NBER WORKING PAPER SERIES

PRECAUTIONARY SAVING AND SOCIAL INSURANCE

R. Glenn Hubbard
Jonathan Skinner
Stephen P. Zeldes

Working Paper No. 4884

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1994

We are very grateful to Eric Engen, Robert Reider, and Charlie Ye for seemingly endless computer programming, and to Keith Dallara, Leslie Jeng, David Kerko, Vasan Kesavan, Richard McGrath, and Dan Zabinski for research assistance. We have benefitted from comments by Henry Aaron, Anne Alstott, Douglas Bernheim, Michael Cragg, Eric Engen, Martin Feldstein, Edward Green, Kevin Hassett, Janet Holtzblatt, Anil Kashyap, John Laitner, James Mirrlees, Robert Moffitt, James Poterba, Elizabeth Powers, Edward Prescott, Robert Shiller, Harrison White, David Wilcox, and seminar participants at Boston College, Carnegie-Mellon University, University of Chicago, Columbia University, Dartmouth College, Harvard University, Johns Hopkins University, London School of Economics, University of Maryland, University of Minnesota, NBER, New York University, Nuffield College (Oxford), Ohio State University, Princeton University, University College (London), the University of Toronto, the University of Virginia, Yale University, the Board of Governors of the Federal Reserve System, Federal Reserve Bank of New York, Federal Reserve Bank of Philadelphia, Federal Reserve Bank of St. Louis, and the Macro Lunch Group at the University of Pennsylvania. This research was supported by the Harry and Lynde Bradley Foundation and the National Science Foundation. Hubbard acknowledges financial support from the Faculty Research Fund of the Graduate School of Business of Columbia University, the Federal Reserve Bank of New York, and a grant from the John M. Olin Foundation to the Center for the Study of the Economy and the State of the University of Chicago; Skinner acknowledges financial support from the National Institute on Aging; and Zeldes is grateful for financial assistance under an Alfred P. Sloan Fellowship. Major computational work was conducted using the Cornell National Supercomputer Facility, and was funded by the Cornell Theory Center and the National Science Foundation. This paper is part of NBER's research program in Public Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4884
October 1994

PRECAUTIONARY SAVING AND SOCIAL INSURANCE

ABSTRACT

Microdata studies of household saving often find a significant group in the population with virtually no wealth, raising concerns about heterogeneity in motives for saving. In particular, this heterogeneity has been interpreted as evidence against the life-cycle model of saving. This paper argues that a life-cycle model can replicate observed patterns in household wealth accumulation after accounting explicitly for precautionary saving and asset-based means-tested social insurance. We demonstrate theoretically that social insurance programs with means tests based on assets discourage saving by households with low expected lifetime income. In addition, we evaluate the model using a dynamic programming model with four state variables. Assuming common preference parameters across lifetime-income groups, we are able to replicate the empirical pattern that low-income households are more likely than high-income households to hold virtually no wealth. Low wealth accumulation can be explained as a utility-maximizing response to asset-based means-tested welfare programs.

R. Glenn Hubbard
Graduate School of Business
Columbia University
609 Uris Hall
New York, NY 10027
and NBER

Jonathan Skinner
Department of Economics
University of Virginia
114 Rouss Hall
Charlottesville, VA 22901
and NBER

Stephen P. Zeldes
Finance Department
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104
and NBER

I. INTRODUCTION

In 1990, Grace Capetillo, a single mother receiving welfare assistance, was charged by the Milwaukee County Department of Social Services with fraud. Her crime: Her saving account balance exceeded $1000, the allowable asset limit for welfare recipients (Rose, 1990).[1] How do programs with asset restrictions, such as Aid to Families with Dependent Children (AFDC), Medicaid, Supplemental Security Income (SSI), and food stamps, affect the incentive to accumulate wealth? This paper addresses the interaction of certain social insurance programs and saving, first in simple theoretical models and later in a dynamic programming model with multiple sources of uncertainty.[2] We find that the interaction of a social insurance "safety net" with uncertainty about earnings and out-of-pocket medical expenses implies behavior that contrasts sharply with simplified models that ignore uncertainty or social insurance programs, or focus only on static incentive effects of these programs. The prospect of bad realizations in future earnings or out-of-pocket medical expenses can influence saving behavior even if the individual never actually encounters the downturn or catastrophic medical expense and never receives transfer payments. Hence the impact of social insurance programs on saving behavior extends to

---

[1] More recently, the Connecticut case of Cecilia Mercado and her daughter Sandra Rosado attracted widespread media attention. Sandra saved $4900 from part-time jobs during high school with the goal of going to college. When officials learned of the accumulated assets, they urged Sandra to spend the money quickly and ordered her mother to repay $9342 in AFDC benefits that she had received while the money was in the bank (Hays, 1992).

[2] Strictly speaking, by "social insurance" we mean welfare programs as opposed to such entitlement programs as Social Security or Medicare.

saving behavior of potential, as well as actual, recipients.

We use a model of consumption and saving subject to uncertainty to address an empirical "puzzle" of wealth accumulation: As we document below, many households accumulate little wealth over their life cycle. For those with low lifetime earnings (represented by educational attainment), wealth accumulation is inconsistent with the orthodox life-cycle model; even prior to retirement, during what are normally considered peak years of wealth holding, many families hold little wealth. By contrast, households with higher lifetime earnings exhibit saving behavior that is broadly consistent with the orthodox life cycle model, in the sense that nearly every households in this group has significant wealth accumulation near retirement.

A number of authors have examined the effects of uncertainty on optimal intertemporal consumption and saving decisions.[3] Deaton (1991) and Carroll (1992) examine the implications for wealth accumulation in these precautionary saving models, and argue that they imply too large an accumulation of wealth. They reconcile the empirical finding that most households accumulate little wealth with the predictions of the life-cycle model

---

[3] Studies of precautionary saving in response to earnings risk include Cantor (1985), Skinner (1988), Zeldes (1989), Kimball (1990a,b), and Caballero (1991), among others; for comprehensive reviews of the literature see Deaton (1992) and Hubbard, Skinner, and Zeldes (1994b). Studies of precautionary saving in response to lifetime uncertainty include Yaari (1965), Davies (1981), Skinner (1985), Abel (1986), Hubbard and Judd (1987), Hurd (1989), and Engen (1992). Kotlikoff (1988) suggested that uncertainty about medical expenses could have a large impact on precautionary saving behavior.

by assuming that the rate of time preference for most people is high relative to the real interest rate, so that in a certainty model families would prefer to borrow against future income. Earnings uncertainty (and in some cases borrowing constraints) leads individuals to maintain a "buffer stock" or contingency fund against income downturns, but the impatience keeps these buffer stocks small. This approach offers one explanation of why so many families save little throughout their life.

While the buffer stock model of wealth accumulation can explain low levels of wealth, it encounters difficulty in explaining saving behavior of those who do accumulate substantial assets. In particular, the buffer stock explanation must assume that these families have lower rates of time preference than families that do not accumulate wealth. We take an alternative approach, assuming that all individuals have the same preferences, and show that the differences in wealth of different groups can be explained by the interaction of uncertainty and social insurance programs with asset-based means testing. We develop simple analytical models to demonstrate the effects on optimal consumption of a social insurance program whose eligibility depends on current wealth -- *i.e.*, one that involves asset-based means testing. While much work has been done examining the effects on economic decisions (*e.g.*, labor supply) of *earnings-based* means tests, little has been done examining the effects of *asset-based* means testing.[4]

Under uncertainty, asset-based means tested social insurance programs depress saving for two distinct reasons. First, the provision of support in the bad states of the world reduces the uncertainty facing households, and therefore decreases precautionary saving (this effect would be present even in the absence of the asset test). Second, the restriction on asset holdings implies an implicit tax of 100 percent on wealth in the event of an earnings downturn or large medical expense. The possibility of facing this implicit tax further reduces optimal saving.

We next show that the nonlinear budget constraint implied by these programs leads to a nonmonotonic relationship between wealth and consumption over certain ranges of wealth, so that an increase in wealth can lead to a *decline* in consumption; in other words, the marginal propensity to consume (MPC) out of wealth can actually be

---

[4] Sherraden (1991) is the only analysis we could find examining asset-based means tests in welfare programs. He argues that the main effects are to reduce participating households' ability to obtain education or training or finance the purchase of a home, which limits the ability of these households to improve their social standing. He also argues that the opportunity to accumulate assets has important effects beyond the consumption that it enables, by creating an orientation towards the future and reducing the isolation of the poor from the economy and society.

There is also some recent work on other types of asset-based means testing. Feldstein (1992) has shown how college financial aid scholarship rules, which depend negatively on existing family assets, create an implicit tax on saving. He finds empirical evidence that such rules have a significant negative impact on wealth accumulation for eligible families. Another example occurs when parents expect to be supported by their children in old age; each additional dollar of wealth accumulated by parents reduces the amount of support given to them by their children (for a theoretical model, see, for example, O'Connell and Zeldes, 1993).

negative over certain ranges. This result is in sharp contrast to standard models in which consumption is always increasing in wealth.

In general, the model cannot be solved analytically, so we use the dynamic programming model developed in Hubbard, Skinner, and Zeldes (1994b) in which households face uncertainty about earnings, out-of-pocket medical expenses, and length of life.[5] We separate the population into three education (as a proxy for lifetime-income) groups and use the empirical parameters for earnings and out-of-pocket medical expenditures processes for each group estimated in Hubbard, Skinner, and Zeldes (1994b). After solving numerically for the optimal state-contingent consumption, we draw randomly from the probability distribution of uncertain health and earnings in each year and generate a time-series for several thousand simulated families.

We find that the presence of means-tested social insurance has a disproportionate impact on saving behavior of lower-lifetime-income households. For example, suppose that we denote families with total net wealth less than current income as "low wealth." Our model predicts that, for households aged 50-59, raising the minimum

---

[5] In that paper, we focus on aggregate saving rather than the distribution of wealth. We find that precautionary saving is large in a realistically parameterized life-cycle model; that is, the precautionary motive plays an important role in determining aggregate saving. We also show that our model better replicates empirical regularities in (1) aggregate wealth and the aggregate saving rate, (2) cross-sectional differences in consumption-age profiles by lifetime-income group, and (3) short-run time-series properties of consumption, income, and wealth.

government-guaranteed level of consumption (which we call the consumption "floor") from $1000 to $7000 increases the percentage of families with low wealth by 22.9 percent for low-lifetime-income households, but by only 4.4 percent among high-lifetime-income households. That is, social insurance policies designed to maintain consumption have the greatest negative effect on saving for lower income groups. This is because the guaranteed consumption floor of $7000 (identical for all education groups) represents a significantly larger fraction of lifetime income for the population with low lifetime income. We find the simulated distributions of wealth by age match in many respects the actual distributions of wealth by age documented in the United States.

The paper is organized as follows. In section II, we present empirical evidence from the Panel Study of Income Dynamics (PSID) on the distribution, by age and education, of U.S. household wealth. Section III presents simple models of consumption in the presence of social insurance and asset-based means testing. Section IV describes our multiperiod dynamic programming model and the empirical specification of the parameters of model. In section V, we present the numerical results, and the simulated age-wealth patterns are shown to mimic in certain important ways the empirical wealth patterns discussed in section II. Section VI concludes the paper.

II. THE DISTRIBUTION OF WEALTH BY AGE IN THE UNITED STATES.

In traditional life-cycle models, asset accumulation by the wealthy is essentially a scaled-up version of asset accumulation by the poor. To see this, consider a life-cycle model under certainty with time-separable homothetic (constant relative risk aversion) preferences. Let two types of families each begin with zero assets, have the same preferences, face the same interest rates and face age-earnings profiles that are proportional to one another. Income in any year for the first type of family ("high earnings") is $\alpha > 1$ times as great as it is for the other type ("low earnings"). Under these assumptions, in every period consumption and accumulated assets of the high-earnings type will be α times as great as those of the low-earnings type; the ratio of assets to income for the high-earnings family will be identical to that of the low-earnings family.

Adding earnings uncertainty to the above model does not necessarily change this result. If the probability distribution for all future incomes is such that every possible realization of income for the high-earnings type is α times as great as for the low-earnings type (but the corresponding probabilities are identical), then for given realizations of earnings (appropriately scaled by α) over the life cycle, both consumption and assets will be α times as great (see Bar-Ilan, 1991). In this case, the distribution of the ratio of accumulated assets to income will be the same for the two types of individuals.

Suppose that unobservable lifetime earnings are related to educational attainment. The simple example above assumed that the earnings of groups with high or low levels of educational attainment are proportional to one another at every age and state of the world. In reality, the age-earnings path for college-educated workers is more steeply sloped, and the variance of log earnings differs across education groups, issues we discuss in more detail in section IV. Still, the implication of the traditional life-cycle model is that saving behavior of the poor and the rich should differ only to the extent that the distribution of earnings and the age-earnings profile differ across lifetime-earnings groups.

As we show below, the actual pattern of wealth holdings for many households is quite different from the simple prediction of the life-cycle model. Empirically, the wealth accumulation patterns for families with lower education levels are not scaled-down versions of the wealth patterns of families with higher levels of education. The cross-sectional age-wealth patterns for many lower-income families does not exhibit the "hump-shaped" profiles of wealth accumulation predicted by the life-cycle model. By contrast, wealth-age profiles for college-educated families display, to a greater extent, the hump-shaped wealth-age profile consistent with life-cycle predictions.

We examine wealth holdings using the full sample of the 1984 Panel Study of Income Dynamics (PSID). We use the 1984 PSID population weights to make the sample representative of the

U.S. population.[6] Measured wealth is equal to the sum of assets -- including stocks, bonds, checking accounts, and other financial assets; real estate equity; and vehicles -- minus liabilities that include home mortgages and personal debts. This measure includes Individual Retirement Accounts (IRAs) but excludes pension and Social Security wealth. Wealth is generally positive, though a small proportion of respondents reported negative wealth.[7] To control for differences in lifetime income, the sample was stratified into three categories of education of the family head: less than twelve years (no high school degree), comprising 28 percent of the weighted sample; between twelve and fifteen years (with a high school degree), comprising 52 percent; and sixteen years or more years (college degree) comprising the remaining 20 percent.[8]

---

[6] An alternative source would have been the Federal Reserve Board's 1983 Survey of Consumer Finances. The PSID survey was not as comprehensive as the Survey of Consumer Finances because it did not oversample the wealthy. According to Curtin, Juster, and Morgan (1989), however, the PSID was surprisingly close in accuracy to the SCF except among the very wealthy.

