NBER WORKING PAPER SERIES

AN EXAMINATION OF EMPIRICAL TESTS OF SOCIAL SECURITY AND SAVINGS

Alan J. Auerbach

Laurence J. Kotlikoff

Working Paper No. 730

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge MA 02138

August 1981

We are grateful to the National Bureau of Economic Research for financial support and to David Reitman for excellent research assistance. The research reported here is part of the NBER's research program in Social Insurance and 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 #730
August 1981

An Examination of Empirical Tests of Social Security and Savings

ABSTRACT

The effect of social security and other forms of government debt on national savings is one of the most widely debated policy questions in economics today. Some estimates suggest that social security has reduced U.S. savings by almost forty percent. This paper examines recent cross-section and time series empirical tests of the social security-savings question and argues that, given current data, neither type of test has much potential for settling the controversy. In particular, there are a number of specification problems relating to social security time series regressions that can easily lead to highly unstable coefficients and to rejection of the hypothesis that social security reduces savings, even if it is actually true.

These points are demonstrated by running regressions on hypothetical data generated by a perfect foresight life-cycle growth model developed previously by the authors. While the data is obtained from a model in which social security reduces the nation's capital stock by almost twenty percent, time series social security regression coefficients vary enormously depending on the specified level of the program, the preferences of hypothetical households, the level of concommitant government policies, and the time interval of the data.

Alan J. Auerbach
Department of Economics
Harvard University
Cambridge, Mass. 02138
(617) 495-3934

Laurence J. Kotlikoff
Department of Economics
Yale University
New Haven, Conn. 06520
(203) 436-1937

The effect of social security and other forms of government debt on national savings is perhaps the most widely debated policy question in economics today. The issue has intrigued economics since David Ricardo, but is receiving increasing attention as U.S. economists seek explanations for low rates of capital formation and productivity growth. The numbers at stake in the social security debate are quite large. Martin Feldstein has pioneered theoretical and empirical research on this subject over the last decade. His estimates (Feldstein (1974)) suggest that social security has reduced U.S. savings by almost 40 percent. General equilibrium simulations of the long run effects of unfunded social security on capital formation (Kotlikoff (1979)) confirm a potentially large effect.

Over the years the issue has been joined at both the thoretical and empirical levels. (Barro (1974, 1978), Darby (1978), Kotlikoff (1979), Leimer and Lesnoy (1980), Feldstein (1974, 1980)). This paper examines recent cross-section and time series empirical tests of the social security-savings question and argues that given current data neither type of test has much potential for settling the controversy. The paper focuses in particular detail on the specification of social security time series regressions. There are a number of problems of specification that can easily lead to highly unstable coefficients and to rejection of the hypothesis that social security reduces savings when this is in fact true. These points are demonstrated by running regressions on hypothetical data simulated from a perfect foresight life cycle growth model (Auerbach and Kotlikoff (1980)). While the data is obtained from a model in which social security reduces the nation's capital stock by almost 20%, time series social security regression coefficients vary enormously depending on the specified level of the program, the preferences of hypothetical households, the level and type of concomitant government policies, and the time interval of the data. Indeed social security coefficients for certain regressions are negative, despite the significant burden of the social security debt.

The next section discusses theoretical problems with cross-section and time series tests of the social security-savings hypothesis. Section II points out significant errors of specification in time series social security regressions. Section III describes dynamic simulations of the effect of unfunded social security in a perfect foresight life cycle growth model. Section IV uses data simulated from this model to examine the predictions of the life cycle model for the sign, magnitude, and stability of social security time series regression coefficients. The time series coefficients obtained from this simplest of life cycle models are sufficiently unstable to cast major doubt on the value of the entire exercise.

I. Empirical Tests of Social Security and Savings—Theoretical Difficulties

The life cycle (Modigliani-Brumberg (1954)) and infinite horizon (Barro (1974)) models of intertemporal consumption choice provide the theoretical underpinning of most tests of social security's effect on savings. Both models posit rational utility maximizing behavior; the life cycle model assumes that economic agents are selfish and care only about their own current and future consumption. The infinite horizon model assumes that agents are alturistic and care about their descendent's consumption as well as their own. In a life cycle model since individuals care only about consumption over their remaining expected lifetimes, their marginal propensities to consume out of full available resources will depend crucially on their age. Intergenerational redistribution of resources in a life cycle model will affect aggregate consumption and, thus, national savings. In a Barro-type infinitely lived family model, consumption of family members of a particular age depends not on their own resources, but rather on total family resources. In a Barro model cohort consumption and thus aggregate consumption is independent of the cross-cohort distribution of resources.

The theory of unfunded social security is identical to the general theory of deficit policy; holding the path of government expenditure constant, any reduction in the net tax liabilities of one generation that are associated with a deficit must be made up by a higher net tax liability on some other generation. Thus social security, as well as any other government deficit, constitutes an intergenerational transfer. This transfer will have a real effect on aggregate consumption and national savings if different cohorts have different marginal propensities to consume out of transferred resources. Ideally, one would think that the best way to test the Feldstein-Barro government deficit question is to regress a cohort's consumption on that cohort's own resources as well as the resources of other cohorts. If resources of other cohorts significantly affect the consumption of a particular cohort, the Barro view would "seem" to be supported. If the resources of other cohorts did not affect a particular cohort's consumption, Feldstein's view would "seem" to be supported. We use the word "seem" because a finding that resources of other cohorts did not affect the consumption of a particular cohort could be consistent with the Barro view according to the following scenario: suppose that the family head at each moment in time allocates total resources to the various family members and tells them to consume just like life cycle consumers. The Barro family head could, then, always redistribute family assets in such a way that each family member at each age looked like he was behaving independently from the others. Such an exact year-by-year reshuffling of family resources seems unlikely, but is, none-the-less, a theoretical possibility.

It is also possible to provide theoretical arguments consistent with both the Feldstein life cycle view of savings and an empirical finding that the consumption of one cohort depends on the resources of other cohorts. Imagine a life cycle world in which unexpected changes in government tax-transfer policies as well as periodic macro economic fluctuations can, ex ante, be expected to significantly alter the cohort distribution of lifetime full resources and utility. In such a world completely selfish, but none-the-less risk averse life cycle savers might seek to pool these risks with members of other cohorts. In the limit perfect insurance arrangements between different cohorts concurrently alive will link the consumption of one cohort to the consumption of other cohorts and the resources of one cohort to the resources of other cohorts. Consider the following simple example of inter-cohort risk pooling. Suppose life cycle agents live for two periods, and suppose the government indicates at the beginning of one period that it will announce with probability P a lump sum transfer from the young to old at the end of the period, and with probability 1-P a lump sum transfer from the old to the young at the end of the period. At the beginning of the period risk averse old and young cohort members will seek to pool this risk among themselves. Since, by assumption, there is no social risk to be absorbed collectively by the young and the old, insurance arrangements permitting, these cohorts will arrange to consume exactly the same amount independent of the eventual government transfer. Each cohort's consumption will depend on aggregate resources rather than the cohort distribution of resources after the government enacts its policy.

