NBER WORKING PAPER SERIES

SOCIAL SECURITY, BEQUESTS, AND THE LIFE CYCLE THEORY OF SAVING: CROSS-SECTIONAL TESTS

Alan S. Blinder

Roger Hall Gordon

Donald E. Wise

Working Paper No. 619

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge MA 02138

January 1981

The authors thank Michael Ransom for extensive research assistance, and John Bossons, Laurence Kotlikoff, W.H. Somermeyer, Lawrence Summers, Toshiaki Tachibanaki, and participants at the IEA Conference in Bergamo, Italy in June 1980 for comments on an earlier draft. Financial support from the U.S. Department of Labor and the National Science Foundation is gratefully acknowledged. The research reported here is part of the NBER's research programs in Economic Fluctuations and Social Insurance, and part of the NBER's research project in Capital Formation. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #619
January 1981

Social Security, Bequests, and the Life Cycle Theory of Saving: Cross-Sectional Tests

ABSTRACT

This paper studies the asset holdings of white American men near retirement age. Assets as conventional defined show no tendency to decline with age, in apparent contradiction of the life-cycle theory of saving. However, a broadened concept of assets which includes expected future pension benefits (both public and private) and expected future earnings ("human wealth") does decline more or less as predicted by the theory. No matter how they are defined, assets are a decreasing function of the number of children--which casts doubt on the strength of the bequest motive. Finally, financial assets and social security wealth fail to exhibit the inverse relationship suggested by Feldstein's displacement hypothesis.

To investigate these issues econometrically, an equation for assets is developed from the strict life-cycle theory. The specification is generalized to allow for (a) a bequest motive, proxied by the number of children; (b) displacement of private wealth by social security wealth that is not exactly dollar-for-dollar; (c) a level of consumption late in life that differs systematically from what the strict life-cycle theory implies.

The equation is estimated by nonlinear least squares on a rich cross-sectional data set containing over 4300 observations. The results show that the life-cycle model has little ability to explain cross-sectional variability in asset holdings. The model's key parameters are poorly identified, despite the large sample size and considerable cross-sectional variation in most variables. According to the estimates, consumption late in life is on average only about half of what the strict life-cycle theory predicts; each dollar of social security wealth displaces about 39¢ (with a large standard error) of private wealth; and the bequest motive, while present, is quite weak.

Alan S. Blinder
National Bureau of Economic Research
1050 Massachusetts Avenue
Cambridge, Massachusetts 02138

(617) 868-3929

Dr. Roger Hall Gordon
Bell Laboratories
600 Mountain Avenue
Murray Hill, New Jersey 07974

(212) 582-6472

Donald E. Wise
Mathtech
Post Office Box 2392
Princeton, New Jersey 08540

(609) 799-2245

SOCIAL SECURITY, BEQUESTS, AND THE LIFE CYCLE THEORY OF SAVING: CROSS-SECTIONAL TESTS

1. Introduction

This paper started out as an effort to use an unusually good source of cross-section data to address two current controversies about savings behavior:

(1) Does "social security wealth," that is, the actuarial present value of future social security benefits, displace ordinary assets in individual portfolios?

(2) Do people save mostly to finance retirement consumption (or, more generally, to shift consumption over time) or mostly to make intergenerational transfers?

The first question, which was raised originally by Feldstein (1974) and Munnell (1974), is of rather obvious importance for public policy toward the social security system, tax incentives for saving, and related issues. Most empirical work on the question to date has been based on time series data and, perhaps as an inevitable result, has been somewhat inconclusive. This is unfortunate, because the issue basically is a time series question rather than a cross-sectional one: once general-equilibrium adjustments have taken place, is total national saving reduced by the social security system? Cross-sectional evidence, however, is not irrelevant. If, as Feldstein (1974) claims, social security wealth does displace private fungible wealth in individual portfolios, this effect should show up at the micro level.[1] One

---

[1] A demonstration that social security wealth displaces private wealth at the micro level, however, does not establish that a similar displacement takes place at the macro level. Even if social security has no effect in the aggregate, those who receive unusually high benefits may save less while those who receive unusually low benefits may save more.

purpose of our research was to add to the rather limited supply of cross-sectional evidence on the social security displacement issue.[1]

An answer to the second question--that most saving is for intertemporal reallocation of consumption over the life cycle, not for bequests--was assumed in the original Modigliani - Brumberg (1954) paper, and has characterized most work on the life-cycle model ever since. This answer, of course, is not inherent in the model, which is easily extended to include bequests. Tobin (1967), we believe, was the first to ask whether the pure life-cycle theory with no bequests could account for aggregate savings in the U.S.. His answer was in the affirmative; but more recently White (1978) reached the opposite conclusion from a similar simulation model. Diamond's (1977) analysis of cross tabulations between assets and income led him to question whether the observed wealth holdings of older people would be sufficient for them to maintain their normal level of consumption--suggesting negative bequests (children supporting their elderly parents). But Mirer (1980) concluded that wealth holdings late in life do not decline as fast as they should if no bequests are intended. The most thorough study of this issue was by Kotlikoff and Summers (1980).[2] Using estimated life cycle patterns of consumption and earnings for individuals, they constructed an estimate of the wealth that those currently alive could have accumulated from their own savings and found that this amount was only a small fraction of total national wealth. The remainder presumably had been inherited--implying that an equivalent amount would have to be bequeathed in equilibrium.

---

[1] See Feldstein and Pellechio (1979), and Feldstein (1980).

[2] See also related work by Darby (1979) and Atkinson (1971).

The importance of bequests is relevant to several issues. Bequest behavior is one important determinant of intergenerational mobility in the distributions of income and wealth,[1] and plays a critical role in the debate over whether or not government bonds are net wealth.[2] Both the distributive and efficiency aspects of the choice between a consumption tax and an income tax depend, in part, on the importance of bequests in life-cycle saving. Since no concensus has emerged on the relative importance of bequests versus intra-life-cycle saving, a second major purpose of this paper was to see what could be learned from cross-sectional data.

Our original research strategy was to investigate these two questions within the framework of the life-cycle theory of saving, which has established itself as the preeminent theory of saving.[3] However, as the work progressed, our trust in the ability of the life cycle theory to organize the data progressively diminished. In the end, we learned less about the questions that originally motivated the study than we did about the limitations of the life-cycle theory itself. And it is the latter message that we now think is the paper's most important contribution.

The life-cycle theory is held up in macroeconomics as an exemplary piece of economic analysis. A model with sound theoretical foundations (thanks to Irving Fisher) was given empirical life (through the ingenuity of Ando, Brumberg, Friedman and Modigliani) and subsequently validated many times in empirical tests. However, most of these tests have been conducted

---

[1] Blinder (1976a), Shorrocks (1979), Menchik (1979).

[2] Barro (1974, 1976), Feldstein (1976)

[3] For our purposes, there is no important difference between the life-cycle theory of Modigliani and Brumberg (1954) and the permanent income theory of Friedman (1957).

on time series data which, as is well known, often have trouble in distinguishing among competing hypotheses.[1] To our knowledge, the precise implications of the strict form of the life cycle theory (defined specifically below) have never been tested. When we imposed the tight structure of the life-cycle model on the data, we found out that:

(1) The model explains very little of the cross-sectional variation in savings behavior. This, in itself, may not be a serious condemnation of the theory since it might still correctly isolate one of many influences on savings. However, we also found that:

(2) The critical parameters of the life-cycle model are very poorly identified, even in a large cross-sectional sample (over 4000 observations) with unusually good data on wealth.

(3) The data are consistent with the life-cycle theory only if it is assumed that people's utility functions shift systematically by age in such a way as to produce an optimal consumption stream with quite low consumption levels late in life.

Each of these is explained more fully later in the paper. Together they suggest to us that the empirical foundations of the life-cycle theory are far shakier than was previously thought. More specifically, while the data do not deliver a strong rejection of the life-cycle theory, they provide very little support for it.

The paper is organized as follows. The next section explains how we

---

[1] In time series, the characteristic implication of the life-cycle theory is that assets, in addition to income, influence consumption. For the permanent income version, as usually implemented, it is that lagged as well as current income influences consumption. It is not hard to think up entirely different models of consumption that have the same implications. Hall (1978) offers a more stringent test which, interestingly enough, supports the theory on time series data but rejects it in cross-section (Hall and Mishkin (1980)).

parameterized the strict form of the life-cycle theory, and then modified the specification to make it suitable for estimation. After Section 3 discusses the data source, Section 4 takes a naive impressionistic look at the data without trying to impose any theoretical structure at all. As we will see, this preliminary look at the data shows broad consistency with the life-cycle theory, does not suggest that social security displaces private financial assets, and does not suggest an important bequest motive for saving. This of course, raises a question of whether serious empirical modelling would reverse these rough impressions. The regression results reported in Section 5 turn out to be rather unfavorable to the life-cycle hypothesis. However, the other two expectations--that the model would suggest neither a quantitatively important bequest motive nor a strong displacement effect of social security wealth--are confirmed. Section 6 is a brief summary of the major conclusions.

2. The Life-Cycle Model: Theory and Empirical Implementation

2.1 The Pure Life-Cycle Theory

The standard life-cycle theory envisions an individual with a flow of earnings through time, E_t, and an initial endowment of wealth, A_0, choosing a consumption path, C_t, to maximize lifetime utility:

$$(2.1) \quad \sum_{t=0}^{T} \theta(t)U(C_t) + B(A_T)$$

where $U(C_t)$ is the one-period utility function, $\theta(t)$ is a weighting factor, and $B(A_T)$ is the utility of terminal assets. (Throughout the paper t is the individual's age and is assumed to run from 0 to T.) Normally $\theta(t)$ is specified as:

$$(2.2) \quad \theta(t) = \frac{1}{(1+\rho)^t} ,$$

and $U(C_t)$ and $B(A_t)$ are assumed to be isoelastic:

$$(2.3) \quad U(C_t) = \frac{C_t^{1-\delta}}{1-\delta} , \quad B(A_T) = b \frac{A_T^{1-\delta}}{1-\delta}$$

While these specific assumptions are not inherent in the principle that individuals formulate life-cycle plans, some such assumptions are necessary if the theory is to be made operational. If $\theta(t)$ is allowed to follow <u>any</u> arbitrary path, then no data can refute the theory. In this study we make the standard assumptions (2.2) and (2.3).