[7] Negative wealth was truncated at -$20,000 for three individuals. In a number of cases, respondents did not reply to questions about wealth holdings of specific assets. In these cases, the interviewer attempted to bracket the amount of assets by asking sequential questions: *e.g.*, are your stock holdings $10,000 or more; if not, are they $1000 or more, etc. We estimated the assets of those who fell within particular brackets to be equal to the average holdings within the same bracket of those who provided exact answers. Note that because the sample was linked to earnings data during 1983-87, we exclude from the sample families who experienced major compositional changes during this period.

[8] An alternative approach to using education as a proxy for lifetime income would be to stratify by average earnings during the sample period. Such an approach is probably less accurate than using education; current earnings may

(continued...)

Scatter diagrams of the wealth holdings by age for these three groups, presented in Figures 1(a) - 1(c), emphasize the sharp differences in wealth accumulation patterns. To adjust for differences in population weighting, each observation is "jittered" by placing dots (equal in number to the population weight) randomly around the family's reported wealth.[9] Quintile regressions that estimate the 20th, 40th, 60th, and 80th percentiles of wealth holdings as a cubic function of age are superimposed on each of the graphs.

Under the simple homothetic model above, wealth holdings will be proportional to lifetime income. To evaluate this hypothesis, we have adjusted the vertical axis in Figures 1(a) - 1(c) to correct for differences in lifetime resources. To do this, we calculate a simple measure of "permanent income": the constant annual real flow of consumption that the average life-cycle household could afford given the education-specific profile of after-tax earnings, Social Security payments, and pensions between age 21 and 85 (assuming a real rate of interest of 3 percent).[10] For those with the

---

[8](...continued)
not be a good predictor of future earnings, nor is information on past earnings always available for retirees.

[9] For example, an observation with a weight of unity would yield a single dot in the graph, while an observation with a weight of ten would result in ten dots randomly arrayed around the sample observation. The graphs are produced using the "jitter" option in STATA.

[10] Equivalently, this number may be viewed as "amortized" lifetime income, since it has the same present value as actual earnings and retirement income. The estimating equations used to calculate average earnings and retirement income can be found in *Appendix A* of Hubbard, Skinner, and Zeldes (1994b).

Figure 1a: Net Wealth by Age, 1984 PSID: No High School Degree

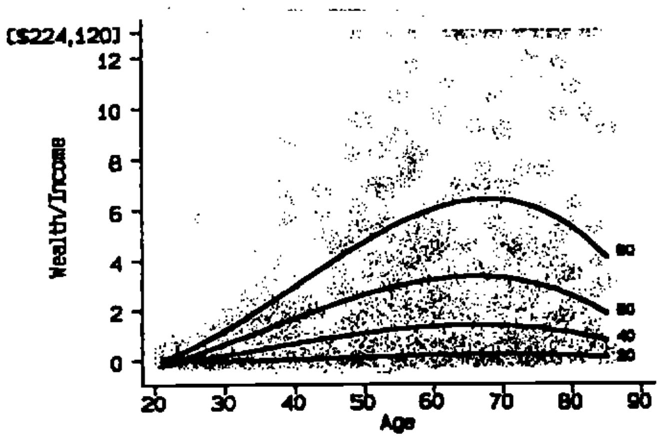

Figure 1b: Net Wealth by Age, 1984 PSID: High School Degree

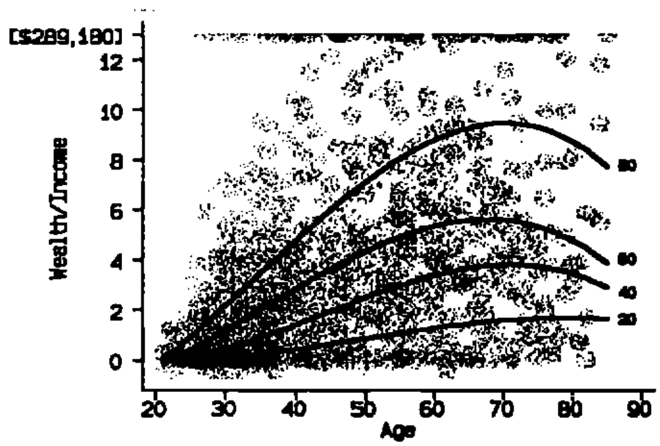

Figure 1c: Net Wealth by Age, 1984 PSID: College Degree

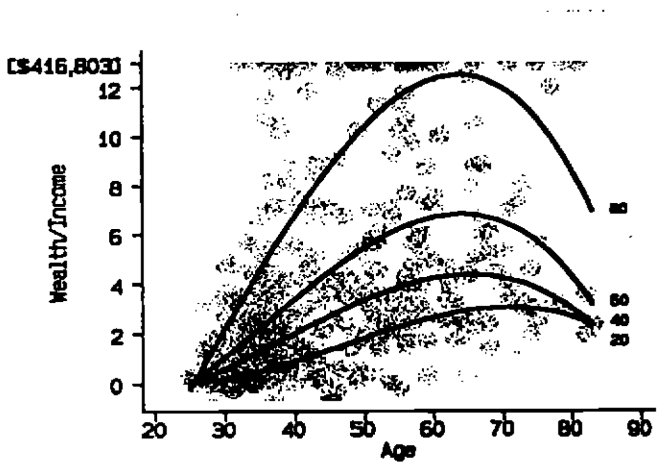

Notes to Figures 1(a) - 1(c): Predicted 20th through 80th percentiles of the wealth distribution, expressed as cubic polynomials in age, are also shown. The vertical axis measures the ratio of reported individual net wealth to (education-specific) average permanent income. Average permanent income for those without high school degrees is $17,241, for high school graduates $22,244, and for college graduates $32,062. The maximum (dollar) wealth level shown at the top of the vertical axis is thirteen times permanent income.

lowest educational attainment, the level of "permanent income" is $17,241, for high school graduates, $22,244, and for college graduates, $32,062. Thus, lifetime earnings are approximately twice as high for college graduates as for those with no high school degree. Wealth is plotted as a multiple of this measure of permanent income. A wealth corresponding to 3.0 among college graduates, for example, is equivalent to $96,186 in net wealth. We truncate the graphed wealth distribution at 13 times the benchmark income level for each education group to promote legibility of the graphs (the truncated values are shown, also jittered, along the top of the respective graph).[11]

Beginning first with Figure 1(a), the cross-sectional evidence indicates that, over the life cycle, many households without high school degrees own very little wealth, even during the ten years prior to retirement that would normally correspond to years in which wealth is highest. The 40th percentile of net wealth for this group is less than $20,000 at all age groups. High school graduates, in Figure 1(b), accumulate a moderate amount of wealth. The wealth accumulation pattern of college-educated households appears most consistent with the life-cycle model; by ages 50 and beyond, very few households hold less than $50,000 in net household wealth. Of course, inferring life-cycle patterns from cross sectional data is speculative, but Figures 1(a)-1(c) lend support to the notion that

[11]Thus, as marked in brackets on the vertical axis, the highest level of wealth graphed for those without a high school degree is $244,120 (13 x $17,241), while the highest level of wealth graphed for those with a college degree is $416,803.

typical wealth accumulation patterns differ substantially by lifetime income.[12]

Detailed wealth holdings by age and by education are shown in Table 1, with all averages weighted by the PSID family weights. Median household wealth is shown both inclusive of and exclusive of housing equity, where housing equity is calculated as the market value of the house less the outstanding mortgage balance. Median income measures labor income, transfer income (including food stamps), pension income, and Social Security benefits for the family head and spouse. Simple ratios of median wealth or median non-housing wealth to median income suggest sharp differences in asset accumulation patterns across educational groups for older age groups. For the lowest education group at ages 50-59, for example, median nonhousing wealth is only about half of the median income. By contrast, median nonhousing wealth is twice median income for households headed by college graduates.

To examine the wealth distribution further, we calculate the percentage of households with net total wealth less than one year's income. This is an arbitrary but convenient measure of "low wealth" households. Table 1 shows that, for younger households, households with less net wealth than current income constitute the vast majority of each education group, ranging from three-fourths among college-educated households to nearly seven-eighths among households without a high school degree. For older cohorts, the differences in wealth holdings become more apparent. Virtually all

---

[12] This result is consistent with the findings of Bernheim and Scholz (1993).

Table 1: Median Wealth and Income, By Age and Education, 1984

	Age					
	< 30	30-39	40-49	50-59	60-69	70 +
			No High School Degree			
Median Wealth ($)	650	13,450	20,000	44,000	36,800	28,000
Median Nonhousing Wealth ($)	605	3,000	5,500	11,500	7,500	7,800
Median Income ($)	10,800	17,000	19,954	20,792	8,860	5,936
Wealth < Income (%)	86.3	68.3	50.7	30.0	29.6	25.0
Nonhousing Wealth < Income/2 (%)	86.1	79.9	75.2	49.8	40.7	39.7
Number of Households	132	161	155	217	211	198
			High School Degree			
Median Wealth ($)	6,855	28,300	62,600	90,300	88,506	92,500
Median Nonhousing Wealth ($)	4,500	8,100	15,500	39,000	38,068	17,700
Median Income ($)	21,360	27,000	30,000	26,808	15,840	9,028
Wealth < Income (%)	81.6	51.4	27.3	15.8	13.7	7.4
Nonhousing Wealth < Income/2 (%)	80.7	66.5	45.4	31.0	20.3	12.0
Number of Households	346	604	238	205	148	101
			College Degree			
Median Wealth ($)	11,000	54,700	113,000	179,000	157,000	115,500
Median Nonhousing Wealth ($)	8,300	17,600	41,000	96,000	83,000	57,760
Median Income ($)	26,000	37,000	47,476	48,000	29,264	18,200
Wealth < Income (%)	74.9	38.4	22.9	4.6	0.4	0
Nonhousing Wealth < Income/2 (%)	67.6	50.4	31.6	22.0	6.4	6.6
Number of Households	39	227	71	86	41	27

Source: Panel Study of Income Dynamics, 1984.

Note: "Wealth < Income" reports the weighted percentage of the sample with net worth (including housing equity) less than after-tax income net of asset income. Similarly, "Nonhousing wealth < Income/2" reports the weighted percentage of the sample with nonhousing wealth less than one-half of income as defined above. All figures are in 1984 dollars.

households aged above 50 and with a college degree hold wealth greater than or equal to one year's income. For those without a high school degree, at least 25 percent of every age group hold net wealth less than current income. An intermediate pattern holds for high school graduates. The percentage of households with nonhousing wealth below one-half of current income (a measure that abstracts from illiquid home ownership) follows much the same pattern, although there are a larger absolute number of households who hold less than half of a year's income in nonhousing wealth.

To summarize, the empirical evidence suggests strongly that wealth accumulation patterns differ by lifetime income. We briefly consider four potential explanations for this differential pattern of wealth accumulation.

First, with a bequest motive it is plausible that households with higher lifetime income hold more assets, especially later in life, because they plan to leave bequests. Those with lower lifetime income are more likely to find the bequest motive inoperative since they expect their children to do better economically (Feldstein, 1988; Laitner, 1990). The absence of "negative bequests" for currently low-income households introduces a corner solution, and hence skewness in the distribution of bequests. In addition, individuals with higher levels of educational attainment may receive greater inheritances to the extent that lifetime income is correlated across generations.

The problem with this explanation is that for those with low lifetime income, wealth accumulation is far below even that predicted by traditional (certainty) life-cycle models. As we discuss below, the life-cycle model predicts at least a modest degree of wealth accumulation to provide for retirement. However, the fact that median nonhousing wealth for the lowest education group is only one-half of income for households prior to retirement (those aged 50-59) suggests that the life-cycle model does not fully capture the saving patterns of this group.

Second, wealth accumulation across education groups may also differ because of differences in the shape of the earnings profile, or in the degree to which Social Security, private pensions and other transfers replace earnings in retirement (as mentioned above). For example, since Social Security benefits equal a larger fraction of average earnings for lower-income workers, such families would not need to save as much relative to higher income workers to ensure adequate consumption during retirement.

As we show below, this explanation alone cannot explain more than a small fraction of the difference in wealth distributions. While lower-income households benefit from the higher earnings replacement rates in Social Security benefits, higher-income (in our case, college graduate) families are more likely to receive private pensions. College-educated households should, moreover, save less relative to income in early years in a life-cycle model because of their more steeply sloped earnings path.

The third possible explanation for the difference in the wealth distribution is variation in rates of time preference by education group. Lawrance (1991), for example, has estimated that college-educated households have lower rates of time preference than lower-income, non-college-educated households. Hence the difference in wealth accumulation could just be the result of different preferences. The lower-income households save little because of their higher rate of time preference, while the higher-income households (or those who are sufficiently patient to attend college) save more.

The Lawrance estimates are based on (food) consumption growth in the PSID during the 1970s and early 1980s. She found that consumption of college-educated households grew faster than that of non-college-educated households, leading her to conclude that college-educated-households have lower rates of time preference. However, Dynan (1993) has shown that this faster growth may have been the consequence of the rapid rise of income for college-educated relative to non-college-educated households. Dynan finds little difference across education groups in the estimated rate of time preference once income changes have been accounted for.[13] While we view differences in rates of time preference as a potentially important factor in wealth accumulation, it seems unlikely that variation in preferences alone can explain the large cross-sectional differences in wealth accumulation.[14]

---

[13] Fuchs (1982) attempted to discern differences in time preference rates by direct survey methods, but he did not find any consistent patterns across education groups.

[14] There are some additional explanations which we have not fully explored. The first is a more general (nonhomothetic) utility function, such as one that (continued...)

The fourth possible explanation is that, in the presence of significant uncertainty about earnings and medical expenditures, lower-income households may rationally accumulate proportionately less than higher-income households because of the existence of an asset-based means-tested social insurance "safety net." This approach follows two strands in the previous literature. Kotlikoff (1988) used simulations to show that a Medicaid program reduced precautionary saving against uncertain medical expenses, while Levin (1990) focused on the impact of Medicaid on the demand for health insurance depending on initial wealth or income.[15] Our work builds on these two insights in a general dynamic programming model of uncertainty, and we pursue it below.

III. OPTIMAL CONSUMPTION WITH TRANSFER PROGRAMS

We begin this section by writing down our general multiperiod model with multiple sources of uncertainty. We then examine simplified versions including a two-period model under

---

[14](...continued)
includes a subsistence level of consumption or a varying intertemporal elasticity of substitution (Atkeson and Ogaki, 1991). Second, length of life and/or age of retirement may differ across education groups. Third, attainable rates of return may be higher for high education or income groups (Yitzhaki, 1987). For further discussion, see Masson (1988).

