

ORS Working Paper Series

Number 40

SOCIAL SECURITY, UNCERTAINTY ADJUSTMENTS,
AND THE CONSUMPTION DECISION

Dean R. Leimer and David H. Richardson^{*}

Division of Economic Research

November 1989

Social Security Administration
Office of Policy
Office of Research and Statistics

* Division of Economic Research, Social Security Administration, 4301 Connecticut Ave., N.W., Washington, D.C. 20008, and Department of Economics, St. Lawrence University, Canton, NY 13617, respectively.

This project was initiated when David Richardson was a visiting scholar at the Social Security Administration. The authors wish to thank Benjamin Bridges, John Hambor, David Harrington, Michael Leonesio, Selig Lesnoy, Daniel Radner, and John Straka for helpful comments.

I. Introduction

This paper reports on an analysis of the consumption decisions of individuals. A consumption function is developed that can be viewed as an extension of the traditional life cycle-permanent income specification, with consumption determined as an age-specific proportion of current and prospective wealth. Special attention is focused on the degree of substitutability between current and prospective wealth and on the differential effects of the various types of prospective income flows on the consumption decision.

The various types of prospective income flows may have different effects on consumption because of differences in the uncertainty attached to these flows and differences in the effect that these flows have on private transfer expenditures and receipts. Both of these factors must be considered in analyzing the potential effect of social security wealth on the consumption decision. Previous empirical analyses of this effect typically adopt perception assumptions that treat prospective social security taxes and benefits as certain and discount them at the riskless market rate of interest.¹ One of the innovations of our approach is that uncertainty is introduced explicitly into the evaluation of prospective wealth. Each of the different types of prospective wealth considered in this paper is assumed to be derived from uncertain income flows, and a subjective risk adjustment associated with each flow is estimated from the data.

The consumption function specification developed in this paper incorporates both life cycle and intergenerational consumption behavior.² Our specification allows for a broad range of private transfer offset behavior in response to social security and other government transfers and their

associated taxes. In particular, our specification expands on previous cross-section analyses by incorporating changes in both bequests and *inter vivos* private transfers as possible offset mechanisms used by intergenerational consumers.

The empirical results presented in this paper have important policy implications for the Social Security program. The relevant parameter estimates provide support for the view that reductions in social security wealth would have a positive and sizable effect on private saving. The results also suggest, however, that the increased saving implied by program reductions may come at a substantial cost in economic efficiency and the welfare of present and future consumers. Under certain interpretations, the results also provide a means of identifying the value to consumers of the insurance aspects of the Social Security program.

The outline of the article is as follows. A few comments on our general approach to modeling the consumption function are given in Section II. Our definition of prospective wealth and treatment of uncertainty are detailed in Section III, and our approach to modeling life cycle and intergenerational behavior is discussed in Section IV. Section V describes the cross-section data from the 1982-83 Consumer Expenditure Survey that are used in the study. Estimates of the key parameters are given in Section VI and the implications of the results are discussed in Section VII. The final section presents some concluding remarks.

II. General Approach

The generally accepted procedure in estimating a consumption function is to begin by deriving the consumption function from utility maximization. The utility function is usually specified as the present value of time-separable single period utility functions, and the budget constraint reflects the equality between the present value of lifetime resources and lifetime consumption and bequests. This procedure not only provides a functional form for the consumption function but also serves as a means of linking the parameters of the consumption function to the more fundamental parameters of the utility function, such as the rate of time preference and the degree of risk aversion.

Analysis of consumption behavior is complicated considerably if uncertainty in the prospective income flows is considered. One approach is that developed by Hall (1978), which invokes the Euler equations derived from the first-order conditions for optimal consumption under uncertainty. This technique can be applied directly using intertemporal data and has the advantage that it does not require any assumptions about how prospective wealth is perceived. The disadvantage of this approach for our purpose is that it precludes estimation of the full set of consumption function parameters.

Another approach explores the conditions under which the life cycle-permanent income proportionality hypothesis holds when the uncertain income flows are replaced by their expectations. The conditions for these certainty equivalence results are very stringent, however, amounting to either strong restrictions on the distribution functions of the uncertain income³ or implausible forms for the utility function.⁴ To avoid making these

restrictive assumptions, some researchers have derived consumption functions directly from utility maximization under more general assumptions, e.g., when the utility function exhibits constant absolute risk aversion.⁵

Unfortunately, derivation of the consumption function for the intuitively more appealing case of constant relative risk aversion is not mathematically tractable.

We recognize the desirability of deriving the consumption function from an underlying utility function under the assumption of maximizing behavior. We have not adopted the rigorous form of this approach under uncertainty, however, because we wish to incorporate the effect of uncertain income on consumption behavior without imposing implausible forms for the utility function or placing strong restrictions on the distribution functions of uncertain future income. Instead, we assume that consumers simplify the utility maximization problem under uncertainty in order to obtain tractable decision rules. Specifically, we assume that consumers first convert uncertain future income flows to their subjective fungible equivalents (i.e., the amount of fungible assets considered by consumers to be equivalent to the uncertain future income flows) and then solve the utility maximization problem as in the certainty case. This allows tractable consumption functions of the general form

$$c = k W \tag{2.1}$$

to be derived from plausible forms for the utility function, where c denotes current consumption, k is an age-specific factor of proportionality, and W denotes the subjective fungible equivalence measure of expected lifetime wealth available for remaining lifetime consumption, bequests, and contingencies. Factors that may influence the proportion k include the rate of time preference, the taste for contingency or precautionary saving, the

intensity of the bequest motive, and the effect of borrowing constraints. Although not estimated separately in our consumption function specification, each of these factors may change over the life cycle. Consequently, we adopt a specification that imposes no constraints on the life cycle pattern of the consumption-wealth proportions.

This specification of the consumption function is also guided by our intended application, a long-run simulation model of the U.S. economy. Our development and estimation of the consumption function was prompted by the failure of traditional specifications to capture all of the features that we felt were essential for our application. Included among these features was the ability of the consumption function to reflect the potential effects of changes in the population age structure. The large changes expected in the population age structure between now and the middle of the next century may have important effects on aggregate consumption if consumption and saving behavior change systematically with age, as suggested by the life cycle model. The potential importance of such effects led us to consider consumption function specifications that were explicitly age-specific and sufficiently flexible to allow changes over the life cycle in the parameters that determine the age-specific consumption to wealth ratios.

III. Prospective Wealth Perceptions

Because of the uncertainty associated with expected future income flows, the actuarial present value of these flows is not a perfect substitute for fungible wealth, i.e., for assets whose present market value is known with certainty. As such, consumers are likely to discount prospective wealth relative to fungible wealth in developing subjective fungible equivalents; e.g.,

$$W = A + \lambda P, \tag{3.1}$$

where A represents fungible assets, P represents uncertain prospective wealth, and λ is an imperfect substitutability parameter constrained on the unit interval.

This specification is also consistent with several approaches suggested by Zeldes (1989) for specifying an empirical consumption function. Based on numerical solutions for the mathematically intractable case of constant relative risk aversion with uncertain future income flows, he suggests using an interest rate above the riskless rate to discount uncertain expected income flows (in the derivation of P) or applying a relative weight of less than one (λ) to the uncertain prospective wealth measure. His preferred approach is to specify the relative contribution of uncertain prospective wealth (λ) as an increasing function of the ratio of certain assets to uncertain prospective wealth. All of these approaches are tested in our empirical implementation. Although Zeldes' numerical solutions provide technical support for equation (3.1), there are also economic rationales for employing such a specification. Borrowing constraints on prospective wealth, for example, might also result in the specification of λ as an increasing function of the ratio of certain to uncertain wealth.

The different forms of prospective wealth considered in this paper include human wealth, derived from expected labor income net of labor income taxes other than the Social Security payroll tax; social security wealth, derived from expected Social Security benefits net of Social Security taxes; other government transfers wealth, derived from expected government transfers from programs other than Social Security; and private transfers wealth, derived from private *inter vivos* transfers and inheritances.⁶ The different forms of prospective wealth will generally have differing effects on the total portfolio risk of the individual consumer, leading consumers to apply different risk premiums or uncertainty adjustments in evaluating the contribution of these prospective flows to consumption wealth.

