

Volume 84
Number 2
May 1994

WHY DO PEOPLE SAVE?

Expanding the Life-Cycle Model:
Precautionary Saving and Public Policy

By R. GLENN HUBBARD, JONATHAN SKINNER, AND STEPHEN P. ZELDES

Reprinted from

THE AMERICAN ECONOMIC REVIEW

© The American Economic Association

WHY DO PEOPLE SAVE?†

Expanding the Life-Cycle Model: Precautionary Saving and Public Policy

By R. GLENN HUBBARD, JONATHAN SKINNER, AND STEPHEN P. ZELDES*

One of the key puzzles in understanding saving behavior is not so much why people save—the title of this session—but why people *don't* save. According to the familiar life-cycle model, households should accumulate wealth to provide for their retirement consumption. The surprising result from the data is the sizable fraction of the population who have accumulated so little, even among those nearing retirement. Given that earnings will almost surely decline when households retire, such behavior can imply poor living standards for the elderly.

Some might interpret the low wealth accumulation as being evidence of myopia, irrationality, or a failure of households to enforce “mental accounting” (Richard Thaler, 1994), while others might view the low level of wealth accumulation as evidence of high individual rates of time preference. Determining the underlying causes of low wealth is crucial for public policies that seek to alleviate low aggregate saving rates in the United States, as well as policies that seek to buttress the adequacy of financial resources for the elderly. In this paper, we outline what we believe to be the causes of why many people do not save. We

conclude by speculating about government policies that may be most effective at encouraging saving.

Much of the research examining levels of consumption, saving, and wealth, as well as their responsiveness to policy, has been done using a life-cycle model with the simplifying assumption of perfect certainty. Alan Auerbach and Laurence Kotlikoff (1987), for example, developed a model with 55 overlapping generations of individual life-cycle “households,” each with empirically plausible age-earnings profiles and utility parameters, and used the model to address tax policy and demographic issues in a regime in which all households are identical within a generation, and all generations know future earnings and interest rates.

More recently, a line of inquiry has examined the effects of uncertainty on saving, generally in the context of highly stylized models. This research has shown that, in these models, uninsured earnings uncertainty can alter optimal saving behavior in a variety of important ways. (For a partial review of this “precautionary saving” literature, see Hubbard et al. [1994].)

In two recent papers (Hubbard et al., 1993, 1994), we have combined these two strands of the literature by examining the implications of a life-cycle model of consumption, saving, and wealth accumulation subject to what we think are the three most important sources of uninsured idiosyncratic risk facing households: uncertainty about earnings, medical expenses, and length of life. Our intent has been to create a realistic model in which families live for many periods, working for part of their lives and retiring later in life. To parameterize the uncertainty facing families, we estimate

†*Discussants:* Christopher Carroll, Federal Reserve Board; John B. Shoven, Stanford University; Laurence Kotlikoff, Boston University.

*Columbia University, New York, NY 10027, University of Virginia, Charlottesville, VA 22901, and The Wharton School, University of Pennsylvania, Philadelphia, PA 19104-6367, respectively. This research has been supported by the Harry and Lynde Bradley Foundation and the National Science Foundation. Major computational work was conducted using the Cornell National Supercomputer Facility.

the stochastic processes using available cross-section and panel data sets on households. In addition, we include asset-based means-tested public welfare programs (e.g., Aid to Families with Dependent Children, Food Stamps, Supplemental Security Income, and Medicaid). In this research program, we are looking at how close a model with rational optimizing consumers can come to matching a wide range of empirical facts about consumption and wealth.¹

The improvements to the life-cycle model under certainty or perfect markets come in three forms. First, adding uncertainty, even in cases where it does not change the optimal decision rule for saving, introduces heterogeneity in wealth and saving when households differ in the realizations of serially correlated earnings and health expenses. Second, we assume household preferences that do not generate "certainty equivalence," so that the introduction of uncertainty changes household decision rules in important ways. Third, many households face a high implicit tax on saving should they become eligible for AFDC, Medicaid, or Food Stamps. These social insurance programs with asset-based means-testing alter in significant ways the incentives to save.

In Hubbard et al. (1993), we show that the combination of these factors can explain much of the observed cross-sectional heterogeneity in wealth holdings in the 1984 Panel Study of Income Dynamics (PSID), including the presence of a substantial number of households with low levels of assets at retirement. In this paper, we follow up on these households; we use data from the 1989 PSID to examine transitions in wealth holdings between 1984 and 1989. Our model

implies that in many cases saving rates will be low among those with low initial levels of wealth (i.e., that low wealth may be an "absorbing state" over a lengthy period of time). We find that this is consistent with evidence from the PSID. More broadly, the results of our 1993 and 1994 papers indicate that our model can explain the saving behavior of "savers" and "nonsavers" alike.