[15] Levin (1990) studied how uncertainty about medical expenses and the Medicaid program affected the demand for health insurance rather than saving. His empirical results provide evidence on the demand for insurance a (function of the second derivative) rather than on precautionary saving (a function of the third derivative).

certainty and under uncertainty. In these examples, we show how the existence of a minimum level of consumption guaranteed by (asset-based) means-tested social insurance programs affects the optimal consumption choice. Later in the paper, we use numerical methods to examine optimal consumption and wealth accumulation in the general multiperiod model.

A. The Consumer's Optimization Problem

We assume that the household maximizes expected lifetime utility, given all of the relevant constraints. At each age *t*, a level of consumption is chosen which maximizes:

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

subject to the transition equation:

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

plus the additional constraints that:

$$A_s \geq 0, \quad \forall\ s. \tag{3}$$

Equation (1) indicates that consumption excluding medical spending C_s is chosen 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 account for random date of death D_s is a state variable that is equal to one if the individual is alive and zero otherwise, and T is the maximum possible length of life. The family 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 interest. It then receives exogenous earnings E_s, pays out exogenous necessary medical expenses M_s, and receives government transfers TR_s. It is left with:

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

which, following 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 *per se* from medical expenditures; the costs are required only to offset the damage brought upon by poor health.[16] The borrowing and terminal constraints in equation (3) prevent negative assets in any period.[17]

Transfers received depend on financial assets, earnings, and medical expenses:

$$TR_s = TR(E_s, M_s, A_{s-1}(1+r)). \tag{5}$$

This general form allows transfer programs to include earnings-based and wealth-based means testing, as well as payments tied to medical expenses. For simplicity, we consider the following

---

[16] Kotlikoff (1988) considers alternative models of health expenditures.

[17] In 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.

parameterization:

$$TR_s = \max \left[0,\ (\bar{C} + M_s) - (A_{s-1}(1+r) + E_s)\right] \qquad (6)$$

We define $\bar{C}$ as the minimum level of consumption guaranteed by the government, and will refer to this as the consumption "floor." Transfers equal this consumption floor $\bar{C}$ plus medical expenses minus all available resources, if that amount is positive, and zero otherwise. In other words, transfer payments, if made, guarantee a minimum standard of living $\bar{C}$ after medical expenses are paid. However, transfer payments are reduced one-for-one for every dollar of either assets or current earnings. The transfer function captures, in a simplified way, the penalty on saving behavior of asset-based means-tested programs such as Medicaid, AFDC, and food stamps. Because eligibility is conditional on having assets less than a given level, such programs place an implicit tax rate of 100 percent on wealth above that limit. While in the model we restrict social insurance to those with no assets at all, in practice, asset limits range between \$1000 and \$3000.[18]

---

[18] The asset limit for AFDC is \$1000 in almost all states (it is less than \$1000 in a few states). Excluded from the assets subject to this limit are housing equity (up to a certain limit), automobile equity (up to \$1500), and, in some states, burial insurance and plot, farm machinery and livestock and household furnishings. The limit for food stamps is \$2000 for non-elderly and \$3000 for elderly households, with somewhat more liberal exclusions, while for SSI the limits are \$2000 for single households and \$3000 for married couples, again with somewhat less stringent exclusions on automobile equity and other types of wealth. Eligibility for SSI or AFDC is usually a necessary precondition to qualify for Medicaid. See Committee on Ways and Means (1991). For simplicity, we assume that the wealth limit is zero over the entire year.

Before we examine the effects of uncertainty and social insurance programs on wealth accumulation in the general model, we present some two-period models to provide intuition. Consider first a two-period certainty model, with all medical expenses as well as initial assets set equal to zero. Suppose that $E_1 > \bar{C}$, so that the household is not eligible for transfers in the first period, but that $E_2 < \bar{C}$, so it is at least potentially eligible for transfers in the second period. To see the effect of the consumption floor, consider the expression for second period consumption:

$$C_2 = (E_1 - C_1)(1 + r) + E_2 + TR_2 \ . \tag{7}$$

Substituting in the expression for transfers in (6) yields:

$$C_2 = \text{Max}\left[\bar{C},\ (E_1 - C_1)(1 + r) + E_2\right] \ . \tag{8}$$

Differentiating equation (7) or (8) with respect to C_1 gives:

$$\begin{aligned} \frac{dC_2}{dC_1} &= 0, && \text{if} \quad TR_2 > 0 \\ &= -(1 + r), && \text{otherwise.} \end{aligned} \tag{9}$$

Thus, consuming one less unit today yields $(1+r)$ extra units tomorrow if the household is not participating in the transfer program tomorrow (the usual intertemporal tradeoff), but zero extra units if it is.

The indifference curves and budget constraints for two different levels of initial resources E_1 (including any initial assets) are shown in Figure 2.[19] For this example, we assume homothetic utility and $r=\delta=E_2=0$.[20] First consider the case of a lower-wealth household with initial resources of E_1^*. The budget constraint when the consumption floor equals $\bar{C}$ is given by $mnb^*E_1^*$. An interior solution leads to a^*, where $C_1=C_2$. Because of the nonconvexity of the budget constraint, there exists another possible solution to the problem: the household could consume all of income today so that the guaranteed consumption level is received in the second period. This possible solution is indicated by b^*. Since b^* is preferred to a^* ($\hat{U}^* > U^*$), the global optimum is b^*. Individuals with low initial resources will save nothing and instead rely on the consumption floor in the second period.

At the higher level of income, E_1^{**}, however, the budget line is $rsb^{**}E_1^{**}$, and the interior solution a^{**} dominates the alternative of b^{**} since $U^{**} > \hat{U}^{**}$. Thus, individuals with somewhat higher initial resources choose not to rely on the consumption floor and therefore must save to finance future desired consumption.

The solution to this two-period model is as follows: For levels of wealth (or earnings) that are low, but greater than $\bar{C}$, the slope of the consumption-wealth profile is one -- all wealth is consumed. At some critical level of wealth, consumption drops sharply, so that at higher wealth, the consumption function reverts to a straight line through the origin with a slope of 0.5; that is, half of

---

[19]We thank Eric Engen for pointing out this graphical interpretation.

[20]The maximization problem and solution are described in the Appendix.

Figure 2: Optimal Consumption with a Welfare "Floor"

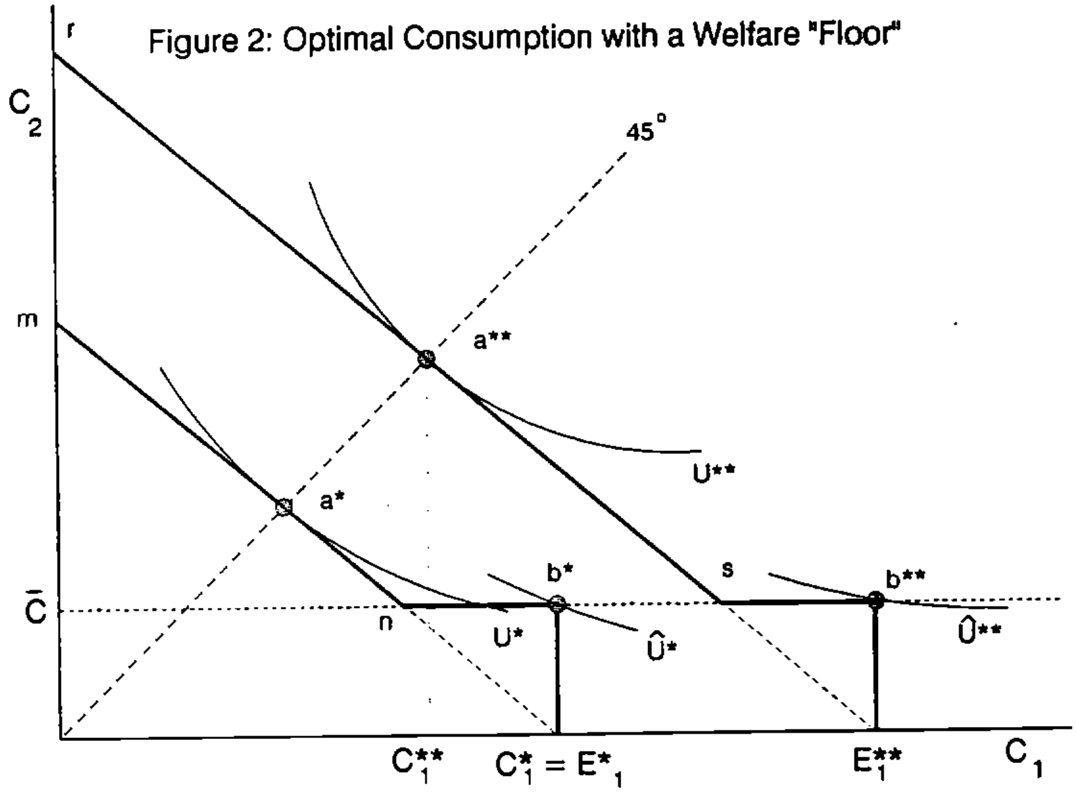

the wealth is consumed in the first period, half in the second, just as it would be in the absence of the transfer program.

This example thus has implications for the marginal propensity to consume out of wealth. As shown in Figure 2, a low level of initial resources E_1^* implies consumption C_1^*. A rise in initial resources to E_1^{**}, however, causes consumption to *decline* to C_1^{**}. That is, over this range of wealth, the marginal propensity to consume out of wealth can actually be negative as the household switches from a regime of consuming all income to one in which it saves for the future.[21] This is in sharp contrast to standard models in which consumption is always increasing in wealth.

One way to generalize this result is to expand the time horizon to three or more periods.[22] A second way to generalize the

---

[21] There is a clear parallel here with the studies of labor supply with nonconvex budget constraints by Burtless and Hausman (1978), Hausman (1981), Moffitt (1986), and Moffitt and Rothschild (1987). As they noted, in a static choice model of leisure and market goods, transfer programs often create kinked budget constraints and can generate multiple local maxima.

[22] Assume that both second- and third-period earnings are less than the consumption floor, but that first-period earnings exceed the floor. In this case, there are three local optima. The individual can: (i) forego transfer payments altogether and choose the traditional interior solution (so that the MPC out of resources is ⅓), (ii) receive transfers only in the third period, so that an interior Euler equation solution holds between first- and second-period consumption (so that the MPC out of resources is ½), or (iii) receive transfers in both the second and third periods (so that the MPC out of resources is unity). Finding the global solution to this model involves choosing the one of these three potential solutions that maximizes utility. For the details of this, see the Appendix.

The result that the MPC depends on the effective horizon of the consumer also appears in model with a borrowing constraint. However, the important difference between the two models is the motivation for consuming

(continued...)

two-period model is to add uncertainty about second-period resources E_2. For now, think of E_2 as earnings less out-of-pocket medical expenses, so the uncertainty may be attributable to either source. Suppose that there were a 50-percent chance of a "good" realization, E_{2g}, and a 50-percent chance of a "bad" realization E_{2b}. Continue to assume that $r=\delta=0$ and that utility is homothetic. Let $E_1 + E_{2b} > 2\bar{C}$, so the individual could save enough to avoid the floor even in the "worst case," if so desired. The maximization problem with respect to C_1 becomes:

$$\max_{C_1} U(C_1) + \frac{1}{2}U[(E_1 - C_1 + E_{2g})(1-Q_{2g}) + \bar{C}Q_{2g}] \qquad \textbf{(10)}$$
$$+ \frac{1}{2}U[(E_1 - C_1 + E_{2b})(1-Q_{2b}) + \bar{C}Q_{2b}] + \mu_1[E_1 - C_1] ,$$

where the first two expressions in brackets are consumption in the good state, C_{2g}, and consumption in the bad state, C_{2b}. The indicator values Q_{2b} and Q_{2g} take on the value of one when income transfers are received under the bad and good scenario, respectively. The first-

[22](...continued)
all of one's wealth. In the model with borrowing constraints, one saves nothing because of high anticipated future earnings. In this model, one saves nothing because of the low anticipated future earnings relative to the consumption floor.

Including more (and thus shorter) time periods leads to smoother consumption-wealth functions. However, at least in the case of a continuous time certainty model, one can show that the marginal propensity to consume wealth may still be (smoothly) negative in the presence of means-tested social insurance.

order condition is written:

$$U'(C_1) = \frac{1}{2}\left[U'(C_{2g})(1-Q_{2g}) + U'(C_{2b})(1-Q_{2b})\right] + \mu_1. \quad \textbf{(11)}$$

Under uncertainty, the first-order conditions indicate a tradeoff between the marginal utility of consuming an extra dollar today and the expected marginal benefit of saving the dollar for the future. In future states of the world in which the household receives a transfer, an extra dollar carried over from the previous period is worthless to the household, because it leads to a one dollar reduction in transfers, leaving future consumption unchanged at $\bar{C}$. In the *Appendix*, we describe the solution to this problem, and show that there exist three local maxima, two of which are interior solutions that satisfy the Euler equation. We also show that households with higher initial resources are more likely to choose the solution that involves much higher saving and a lower probability of receiving transfers. Thus, optimal consumption can again decline as wealth increases over some ranges. Finally, we show that the welfare program affects the saving of those households who have some probability of receiving transfers, even if, *ex post*, they never receive transfers.

Before proceeding, it is worth emphasizing that, if we assume that the period utility function has a positive third derivative (which induces precautionary saving in the presence of earnings uncertainty), there are two distinct effects of introducing an asset-based means-tested social insurance program. One effect comes from the provision of the transfer, and would be present even if the program involved no asset-based means testing. The government is providing a transfer that raises income in the bad states of the world. This serves to reduce the precautionary motive and causes households (particularly low lifetime income households) to save less. The second effect comes from the asset test itself. The government effectively imposes a 100-percent tax on assets in the event that the household receives a health-expense or earnings shock large enough to make it eligible to receive the transfer. This tax further reduces desired saving, again primarily for low-lifetime-income households. In the model used in this paper, we consider the joint effect of these channels on households' consumption.

IV. PARAMETERIZATION AND SOLUTION OF THE MULTIPERIOD MODEL

In this section, we begin by describing the utility function and parameterization of the model. These are described more fully in Hubbard, Skinner, and Zeldes (1994b). We then examine the empirical magnitude of the consumption "floor," and close with a discussion of the numerical solution to the dynamic programming problem.