Social security wealth, in particular, has unique properties that affect the overall risk of the individual portfolio. The insurance aspects of the Social Security program allow the diversification of certain mortality, disability, involuntary early retirement, and human capital risks that the individual would otherwise be unable to diversify due to imperfect private annuity and insurance markets. The compulsory nature of the Social Security program avoids the problem of adverse selection faced by private insurance coverage of mortality, disability, and involuntary early retirement risks. A small theoretical literature has developed illustrating the potential for the insurance aspects of Social Security to lead to welfare gains. Abel (1985) and Hubbard and Judd (1987) are two recent examples focusing on the diversification of mortality risk effected by Social Security. Merton (1983) demonstrates that a tax and transfer system not unlike the present Social Security system can reduce certain economic inefficiencies resulting from the nonmarketability of human capital. The automatic inflation-adjustment of Social Security benefits provides the individual with protection against

unexpected inflation for that part of the portfolio, and net social security wealth represents an asset with less expected intertemporal volatility in its rate of return than even that associated with long-term government bonds.⁷

Because of imperfections in private annuity and insurance markets, then, so-called "riskless" assets fail to cover certain types of risk, such as the risk of living longer than expected. Consequently, inclusion of social security wealth in the individual portfolio may actually reduce the overall risk of the portfolio, even relative to riskless assets, implying a negative risk premium or uncertainty adjustment. Individual uncertainty concerning future Social Security taxes and benefits, including the effect of possible changes in Social Security program provisions, provide an opposing influence, so that the direction of the uncertainty adjustment is not determinate *a priori*.

Similar comments apply to the effect of other government transfers wealth on the overall risk of the individual portfolio. Some programs included in this category, such as unemployment insurance and workmen's compensation, have explicit insurance components. Other programs in this category, such as public assistance and welfare, can be interpreted as providing implicit insurance against the possibility of various types of need arising. The insurance aspects of these programs allow individuals to diversify certain risks not covered perfectly by private insurance markets and suggest the possibility of an associated negative uncertainty adjustment. As in the case of social security wealth, the direction of the uncertainty adjustment for other government transfers wealth is not known *a priori* because of opposing influences, such as uncertainty concerning future program provisions.

The different uncertainty characteristics associated with the various types of prospective wealth suggest an approach that recognizes the general

imperfect substitutability of prospective wealth for fungible assets, but also allows for different uncertainty adjustments for each of the various types of prospective wealth. In equation (3.1) above, the proportional adjustment factor λ reflects the general imperfect substitutability of prospective wealth for fungible assets.

The perception assumptions adopted for the various components of prospective wealth also allow for a separate, exponential, uncertainty adjustment for each component. Consumers were assumed to be aware of the current cross-section average age profiles of the proportion of recipients and income per recipient for each type of prospective income flow. In particular, we assumed that consumers perceive these average age profiles as constant over time in relative terms; i.e., the ratio of the cross-section profile values for any two ages is perceived by consumers as time-invariant. Let N_a represent the total population of age a , N_a^x the number of age a receiving a particular income type x , and X_a^t the total income flow of that type to individuals of age a . The current "reciency rate" profile for that income type is then given by

$$p_a = N_a^x / N_a, \quad (3.2)$$

and the current per recipient income profile for that type is given by

$$X_a = X_a^t / N_a^x. \quad (3.3)$$

For each component of prospective wealth, we assumed that individuals perceive their future reciency status for the associated income flow to be dependent on their present status. Specifically, let $P_{a,i}^{rx}$ represent the probability of reciency of a future income flow of type x at age i , given reciency at current age a , where

$$\begin{aligned} P_{a,i}^{rx} &= 1 && \text{if } p_i > p_a \\ &= p_i / p_a && \text{if } p_i \leq p_a. \end{aligned} \quad (3.4)$$

Analogously, let $p_{a,i}^{nx}$ represent the probability of reciprocity of an income flow of type x at age i , given non-reciprocity at age a , where

$$\begin{aligned} p_{a,i}^{nx} &= (p_i - p_a) / (1 - p_a) && \text{if } p_i > p_a \\ &= 0 && \text{if } p_i \leq p_a. \end{aligned} \quad (3.5)$$

In effect, when reciprocity rates increase over the remaining life cycle, new recipients are assumed to be drawn from current non-recipients, with no current recipients becoming non-recipients. Conversely, when reciprocity rates decrease over the remaining life cycle, new non-recipients are assumed to be drawn exclusively from current recipients, with no current non-recipients becoming recipients. Although arbitrary, these perception assumptions are clearly superior to the alternatives of assuming that future reciprocity is either completely uncorrelated or perfectly correlated with current reciprocity status. The former alternative ignores the strong correlation between current and future reciprocity status, while the latter ignores typical life cycle patterns of reciprocity and therefore generates reciprocity rate expectations that are unrealistic on average.

Using equations (3.4) and (3.5), the assumed human wealth perception for those with current labor income can be represented as

$$HF_{a,j} = \sum_{i=a}^D S_{a,i} p_{a,i}^{rh} H_{a,j} (H_i/H_a) (1+r_H)^{a-i}, \quad (3.6)$$

where D denotes the maximum attainable age (assumed to be age 100), $S_{a,i}$ the probability of surviving from age a to age i , $p_{a,i}^{rh}$ the probability of receiving labor income at age i for a current worker aged a , $H_{a,j}$ the current gross labor income of individual j at age a , H_i the current gross labor income per recipient profile at age i ,⁸ and r_H the effective discount rate applied to prospective labor income. Thus, current workers were assumed to perceive that, when working, their future earnings would be maintained at the same

position relative to the average earnings profile. The term $(1+r_H)^{a-i}$ can be envisioned as encompassing three effects: the expected rate of upward shift in the cross-section average earnings profile over time, the after-tax riskless discount rate, and the specific exponential risk adjustment applied to this type of prospective income flow.

Since we did not have information on potential earnings levels for current non-workers, we simply assumed that future earnings, if positive, were perceived to be on the average earnings profile. The prospective human wealth for current non-workers, then, is given by

$$HF_{a,j} = \sum_{i=a+1}^D S_{a,i} P_{a,i}^{nh} H_i (1+r_H)^{a-i}, \quad (3.7)$$

where $P_{a,i}^{nh}$ denotes the probability of receiving labor income at age i for a current non-worker aged a , and, as before, r_H encompasses the expected rate of upward shift in the cross-section average earnings profile over time, the after-tax riskless discount rate, and the specific exponential risk adjustment applied to future earnings.

The assumed gross social security wealth perception for those currently receiving social security benefits can be represented as

$$SFB_{a,j} = \sum_{i=a}^D S_{a,i} P_{a,i}^{rb} B_{a,j} (1+r_S)^{a-i}, \quad (3.8)$$

where $P_{a,i}^{rb}$ denotes the probability of receiving social security benefits at age i for a current beneficiary aged a , and $B_{a,j}$ represents the current benefit of individual j at age a . Current beneficiaries, then, were assumed to perceive their benefits to change at a constant rate over time, if they remained beneficiaries at future ages. This constant rate of change in expected benefits is reflected in the r_S parameter, along with the after-tax

riskless discount rate and an exponential risk adjustment for social security wealth.

The gross social security wealth for current non-beneficiaries was broken into two groups, corresponding to individuals paying and not paying current social security taxes. Current social security taxpayers were assumed to perceive their future benefits, if positive, as equal to the average benefit per beneficiary, adjusted for their current social security tax payment relative to the average tax payment; i.e.,

$$SFB_{a,j} = \sum_{i=a+1}^D S_{a,i} P_{a,i}^{nb} B_i (T_{a,j}/T_a) (1+r_S)^{a-i}, \quad (3.9)$$

where $P_{a,i}^{nb}$ denotes the probability of receiving social security benefits at age i for a current non-beneficiary aged a , B_i the current benefit per beneficiary profile at age i , $T_{a,j}$ the current social security tax payment of individual j at age a , and T_a the current social security tax payment per taxpayer profile at age a . Current social security non-beneficiaries who are not current social security taxpayers were assumed to perceive their future benefits, if positive, as simply equal to the average benefit per beneficiary; i.e.,

$$SFB_{a,j} = \sum_{i=a+1}^D S_{a,i} P_{a,i}^{nb} B_i (1+r_S)^{a-i}. \quad (3.10)$$

In both equations (3.9) and (3.10), then, the r_S parameter encompasses the expected rate of upward shift over time in the cross-section average benefit profile, the after-tax riskless discount rate, and the specific exponential risk adjustment applied to prospective social security benefits.

