, the change in preretirement asset accumulation relative to the benchmark case is most pronounced for the no-high-school education group, followed by the high-school group, followed by the college group.

The average consumption-age profiles under various scenarios for the benchmark case of parameter assumptions are presented for the three education groups in Figure 3(a)-(c). The simulated profiles under certainty are straightforward in this case. Consumption rises with earnings initially while the borrowing constraint binds;²⁸ thereafter, since the real rate of interest is set equal to the rate-of-time preference, the consumption-age profile is flat. The introduction of lifetime uncertainty alone reduces consumption over time relative to the certainty profile, giving rise to a form of precautionary saving. When we add uncertain earnings and out-of-pocket medical expenses, average consumption starts out at a lower level than in the certainty and lifetime-uncertainty-only cases (recall that borrowing constraints no longer bind). Preretirement consumption growth is more rapid in the general uncertainty case than under certainty (*c.f.*, Zeldes, 1989b, under

²⁷The case in which only lifespan is uncertain maintains the binding borrowing constraints in the early years of life; in old age, as the conditional probability of dying rises, wealth decumulation is slower than would be predicted under certainty. The increase in wealth and slower dissaving in retirement are present in the results for all three education groups.

²⁸Consumption and earnings are not equal because consumption excludes medical expenditures. The sum of consumption and medical expenditures equals earnings in these periods.

Figure 2a
Average Assets by Age
No High School Degree



gamma=3, delta=.03

Figure 2b
Average Assets by Age
High School Degree




Figure 2c
Average Assets by Age
College Degree






Figure 3a
Average Consumption and Earnings by Age
No High School Degree




gamma=3, delta=.03

Figure 3b
Average Consumption and Earnings by Age
High School Degree



gamma=3, delta=.03

Figure 3c Average Consumption and Earnings by Age



gamma=3, delta=.03

earnings uncertainty); in old age, the effect of lifetime uncertainty dominates (since earnings are no longer uncertain), and the consumption-age profile is downward-sloping. These patterns hold for all three education groups. While the addition of a consumption floor of \$7000 has a relatively minor effect on the consumption-age profile for college graduates (Figure 3(c)), for the no-high-school (Figure 3(a)), and high-school groups (Figure 3(b)), initial consumption is significantly higher on account of the consumption floor, and the growth in preretirement consumption is correspondingly lower.

Carroll (1992) argues that a model of precautionary saving (with a high rate of time preference) might well explain the parallel movements in consumption and income. The findings in our model, that consumption tracks income more closely over the life cycle than would be predicted by the standard life-cycle model and that the average consumption-age profile of college graduates is more humped-shaped than those for high-school graduates or high-school dropouts, are also consistent with actual observations made by Carroll and Summers (1991) described earlier.

Carroll and Summers (1991, p. 321) note that it is difficult to reconcile the basic life-cycle model with the observation that individuals in education groups whose income peaks late in life (college graduates in our analysis) do not appear to borrow very much when they are young against their future resources. In our benchmark case (in which $\delta = 0.03$ and $\gamma = 3$), a “certainty” version of the life-cycle model with borrowing constraints does not generate a hump-shaped age-consumption profile. However, if we assume a higher rate of time preference, $\delta = 0.10$, while maintaining $\gamma = 3$, the age-consumption profile is hump-shaped for all education groups, with the most pronounced hump-shaped profile for college graduates.²⁹ In addition, if we incorporate uncertainty over lifespan, earnings, and out-of-pocket medical expenses, we find that the high rate-of-time-preference case ($\delta = 0.10$) generates a hump-shaped age-consumption profile. This profile peaks at an earlier age than the profile generated in our benchmark case summarized in Figures 3(a)-(c).

To summarize, wealth-age and consumption-age profiles predicted by the life-cycle model under uncertainty differ in important ways from those generated by standard certainty models. While our “uncertainty” version of the life-cycle model generates hump-shaped age-consumption profiles consistent with empirical observations from the United States under assumptions of a modest rate-of-time preference (as in our benchmark case) or a high rate-of-time preference (as in a buffer stock model), the assumption of a high rate-of-time preference leads to counterfactually small levels of individual and aggregate asset accumulation prior to retirement.³⁰

²⁹Decreasing the rate of time preference δ to 0.015 (which is less than the assumed real interest rate of 0.03) yields an upward-sloping age-consumption profile throughout life.

³⁰The observed hump-shaped consumption profile may also reflect changes over time in

VI. Can the augmented life-cycle model explain the short-run time-series properties of consumption, income, and wealth?

The final issue we consider is the extent to which our model can explain the relationship between observed short-run changes in consumption and income. A number of tests have been performed that examine this relationship. One approach is to test for the violation of the orthogonality conditions implied by the Euler equation (see, for example, Hall, 1978, and Flavin, 1981, using aggregate time series data, or Zeldes, 1989a, using panel data).

In our model, the Euler equation could fail to be satisfied in two ways. First, the borrowing constraint could bind: households could be at a corner, consuming all of their cash on hand and desiring to borrow to raise consumption. Second, the nonlinear Euler equation could be satisfied, but the log-linear approximation to that Euler equation could generate an apparent rejection.³¹

Under certainty, the life cycle model with empirically accurate age-earnings profiles predicts that borrowing constraints should bind only in youth. With uncertain earnings, out-of-pocket medical expenses, and lifespan, this will no longer be the case. In the presence of a consumption floor, borrowing constraints can bind at any time in the life cycle. The presence of means-tested social-insurance programs adds to the prevalence of Euler equation violations. The intuition is that in the presence of means-tested public-assistance programs, low-wealth and low-saving behavior can be an “absorbing state,” in the sense that there is little incentive to build up even a buffer stock of wealth to protect against uncertain contingencies (see Hubbard, Skinner, and Zeldes, 1993). Furthermore, these constrained families tend to remain constrained over time in our model. In contrast, the buffer stock model in Carroll and Samwick (1992b), for example, predicts that few, if any, families will allow themselves to become constrained in the sense that they hold little or no wealth.

In this paper, we do not attempt to disentangle the various reasons why consumption may “track” income. Rather, we examine how well our model, with its characterizations of borrowing constraints, uncertainty, and social insurance programs, matches the data. We do this in two ways. First, we examine the fraction of the observations in the model and in the PSID that sets average consumption approximately (within 0.5 percent) equal to average income over a five-year period. Second, we use simulated data to estimate an equation similar to the ones estimated by Campbell and Mankiw (1989) (using aggregate time-series data) and Lusardi (1993) (using panel data).

the demographic composition of the household (see, *e.g.*, Blundell, Browning, and Meghir, 1992).

³¹See, *e.g.*, the suggestion to this effect in Carroll (1992) and Hahm (1993).

and compare the resulting coefficients. We also examine how well the “buffer stock” version of our model (with the annual rate of time preference set to 10 percent and the consumption floor set to \$1) generates results that match the data.

In Table 2, we present the fraction of individuals at various ages (less than 29, 30–39, 40–49, 50–59, and 60–69) with saving less than 0.5 percent of current income. To do so, we use simulated panel data — ten-year histories of 2840 families in each education group, assuming uncertain lifespan, earnings, and out-of-pocket medical expenses. Calculations based on these data are compared with calculations based on the data from the PSID.

Table 2:
Percent of Actual (PSID) and Simulated Households With
Consumption Approximately Equal to Income
(Absolute Average Saving Rate < 0.5 Percent of Income)

Age	PSID			Simulated $\delta = .03, \text{Floor} = \7000			Simulated $\delta = .10, \text{Floor} = \1		
	NHS	HS	Col.	NHS	HS	Col.	NHS	HS	Col.
<29	.362	.059	.060	.332	.163	.263	.000	.000	.019
30-39	.157	.064	.040	.231	.093	.167	.000	.000	.073
40-49	.067	.017	.025	.212	.045	.035	.023	.005	.018
50-59	.103	.032	.011	.148	.027	.002	.023	.026	.007
60-69	.095	.020	.038	.131	.038	.015	.018	.005	.031
Total	.116	.038	.031	.204	.067	.086	.013	.007	.033
For Households with Initial Assets < 0.5 × Average Income									
<29	.389	.080	.069	.494	.226	.377	.000	.000	.040
30-39	.205	.101	.052	.484	.231	.314	.000	n/a	.070
40-49	.133	.045	.031	.613	.310	.189	.000	n/a	.000
50-59	.272	.114	.135	.565	.331	.000	n/a	n/a	.000
60-69	.302	.083	.381	.735	.362	.360	n/a	n/a	n/a
Total	.252	.087	.060	.545	.250	.303	.000	.000	.043

Source: PSID and authors’ calculations.

Notes: Average saving rates equal average annual saving during a five-year period (1984–89 for data from the PSID) divided by the average annual real income over the period. The tabulations reported in the second table are based on households whose initial assets (in 1984 in the PSID data) are less than half their average income. The symbol n/a denotes no simulated household in this cell.

The top panel of Table 2 reports fractions of households with an average absolute saving rate over a five-year period of less than 0.5 percent, by education group. The first three columns represent tabulations of data from the PSID. The 1989 wave of the Panel Study of Income Dynamics reports both net wealth in 1984 and net wealth in 1989. The Haig-Simons definition of saving during this period is just the difference in wealth between the two years. However, some part of the change in wealth may be the consequence of the receipt of capital gains, inheritances, or cashed-out pensions. To abstract from such “passive” saving, the Panel Study has constructed a measure of “active” saving that subtracts from overall saving these components of passive saving.³² One advantage of this measure is that it is not contaminated by measurement error in income since it is based entirely on changes in wealth. We define “constrained” households to be those for whom average active saving over the five-year period is within 0.5 percent of average money income (in 1984 dollars) during the same period.³³ The sample was restricted to households with heads aged less than 70 with no major compositional change during the period, and with total money income in excess of \$2000 in each of the relevant years. The occurrence of very low saving rates is greatest for the no-high-school group of households in the PSID at all ages, particularly at younger ages. In addition, a smaller fraction of households aged 40–60 have very low saving rates relative to the fraction of households at younger ages. The second set of three columns reports results from our simulated benchmark case with a consumption floor of \$7000. The simulated fraction of households with very low saving rates by age and education group mimics closely the patterns in the PSID data, though our simulated values overpredict somewhat the proportion of households with virtually no current saving.³⁴ The third set of three columns reports results from our simulated version of a buffer-stock model (with a rate of time preference δ of 0.10 and a minimal floor). With few exceptions, these simulated values grossly underpredict the fraction of households with very low current saving rates by age and education group.