When we estimate empirical parameters characterizing uncertainty, our primary interest is in uninsured risk -- that is, the risk faced by households conditional on existing insurance coverage. In the model, for example, the effect of uncertainty in lifespan on saving is conditional on a preexisting pension and Social Security payment that acts as a partial annuity. Similarly, our estimates of uncertainty with respect to health expenses condition on preexisting private insurance and Medicare, and are therefore based only on the uninsured out-of-pocket risk.

A. Parameterization of the Model

The Utility Function. We assume that the period utility function in (1) is isoelastic[23]:

$$U(C_s) = \frac{C_s^{1-\gamma} - 1}{1 - \gamma} . \tag{12}$$

We assume a value for γ of 3, which is consistent with many empirical studies. The rate of time preference δ is assumed to be 3 percent per annum for all education groups, and the real after-tax rate of interest is assumed to be 3 percent per annum.[24]

Lifespan Uncertainty. We use mortality probabilities based on mortality data (from 1980) as a function of sex and age from the National Center for Health Statistics and the Social Security Administration (Faber, 1982). Calculating mortality probabilities for a representative family is problematic, given the mixture of married and single households. We use the mortality probabilities for

---

[23]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, 1990b). 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.

[24]We review empirical estimates of γ and present sensitivity analyses using alternative values of γ and δ in Hubbard, Skinner, and Zeldes (1994b).

women. These capture both the expectations of life for single women, and the expectation of life for a currently married family in which the husband dies first. The maximum possible age in the model is set to 100; since we assume economic life begins at age 21, there are a maximum of 80 periods in the model.

Earnings Process. Time-series patterns of earnings and wages have been the subject of many studies (see *e.g.*, Lillard and Willis, 1978; MaCurdy, 1982; and Abowd and Card, 1989). Our measures of earnings risk differ in two general respects. First, we include unemployment insurance and subtract taxes in our measure of "earnings"; these adjustments are likely to reduce earnings variability.[25] Second, we separate our sample into three educational categories.

Earnings during working years are uncertain and correlated over time and follow :

$$\begin{aligned} y_{it} &= Z_{it}\beta + u_{it} + \nu_{it} \\ u_{it} &= \rho_e u_{it-1} + e_{it}, \end{aligned} \tag{13}$$

where y_{it} is the log of earnings, Z_{it} is a cubic polynomial in age and year dummy variables (included to control for cohort productivity growth) and β is a vector of coefficients. The error term u_{it} follows an AR(1) process, where e_{it} is a white-noise innovation. The variable υ_{it} is a combination of i.i.d. transitory variation in earnings

---

[25] Carroll (1992) also included transfer payments in his measure of earnings for the same reason. In our model, means-tested transfers such as AFDC and food stamps are excluded from the definition of earnings because they are received only if assets are sufficiently low. Instead, they are included in the consumption floor.

and measurement error. To simplify the dynamic programming model, we assume that v_{it} is entirely measurement error and ignore it in our parameterization of the model. Hence our measure of earnings uncertainty is conservative because it excludes all transitory variation in earnings. We assume in the model that the head of the household retires at age 65, at which point the family receives Social Security, pensions, and other non-asset income with certainty.

Estimates of the uncertainty parameters are summarized in the top panel of Table 2.[26] The results imply substantial persistence in shocks to earnings, a result that is consistent with many of the studies cited above. In addition, the log of labor income is more variable for non-high-school graduates than it is for the two other educational groups.

Out-of-Pocket Medical Expenses. We use data from a merged sample of observations from the 1977 National Health Care Expenditure Survey and the 1977 National Nursing Home Survey to calculate a cross-sectional distribution of out-of-pocket medical expenses. Our measure of medical costs includes expenses paid by Medicaid, because Medicaid payments are determined endogenously in our model as the difference between total medical costs and available financial resources of the family.

We assume a model of medical spending of the following

---

[26]When we estimate the uncertainty parameters in (13) (σ^2_ϵ, ρ_ϵ, σ^2_v), we exclude households with very low earnings realizations. When we estimate the mean age-earnings profile, we estimate the equations in levels rather than logs, and include all households. Details of the estimation approach are given in Hubbard, Skinner, and Zeldes (1994b).

Table 2: Parameters for Uncertain Earnings and
Uncertain Medical Expenses for Dynamic Programming Model

	No High School	High School +	College +
Earnings			
AR(1) coefficient (ρ)	0.955	0.954	0.959
Variance of the innovation e	0.033	0.026	0.020
Variance of combined measurement error and transitory shock v	0.040	0.028	0.018
Out-of-Pocket Medical Expenses[1]			
Total Medical Expenses (including Medicaid)	\$2023	\$1974	\$2149
AR(1) coefficient (ρ_m)	0.901	0.901	0.901
Variance of the innovation ε	0.175	0.156	0.153

Notes: See equation (12) for the time-series model of earnings and equation (13) for the time-series model of out-of-pocket medical expenses.

[1] Based on 1977 cross section study of the National Health Care Expenditure Survey and the 1977 National Nursing Home Survey, and data and estimation methods in Feenberg and Skinner (1992). See Hubbard, Skinner, and Zeldes (1994b) for more detail.

form:

$$m_{it} = G_{it}\Gamma + \mu_{it} + \omega_{it}$$
$$\mu_{it} = \rho_m \mu_{it-1} + \varepsilon_{it} \quad , \qquad \textbf{(14)}$$

where m is the log of medical expenses, ω_{it} is the purely transitory component, assumed to be entirely measurement error, μ_{it} follows an AR(1) process (where ε_{it} is a white-noise innovation), and G_{it} is a quadratic in age and an individual fixed effect. We estimate separately for elderly individuals aged 65 years or over and the nonelderly. The estimates are presented in the bottom part of Table 2. The merged cross-section data set enables us to estimate more accurately the cross-sectional distribution of medical spending by education group and by age, but not the time-series properties of medical expenses. Instead, we use estimates of ρ_m from Feenberg and Skinner (1992), who use a quadrivariate tobit procedure with a panel of tax data from 1968 to 1973 to measure the time-series pattern of declared medical spending (in excess of 3 percent of adjusted gross income). There is surprisingly little difference in the overall level of medical spending by education group, implying that average medical expenses are a larger fraction of lifetime income for low education groups. This is in part because of the much higher Medicaid spending for the lower education groups.

Consumption Floor. Finally, the consumption floor is defined as the level of consumption guaranteed by the government above and beyond medical expenses. Measuring the means-tested consumption floor is difficult, since potential payments from social insurance programs differ dramatically according to the number of children, marital status, age, and even the recipient's state or city. Nevertheless, we make a first approximation by calculating separate consumption floors for "representative" families both under age 65 and over 65. Details of the calculation are in *Appendix A* of Hubbard, Skinner, and Zeldes (1994b) and are largely based on figures in Committee on Ways and Means (1991).

We include in our estimate of the floor only means-tested transfer payments such as AFDC, food stamps, and Section 8 housing assistance for those under age 65, and SSI, food stamps, and Section 8 housing assistance for those over age 65.[27] Unemployment insurance is not included in these transfers because it is not means-tested; instead, it is included in net earnings. Medicaid is also not included as part of the floor because it is used exclusively to pay for medical expenses.

We distinguish between entitlement and non-entitlement programs. Under entitlement programs, everyone who is eligible may sign up. Despite the fact that many who are eligible do not take advantage of the program, the money is at least potentially available to them. Housing subsidies are not entitlements, since there are often waiting lists. In such cases, we include the expected value of benefits -- *i.e.*, the probability of receiving the benefits times the dollar amount, in our estimate of the floor.

---

[27] We assume these benefits are valued by recipients at their dollar cost. Moffitt (1989) estimates that food stamps can largely be valued as cash, and Section 8 housing subsidies are unlikely to distort consumption behavior given that the vouchers are generally for an amount less than market rent.

For the nonelderly, the median AFDC and food-stamp transfers to a female-headed family in 1984 with two children and no outside earnings or assets was \$5764. The representative family is assumed to include a single parent with children; if the father were present in the household, or married to the mother, then benefits would be reduced in some states of residence. We assume that housing subsidies are received entirely from the Section 8 housing program, which provides housing vouchers for existing rental property. The mean housing subsidy paid is multiplied by 0.35 to adjust for the fact that only 35 percent of the eligible population who actually receive the Section 8 housing subsidy. Hence the net (expected) housing subsidy is \$1173. Summing AFDC and housing subsidies yields a combined "safety net" for the non-elderly of \$6937.

For the elderly, a weighted average of single and married families implies that combined SSI and food stamp annual payments in 1984 were \$5400, inclusive of median state supplements. Adding Section 8 housing benefits for elderly families yields a net total "safety net" of \$6893. Because the measures for the elderly and nonelderly are close to \$7000, we adopt a common value for both groups of \$7000 for the consumption floor, $\bar{C}$.

For a number of reasons, this estimate should be treated with caution. Calculating the consumption floor for individuals in nursing homes, for whom SSI is reduced to only \$30 per month for spending money, is difficult because it involves valuing the room and board provided by the nursing home. The "safety net" for a couple with grown children in their fifties, before they are eligible to receive SSI, is likely to be much less than the $7000 floor assumed above. Furthermore, in using expected values of housing subsidies, we ignore the more complicated problem of uncertainty about the value of the consumption floor faced by potential recipients.

B. Numerical Solution of the Dynamic Model

Because we cannot solve the household's multiperiod problem analytically, 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.

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 discrete nodes, with a maximum and minimum equal to plus and minus 2.5 standard deviations of the unconditional distribution. 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 current 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.[28] The problem is solved by starting in the last possible period of life (T) and solving backward. In 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.[29]

Once we determine the optimal consumption function for all possible nodes, we simulate a history for each of a large number of families (16,000). For each family, we use the following procedure.

---

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

[29] In total, optimal consumption is calculated at more than 230,000 individual wealth-health-earnings-age nodes. Each optimal consumption calculation involves searching over a large number of consumption choices, and the expected marginal utility and value function must be calculated for each of these possible choices. All 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.

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 medical expenses from the appropriate distributions.[30] 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.

V. SIMULATED DISTRIBUTIONS OF AGE-WEALTH PROFILES

We begin by presenting the wealth accumulation pattern of a model in which the mean values of medical expenditures and earnings are anticipated with certainty, and lifespan is also certain. In this certainty benchmark, consumption and wealth paths differ across education groups, but are identical within each educational group. We examine whether differences in the age profile of medical expenses, earnings, and retirement income can explain the observed (average) differences in wealth accumulation.

---

[30] We draw a starting value for earnings and medical expenses for period 1 from a log-normal distribution with variance equal to the unconditional variance of the distributions. Subsequent draws for medical expenses and for earnings (through retirement) are drawn from the conditional distributions.

The earnings, health, consumption, and wealth profiles for the lowest and highest education groups are shown in Figures 3(a) and 3(b). Again, the graphs are scaled to adjust for differences in lifetime income across education groups. Consider first the lowest education group. The household's consumption is limited by borrowing constraints until its mid-thirties. After that point, it accumulates wealth, arriving at a level of wealth at retirement of about five times peak earnings, and then gradually spends down accumulated wealth. Next consider the highest education group. The wealth-age path is very similar to that for the lowest education group. That is, differences in the profile of earnings and retirement income cannot explain the differences in mean wealth-income ratios between the lowest and highest education groups. While households with lower levels of income may experience higher replacement rates from social security benefits (and hence less need to save for retirement), they are also less likely to receive pension income. On balance, pension plus Social Security income yield a similar fraction of pre-retirement earnings for the two education groups, leading to similar wealth-income profiles.

In order to analyze not just *mean* wealth profiles, but the *distribution* of wealth for different groups (given our assumption of homogenous preferences), we need to examine a model with uncertainty. Therefore, we next examine the predictions of the dynamic programming model subject to income, health, and lifespan

Figure 3a: Consumption and Wealth by Age
No High School Degree, Everything Certain

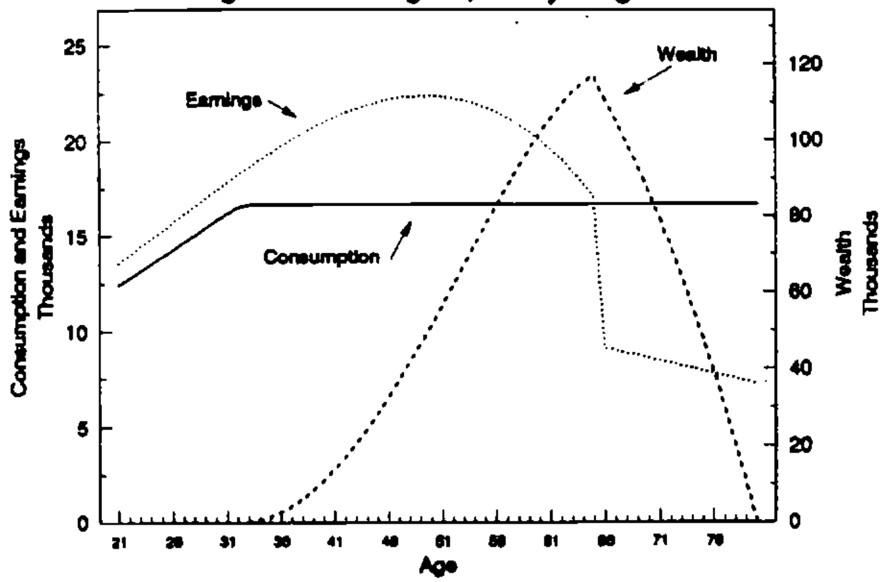

Figure 3b: Consumption and Wealth by Age
College Degree, Everything Certain

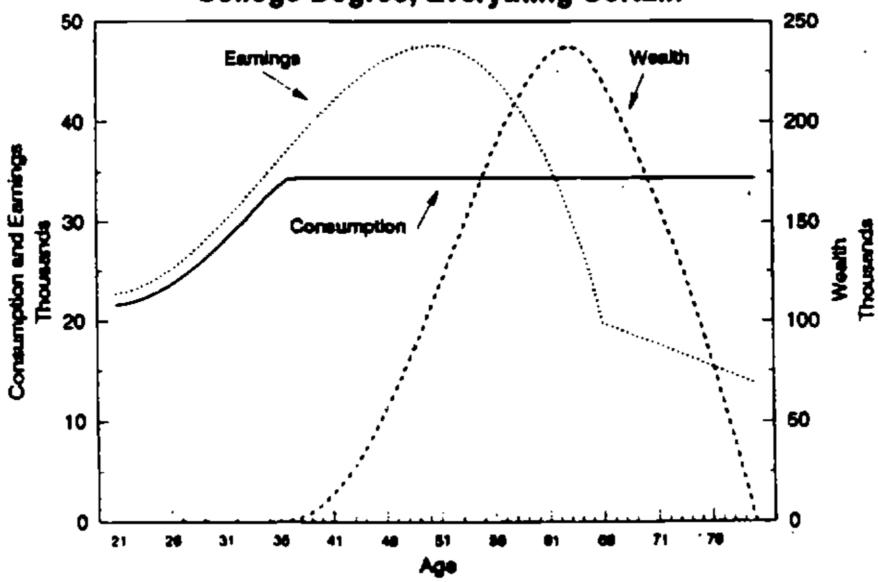

uncertainty,[31] but with a minimal guaranteed consumption floor of $1000. In this case, there is little difference in the wealth accumulation patterns of the lowest and the highest lifetime-income groups. Tabulations in Table 3 compare the fraction of families with wealth less than income in the PSID (the first column) and the simulated data (the second column). For the simulated data based on a $1000 consumption floor, each educational group has virtually the same small fraction of "low wealth" households. Furthermore, this precautionary saving model dramatically underpredicts the proportion of "low wealth" households, especially for those without college degrees.[32]

Finally, we consider the most realistic specification: a social insurance program that guarantees a $7000 consumption floor.[33] Figures 4(a) and 4(b) depict the predicted wealth accumulation patterns for the two educational categories. These wealth profiles are taken from 16,000 households simulated for each education

---

[31]In this version of the model, accidental bequests arising from lifespan uncertainty are effectively confiscated, since no other generation receives them. Experiments in which the average (education-group-specific) bequest was given to members of the next generation at the beginning of their working lives yielded higher steady-state asset-income ratios. However, this approach provides younger generations with an unrealistically large initial stock of assets. An alternative approach would have younger generations face uncertain future inheritances. This more general model is a topic for future research.