In a certain world with a perfect capital market, the only constraint on the maximization of lifetime utility is the lifetime budget constraint:

$$(2.4)\quad \sum_{t=0}^{T} \frac{C_t}{(1+r)^t} + \frac{A_T}{(1+r)^T} = A_0 + \sum_{t=0}^{T} \frac{E_t}{(1+r)^t} \equiv A_0 + Y_0,$$

where r is the rate of interest (assumed constant over time) and Y_0 is lifetime discounted earnings, expressed in the dollars of t=0. Maximizing (2.1) with respect to (2..4) under the specific assumptions (2.2) and (2.3) leads as is well known to:

$$(2.5)\quad C_t = C_0(1+g)^t$$

$$(2.6)\quad A_T = \hat{\beta}\, C_0$$

where g is positive if the rate of interest exceeds the rate of time discounting and negative otherwise,[1] and where $\hat{\beta}$ is a taste parameter.[2]

To adapt this to the case of a family consisting of N_t adult equivalents when the head is of age t, it is convenient to assume that family utility is:

$$N_t U(C_t/N_t)$$

where C_t now stands for <u>family</u> consumption. In words, family utility is the utility of the family's consumption per adult equivalent multiplied by the number of adult equivalents. In this case, it is readily shown that (2.5) and

---

[1] Specifically, $1+g \equiv \left(\frac{1+r}{1+\rho}\right)^{\frac{1}{\delta}}$

[2] Specifically, $\beta = (b)^{\frac{1}{\delta}}\ (1+r)^{\frac{T}{\delta}}$.

(2.6) are replaced by:

$$(2.7) \quad C_t = \left(\frac{N_t}{N_0}\right) C_0(1+g)^t$$

$$(2.8) \quad A_T = \hat{\beta}\left(\frac{C_0}{N_0}\right).$$

Almost no cross-sectional data set has good data on consumption by age; ours is no exception. To move toward a testable equation, we must draw out the implications of the model for current holdings of assets A_t. Let Y_t be the present discounted value (in time t dollars) of earnings from time t forward, viz.:

$$Y_t = \sum_{s=t}^{T} \frac{E_s}{(1+r)^{s-t}} ,$$

where T, the length of life, is assumed for the present to be known. The budget constraint from time t forward implies that the sum $A_t + Y_t$ must be equal to the discounted present value of future consumption plus the planned bequest. Thus:

$$A_t + Y_t = \sum_{s=t}^{T} \frac{C_s}{(1+r)^{s-t}} + \frac{A_T}{(1+r)^{T-t}} .$$

Using (2.7) and (2.8) this can be written:

$$A_t + Y_t = \sum_{s=t}^{T} \left(\frac{C_0}{N_0}\right) N_s \frac{(1+g)^s}{(1+r)^{s-t}} + \frac{\hat{\beta}}{(1+r)^{T-t}} \frac{C_0}{N_0}$$

or

$$(2.9) \qquad A_t + Y_t = (1+r)^t \frac{C_0}{N_0} \left[\sum_{s=t}^{T} (1+\mu)^s N_s + \beta \right]$$

where $1 + \mu \equiv \frac{1+g}{1+r}$ and $\beta \equiv \hat{\beta}(1+r)^{-T}$. By setting t = 0 in (2.9) we derive the lifetime budget constraint:

$$(2.10) \qquad A_0 + Y_0 = \frac{C_0}{N_0} \left[\sum_{s=0}^{T} (1+\mu)^s N_s + \beta \right] ,$$

which can be used to solve for C_0/N_0. Substituting the result back into (2.9) gives the basic life-cycle equation that we would like to estimate:

$$(2.11) \qquad A_t + Y_t = \frac{\left[\sum_{s=t}^{T} (1+\mu)^s N_s + \beta \right]}{\left[\sum_{s=0}^{T} (1+\mu)^s N_s + \beta \right]} (\hat{A}_0 + \hat{Y}_0)$$

where $\hat{A}_0 + \hat{Y}_0$ is initial (nonhuman plus human) wealth expressed in time t dollars, that is, $\hat{A}_0 + \hat{Y}_0 = (A_0 + Y_0)(1 + r)^t$.

The ratio of the two sums in equation (2.11) has a straightforward intuitive interpretation. The denominator is the number of adult equivalent years of consumption (properly discounted, and embodying any desired trend in consumption) in the family's entire life cycle, including an allowance for planned bequests. The numerator is similarly interpreted as the number of adult equivalent years of consumption still remaining when the head is age t. The equation then states that the fraction of total lifetime resources still available, $\frac{A_t + Y_t}{\hat{A}_0 + \hat{Y}_0}$, is equal to the fraction of adult equivalent years of life still remaining.

Notice two features of this specification. First, the units for β are years of adult equivalent consumption. Thus an estimated β of 5, say, would imply that the average family leaves a bequest equivalent to five adult-years of consumption. Second, (2.11) is <u>very</u> tightly parameterized according to the life-cycle theory. Apart from whatever parameters constitute β (which we specify below), there is only <u>one</u> parameter to be estimated in (2.11). In particular, the dependence of asset holdings on age follows the very specific (and highly nonlinear) functional form dictated by the strict form of the life-cycle theory. This is very different, e.g., from regressing assets on age and age squared.

We can make the specification a little more flexible, and at the same time provide a test of the rational planning calculation inherent in the life-cycle model, by adding a parameter γ in front of the sum in the numerator of (2.11). That is:

$$(2.12) \qquad A_t + Y_t = \frac{\left[\gamma \sum_{s=t}^{T} (1+\mu)^s N_s + \beta\right]}{\left[\sum_{s=0}^{T} (1+\mu)^s N_s + \beta\right]} (\hat{A}_0 + \hat{Y}_0) .$$

If people arrive at age t (which in our sample ranges from 60 to 65) with the level of assets implied by the life-cycle theory, then γ should be equal to unity, as in (2.11). If, however, people systematically underprovide for their consumption at older ages, then γ will be less than unity. Thus testing whether γ differs from unity is a way of testing the life-cycle model.[1]

2.2 Empirical Implementation

Several additional assumptions were required before (2.12) could be confronted with data.

Uncertainty and Family Composition

The theoretical model just sketched assumes that the length of life, T, and the size of the family at each t, N_t, are known with certainty from the beginning. To point out that such knowledge is not available is trite. Knowing what to do about this uncertainty is less obvious, since a full blown theoretical development of the model under uncertainty is very difficult.

---

[1] An alternative interpretation of $\gamma < 1$, which we do not find very satisfying is that θ_t shifts systematically with age in such a way that average consumption at ages beyond 60-65 is lower than $C_0(1+g)^t$. As stated earlier, if arbitrary changes in θ_t over time are permitted, any data are consistent with the life-cycle theory.

We have therefore adopted several simplifying assumptions.

First, we assume that individuals have access to a competitive annuities market--either directly or through pension plans (including, importantly, public pensions through social security). As Yaari (1965) has shown, the optimal life-cycle plan in such circumstances is basically the same as in a world of certain lifetimes: the individual just needs to provide for his expected lifetime consumption by purchasing the requisite annuities.[1]

Second, we assume that future wage income and the dates of marriage and having children are known with certainty at the outset. Uncertainty over future income and family size would probably raise, rather than lower, saving. So if we find (as we do) that savings are lower than called for by the life-cycle theory with certainty, then they must be lower than what an uncertainty model would predict.

Under these assumptions, it makes sense to adopt the following empirical proxies for the variables N_s in the theoretical model. Let Δ denote the husband's age minus the wife's age, and define:

$P_m(s,t)$ = probability that a male of age t survives to age s.

$P_f(s,t)$ = probability that a female of age t survives to age s.

Then our empirical proxy for N_s in the denominator of (2.12) was:

$$N_s = \begin{cases} 2 \text{ with probability } P_m(s,17)P_f(s-\Delta,17) \text{ (both partners living)} \\ 1 \text{ with probability } [1-P_m(s,17)]P_f(s-\Delta,17) + P_m(s,17)[1-P_f(s-\Delta,17)] \\ \quad \text{(one partner living)} \\ 0 \text{ with probability } [1-P_m(s,17)]\,[1-P_f(s-\Delta,17)] \text{ (neither partner living)} \end{cases}$$

---

[1] Unfortunately, the requisite annuities are indexed annuities, which normally are not available apart from social security.

where survival probabilities are drawn from standard life tables.[1] Similarly, we replaced N_s in the numerator by:

$$N_s = \begin{cases} 2 \text{ with probability } P_m(s,t)P_f(s-\Delta,t-\Delta) \\ 1 \text{ with probability } [1-P_m(s,t)]P_f(s-\Delta,t-\Delta) + P_m(s,t)[1-P_f(s-\Delta,t-\Delta)] \\ 0 \text{ with probability } [1-P_m(s,t)]\ [1-P_f(s-\Delta,t-\Delta)]. \end{cases}$$

Note the updating of the survival probabilities to ages t and t-Δ.

Children had to be treated differently since, except for those still supported by their parents (a rarity in a sample where fathers ranged in age from 60 to 65 years), we did not know how many years of support each child had claimed. We decided arbitrarily that each child had been supported for 18 years, and that this support cost the family 18α adult equivalent years of consumption, where α is a parameter to be estimated. Thus, if we define:

NKIDS = number of children ever born

NSUP = sum of (18-age), summed over any children still supported,

our approach was to add α_1NSUP to the numerator of (2.12) and $18\alpha_2$NKIDS to the denominator. We allowed α_1 to differ from α_2 on the grounds that, if intergenerational transfers are made inter vivos, adult sons and daughters might have received some of their inheritances but minor children would not have. This suggests that α_2 might exceed α_1.

Taste for Bequests

An obvious determinant of the taste for bequests, and the only one we considered, is the number of children ever born. We therefore specified that

[1] Age 17 is assumed to be the age of economic adulthood. We assumed marriages took place when the younger partner was 17; for years in which only one partner was over 17, N_s was therefore set equal to 1, not 2.

β in (2.12) was given by:

$$\beta = b_0[b_1 + b_2\text{NKIDS} + b_3(\text{NKIDS})^2].$$

Thus, we look for a bequest motive by looking for an association between asset holdings late in life and number of children.

The seemingly redundant parameter b_0 was included originally to allow for possible nonlinearities in bequest behavior (i.e., a wealth elasticity of bequests different from unity). To this end, we specified that:

$$b_0 = 1 + b_4(\hat{A}_0 + \hat{Y}_0),$$

so that a negative b_4 would signify a wealth elasticity of bequests below one and a positive b_4 would signify a wealth elasticity above one.[1] In the strict life-cycle model, b_4 is zero (i.e., indifference curves are homothetic).

Finally, lacking any information on inheritances received ($\hat{A}_0$ in equation (2.12)) we simply assumed $A_0 = 0$ for everyone in the sample. If A_0 were typically positive, our equation should systematically understate asset holdings. In fact, it overstates them.