I. Life-Cycle Models of Saving Decisions

We begin by presenting a simple description of our general multiperiod model with multiple sources of uncertainty (readers interested in a more comprehensive description should see our 1994 paper). Consumers maximize expected lifetime utility, given all of the relevant constraints. At each age t , the consumption chosen maximizes

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

subject to the transition equation

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

plus the additional constraints that

$$(3) \quad A_s \geq 0 \quad \forall s.$$

The first expression describes the optimization in which consumption excluding medical expenses, C_s , is selected to maximize expected lifetime utility (where E_t is the expectations operator conditional on information at time t), discounted based on a rate of time preference δ . To allow for a random date of death, D_s is a state variable equal to unity if the individual is alive and zero otherwise, and T is the maximum possible length of life. The household begins period s with assets from the previous period plus accumulated interest, $A_{s-1}(1 + r)$, where r is the nonstochastic real after-tax rate of return. It then receives exogenous

¹In Hubbard et al. (1994), we show that, using realistic parameter values, our model replicates empirical regularities in (i) aggregate wealth and the aggregate saving rate, (ii) cross-sectional differences in consumption-age profiles by lifetime-income group, and (iii) short-run time-series properties of consumption and income. In Hubbard et al. (1993), we analyze the effects of a social insurance program with asset-based means testing and show that our model helps explain the observed cross-sectional distribution of wealth.

earnings E_s , pays out necessary medical expenses M_s , and receives government transfers TR_s . It is left with

$$(4) \quad X_s = A_{s-1}(1+r) + E_s - M_s + TR_s$$

which, following Angus Deaton (1991), we denote as "cash on hand."

Given X_s , consumption is chosen, and what remains equals end-of-period assets, A_s . We assume that no utility is derived from medical expenditures, as these costs only offset the damage inflicted by ill health. The borrowing and terminal constraints in equation (3) prevent negative assets in any period. Since the setup includes a government-guaranteed level of consumption, borrowing constraints rule out borrowing in one period, defaulting, and receiving the guaranteed consumption in the subsequent period.

For simplicity, we specify a generic transfers function including income-based and asset-based means testing, as well as payments tied to medical expenses:

$$(5) \quad TR_s = \max\left[0, (\bar{C} + M_s) - (A_{s-1}(1+r) + E_s)\right].$$

The minimum level of consumption guaranteed by the government is \bar{C} (the "consumption floor"). Transfers equal this consumption floor plus medical expenses minus all available resources (if that amount is positive, and zero otherwise). Simply put, transfer payments guarantee a minimum standard of living after medical expenses equal to \bar{C} . The one-for-one reduction in transfer payments in response to increases in assets or current earnings captures in a stylized fashion the penalty on saving of asset-based means-tested programs. Because eligibility is conditional on having assets less than some specified amount, such programs place an implicit tax rate of 100 percent on assets above the limit. We discuss these incentives in greater detail in our 1993 paper.

Given our interest in individual as well as aggregate saving, we allow for heterogeneity in life-cycle saving decisions across different lifetime-income groups. Using educational attainment as a proxy for lifetime income, we specify three groups: households whose head does not have a high-school degree, households whose head is a high-school graduate, and households whose head has a college degree.

Solving the model requires a functional form to describe household utility, an empirical characterization of the sources of uncertainty, a description of the consumption floor, and assumptions regarding the rate of time preference and real interest rate. We assume that the period utility function is isoelastic, and we experiment with alternative values for both the coefficient of relative risk aversion and the rate of time preference. Mortality probabilities are taken from official data. We estimate functions for earnings (and Social Security and private pension receipts) and out-of-pocket medical expenses using micro data. Residuals from log-earnings and log-medical-expenses regressions are used to estimate education-group-specific AR(1) processes to account for exposure to risk of fluctuations in earnings and medical expenses. We estimate a consumption floor of \$7,000 in 1984 dollars. We solve numerically the dynamic-programming problem; this yields the optimal state-contingent consumption function.