[32] Similarly, Carroll and Samwick (1992) have shown that wealth accumulation in the conventional precautionary saving model is implausibly high for individuals with low time preference rates.

[33] Note that we are varying the minimum guaranteed level of consumption (the consumption "floor"). In this paper, we do not consider changes in the asset limit, which is assumed for simplicity to be zero in our model.

Table 3: Percentage of Families With Wealth < Income, Actual and Simulated

Age	Education	Actual (PSID)	Simulated $1000 Consumption Floor	Simulated $7000 Consumption Floor
< 30	No High School	86.3	43.7	80.9
	College	74.9	90.8	93.5
30-39	No High School	68.3	8.0	50.2
	College	38.4	49.8	66.2
40-49	No High School	50.7	3.7	34.1
	College	22.9	11.0	25.8
50-59	No High School	30.0	1.6	24.5
	College	4.6	0.5	4.9
60-69	No High School	29.6	2.3	19.9
	College	0.4	0.5	3.0
70-80	No High School	25.0	0.5	25.2
	College	0.0	0.0	1.3

Source: Data are from the 1984 PSID and authors' calculations.

Figure 4a: Simulated Net Wealth by Age: No High School Degree

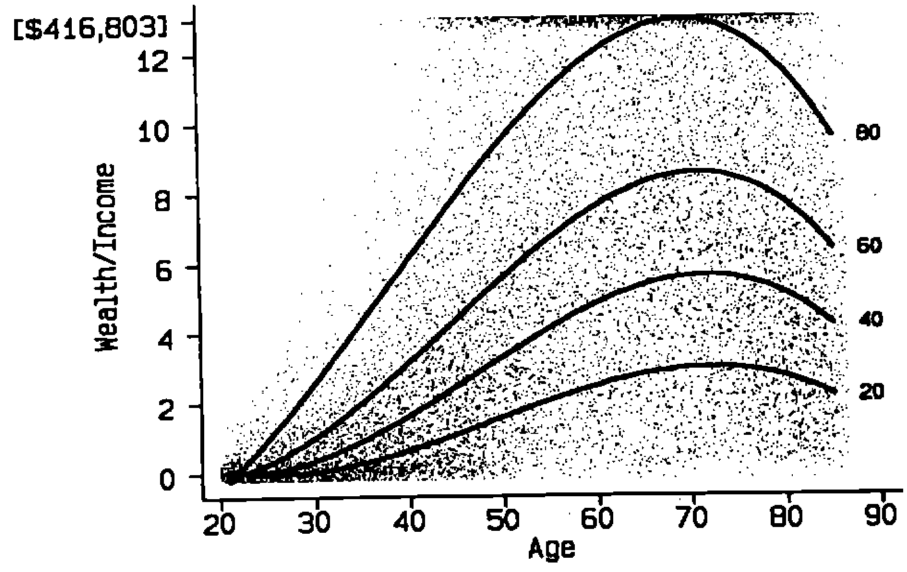

Figure 4b: Simulated Net Wealth by Age: College Degree

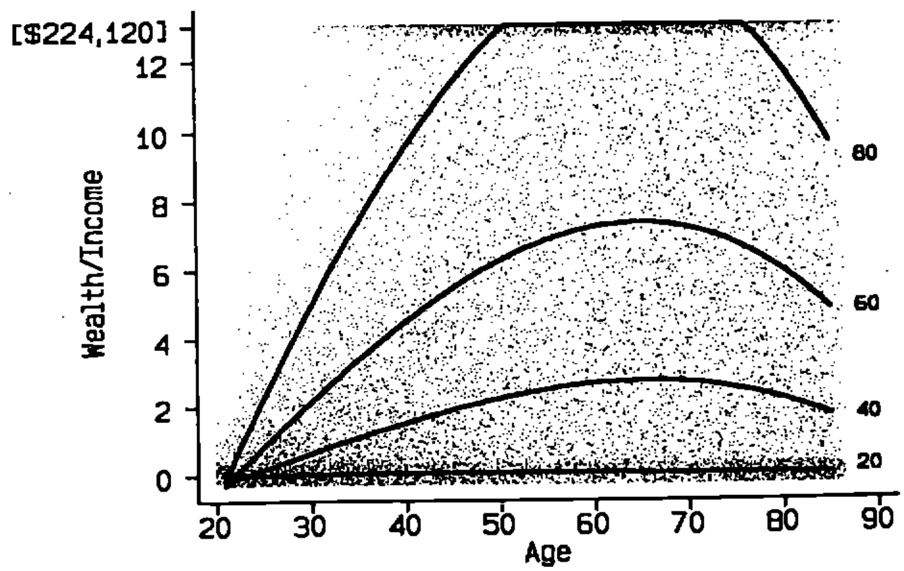

Notes to Figures 4(a) - 4(b): Predicted 20th through 80th percentiles of the wealth distribution, expressed as cubic polynomials in age, are also shown. The vertical axis measures the ratio of reported individual net wealth to (education-specific) average permanent income. Average permanent income for those without high school degrees is $17,241, and for college graduates $32,062. The maximum (dollar) wealth level shown at the top of the vertical axis is thirteen times permanent income. Wealth data are simulated using the dynamic programming model described in the text.

group and are drawn to the same scale, and with the same quintile regressions, as the graphs in Figure 1.[34] Consider first the graph for college graduates, Figure 4(b). The quintile regressions for the simulated age-wealth profiles match closely the actual wealth profiles in Figure 1(c), for all of the quintiles.[35] Note in particular that in both the actual data and the data simulated by the model, there is substantial wealth accumulation for the bottom quintile. For example, simple tabulations show that the 20th-percentile level of wealth among those aged 50-59 is 2.4 years of (permanent) income in the PSID, and 2.8 years in the simulated data. In general, this model with uncertainty about earnings, medical expenses, and length of life does a good job at explaining the distribution of wealth for this group.

Next consider the graph for those with no high school degree, Figure 4(a). In the simulated data, wealth for the bottom 20th percentile of this group is bunched near zero for all ages, just as it is in the actual PSID data in Figure 1(a). For example, the

---

[34] Although we have calculated the entire lifetime wealth profile for each of these households, we chose only one randomly selected wealth-age combination per household to replicate a cross-sectional sample.

[35] We have used cubic polynomials in age to summarize the quantile distributions. These cubic approximations, however, may be inadequate in summarizing wealth distributions for given age groups, which may be better revealed using nonparametric approaches. These more detailed comparisons (shown in the Appendix) suggest that at the ages of peak wealth, the simulation model tends to overpredict wealth accumulation for the higher quantiles. For example, the actual 60th-percentile level of net wealth (from the PSID) among college-educated households at age 50-59 is $216,000, while the simulated 60th-percentile wealth level for the same age group is $278,000.

tabulated 20th-percentile level of wealth among those aged 50-59 is very low: 0.25 years of (permanent) income in the PSID data and 0.35 years in the simulated data. That is, the model is capable of explaining one of the key "puzzles" in the data -- unlike the high-lifetime-income group, a significant fraction of the middle-aged low-lifetime-income group has virtually no wealth.

In the third column of Table 3, we present the fraction of households with wealth less than income in the two education groups for the higher value of the consumption floor. The entries generally correspond closely to figures tabulated from the PSID. For example, the simulated percentages of "low wealth" households at age 50-59 are 24.5 percent and 4.9 percent for no-high-school and college-educated households, respectively, compared with the corresponding actual PSID tabulations of 30.0 and 4.6 percent. To summarize, the simulation model replicates well the wide disparity by lifetime-income group in the fraction of households with low levels of wealth.

Finally, Table 4 documents the fraction of households receiving means-tested transfers, based on 1984 data from the PSID, by age and by education group. The tabulations from the PSID data, in the first column, are contrasted with the simulated percentages given a consumption floor of $1000 (the second column) and $7000 (the third column). Assuming a consumption floor of $1000 implies that few households in either education group receive means-tested transfers. By contrast, a $7000 consumption floor implies that a much larger percentage of households with lower levels of

Table 4: Percentage of Families Receiving Transfer Payments: Actual and Simulated

Age	Education	Actual (from PSID)	Simulated $1000 Consumption Floor	Simulated $7000 Consumption Floor
< 30	No High School	48.4	1.6	25.0
	College	0.0	0.0	2.5
30-39	No High School	24.9	0.7	19.5
	College	0.9	0.1	0.9
40-49	No High School	23.7	0.9	12.5
	College	2.3	0.1	0.7
50-59	No High School	12.7	0.2	10.7
	College	0.0	0.0	0.1
60-69	No High School	19.0	0.3	9.1
	College	0.0	0.0	0.3
70-80	No High School	23.0	0.1	11.8
	College	0.0	0.0	0.0

Note: "Positive transfers" means that the family received AFDC, SSI, or food stamps.
Source: Data are from the 1984 PSID and authors' calculations.

educational attainment receive transfers, with little effect on college-educated households. For example, at ages 50-59, the actual percentage of households without a high school degree receiving transfers is 12.7. The simulated percentage is 10.7 with a $7000 floor, but only 0.2 percent with a $1000 floor. Few college-educated households receive transfers at any age. Overall, the simulated model with a $7000 floor closely matches age- and education-related patterns of income transfer receipts.

The dynamic programming model with a $7000 floor generally predicts accurately differences in wealth accumulation patterns across education groups. However, it performs poorly in two respects. First, the model overpredicts the fraction of "low wealth" college-educated households at younger ages (Table 3). Because of the more steeply sloped earnings profile for college-educated households, the simulation model predicts that many of these households will possess very little wealth prior to age 40. This contrasts with the actual patterns from the PSID, perhaps because of *inter vivos* transfers. Second, the simulated 60th and 80th percentile age-wealth profile for households with low education levels are considerably higher than the corresponding actual profile from the PSID. For example, for ages 50-59, the 60th percentile of wealth in the PSID is $59,000, compared to the 60th-percentile value in the simulated data of $147,000.

With a conventional utility function and empirically consistent parameters for earnings and health expenses, our simulation model predicts a large impact on wealth accumulation of means-tested welfare programs. We have presented evidence that our model is consistent with important features of the empirical distribution of wealth. Is there additional direct empirical evidence that can shed light on whether differences in the structure of government-provided assistance programs can predict empirical differences in saving behavior as our model suggests?

A formal statistical test of how government social insurance programs affect saving behavior is beyond the scope of this paper. Nevertheless, we consider below two types of evidence that may bear on the empirical issue of how social insurance affects wealth accumulation: the first based on historical trends in social insurance policy in the United States, and the second based on cross-sectional differences in saving behavior, either by states or by income groups.

In our approach, all else equal, an expansion in the magnitude of means-tested social insurance programs (measured by an increase in $\bar{C}$) should reduce wealth holdings of low-lifetime-income households, while having little effect on wealth holdings of high-lifetime-income households. The reason is that, on account of the increase in $\bar{C}$, low-lifetime-income households face a greater likelihood of participating in the government consumption-maintenance programs and reduce their saving accordingly.

To examine this prediction, one would need to examine differences in the distribution of assets by lifetime-income groups in periods with "low" values of the consumption floor $\bar{C}$ and periods with "high" values of the consumption floor. One might think that a good natural experiment would be a comparison of the early 1960s with a more recent period such as the 1980s. Detailed wealth data are available in the 1962 Survey of Financial Characteristics of Consumers and the 1983 Survey of Consumer Finances. The size of means-tested programs expanded substantially between 1962 and 1983, with expenditures more than one and one-half times their 1962 level by 1983. Real spending on means-tested in-kind transfers (food stamps, housing subsidies, and Medicaid) rose even more dramatically over the 1962-1983 period (see Burtless, 1986; and Ellwood and Summers, 1986).

However, there are at least two problems with this as a natural experiment. First, there are a large number of other factors that have changed between the 1960s and the 1980s.[36] Second, the real benefits from AFDC and food stamps for a single mother with a family of four rose by only 5.2 percent, from $6612 to $6957 (in 1984 dollars), between 1964 and 1984. The increase in total

---

[36]Factors other than the consumption floor were not constant over the 1962-1983 period. For example, average real out-of-pocket medical expenses for the elderly has risen from $962 in 1966 to $1562 in 1984, which was also likely accompanied by an increase in the *variance* of such out-of-pocket expenses. (See U.S. Congress, Select Committee on Aging, House of Representatives, "Emptying the Elderly's Pocketbook - Growing Impact of Rising Health Care Costs," Comm. Pub.No. 101-76, page 25. Our calculation is expressed in 1984 dollars; we adjust from 1966 data using the CPI-U.) Increased out-of-pocket health expenses could lead to greater saving while young in anticipation of future medical expenses, but could also discourage saving by those with greater potential eligibility for Medicaid. In addition, there may be greater uncertainty about the growth rate of earnings across education groups, especially given the divergence during the 1980s in earnings for those without high school education relative to those with a college education (see Levy and Murnane, 1992). Finally, the asset limits for the programs we examine were changed significantly between 1962 and 1983 (see Powers, 1993).

expenditures arose from a rapid growth in enrollment rather than a rapid increase in benefits conditional on receiving them. Unfortunately, in its present form our model does not incorporate the changes in family composition, eligibility requirements, or welfare "stigma" that may account for the rapid rise in enrollment in welfare programs and hence the greater likelihood of receiving welfare payments.