The bottom panel of Table 2 carries out similar exercises, but restricts

³²The measure also adds contributions to annuities, and adjusts for assets of household members entering and leaving the household. An ultra-rational family might view such “passive” changes as partially predictable, and might offset passive saving by active dis-saving. Results are similar, however, when saving is defined simply as the real change in net worth between 1984 and 1989.

³³Even if the Euler equation is satisfied, it may be possible for consumption to move closely with income if the income process exhibits sufficient persistence. It is also possible that saving rates over the five-year period just happened to be close to zero for some unconstrained households. However, it is unlikely that the proportion of households falling into this group would differ systematically by education group.

³⁴Reasons that the PSID tabulations might diverge from our simulated data include measurement error and random fluctuations in the return to capital in the PSID.

the sample to households with low initial assets (less than one-half of current income in 1984, the beginning of the five-year period). The PSID tabulations suggest that very low levels of assets are an absorbing state; many households with low levels of wealth in 1984 had virtually no saving over the five-year period. As the bottom panel of Table 2 shows, the simulated data from our model are consistent with this behavior. For example, households in the no-high-school group with low initial wealth save very little over the period. The buffer-stock model, in contrast, implies that, if households' wealth falls below the buffer stock, they will replenish wealth up to the optimal buffer stock. That is, households with low initial wealth (in 1984) should have had high saving rates over the 1984–1989 period. Simulated values from the buffer-stock case generate almost no households who virtually consume their current income. The observed saving behavior of lower-wealth households appears more consistent with our model than with a buffer-stock model.

For the simulation exercises reported in Table 2, our model predicts that only about 10 percent of the population as a whole approximately consumes its current income. Can the interaction of uncertainty and borrowing restrictions explain the correlation of changes in consumption and income observed in aggregate data sometimes ascribed to “liquidity constraints”? Campbell and Mankiw (1989), for example, relate contemporaneous changes in consumption and income in a model nesting permanent-income consumers and “rule of thumb” consumers who consume their current income. When they assume that the real interest rate equals the rate of time preference and that the marginal utility function is linear, Campbell and Mankiw can identify the fraction of consumers who consume their current income from the coefficient in a regression of the change in consumption on the contemporaneous change in income (using instrumental variables). While the Campbell-Mankiw model is estimated using aggregate time-series data, Lusardi (1993) estimates a version using micro data from the PSID and the Consumer Expenditure Surveys and obtains similar results for nondurable consumption.³⁵

Using simulated data from our benchmark uncertainty case (in which $\delta = 0.03$ and $\gamma = 3$) with a consumption floor of \$7000, we estimate the Campbell-Mankiw model in levels and logs; coefficient estimates are reported in Table 3.³⁶ We estimate coefficients on the anticipated change in income of

³⁵Lusardi uses as instrumental variables dummy variables for marital status, sex, race, presence of children, composition of earners in the household, educational attainment, and occupation, as well as interactions of educational and occupational dummy variables with age. Because Lusardi estimates a log linear version of the Campbell-Mankiw specification, it is not possible to interpret her coefficient on (expected log) income as the fraction of households who are “rule of thumb” consumers.

³⁶We use two sets of instrumental variables. The first set includes second- and third-lagged values of consumption and income, as in Campbell and Mankiw (1989). The second set includes the first set plus the first lag of income (legitimate in our case owing to the

0.411 in levels and about 0.5 in logs, which is very similar to the aggregate-time-series estimate of 0.469 reported in Campbell and Mankiw (1989) and the estimate of 0.409 reported in Lusardi (1993). Hence, the benchmark case of the model replicates another feature of U.S. data, the relationship between contemporaneous changes in consumption and current income analyzed by Campbell and Mankiw and by Lusardi. The alternative version of the model in which the rate of time preference is much higher (*i.e.*, in which $\delta = 0.10$, but with no consumption floor, our analogue of a buffer-stock model) fares poorly in this respect; the (precisely) estimated coefficient on the anticipated change in income is about unity.

A puzzle remains, however: since only about 10 percent of the simulated population approximately consumes its current income, the estimated coefficients must reflect something other than the occurrence of binding borrowing constraints. We plan to explore this result in future research.

VII. Conclusions

This paper is one of two papers in which we examine predictions for wealth accumulation of a life-cycle simulation model in which individuals face uncertainty regarding their length of life, earnings, and out-of-pocket medical expenditures, and imperfect insurance and lending markets. In Hubbard, Skinner, and Zeldes (1993), we focus on the asset-based means-testing features of social insurance programs and their interaction with precautionary saving, and we find that the augmented life-cycle model can explain the low wealth accumulation of a significant fraction of the population. In this paper, we demonstrate that the aggregate wealth accumulation predicted by our augmented life-cycle model is more consistent with observed household asset holdings in the United States than are the predictions of conventional life-cycle models or buffer-stock models.³⁷ The perfect-certainty life-cycle model is unable to explain overall capital-income ratios, age-consumption profiles, and short-term comovements in income and consumption. The buffer-stock model can rationalize observed low levels of wealth among some parts of the population. However, it cannot easily account for the saving behavior of individuals who accumulate significant assets during their lifetime, except by assigning much lower rates of time preference to this group. By contrast,

absence of measurement error).

³⁷Our conclusion that precautionary motives are important for explaining U.S. saving suggests that differences across countries in uncertainty, public insurance, and private capital and insurance-market imperfections might help explain cross-country variations in saving rates. We do not claim, however, that these differences are the only or even dominant factors. Therefore, these differences alone are unlikely to explain why the saving rate in Japan is so much higher than in the United States, given the lower labor turnover and higher Social Security replacement rates in Japan (see Kitamura and Takayama, 1993).

Table 3:
Campbell-Mankiw-Lusardi Euler Equations
Using Simulated Data from the Dynamic Programming Model

Dependent Variable	ΔC^*	$\Delta \ln(C)^*$	$\Delta \ln(C)^{**}$	$\Delta \ln(C)^{**}$
$\delta = .03, \text{Floor} = \7000				
ΔY	0.411 (38.74)			
$\Delta \ln(Y)$		0.512 (40.11)	0.498 (41.29)	0.489 (14.63)
Age				0.109 (0.61)
Age ²				-0.163 (1.10)

Source: Simulated data (38520 observations) from the dynamic programming model under the benchmark case ($\delta = 0.03, \gamma = 3$) with a consumption floor of \$7000. Absolute values of t -statistics are in parentheses. The Campbell and Mankiw (1989) coefficient (in levels, corresponding to the first column) is 0.469, and the Lusardi (1993) coefficient (in logs, corresponding to columns 2 through 4) is 0.409.

* Instruments are ΔC and ΔY , each lagged two and three years, age, and age².

** Instruments are ΔC and ΔY , each lagged one, two, and three years, age, and age².

our approach assumes that all individuals have identical preferences and are subject to the same economic environment, and broadly matches empirical regularities in consumption and income. Nonetheless, the buffer-stock approach and our approach are complements in a research program to establish to what extent optimizing models of consumption under uncertainty can explain empirical patterns in consumption and wealth.

Based on the findings in our two papers, we conclude that a well-specified optimizing life-cycle model with precautionary saving can explain many important features of the world, including the saving behavior of much of the population. It is therefore likely to be an important tool for studying household saving decisions and the effects of taxation and social insurance on those decisions.³⁸ Our results suggest that the structure of many social-insurance programs depresses saving by some groups in the population. While beyond the scope of this paper (which has not addressed the efficiency consequences of those incentives), the finding may suggest the desirability of altering the saving incentives in these programs to the extent that the United States is believed to save too little.

Two other policy issues are raised by our findings. First, models of capital accumulation generally consider effects of government policy on saving through its impact on the after-tax return to saving or work. Our results suggest that tax policy may affect saving as much by shifting the variance of income as by changing the after-tax rate of return (see, *e.g.*, Barsky, Mankiw, and Zeldes, 1986; and Engen, 1992a). Our model also suggests that the state-contingent structure of government *expenditures* (transfers) can have a significant impact on household saving, especially for those households whose saving decisions might not be sensitive to changes in the after-tax rate of return. Second, the simulations reported by Bernheim and Scholz (1993) suggest that the saving of college graduates is likely to be more sensitive to tax-policy incentives suggested by the life-cycle model than the saving of less educated households. Their results indicate that the standard prescription for increasing household saving — increasing the after-tax rate of return — will not significantly alter saving rates for the no-high-school and high-school groups. Our results also suggest that the effects on saving of changes in marginal tax rates on capital income are likely to be different for high-lifetime-income households than for low-lifetime-income households. These policy implications must be viewed as preliminary, however. We would not, for example, rule out the possibility that other factors, such as a bequest motive or a simple desire to accumulate, influence the saving decisions of the very wealthy, thereby altering policy analysis.

³⁸Our findings cast significant doubt on the usefulness of the standard, perfect-certainty version of the life-cycle model for evaluating effects of tax policy and social insurance on household saving.