The assumed perceptions for social security tax liability are exact analogues of the human wealth perceptions. For current social security taxpayers,

$$SFT_{a,j} = \sum_{i=a}^D S_{a,i} P_{a,i}^{rt} T_{a,j} (T_i/T_a) (1+r_S)^{a-i}, \quad (3.11)$$

where $P_{a,i}^{rt}$ denotes the probability of paying social security taxes at age i for a current taxpayer aged a . Thus, current taxpayers were assumed to perceive that, when positive, their future social security tax payments would be maintained at the same position relative to the average tax payment profile. The assumed social security tax liability perception for current non-taxpayers is given by

$$SFT_{a,j}^{nt} = \sum_{i=a+1}^D S_{a,i} P_{a,i}^{nt} T_i (1+r_S)^{a-i}, \quad (3.12)$$

where $P_{a,i}^{nt}$ denotes the probability of paying social security taxes at age i for a current non-taxpayer aged a . Current non-taxpayers, then, were assumed to perceive future tax payments, if positive, as equal to average tax payments.⁹

The assumed perceptions for other government transfers wealth and private transfers wealth are also exact analogues of the human wealth perceptions. Specifically, other government transfers wealth can be represented as

$$OF_{a,j} = \sum_{i=a}^D S_{a,i} P_{a,i}^{ro} O_{a,j} (O_i/O_a) (1+r_O)^{a-i} \quad (3.13)$$

for current recipients of other government transfers, and as

$$OF_{a,j} = \sum_{i=a+1}^D S_{a,i} P_{a,i}^{no} O_i (1+r_O)^{a-i} \quad (3.14)$$

for current non-recipients. Similarly, current recipients of private transfers were assumed to perceive their private transfers wealth as

$$IF_{a,j} = \sum_{i=a}^D S_{a,i} P_{a,i}^{ri} I_{a,j} (I_i/I_a) (1+r_I)^{a-i}, \quad (3.15)$$

while current non-recipients of private transfers were assumed to perceive their private transfers wealth as

$$IF_{a,j} = \sum_{i=a+1}^D S_{a,i} P_{a,i}^{ni} I_i (1+r_I)^{a-i}. \quad (3.16)$$

IV. Life Cycle vs. Intergenerational Behavior

Previous investigations of the effect of social security on consumption have contrasted two types of consumption behavior, which we refer to as life cycle and intergenerational. Life cycle consumers consider social security wealth to be equivalent to other forms of consumption wealth, with a possible adjustment for imperfect substitutability. The private transfers behavior of life cycle consumers is assumed to be unaffected by changes in social security wealth, so that increases in social security wealth, such as those occasioned by the start up of a pay-as-you-go social security system, tend to increase the consumption of life cycle consumers and reduce private saving. In contrast, intergenerational consumers consider the effect on other generations of the unfunded liability associated with social security wealth in a pay-as-you-go system, and altruistically adjust their private transfers accordingly. As such, a pay-as-you-go social security system does not alter the consumption of intergenerational consumers, as any associated changes in lifetime resources are passed on to other generations rather than consumed.

The consumption function specification adopted in this analysis extends the investigation of these two types of consumption behavior in several ways. First, we incorporate estimates of other government transfers wealth into the consumption function in addition to social security wealth. We assume that life cycle and intergenerational consumers react to other government transfers wealth in the same way that they react to social security wealth. Second, as indicated above, we estimate separate discount rate parameters for each of the prospective wealth components, allowing the data to dictate the growth rates and risk adjustments appropriate for each prospective wealth component. Third, we attempt to model several alternative private transfer mechanisms

that might be used by intergenerational consumers to transfer to other generations the changes in lifetime resources associated with their social security and other government transfers wealth. These offsetting intergenerational transfers may be evident as portions of fungible assets and prospective wealth that have been set aside for transfer to other generations and are therefore not available for the lifetime consumption of the generation currently holding that wealth.

Equation (3.1) can be revised and expanded to illustrate these effects. First, for pure life cycle consumers, the prospective wealth component can be written as

$$P = HF(1-t_H) + SF + OF + IF, \quad (4.1)$$

where t_H denotes the tax rate (assumed constant over time) on labor income, exclusive of Social Security payroll taxes, and HF, SF, OF, and IF represent gross human wealth, net social security wealth (gross social security benefit wealth, SFB, less social security tax liability, SFT), other government transfers wealth, and private transfers wealth, respectively.

The simplifying principle underlying our conceptualization of intergenerational consumption behavior is that some portion of government transfers and associated taxes are offset by changes in private transfers and thus do not alter the resources available for own lifetime consumption. The key, then, is to identify what the lifetime budget constraint would be in the absence of these government transfers and taxes. The private transfer offsets envisioned under our conceptualization include *inter vivos* offsets in addition to the bequest offsets considered in previous studies.¹⁰ These different types of offsets have different effects on observed assets and private transfers. In our descriptions of these effects, we assume the absence of

capital market imperfections and focus on social security taxes and transfers to simplify the exposition.

In the bequest offset case, both observed assets and private transfers may be affected. Consumers under this case can be conceptualized as altering their bequests to succeeding generations by drawing down assets (or borrowing) to make their social security tax payments and saving their social security benefits. Thus, observed assets will be less than they would be in the absence of social security by the amount of past accumulated taxes less benefits. If a positive lifetime wealth increment is experienced, for example, this increment is bequeathed to subsequent generations to recompense those generations that experience negative lifetime wealth increments. Each succeeding generation may add to, subtract from, or simply pass through the induced inheritance that they receive from the preceding generation, depending on whether their own lifetime wealth increment from social security is positive, negative, or zero, respectively. The observed assets of intergenerational consumers, then, may be affected by this "pass through" component in addition to the component determined by their own past accumulated taxes less benefits. Similarly, prospective private transfer receipts may include "pass through" components and recompense components in addition to the prospective private transfer receipts that would obtain in the absence of social security.

In the *inter vivos* offset case, the effect of social security taxes and benefits on observed assets depends on the timing of the private transfer offsets. If the private transfer offsets are immediate, observed assets are not affected; if the private transfers offsets are not immediate, observed assets will be affected until the offsetting private transfers are made. Prospective private transfer receipts, of course, will be affected by

prospective social security taxes and benefits. Two types of *inter vivos* offsets can be envisioned. First, the social security program may simply replace a system of private transfers that would otherwise have occurred; e.g., instead of making direct, private, payments to support their retired parents, children now make tax payments to the government, which transfers the taxes to their parents, perhaps providing a more politically efficient or socially acceptable transfer mechanism. In this case, observed private transfer receipts to beneficiaries are smaller than they would be in the absence of social security by the amount of the social security benefits. A second type of *inter vivos* offset may occur if intergenerational consumers negate public transfers that would not have occurred privately by making offsetting private transfers. In this case, observed private transfers to taxpayers will be larger than in the absence of social security, as benefit recipients return to taxpayers that portion of the transfer that would not have occurred privately.

None of these intergenerational offset models is likely to hold in the extreme, of course, but it does seem likely that each has some validity. Social Security undoubtedly has reduced the need for private transfers from children to aged parents, for example, thereby replacing to some extent the offsetting *inter vivos* private transfers that would have occurred in the absence of Social Security. Similarly, the lifetime wealth increment experienced by some retirees under Social Security has undoubtedly resulted in an increase in their *inter vivos* transfers and bequests to offspring. Thus, regardless of whether or not one adheres to the rigorous overlapping generations utility function interpretation of intergenerational consumption behavior, it does appear likely that Social Security has affected the size and nature of intergenerational private transfers. The omission of these effects

from the consumption function specification may bias the remaining parameter estimates and lead to misinterpretations of the effect of social security on consumption behavior.

Several modifications of the consumption function are required under our conceptualization of pure intergenerational behavior. First, social security wealth and other government transfers wealth do not appear as arguments in the consumption function for intergenerational consumers, who are assumed to effectively ignore such wealth in the consumption decision. Second, some proportion (f) of prospective private transfers wealth may be intended as an offset for social security taxes and transfers or for other government transfers and their associated taxes; only the remaining proportion ($1-f$) is considered by intergenerational consumers as available for their own lifetime consumption and bequests. As implied in the discussion above, f may be negative in the case where social security transfers replace previously-existing private transfers, so that observed transfers are less than would obtain in the absence of social security, but otherwise is bounded on the unit interval.

An analogous offset proportion (δ) may apply to observed fungible assets; i.e., this proportion is intended for subsequent private transfer to other cohorts and therefore is unavailable for the normal lifetime consumption and bequests of the intergenerational consumers currently holding the assets. As implied by the discussion above, δ is bounded above by one, but may be negative at ages for which past accumulated taxes less benefits are positive.

Although not noted in previous analyses, the effective tax rates on labor and capital income faced by intergenerational consumers may also differ from those actually observed (and faced by life cycle consumers); i.e., some portion of observed tax payments for intergenerational consumers may be

effectively negated by offsetting private transfers. In effect, then, these taxes are not paid by intergenerational consumers and have no effect on their budget constraint or optimizing behavior. In the estimates presented below, pure intergenerational consumers were assumed to effectively offset all government taxes and transfers except for the taxes implicitly required to finance current government consumption expenditures.