Though not reported in detail here, we compared patterns of wealth holdings using the 1962 Survey of Financial Characteristics of Consumers (SFCC) and the 1983 Survey of Consumer Finances (SCF).[37] To control for differences in educational attainment between 1962 and 1983, we defined the low-lifetime-income group to be the bottom quintile of educational attainment (in 1983, the group who had not completed high school) and the high lifetime-income group to be the top quartile of educational attainment (in 1983, college graduates). The data did not show large differences in wealth between the two periods. For example, among households with heads aged 46-60, median wealth fell from 3.8 percent of household income in 1962 to 1.9 percent in 1983. For households of the same age with high lifetime income, median wealth as a percentage of income rose from 34.8 percent to 36.3 percent. In sum, changes in median wealth accumulation between 1962 and 1983

[37] For a description of the 1962 Survey of Consumer Finances, see Projector and Weiss (1963); for a description of the 1983 Survey of Consumer Finances, see Avery and Kennickell (1987).

were not large.[38]

Preliminary cross-section evidence on how asset-based means testing affects wealth accumulation is more supportive of our model. Powers (1994) used data on female-headed households in the National Longitudinal Survey of Women to exploit cross-sectional (state-level) variation in AFDC policy to identify effects of asset limits on wealth levels. In particular, Powers finds that, for two otherwise identical female-headed households who reside in different states, a one-dollar differential in the AFDC asset limit is associated with a 30-cent difference in assets.[39] Moreover, the size of this estimated effect is qualitatively robust to a number of alternative specifications.

Another implication of our analysis is that low wealth holdings by low-lifetime-income households are likely to be an "absorbing state" because of asset-based means testing of welfare programs. In Hubbard, Skinner, and Zeldes (1994a), we compare the persistence of wealth holdings for households in the 1984 and 1989 samples of the PSID. Simulated five-year transition probabilities from our model with uninsured idiosyncratic risks and a means-tested consumption floor of $7000 replicate very closely the

[38]One potential problem with comparing the 1962 and the 1983 Surveys is changes in the accuracy of wealth reporting. For example, Wolff (1987) detailed substantial deviations between the aggregates in the 1962 SFCC and the aggregate household balance sheets.

[39]Powers includes lagged assets in her model with an estimated coefficient not statistically significantly different from unity. Hence, one might interpret her results as corroborating an important effect of asset limits on saving.

observed transition probabilities in the PSID. Results of alternative simulations with a high annual rate of time preference (10 percent) and no consumption floor -- designed to mimic a "buffer stock" approach -- greatly overpredicted the likelihood of a recovery from low levels of wealth.

VI. CONCLUSIONS AND DIRECTIONS FOR FUTURE RESEARCH

Empirical studies using micro data often find a significant group in the population with virtually no wealth, raising concerns about heterogeneity in motives for saving. In particular, this heterogeneity has been interpreted as evidence against the life-cycle model of saving. This paper argues that a life-cycle model can replicate observed patterns in household wealth accumulation once one accounts for precautionary saving motives and social insurance programs. This suggests that a properly specified life-cycle model with precautionary saving and social insurance can be useful for analyzing determinants of household saving and particularly for assessing effects of certain social insurance programs on saving.

Our reconciliation of the generalized life-cycle model with observed patterns of household wealth accumulation proceeds in two steps. First, we show how social insurance programs with asset-based means testing can discourage saving by households with low expected lifetime incomes. The implicit tax bias against saving in this context is significant relative to other areas of tax and expenditure policy, since saving and wealth are subject to an implicit tax rate of 100 percent in the event of a sufficiently large earnings downturn or medical expense.

Second, we evaluate this model of saving and social insurance using a large dynamic programming model with four state variables. Assuming common preference parameters across education groups, we are able to replicate along important dimensions actual wealth accumulation patterns for both lower-and higher-lifetime-income families. The results presented here complement those presented in Hubbard, Skinner, and Zeldes (1994b), in which we argue that a life-cycle model with precautionary saving motives and social insurance can explain aggregate wealth accumulation and observed co-movements of changes in consumption and current income.

In particular, we find that the presence of asset-based means-testing of welfare programs can imply that a significant fraction of the group with lower lifetime income will not accumulate wealth. The reason is that saving and wealth are subject to an implicit tax rate of 100 percent in the event of a earnings downturn or medical expense large enough to cause the household to seek welfare support. This effect is much weaker for those with higher lifetime income for two reasons. First, the consumption floor is a much smaller fraction of their lifetime income and normal consumption levels, and hence represents a less palatable support program. Second, the uninsured risks of medical spending are a smaller fraction of lifetime resources. These results suggest that observed empirical behavior of lower income groups that might appear inconsistent with the life-cycle model (Bernheim and Scholz, 1993), may in fact be consistent with optimizing behavior.

We have made a number of simplifying assumptions in the model that may affect the results we present here. First, we do not control for family compositional changes. Children are likely to increase levels of consumption at middle age, which can generate low levels of wealth accumulation independent of means-tested social insurance programs. For example, Blundell, Browning, and Meghir (1994) suggest that household demographics are a significant explanation of the hump-shaped consumption profile commonly observed in cohort and cross-section data. However, their data also suggest that the average number of children in a family peaks past age 35. Hence, households anticipating future child-rearing expenses (and college expenses) might actually save more while young, which would explain why the empirical data indicates more saving at young ages than that implied by our simulation model. Our model also does not account for life-cycle changes at older ages, and in particular the role of self-insurance against lifespan uncertainty by married elderly couples and their children (see, *e.g.*, Kotlikoff and Spivak, 1981). Allowing for a richer demographic model of consumption might therefore reduce the predicted level of overall wealth accumulation because of greater demand for consumption while middle-aged and less demand while retired.

Second, we ignore bequests in the model. Allowing for bequests is likely to increase the overall level of wealth accumulation in the simulation model (see for example, Hubbard, Skinner, and

Zeldes, 1994b), and may allow a better explanation of saving behavior of the very wealthy. However, including bequests is unlikely to affect our fundamental conclusions about the nature of wealth accumulation at lower income levels. Most people who are potentially eligible for means-tested welfare programs are unlikely to be leaving substantial bequests.

To conclude, the economically significant role in saving decisions by low-income households played by asset-based means testing of many social insurance programs suggests its relevance for public policy discussions of welfare and social insurance. A model such as this can be particularly helpful in evaluating the effects of welfare reform (such as changing the guaranteed level of consumption or the size of the asset limit) on saving by both current and potential future recipients. More broadly, deliberation of the consequences of introducing asset-based means testing for Social Security should also focus on the incentive effects emphasized here.

APPENDIX

Optimal Consumption in Two-Period and Three-Period Models With Certain Earnings

The Lagrangian for the basic two-period problem outlined in the text can be written: where Q_2 is an indicator variable that equals

$$\mathcal{L} = U(C_1) + \frac{U(C_2)}{1+\delta} + \lambda\left\{\left[E_1 + \frac{E_2}{1+r} - C_1 - \frac{C_2}{1+r}\right](1-Q_2) + \left[\frac{\bar{C}-C_2}{1+r}\right]Q_2\right\} + \mu_1(E_1 - C_1), \quad \text{(A1)}$$

unity when the individual is receiving a transfer, and zero otherwise, λ is the marginal utility of income, and μ_1 is the shadow price of the borrowing constraint in the first period. The first-order conditions are:

$$\begin{aligned} U'(C_1) - \lambda(1 - Q_2) - \mu_1 &= 0 \\ U'(C_2) - \lambda\left[\frac{1+\delta}{1+r}\right] &= 0, \end{aligned} \quad \text{(A2)}$$

where $U'(C_s)$ is the marginal utility with respect to period s consumption.

Because of the nonlinearity of the budget constraint, there exist two local maxima for the expression in (A1), one corresponding to $Q_2 = 1$, and the other to $Q_2 = 0$. Whether Q_2 is positive is clearly endogenous; to find the global maximum, we find the two local maxima (corresponding to $Q_2 = 0$ and $Q_2 = 1$) and then choose the larger.

Begin with $Q_2 = 0$. Because we have assumed that $E_2 < \bar{C}$, in order to not receive the transfer the household must have saved resources from period 1; *i.e.*, $Q_2 = 0$ implies $\mu_1 = 0$. Thus, the first-order conditions have the standard interior solution: $U'(C_1) = U'(C_2)(1+r)/(1+\delta)$. When $Q_2 = 1$, so that the household receives a transfer in period 2, the first-order condition is $U'(C_1) = \mu_1$. The household will consume all of its resources in the first period and rely on the consumption floor in the second period.

Figure 2 in the text showed the budget constraints and indifference curves for this two-period problem. Figure A1 provides a different view of the problem facing the consumer. On the horizontal axis is first-period consumption, and on the vertical axis is marginal utility of consumption. The downward sloping curve measures the marginal utility of C_1, which is continuous everywhere. The other, initially upward sloping, curve FF' is equal to $U'(C_2)(1-Q_2)/(1+\delta) + \mu_1$.[40] The two intersections of these curves correspond to the two local maxima described above. Point a again corresponds to the interior solution (with $Q_2 = 0$); at this point FF'

[40] The analytic derivation of FF' comes by substituting $U'(C_2)(1+r)/(1+\delta)$ for λ in the first line of equation (13). This optimality condition can then be broken into two parts; $U'(C_1)$ and (minus) the remainder, which is denoted by FF'; when the two components are equal, of course, the first-order condition is satisfied.

Figure A1: Two Local Maximums for Optimal Two-Period Consumption

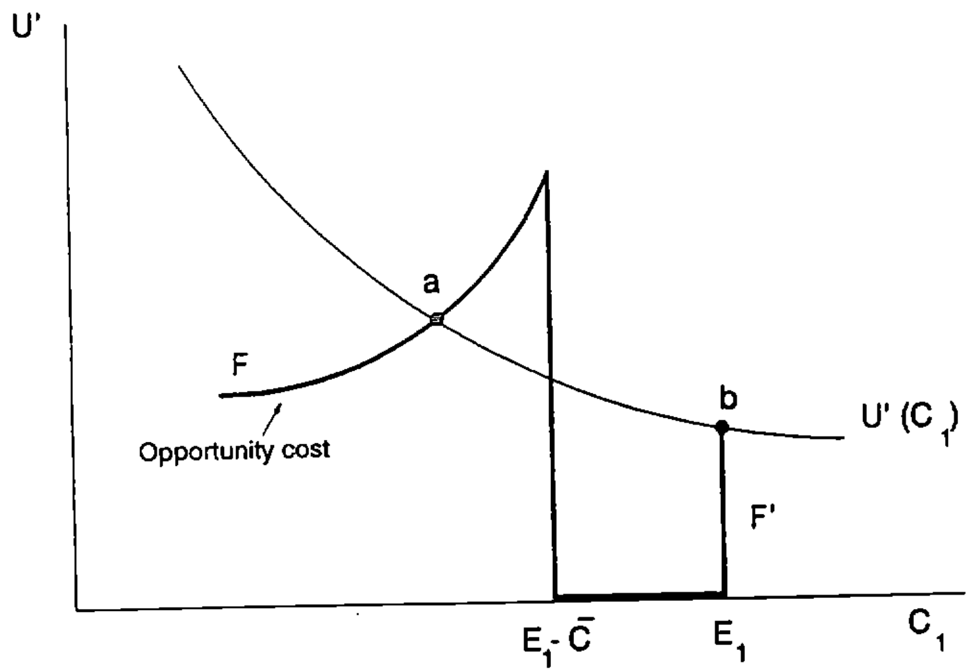

is upward-sloping, since as C_1 rises, C_2 falls, so that $U'(C_2)$ rises. When C_1 is high enough that saving the remainder would cause C_2 to equal $\bar{C}$, the "opportunity cost" of additional first-period consumption drops to zero. At this point $dC_2/dC_1 = 0$, so any additional saving is effectively completely confiscated due to the consumption maintenance program, and it makes no sense for an individual to save any additional dollars. Finally, at $C_1 = E_1$, the borrowing constraint binds, so that $U'(C_1)$ is equal to the shadow price μ_1. The alternative solution is therefore b. Whether point b is preferred to point a cannot be answered without knowing the shape of the utility function over the entire range of consumption.

The solution to this model is shown in Figure A2(a). For levels of wealth (or earnings) that are low, but greater than $\bar{C}$, the slope of the consumption-wealth profile is one -- all wealth is consumed. Above E^*, the consumption function reverts to the traditional interior solution with a slope of 0.5; that is, half of the wealth is consumed in the first period, half in the second.

One way to generalize this result is to expand the time horizon to three periods.[41] Assume that both second- and third-period earnings are less than the consumption floor, but that first-period earnings exceed the floor. Formally, the three-period model can be posed as follows. The Lagrangian can be expressed as: with the obvious generalizations from the two-period model. The condition for $Q_2 = 1$ is similar to the two-period model, while the condition for $Q_3 = 1$ is somewhat more complicated since it relies

[41] It is straightforward but tedious to generalize to more than three periods.