After all these modifications, (2.12) became:

$$(2.13) \quad A_t + Y_t = \left[\frac{\gamma P_t + \alpha_1 \text{NSUP} + \beta}{P_{17} + 18\alpha_2 \text{NKIDS} + \beta} \right] \hat{Y}_0$$

where

$$(2.14) \qquad P_t = \sum_{s=t}^{110+\Delta} (1+\mu)^s (2P_m(s,t)P_f(s-\Delta,t-\Delta)$$

$$+ [1-P_m(s,t)]P_f(s-\Delta,t-\Delta) + P_m(s,t)[1-P_f(s-\Delta,t-\Delta))$$

$$(2.15) \qquad \beta = (1 + b_4\hat{Y}_0)[b_1 + b_2\text{NKIDS} + b_3(\text{NKIDS})^2] .$$

Notice that apart from the parameters γ and b_4, this equation adheres scrupulously to the strict parameterization implied by the life-cycle theory. It has a very specific (and highly nonlinear) structure, and includes only eight parameters.

So far, however, we have said nothing about one of the critical variables

---

[1] For reasons why this is an important parameter, see Blinder (1976b), or Menchik and David (1979).

of our study: social security wealth (henceforth, SSW). As written in (2.13), the model assumes the validity of Feldstein's hypothesis that SSW displaces private fungible wealth on a dollar-for-dollar basis, that is, if total lifetime resources are held constant, (2.13) implies that:

$$\frac{\partial A_t'}{\partial SSW_t} = -1,$$

where A_t' is financial assets. But there are a number of reasons why this might not be true: social security provides an annuity that cannot be bequeathed, SSW is an illiquid asset that is not fungible, SSW cannot be used as collateral on loans, social security benefits are indexed, and so on. For similar reasons and some additional ones (e.g., lack of vesting, worries about the financial reliability of the employer), private pension wealth may also be a less-than-perfect substitute for other wealth (and also for SSW). Similarly, there are many reasons why the net asset value of one's house may not be a perfect substitute for financial assets. To allow for these phenomena, we added the following three parameters to the model:

$$\lambda_1 \equiv -\frac{\partial A_t'}{\partial SSW_t}, \quad \lambda_2 \equiv -\frac{\partial A_t'}{\partial PPW_t}, \quad \lambda_3 \equiv -\frac{\partial A_t'}{\partial RE_t}$$

where SSW_t is the actuarial discounted value (in time t dollars) of the future social security benefits to which the individual is entitled;[1] PPW_t is private pension wealth, the analogous concept for private pension benefits; and RE_t is the net value of real estate owned. This led us to amend (2.13) as follows:

$$(2.16)\quad A_t' + \lambda_1 SSW_t + \lambda_2 PPW_t + \lambda_3 RE_t + Y_t = \left[\frac{\gamma P_t + \alpha_1 NSUP_t + \beta}{P_{17} + 18\alpha_2 NKIDS + \beta}\right](\hat{Y}_0 + \lambda_1 SSW_t + \lambda_2 PPW_t + \lambda_3 RE_t) .$$

Notice that the concept of lifetime resources on the righthand side of (2.16)

---

[1] Our SSW measure is gross social security wealth since payroll taxes are deducted before computing earnings.

is augmented in the same way as the assets variable on the lefthand side--each different type of asset gets its own λ. Equation (2.16), along with the definitions in (2.14) and (2.15), constitutes the model we originally set out to estimate.[1] Unfortunately, this task turned out to be surprisingly difficult.

3. Data and Construction of Variables

3.1 Data Source and Sample Selection

The basic sampling frame was white men other than the self-employed in the 1971 wave of the Longitudinal Retirement History Survey (LRHS). This survey began with 11,153 individuals (both sexes, all races) between the ages of 58 and 63 in 1969, and reinterviewed those who survived and could be located at two-year intervals through 1977.[2] Detailed information on asset holdings, earnings histories, pensions, and various socio-economic traits was collected from each individual. Though our assets data came from the 1971 wave, information from the 1969, 1973, and 1975 interviews and from the related social security earnings histories was used in constructing several of the variables.

The potential sample of over 6,700 white men still living in 1971 was pared to only 4,130 observations for a variety of reasons:

(1) When the asset data was incomplete, the observation was dropped.

(2) All individuals who reported having an investment in a business were dropped. Due to some unfortunate wording in the questionnaire, the reported value of this investment was unreliable.

(3) Lifetime earnings were constructed for each individual and his spouse by a procedure sketched below and described in detail in the appendix.

---

[1] This is not quite true. When we started the work we had more confidence in the life-cycle theory, and hence constrained $\gamma=1$. We also constrained $\alpha_1=\alpha_2$ at first.

[2] Survival is assumed random, so that no selectivity problems arise.

When respondents did not supply enough information to construct it, they were dropped from the sample.[1]

(4) When the husband was reported to be 20 or more years older than the wife, or more than 15 years younger, we dropped the observation. Most of these cases seem to be clear errors in the data.[2]

(5) For a few cases, the reported value of the house net of the mortgage was negative. Skeptical about such a possibility, we dropped these cases.

3.2 Lifetime Earnings Data

The variable $\hat{Y}_0$ in equation (2.16) is the actuarial discounted present value from age seventeen forward (converted to 1971 dollars) of the earnings of both the husband and the wife. The variable Y_t is the same variable, but including only earnings from age t forward. To construct these variables, it was necessary to generate for each husband and wife a lifetime profile of both wages and hours. Since this is a complex calculation, we sketch it here and relegate the details to an appendix.

Wage Rates of the Husband

Our procedure for generating a wage profile for men began by gathering all possible observations on past wages available in the LRHS surveys. This provides us with at most four wage observations: the wage reported for the first job, the wage reported for the "last job" as of the 1969 interview, and

---

[1] We also dropped a few observations where the various reported wage rates were dramatically at odds with one another suggesting an error on the original data tape.

[2] Our model of asset holdings depends heavily on the ages of the couple, so would be very sensitive to reporting errors in these ages.

the wages reported for the "current jobs" in the 1969 and 1971 interviews.[1] The gaps between observed wage rates were "filled in" by using estimates derived from a wage equation discussed in Gordon and Blinder (1980), as explained in the appendix.

<u>Hours of Work for the Husband</u>

There is a fairly obvious endogeneity problem between assets and hours of work late in life. Differences in assets can cause differences in labor supply and/or retirement age. To purge the data of this potential endogeneity, we used an *estimated* labor supply profile, rather than the *actual* profile observed in the data. From an econometric point of view, this can be thought of as using an instrument for $\hat{Y}_0$.

Specifically, we assumed that typical hours are always 2,000 per year until retirement and zero thereafter,[2] which reduces the problem of constructing a typical hours profile to one of estimating the probability of being retired at each age. For this purpose, we used the retirement probabilities derived and presented in Gordon and Blinder (1980) to generate for each individual a predicted probability of being retired at each age between 58 and 85 based on his wage rate and personal characteristics. Retirement before age 58 was assumed never to occur, and every individual was assumed retired after age 85 even if our equations gave them some small probability of still being at work. Hours for individual i in year t were set equal to $2{,}000(1-p_{it})$, where p_{it} is the estimated probability that individual i is retired in year t.

Having thus estimated both wages and hours in each year of life,

---

1 We had in addition wage information on the current jobs reported in 1973 and 1975, but did not use these data because of potential endogeneity between current assets in 1971 and future labor supply decisions.

2 This may seem an extreme assumption, and perhaps it is. However, evidence presented in Gordon and Blinder (1980) suggests that it is a tolerably accurate description of the labor-supply behavior of men in this sample.

the implied earnings for each year were discounted or accumulated to 1971 using a 2% real interest rate to arrive at $\hat{Y}_0$ for the husband. Predicted wages and hours from 1971 forward were used to construct Y_t in precisely the same way. Table 1 reports the means and standard deviations of both of these constructed variables. In terms of expected discounted present value, these men have on average less than 6% of their lifetime earnings still ahead of them. Notice also that while Y_t is very small relative to Y_0, it is much more widely dispersed. (The coefficient of variation of Y_0 is 0.9, while the coefficient of variation of Y_t is 1.8.) This is as expected because of individual differences in retirement decisions.

Earnings of the Wife

Most of the men in the sample were married. The presence of a wife is relevant both as an additional claim on the resources of the family (as explained earlier) and as an additional contributor to the family's resources through her own earnings. Lifetime earnings measures for wives were constructed in much the same way as for husbands. However, a few extra problems arose, and are explained in the appendix.

Table 1 shows that the computed lifetime earnings figures for women are on average about 10% of the corresponding figures for men, and are also much more variable. Both of these are as expected, given the spotty work histories of women of this age cohort. It is also not surprising that the wives have a much larger percentage of their total lifetime earnings still ahead of them--over 13% versus less than 6% for husbands.

The variable $\hat{Y}_0$ used in the regression consisted of the sum of the estimated lifetime earnings of the husband and wife. Its mean is about \$482,000.

It is worth comparing the numbers in Table 1 with the corresponding figures from a recent paper by Feldstein (1980), which uses the 1969 wave of the LRHS,

Table 1

Selected Data on Lifetime Earnings
(in 1971 dollars)

Husbands	Mean	Std. Deviation
Y_0	$ 445,514	$ 493,057
Y_t	$ 24,207	$ 48,183
Wives		
Y_0	43,666	64,569
Y_t	5,772	13,521

since a sharp difference emerges. The mean of lifetime earnings in Feldstein's sample is only $244,566, barely more than half of ours. That Feldstein uses 1969 dollars while we use 1971 dollars requires some small adjustment to put his numbers on an equal footing with ours, and his numbers should also differ on account of excluding single men (which should make his mean higher) and government employees. But we imagine that the largest source of difference is that Feldstein's estimates of lifetime earnings were derived from social security records. The earnings figures he uses therefore will be systematically too low because:

(1) Earnings prior to 1936 are excluded, and people in this sample were already 25-30 years old by 1936.

(2) Social security records include only covered earnings, and coverage of the system was far from complete in its early years. For example, in 1939 and 1955 only 55% and 76% respectively of all civilian employees were in covered employment.[1]

(3) Only the simple sum of earnings between 1936 and 1950 appears in the social security records. By necessity, the accummulated (at compound interest) earnings over this 15-year period must exceed this simple sum.

3.3 Assets Data

The LRHS offers unusually detailed information on asset holdings by type of asset. The variable A_t' in equation (2.16) was constructed as the sum of the following components of net worth (each expressed in 1971 dollars):[2]

1. Net financial assets: the sum of all bank account balances, plus the market value of any stocks, bonds, or life insurance, minus any outstanding

---

[1] See Social Security Bulletin, Statistical Supplement, 1975, Table 7, page 44.