In what follows, we will compare our results to those from an alternative approach in the same spirit that also examines the effects of uncertainty on optimal intertemporal consumption decisions. Deaton (1991) and Christopher Carroll (1992) reconcile the life-cycle model's predictions with the empirical finding of low levels of wealth accumulation by many households by assuming a high rate of time preference (relative to the real interest rate). In this case, absent uncertainty, households would like to borrow against future income. With earnings uncertainty (and in some cases borrowing constraints), households maintain a "buffer stock" or contingency fund against income downturn, but households' impatience keeps buffer stocks small.

II. Why Don't People Save?

A central prediction of the life-cycle framework is that households accumulate substantial assets for retirement. An important empirical puzzle, however, is that many households have very low levels of accumulated wealth over their life cycle. In Hubbard et al. (1993), we use 1984 data from the PSID to show that a significant fraction of households with low lifetime earnings (represented by educational attainment) have preretirement wealth accumulation much too small to be consistent with the perfect-markets version of the life-cycle model. However, households with higher lifetime earnings (as proxied by a college degree) exhibit saving behavior more consistent with the predictions of the life-cycle model in the sense that most households have substantial assets near retirement.

In this section, we extend these results to consider what happened to these "low wealth" households in 1984 over the subsequent five years. The empirical evidence comes from the 1989 wave of the PSID, which asked detailed questions about wealth holdings similar to those in the 1984 survey. The simulated evidence comes from drawing random shocks from the distributions of earnings and out-of-pocket medical expenses, which, together with the optimal consumption function, are used to generate artificial "life histories" of consumption and wealth for a large number of simulated households. To contrast the predictions of our generalized life-cycle model with the "buffer-stock" alternative, we also generate simulated consumption, wealth, and earnings panels using the assumption of a high time-preference rate (10 percent) and a very low consumption floor.

We analyze movements of household assets between 1984 and 1989 for households headed by an individual aged 65 years or younger in 1984. We grouped assets into four classes according to whether the level of assets A is such that $A \leq \$1,000$; $\$1,000 < A \leq \$5,000$; $\$5,000 < A \leq \$25,000$; or $A > \$25,000$ (all figures are in 1984 dollars). The top entry in each cell in Table 1 reports the proportion of households with a

TABLE 1—TRANSITIONS BETWEEN ASSET CATEGORIES OVER A FIVE-YEAR PERIOD, ACTUAL AND SIMULATED DATA

Initial assets	Final assets			
	< \$1,000	\$1,000–5,000	\$5,000–25,000	> \$25,000
< \$1,000	56.16	15.38	18.83	9.63
	60.23	20.15	16.43	3.19
	<i>6.76</i>	<i>33.71</i>	<i>51.09</i>	<i>8.44</i>
\$1,000–5,000	24.80	22.35	35.77	17.08
	15.11	20.24	48.39	16.26
	<i>0.90</i>	<i>15.18</i>	<i>71.34</i>	<i>12.58</i>
\$5,000–25,000	10.98	12.70	38.38	37.95
	4.62	5.20	32.45	57.73
	<i>0.14</i>	<i>0.81</i>	<i>35.63</i>	<i>63.43</i>
> \$25,000	1.19	1.27	6.06	91.48
	0.16	0.07	1.47	98.30
	<i>0.09</i>	<i>0.15</i>	<i>1.24</i>	<i>98.52</i>

Notes: The numbers reported above are the percentage of those households with given initial assets ending up after five years in the designated final asset bracket (each row, therefore, sums to 100). The top entry, in bold type, is calculated from the Panel Study of Income Dynamics using the 1984 and 1989 wealth supplements. The middle entry is drawn from our simulation model with uncertainty about lifespan, earnings, and out-of-pocket medical expenses, and a consumption floor of \$7,000. The bottom entry, in italics, comes from simulations of our uncertainty model with a rate of time preference equal to 10 percent and a \$1 consumption floor. All figures are in 1984 dollars.

given level of assets in 1984 with given levels of assets in 1989. These transition calculations reveal that, even over a five-year period, there is substantial persistence of low wealth in households. For example, of those with less than \$1,000 total wealth in 1984, 56 percent still had less than \$1,000 total wealth in 1989.

The second entry in each cell of Table 1 is drawn from simulations of our model with uncertainty over lifespan, earnings, and out-of-pocket medical expenses under the assumptions that the annual rate of time preference and the interest rate are both 3 percent and the consumption floor is \$7,000. The transition matrix from the simulated data matches the actual data quite closely. For example, 60 percent of simulated households (compared to 56 percent

of actual households) who initially hold less than \$1,000 in assets still have less than \$1,000 in assets five years later. The simulated transition probabilities at other levels of wealth also match the actual data reasonably well.