Two promising theoretical extensions of the model are the incorporation of a production sector to permit analysis of general equilibrium effects of precautionary saving³⁹ and the specification of multiple assets to examine effects of precautionary saving on portfolio allocation.⁴⁰ Some possible future applications include analyses of effects on saving of such government policies as changes in means testing of transfer programs, the introduction of tax-financed health insurance, and shifts in the mix of income and consumption taxes. Analysis of such policies to affect saving is particularly important if exogenous changes in saving rates affect long-run economic growth.

³⁹Aiyagari (1992) shows, for example, that the general equilibrium effect of precautionary saving motives (in particular, saving against uninsurable idiosyncratic fluctuations in earnings) on aggregate saving can be much smaller than the impact suggested in a partial equilibrium model. A similar point is made in the analysis of lifetime uncertainty and Social Security in Hubbard and Judd (1987). Nevertheless, both studies demonstrate potentially significant effects of precautionary saving on total saving in general equilibrium.

⁴⁰Recent theoretical analyses of this topic include Elmendorf and Kimball (1991) and Kimball (1993).

Appendix A:**Empirical estimates of lifespan, earnings, and medical expenses, and the consumption floor**

In this section, we first consider lifespan uncertainty. We then discuss the estimation of uncertain earnings and medical expenses, and conclude with a discussion of our calibration of the consumption floor \bar{C} .

Lifespan uncertainty

We consider that the family “dies” when the last member, assumed to be the wife, dies; we use mortality data on women for 1980 (Faber, 1982). The model initiates household decisions at age 21, and the maximum length of life is set equal to 100. The conditional mortality probabilities are presented in Table A.1.

Table A.1
Conditional Probability of Dying Before Attaining Age i ,
Having Attained Age $i - 1$

Age i	Probability	Age i	Probability	Age i	Probability	Age i	Probability
22	0.00058	42	0.00180	62	0.01033	82	0.06214
23	0.00061	43	0.00200	63	0.01124	83	0.06885
24	0.00062	44	0.00221	64	0.01223	84	0.07631
25	0.00064	45	0.00242	65	0.01332	85	0.08455
26	0.00065	46	0.00266	66	0.01455	86	0.09352
27	0.00067	47	0.00292	67	0.01590	87	0.10323
28	0.00069	48	0.00320	68	0.01730	88	0.11367
29	0.00070	49	0.00349	69	0.01874	89	0.12484
30	0.00072	50	0.00380	70	0.02028	90	0.13677
31	0.00075	51	0.00413	71	0.02203	91	0.14938
32	0.00078	52	0.00450	72	0.02404	92	0.16289
33	0.00082	53	0.00490	73	0.02623	93	0.17721
34	0.00086	54	0.00533	74	0.02863	94	0.19234
35	0.00091	55	0.00581	75	0.03128	95	0.20828
36	0.00098	56	0.00632	76	0.03432	96	0.22418
37	0.00105	57	0.00689	77	0.03778	97	0.23980
38	0.00115	58	0.00749	78	0.04166	98	0.25495
39	0.00128	59	0.00811	79	0.04597	99	0.26937
40	0.00144	60	0.00878	80	0.05078	100	0.28284
41	0.00161	61	0.00952	81	0.05615	101	1

Source: Faber (1982)

Note: We set the probability at age 101 equal to 1, rather than the actual value given in the mortality tables.

The age-earnings profile and earnings uncertainty

The typical approach to estimating age-earnings profiles is to regress the log of earnings on a polynomial of age and year-dummy variables. Observations with low (or zero) earnings are excluded from such an analysis so that the implied coefficients (and variances) are not excessively influenced by the extremely low values of the log transformations. In our analysis, we do not want to exclude these low earnings draws.⁴¹ To estimate the conditional expectation of earnings at a given age, we first regress the level of earnings on age. While the standard errors of the coefficients may be biased owing to heteroscedasticity, the purpose of this first-step regression is to characterize the age-earnings profile accurately. Note that we include in earnings the combined labor income of the husband and wife (if married) plus unemployment insurance. We adjust for taxes by multiplying earnings by one minus the average tax rate for families in that income class.

The first-step regressions of earnings on a cubic polynomial in age and year dummy variables are reported, for each education group, in Table A.2.⁴² Earnings are more peaked for college-educated families, a result that has been found in many studies. The fitted values for earnings (adjusted for the year dummy) and the actual values averaged by age are presented in Figures A.1(a) - A.1(c).⁴³

The next step is to estimate the expected value of pension, Social Security, and non-means-tested income at retirement. Because we want to focus on the expectation of such income at a given age, we estimate the level of this retirement income in levels, including just age and year dummy variables.⁴⁴ These results are presented in Table A.3, and also graphed in Figures A.1(a)-(c). The implied income path decreases with age, which may reflect pensions that are imperfectly indexed for inflation. It may also reflect cohort effects (although these were at least partially controlled for in using year-dummy variables).

⁴¹Carroll (1992) and Hansen and Imrohorglu (1992) address this problem by modelling explicitly the probability of unemployment or a low-income event.

⁴²Murphy and Welch (1990) suggested that a quartic polynomial in age is a more accurate representation of the earnings profile. For our much smaller sample sizes, including a fourth-order term yielded unstable estimates at low ages.

⁴³Note that the fitted values of earnings are not smooth across age groups. This variation reflects the fact that the "cross-section" earnings profile includes data from five different years. For example, suppose 70 percent of the age s earning group were sampled in 1982, but 70 percent of the age $s + 1$ earnings group were sampled in 1986, a year with higher average earnings. The predicted earnings for age $s + 1$ would reflect the higher aggregate earnings (through the year-dummy effects), causing a slight jump in the predicted (and actual) value between age s and $s + 1$.


⁴⁴Because of small sample sizes and the small number of people over age 80, higher-order polynomials of age implied unstable retirement income at higher ages.

Table A.2
Parameter Estimates of Age-Earnings Profile
Age 65 and Under: 1982–86

Parameter/ Variable	< 12 Years of Education	12-15 Years of Education	16 + Years of Education
Constant	10,993 (0.72)	-11,833 (1.23)	72,270 (1.61)
Age	-196 (0.17)	1421 (1.91)	-5579 (1.69)
Age ² /100	2167 (0.81)	-137 (0.07)	19,200 (2.42)
Age ³ /10,000	-2655 (1.30)	-2186 (1.45)	-18,076 (2.93)
Year = 1982	-1054 (1.30)	-754 (1.37)	-1370 (0.86)
Year = 1983	-596 (0.73)	-585 (1.06)	-1380 (0.87)
Year = 1985	391 (0.50)	369 (0.67)	346 (0.22)
Year = 1986	-75 (0.09)	620 (1.12)	2100 (1.31)
R^2	0.033	0.072	0.091
N of Cases	1800	5291	2005


Notes: Dependent variable is the level of after-tax, after-unemployment-insurance-benefits labor income of the household head and spouse, if present. All regressions are weighted with family population weights. Asymptotic t -statistics are reported in parentheses.

Figure A1(a)
Predicted and Actual Labor and Retirement Income
No High School Degree



Source: PSID and authors' calculations

Figure A1(b)
Predicted and Actual Labor and Retirement Income
High School Degree




Source: PSID and authors' calculations

Figure A1(c)

Predicted and Actual Labor and Retirement Income

College Degree



Source: PSID and authors' calculations

Table A.3
 Predicted Pension and Social Security Income
 (Age > 65)

	< 12 Years of Education	12-15 Years of Education	16 + Years of Education
Constant	17,861 (7.73)	29,733 (6.36)	48,123 (4.29)
Age	-133 (4.34)	-245 (3.89)	-429 (2.82)
Year = 1982	-544 (1.01)	-645 (0.59)	1403 (0.52)
Year = 1983	-320 (0.61)	-74 (0.07)	1159 (0.46)
Year = 1985	151 (0.29)	-10 (0.01)	578 (0.23)
Year = 1986	446 (0.87)	780 (0.75)	3622 (1.45)
N of Cases	709	451	108
M.S.E.	4378	7159	8312
R^2	.029	.034	.088

Note: The models are estimated in levels using data from the PSID. Family weights are used. Asymptotic *t*-statistics are reported in parentheses.

Table A.4
Parameter Estimates of Earnings Uncertainty

Parameter/ Variable	Earnings: < 12 Years of Education	Earnings: 12-15 Years of Education	Earnings: 16 + Years of Education
Estimated Error Structure (Optimally-Weighted MDE)			
AR(1) term (ρ)	0.955 (8.94)	0.946 (7.30)	0.955 (7.91)
i.i.d. error variance (σ_e^2)	0.033 (0.43)	0.025 (0.40)	0.016 (0.40)
measurement error plus transitory income variance (σ_v^2)	0.040 (0.53)	0.021 (0.39)	0.014 (0.42)

Notes: The equations estimated are: $y_{it} = Z_{it}\beta + u_{it} + v_{it}; u_{it} = \rho u_{it-1} + e_{it}$. The dependent variable, y_{it} , equals the log of after-tax, after-unemployment-insurance-benefits labor income of the household head and spouse, if present.

Year dummy variables included in earnings regressions (coefficients not reported). Absolute values of t -statistics are reported in parentheses. The estimates use *unweighted* data.

We return now to earnings to characterize its time-series properties and variance by education group. Rather than adopting an approach that emphasizes the discrete nature of unemployment (as in Carroll, 1992), we instead sample only those whose combined earnings and unemployment insurance benefits are at least \$3000 per year. Obviously, this second data set, which does not include those families who are not in the labor force, unemployed or low earners, understates the true degree of uncertainty in earnings. We estimate the age-log-earnings profile using a cubic polynomial in age and dummy variables for years (not reported), with residuals calculated conditional on the age-earnings profile. The resulting matrix of log residuals used to estimate subsequently the time-series covariance structure is of dimension $N_j \times 5$, since there are five years (1982, ..., 1986) and N_j families in the j th education class.