While it is hard to point to explicit real world insurance markets that are in the business of pooling these types of risks across cohorts, implicit cohort risk-sharing arrangements may be operative within real world institutions such as firms and families. Robert Hall (1980) presents evidence that a sizeable fraction of American workers experience a more or less permanent relationship with a single firm over their working lifes. If one thinks of firms as competing for new young hires by offering implicit lifetime contracts that offer the highest level of expected utility, then firms will certainly have an incentive to provide implicit contracts that pool cohort risks between the firm's current old workers and the firm's new young hires. Firms could, for example, offset government tax-transfer policies that redistribute across selfish, but risk averse, cohorts by adjusting their age compensation profile. Presumably selfish old and young family members can also devise acceptable cohort risk sharing schemes, since such an implicit contract could make each member better off.

We do not mean to suggest here that perfect insurance arrangements between living members of different cohorts would completely undo the long run effects of social security on savings. Such an argument for debt neutrality would only be valid if all current and future cohorts could come together ex-ante and sign an insurance contract protecting themselves against redistributive government behavior. Our point here is that existing cohorts can collectively hedge themselves against government redistribution among themselves, and that this behavior could make them act as if only total resources mattered. Introduction of social security in such a world would raise the consumption of all currently alive cohorts, since the program would still constitute redistribution against cohorts not yet born.

While this discussion suggests that both the life cycle and infinitive horizon models can rationalize a wide range of cohort specific consumption behavior, actual knowledge of such behavior would greatly narrow the range of theories that are consistent with the facts. Unfortunately, time series data on cohort consumption and after tax future resources is not available. In its absence Feldstein and others have used time series data on aggregate consumption. Section II examines a host of difficulties with conventional time series tests of social security and savings.

If time series data on cohort consumption is not available and aggregate time series consumption data is insufficient, is there hope that cross-sectional micro evidence could resolve the controversy? A key problem with using cross section data is that the data is based on households rather than families and again certain results that might at first glance be thought to verify the life cycle theory are also perfectly consistent with the Barro theory and visa-versa. In this case problems arise because the cross section data does not distinguish transfers across generations from transfers across families.*

To give an example, suppose one found a positive and significant effect of social security wealth on household consumption in a cross section regression. While such a finding is predicted by the life cycle theory, it is also consistent with Barro's view for the following reason: Those households with greater than average social security wealth are presumably members of families that receive greater than average social security wealth; thus the social security variable is contaminated with inter-family transfers and these inter-family transfers will increase the consumption of Barro-type family members in those families receiving these greater than average transfers.

While a social security wealth coefficient would not help resolve the issue, other cross section regression experiments might. Under the life cycle theory, for example, the marginal propensity to consume this period out of all future resources increases with age. This prediction holds for conventional specification of preferences independent of whether

---

*Robert Barro suggested this point in a recent conversation with one of the authors.

the date of death is taken to be uncertain and independent of whether annuity insurance is available. In a Barro model the marginal propensity to consume out of total family resources could increase or decrease by age. Hence, given the problem that greater household resources may proxy for greater family resources, the Barro model could not be refuted by a finding that the marginal propensity to consume increased with age. But suppose one found that the marginal propensity declined with age. Such a finding would, we believe, be hard to reconcile with the Feldstein life cycle view. This is an example of a test that could possibly refute the life cycle theory, but that could not support the strict life cycle theory over the Barro view.

Unfortunately the paucity of U.S. consumption data is not confined to the macro level; there exists no U.S. micro data set that simultaneously reports household consumption and household resources in sufficient detail to conduct the above described experiment.

In the absence of micro consumption data, economists have naturally turned to data on net worth and bequests. Net worth reflects in part consumption, but in addition to accumulated past flows of consumption, current net worth reflects accumulated flows of net intergenerational transfers as well as patterns of earnings and rates of return that are not directly related to consumption.

Insofar as net worth data reflects past patterns of consumption, regression coefficients of net worth on social security variables face the same problems of interpretation as regression coefficients of consumption on social security variables. Thus a finding that social security wealth depresses private household net worth accumulation is consistent with both the life cycle and the infinitely-lived family theories of consumption. Again cross-family transfers could make families with greater than average social security wealth accumulate less private assets than would otherwise be the case.

Cross section bequest regressions may provide refutable tests of at least the life cycle theory, but caution is advised here as well. At first glance the Feldstein view would seem to be refutable by finding that the elasticity of bequests with respect to full future resources was significantly greater than zero both in the statistical sense and in the sense of a large number, and that this large elasticity held true at all levels of resources; that is, for the rich and poor. If bequests represented the only form of intergenerational transfers from the old to the young and there were no transfers from the young to the old, the issue would be settled. One could simply take this elasticity, find the percentage change in full resources of the aged represented by social security intergenerational transfers, and compute how much they transferred back.

One might think that simply the real world observation that the elderly at all lifetime income levels leave bequests by itself refutes the life cycle theory. The answer is not necessarily. First of all the theory is about net transfer flows. Our data generally details gross bequests and these gross bequests represent only one component of the gross flow of transfers from the old to the young. Very little information is available concerning transfers from the young to the old. These transfers need not, by the way, be pecuniary. Care of elderly parents in terms of housing and assistance in chores etc. constitute a transfer of resources from the young to the old.

In general, resource transfers from the old to the young and from the young to the old need not reflect altruism; these transfers can occur and still be perfectly consistent with the selfishness assumed in the life cycle model. Kotlikoff and Spivak (1981) demonstrate that family risk pooling behavior, where the risk is the uncertainty of the date of death, could give rise to bequests as well as transfers from the young to the old. These transfers can occur even though family members are completely selfish. Indeed, family members could hate each other and still bequeath and transfer resources to each other simply because of the desire to mitigate the risks associated with the date of death. The way these implicit family annuity contracts work is that old family members use their promise to bequeath in the event that they die to extract consumption support from their children as they continue to live. If the government now comes to this selfish, but risk pooling family and uses social security to transfer a dollar from the young members to the old members, the family risk sharing deal may be renegotiated with the old extracting more support from the young in exchange for their now larger promised bequest.* The fact that social security

*This presumes that the family members did not ex-ante insure among themselves against such government behavior.

comes in the form of an annuity is not fundamental, since old family members can purchase life insurance to, in effect, transfer those future streams into a current contingent bequest. In such a Feldstein world one could find a positive cross sectional coefficient of social security on bequests despite the fact that the old people end up consuming more than they would otherwise have consumed. The old may thus increase their gross transfers to the young in response to social security while at the same time expected net transfers from the old to the young decline.

To summarize these general comments, one must be cautious of any macro or micro findings that proport to support one theory and refute the other because both theories are quite flexible in terms of the savings and transfer behavior they can explain and the regression coefficients they can rationalize. In addition, in our pre-occupation with Social Security coefficients we may be overlooking some more fundamental tests of the life cycle theory.