[2] Throughout the study, whenever it was necessary to discount or accumulate sums at compound interest we used a 6% nominal or a 2% real rate of interest.

loans. The market value of life insurance was estimated as the reported face value of the policy multiplied by $\sum_{s=t}^{100} \frac{p_m(s,t)}{(1+r)^{s-t}}$.

2. Net value of real estate: the market value of an owner-occupied house or any other real estate owned, minus the outstanding mortgage balance.

The other two assets considered in the study were:

3. Social Security Wealth: the actuarial discounted present value of future social security benefits. This is an important asset for older people, and earlier studies have been plagued by the absence of a reliable estimate of its value.[1] To estimate an individual's social security benefits at retirement, we applied the social security benefit formula to the earnings history that we estimated for him using the procedure described in the appendix. Estimated rather than actual earnings were used to avoid any endogeneity problems between hours of work and assets. However, in any year for which an individual's covered earnings were much smaller than our estimated earnings, we assumed his main job was not covered by social security, and set his covered earnings to zero. The individual's yearly benefits, assumed to be indexed to prices and to grow in real terms by 2% per annum, were then actuarially discounted back to the present to arrive at our measure of social security wealth. The effect of using estimated rather than actual, social security benefits was to change the mean of SSW very little, but to reduce its standard deviation by about 25%.

4. Private Pension Wealth: a similarly constructed measure of the actuarial present value of expected future benefits from private and/or government pensions. For individuals reporting their expected pension benefits, we treated the reported flow as a fixed nominal amount (unlike social security benefits, which are indexed), and discounted by the nominal interest rate. For people

---

1 See, for example, Munnell (1976), Feldstein and Pellechio (1979).

reporting a pension but not knowing the amount of their future benefits, we made an imputation based on regressing pension benefits on wages for those who reported their future benefits.

Even before any statistical analysis is performed, it is interesting to look at the amount of assets of various types held by people in our sample.[1] Table 2 shows the means and standard deviations for each category of net worth. A few striking facts emerge immediately. First, social security wealth is by far the most important source of wealth for the average person in the sample. Second, however, in accounting for the variance of wealth across people, SSW is probably of minor importance because it is so equally distributed as compared with any of the other sources of wealth. Third, mean holdings of financial wealth--which here include the market value of life insurance--look so small compared to lifetime resources (whose mean is over $500,000) as to suggest irrationally low savings. However, total nonhuman wealth is about four times as large as financial wealth--about $70,000 for the average family, or one-seventh of total lifetime resources. In addition, most of these families still have some unused "human wealth" (discounted present value of expected future earnings) as well--amounting on average to about 6% of lifetime resources according to Table 1. It is far from obvious that reserving one-fifth of lifetime resources for use after age 60-65 (including any planned bequests) constitutes "too little" saving.

Once again, comparisons with Feldstein's (1980) tabulations with the same data source are instructive. The per capita (or per couple) wealth totals are very similar (his is roughly $69,000) but the distribution across components

---

[1] Our data set does not report information on the value of consumer durables, hence this asset is ommitted.

Table 2

Selected Data on Wealth Holdings, 1971

Asset Type	Mean	Std. Deviation
Financial	$17, 159	$36, 116
Real Estate	15, 888	26, 662
Social Security	34, 527	8, 270
Pensions	6, 383	12, 970

is quite different. His estimated social security wealth per couple--about \$45,000--is far higher than ours, mainly because Feldstein excluded government employees, whose SSW is zero. Offsetting differences appear in nonpension wealth, where Feldstein's mean is \$23,682 and our sum of mean financial wealth plus mean real estate is \$33,062, even though we eliminated from the sample all business assets. Finally, Feldstein does not attempt to estimate private pension wealth for the individuals in his sample.

4. Determinants of Asset Accumulation: A Preliminary Look

Before proceeding to explain the difficulties that arose in attempting to estimate (2.16) and to present the estimates, it is worth pausing to consider what we are looking for in the data. Such an initial look will teach us much that maximizing a likelihood function with 8 parameters and 4130 observations conceals.

Our version of the life-cycle theory makes asset accumulation depend essentially on the ages of the husband and the wife and on the number of children (which we use as a proxy for the bequest motive). The theory implies a very particular rate of decline with age of a comprehensive concept of wealth, $A_t + Y_t$. Specifically, denote the ratio of $A_t + Y_t$ to $\hat{A}_0 + \hat{Y}_0$ by the symbol R_t. Then, in the case $\gamma = 1$, the theory implies:

$$\frac{R_{t+1}}{R_t} = \frac{P_{t+1} + \beta}{P_t + \beta}$$

when there are no dependent children in the household.[1] If the parameter μ is close to zero (as it should be), and both husband and wife are alive both in

---

[1] The definition of P_t is provided in equation (2.14).

period t+1 and period t, then $P_{t+1} = P_t - 2$, so the rate of decline of R_t is:

$$(4.1) \quad \frac{R_{t+1} - R_t}{R_t} = \frac{-2}{P_t + \beta} \quad .$$

In our sample, a value of P_t near 30 is about average, so (4.1) implies that R_t should fall about 6.7% per year if $\beta = 0$ or about 5% per year if $\beta = 10$.

Do the data display this predicted behavior? On the surface, they seem not to, as Figure 1 indicates. The bottom line in Figure 1 plots the age profile of financial assets, A_t', relative to lifetime income. There is not a hint that this ratio declines with age. However, it is worth stressing that the theory predicts a declining age profile only of a much more comprehensive measure of wealth: $A_t' + RE_t + SSW_t + PPW_t + Y_t$. Since the last three items on this list must decline with age, we might see a decline in a broader wealth concept that does not appear in A_t'. The remaining lines in Figure 1 investigate successively more comprehensive definitions of wealth by adding additional assets to A_t', more or less in order of liquidity. Thus, the second line adds the value of real estate to financial assets; still there is no trace of a negative slope. When pension assets (both SSW and PPW)--which, by their method of construction are virtually guaranteed to decline with age[1]--are added to wealth, a slight tendency toward a declining wealth profile does emerge. (The ratio of all nonhuman assets to lifetime income drops only 8% between ages 60 and 65, and is almost the same at age 64 as it is at age 61.) However, when the present value of future earnings is added to nonhuman assets, a very pronounced declining age profile appears. On average, the ratio of (human plus nonhuman) assets to lifetime income falls by 6.3% per year between ages 60 and 65 for the people in our sample. As noted, this is broadly consistent with the prediction embodied in (4.1).

---

[1] In calculating SSW and PPW, we assumed individuals initially claimed benefits at age 65, so that no one in our sample would yet have drawn on their benefits. We thus underestimate any decline in the value of this component of wealth with age.

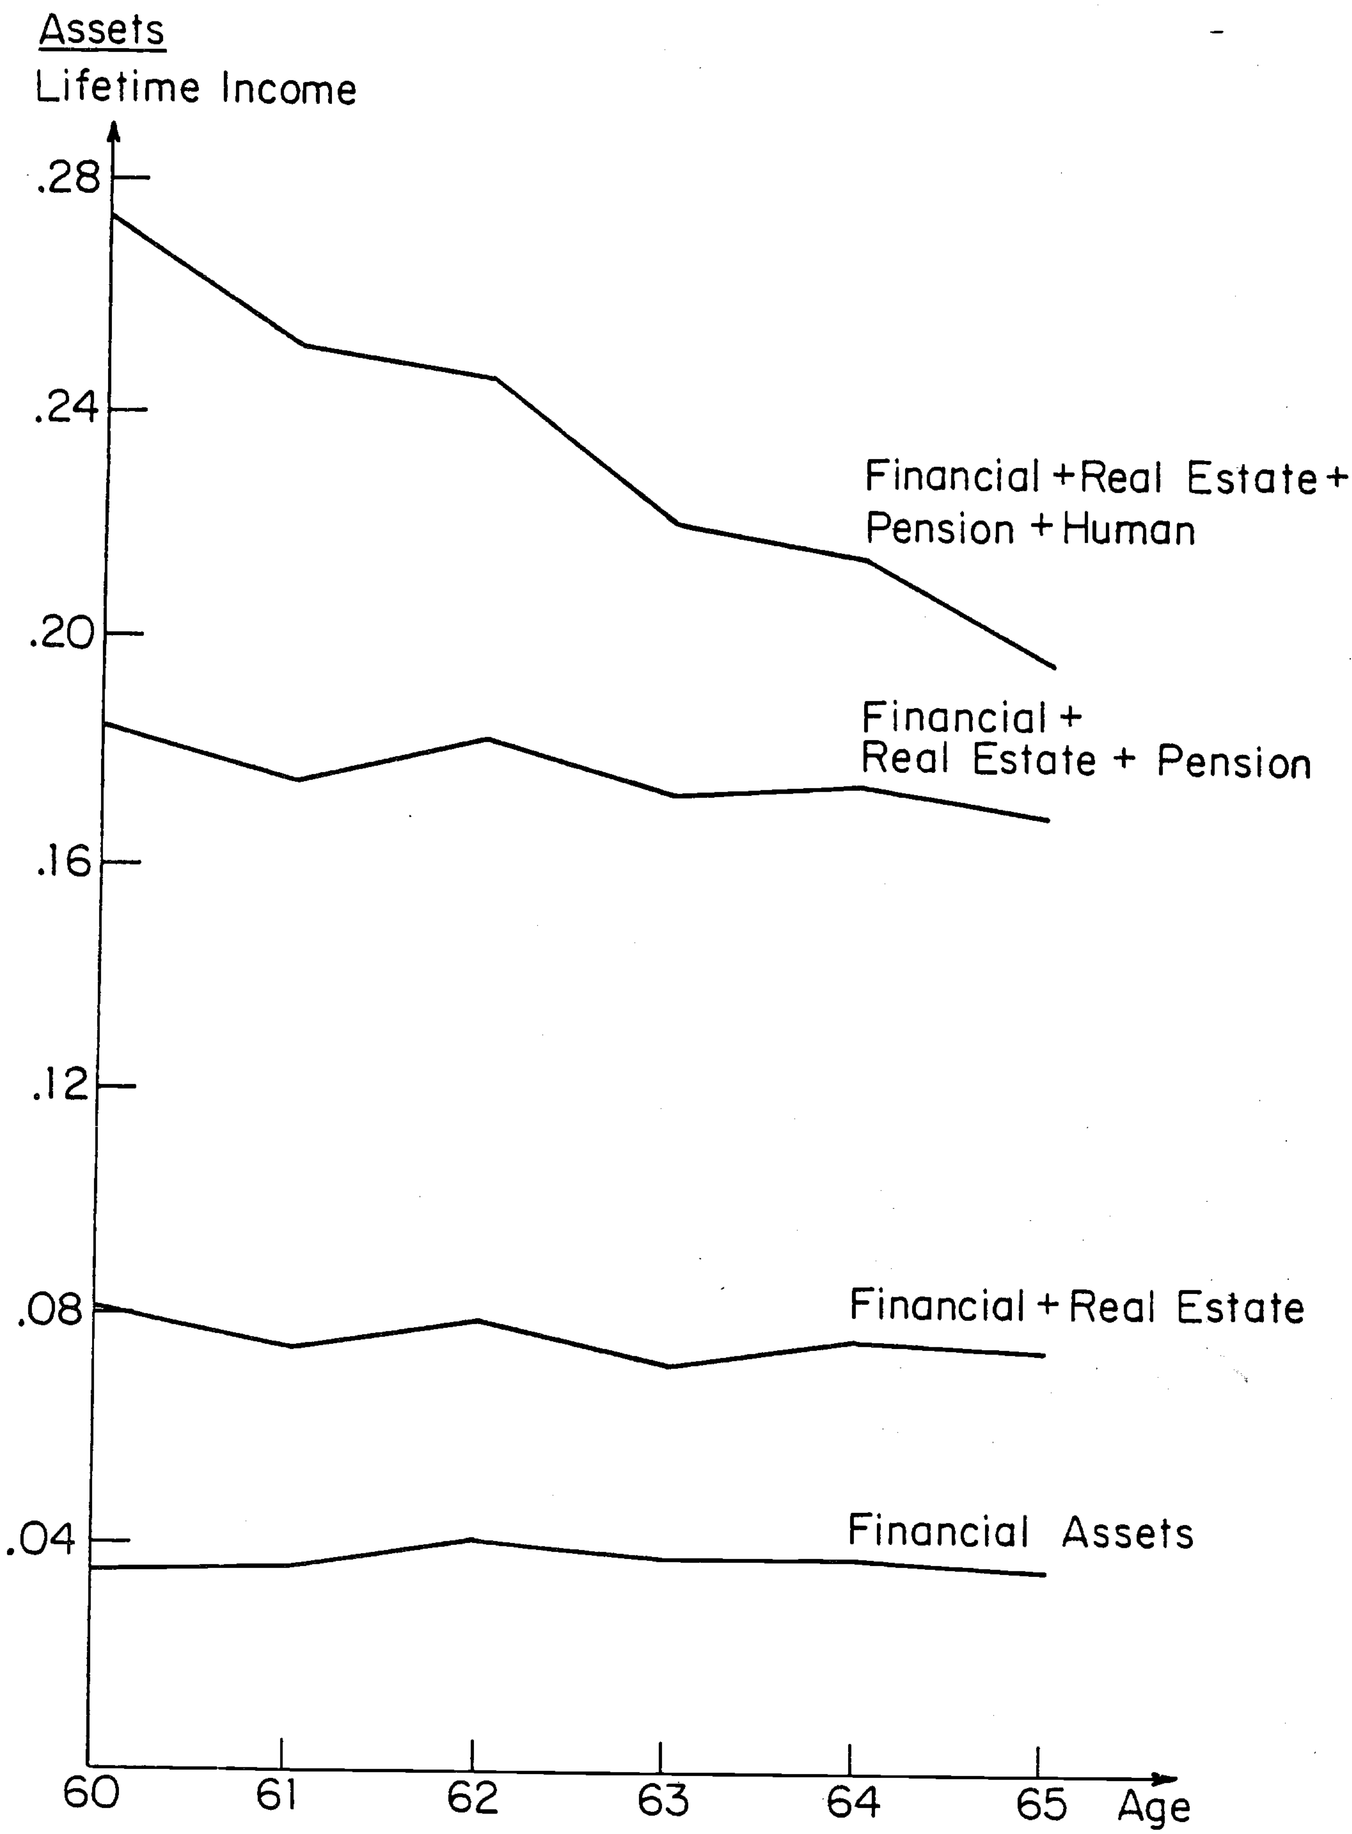
Assets
Lifetime Income
.28
.24
.20
.16
.12
.08
.04
Financial + Real Estate + Pension + Human
Financial + Real Estate + Pension
Financial + Real Estate
Financial Assets
60
61
62
63
64
65
Age