The third step is to use the matrix of residuals to estimate the model described in the text. As noted above, while we estimate the variance of ν , we do not use it in the simulation, which tends to bias downward the estimated uncertainty.⁴⁵ The parameters ρ and σ_c^2 are estimated using optimally weighted GMM, where the elements of the weighting matrix are the fourth moments of the 5×5 covariance matrix. Parameter estimates are reported in Table A.4.

Including an individual fixed effect in the earnings equation process would tend to reduce the persistence and magnitude of the earnings shocks. We do not pursue this parameterization of the model for two reasons. First, by separating our data by educational group, we have already controlled for a large part of permanent differences in earning ability. Second, the entire dynamic programming model would have to be calculated for *each* fixed effect, a task that at this point is computationally infeasible.

The final step is to merge the age-earnings profile, estimated in levels, with the uncertainty parameters estimated in logs. To ensure that the expected value of a log-normal earnings distribution is held constant in the simulations, we adjust the predicted *level* of earnings to account for σ_u^2 , the unconditional variance of earnings. Hence for individual i at time t :

$$W_{it} = \exp\{\log[Z_{it}\beta - 0.5\sigma_u^2 + u_{it}]\}.$$

This specification ensures that when we compare the certainty case with the earnings uncertainty case, we hold the age-conditional means of earnings constant. Of course, when there is no uncertainty ($\sigma_u^2 = 0$), the expression collapses to the level regression estimated in the first step.

⁴⁵On the other hand, we assume that the individual has no information aside from current earnings to forecast future earnings. Additional information about future earnings that is not observed by the econometrician biases upward our estimate of overall earnings uncertainty.

The age profile of and uncertainty about out-of-pocket expenses for medical care

In this section, we describe the procedure for estimating empirical measures of health costs to be used in the dynamic programming model.⁴⁶ The 1977 National Health Care Expenditures Survey (NHCES) provides the most recently available detailed information on out-of-pocket health costs in a large sample of noninstitutionalized families. We first use this survey to measure the cross-sectional distribution of health-care costs, and update these measures to the 1984 benchmark year. We then combine both data on nursing-home costs (which are entirely excluded from the NHCES) and estimated parameters from panel data discussed in Feenberg and Skinner (1992) to completely characterize the health-care uncertainty faced by families.

Out-of-pocket medical expenditures in the NHCES include hospital, physician, drug, dental, and related expenses, as well as insurance premiums. We also include Medicaid-financed medical costs to ensure consistency with the dynamic programming model. Recall from the third section in the paper that total uninsured medical costs are the relevant state variable; if the household is eligible for Medicaid, then it need not pay for those costs. Hence the appropriate measure of the uncertain health-care cost includes charges subsequently covered by Medicaid.

One problem with these data is that they specifically exclude nursing-home expenditures. Institutionalized individuals were not surveyed, and even if the surveyed individual had just been discharged from a nursing home, he or she was not asked about those expenses.⁴⁷ Nursing-home expenses are clearly a large fraction of out-of-pocket expenses since they are rarely covered by Medicare or by private health insurance. According to the Health Insurance Association of America (1988), 81 percent of all out-of-pocket expenses for the elderly in excess of \$2000 were spent for nursing-home care. Hence it is important to include such expenses.

Data from the National Nursing Home Survey in 1977 provided measures of length-of-stay and nursing-home costs for the entire population of institutionalized individuals. The yearly cost in the nursing home can be approximated by the product of the cost per diem and average length of stay. Based on detailed information in the Nursing Home Survey summary, we calculate average length of stay as the midpoint of the length-of-stay categories (so that those who have been in the nursing home between three and six months are assigned an average length of stay of $4\frac{1}{2}$ months); those with lengths of stay in excess of one year are assigned a length of stay of one year.

⁴⁶While out-of-pocket medical expenses are small on average (see, *e.g.*, Rossiter and Wilensky, 1982), they are highly skewed (see Kotlikoff, 1988a). In addition, lengthy periods of nursing-home care could easily dissipate the financial assets of most households.

⁴⁷We are grateful to Michael Palumbo for pointing this out to us.

The assumed nursing-home yearly cost is simply the average length of stay times the per diem rate, assumed constant across individuals at \$44.49 per day.⁴⁸

Using cross-sectional measures of length of stay will understate the number of short-term nursing-home stays, because the point-in-time survey misses all those individuals who spent time in a nursing home during 1977, but were not present when the survey was compiled.

It is difficult to determine whether these nursing-home patients are separate “households” entirely missed by the NHCES, or whether they are family members in extant households included in the NHCES survey universe. The Nursing Home Survey includes information on whether the formerly noninstitutionalized patient had been living alone or with others. The fraction of those who had been alone ranged from 29.8 percent for those with lengths of stay less than three months to 34.8 percent for those with lengths of stay in excess of one year. (It was assumed that those transferred to the nursing home from other institutions exhibited similar household compositions.)

An example is instructive: There were 12,900 nursing-home residents aged 70–74 experiencing an average length of stay of 135 days in 1977. Of these patients, 35.57 percent had previously been living alone. Hence the NHCES “missed” 4589 households ($0.3557 \times 12,900$), with each individual experiencing an out-of-pocket nursing-home expense equal to \$6006. We assume a uniform distribution within each age group, so that there were $918 = 4589/5$ individuals at each age 70–74.

The nursing-home expense ignores the likely out-of-pocket expenses for these patients during the remaining 225 days of the year. We added to this figure the average out-of-pocket expense for noninstitutionalized families during the remaining 60 percent of the year. The estimated overall annual spending for each of the 12,900 nursing-home patients experiencing a length of stay equal to 135 days was therefore \$6439. As noted above, the 4589 “missing” institutionalized households were added to the NHCES universe of noninstitutionalized households.

The remaining problem is to assign expenses for households with members in nursing homes. Since there is no way to link nursing-home residents with particular households, we took the remaining nursing-home patients, and assigned them randomly to existing NHCES households with a household head over age 54. A random-number generator was used to determine first whether the existing household might have a member in the nursing home. Conditional on having a family member in the nursing home, the family was assigned a random value for the length of stay consistent with the measured

⁴⁸This was based on the 1985 Nursing Home Survey average monthly charge of \$1456, deflated to 1984 dollars using the medical care CPI and dividing by 30 days for a per diem rate of \$44.49.

probabilities.

Table A.5 provides summary statistics of these medical expenses, with appropriate population weights. Owing to roughly 3000 observations with missing data entries, the weights were scaled up to represent the total household population of 74 million families. All costs were adjusted using the medical-care price deflator to the 1984 benchmark year, while income brackets were adjusted to 1984 dollars using the CPI-U index.

Table A.5
 Summary Statistics from the National Health Care
 Expenditure Survey (1977) and National Nursing Home Survey (1977)

	Sample Size	Population Weights ¹	Medical Costs ²	Insurance Premium	Medicaid	Total
Entire Sample	10,427	74.07	912	897	211	2020
AGE						
Less than 31	2214	16.10	520	859	179	1558
31-45	2660	20.00	813	1086	238	2137
46-60	2638	18.96	1076	1117	201	2394
61-75	2194	14.33	1175	579	207	1961
Over age 75	721	4.69	1218	310	254	1782
EDUCATION						
< 12 Years Educ.	4401	27.95	905	704	413	2023
12-15 Years Educ.	4583	34.19	868	994	112	1974
16 + Years Educ.	1443	11.93	1055	1074	21	2149
INCOME						
< \$5,000	1074	6.69	733	187	470	1390
\$5-10,000	1349	8.53	811	275	581	1667
\$10-20,000	2452	16.62	835	595	231	1661
\$20-35,000	2754	20.26	910	1083	115	2109
\$35-50,000	1437	11.00	978	1425	90	2493
\$50-75,000	842	6.76	1100	1499	12	2612
> \$75,000	519	4.21	1242	1243	60	2545

Notes: ¹Figures are in millions. ²Figures are for total medical expenses less the sum of Medicaid and insurance premiums. Figures are expressed in 1984 dollars; medical costs are adjusted by the medical care CPI; income is adjusted by the average CPI.

Source: NHCES (1977) and National Nursing Home Survey (1977).

Out-of-pocket total medical expenses per household show a pattern by age that peaks between the ages of 45 and 60. There are two reasons for this pattern. The first is that, not surprisingly, larger households spend more on medical care than smaller households, and during retirement households tend to be smaller. The second is the fall in insurance premiums by the elderly as they become eligible for Medicare coverage; household insurance premiums

decline from \$1117 at age 45–60 to \$310 for those over age 75.

Finally, out-of-pocket medical-care costs are shown to rise with income level, from \$1390 for those with income level under \$5000 to \$2545 at income levels over \$75,000. Much of this variation in spending by income is a consequence of insurance premiums. Taking just medical costs plus Medicaid and excluding insurance premiums, spending by families with income between \$5000 and \$10,000 is \$1392, while for families with income between \$50,000 and \$75,000, the comparable figure is \$1113. That is, aside from the insurance component of medical expenses, combined out-of-pocket plus Medicaid payments decline across income groups.⁴⁹

Table A.6
Log Out-of-Pocket Medical Spending

Education	< 12 Years of Education		12-15 Years of Education		16 + Years of Education	
	< 65	65 +	< 65	65 +	< 65	65 +
Age of Head						
Constant	5.373 (17.77)	14.441 (3.12)	4.749 (20.24)	11.371 (1.66)	4.816 (10.31)	9.553 (0.74)
Age	0.073 (4.87)	-0.200 (1.63)	0.106 (8.77)	-0.101 (0.55)	0.109 (4.68)	-0.054 (0.16)
Age ² /1000	-0.753 (4.32)	1.288 (1.58)	-1.084 (7.48)	0.540 (0.44)	-1.090 (3.99)	0.297 (0.13)
Number of Observations	2716	1391	3845	522	1182	198
M.S.E.	1.368	1.305	1.245	1.147	1.105	1.236
R ²	0.013	0.001	0.040	0.007	0.039	0.008
Variance of i.i.d. error term, σ_ϵ^2	0.175		0.156		0.153	
Variance of measurement error + transitory expenses, σ_ω^2	0.220		0.220		0.220	

Notes: The equations estimated are: $m_{it} = G_{it}\Gamma + \mu_{it} + \omega_{it}$; $\mu_{it} = \rho_m\mu_{i,t-1} + \epsilon_{it}$. The dependent variable m_{it} , is the log of out-of-pocket medical expenses. Absolute values of t-statistics are reported in parentheses.