II. Specification Bias in Social Security Time Series Regressions

The major difficulty in evaluating time series results suggesting that social security does or does not reduce private saving is that the regression specifications do not correspond very closely to the life cycle model of consumer choice. Though this criticism may also be applied to standard attempts at estimating the basic life-cycle model using aggregate time series data (Ando and Modigliani, 1963), it is a more serious problem in the "extended" life-cycle model because the particular phenomonon of interest, social security, introduces new difficulties to the interpretation of regression results.

To begin our discussion, we derive the expression for optimal consumption behavior by life-cycle consumers who have no bequest motive or labor supply decision. Lifetime utility is assumed to take the form:

$$U = \begin{cases} \sum_{t=1}^{T} (1+\rho)^{-(t-1)} \dfrac{C_t^{1-\gamma}}{1-\gamma} & \gamma > 0 \, , \; \gamma \neq 1 \\[2ex] \sum_{t=1}^{T} (1+\rho)^{-(t-1)} \log C_t & \gamma = 1 \end{cases} \tag{1}$$

where C_t is the household's consumption at age t, T is the age at death, and ρ and γ are taste parameters characterizing respectively, the pure rate of time preference and the inverse of the partial elasticity of substitution between any two years' consumption.

For a household of age i, the appropriate objective function is

$$U(i) = \begin{cases} \sum_{t=i}^{T} (1+\rho)^{-(t-1)} \dfrac{C_t^{1-\gamma}}{1-\gamma} & \gamma>0, \gamma \neq 1 \\ \sum_{t=i}^{T} (1+\rho)^{-(t-1)} \log C_t & \gamma = 1 \end{cases} \tag{2}$$

Under a regime with a proportional income tax and a social security system supported by payroll taxes, the budget constraint facing a household of age i is:

$$W_i + HW_i + SSW_i \geq PVC_i \tag{3}$$

where W_i, HW_i, SSW_i and PVC_i are, respectively, real assets, human wealth, social security wealth and the present value of current and future consumption of the household. Real wealth is directly observable. Human wealth equals the present discounted value of labor income, net of income and social security taxes. Social security wealth is the discounted value of future benefits.

Solving for the optimal consumption path yields the following expression for current consumption for such a household:

$$C_1 = \alpha_1 (W_1 + HW_1 + SSW_1) \qquad (4)$$

where

$$\alpha_1 = \left[\sum_{t=1}^{T-1} (1+\rho)^{-\frac{(t-1)}{\gamma}} \left[\prod_{s=2}^{t} (1+r_s(1-\tau_s)) \right]^{-(1-\frac{1}{\gamma})} \right]^{-1} \qquad (5)$$

and r_s and τ_s are the gross interest rate and rate of income tax that prevail when the household is age s.

From (4) and (5), it is apparent that unless γ equals one, each household's behavior will depend on future interest and tax rates, <u>given</u> the correct computation of HW_1 and SSW_1. Thus, even at the individual level, the life-cycle expression (4) will not, in general, exhibit stable coefficients as interest tax rates vary over time. This problem is particularly acute when one considers as important a change in policy as the introduction of unfunded social security. If the hypothesis that social security drastically reduces savings is correct, it will obviously lead to changes in interest rates and income tax rates along the transition path to the economy's new steady state equilibrium. Moreover, the presence of active fiscal policy will also make interpretation of regression results difficult, because various deficit and tax policies that need have nothing to do with social security will also change tax rates and rates of return and hence the estimated time series coefficients. This is the Lucas' (1974) critique.

Even assuming that $\gamma=1$, in which case α_i is a constant that depends only on ρ and i, aggregation of (4) into an estimable equation is not straightforward. Aggregation of these consumption relations over all households yields the expression:

$$C = \bar{\alpha}W + \tilde{\alpha}HW + \hat{\alpha}SSW \qquad (6)$$

where $\bar{\alpha}$, $\tilde{\alpha}$ and $\hat{\alpha}$ are the average values of α in the sample weighted by each cohort's real wealth, human wealth and social security wealth, respectively, and C, W, HW and SSW are aggregate measures. At any point in time, these three coefficients will differ because the three forms of wealth are distributed differently across age cohorts in the population. Among young individuals, for example, real wealth will be small (perhaps negative) relative to human wealth. The opposite is true for retirees. Thus, the aggregate coefficient on real wealth should be higher than that on human wealth, since α_i increases with i.

This relationship would not pose a problem in itself if it were fairly stable over time, i.e., if there were no important demographic shifts or other changes in the age distribution of these variables. However, such can hardly be claimed for a period during which social security is introduced; the age distribution of human and non-human wealth will change, and the aggregate coefficients $\bar{\alpha}$ and $\tilde{\alpha}$, and, as a result, the social security wealth coefficient $\hat{\alpha}$ should be unstable.

Further difficulties arise in actual estimation attempts because the values of HW and SSW are represented by inexact proxies. Human wealth typically is replaced by some measure of disposable income. This involves several problems. First, the appropriate income measure would not include capital income, for this constitutes double-counting. Second, the relationship between net labor income and human wealth changes by age, with the ratio of the latter to the former decreasing over time. Again, these ratios will depend on the interest and tax rates expected to prevail in the future, and these surely cannot be assumed constant.

Finally, social security wealth as it appears in equations (4) and (6) should be calculated using the actual path of future net interest rates; standard empirical analyses assume some constant discount rates to estimate social security wealth. Since rates of return will change during the economy's transition response to social security, this procedure obviously introduces errors in variables bias.

In summary there appear to be severe difficulties in estimating a stable relationship between aggregate consumption and the various forms of aggregate wealth when variations in the data reflect changes in tax rates or savings incentives that disturb net rates of return and the age distribution of different forms of wealth. The theory of unfunded social security in a life cycle model unequivocably predicts significant changes in the time path of interest rates and the distribution of resources by age.

III. Simulation of a Life-Cycle Economy with Social Security

The previous section detailed several problems of specification that make the interpretation of social security time series regression results difficult, if not impossible. Since the specification errors interact in a complex way, it is hard to know a priore what social security regression coefficient one should expect to find under the hypothesis that the data was actually obtained from a life cycle world. To investigate how serious this problem is in practice requires, then, data for which the social security life cycle savings hypothesis is actually true. The purpose of this section is to describe simulation experiments that were run to, in effect, manufacture hypothetical life cycle social security data. These data are used in Section III in time series regressions to investigate the actual predictions of the life cycle model for standard time series regression coefficients.

The simulation model is based on the assumption of perfect foresight or, equivalently, since there is no uncertainty in the model, rational expectations. The interrelated supplies of capital of all cohorts living during the transition from the economy's initial steady state to the new steady state are solved simultaneously with firms' demands for capital along the transition path to compute the economy's equilibrium growth path. We next present additional equations of the model and then describe the algorithm used to solve for the economy's equilibrium growth path.