There are two reasons why our model implies a high degree of persistence in low wealth. First, the simulated earnings and medical expenses exhibit considerable persistence (based on empirical parameters), so that a significant adverse shock to either of these two variables signals an adverse realization in the future and generates continued low consumption and low wealth. Preliminary results using a certainty-equivalent life-cycle model (i.e., one in which the household recalculates optimal consumption each period given the shocks to earnings and medical expenses but makes decisions as if all future variables were known with certainty) suggests that, even near retirement, roughly one-tenth of the population may exhibit low levels of wealth because of bad earnings or health shocks.

The second reason for the strong persistence in low levels of wealth is the existence of asset-based means testing of welfare programs. Low wealth holdings are likely to be an "absorbing state," in the sense that low levels of wealth increase the likelihood of receiving the consumption floor; saving while receiving transfers is discouraged by the implicit tax associated with means testing. By contrast, households with higher levels of wealth are less likely to qualify for social insurance programs, so that saving is less discouraged for this group.

The third entry in each cell of Table 1 is drawn from simulations of a "buffer-stock" model with a time preference rate of 10 percent, interest rate equal to 3 percent, and a negligible consumption floor of just \$1. The buffer-stock model predicts much less persistence of low levels of wealth than experienced by households in the PSID. The intuition is that, in the buffer-stock model, wealth accumulation serves primarily to insulate consumption against a bad draw in disposable income. Should a household experience low wealth because of an unusually bad draw, the household is careful to

rebuild the buffer stock. Under the buffer-stock model, there are fewer than 7 percent of households with initial wealth less than \$1,000 in "1984" who still have low wealth five years later.

III. Conclusions and Implications for Public Policies Toward Saving

The perfect-markets version of the life-cycle model has directly or indirectly contributed much of our intuition about effects of public policies on household saving decisions. While the findings of our research project cast doubt on the applicability of the standard, perfect-certainty version of the life-cycle model for policy analysis, we conclude that a well-specified optimizing life-cycle model with uninsured idiosyncratic risks and social insurance can explain many empirical observations, including the saving behavior of much of the population. It is therefore likely to be a useful tool for analyzing the effects of government policy on saving behavior. We briefly consider two examples below.

In the traditional perfect-markets version of the life-cycle model, the taxation of capital income has a significant impact on aggregate saving. The policy implications stemming from such models—that consumption taxes are more efficient than income taxes, for example—are largely predicted on the high interest elasticity of saving in the basic model. Recent research by others (see e.g., Eric Engen, 1993) has shown that, in a setting with uninsurable idiosyncratic risk, the after-tax rate of return exerts a smaller impact on saving. While this is a topic for future research, taking seriously uncertainty and imperfect insurance and lending markets is likely to yield tax-policy implications quite different from conventional life-cycle models.

The failure of perfect-markets versions of the life-cycle model to explain the negligible wealth accumulation by many households, particularly those with low lifetime incomes, has stimulated interest in models of saving that can account for such behavior. Our approach has addressed this question and has emphasized that expenditure policy,

such as the design of social insurance programs, may exert as large an effect on saving behavior as tax policy. The asset-based means testing of AFDC, Supplemental Security Income, and Medicaid may ultimately affect the consumption and saving choices of low-income households by more than even extreme variations in explicit marginal income tax rates. Raising the asset limit for AFDC and food stamps from below \$3,000 to \$10,000, a proposal put forth by Jack Kemp, the former Secretary of Housing and Urban Development, could increase saving among those least likely to save for retirement. Whether the increased saving justifies the revenue cost of the enhanced program is, of course, a question for future research.

REFERENCES

- Auerbach, Alan J. and Kotlikoff, Laurence J. *Dynamic fiscal policy*. Cambridge: Cambridge University Press, 1987.
- Carroll, Christopher D. "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity*, 1992, (2), pp. 61-135.
- Deaton, Angus. "Saving and Liquidity Constraints." *Econometrica*, September 1991, 59(5), pp. 1221-48.
- Engen, Eric M. "Precautionary Saving and the Structure of Taxation." Mimeo, University of California-Los Angeles, 1993.
- Hubbard, R. Glenn; Skinner, Jonathan and Zeldes, Stephen P. "Precautionary Saving and Social Insurance." Mimeo, University of Virginia, 1993.
- _____. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy*, 1994 (forthcoming).
- Thaler, Richard H. "Psychology and Savings Policies." *American Economic Review*, May 1994 (*Papers and Proceedings*), 84(2), pp. 186-92.