Incorporating all of these factors into equation (3.1), then, the corresponding consumption function for pure intergenerational consumers can be written as

$$c = k [(1-\delta)A + \lambda P], \quad (4.2)$$

where

$$P = HF(1-t_G) + (1-f)IF, \quad (4.3)$$

δ and f represent the fungible assets and private transfers wealth offset parameters, and t_G represents the effective tax rate on labor income for intergenerational consumers.¹¹

Since it is not known *a priori* which individuals in the sample are life cycle consumers and which are intergenerational consumers, we introduce an additional parameter ψ to represent the proportion of life cycle consumption behavior. In the consumption function estimates presented below, the parameters ψ , r_H , r_S , r_O , and r_I are assumed to be identical across all individuals in the sample. As indicated above, the parameter λ is initially assumed to be an increasing function of the ratio of fungible assets, A , to prospective wealth, P . The consumption-wealth ratios, k , are assumed to be identical across consumers within relatively narrow (generally three year) age groups, but are free to vary across age groups. This general specification allows the consumption-wealth ratios to follow virtually any pattern across age groups, permitting wide flexibility in reflecting possible changes over

the life cycle in underlying factors, such as rates of time preference and tastes for contingency and bequest saving. The intergenerational offset parameters δ and f are assumed identical across consumers of a given age, and a variety of functional forms were tested specifying the pattern of these parameters across single years of age. The general form of the consumption function specification, then, combining both life cycle and intergenerational consumption behavior, is given by

$$c_a = k_a W_a, \quad (4.4)$$

where

$$W_a = A_a [\psi + (1-\psi)(1-\delta_a)] + \lambda P_a, \quad (4.5)$$

and

$$P_a = HF_a [1-t_G + \psi(t_G - t_H)] + \psi(SF_a + OF_a) + (1-f_a + \psi f_a) IF_a. \quad (4.6)$$

V. Data

The data used in this study were drawn from the 1982-83 Consumer Expenditure Survey (CES). Although the basic survey unit in the CES is the family, the requirements of our simulation model application dictated the choice of the individual as the consumption unit. Consequently, some of the variables reported in the CES for the family unit, including consumption expenditures, assets, welfare and public assistance benefits, and private transfer receipts, were allocated equally among adult members of the consumer unit for our estimation purposes.¹² Workers' compensation and unemployment compensation benefits were allocated equally among adult members of the CES unit aged less than 60. Other variables were generally available on an individual basis. These variables include wages and salaries, self-employment income, social security benefits and tax payments, and Supplemental Security Income benefits. Unfortunately, data were not available to permit the construction of a reliable pension wealth variable for persons not currently receiving pension income, so pension wealth is excluded from our consumption function specification.

We adopted the CES definition of total consumption expenditures with a number of exclusions that were required for our simulation model application. Some of the CES expenditure categories, such as deductions for retirement plans, private pensions, and social security, were excluded from our measure to achieve greater compatibility with the National Income and Product Accounts definition of personal consumption expenditures. In addition, the CES data on private transfer expenditures and charitable contributions were treated as a separate expenditure category in our application and consequently excluded from our definition of consumption expenditures.

Asset data in the 1982-83 CES includes information on checking, saving, brokerage, and similar accounts; stocks, bonds, mutual funds, and similar securities; amounts owed by and to the CES unit; and the market value of owned home. The remaining mortgage balance on owned home was not available in the 1982-83 CES, so it had to be imputed. Information on the reported current interest component of the mortgage payment was used to infer the date, term, and interest rate of the mortgage, using a historical time series of average contract mortgage interest rates and terms of new home mortgages. The implied remaining mortgage balance was then checked for consistency with the reported market value of the owned home.

Labor income was calculated as wages and salaries plus the estimated labor component of self-employment income. The labor component was estimated by assuming that labor's share of self-employment income was identical to the share of labor compensation in aggregate output other than proprietor's income in 1983.¹³ Other government transfers data in the 1982-83 CES include unemployment compensation, workmen's compensation, veteran's payments, public assistance and welfare benefits, and Supplemental Security Income payments. Private transfers receipts reported in the 1982-83 CES include "regular contributions received from alimony or child support and other sources" and "lump sum payments from estates, trusts, royalties, alimony, child support, prizes or games of chance or from persons outside the consumer unit."¹⁴

Two types of data were drawn from the 1982-83 CES for use in this analysis — average age profiles and individual regression observations. Average age profiles were developed for the reciprocity rate and per recipient income flow variables used in the prospective wealth estimates, as discussed in Section III.¹⁵ Each age profile was developed by computing weighted averages across all valid observations for the four CES Quarterly Interview

surveys for calendar year 1983. Each profile was then twice smoothed by taking five year moving averages, and extrapolated to age 100 (the assumed maximum age) for ages beyond 89, the last single year of age reported in the CES.

The individual regression observations were drawn from three panels in the 1982-83 CES whose last interview quarter ended in either the third or fourth quarter of calendar year 1983 or the first quarter of calendar year 1984. These panels were chosen because their final four interview quarters were most closely centered on calendar year 1983. A total of 4423 observations in these three panels had at least one quarter of valid responses to all of the CES questions used to construct the regression observations. To improve the reliability of the data, only observations with valid responses to the relevant questions in all four interview quarters were selected for our sample, leaving a total of 2884 observations. These observations were further checked, where possible, for clearly erroneous or apparent outlier data values. A total of 13 observations were arbitrarily excluded because of the large negative net assets reported (net debt exceeding \$25,000) and an additional 156 observations were excluded because reported social security tax payments exceeded the legal maximum. The remaining sample used in the regression analysis contained 2715 individual observations.

VI. Estimation

The final form of the consumption function specification corresponding to (4.4) can be written as

$$c_i = g(X_i, \theta) + \epsilon_i \quad i = 1, \dots, N \quad (6.1)$$

where c_i is consumption expenditure for the i th individual, X_i is a vector of exogenous variables for the i th individual, θ is a vector of unknown parameters, and ϵ_i is an unobservable error term. We assume that the error terms are independently distributed with mean zero and variance σ^2 . In our application, g is the nonlinear function of the parameters defined by (4.4)-(4.6), and many of the parameters are constrained from above or below *a priori* by constants.

We experimented with several alternative methods of constrained nonlinear estimation. The method we finally adopted is the technique introduced by Marquardt (1963). The algorithm begins with a set of arbitrary seed points for the parameters and then proceeds in an iterative fashion to adjust the parameter estimates until there is no further reduction in the error sum of squares. At each step of the process the parameters are updated by computing a matrix of partial derivatives of g with respect to the parameters and evaluating this matrix at the current parameter values. The adjustment to the parameters is obtained by computing a ridge regression in which the matrix of partial derivatives is regressed on the vector of residuals from the previous step. At each step the parameter estimates are checked against the theoretical bounds, and the estimate is set equal to the bound if it otherwise falls outside the bound. If a parameter remains on a bound after the procedure has converged, the estimation is recalculated with that parameter

set equal to the bound; this allows us to estimate conditional standard errors for the remaining parameters.

We begin by presenting the estimation results using a general specification of the λ function. Recall that λ is the proportional adjustment suggested by Zeldes (1989) that distinguishes uncertain prospective wealth from certain assets. Initially, we adopt the following general quadratic specification for the λ function

$$\lambda = av^2 + bv + c \quad (6.2)$$

where v is the ratio of certain assets to uncertain prospective wealth. The parameters a , b , and c are expressed in terms of λ by an arbitrary selection of three values for v . The parameters λ_0 , λ_1 , and λ_2 are defined by (6.2) for v equal to 0, .15, and .30, respectively.¹⁶

In our initial specification, the intergenerational asset offset parameter δ_a is a quadratic function constrained so that $\delta_{18}=0$, the minimum occurs at age 50 ($\delta_{50}\leq 0$), and $\delta_{100}\leq 1$. Given these constraints, the entire δ_a profile is determined by the estimation of the single parameter δ_{100} . This specification was chosen because net social security transfers typically dominate other government transfers, and net social security transfers (taxes) tend to be negative early in the life cycle, become increasingly positive as retirement age approaches, and are positive towards the end of the life cycle. If such net transfers are not immediately offset by intergenerational consumers, a δ_a profile similar to that suggested by this initial specification might result. This specification is also flexible in allowing a zero estimate for δ_{100} , suggesting no intergenerational offset of fungible assets at any age.

A wide variety of age patterns are possible for the intergenerational private transfers wealth offset parameter f_a . As noted in Section IV, the age

pattern of this parameter depends on the relative importance of bequests and the two types of *inter vivos* transfers as means of effecting intergenerational offsets. In our initial specification, f_a is defined as a linear function of age, with the end points, f_{18} and f_{100} , bounded above by one. This specification encompasses the full range of possible bequest and *inter vivos* offset behavior described in Section IV and also allows for the case where $f_a=0$ for all a , suggesting no intergenerational offset of private transfers wealth.