To the household behavior already described, we add a single production sector characterized by the Cobb-Douglas production function:

$$Y_t = A\, K_t^{\varepsilon}\, ((1+g)^t L_t)^{1-\varepsilon} \qquad (7)$$

where Y_t , K_t and L_t are output, capital and labor at time t , A is a scaling constant, g is an exogenous productivity growth rate that is set equal to .02 and ε is the capital share of output that is assumed to equal .25 . L_t is simply equal to the sum of labor endowments of all individuals in the work force. In our analysis we assume that each individual's endowment of human capital grows at an annual rate of .007 and that general population growth is one percent per year. K_t is generated by a recursive equation that dictates that the change in the capital stock equals private plus public savings. Competitive behavior on the part of producers insures that the gross factor returns r_t and w_t are equated to the marginal products of capital and labor at time t:

$$r_t = \varepsilon A(K_t/(1+g)^t L_t)^{-(1-\varepsilon)} \qquad (8a)$$

$$w_t = (1-\varepsilon)A(1+g)^t (K_t/(1+g)^t L_t)\varepsilon \qquad (8b)$$

The assumption that the return to capital equals its marginal product implies that the market value of capital goods always equals their reproduction cost; i.e., adjustment of capital to desired levels is instantaneous.

Leaving aside social security, the government in our model is assumed to have a stream of consumption expenditures, G_t, that it must finance using the income tax and one-period debt. Debt is a perfect substitute for capital in household portfolios and enables the government to save (run surpluses) and dissave (run deficits). If Ag_t is defined as the value of government's assets (taking a negative value if there is a national debt), government tax revenue at the end of period t is:

$$R_t = \tau_t[W_tL_t + r_t(K_t - Ag_t)] \qquad (9)$$

where τ_t and r_t are the rates of income tax and interest at time t. Given the government's ability to issue and retire debt, its budget constraint relates the present value of its expenditures to the present value of its tax receipts plus the value of its initial assets:

$$Ag_0 + \sum_{t=0}^{\infty} [\prod_{s=0}^{t} (1+r_s)]^{-1} R_t = \sum_{t=0}^{\infty} [\prod_{s=0}^{t} (1+r_s)]^{-1} G_t \qquad (10)$$

(Note that G_t corresponds to a different concept from that reported in the National Income Accounts, which includes government purchase of capital goods.) Equation (10) is required to hold in all our simulations.

Social security is added to this government sector by introducing an announced path of social security benefits. The path of social security benefits is calculated by assuming retired individuals receive a certain replacement rate of average lifetime wages, adjusted for productivity growth. Our procedure corresponds to the lifetime wage indexing formula of the U.S. social security system. The replacement rate chosen, in combination with the path of wages for each cohort, yields a pattern of liabilities of the social security system, SSB_t. To pay for those benefits, the government uses strict "pay-as-you-go" financing, assessing a payroll tax at rate ϕ_t in year t to yield receipts equal to SSB_t.

The actual solution for the economy's behavior over time always begins with a characterization of the initial steady state, given initial tax structure and government debt. We assume that individuals of different generations alive during this steady state correctly perceive the tax schedule and factor prices they will face over time, and behave optimally with respect to these conditions. We utilize a Gauss-Seidel iteration technique to solve for this equilibrium; we start with an initial guess of the capital labor ratio (K/L), derive in each iteration a new estimate to update our guess, and continue the procedure until a fixed point is reached. Given the method of deriving new estimates of K/L, such a fixed point corresponds to a steady-state rational expectations equilibrium.

The following is a description of how the iteration technique proceeds. A schematic representation is provided in Figure 1. In the first stage, a guess is made of the capital-labor ratio (equivalent to a guess of the capital stock, since labor supply is fixed). Given the marginal productivity equations (8a) and (8b), this yields values for the wage w and interest rate r. By using the estimate of w, we may then solve for the steady state level of benefits, SSB, as well as the payroll tax, ϕ, necessary to support those benefits.. Since income and required government revenues are given a first guess of τ is obtained. From the values of r and τ, we then calculate the present values of social security benefits (SSW) as well as future earnings (HW). Assuming initial assets to be zero, we then apply equations (4) and (5) to obtain the life-cycle consumption path of the representative individual, $\underset{\sim}{C}$. From the definition of savings, this yields the age-asset profile, A, which may be aggregated (subtracting any national debt assumed to exist) to provide a new value of the capital stock and capital-labor ratio. When the initial and final values of K/L and the tax rate are the same, this implies that the steady state has been reached.