Log medical expenses estimated using NHCES (1977) with population weights. σ_ϵ^2 and $\rho_m = 0.901$ are estimated from data in Feenberg and Skinner (1992).

⁴⁹To measure accurately the true moments of medical-care expenses, we have excluded those reporting no medical expenses, roughly six percent of the sample, from the subsequent regression analysis as if they were nonrespondents. This might be expected to bias upward the predicted values of medical-care spending if in fact they had spent only a few dollars over the year on medical-care services. Different measures of medical-care costs compiled by the U.S. House Select Committee on Aging (1990) suggest, however, that, even with the exclusion of this group of apparently very healthy people, we are still understating the average out-of-pocket costs to the elderly.

Estimated coefficients from regressions of the log of total out-of-pocket expenses inclusive of Medicaid for both those 65 and over and for those under age 65 are shown by education group in Table A.6. Little of the total variation in medical expenses is explained just by age and age². These regressions are used in the dynamic programming model to benchmark the age pattern of medical expenses.

The structural model assumed for the purpose of the dynamic programming model is AR(1). Without time-series data, of course, we are unable to identify the AR parameter. Instead, we draw on information from out-of-pocket medical expenses taken as deductions, which were documented in a panel of taxpayers during the years 1968, 1970, 1972, and 1973. During these years, taxpayers that itemized their deductions could deduct medical expenses in excess of 3 percent of adjusted gross income. More than half of all (exogenous) itemizers declared medical expenses in a given year. Feenberg and Skinner (1992) have estimated a general quadrivariate tobit model of medical spending with an ARMA(1,1) error structure to characterize the dynamic properties of these tax data. The tobit procedure adjusts for the possible censoring of medical expenditures below three percent of adjusted gross income. We use their quadrivariate tobit estimator to consider a simpler model of medical expenses described in equation (9). The estimated parameter coefficients, reported in Table A.6, yield both the AR(1) parameter, 0.901, and the share of the residual variance due to transitory or measurement error (σ_ω^2). Although the data from Feenberg and Skinner are restricted to those whose age exceeds 55 in 1968, we apply the estimated AR(1) parameter, 0.901 to all ages. We allow for an individual fixed effect in medical spending, to reflect differences in insurance coverage or in underlying health status, but the estimated variance of permanent individual effects is near zero.

We suspect that the older tax data understates the overall variation in out-of-pocket medical expenses. Furthermore, although the NHCES data contain information on medical spending by education, the tax data do not. We therefore use the unconditional error term from the NHCES and nursing-home data to impute the “true” variance of the white noise error term.⁵⁰

The consumption floor

The consumption floor for an individual depends on many factors such as age, number of children, and state of residence. We attempt to calculate the value of the floor for a representative (eligible) individual, and assume that

⁵⁰That is, we know that the unconditional variance is equal to $\sigma_\mu^2 + \sigma_\omega^2$. We know the ratio $\kappa = \sigma_\mu^2 / (\sigma_\mu^2 + \sigma_\omega^2)$ from the tax data, and we know that $\rho_m = 0.901$. For the case of no-high-school-degree families, $\sigma_\mu^2 + \sigma_\omega^2 = 1.24$ (averaging across elderly and nonelderly families), so that $\sigma_\mu^2 = 1.24\kappa(1 - \rho_m^2)$.

recipients treat the transfers as if they were cash.⁵¹ Below, we describe the construction of consumption floors for the nonelderly and elderly populations, respectively.

The non-elderly: AFDC and food stamps. Committee on Ways and Means (1991, p. 1078) has calculated the net transfers from food stamps and Aid to Families with Dependent Children (AFDC) benefits for a female-headed family with two children in 1984 as \$7521 in 1990 dollars. (The median number of children for AFDC households was two in 1984.) Converting the 1990 dollar figure to 1984 dollars using the CPI-U index yields annual benefits of \$5979. The representative “family” receiving AFDC and food stamps is assumed to be one with a female head and two children. An adult male who is not working would be eligible only for food stamps.

The non-elderly: housing subsidies. Here we concentrate on the Section 8 housing program, which provides housing vouchers for existing rental property. Local housing conditions determine a “Fair Market Rent” (FMR) determined by the government. The Department of Housing and Urban Development will pay the difference between the apartment rent (up to the FMR) and 30 percent of eligible income; nearly every apartment rent is at least the FMR.

The mean “fair market rent” in the United States for 1984 was \$379 per month (based on unpublished calculations by the Economic and Market Analysis Division, U.S. Department of Housing and Urban Development, HUD). Hence, the overall subsidy at zero income is potentially \$4548. This number is likely to be too large for two reasons. First, AFDC benefits are included in the HUD definition of “income” for the purpose of the 30 percent of income paid as rent. Second, the housing subsidy is not an entitlement – not everyone who wants it can get it. As a result, we calculate the expected value of the benefit: the mean value of Section 8 subsidies times the probability of getting the subsidy. The probability is assumed to be the available “units” as a percentage of the target group of “very poor,” with an assumed 10 percent of the units unavailable (because they are being used for people who are not very poor). These calculations were presented by the Congressional Budget Office (CBO) in *Current Housing Problems and Possible Federal Responses* (December 1988).

From the Committee on Ways and Means (1991), median maximum AFDC benefits by state in January 1985 averaged \$332 per month. Annual rental payments are therefore expected to be 30 percent of that, or $332 \times 12 \times 0.3 = \1195 . Hence, the net housing subsidy is $379 \times 12 - 1195 = \$3353$, conditional upon actually receiving it.

From the CBO study (p. 55), we take the percentage of eligible families

⁵¹While this assumption is reasonable for food stamps and for Aid to Families with Dependent Children (AFDC), it is perhaps less so for Section 8 housing assistance.

with one or two children who are receiving Section 8 to be 35 percent. Hence, the expected value of the subsidy equals $3353 \times 0.35 = \$1173$.

Combined with the AFDC and food stamp subsidy, we estimate an average "consumption floor" for nonelderly households of \$7152 ($\$5979 + \1173).⁵²

The elderly: SSI and food stamps. Once again, there is an interaction between SSI and food stamps that must be accounted for in calculating total benefits. The Committee on Ways and Means (1991) calculates these cumulative benefits. Calculations for benefits in 1984 are reported (on p. 748) in Table A.7.

Table A.7
SSI, Social Security, and
Food Stamp Benefits for the Elderly

	Single	Couples
SSI Benefits	\$3738	\$5664
SSI and Social Security	4008	5904
SSI, Social Security, and Food Stamps	4294	6393

Source: Authors' calculations.

The Ways and Means publication does not report the concept we need for our calculations: SSI plus food stamps (without Social Security). To obtain that measure, we add the incremental income from food stamps given Social Security ($\$286 = \$4294 - \$4008$) to the SSI benefits, for a combined benefit of \$4024. (This is likely to be a lower bound, since food-stamp benefits are reduced as SSI plus Social Security income is increased.)

We also adjust this to account for state supplements to SSI benefits. The median state supplement in January 1985 was \$36 for individuals in 1990 dollars (p. 741). However, only 42 percent received the supplement in 1990 (Committee on Ways and Means, 1991, p. 740). Correcting for eligibility and converting to 1984 dollars yields an annual supplement of $(36 \times 12) \times 0.42 \times (103.9/130.7) = \144 ; the sum for individuals is \$4168 ($\$4024 + \144). The federal SSI benefit plus food stamps totals \$6153 for couples. Noting

⁵²For purposes of comparison, consider similar numbers for representative high, medium, and low states from Keane and Moffitt (1991): \$14,456 (California), \$10,296 (New York), \$9932 (Ohio), \$9464 (Kansas), \$8580 (Alabama), and \$9412 (Texas). The Keane and Moffitt calculations exceed our estimates of the average value of the consumption floor. Virtually all of this difference is accounted for by the fact that we weight the housing subsidy by the probability of getting it. In addition, we abstract from public-housing subsidies.

that the (1990\$) median state SSI supplement was \$66 in January 1985, an application of both the CPI adjustment and the 42-percent participation yields total benefits of \$6417.

To weight singles and couples for an average, we use information from the Current Population Reports Series P-60, No. 151 (U.S. Department of Commerce, Bureau of the Census, *Money Income of Households, Families, and Persons in the United States: 1984*, which reports the fraction of married to unmarried families was 54.8 percent for those with money income less than \$5000 (4.37 million/7.97 million). Assigning these weights yields a final average of \$5400.

The elderly: housing subsidies. We follow the same procedure for the elderly as for the nonelderly to estimate housing subsidies. We include SSI and food stamps in income, so that the rent that people are expected to pay equals $\$5400 \times 0.3$, or \$1620. Hence, total housing subsidies, weighted by the 51-percent provision to elderly families (and using the 90-percent commitment) yields $(\$4548 - \$1620) \times 0.51$, or \$1493.

A summary of the consumption floor. For the elderly population, the combined non-Social-Security consumption floor is \$6893 ($\$5400 + \1493). The average of the two measures (young and old) – our proxy for the consumption floor – is \$7022. We round this to \$7000 for use in our numerical model. As an estimate of the “true” consumption floor, this measure must be interpreted cautiously. We have not accounted for differences in eligibility status among those under age 65. For example, a married couple under age 65 not in the labor force, and with grown children, will receive very little in welfare benefits.