In our full consumption function specification, then, there are 33 parameters: 22 consumption-wealth ratios (k_a), one for each of 22 age groups;¹⁷ three parameters of the quadratic function in the certain-uncertain wealth ratio (λ_0 , λ_1 , and λ_2); the proportion of life cycle behavior (ψ); one parameter of the intergenerational asset offset function (δ_{100}); two parameters of the intergenerational private transfers wealth offset function (f_{18} and f_{100}); and four discount rates, one corresponding to each type of prospective wealth (r_H , r_S , r_O , and r_I).

The results of the estimation for several alternative specifications are presented in Table 1. For brevity, only three of the consumption-wealth coefficient estimates are displayed, for the age groups 18-20, 48-50, and 85-89. The relative age pattern of these estimates was similar across specifications and is illustrated in Figure 1 and discussed more fully in the following section.

Column 1 of Table 1 shows the parameter estimates for the case in which the discount rates are specified exogenously. For each of the components of prospective wealth, we assumed a real riskless discount rate of .030. The real expected rate of growth over time in the appropriate cross-section or individual income variable was assumed to be .02 for all but gross social

security wealth for current beneficiaries; current beneficiaries were assumed to perceive their current benefits to remain constant in real terms. Under this specification, the λ function was estimated to be very nearly linear and close to 1 throughout its range. The consumption-wealth coefficients k_a increased with age, as expected, but were all relatively small. The proportion of life cycle behavior ψ was estimated at its upper bound of 1, suggesting perfect substitutability of social security and other government transfers wealth for human wealth and implying the absence of intergenerational behavior. As such, the intergenerational offset behavior parameters, δ_{100} , f_{18} , and f_{100} , effectively drop out of the specification.

The results for the specification in which the four discount rates were estimated is given in Column 2 of Table 1. The fit of this equation is much better than the corresponding equation with exogenous discount rates, as indicated by the error sum of squares (SSE). We tested the hypothesis that the discount rates were equal to those used in Column 1 using the likelihood ratio procedure suggested by Gallant (1975). The hypothesis is rejected at a .1% level of significance.¹⁸ It is interesting to note the estimated values and precision with which the discount rates are estimated. The estimated discount rates r_H , r_S , r_O , and r_I are .081, -.020, -.022 and .051, respectively. Discussion of the relative magnitudes of these estimates, including the interesting result that the estimates of r_S and r_O are significantly negative, is presented below.

The estimates of λ_0 and λ_1 given in Column 2 of Table 1 (.992 and .996, respectively) are very close to one, and the estimate of λ_2 is on its theoretical upper bound of 1. This led us to consider the computationally-simpler specification in which λ was set equal to one, independently of the ratio of certain assets to uncertain prospective wealth. The results of this

estimation are presented in Column 3 of Table 1. The likelihood ratio test of the joint hypothesis $\lambda_0 = \lambda_1 = \lambda_2 = 1$ cannot be rejected at significance levels as high as 50%,¹⁹ and the estimates of the other parameters were not sensitive to this change in specification. We therefore conclude that it is not necessary to distinguish between certain assets and uncertain prospective wealth, as suggested by Zeldes, at least as long as the specification allows the data to dictate the appropriate discount rates applied to the various components of prospective wealth.

One of the more questionable aspects of our specification is the inclusion of other government transfers and private transfers as part of prospective wealth. We consider these two categories of prospective wealth as the most likely to suffer from measurement errors.²⁰ We thus estimated a consumption function in which other government and private transfers wealth were excluded in order to test the sensitivity of our qualitative conclusions to their presence. The results of this estimation are reported in Column 4 of Table 1. Although the estimates of the key parameters are not qualitatively affected by the exclusion of other government and private transfers wealth, the SSE is increased by nearly 5% when these variables are excluded. This suggests that these variables add significant explanatory power to the consumption function.

We performed the test suggested by Davidson and MacKinnon (1981) to formally distinguish between the non-nested hypotheses represented by Columns 3 and 4 of Table 1. When the specification of Column 3 is the hypothesis being tested and the specification of Column 4 is the alternative hypothesis, the Davidson-MacKinnon P test statistic is 2.043; the test statistic when the hypotheses are reversed is 11.44. Since these test statistics are asymptotically distributed according to the standard normal distribution when

the tested hypothesis is true, these tests lead to a rejection of both specifications at a 5% level of significance. Nevertheless, the relative magnitudes of the test statistics favor the specification in Column 3 over that in Column 4; at a 1% level of significance, for example, this test would accept the model in Column 3 and reject that of Column 4. We therefore consider the specification in Column 3 of Table 1 to be our preferred specification.

The Column 3 specification provides some evidence of intergenerational *inter vivos* offset behavior in that both f_{18} and f_{100} are estimated at their upper bound of 1. Although these parameter estimates are extreme, they are consistent with social security and other government transfer recipients increasing their private transfers as offsets to taxpayers. This evidence must be regarded as weak, however, since the standard errors for both f_{18} and f_{100} were relatively large (8.27 and 2.07, respectively) at the point that these parameters approached their upper bound. There is no evidence of intergenerational asset offset behavior in the Column 3 specification, as the estimated δ_a function has a zero value for all ages. The standard error of the δ_{100} estimate, however, was relatively large (.64) as this parameter approached its lower bound.

Because the standard errors of the δ_{100} , f_{18} , and f_{100} estimates were relatively large as they approached their bounds, seven additional regressions were estimated to test the sensitivity of the other parameter estimates in our preferred specification to extreme values assumed for the asset and private transfers wealth offset parameters. We considered the following combinations of extreme values for the δ_{100} , f_{18} , and f_{100} parameters: (0,1,0), (0,-10,-10), (0,0,-10), (1,1,1), (1,1,0), (1,-10,-10), and (1,0,-10). The general effect of these alternatives is to increase the ψ estimate, decrease

the r_H estimate, and increase the other discount rate estimates (r_S , r_O , and r_I). Not surprisingly, the parameter estimates were most affected, particularly for ψ and r_I , in the cases where f_{18} or f_{100} were set to the -10 extreme; in these cases, the estimate of ψ ranged from .702 to .999, and the estimate of r_I ranged from .090 to 1.604. In the remaining cases, the changes were less extreme, with ψ ranging from .652 to .720 and r_I ranging from .049 to .071. Across all of the cases, r_H ranged from .056 to .085, r_S from -.021 to -.006, and r_O from -.023 to -.004. Standard errors were not substantially affected except for the r_I estimates, whose standard errors generally increased commensurate with the point estimates. In none of the cases did the SSE change appreciably. These tests suggest that the r_H , r_S , and r_O parameter estimates do not appear qualitatively sensitive to extreme values assumed for the intergenerational offset parameters. Even the ψ parameter estimates in the most extreme cases we tested would leave our qualitative conclusions unchanged. Clearly the most sensitive parameter estimate in these tests is r_I , which is not central to our policy conclusions.

Although the functional forms chosen for the intergenerational offset functions δ_a and f_a are both reasonable and flexible, they are arbitrary impositions on the consumption function specification. This raised the concern that the specific nonlinear functional form chosen for the δ_a function, in particular, might bias the other parameter estimates and possibly account for the lack of intergenerational asset offset evidence in our preferred specification. This concern prompted us to estimate an additional consumption function specification using an alternative functional form for δ_a . Recall that for the results presented thus far, the δ_a function is a quadratic with δ_{18} equal to 0 and a minimum at age 50. Column 5 of Table 1 presents estimates under the alternative specification of the δ_a function as

bilinear (two linear segments) with a bend point at age 50 and each linear end point estimated as a parameter. These results show evidence of asset offset behavior with δ_a increasing with age. The estimates of the other parameters are not qualitatively affected, although the estimate of ψ is slightly lower and the estimate of r_H is slightly higher than the corresponding estimates in Column 3. While we offer Column 5 as evidence of intergenerational asset offset behavior, we note that the improvement in fit is very slight and that the precision with which the asset offset parameters are estimated is not very good. Moreover, the Davidson-MacKinnon P test indicates a preference for the Column 3 specification over the Column 5 specification; the P test statistic is 1.878 when the Column 3 specification is the tested hypothesis and the Column 5 specification the alternative hypothesis, while the P test statistic is -2.312 when the hypotheses are reversed. This leads to acceptance of the Column 3 specification, and rejection of the Column 5 specification, at the 5% level of significance. Consequently, we retain the model of Column 3 as the preferred specification.

VII. Implications

A. Consumption-wealth ratios.

The consumption-wealth ratios k_a for each of the 22 age groups that were estimated in the preferred specification (Column 3 of Table 1) are plotted (as the diamond points) in Figure 1. The ratios are estimated fairly precisely and, as anticipated, have the property of generally increasing with age. The smallest ratio (.009) is for the youngest age group (ages 18-20) and the largest ratio (.092) is for one of the oldest age groups (ages 75-79). One striking feature of the profile is the low level of the consumption-wealth ratios at even the oldest ages. The small magnitudes of these ratios suggest strong bequest or contingency motives for saving.