Figure 1

Iteration Procedure:
Progressive Income Tax

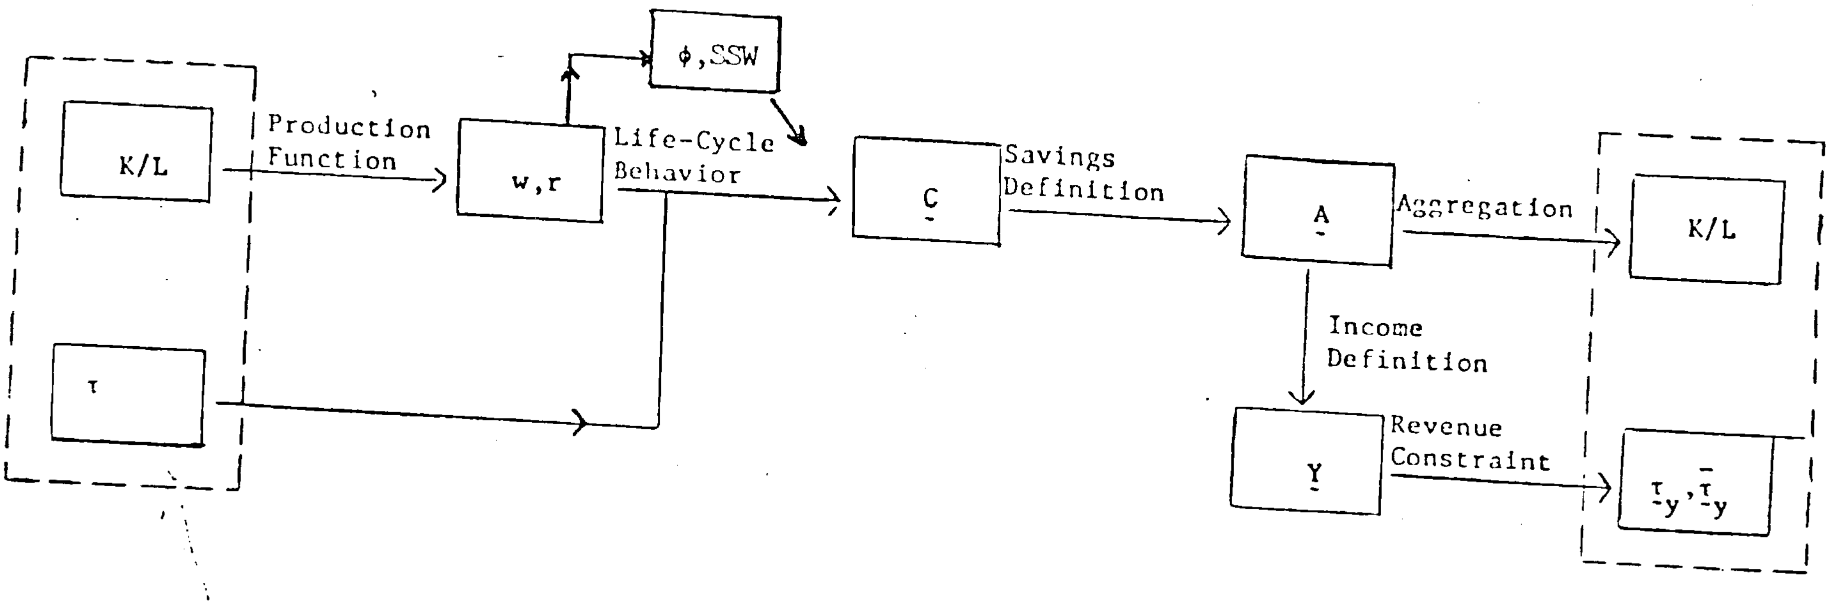

Solution for the final steady state proceeds in a similar manner. The transition is solved for in the following way. We assume the transition to the new steady state takes 150 years, then solve simultaneously for equilibrium in each of the 150 years of the transition period under the assumption that everyone believes the new steady state will obtain. This solution method is necessary because each household is assumed to take the path of future prices into account in determining its behavior. Hence, the equilibrium that results in later years will affect the equilibrium in earlier years. Specifically, we assume that individuals born after the transition begins know the transition path immediately, and that those born before the beginning of the transition behaved up to the time of the change in government policy as if the old steady state would continue forever. At the time of the announcement of a new policy existing cohorts are "born again," they behave like members of the new generation except that their horizon is less than fifty-five years, and they possess initial assets as a result of prior accumulation. An iteration technique is used again, but here we must begin with vectors of capital stocks and tax rates (one pair for each year). Further, we cannot simply solve for the behavior of a representative cohort, but rather must calculate the behavior of each cohort alive during the transition. This procedure, while conceptually no more difficult than that used to find the steady states, requires considerably more computation. As the ultimate paths converge to the final steady state well before year 150, the assumption about conditions after year 150 does not influence our results.

The solution technique just described was used to simulate the introduction of an unfunded social security system similar in size and character to that currently in force in the U.S. A number of different types of simulations are considered. In each simulation exercise we assume that there is no social security program in the economy's initial steady state; social security is introduced unexpectedly; there is no prior announcement of the system before benefits start being paid.

In our basic simulation the utility taste parameters ρ and γ are set equal to .02 and 1, respectively. There is no national debt aside from that implicit in the social security system, the income tax rate is set initially at .30, and the replacement rate is immediately set at its long run value of .45.

In the initial steady state without social security, the capital-output ratio is 3.11 and the gross interest rate is 8.04 percent. The introduction of social security ultimately leads to a new steady state where the capital stock is lowered by 17 percent, with the gross interest rate rising to 9.24 percent. Because of the lower output levels the income tax rate must rise to 31.4 percent to raise the required revenue. The solid black line in Figure 2 describes the transition path of the economy's capital-effective labor ratio. Year 1930 is the hypothetical year of enactment. The value of 100 indicates the capital-effective labor ratio that would have prevailed with no social security. Suprisingly, the transition to this final steady state occurs quite rapidly; it is virtually complete only thirty years after social security is first introduced.