Appendix B

Numerical Solution and Aggregation

Numerical solution of the dynamic model

The household's problem. As outlined, we do not believe that the individual household's problem can be solved analytically, so we use numerical stochastic dynamic programming to approximate closely the solution. We calculate explicit decision rules for optimal consumption as well as the consumer's value function. (The program was written in FORTRAN and is about 1300 lines of code, exclusive of comments). We run the program on the Cornell National Supercomputer Facility. For reasons that will become evident below, the computer resources required to solve for optimal consumption are substantial.

Consider first the state variables for this problem. Cash on hand summarizes the financial resources currently available to the household. Earnings and medical expenses are assumed to follow first-order autoregressive pro-

cesses, so that current earnings and current medical expenses summarize the information known about future earnings and medical expenses. The problem therefore has three state variables (in addition to age): cash on hand, current earnings, and current medical expenses.

We first describe how we discretize the exogenous variables and random variables: earnings and medical expenses. The residuals around the deterministic trend are discretized into n different outcomes, where n equals 9 in our calculations. The n outcomes range between 2.5 standard deviations (of the unconditional distribution) above and below the age-specific mean earnings or medical expense outcome. Hence if a family remains on the same discrete node from one year to the next, its earnings or medical expenses will still change due to the trend. Given that the assumed dynamic process of earnings and medical expenses is AR(1), this discretized distribution corresponds to a first-order Markov process, with an $n \times n$ matrix of transition probabilities.⁵³

We discretize the endogenous variable cash on hand as follows. We first find the maximum feasible cash on hand for each period by calculating how much wealth would be accumulated if consumption were set equal to zero in every period (*i.e.*, if all income were saved) and the household received the maximum realization of earnings and the minimum realization of medical expenses in every period. We divide the maximum feasible range for cash on hand in each period into 61 nodes. These nodes are spread evenly on the basis of the log of cash on hand in order to get finer intervals at lower absolute levels of cash on hand, at which nonlinearities in the consumption function are most likely.

The dynamic programming problem is solved by starting in the last possible period of life (T) and solving backwards. In period T , C_T equals X_T . From the beginning of retirement until period T , individuals face no uncertainty from earnings but do face uncertainty from medical expenses and lifespan. Accordingly, we calculate optimal consumption for each of the (9 x 61) possible combinations of values for medical expenses and cash on hand. In the periods prior to retirement, households face uncertainty about medical expenses, lifespan, and earnings, and therefore we calculate optimal consumption for each of the (9 x 61 x 9) possible combinations of values for medical expenses, cash on hand, and earnings.

At each combination of state variables, we search for the optimal value

⁵³For each node i , the probability of moving to j is equal to the area under the (conditional on node i) log-normal distribution from the midpoint between node j and $j - 1$ to the midpoint between node j and $j + 1$. For node 1, the lower bound is 0, and for node n , the upper bound is infinity.

of consumption. The Euler equation at each node is:

$$U'(C_s) - \frac{(1+r)p(\text{live}_{s+1})}{1+\delta} \sum_{i=1}^n \sum_{j=1}^n p(W_{s+1}^i)p(M_{s+1}^j) \times \\ U'(C_{s+1}^*(X_{s+1}|W_{s+1}^i, M_{s+1}^j)) = 0, \quad (B.1)$$

where $p(\text{live}_{s+1})$ is the conditional one-period survival probability, $p(W_{s+1}^i)$ and $p(M_{s+1}^j)$ are the conditional probabilities of the age $s+1$ discrete earnings outcome i and the discrete medical expense outcome j occurring, and $C_{s+1}^*(X_{s+1}|W_{s+1}^i, M_{s+1}^j)$ is the optimal age $s+1$ consumption given cash on hand X_{s+1} , and earnings and health outcomes, (W_{s+1}^i) and (M_{s+1}^j) , respectively. Cash on hand in period $s+1$ is obviously a function of the choice of C_s . We try a range of consumption values C_s from \bar{C} to X_s . At each possible choice of C_s , we calculate the left-hand side of the Euler equation above. We also calculate the value function:

$$U(C_s) + \frac{p(\text{live}_{s+1})}{1+\delta} \sum_{i=1}^n \sum_{j=1}^n p(W_{s+1}^i)p(M_{s+1}^j)V(X_{s+1}|W_{s+1}^i, M_{s+1}^j). \quad (B.2)$$

At the optimum, the left-hand side of (B.1) should equal zero (unless the borrowing constraint binds, in which case $C_s = X_s$), and the expression in (B.2) should be maximized. The search process is made difficult, however, by the presence of the means-tested consumption floor \bar{C} . As we show in Hubbard, Skinner, and Zeldes (1993), such a consumption floor introduces nonconvexities in the intertemporal budget constraint. These nonconvexities can cause multiple values of consumption that satisfy the Euler equation above, only one of which corresponds to the global optimum. Therefore, for each value of C_s that we try, we evaluate the left-hand side of (B.1) and the expression in (B.2), and we make sure that, at the optimal consumption choice, (B.1) is satisfied (or else $C_s = X_s$) and that (B.2) is maximized. This complicates the solution technique considerably.

For each possible choice of C_s , we must calculate the expected marginal utility and the expected value function in period $s+1$. This is the innermost loop in our computer program and uses the bulk of the computer time. We do the calculation as follows. Given a possible choice of consumption, we first calculate the exact level of next-period cash on hand for a combination of realizations of earnings and medical expenses.⁵⁴ We then use the previously computed optimal consumption function to calculate the optimal next-period consumption, marginal utility, and value function. To calculate marginal utility, we use a linear interpolation of the consumption function between

⁵⁴For periods after retirement, income is nonstochastic, and the problem is somewhat simpler.

the two nearest cash-on-hand nodes, and then evaluate the marginal utility function at this consumption level. To calculate the expected value of the value function, we use a log-linear interpolation of the previously computed value function between the two nearest cash-on-hand nodes.^{55,56} To calculate the expected value of next period's marginal utility, we sum the product of the probability of each realization of earnings and medical expenses and the marginal utility of next period's consumption given that realization. To calculate the expected value of next period's value function, we sum the product of the probability of each realization of earnings and medical expenses and the value function given that realization.

Each calculation of optimal consumption involves searching over a large number of consumption choices. We first evaluate expressions (B.1) and (B.2) for a grid of 30 consumption choices equally spaced between \bar{C} and X_s . The left-hand side of (B.1) equals (except for approximation error) the slope of (B.2). We find local maxima to (B.2) as follows. If the left-hand side of (B.1) is positive at one point on the consumption grid and negative at the next higher one, or if the left-hand side of (B.1) is positive at both but (B.2) is lower at the second than at the first, or if the left-hand side of (B.1) is negative at both nodes but (B.2) is higher at the second than at the first, we use a hill-climbing technique to find the local maximum. Thus, we do not discretize the choice variable (consumption) but rather allow it to take on any value. This is likely to improve the accuracy of the approximation. After finding all local maxima, we compare (B.2) at each of these and at the corner solution $C_s = X_s$, and choose the global maximum. The corresponding level of consumption and value function are stored for use in the previous period.

Constructing aggregate consumption, earnings, and assets. We begin by calculating sample means for each age for consumption, earnings, and assets:

$$\bar{C}_{k,s} = \sum_{i=1}^{N_{k,s}} C_{k,i,s}/N_{k,s}, \bar{W}_{k,t} = \sum_{i=1}^{N_{k,s}} W_{k,i,s}/N_{k,s}, \text{ and } \bar{A}_{k,s} = \sum_{i=1}^{N_{k,s}} A_{k,i,s}/N_{k,s},$$

where s denotes adult age (equal to actual age minus 20), k denotes the lifetime-income group, i denotes individuals, and $N_{k,s}$ denotes the number of individuals in group k of age s . These sample means make it much easier to calculate total assets within the overlapping-generations framework of the life-cycle model. To accomplish this, we must sum over the different cohorts (of different lifetime-income groups) alive at a given point in time, taking into account mortality rates, population growth, and productivity growth. For

⁵⁵Recall that the highest node for cash on hand is larger than any feasible realization for cash on hand. This was done to be able to use interpolation and avoid ever using extrapolation, thereby improving the accuracy of the approximations.

⁵⁶We wrote the interpolation routine so that it would vectorize, in order to take advantage of the vector features of the Cornell Supercomputer.

each time t , we can aggregate consumption, earnings, and assets as follows:

$$C_{k,t}^* = \sum_{s=1}^T \left[\sum_{i=1}^{N_{k,s}} C_{k,i,s} \right], \quad (B.3a)$$

$$W_{k,t}^* = \sum_{s=1}^T \left[\sum_{i=1}^{N_{k,s}} W_{k,i,s} \right], \quad (B.3b)$$

and

$$A_{k,t}^* = \sum_{s=1}^T \left[\sum_{i=1}^{N_{k,s}} A_{k,i,s} \right], \quad (B.3c)$$

where $C_{k,t}^*$, $W_{k,t}^*$, and $A_{k,t}^*$ represent aggregate consumption, earnings, and assets, respectively, for lifetime-income group k ; $N_{k,s}$ represents the number of individuals in cohort s of lifetime-income group k ; and T now equals the maximum adult age of 80 (corresponding to an actual age of 100).

Note that we can simplify these expressions as follows:

$$C_{k,t}^* = \sum_{s=1}^T N_{k,s} \left[\sum_{i=1}^{N_{k,s}} C_{k,i,s} / N_{k,s} \right] = \sum_{s=1}^T N_{k,s} \bar{C}_{k,s}, \quad (B.4a)$$

$$W_{k,t}^* = \sum_{s=1}^T N_{k,s} \bar{W}_{k,s}, \quad (B.4b)$$

$$A_{k,t}^* = \sum_{s=1}^T N_{k,s} \bar{A}_{k,s}, \quad (B.4c)$$

where $\bar{C}_{k,s}$, $\bar{W}_{k,s}$, and $\bar{A}_{k,s}$ are the cohort means defined previously.