FIGURE 1

In sum, if a broadly-defined concept of wealth that includes expected future earnings is used to measure current assets, the data do show the sort of pattern prescribed by life-cycle theory. In terms of the specific model in equation (2.16), then, there is nothing in this crude look at the data to warn us that the estimated γ will stray far from its theoretical value of unity.[1]

The second thing we are looking for in these data is a systematic dependence of assets on the number of children, either increasing with children due to bequest, or decreasing due to the expenses of child rearing and/or intergenerational transfers given inter vivos rather than at death. Figure 2 plots the same four concepts of wealth (all normalized by lifetime income) against the number of children.[2] A clear relationship does emerge: with few exceptions, assets (no matter how defined) are decreasing in the number of children. Now this, of course, does not disprove the existence of a bequest motive. But it does indicate that the factors leading assets to decline with the number of children dominate any bequest motive that may exist. Our statistical specification in (2.16) is designed precisely to separate out the effects of the bequest motive (which is captured by the b's) from the effects of child-rearing expenses (which are captured by the α's). But even before estimation, Figure 2 should give pause to believers in a strong bequest motive. Of the assets considered here, financial assets and real estate seem most suited for bequests, yet the pattern of decline of $A_t' + RE_t$ with the number of children is marked.

Let us next turn the glare of crude empiricism on the hypothesis suggested by Feldstein: that families with a high ratio of SSW to lifetime income should have a correspondingly low ratio of A_t' to lifetime income. Figure 3 plots

---

[1] When the parameter γ is introduced, (4.1) becomes: $\frac{R_{t+1} - R_t}{R_t} = -\frac{2\gamma}{\gamma p_t + \beta}$ so an estimated γ well below unity would imply that assets were not falling "fast enough."

[2] The graph is truncated at 12 children even though a few families in our sample had more because the tiny sample sizes lead to erratic behavior in the upper tail.

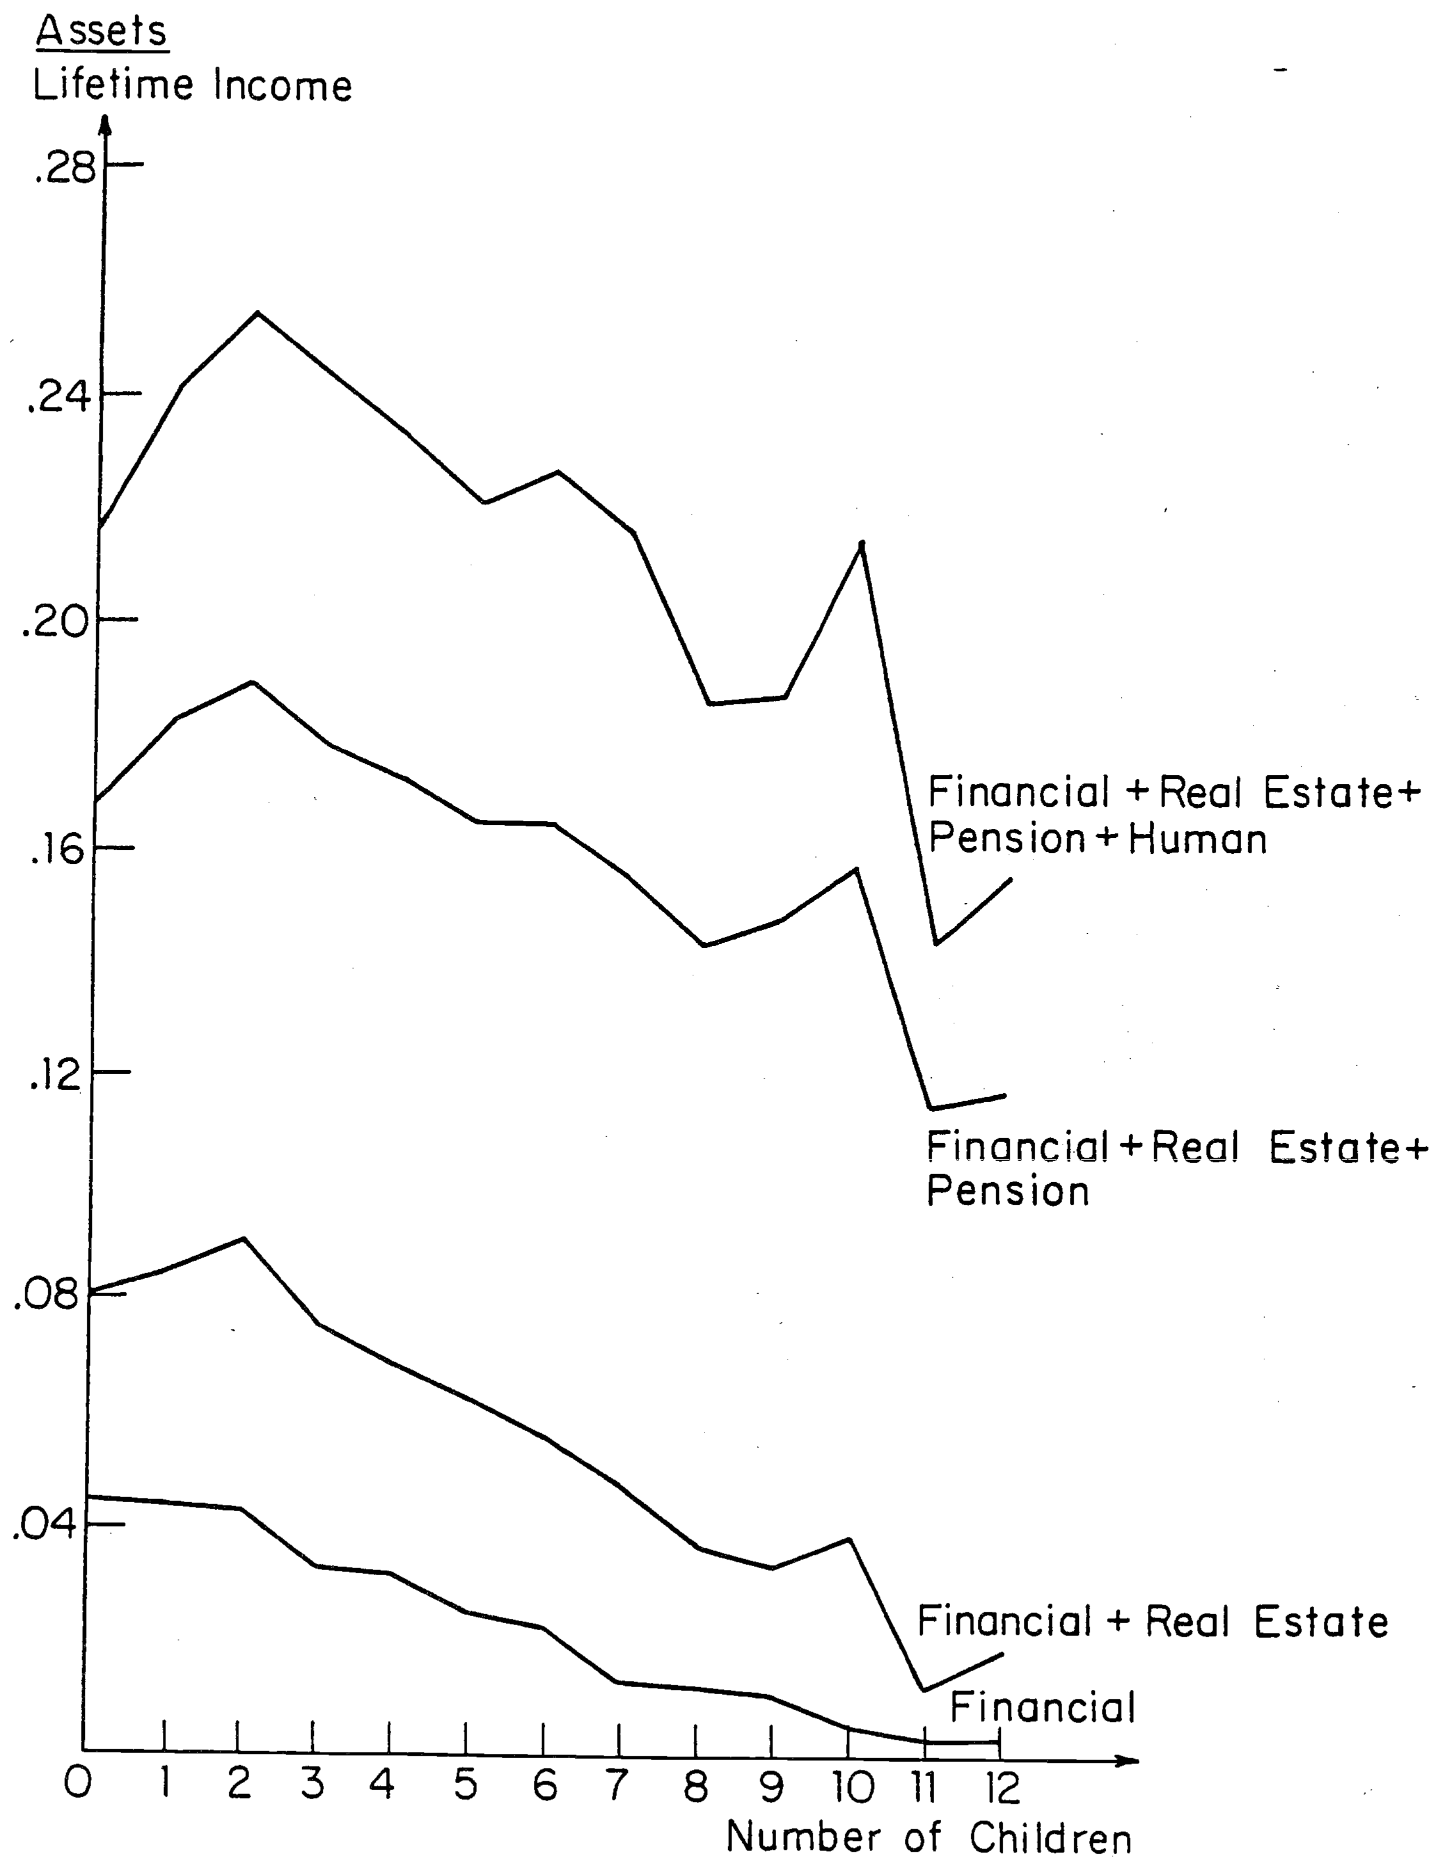
Assets
Lifetime Income
.28
.24
.20
.16
.12
.08
.04
Financial + Real Estate+
Pension + Human
Financial + Real Estate+
Pension
Financial + Real Estate
Financial
0 1 2 3 4 5 6 7 8 9 10 11 12
Number of Children