Figure 2

CAPITAL STOCK IN TRANSITION*

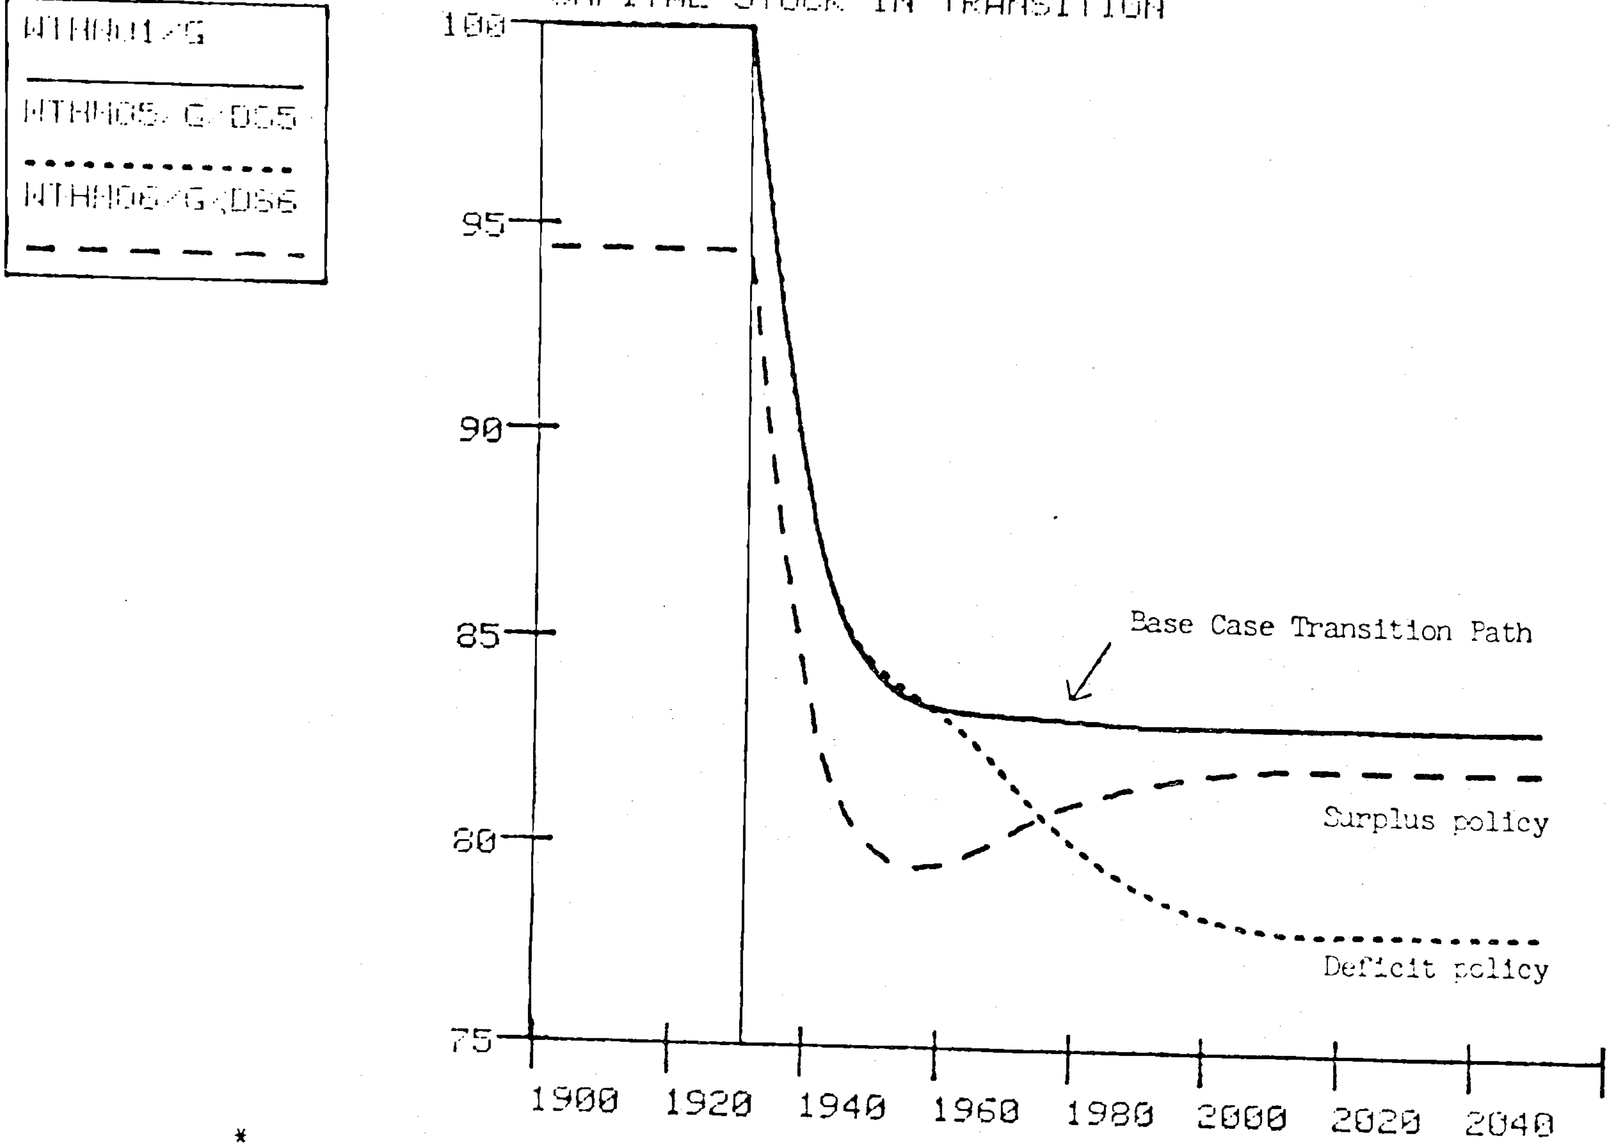

*A 40% replacement rate is assumed.

As expected, the introduction of pay-as-you-go social security is beneficial to initial retirees under the new system since they receive benefits and are required to pay no payroll taxes. However, ultimate crowding out of capital lowers the welfare of individuals in the long run. Figure 3 depicts the path of welfare changes by birth date of particular cohorts; the vertical line in the diagram marks the year of birth of the cohort born in the year of enactment. 1900 is taken as the hypothetical year of enactment. The numbers on the vertical axis measure the percentage increase (decrease) in each cohort's lifetime resources that would leave them with the same lifetime level of utility in the absence of social security as they receive under social security. Thus a cohort born in 1900 at the time of enactment suffers a loss in welfare due to social security that corresponds to a 7 percent reduction in lifetime resources in a world of no social security.

All retirees at the time of enactment gain; those retiring just at the time of the enactment gain the most, the equivalent of a 4 percent increase in lifetime resources; these retirees receive benefits for their entire retirement period, but pay no taxes whatsoever. The gain to initial middle-age generations is smaller, since they must pay payroll taxes, and continues to decline for successive younger cohorts because of the gradual crowing out of capital and reduction in wages. The eventual loss in welfare for generations in the new steady state is quite large—about 9 percent. Moreover, this loss comes on quite rapidly. Even cohorts fifteen years old when social security is introduced are net losers, and individuals born in the year of enactment lose close to the full long-run amount.