Figure A2(a): Consumption - Wealth Profile: Two-period model

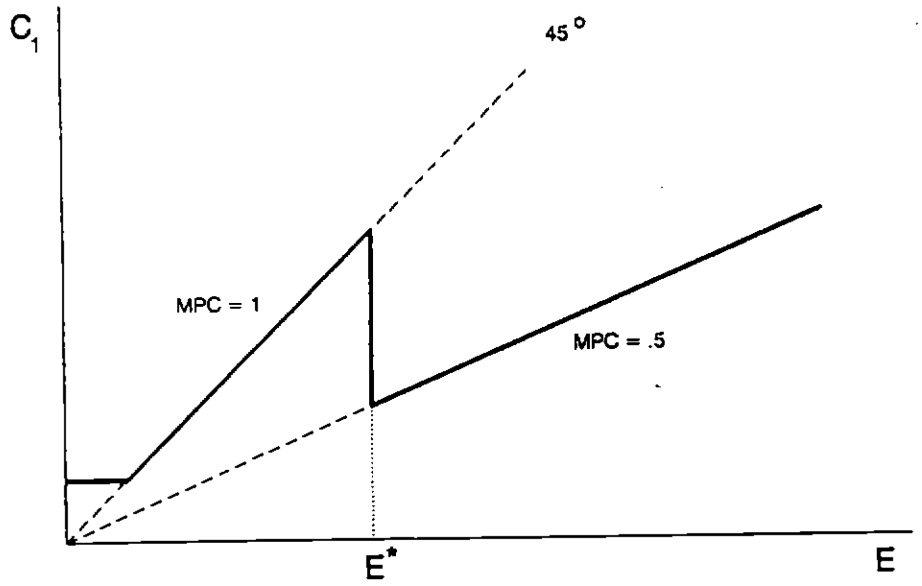

$$
\begin{aligned}
\mathcal{L} = {} & U(C_1,C_2,C_3) \\
& + \lambda \left\{ \left[E_1 + \frac{E_2}{1+r}\right](1-Q_2)(1-Q_3) + \frac{E_3(1-Q_3)}{(1+r)^2} \right. \\
& \qquad \left. - \left[C_1(1-Q_2) + \frac{C_2}{1+r}\right](1-Q_3) - \frac{C_3}{(1+r)^2} \right\} \\
& + \mu_1\{E_1 - C_1\} \\
& + \mu_2\{[(E_1 - C_1)(1+r) + E_2](1-Q_2) + Q_2\bar{C} - C_2\}
\end{aligned}
$$

also on past saving behavior:

$$
\begin{aligned}
Q_3 &= 1 \ , \quad \text{if} \qquad C_2 > X_2 + \frac{E_3 - \bar{C}}{1+r} \\
&= 0 \ , \quad \text{otherwise} \ ,
\end{aligned}
$$

where $X_2 = ((E_1\text{-}C_1)(1+r)+E_2)(1\text{-}Q_2) + Q_2\bar{C}$ is accumulated wealth plus earnings plus transfer payments in period 2. Note then that choosing C_1 such that a transfer program in period 2 will occur (*i.e.*, $Q_2 = 1$) makes it more likely that the individual will also choose $Q_3 = 1$. As will be shown below, under the assumptions of the model, once payments from the income transfer program are accepted in period 2, payments will also be accepted in period 3.

The first order conditions can be written:

$$U_1 - \lambda(1-Q_2)(1-Q_3) - \mu_1 - \mu_2(1+r)(1-Q_2) = 0 ,$$

$$U_2 - \frac{\lambda(1-Q_3)}{1+r} - \mu_2 = 0 ,$$

$$U_3 - \frac{\lambda}{(1+r)^2} = 0 .$$

There are a number of possible strategies that involve either receiving no income transfer payments, or receiving payments in one or two periods. Consider the four possible combinations of Q_2 and Q_3:

(i) $Q_2 = Q_3 = 0$ corresponds to the standard interior solution without any transfer payments. The Euler equation is given by $U_i = (1+r)U_{i+1}$ for $i = 1,2$. It is possible, but unlikely, that the individual's borrowing constraint binds in the first period ($\mu_1 > 0$), and impossible for the borrowing constraint to bind in the second period.[42]

(ii) $Q_2 = 1$, $Q_3 = 0$ is ruled out by assumption. If the individual receives transfer payments in the second period, then he or she will have no savings for the third period. Since $E_3 < \bar{C}$ by assumption, the individual will prefer $\bar{C}$, leading to $Q_3 = 1$ and a contradiction.

(iii) $Q_2 = 0$, $Q_3 = 1$ occurs when an interior solution is

---

[42] Since $E_3 < \bar{C}$ by assumption, if the borrowing constraint is binding in the second period, the household will save nothing for the third period and will prefer $\bar{C}$ to E_3, so that $Q_3 = 1$.

chosen in the second period, but transfers are received in the third period. In this case, the marginal conditions between the first and second periods generate interior solutions; $U_1 = U_2(1+r)$, but $U_2 = \mu_2$. That is, consumption in the first two years are chosen as if in a two-period model, because it makes no sense to carry over a small amount of saving to the final period.

(iv) $Q_2 = Q_3 = 1$ implies that the individual will consume all of his or her initial wealth in the first period ($\mu_1 > 0$) and rely on income transfers thereafter.

Finding the global solution to this model involves choosing the one of these three potential solutions that maximizes utility. Which of the three feasible solutions is a global maximum depends on the form of the utility function, the level of wealth, and the size of the guaranteed consumption floor.

The solution to the model for the case that $E_2 = E_3 = 0$ is presented in Figure A2(b). For lower levels of initial resources $E_1 < E^*$, the MPC out of resources is 1.0 and all wealth is consumed (i.e., case (iii) above). When wealth exceeds E^* but is less than E^{**}, case (ii) is chosen; wealth is split evenly between period 1 and 2, so the MPC of wealth is 1/2, and the consumption floor $\bar{C}$ is consumed in the third period. Finally, for wealth greater than E^{**}, case (i) is chosen; the individual chooses never to be on the floor and an interior solution holds with the MPC out of wealth is equal to 1/3.

Figure A2(b): Consumption - Wealth Profile: Three-period model

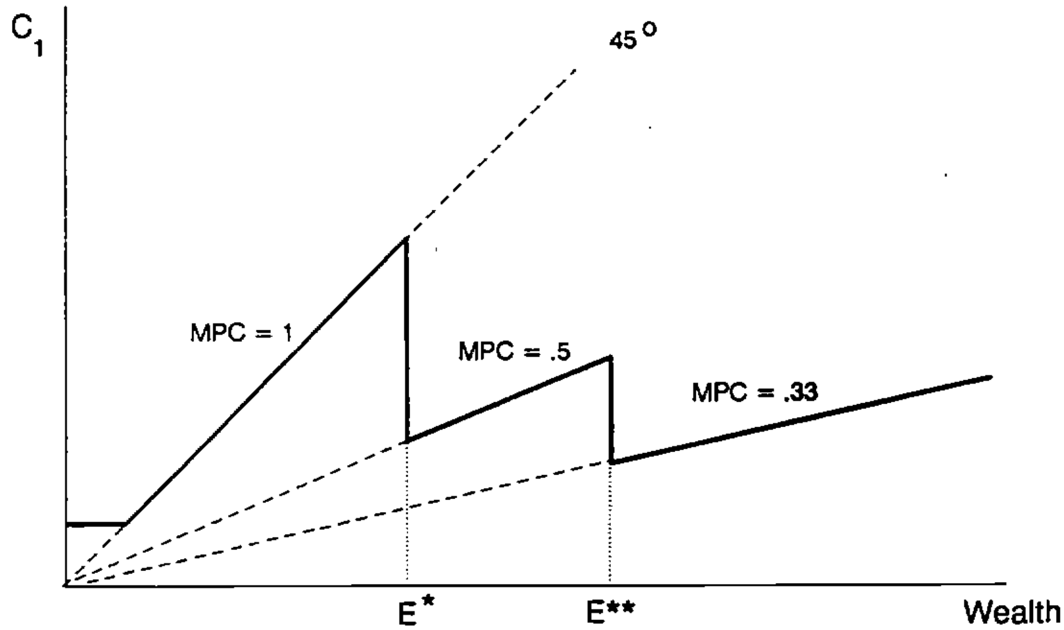

Optimal Consumption in a Two-Period Model With Uncertain Earnings

Figure A3 offers a graphical description of the effect of uncertain second-period resources and social insurance on consumption. On the horizontal axis is first-period consumption, and on the vertical axis is the marginal utility of consumption. The downward-sloping curve, equal to the left hand side of equation (11) in the text, measures the marginal utility of C_1, which is continuous everywhere. The other, initially upward-sloping, curve labeled "opportunity cost" is equal to the right-hand side of equation (11) in the text.

The intersections of these curves represent local maxima. Point *d* corresponds to the interior solution at which the household saves enough to avoid welfare even in the worst earnings outcome ($Q_{2g} = Q_{2b} = 0$). At the point C^*, the amount of saving provides exactly $\bar{C}$ in the bad state of the world in which E_{2b} is realized. This is not an optimal choice, because the household could consume more C_1 today, and still receive $\bar{C}$ in the bad state of the world, owing to the existence of the consumption floor. Hence, the "opportunity cost" curve drops suddenly, as the value of Q_{2b} switches from zero to one. That is, increasing consumption today by $1 causes a reduction in next period's consumption only if the good outcome is realized, so the opportunity cost of $1 consumed today is just the marginal utility of second-period consumption C_{2g}, weighted by the probability that the good state occurs. Point *e* corresponds to the interior solution at which the household receives the consumption

Figure A3: Marginal Utility and Optimal Consumption:
Uncertain Earnings

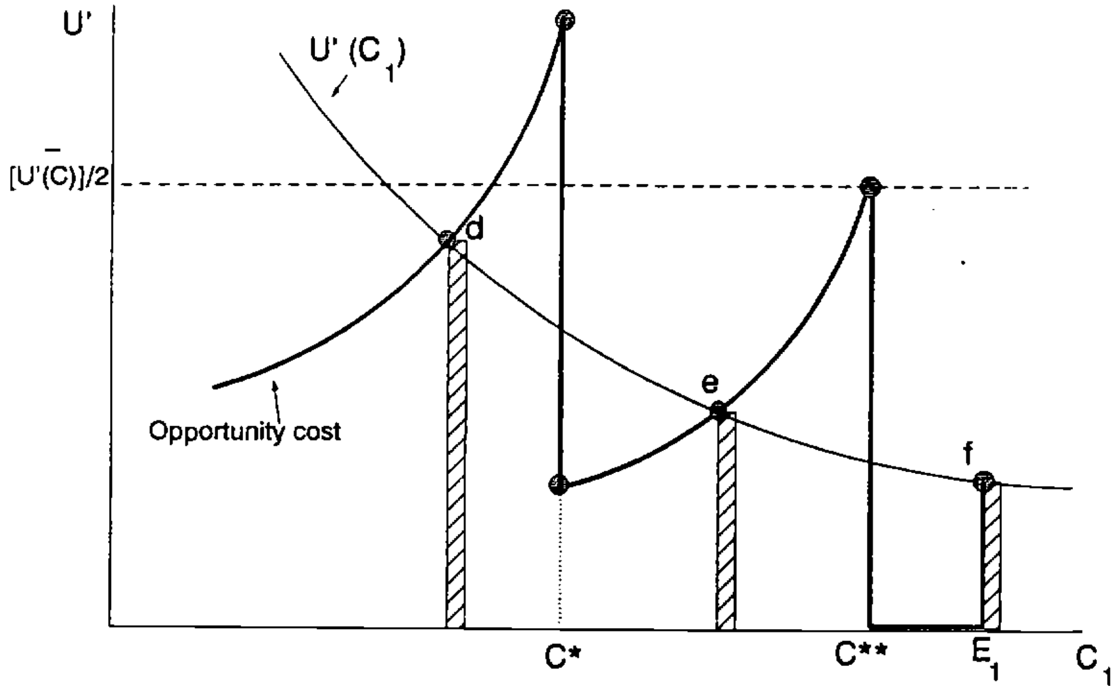

floor in the bad state of the world, but not in the good state of the world. At $\dot{C}^{**}$, first-period consumption is sufficiently high that in either state of the world, the family will be eligible for the consumption floor. The opportunity cost curve drops to zero, because increasing C_1 today by a dollar does not reduce C_2. Finally, at point *f*, the household is consuming all of its resources. The optimal consumption choice corresponds to the global utility maximum that corresponds either to point *d*, *e*, or *f*.

Figure A3 can be used to analyze how an increase in resources E_1 affects the relative value of points *d*, *e*, and *f*. Because of an envelope condition, the increase in utility conditional on choosing C_1 at points *d*, *e*, and *f* equals (approximately) the shaded areas to the right of and below points *d*, *e*, and *f*, respectively. Clearly, the value of *d*, saving against both outcomes, rises by more than *e*, saving against just the good outcome, and the value of *e* rises by more than *f*, making it more likely that the individual will save. In other words, households with higher initial resources are more likely to save for future contingencies, and hence less likely to rely on the consumption floor. In this case of uncertainty about earnings net of medical expenses, the wealth-consumption profile can again be nonmonotonic.

Details by age of the distribution of wealth

The following graphs show the weighted 20th, 40th, 60th, and 80th percentiles of wealth (actual PSID and simulated by the model) for each age. These therefore show unsmoothed data, as opposed to figures 1 and 4 in the text which showed the fitted values for the quintile regressions.