FIGURE 2

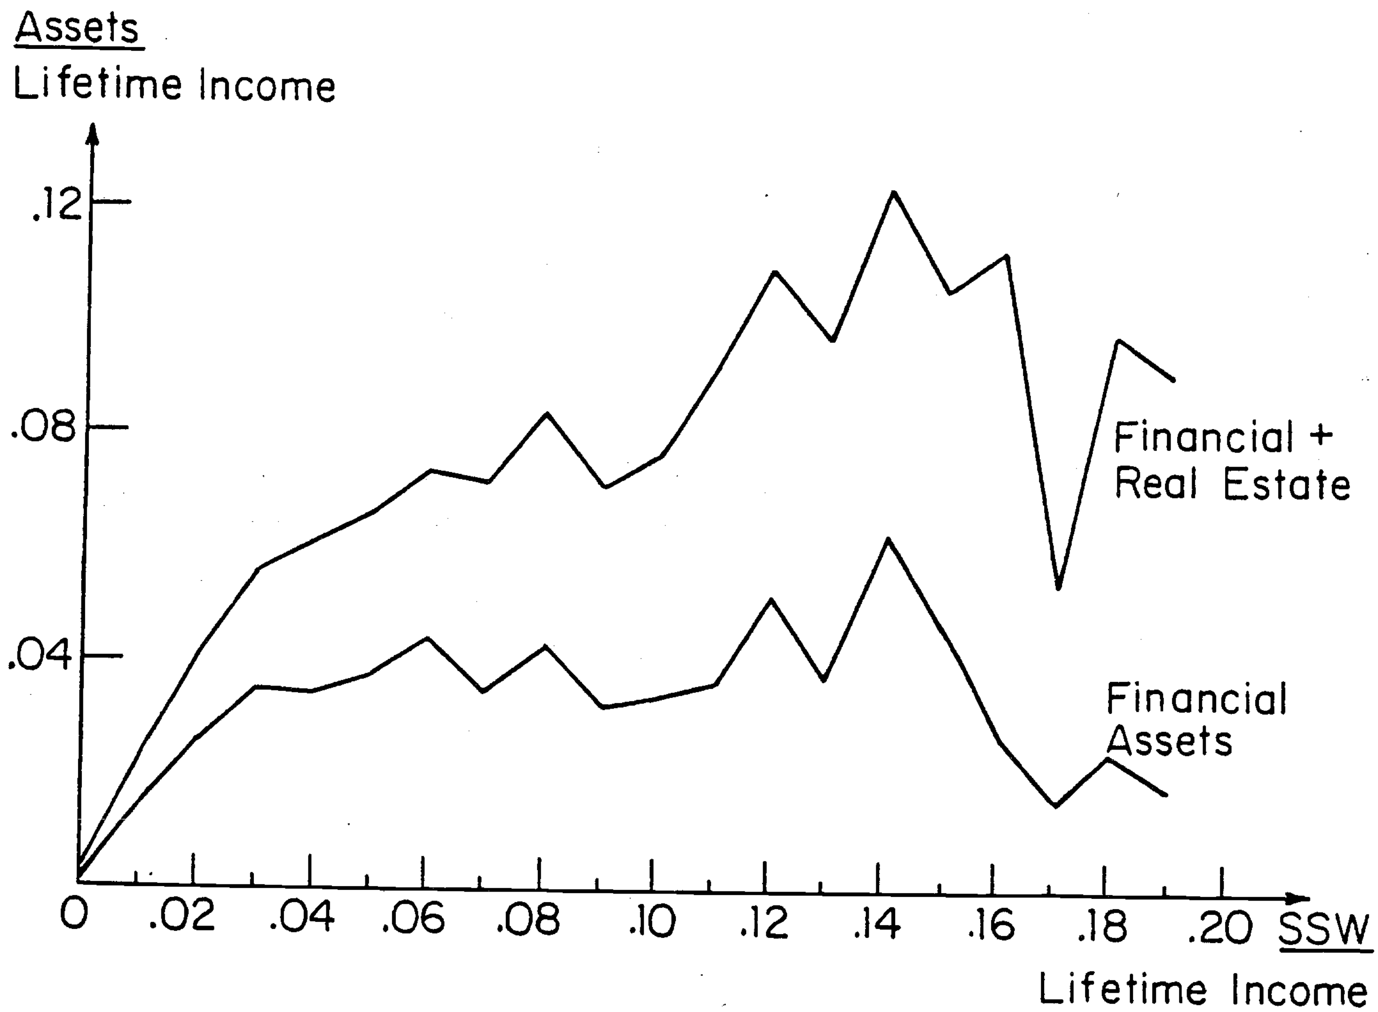
Assets
Lifetime Income
.12
.08
.04
0
.02
.04
.06
.08
.10
.12
.14
.16
.18
.20
SSW
Lifetime Income
Financial +
Real Estate
Financial
Assets

FIGURE 3

these two ratios against each other[1] on a graph that should produce a slope of -1 if the Feldstein hypothesis is correct. This pattern most certainly does not appear in the data. If anything, both A'_t and $A'_t + RE_t$ relative to lifetime income look to be increasing with SSW_t relative to lifetime income over most of the range. The dotted lines in Figure 3 indicate where most of the people are located, with social security wealth between 3% and 11% of lifetime income. In this range, A'_t looks relatively independent of SSW_t, while RE_t apparently is strongly increasing.

Once again, we must emphasize that these plots are only indicative of simple correlations. It could be that once we control for other pertinent influences (such as age and number of children) evidence that SSW displaces A' will emerge. This is why we need to estimate a statistical model. Nonetheless, in view of Figure 3 it would be quite surprising if the data provided strong support for Feldstein's displacement hypothesis.

To sum up, this preliminary look at the data makes it seem reasonable to try to use the life-cycle model of consumption to organize the data, but does not give any reason to expect the estimated model to turn up a strong bequest motive or a strong displacement effect of social security wealth.

5. Estimation Problems and Empirical Results

As everyone knows the distribution of wealth is incredibly skewed. This makes it important to estimate (2.16) by weighted least squares, for otherwise a handful of very rich families may dominate the results.[2] We assumed that the error variance was proportional to the square of $(Y_0 + SSW + PPW)$ and hence

---

[1] Specifically, the ratios of SSW to lifetime income for each person are grouped as follows: numbers between 0 and .009 are considered as zero in the figure; numbers between .01 and .019 are considered as .01, and so on. For each such class, the mean of assets divided by lifetime income is calculated and plotted in Figure 3.