If we consider the special case in which there is no uncertainty, a constant rate of time preference ρ , a constant taste for bequests parameter α , and logarithmic utility, the consumption-wealth ratios can be derived as

$$k_a = [\alpha + \sum_{i=a}^D S_{a,i} (1+\rho)^{a-i}]^{-1}, \quad (7.1)$$

where, as before, $S_{a,i}$ is the probability of surviving from age a to age i . We obtained estimates of α and ρ by applying nonlinear regression to equation (7.1) using the estimates of k_a from our preferred specification as the dependent variable. The resulting estimates of α and ρ were 5.4549 and .0381, respectively, and the fitted equation is plotted (as the solid line) in Figure 1. This example illustrates the implication of substantial bequest or contingency saving motives in our estimates. The consumption-wealth ratios predicted by this approach at ages 95 to 100 in our preferred specification only range from .12 to .15; in the absence of bequest and contingency saving motives, the consumption-wealth ratio would be 1 at age 100. The strength of

this estimated bequest/contingency saving effect extends to the early ages as well; using the estimated time preference rate of .0381, for example, but assuming no bequest or contingency saving motive ($\alpha=0$), equation (7.1) predicts consumption-wealth ratios of about .043 for the 18-20 age group, compared to the .009 estimate for that age group in our preferred specification. Even if a time preference rate of zero ($\rho=0$) is assumed in the absence of the bequest/contingency saving motive ($\alpha=0$), equation (7.1) predicts a consumption-wealth ratio of over .017 for the 18-20 age group.

There are other explanations for the low consumption-wealth ratio estimates in our preferred specification. First, since our data measure consumer expenditures and not the value of consumption services, there is a measurement error in our consumption data that could result in a systematic underestimation of the consumption-wealth ratios at older ages. Second, it is possible that consumption-wealth ratios are lower, especially at the early stages of the life cycle, because of capital market imperfections. There have been several studies of consumption behavior in which borrowing constraints have been shown to play a central role.²¹ Borrowing constraints may explain the apparent systematic pattern of the age group point estimates below the constant alpha-rho profile in Figure 1 at the earliest ages. Arguing against this conclusion, however, is the failure of the specifications with λ as an increasing function of the certain-uncertain wealth ratio (Columns 1 and 2 in Table 1) to add significant explanatory power or support a distinction between fungible assets and prospective wealth.

B. Relative uncertainty adjustments and risk premiums.

One of the novel features of our specification is the treatment of uncertainty in the evaluation of prospective wealth. We have estimated

discount rates for each of the four different types of prospective wealth. If we let r_x denote the estimated discount rates, r denote the real after-tax riskless rate of interest, and g denote the real rate at which the appropriate cross-section or individual income variables are expected to increase, then the risk premium for each type of prospective wealth can be represented as

$$\pi_x = (1+g)(1+r_x) - (1+r), \quad x = H, S, O, I. \quad (7.2)$$

The risk premiums corresponding to the estimated discount rates in our preferred specification (Column 3 of Table 1) are given in Table 2 for selected values of r and g .

The estimated discount rates for social security wealth and other government transfers wealth are significantly negative. The corresponding risk premiums in Table 2 range from about $-.05$ to $-.02$ for both social security wealth and other government transfers wealth, depending on the values assumed for r and g . This implies that the insurance aspects of the corresponding government programs outweigh the uncertainty associated with possible future program changes. The net effect of including social security wealth and other government transfers wealth in the individual portfolio, then, is to reduce overall portfolio risk, even relative to riskless assets. In the case of social security wealth, this result contrasts sharply with public opinion surveys that consistently reveal a lack of confidence in the ability of the Social Security program to pay benefits in the future. Our finding is consistent, however, with the broad public support for the Social Security program evidenced in the same public opinion surveys as well as in the political arena.²²

The risk premium for human wealth ranges between $.054$ and $.085$ in Table 2, depending on values assumed for r and g . This magnitude is not unreasonable, given the variability in expected wages and the uncertainties

associated with possible unemployment. The most significant point, however, is the relationship between the human wealth, social security wealth, and other government transfers wealth risk premiums. Human capital tends to increase the overall risk of the portfolio relative to riskless assets, whereas social security wealth and other government transfers wealth tend to reduce it. The important policy implications of this finding are discussed more fully below.

The risk premiums for private transfers displayed in Table 2 range from .022 to .053; these values suggest that such transfers are perceived as more certain than prospective labor income but less certain than prospective social security and other government transfers. As noted above, data measurement problems associated with the other government transfers wealth and private transfers wealth variables give us less confidence in the corresponding discount rate estimates, compared to those for human wealth and social security wealth. Nevertheless, all of our discount rate estimates can be rationalized on the basis of economic considerations and are relatively precise, as indicated by their standard errors.

We recognize that the exponential risk adjustment specification that we have estimated is only one of many possible ways to account for differential risk among the four types of prospective wealth. Perhaps the most straightforward alternative is to discount all future flows at the riskless rate and then multiply the present value of each type of wealth by a parameter that measures the relative degree of substitutability among the wealth types.²³ This substitutability parameter can be interpreted in part as reflecting relative risk differences among the types of consumption wealth. We estimated the consumption function under this specification and found the prospective wealth component of the function to be

$$P_a = .585 HF_a (1-t_H) + 1.235 SF_a + 1.123 OF_a + .318 IF_a \quad (7.3)$$

(.055)
(.112)
(.150)
(.057)

where HF_a , SF_a , OF_a , and IF_a are present values of the four types of prospective wealth, and the standard errors are in parentheses.²⁴ The SSE for the estimated consumption function was 58.4897×10^9 , which is somewhat higher than our preferred specification with exponential risk adjustments. We performed the Davidson-MacKinnon P test to distinguish between our preferred specification of Column 3 of Table 1 and the analogous specification with proportional risk adjustments. When our preferred specification was the tested hypothesis, the P test statistic was 2.358. When the specification with proportional risk adjustments was the tested hypothesis, the P test statistic was 6.869. These tests reject both hypotheses at a 5% level of significance, but accept our preferred specification with exponential risk adjustments, and reject the specification with proportional risk adjustments, at a 1% level of significance.

Note that the relative magnitudes of the coefficients in the proportional specification (7.3) are generally consistent with the relative magnitudes of the estimated discount rate parameters in our preferred exponential specification; i.e., the coefficients of the HF_a and IF_a terms are considerably smaller than the coefficients of SF_a and OF_a . The coefficients in excess of 1.0 for social security wealth and other government transfers wealth in (7.3) are also consistent with negative risk premiums on these wealth types; i.e., it implies that social security wealth and other government transfers wealth substitute for fungible assets on more than a one-to-one basis.²⁵ This alternative specification, then, lends support to the qualitative conclusions suggested by our preferred specification.

C. Effects of Social Security Wealth on Consumption

The parameter ψ in our specification can be interpreted as the proportion of life cycle, as opposed to intergenerational, consumption behavior. This parameter is bounded below by zero and above by one. If ψ is equal to zero, the consumer disregards social security and other government transfers wealth in making the consumption decision; if ψ is equal to one, then the consumer views social security and other government transfers wealth as equivalent to other forms of prospective wealth. The estimate of ψ in our preferred specification is .609 and its standard error is .072. This estimate is more precise than the corresponding coefficient of social security wealth typically estimated in asset equation models²⁶ and implies a 95% confidence interval of [.467, .750] for the proportion of life cycle behavior in the population.

It is important to understand the relationship between ψ and the risk premiums on social security and other government transfers wealth. A smaller risk premium on social security wealth, for example, implies that a stream of expected net social security benefits will have a greater effect on consumption for life cycle consumers than an identical stream of expected earnings, whereas a value of ψ less than one simply indicates the proportion of social security wealth that affects current consumption. In effect, the risk premiums for social security and other government transfers wealth are not defined for intergenerational consumers, since they ignore this wealth in the consumption decision.

Pure life cycle behavior would be represented by a specification in which ψ is equal to one and there are no intergenerational offset mechanisms. If we consider our preferred specification of Column 3 of Table 1 as the tested hypothesis and the pure life cycle hypothesis as

$$H_0 : \psi=1, \quad \delta_{100}=0, \quad f_{18}=0, \quad f_{100}=0, \quad (7.8)$$

we can test H_0 using the procedure suggested by Gallant (1975). The computed F test statistic is 5.414, which indicates rejection of H_0 at a .1% level of significance.