No High School Degree

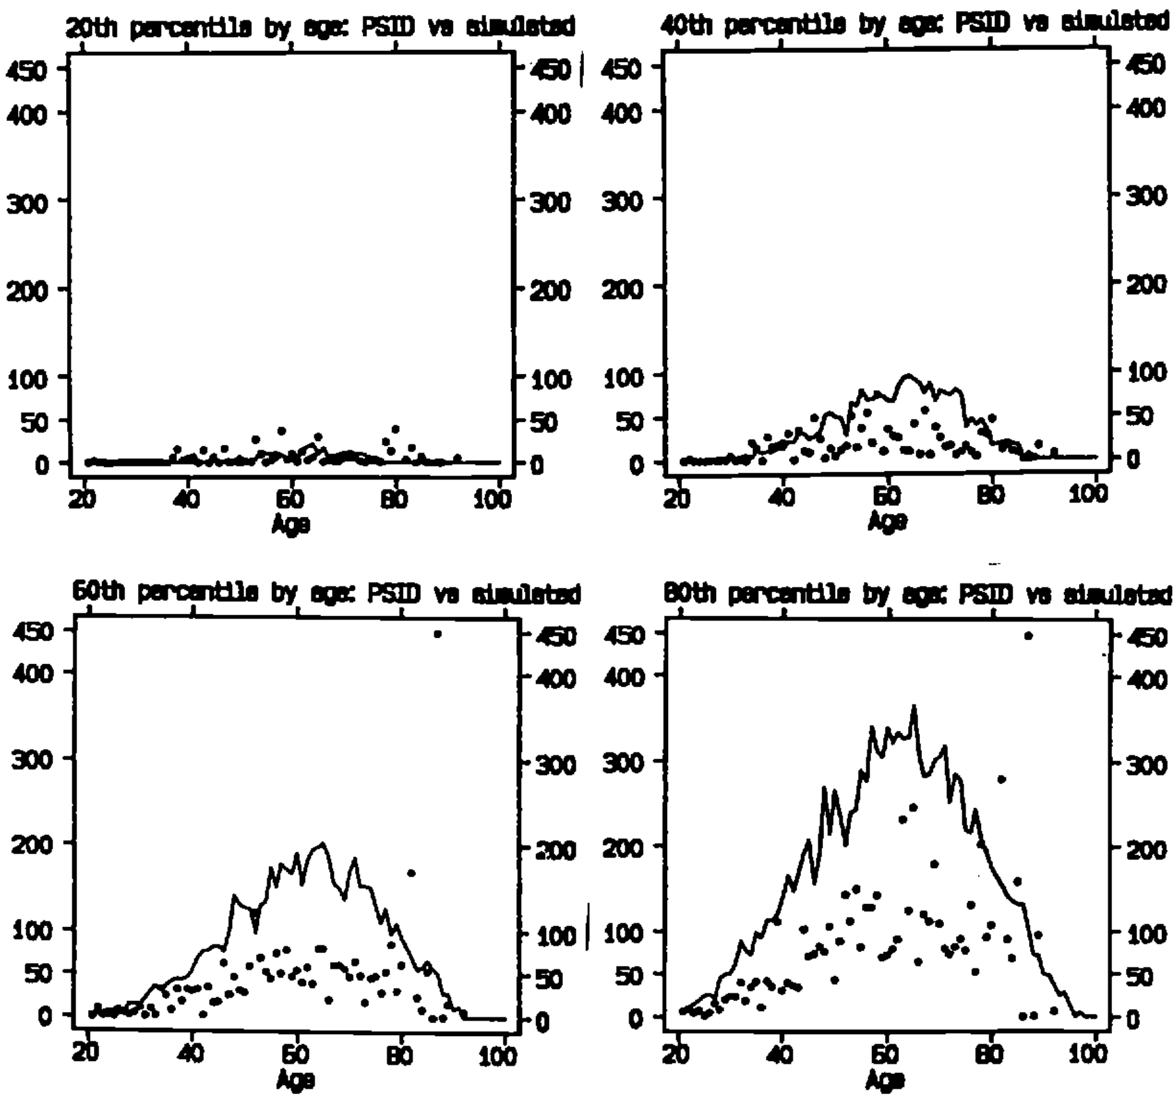

College Degree

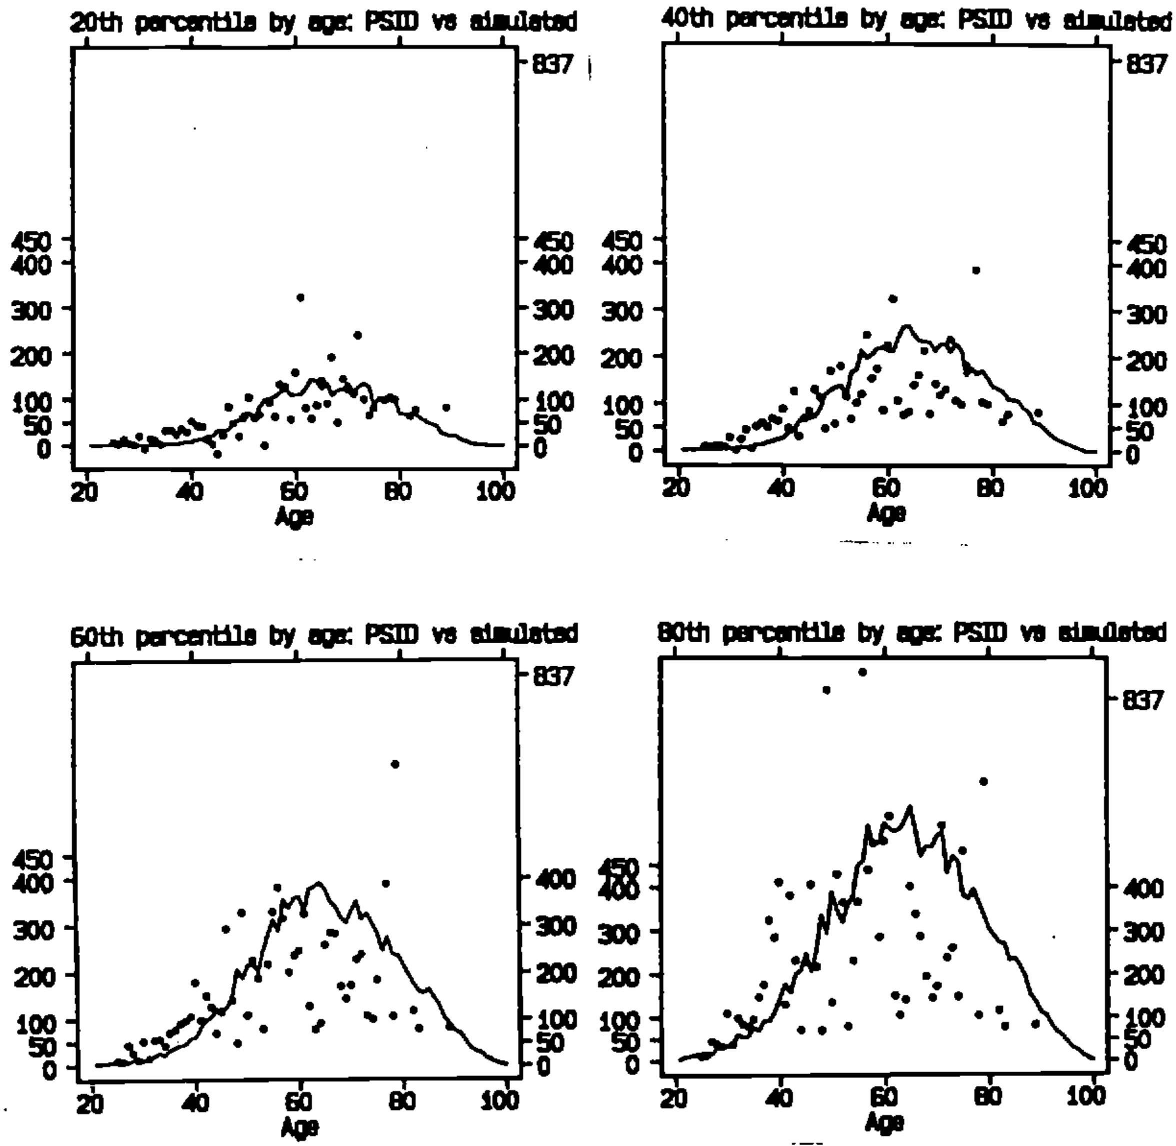

References

Abel, Andrew. "Capital Accumulation and Uncertain Lifetimes with Adverse Selection." Econometrica 54 (September, 1986): 1079-1098.

Abowd, John, and Card, David. "On the Covariance Structure of Earnings and Hours Changes." Econometrica 57 (March 1989): 411-455.

Atkeson, Andrew, and Ogaki, Masao. "Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel and Aggregate Data." University of Rochester Working Paper No. 303, November 1991.

Bar-Ilan, Avner. "On the Proportionality and Homogeneity of Consumption and Income," Mimeograph, University of British Columbia, September 1991.

Bernheim, B. Douglas, and Scholz, John Karl. "Private Saving and Public Policy." In James Poterba (ed.) Tax Policy and the Economy, vol. 7. Cambridge: MIT Press, 1993.

Blundell, Richard, Browning, Martin, and Meghir, Costas. "Consumer Demand and the Life-Cycle Allocation of Household Expenditure." Review of Economic Studies 61 (1994): 57-80.

Burtless, Gary, and Hausman, Jerry A. "The Effect of Taxation on Labor Supply: Evaluating the Gary Negative Income Tax Experiment." Journal of Political Economy 86 (December 1978): 1103-1130.

Burtless, Gary. "Public Spending on the Poor: Trends, Prospects, and Economic Limits." In Sheldon H. Danziger and Daniel H. Weinberg, eds., Fighting Poverty: What Works and What Doesn't. Cambridge: Harvard University Press, 1986.

Caballero, Ricardo J. "Earnings Uncertainty and Aggregate Wealth Accumulation." American Economic Review 81 (September 1991): 859-871.

Cantor, Richard. "The Consumption Function and the Precautionary Demand for Savings." Economics Letters 17 (1985): 207-210.

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

Carroll, Christopher D., and Samwick, Andrew. "The Nature and Magnitude of Precautionary Wealth." Mimeograph, Board of Governors of the Federal Reserve System, 1992.

Committee on Ways and Means. Overview of Entitlements Programs: 1991 Green Book. U.S. House of Representatives, May 7, 1991.

Curtin, Richard T., Juster, F. Thomas, and Morgan, James N. "Survey Estimates of Wealth: An Assessment of Quality." In R. E. Lipsey and H. S. Tice (eds.) The Measurement of Saving, Investment, and Wealth. Chicago: University of Chicago Press, 1989, 473-552.

Davies, James. "Uncertain Lifetime, Consumption and Dissaving in Retirement." Journal of Political Economy 89 (June 1981): 561-578.

Deaton, Angus. "Saving and Liquidity Constraints." Econometrica 59 (September 1991): 1221-1248.

Deaton, Angus. Understanding Consumption. Oxford: Clarendon Press, 1992.

Dynan, Karen. "The Rate of Time Preference and Shocks to Wealth: Evidence from Panel Data." Economic Activity Section Working Paper No. 134, Board of Governors of the Federal Reserve System, July 1993.

Ellwood, David T., and Summers, Lawrence H. "Poverty in America: Is Welfare the Answer or the Problem?" In Sheldon H. Danziger and Daniel H. Weinberg, eds., Fighting Poverty: What Works and What Doesn't. Cambridge: Harvard University Press, 1986.

Engen, Eric. "Precautionary Saving and the Structure of Taxation." Mimeograph, U.C.L.A, 1992.

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

Feenberg, Daniel, and Skinner, Jonathan. "The Risk and Duration of Catastrophic Health Care Expenditures." NBER Working Paper No. 4147, August 1992; forthcoming, Review of Economics and Statistics.

Feldstein, Martin. "The Effects of Fiscal Policies When Incomes Are Uncertain: A Contradiction to Ricardian Equivalence." American Economic Review 78 (March 1988): 14-23.

Feldstein, Martin. "College Scholarship Rules and Private Saving," NBER Working Paper No. 4032 (March 1992).

Fuchs, Victor. "Time Preference and Health: An Exploratory Study." In V.R. Fuchs, ed; Economic Aspects of Health. Chicago: University of Chicago Press, 1982, pp. 93-120.

Hausman, Jerry A. "Labor Supply." In H. Aaron and J. Pechman, eds. How Taxes Affect Economic Behavior. Washington D.C.: Brookings Institution, 1981.

Hays, Constance L. "Girl's Plan to Save for College Runs Afoul of Welfare Rules." The New York Times (May 15, 1992), p.1.

Hubbard, R. Glenn, and Judd, Kenneth L. "Social Security and Individual Welfare: Precautionary Saving, Liquidity Constraints, and the Payroll Tax." American Economic Review 77 (September 1987): 630-646.

Hubbard, R. Glenn, Skinner, Jonathan, and Zeldes, Stephen P. "Expanding the Life-Cycle Model: Precautionary Saving and Public Policy." American Economic Review 84 (May 1994a) : 174-179.

Hubbard, R. Glenn, Skinner, Jonathan, and Zeldes, Stephen P. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." Carnegie-Rochester Conference Series on Public Policy, 1994b, forthcoming.

Hurd, Michael D. "Mortality Risk and Bequests." Econometrica 57 (July 1989): 779-814.

Kimball, Miles S. "Precautionary Saving and the Marginal Propensity to Consume." Working Paper No. 3403, National Bureau of Economic Research, July 1990a.

Kimball, Miles S. "Precautionary Saving in the Small and in the Large." Econometrica 58 (January 1990b): 53-73.

Kotlikoff, Laurence J. "Health Expenditures and Precautionary Savings." In L.J. Kotlikoff, What Determines Savings? Cambridge: MIT Press, 1988.

Kotlikoff, Laurence J., and Spivak, Avia. "The Family as an Incomplete Annuities Market," Journal of Political Economy 89 (April 1981): 372-391

Laitner, John. "Random Earnings Differences, Lifetime Liquidity Constraints, and Altruistic Intergenerational Transfers," Mimeograph, University of Michigan, 1990.

Lawrance, Emily. "Poverty and the Rate of Time Preference: Evidence from Panel Data," Journal of Political Economy 99 (February 1991): 54-77.

Levin, Laurence. "On the Demand for Health Insurance and Precautionary Savings Among the Elderly." Mimeograph, Santa Clara University, August 1990.

Levy, Frank, and Murnane, Richard. "U.S. Earnings Levels and Earnings Inequality : A Review of Recent Trends and Proposed Explanations." Journal of Economic Literature 30 (September 1992) : 323-381.

Lillard, Lee A., and Willis, Robert J. "Dynamic Aspect of Earnings Ability." Econometrica 46 (September 1978): 985-1012.

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

Masson, Andre. "Permanent Income, Age, and the Distribution of Wealth," Annales D'Economie et de Statistique No. 9 (1988): 227-256.

Moffitt, Robert. "Estimating the Value of an In-Kind Transfer: The Case of Food Stamps," Econometrica 57 (March 1989): 385-409.

Moffitt, Robert. "The Econometrics of Piecewise-Linear Budget Constraints," Journal of Business and Economic Statistics 4 (July 1986): 317-328.

Moffitt, Robert, and Rothschild, Michael. "Variable Earnings and Nonlinear Taxation." The Journal of Human Resources 22 (Summer 1987): 405-421.

O'Connell, Stephen A., and Zeldes, Stephen P. "Dynamic Efficiency in the Gifts Economy," Journal of Monetary Economics 31 (1993): 363-379.

Powers, Elizabeth T. "Does Means Testing Discourage Saving?: Evidence from the National Longitudinal Survey of Women." Mimeograph, Federal Reserve Bank of Cleveland, March 1994.

Rose, Robert L. "For Welfare Parents, Scrimping is Legal, But Saving is Out," The Wall Street Journal (February 6, 1990): pp. 1,11.

Sherraden, Michael. Assets and the Poor: A New American Welfare Policy. Armonk, New York: M.E. Sharpe, Inc., 1991.

Skinner, Jonathan. "Variable Lifespan and the Intertemporal Elasticity of Consumption." Review of Economics and Statistics 67 (November 1985): 616-623.

Skinner, Jonathan. "Risky Income, Life Cycle Consumption, and Precautionary Savings." Journal of Monetary Economics 22 (1988): 237-255.

Wolff, Edward N. "Estimate Of Household Wealth Inequality In The U.S., 1962-1983." Review of Income and Wealth 33 (1987): 231-256.

Yaari, Menahem E. "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer." Review of Economic Studies 32 (April 1965): 137-158.

Yitzhaki, Shlomo. "The Relation Between Return and Income," Quarterly Journal of Economics 52 (February 1987): 77-96.

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