Further simplifying, let the population of the youngest cohort in the lowest lifetime-income groups be normalized to unity. Then, if n is the population growth rate and p_s is the probability of surviving to adult age s conditional on having survived to adult age 1 (both assumed to be the same for all lifetime-income groups),

$$N_{k,s} = p_s (1+n)^{1-s} \gamma_k, \quad (B.5)$$

where γ_k (equal to unity for the lowest lifetime-income group) represents the ratio of the population of lifetime-income group k to that of the lowest lifetime-income group. Hence, we can express aggregate consumption, earnings, and assets for each lifetime-income group k as:

$$C_{k,t}^* = \sum_{s=1}^T p_s (1+n)^{1-s} \gamma_k \bar{C}_{k,s}, \quad (B.6a)$$

$$W_{k,t}^* = \sum_{s=1}^T p_s (1+n)^{1-s} \gamma_k \bar{W}_{k,s}, \quad (B.6b)$$

$$A_{k,t}^* = \sum_{s=1}^T p_s (1+n)^{1-s} \gamma_k \bar{A}_{k,s}, \quad (B.6c)$$

where, again, $\bar{C}_{k,s}$, $\bar{W}_{k,s}$, and $\bar{A}_{k,s}$ are the cohort means as defined above.

Total consumption, earnings, and assets in the economy are weighted sums of the education-group-specific aggregates and are given by:

$$C_t^* = \sum_k C_{k,t}^* \quad (B.7a)$$

$$W_t^* = \sum_k W_{k,t}^* \quad (B.7b)$$

$$A_t^* = \sum_k A_{k,t}^*. \quad (B.7c)$$

Note the aggregate assets equal the aggregate capital stock:

$$A_t^* = K_t^*.$$

Finally, let aggregate saving be defined as:

$$S_t^* = \sum_{s=1}^T p_s (1+n)^{1-s} \gamma_k \bar{S}_{k,s}, \quad (B.8)$$

where:

$$\bar{S}_{k,s} = \sum_{i=1}^{N_{k,s}} [rA_{k,i,s-1} + W_{k,i,s} + TR_{k,i,s} - M_{k,i,s} - C_{k,i,s}] / N_{k,s}; \quad (B.9)$$

that is, $\bar{S}_{k,s}$ is the average saving over the period from age $s-1$ to s of those households surviving to age s . When lifespan is certain, aggregate steady-state saving can be expressed as:⁵⁷ $S_t^* = (n/(1+n))A_t^*$. However, when lifespan is uncertain, calculated aggregate saving will be larger, because (accidental) bequests are counted in a given year's saving but not in the subsequent capital stock. To see this, note that individual saving can be written as the change in assets,

$$S_{k,i,s} = A_{k,i,s} - A_{k,i,s-1}.$$

⁵⁷This follows from the definition of saving as the change in the aggregate capital stock: $S_t^* = A_t^* - A_{t-1}^*(1+n)^{-1}$. In steady state, $A_t^* = A_{t-1}^*$, leading to the expression in the text.

Rewriting (B.8),

$$S_t^* = \sum_{s=1}^T p_s (1+n)^{1-s} \gamma_k [\bar{A}_{k,s} - \bar{A}_{k,s-1}]. \quad (B.10)$$

Rearranging and noting that $\bar{A}_{k,0} = \bar{A}_{k,T} = 0$,

$$S_t^* = \sum_{s=1}^T \bar{A}_{k,s} p_s (1+n)^{-s} \gamma_k [n + (1 - \frac{p_{s+1}}{p_s})], \quad (B.11)$$

where $[1 - (p_{s+1}/p_s)]$ is the conditional probability of dying in period $s + 1$. Note that this expression reverts to $(n/(1+n))A_t^*$ in the case of certain lifespan. In the case of uncertain life span, S_t^* is augmented by accidental bequests, equal to (age-specific) assets times the probability of death. In our model, bequests are effectively confiscated, since no other generation receives them. An alternative approach would have younger generations face uncertain future inheritances. This more general model is a topic for future research.

References

- Abel, A.B., (1985). Precautionary Saving and Accidental Bequests. *American Economic Review*, **75**: 777–791.
- Abowd, J. and Card, D., (1989). On the Covariance Structure of Earnings and Hours Changes. *Econometrica*, **57**: 411–455.
- Aiyagari, S.R., (1992). Uninsured Idiosyncratic Risk and Aggregate Saving. Mimeograph, Research Department, Federal Reserve Bank of Minneapolis.
- Aiyagari, S.R. and Gertler, M., (1991). Asset Returns with Transactions Costs and Uninsured Individual Risk. *Journal of Monetary Economics*, **27**: 311–331.
- Ando, A. and Kennickell, A.B., (1987). How Much (or Little) Life Cycle is There in Micro Data?: The Cases of the United States and Japan. R. Dornbusch, S. Fischer, and J. Bossons, (eds.), *Macroeconomics and Finance: Essays in Honor of Franco Modigliani*. Cambridge: MIT Press.
- Ando, A. and Modigliani, F., (1963). The ‘Life Cycle’ Hypothesis of Saving: Aggregate Implications and Tests. *American Economic Review*, **53**: 55–84.
- Attanasio, O.P., (1993). Personal Saving in the U.S. Mimeograph, Stanford University.
- Auerbach, A.J. and Kotlikoff, L.J., (1987). *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press.
- Auerbach, A.J., Kotlikoff, L.J., and Skinner, J., (1983). The Efficiency Gains from Dynamic Tax Reform. *International Economic Review*, **24**: 81–100.
- Barsky, R.B., Mankiw, N.G., and Zeldes, S.P., (1986). Ricardian Consumers with Keynesian Propensities. *American Economic Review*, **76**: 676–691.
- Bernheim, B.D. and Scholz, J.K., (1993). Private Saving and Public Policy. J.M. Poterba, (ed.), *Tax Policy and the Economy*, **7**: Cambridge: MIT Press.
- Bewley, T., (1977). The Permanent Income Hypothesis: A Theoretical Formulation. *Journal of Economic Theory*, **16(2)**: 252–292.
- Blinder, A.S., (1988). Comments on Chapters 1 and 2. D. Kessler and A.

Masson, (eds.), *Modelling the Accumulation and Distribution of Wealth*. Oxford: Clarendon Press.

Blundell, R., Browning, M., and Meghir, C., (1992). Consumer Demand and the Life-Cycle Allocation of Household Expenditures. London: Institute for Fiscal Studies Working Paper No. W93/11; forthcoming, *Review of Economic Studies*.

Boskin, M.J., (1978). Taxation, Saving, and the Rate of Interest. *Journal of Political Economy*, **86**: 3–27.

Caballero, R.J., (1990). Consumption Puzzles and Precautionary Saving. *Journal of Monetary Economics*, **25**: 113–136.

Caballero, R.J., (1991). Earnings Uncertainty and Aggregate Wealth Accumulation. *American Economic Review*, **81**: 859–871.

Campbell, J.Y. and Mankiw, N.G., (1989). Consumption, Income, and Interest Rates: Reinterpreting the Time-Series Evidence. O.J. Blanchard and S. Fischer, (eds.), *NBER Macroeconomics Annual*. Cambridge: MIT Press.

Cantor, R., (1985). The Consumption Function and the Precautionary Demand for Savings. *Economic Letters*, **17**: 207–210.

Carroll, C.D., (1992). The Buffer Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, **2**: 61–135.

Carroll, C.D. and Samwick, A.A., (1992a). How Important is Precautionary Saving? Mimeograph, Economic Activity Section, Board of Governors of the Federal Reserve System.

Carroll, C.D. and Samwick, A.A., (1992b). The Nature and Magnitude of Precautionary Wealth. Working Paper No. 124, Economic Activity Section, Board of Governors of the Federal Reserve System.

Carroll, C.D. and Summers, L.H., (1991). Consumption Growth Parallels Income Growth: Some New Evidence. B.D. Bernheim and J.B. Shoven, (eds.), *National Saving and Economic Performance*. Chicago: University of Chicago Press.

Clarida, R.H., (1987). Consumption, Liquidity Constraints, and Asset Accumulation in the Presence of Random Income Fluctuations. *International*

Economic Review, **28**: 339–351.

Committee on Ways and Means, (1991). *Overview of Entitlement Programs: 1991 Green Book*. U.S. House of Representatives.

Darby, M.R., (1979). *Effects of Social Security on Income and the Capital Stock*. Washington, D.C.: American Enterprise Institute.

Dardanoni, V., (1991). Precautionary Savings Under Income Uncertainty: A Cross-Sectional Analysis. *Applied Economics*, **23**: 153–160.

Davies, J., (1981). Uncertain Lifetime, Consumption and Dissaving in Retirement. *Journal of Political Economy*, **89**: 561–578.

Deaton, A., (1991). Saving and Liquidity Constraints. *Econometrica*, **59**: 1221–1248.

Deaton, A., (1992). *Understanding Consumption*. Oxford: Oxford University Press.

Drèze, J.H. and Modigliani, F., (1972). Consumption Decisions Under Uncertainty. *Journal of Economic Theory*, **5**: 308–335.

Elmendorf, D.W. and Kimball, M.S., (1991). Taxation of Labor Income and the Demand for Risky Assets. Working Paper No. 3204, National Bureau of Economic Research.

Engen, E.M., (1992a). Precautionary Saving and the Structure of Taxation. Mimeograph, U.C.L.A.

Engen, E.M., (1992b). Precautionary Saving, Consumption, and Taxation in a Life Cycle Model with Stochastic Earnings and Mortality Risk. Unpublished Ph.D. dissertation, University of Virginia.

Faber, J.F., (1982). *Life Tables for the United States: 1900–2050*. U.S. Department of Health and Human Services, Social Security Administration, Actuarial Study No. 87.