Figure 3

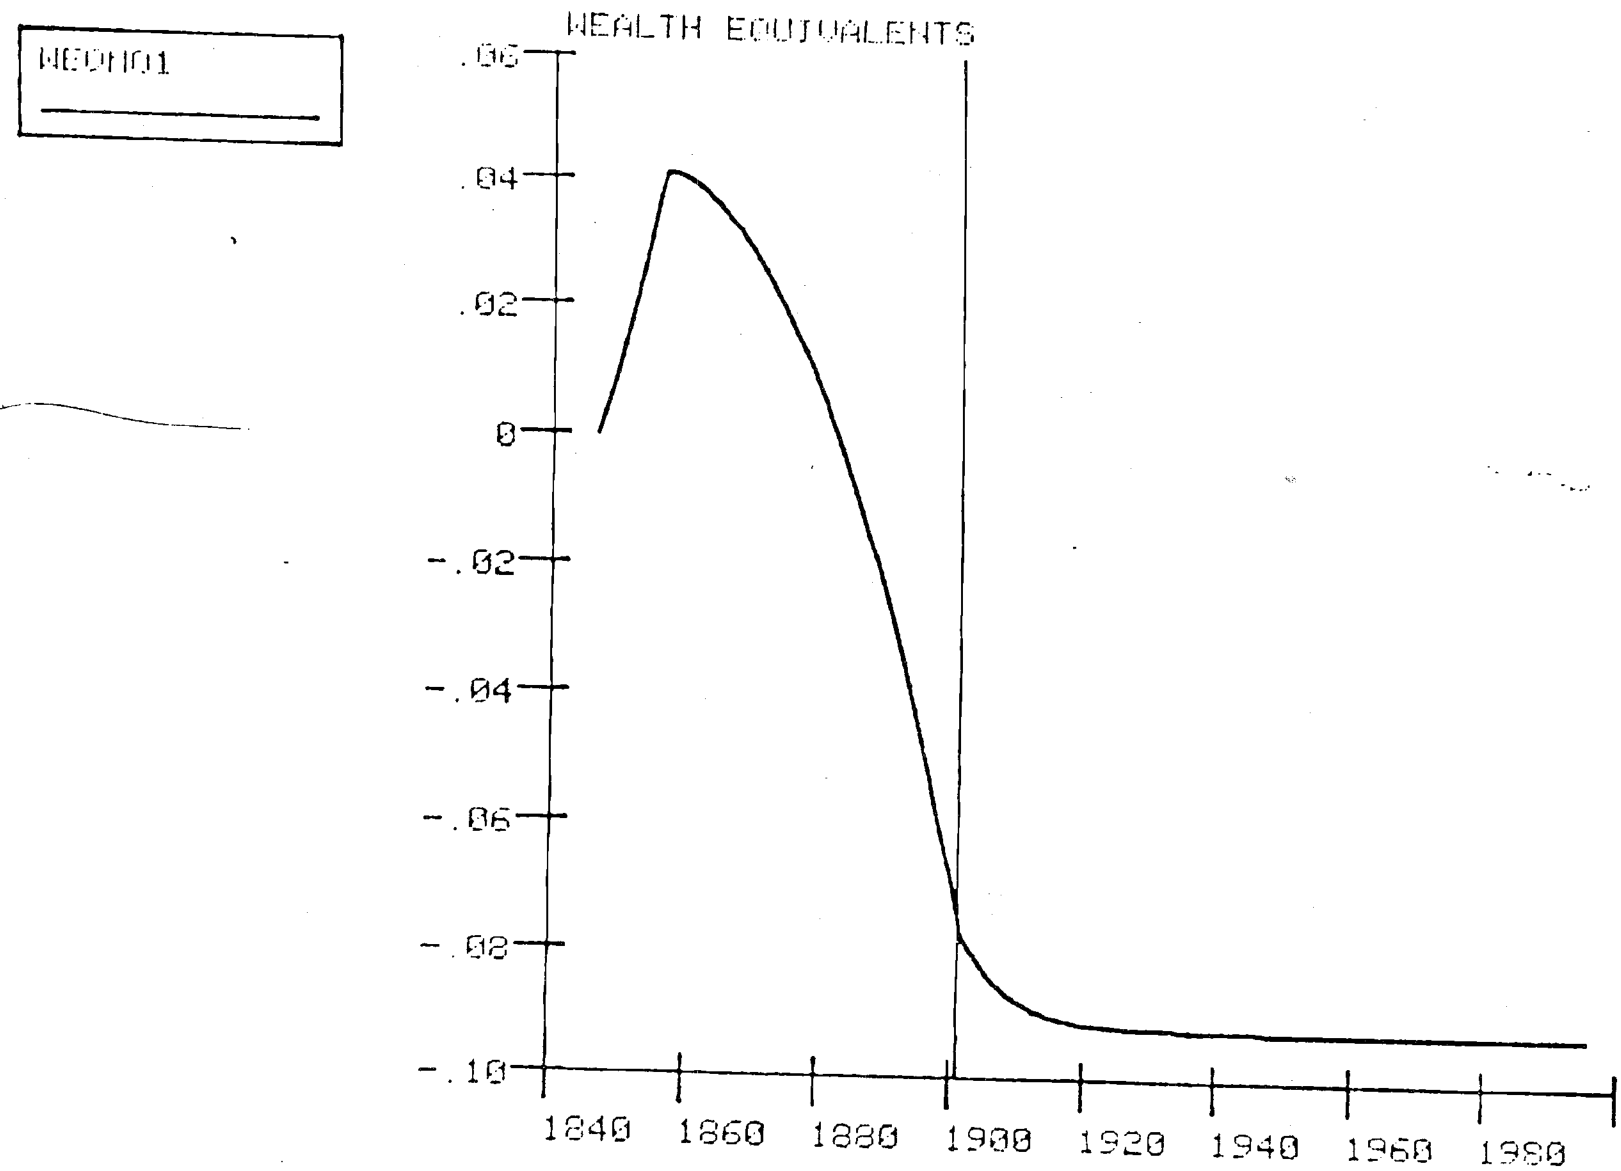
WEALTH EQUIVALENTS
.06
.04
.02
0
-.02
-.04
-.06
-.08
-.10
1840
1860
1880
1900
1920
1940
1960
1980

A more generous social security system would be expected to amplify the results just described. This is indeed the case. In fact, the basic characteristics of the transition are virtually unchanged. Adoption of a 60 percent replacement rate, rather than the 45 percent rate used in the baseline simulation, leads to greater gains and losses by those generations already in these respective categories. The long-run welfare loss is now 12 percent; those just retiring at the time of enactment gain about 55 percent, and cohorts age 15 at the time of enactment again approximately break even. The capital stock declines by 22 percent, and the transition is again quite rapid.

Just as a simple increase in replacement rates has little qualitative effects on the pattern of economic consequences of social security, this is also true of a change in preferences. Simulations carried out for values of γ ranging between .5 and 2 indicated that both the speed of convergence to the new steady state and the ultimate size of capital stock reduction were comparable to those of the baseline simulation.

In reality, we are never able to observe the effect of a single isolated fiscal policy change on the economy. Many policy shocks may be at work at once, and this fact not only confounds our understanding of how any one of these policies affects the economy, but may also influence the coefficients of a time series regression attempting to "hold constant" the effects of all but one policy change, such as social security. This latter problem is present because, as discussed above, such regressions are misspecified and hence are not invariant to policy.

One major government fiscal policy lever other than social security is the national debt.* Figure 2 shows the additional effects on the capital stock of an active debt policy during the period of transition to social security. The light dotted line shows what happens when a national debt equal to 10 percent of long run wealth is introduced over a ten year period, beginning twenty years into the social security transition, through a temporary cut in income taxes. This leads to a further decline in the capital stock equal to about 70 percent of the debt introduced. An alternative debt policy shown by the darker dashed line assumes an initial steady state debt level equal to 20 percent of wealth; in this simulation a quarter of this debt is retired over a ten year period commencing twenty years after the introduction of social security. Here, of course, there is an offsetting increase in the capital stock that comes from the debt retirement.

*See Auerbach and Kotlikoff (1980) for a simulation analysis of debt policy.

IV. Time Series Regressions Using Simulated Data

The data generated by the various simulations described in the previous section are much "cleaner" than any aggregate data we could hope to collect from actual observation. There is no measurement error, and the multitude of other shocks to the system have been assumed away. Thus, it is not possible to state with confidence that our modelling of the effects of social security is or is not realistic. However, these data may be used to estimate what the time series regressions coefficients would be for a Feldstein-type consumption function if indeed the "extended" life-cycle model were exactly correct and social security had a major impact on national saving. If the coefficients of such a controlled regression prove to be highly unstable, this should cast doubt on the assumed stability of time series coefficients in more complicated dynamic economies.

Table 1 presents results of regressions on data obtained from the baseline simulation in which an unfunded social security system with a 45 percent replacement rate is enacted at time zero. The dependent variable in the estimated equation is per capita consumption. The independent variables are disposable income (YD), disposable income lagged (YD_{-1}), real wealth (W), and social security wealth (SSW), all measured per capita. The three regressions in Table 1 differ only with respect to sample period; the first regression is based on 70 years of data starting ten years prior to enactment. Regression 2 covers 60 years of data, but begins with data from the first year of enactment. The third regression is based on data from the year 10 to year 60. The results for the sample period which begins 10 years prior to enactment do not correspond exactly to previous estimates,

Table 1

Social Security Regressions for Different Sample Periods

Base Case

Independent Variable	Interval -10 to +60	0 to +60	+10 to +60
Constant	.247 (3.23)	-.742 (10.43)	.029 (3.43)
YD	.443 (9.00)	-4.401 (14.53)	-10.026 (66.43)
YD_{-1}	-.239 (4.00)	-.132 (4.71)	11.466 (57.16)
W	.146 (29.71)	.205 (47.21)	-.036 (8.27)
SSW	.281 (77.33)	10.863 (16.47	-.523 (4.64)

t - statistics in parentheses

but are of the appropriate sign and order of magnitude. The coefficients on disposable income and wealth are of the "correct" sign; the wealth coefficient is much larger than obtained in standard time series regressions, while the social security wealth coefficient is about ten times larger than that presented by Feldstein (1974).