The other extreme possibility is pure intergenerational behavior, in which we have the same tested hypothesis, and the null hypothesis is given by

$$H'_0 : \psi=0. \quad (7.9)$$

In this case, all social security and other government transfer receipts are completely offset, and consumption does not depend on social security or other government transfers wealth. This hypothesis can be tested using the ratio of the estimate of ψ to its asymptotic standard error, which follows the standard normal distribution under H'_0 . The computed ratio is 8.407, which indicates rejection of H'_0 at a .1% level of significance. We thus reject both the pure life cycle and pure intergenerational hypotheses.

The estimate of .609 for ψ in our preferred specification suggests that social security wealth has a substantial impact on consumption and saving. Moreover, the negative risk premium estimated for social security wealth magnifies the extent to which expected social security taxes and benefits affect consumption behavior. To illustrate the potential importance of these effects, we have used the parameter estimates from our preferred specification to predict the effect on aggregate consumption of eliminating social security wealth.²⁷ The predicted decline in aggregate consumption is 18% when social security wealth is eliminated. This estimate is much larger than previous estimates based on time series analysis. The revised estimate obtained by Feldstein (1982), for example, implies a reduction in aggregate consumption of about 5%.²⁸ Our much larger estimate is interesting but should be regarded with skepticism at this point because of the crudeness of the aggregate consumption estimate, because the partial equilibrium analysis ignores all of

the price changes and associated portfolio, labor supply, and other adjustments that would occur if Social Security were abolished, and because the underlying parameter estimates require confirmation by other studies. Nevertheless, the estimated magnitude of this partial equilibrium effect suggests that Social Security may have had a substantial impact on private consumption and saving.

It should be emphasized that part of the estimated reduction in saving caused by social security wealth represents a gain, rather than a loss, in efficiency; i.e., by providing alternatives to imperfect private annuity and insurance markets, Social Security permits increased consumption, improving the welfare of consumers. This part of the implied increase in consumption might be defined in terms of the effect of the negative risk premium estimated for social security wealth. To illustrate the potential magnitude of this effect, the aggregate consumption exercise described above was repeated under the assumption of a zero risk premium for social security wealth.²⁹ In this exercise, aggregate consumption falls by only 7% if social security wealth is eliminated.³⁰ This suggests that more than half of the implied aggregate consumption increase caused by social security wealth is associated with an efficiency gain, derived from the insurance aspects of the program, and less than half of the implied consumption increase is associated with an efficiency loss, derived from the life cycle asset substitution effect of a pay-as-you-go social security program.

An alternative measure of the relative sizes of the consumption components associated with these efficiency gains and losses can be derived by comparing the effect on aggregate consumption when social security wealth is assumed to be as risky as human wealth; i.e., in the absence of the efficiency-enhancing aspects of the program, the risk characteristics of

social security wealth may be more akin to human wealth than to riskless assets, since prospective earnings determine the level of a worker's participation in the Social Security program. When the aggregate consumption exercise is repeated with the discount rate for social security wealth set equal to the estimated discount rate for human wealth (.084), aggregate consumption falls by only 2% if social security wealth is eliminated. From this perspective, nearly all of the increase in consumption caused by social security wealth could be associated with an efficiency gain.

To repeat, these exercises should only be regarded as illustrative, given the simplistic, partial equilibrium, nature of the analyses. In a broader context, all of these partial equilibrium estimates of the aggregate consumption induced by social security wealth are likely to be overstated because they ignore the smaller unplanned bequests that also result from the insurance aspects of Social Security. That is, for each generation, the consumption gains provided by these insurance aspects are offset to some extent, and possibly even reversed, by the correspondingly smaller inheritances left by the preceding generations. Given their initial endowments, however, and abstracting from general equilibrium effects, each generation experiences a welfare gain from these insurance aspects.

If changes in initial endowments and general equilibrium effects are taken into account, of course, the introduction of Social Security may result in reduced lifetime consumption for many cohorts through its potential negative effect on the capital stock and the rate of economic growth. This is by no means certain, however, even if one accepts our parameter estimates, because the effect of Social Security on aggregate economic activity depends critically on its interaction with the business cycle and stabilization policy; e.g., by encouraging private consumption, the introduction of a pay-

as-you-go social security system may increase aggregate consumption, output, and saving if introduced during a period of excess capacity.³¹ In any event, the consideration of changes in initial endowments and general equilibrium effects on subsequent generations implies that the evaluation of policy can only be conducted in the context of an intergenerational or intertemporal social utility function. In terms of the simpler guideline that the removal of market distortions enhances efficiency in the context of consumer sovereignty, our results suggest that Social Security may have encouraged consumption in part by reducing some market imperfections but also in part by introducing other distortions. Our crude aggregate consumption exercises suggest that the induced consumption component associated with efficiency gains exceeds the component associated with efficiency losses.

VIII. Concluding Remarks

The approach taken in this paper to estimating the effect of social security and other forms of wealth on consumption behavior has several advantages over previous approaches. Our approach allows the estimation of the subjective discount rates attached by consumers to the different forms of prospective wealth. Our estimates of these rates vary widely across the different forms of prospective wealth and have important policy implications. Our approach also incorporates testable hypotheses about the existence of intergenerational consumers and the private transfer mechanisms that such consumers might use to offset the lifetime resource changes effected by government wealth. While our attempts to measure these offset mechanisms met with some success, we believe that subsequent studies should extend this effort in order to obtain more precise estimates.

Nearly all previous cross-section analyses of the effect of Social Security on consumption and saving have estimated an asset equation model,³² in contrast to our approach of estimating the consumption function directly. Both consumption and asset equation models require assumptions about how consumers perceive prospective wealth, but additional assumptions are typically required for the social security wealth coefficient in asset equation models to accurately reflect the displacement of fungible assets. One unrealistic assumption implicit in asset equation models is that past wealth expectations have been realized for the individuals in the sample. If past expectations are not realized, current assets will be affected, and the estimate of the social security wealth coefficient will be corrupted. The interpretation of the social security wealth coefficient is further muddled in previous asset equation models, because these models typically have been

developed in the context of the life cycle model.³³ Consequently, the social security wealth coefficient in these specifications reflects the degree of fungible asset displacement only if one assumes the absence of intergenerational behavior. These specifications provide little information discriminating between the competing life cycle and intergenerational hypotheses, and the models lose validity in the context of intergenerational consumption behavior.

The consumption function approach adopted in this paper avoids these difficulties. The current consumption decision depends on current and prospective wealth and does not require the assumption that past expectations have proved accurate. In addition, we incorporate parameters designed to uncover evidence of the existence and nature of the offset behavior engaged in by intergenerational consumers. If such evidence is difficult to find, it increases the likelihood that the social security wealth coefficient represents the imperfect substitutability of social security wealth for other types of consumption wealth rather than the relative prevalence of life cycle versus intergenerational consumption behavior. In either case, however, positive values for this parameter in our specification unambiguously imply that, *ceteris paribus*, increases in consumption are induced by social security wealth.

The empirical results presented in this paper have profound policy implications for two important issues related to the Social Security program. The first is the effect of Social Security on private saving and the second is the extent to which workers receive their "money's worth" from their Social Security tax payments. On both counts, critics of Social Security have advocated a reduction in the size of the program, changing to a partially or fully funded basis, or moving to private alternatives as ways of reducing or

eliminating the alleged adverse effects of the program. The results presented in this paper provide a much different perspective on these issues.

Our results do provide strong support for the proposition that reductions in social security wealth would have a positive and sizable effect on private saving. Our life cycle proportion parameter ψ is estimated more precisely than in most previous analyses and suggests that a dollar of social security wealth substitutes for about three-fifths of a dollar of fungible assets. Our estimate of the subjective discount rate applied to social security benefits and taxes implies a further amplification of their effect on consumption and saving. We interpret the implied negative risk premium for social security wealth, however, as a reflection of the insurance aspects of the program that provide an alternative to imperfect private annuity and insurance markets. Under this interpretation, our estimates suggest that any reductions in saving occasioned by Social Security largely reflect a gain, rather than a loss, in efficiency.

While we share the concern of others over our relatively low national saving rate, our results suggest that reducing the size of the Social Security program is an inappropriate policy for effecting increases in saving. A program reduction may have the intended effect of reducing consumption and increasing saving, but at a substantial cost in economic efficiency and the welfare of consumers. Program provisions other than the pay-as-you-go financing basis account for much of the risk reduction occasioned by Social Security. As such, moving to a partially or fully funded basis may be one way of increasing national saving without abandoning the welfare gains implied by our results. Preserving the risk-reducing characteristics of Social Security would be harder to achieve under privatization; even if participation were compulsory, for example, the diversification of risk would be less complete

with multiple private providers, and some risks, such as the risk of unexpected inflation, would require some form of government participation, such as the issuance of inflation-indexed bonds. Some of the risk reduction effected by Social Security, however, may be directly related to the pay-as-you-go payroll tax financing of the program. We noted earlier, for example, that social security wealth represents an asset with less expected intertemporal volatility in its internal rate of return than even that associated with long-term government bonds. Thus, changing the financing basis of the program could itself result in a consumer welfare loss. Moreover, changing the financing basis of Social Security, either through funding or privatization, would result in large intergenerational wealth redistributions from present to future generations unless other public dissaving is increased. Our own view is that the goal of increasing saving is better served by other more flexible, measured, and potent policy approaches than by privatizing or changing the financing basis of Social Security.