Feenberg, D. and Skinner, J., (1992). The Risk and Duration of Catastrophic Health Care Expenditures. Working Paper No. 4147, National Bureau of Economic Research.

Feldstein, M., (1974). Social Security, Induced Retirement, and Aggregate Capital Accumulation. *Journal of Political Economy*, **82**: 905–926.

Feldstein, M., (1981). *Hospital Costs and Health Insurance*. Cambridge: Harvard University Press.

Flavin, M.A., (1981). The Adjustment of Consumption to Changing Expectations About Future Income. *Journal of Political Economy*, **89**: 974–1002.

Friedman, B.M. and Warshawsky, M., (1988). Annuity Yields and Saving in the United States. Z. Bodie, J.B. Shoven, and D.A. Wise, (eds.), *Pensions in the U.S. Economy*. Chicago: University of Chicago Press.

Ghez, G. and Becker, G.S., (1975). *The Allocation of Time and Goods Over the Life Cycle*. New York: Columbia University Press.

Guiso, L., Jappelli, T., and Terlizzese, D., (1992). Earnings Uncertainty and Precautionary Saving. *Journal of Monetary Economics*, **30**: 307–338.

Hahm, J.-H., (1993). Consumption Growth and Earnings Uncertainty: Simple Cross-Country Evidence. Mimeograph, Columbia University.

Hall, R.E., (1978). Stochastic Implications of the Life Cycle - Permanent Income Hypothesis: Theory and Evidence, *Journal of Political Economy*, **86**: 971–987.

Hall, R.E. and Mishkin, F.S., (1982). The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. *Econometrica*, **50**: 461–481.

Hansen, G.D. and Imrohoroglu, A., (1992). The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard. *Journal of Political Economy*, **100**: 118–142.

Heaton, J. and Lucas, D., (1992). Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing. Mimeograph, M.I.T.

Health Insurance Association of America, (1988). The Consumer's Guide to Long-Term Care Insurance. U.S. Senate, Hearings Before the Subcommittee on Health, May 27/June 17: 139.

Hubbard, R.G., (1987). Uncertain Lifetime, Pensions, and Individual Saving.

Z. Bodie, J.B. Shoven, and D.A. Wise, (eds.), *Issues in Pension Economics*. Chicago: University of Chicago Press.

Hubbard, R.G. and Judd, K.L., (1986). Liquidity Constraints, Fiscal Policy, and Consumption. *Brookings Papers on Economic Activity*, 1: 1–50.

Hubbard, R.B. and Judd, K.L., (1987). Social Security and Individual Welfare: Precautionary Saving, Liquidity Constraints, and the Payroll Tax. *American Economic Review*, 77: 630–646.

Hubbard, R.G., Skinner, J., and Zeldes, S.P., (1993). Precautionary Saving and Social Insurance. Mimeograph, Columbia University, University of Virginia, and University of Pennsylvania.

Hurd, M.D., (1989). Mortality Risk and Bequests. *Econometrica*, 57: 779–814.

Hurd, M.D. and Mundaca, B.G., (1989). The Importance of Gifts and Inheritances Among the Affluent. R.E. Lipsey and H.S. Tice, (eds.), *The Measurement of Saving, Investment, and Wealth*. Chicago: University of Chicago Press.

Keane, M. and Moffitt, R., (1991). A Structural Model of Multiple Welfare Program Participation and Labor Supply. Mimeograph, Brown University.

Kessler, D. and Masson, A., (1988). Introduction. D. Kessler and A. Masson, (eds.), *Modelling the Accumulation and Distribution of Wealth*. Oxford: Clarendon Press.

Kessler, D. and Masson, A., (1989). Bequest and Wealth Accumulation: Are Some Pieces of the Puzzle Missing? *Journal of Economic Perspectives*, 3: 141–152.

Kimball, M.S., (1990). Precautionary Saving in the Small and in the Large. *Econometrica*, 58: 53–73.

Kimball, M.S., (1991). Precautionary Motives for Holding Assets. J. Eatwell, M. Milgate, and P. Newman, (eds.), *The New Palgrave Dictionary of Money and Finance*. New York: W.W. Norton.

Kimball, M.S., (1993). Standard Risk Aversion. *Econometrica*, 61: 589–612.

Kimball, M.S. and Mankiw, N.G., (1989). Precautionary Saving and the Timing of Taxes. *Journal of Political Economy*, **97**: 863–879.

Kitamura, Y. and Takayama, N., (1993). International Comparisons of Household Saving: Household Data on Saving Paper in Japan. Paper presented at NBER Conference on International Saving Comparisons.

Kotlikoff, L.J., (1988a). Health Expenditures and Precautionary Savings. L.J. Kotlikoff, (ed.), *What Determines Saving?* Cambridge: MIT Press.

Kotlikoff, L.J., (1988b). Intergenerational Transfers and Savings. *Journal of Economic Perspectives*, **2**: 41–58.

Kotlikoff, L.J. and Summers, L.H., (1981). The Role of Intergenerational Transfers in Aggregate Capital Accumulation. *Journal of Political Economy*, **89**: 706–732.

Kotlikoff, L.J. and Summers, L.H., (1988). The Contribution of Intergenerational Transfers to Total Wealth: A Reply. D. Kessler and A. Masson, (eds.), *Modelling the Accumulation and Distribution of Wealth*. Oxford: Oxford University Press.

Leland, H.E., (1968). Saving and Uncertainty: The Precautionary Demand for Saving. *Quarterly Journal of Economics*, **82**: 465–473.

Levhari, D., Mirman, L.J., and Zilcha, I., (1980). Capital Accumulation Under Uncertainty. *International Economic Review*, **21**: 661–671.

Lillard, L.A. and Willis, R.J., (1978). Dynamic Aspect of Earnings Ability. *Econometrica*, **46**: 985–1012.

Lusardi, A., (1993). Euler Equations in Micro Data: Merging Data from Two Samples. Mimeograph, Dartmouth College.

MaCurdy, T.E., (1982). The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis. *Journal of Econometrics*, **18**: 83–114.

Miller, B.L., (1976). The Effect on Optimal Consumption of Increased Uncertainty in Labor Income in the Multiperiod Case. *Journal of Economic Theory*, **13**: 154–167.

Modigliani, F., (1988a). Measuring the Contribution of Intergenerational Transfers to Total Wealth: Conceptual Issues and Empirical Findings. D. Kessler and A. Masson, (eds.), *Modelling the Accumulation and Distribution of Wealth*. Oxford: Clarendon Press.

Modigliani, F., (1988b). The Role of Intergenerational Transfers and Life Cycle Saving in the Accumulation of Wealth. *Journal of Economic Perspectives*, **2**: 15–40.

Modigliani, F. and Brumberg, R., (1954). Utility Analysis and Aggregate Consumption Functions: An Interpretation of Cross-Section Data. K.K. Kurihara, (ed.), *Post-Keynesian Economics*, New Brunswick, New Jersey: Rutgers University Press.

Murphy, K.M. and Welch, F., (1990). Empirical Age-Earnings Profiles. *Journal of Labor Economics*, **8**: 202–229.

Rossiter, L.F. and Wilensky, G.R., (1982). Out-of-Pocket Expenditures for Personal Health Services, 1982. National Health Care Expenditures Study, DHHS Publication No. (PHS) 82–3332, U.S. Department of Health and Human Services.

Rothschild, M. and Stiglitz, J.E., (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *Quarterly Journal of Economics*, **90**: 629–650.

Sandmo, A., (1970). The Effect of Uncertainty on Saving Decisions. *Review of Economic Studies*, **37**: 353–360.

Schechtman, J. and Escudero, V.L.S., (1977). Some Results on ‘An Income Fluctuation Problem.’ *Journal of Economic Theory*, **16**: 151–166.

Sibley, D.S., (1975). Permanent and Transitory Income Effects in a Model of Optimal Consumption with Wage Income Uncertainty. *Journal of Economic Theory*, **11**: 68–82.

Skinner, J., (1985). Variable Lifespan and the Intertemporal Elasticity of Consumption. *Review of Economics and Statistics*, **67**: 616–623.

Skinner, J., (1988). Risky Income, Life Cycle Consumption, and Precautionary Savings. *Journal of Monetary Economics*, **22**: 237–255.

Summers, L.H., (1981). Capital Taxation and Accumulation in a Life Cycle Growth Model. *American Economic Review*, **71**: 533-544.

Tobin, J., (1967). Life Cycle Saving and Balanced Growth. W. Fellner, *et al.* (eds.), *Ten Economic Studies in the Tradition of Irving Fisher*. New York: Wiley.

Tobin, J. and Dolde, W., (1971). Wealth, Liquidity, and Consumption. Federal Reserve Bank of Boston, *Consumer Spending and Monetary Policy: The Linkages*. Federal Reserve Bank of Boston Conference Series No. 5: 99-146.

U.S. House of Representatives Select Committee on Aging, (1990). *Emptying the Elderly's Pocketbook: Impact on Rising Health Care Costs*. Committee Publication No. 101-746, Washington: U.S. Government Printing Office.

Venti, S.F. and Wise, D.A., (1987). IRAs and Saving. M. Feldstein, (ed.), *The Effects of Taxation on Capital Accumulation*. Chicago: University of Chicago Press.

White, B.B., (1978). Empirical Tests of the Life-Cycle Hypothesis. *American Economic Review*, **68**: 546-560.

Wilson, C., (1977). A Model of Insurance Markets with Incomplete Information. *Journal of Economic Theory*, **16**: 167-207.

Yaari, M.F., (1965). Uncertain Lifetime, Life Insurance, and the Theory of the Consumer. *Review of Economic Studies*, **32**: 137-158.

Zeldes, S.P., (1989a). Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*, **97**: 305-346.

Zeldes, S.P., (1989b). Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence. *Quarterly Journal of Economics*, **104**: 275-298.