Shifting the sample period wreaks havoc with these results. For the remaining two intervals, the coefficients are nonsensical. The social security wealth coefficient for the middle period suggests that an additional dollar of SSW will increase concurrent consumption by almost eleven dollars. Yet this same coefficient becomes significantly negative with the next shift in sample period.

One might argue that this is an unfair test because the history of U.S. social security is one of a gradual increase in benefits, while the data underlying Table 1 assume an immediate shift for a social security system with a 40 percent replacement rate. An additional simulation was run that involved a smooth increase in the replacement rate over 40 years from zero to 60 percent. Table 2 provides some regression coefficients based on this simulated data. The coefficients for the -10 to +40 sample for the non-social security variables are in closer compliance with the real world results, but the SSW coefficient is negative and insignificant. The stability problem is even worse for this simulation as one changes sample periods. Social Security coefficients for 0 to +50 and +10 to +60 regressions are large negative numbers and highly significant.

Table 2

Social Security Regressions: 40 Year

Gradual Phase In To A 60% Replacement Rate

Independent Variable	Interval: -10 to +40	0 to +50	+10 to +60
Constant	-9.609 (30.82)	-23.001 (-19.67)	-33.896 (-9.976)
YD	.766 (1.293)	.147 (1.137)	.506 (1.908)
YD-1	-.089 (-.185)	-.172 (-1.369)	-.487 (-2.074)
W	.102 (1.896)	.909 (15.37)	1.428 (9.375)
SSW	-.00136 (-.077)	-6.277 (-11.01)	-11.408 (-7.082)

t - statistics in parentheses

Table 3 depicts the results of other experiments. Here, the sample period is fixed at -10 to +60, as in the first column of Table 1, and differences in debt policy, social security policy and preferences are posited to generate simulated data. Reading from left to right, the first two columns of estimated coefficients correspond to cases where the social security policy follows the baseline simulation, but γ is set at .5 and then to 1.5, rather than 1. The next two columns present results for the simulations described in the previous section of a phase-in and phase-out of national debt, again with the basic social security policy. The last two columns correspond to simulations of a higher level of the social security replacement rate, .6, enacted either immediately or in three stages over a twenty year period.

The most interesting general characteristic of these results is their stability relative to those in the previous table. However, there is still substantial variation in the coefficients where naive theory would not predict any. For example, an increase in the level of social security from a replacement rate of .45 to one of .6 increases the social security coefficient by over 30 percent. However, this instability is quite small relative to the severe instability found above. Indeed, just as the results for alternative policies cluster about those for the baseline simulation, given the same sample period, the same is true for other sample periods. For example, the absurd coefficient on social security wealth in the middle sample period is present regardless of which simulation is used. This suggests that, at least in the simulations compared here, the adoption of unfunded social security dominates other effects, and it is the behavior of the economy around the period of enactment which largely determines the regression results.

Table 3

Social Security Regressions for Different Regimes and Preferences

(Sample Period = -10 to +60)

	Preferences		Debt Policy		Social Security	
Independent Variable	$\gamma = .5$	$\gamma = 1.5$	phase-in	phase-out	rr = .6	rr = .6*
Constant	.291 (4.35)	.193 (2.31)	2.642 (10.44)	.063 (0.48)	.337 (3.31)	-.357 (3.50)
YD	.150 (3.64)	.592 (10.52)	.439 (2.62)	.434 (8.88)	.425 (8.94)	.613 (13.85)
YD_{-1}	-.223 (4.44)	-.239 (3.52)	-.206 (1.06)	-.091 (1.56)	-.231 (-3.98)	.030 (0.561)
W	.145 (35.05)	.149 (23.56)	.121 (9.99)	.107 (20.07)	.148 (28.14)	.052 (8.52)
SSW	.165 (75.74)	.386 (67.91)	.222 (19.07)	.280 (59.72)	.379 (82.15)	.303 (89.06)

t - statistics in parentheses

*gradual adoption

Summary and Conclusion

This paper has raised a number of concerns with respect to actual and potential tests of the effects of social security on savings. Much of the empirical literature on social security and savings has proceeded without a precise statement as to what empirical results eould definitively refute either the life cycle or infinite horizon consumption models. Many actual empirical tests appear to have very little power in deciding the issue they address. Time series regression analysis of social security is a case in point. The simulation regressions presented here suggest that virtually any social security time series coefficient, negative, zero, or positive is potentially consistent with the life cycle hypothesis.

References

Auerbach, A. and L. Kotlikoff, "National Savings, Economic Welfare, and Structure of Taxation," Cowles Foundation discussion paper, no. 570, December, 1980.

Barro, R., "Are Government Bonds Net Wealth?," Journal of Political Economy, Nov./Dec. 1974, 82, 1095-119.

Darby, M., Effects of Social Security on Income and the Capital Stock, American Enterprise Institute, 1979.

Feldstein, M., "Social Security, Induced Retirement and Aggregate Capital Accumulation," Journal of Political Economy, Sept./Oct. 1974, 82, 905-26.

__________, "Social Security, Induced Retirement and Aggregate Capital Accumulation: A Correction and Update," NBER working paper 579, Nov. 1980.

Hall, R., "The Importance of Lifetime Jobs in the U.S. Economy," NBER working paper no. 560, Oct. 1980.

Kotlikoff, L., "Social Security and Equilibrium Capital Intensity," Quarterly Journal of Economics, May 1979, 93, 233-254.

__________, "Testing the Theory of Social Security and Life Cycle Accumulation," American Economic Review, June 1979, 69, 396-410.

__________ and Avia Spivak, "The Family as an Incomplete Annuities Market," forthcoming, Journal of Political Economy, 1981.

Leimer, D. and S. Lesnoy, "Social Security and Private Savings: A Re-examination of the Time Series Evidence using Alternative Social Security Wealth Variables." paper present at AEA meetings 1980.

Lucas, R., "Econometric Policy Evaluation: A Critique," in The Phillips Curve and Labor Markets, edited by Karl Brunner and Allan H. Meltzer. The Carnegie-Rochester Conferences in Public Policy, a suppl. series to the Journal of Monetary Economics, vol. 1, Amsterdam: North-Holland, 1976.

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