[2] We have in our sample, for example, one observation for which assets are 252 times as large as lifetime earnings.

divided both sides of (2.16) by this quantity to arrive at the estimating equation:

$$(5.1)\quad \frac{A'_t + \lambda_1 SSW_t + \lambda_2 PPW_t + \lambda_3 RE_t + Y_t}{\hat{Y}_0 + SSW_t + PPW_t} = \left[\frac{\gamma P_t + \alpha_1 NSUP_t + \beta}{P_{17} + 18\alpha_2 NKIDS + \beta}\right] \times \left[\frac{\hat{Y}_0 + \lambda_1 SSW_t + \lambda_2 PPW_t + \lambda_3 RE_t}{\hat{Y}_0 + SSW_t + PPW_t}\right] + \varepsilon_t .$$

Estimation was by nonlinear least squares, using the Davidon-Fletcher Powell algorithm as programmed in the GQOPT package. However, we quickly learned that the model as originally posited was statistically unidentified, even with over 4,000 observations.

Some of the reasons are easy to explain by referring to (5.1). Consider first the parameter b_4 in β (see equation (2.15) for definition of all the b's), which was designed to estimate the wealth elasticity of bequests. Obviously it is very hard to distinguish b_4 from the other parameters pertaining to bequests (b_1, b_2, and b_3).[1] Similarly, the functional form shows that the growth rate parameter, μ in equation (2.14), will probably be a good substitute for the multiplicative parameter, γ, at least in a sample with limited variation in ages such as ours; that is, raising one while lowering the other will have very little effect on the sum of squared residuals (or on the predictions of the model). Thus it became quite clear very early in the research that neither μ nor b_4 could be estimated with any precision. Fortunately, the initial point estimates of both b_4 and μ were extremely close to zero, so little was lost by constraining them both to be zero. Setting

---

[1] Identification of b_4 hinges precariously on the presence of the constant 1.0 in equation (2.15).

$b_4 = 0$ imposes the usual linearity assumption of the life-cycle model, and setting $\mu = 0$ imposes $g = r$. These constraints raised the sum of squared residuals hardly at all, and improved the precision of the other estimates somewhat (though they were still not estimated very precisely).

What was more surprising, and more distressing, was that the α parameters in (5.1) were so hard (really impossible) to pin down empirically. Recall that these parameters are designed to distinguish between child rearing expenses (plus gifts <u>inter vivos</u>) and bequests at death as conflicting influences on asset accumulation of people with different amounts of children. The reason for the lack of identification of the α_1 in the numerator is clear enough: very few people in this age range still have dependent children. However, the α_2 in the denominator also proved impossible to estimate, despite very considerable cross-sectional variation in NKIDS. The reason in this case is that α_2 whose presumptive sign is positive, cannot be distinguished from b_2, the linear term in the bequest function (whose presumptive sign is also positive). Included in the denominator of (5.1) is a term that is essentially $(b_2 + 18\alpha_2)$NKIDS; and the presence of b_2NKIDS by itself in the numerator is apparently not enough to identify α_2 and b_2 separately. Some experimentation showed that the sum of squared residuals (henceforth) SSR function was essentially flat over a wide range of α_2 values, so we arbitrarily set $\alpha_2 = 1$.

Even after making all these additional restrictions, the estimates we obtained were disappointing. The point estimates of (5.1) along with their asymptotic standard errors are presented in Table 3, and it can be seen that they are regrettably imprecise.

Let us start with γ, the parameter designed to test whether people were following the age-assets pattern dictated by the life-cycle theory. The estimate of γ (.45) is significantly below its theoretical value of unity, which means that assets do not decline fast enough with age. This is not what Figure

Table 3

Parameter Estimates: Constrained Model

Coefficient	Estimate	Standard Error
γ	.45	.19
λ_1	.39	.45
λ_2	-.30	.54
λ_3	-.34	.10
μ	0	constrained
α_1	.10	1.45
α_2	1.0	constrained
b_1	-2.67	4.19
b_2	1.98	2.56
b_3	-.10	.47
b_4	0	constrained

Standard error of the regression = .101

1 suggested to us earlier. Apparently, controlling for other variables (such as wife's age) changes things dramatically.

Turning next to the bequest parameters (the b's), we find very little ability to pin down any systematic effect of the number of children on asset accumulation, even after constraining $\alpha_2 = 1$. The standard errors of b_1, b_2, and b_3 are all very large relative to their point estimates. Taken at face value, the point estimates do suggest a bequest motive, but a rather weak one. Table 4 uses the point estimates to tabulate the total, marginal, and average bequest as a function of the number of children. The point estimates imply that the marginal effect of an additional child on the planned bequest is positive up to 10 children, and that the bequest per child is constant at just under 1 year's worth of adult consumption over much of the relevant range. Thus, despite the fact that Figure 2 showed assets <u>decreasing</u> as the number of children increased, there is some weak evidence here of a bequest motive. But it is weak indeed, and gives no reason to think that the desire to leave a bequest is a major motive for saving. (The estimated mean bequest in the sample is only 1.25 years of consumption.)

The other question of interest was the displacement effect of social security wealth. Our point estimate states that each $1 of social security wealth leads people to decrease their financial wealth by 39¢. This does not seem an unreasonable number, and is in fact better than we might have expected after looking at Figure 3. However, the standard error of this estimate is so large (.45) that we can reject neither $\lambda_1=0$ nor $\lambda_1=1$. One other remark should be made about λ_1. In experimenting with minor changes in the specification (principally constraining different subsets of $(\gamma, \alpha_2, \lambda_2, \lambda_3)$ to be unity), we found that the estimate of λ_1 was highly unstable--as is only to be expected from the large standard errors.[1]

[1] The estimates of the b's were also highly unstable. However, we never obtained a set of point estimates of b_1, b_2, and b_3 that suggested an important bequest motive.

Table 4

Estimated Planned Bequests (in years of consumption)

Number of Children	Fraction of Sample	Total Bequest[1]	Marginal Bequest[2]	Average Bequest[3]
0	.16	-2.7	---	---
1	.18	-0.8	1.9	-0.8
2	.24	0.9	1.7	0.4
3	.17	2.3	1.5	0.8
4	.10	3.6	1.3	0.9
5	.06	4.6	1.0	0.9
6	.03	5.4	0.8	0.9
7	.02	6.1	0.6	0.9
8	.01	6.5	0.4	0.8
9	.01	6.7	0.2	0.7
10	.005	6.7	0	0.6

1/ Computed as $-2.67 + 1.98\,(NKIDS) - .10(NKIDS)^2$

2/ Computed as $1.98 - .10[(NKIDS)^2 - (NKIDS - 1)^2]$

3/ Computed as $\frac{-2.67}{NKIDS} + 1.98 - .10NKIDS$

The other λ's are much worse. The weight for private pension wealth (λ_2) has the wrong sign, though it is not significantly different from zero. The weight on real estate wealth (λ_3) is also incorrectly signed, and is highly significant as well. These estimates suggest that PPW and RE may be capturing some aspects of lifetime income that we missed in our constructed proxy $\hat{Y}_0$.

6. Summary and Conclusions

The pattern of asset holdings by age that is found in our sample of data is not on its face inconsistent with the strict form of the life-cycle theory of saving. However, when the parameters of the life-cycle model are estimated we are forced to conclude either that the taste for consumption is systematically weaker in old age than in earlier years, or that people are underproviding for their retirement consumption. In addition, the parameters of the model appear to be very poorly identified, even with an extensive cross-sectional data set. Some parameters could not be estimated at all while others could be estimated only with very large standard errors.

Though some weak evidence (not statistically significant) was turned up for the existence of a bequest motive, there is nothing in these data to suggest that the desire to leave a bequest is an important motive for saving. A few remarks are needed to put this finding into proper perspective. First, there is an untested hypothesis that underlies this research: that if there is a strong bequest motive, it ought to show up in asset holdings late in life being higher for people with more children. This seems a reasonable hypothesis to us, though it remains untested. What we actually find in the data is that once we control for other pertinent influences, families with more children have only slightly higher assets.

[1] It is apparently important to control statistically for "other things." Without such controls, Figure 2 showed clearly that families with more children have lower assets.

Second, this finding does not say that intergenerational transfers are trivial because transfers *inter vivos* might be important. Third, in the real world, unlike in our model, the actuarially fair indexed annuities of the Yaari (1965) model are not available (except for those households whose social security wealth is large enough to meet their entire demand for annuities). As a result, actual bequests will be more variable than planned bequests, with those who die earlier than expected leaving larger bequests than they had planned.

Finally, even though the data show that people with higher social security wealth (relative to lifetime income) have higher financial assets (relative to lifetime income), our model does suggest that there may be some "displacement" of private savings by social security once other things are held equal. Unfortunately, the point estimate--that each $1 of social security wealth displaces 39¢ of private financial wealth-- is too imprecise (and too unstable) to draw any conclusions.

APPENDIX

THE ESTIMATION OF LIFETIME EARNINGS

As explained in the text, lifetime earnings were generated by modelling separately the time profiles of hours and wages. The former was explained in the text. The latter required a good deal of interpolation and extrapolation using an estimated wage equation reported in Gordon and Blinder (1980). Figure A1 illustrates the construction of Y_0 in the case of an individual whose record supplied the maximum number of wage observations--four. The observations are plotted as points A, D, F, and G. The wage observation from 1928 (point A) is the wage the individual reported receiving on his first full-time job. The wage observation for 1952 (point D) is the wage the individual reported earning when he left his last job (at the time of the 1969 interview). The 1969 observation (point F) is the current wage as of the time of the 1969 interview and the 1971 observation (point G) is the current wage as of the 1971 interview. We now turn to our method of filling in the gaps between these wage observations.

The first gap we need to fill is the one between 1928 and 1942, which is the year that the "last job" began. This time period may cover more than one job. We know (except when data are missing) the starting and ending dates of the first job, the starting date of the last job (1942), and whether the individual's "longest job" intervened between these two. Figure A1 illustrates a case in which the period was entirely spent on the first job. To fill the gap between 1928 and 1942 we adjust the constant term generated by our wage equation for each individual so that we estimate the 1928 wage rate exactly, and then forecast the remainder of the period (through 1942) on the basis of the estimated wage equation. The path from A to B shows the predicted wage profile resulting from this step of the process.

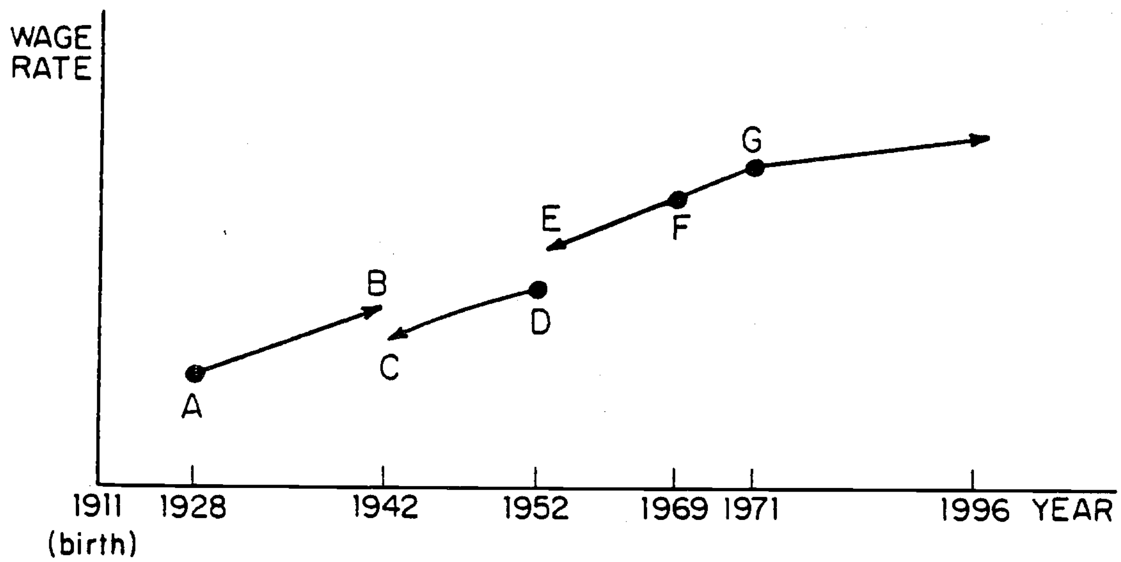
WAGE
RATE
G
E
F
B
D
C
A
1911
(birth)
1928
1942
1952
1969
1971
1996
YEAR

FIGURE A1

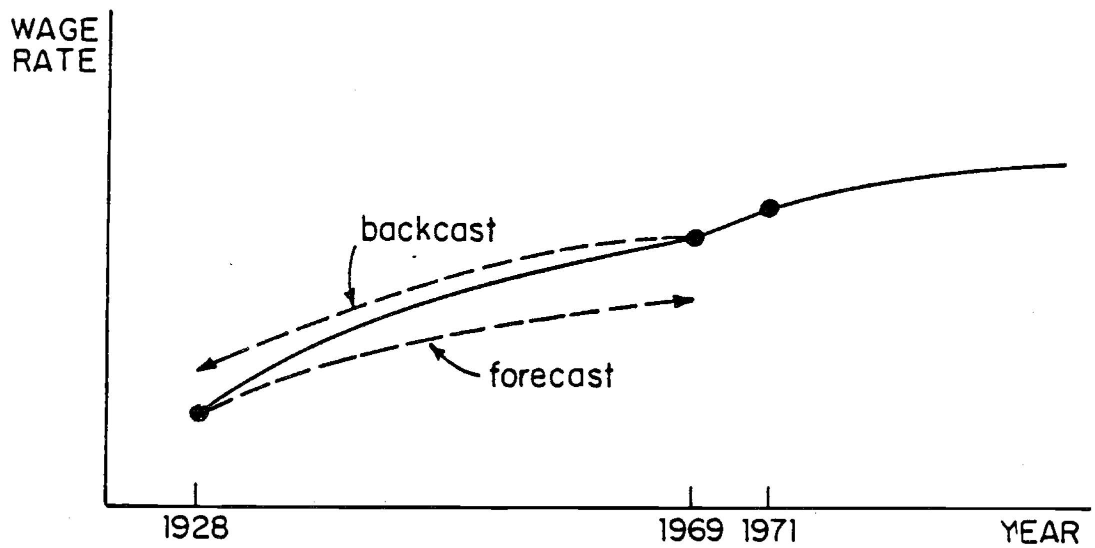

FIGURE A2

The second gap to be estimated is the one from 1942 to 1952--the period of work on the last job reported on the 1969 survey. In this case, we do not know the starting wage, but rather the wage the head received when he left the job (point D). So, instead of having to forecast this gap, we backcast it from point D by a procedure conceptually identical to that used to forecast the first wage gap. The path from D back to C illustrates the result of this projection. (In this example we have predicted that the individual took a reduction in the wage rate when he changed jobs in 1942.)

The final gap, between 1952 and 1969, is handled in the same fashion. We adjust the constant term to make the profile pass through point F exactly and backcast the market wage equation over the period to point E. The wage in 1970 is assumed to be the average of the 1969 and 1971 wages. This gives us an estimate of each individual's lifetime earnings profile from the beginning of his first full time job through 1971. Future earnings are estimated by adjusting the constant of the wage equation to make it pass through point G (the observed 1971 wage), and forecasting wages forward through age 85, the oldest age at which it was assumed anyone could work. All wages are expressed net of payroll tax. Figure A1 illustrates a completed wage profile constructed by this method.

So far, we have assumed that the individual reported the maximum number of wage observations. This, of course, was not always the case. For those individuals who held only one job during their lifetime, and there were a number of them, we backcast their 1969 wage rate to the beginning of their work history. If we also had the wage when they started their first job, which in this case was the same as their current job, we had two wage observations for the same job. This enabled us either to take the 1969 observation and backcast wages to the beginning of the job or to take the starting wage and forecast wages to 1969. We chose to do both and to take a weighted average of the two sets of estimates, using weights that vary linearly with age to force the profile through the two observed wages. Figure A_2 shows an example of an individual for whom this procedure was used.

In other cases, we knew that the individual had a job, but he did not report the wage rate he earned. In this case we constructed an estimate of the wage rate from our wage equation without any adjustment of the constant term.

Since both the data and the work histories were scantier for wives than for husbands, the construction of lifetimes earnings for wives was more problematic. We began by taking whatever wage observations we could from the LRHS surveys, forecasting and backcasting earnings by use of the wage equation estimated for husbands. In doing this, we implicitly assumed that women's earnings profiles are generated by the same mechanism that generates the earnings profile for men. The reader who is incredulous about this assumption should recall that we always adjusted the levels of the wage profiles up or down to conform with the observed wage rates at various points in time. Our assumption therefore introduces error only insofar as women's earnings profiles differ in shape from men's.

An additional difficulty arose in attempting to generate earnings profiles for women: wage rates were frequently unavailable. For example, we might have a wife whom we knew held a job between 1930 and 1935, after which she left the labor force for a protracted period (probably to raise a family) before reentering the labor force in, say, 1955. At the time of the 1969 interview, the woman may not remember the wage she earned during the earlier job. Estimating that wage rate with our wage equation would clearly be incorrect, however, since our equation was estimated for men only. How then do we estimate wage rates for such women? There were too many cases of missing data on women's wage rates to discard these observations from the sample. So instead we imputed a woman's wage as 70% of the average wage in the industry in which she worked.

Hours of work were treated in the same way as for men; in particular, the probability that a wife would be retired in any given year was assumed equal to the probability that her husband would be retired. This seems a reasonable assumption when husband and wife are close in age (the vast majority of cases), but is probably totally unreasonable in (the small minority of) cases in which the two ages differed dramatically. This was another reason why we eliminated such couples from the sample.

REFERENCES

Atkinson, Anthony B., "The Distribution of Wealth and the Individual Life Cycle," Oxford Economic Papers, July 1971, pp. 239-254.

Barro, Robert J., "Are Government Bonds Net Wealth?", Journal of Political Economy, 82, November-December 1974, pp. 1095-1117.

Barro, Robert J., "Reply to Feldstein and Buchanan," Journal of Political Economy, 84, April 1976, pp. 343-349.

Blinder, Alan S., "Inequality and Mobility in the Distribution of Wealth," Kyklos, 29, Fasc. 4, 1976a, pp. 607-638.

Blinder, Alan S., "Intergenerational Transfers and Life-Cycle Consumption," American Economic Review, 66, May 1976b, pp. 87-93.

Darby, Michael R., The Effects of Social Security on Income and the Capital Stock (Washington: American Enterprise Institute), 1979.

Diamond, Peter A., "A Framework for Social Security Analysis," Journal of Public Economics, 8, 1977, pp. 275-298.

Feldstein, Martin S., "Social Security, Induced Retirement, and Aggregate Capital Accumulation," Journal of Political Economy, 1974, pp. 905-926.

Feldstein, Martin S., "Social Security Benefits and the Accumulation of Preretirement Wealth," NBER Working Paper No. 477, May 1980.

Feldstein, Martin S. and Anthony Pellechio, "Social Security and Household Wealth Accumulation: New Microeconomic Evidence," Review of Economics and Statistics, 1979, pp. 361-368.

Friedman, Milton, A Theory of the Consumption Function (Princeton, N.J.: Princeton University Press), 1957.

Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," Journal of Political Economy, 86 (December 1978), pp. 971-987.

Gordon, Robert H. and Alan S. Blinder, "Market Wages, Reservation Wages, and Retirement Decisions," Journal of Public Economics, 14 (1980), pp. .

Hall, Robert E. and Frederic S. Mishkin, "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households," NBER Working Paper No. 505, July 1980.

Kotlikoff, Laurence J. and Lawrence H. Summers, "The Role of International Transfers in Aggregate Capital Accumulation," NBER Working Paper No. 445, February 1980.

Menchik, Paul L., "Inter-generational Transmission of Inequality: An Empirical Study of Wealth Mobility," Economica, 46, November 1979, pp. 349-362.

Menchik, Paul L. and Martin David, "The Effect of Income Distribution and Redistribution on Lifetime Savings and Bequests," mimeo, University of Wisconsin, November 1979.

Mirer, Thad W., "The Dissaving Behavior of the Aged," Southern Economic Journal, April 1980, pp. 1197-1205.

Modigliani, Franco and Richard Brumberg, "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in K.K. Kurihara (ed.), Post Keynesian Economics (Rutgers University Press: New Brunswick, N.J.), 1954.

Munnell, Alicia H., The Effect of Social Security on Personal Savings Behavior (Cambridge, Mass: Ballinger), 1974.

Shorrocks, Anthony F., "On the Structure of Inter-generational Transfers between Families," Economica, 46, November 1979, pp. 415-425.

Tobin, James, "Life-Cycle Saving and Balanced Growth," in W. Fellner (ed.), Ten Economic Studies in the Tradition of Irving Fisher, New York, 1967, pp. 231-256.

White, Betsy B., "Empirical Tests of the Life Cycle Hypothesis," American Economic Review, 68, September 1978, pp. 547-560.

Yaari, Menachem E., "Uncertain Lifetime, Life Insurance and the Theory of the Consumer," Review of Economic Studies, 32, April 1965, pp. 137-150.