Our results also have important implications for Social Security money's worth issues. Current projections suggest a real internal rate of return to OASI tax payments of roughly two percent for new labor force entrants.³⁴ Consequently, studies that use real market interest rates of more than two percent to discount Social Security taxes and benefits are likely to conclude that new entrants, on average, can no longer expect to receive their money's worth from Social Security. One problem with all such analyses has been the difficulty of valuing the inflation-indexing and other insurance features of the Social Security program. Because market prices are not available for these features, their value has not been incorporated into the money's worth calculations. The subjective social security wealth discount rate estimated in this paper, however, provides one approach for incorporating the valuation

of the insurance features of Social Security into money's worth analyses. For reasonable growth rate estimates, Table 2 suggests that the real market discount rate appropriate for comparison with the Social Security internal rate of return is close to zero or negative; i.e., the value to consumers of the insurance features of the program implies the use of a discount rate 2 or 3 percentage points below the real riskless market interest rate. These results suggest, then, that even young workers can expect to receive more than their money's worth from Social Security when consumer valuation of the insurance features of the program are taken into account.

While our empirical results are provocative and have important policy implications, we caution that they should only be regarded as suggestive at this point. The general tenor of our results is reasonable, but we are skeptical about the magnitude of the implied negative risk premiums estimated for social security and other government transfers wealth. Before much confidence can be placed in these results, they must be replicated using other data, and the sensitivity of conclusions to changes in specification and perception assumptions should be tested.

Table 1. Parameter Estimates Under Alternative Specifications


Parameter	(1)	(2)	(3)	(4)	(5)
k_{18-20}	.004 (.000)	.009 (.002)	.009 (.002)	.008 (.002)	.010 (.003)
k_{48-50}	.045 (.003)	.050 (.005)	.051 (.004)	.049 (.007)	.053 (.008)
k_{85-89}	.070 (.010)	.084 (.014)	.086 (.011)	.088 (.012)	.108 (.028)
δ_{18}	-	-	-	-	-.843 (1.31)
δ_{50}	-	-	-	-	.195 (.424)
δ_{100}	-	[0]	[0]	.251 (.705)	.728 (.666)
λ_0	.934 (.024)	.992 (.037)	-	-	-
λ_1	.967 (.012)	.996 (.018)	-	-	-
λ_2	[1]	[1]	-	-	-
ψ	[1]	.614 (.093)	.609 (.072)	.578 (.085)	.563 (.084)
f_{18}	-	[1]	[1]	-	[1]
f_{100}	-	[1]	[1]	-	[1]
r_H	-	.081 (.023)	.084 (.014)	.089 (.024)	.089 (.028)
r_S	-	-.020 (.005)	-.020 (.005)	-.027 (.005)	-.021 (.006)
r_O	-	-.022 (.008)	-.022 (.008)	-	-.024 (.009)
r_I	-	.051 (.018)	.052 (.018)	-	.050 (.020)
SSE(10^9)	61.3948	57.5427	57.5573	60.3390	57.4849

Numbers in brackets [] are parameter values estimated on the bound.
Numbers in parentheses () are estimated standard errors.

Table 2. Risk Premiums for Different Wealth Types

Wealth Type	g=0		g=.02	
	r=.02	r=.03	r=.02	r=.03
Human	.064	.054	.085	.075
Social Security	-.040	-.050	-.020	-.030
Other government	-.042	-.052	-.022	-.032
Private transfers	.032	.022	.053	.043

Figure 1. Consumption-Wealth Ratios for the Preferred Specification:
Age Group Point Estimates vs. Constant α -rho Profile



Footnotes

1. These studies include the cross-section analyses of Feldstein and Pellechio (1979), Kotlikoff (1979), Feldstein (1983), Blinder, Gordon, and Wise (1983), Hubbard (1986), Mariger (1986), Bernheim (1987), and Kurz (1987). Time series analyses include Munnell (1974), Feldstein (1974), Barro (1978), Darby (1979), and Leimer and Lesnoy (1982). All of these time series studies except Leimer and Lesnoy used incorrect data for the social security wealth variable; Lesnoy and Leimer (1981) reestimated those specifications using correct data.
2. We use the term "intergenerational" to refer to the type of consumption behavior, characterized by Barro (1974), that offsets government transfers and associated taxes through changes in private transfers.
3. See, for example, Sibley (1975) and Yaari (1976).
4. A quadratic utility function yields certainty equivalence but has the undesirable property of implying increasing absolute risk aversion.
5. See Blanchard and Mankiw (1988) and Zeldes (1989) for references.
6. We were unable to include private pension wealth because of data limitations.
7. The internal rate of return to a mature, payroll-financed, pay-as-you-go, social security program is determined by the rate of growth in its payroll tax base (for example, see Aaron (1966)). Historically, the variance in this annual growth rate has been substantially less than the variance in annual yields of long-term government bonds or riskier assets. The standard deviation of the real change in aggregate wages and salaries over the period 1926-85 was 5.8 percent, for example, while the standard deviation of the average real yield on long-term government bonds over the same period was 9.9 percent (see Ibbotson Associates (1989)).
8. Our treatment of prospective earnings requires the implicit assumption that the determination of labor supply is exogenous to the consumption decision.
9. For simplicity, these perceptions for social security tax liability ignore the annual taxable maximum. A similar comment applies to the social security benefit perceptions for current non-beneficiaries in equations (3.9) and (3.10), which ignore the progressivity of the benefit formula.
10. Although generally couched in the context of the life cycle model, some previous cross-section specifications can be interpreted as consistent with bequest offset behavior. For example, Kotlikoff's (1979) description of the effect of past social security taxes and benefits on the life cycle asset profile is consistent with bequest offsets, but not with *inter vivos* offsets, which have a much different effect on the asset profile, as discussed below. Mariger (1986) explicitly considers intergenerational behavior, but his conceptualization of the effect of past social security taxes and benefits on the asset profile follows that of Kotlikoff.

11. The average tax rate applicable to asset income for life cycle consumers in 1983, including personal income taxes and corporate profits taxes, was estimated as 17.9 percent. The corresponding tax rate applicable to the asset income of intergenerational consumers was estimated as 13.7 percent. So that the effect of all other taxes could be represented simply in the lifetime budget constraint, the incidence of all other taxes was assumed to be on labor income. Under this assumption, the tax rate applicable to labor income for life cycle consumers in 1983 (t_H) was estimated as 33.3 percent, while the corresponding tax rate for intergenerational consumers (t_G) was estimated as 25.6 percent.
12. One consequence of this allocation is that the consumption of children under 18 is captured in the consumption observations of consumer unit adults. To test the effect of this consumption allocation on our parameter estimates, a variable (Z) for number of children under 18 was added to our preferred specification (described below) with a multiplicative parameter; i.e., $c = \alpha Z + kW$. The inclusion of the number of children variable had very little effect on the remaining parameter estimates, leaving our qualitative conclusions unchanged.
13. This share was estimated as .631 for 1983.
14. These definitions are from the documentation included on the 1982-83 CES computer tape, available from the Bureau of Labor Statistics.
15. The survival probabilities required for the prospective wealth estimates were derived, by single year of age, from cross-section mortality rates for 1983. These rates were provided by the Office of the Actuary of the Social Security Administration.
16. In this specification, the three arbitrary values chosen for v are not critical, as long as they are reasonably dispersed within the sample range of v .
17. The age range 18-74 was divided into 19 three-year age groups, and the age range 75-89 was divided into 3 five-year age groups.
18. The computed F test statistic was 44.885, which is significant at considerably less than a .1% level of significance.
19. The computed F statistic was .227; the probability that the F statistic will exceed this value under the null hypothesis is .878.

20. As a rough check, per capita profiles by age, estimated from the 1982-83 Consumer Expenditure Survey, were weighted by an estimate of the 1983 population by age, adjusted for consistency with the National Income and Product Account (NIPA) aggregate population for that year. These aggregate estimates were then compared to corresponding aggregates from the NIPA. The implied proportional adjustment for the other government transfers variable was nearly twice that for any of the other flow variables for which aggregates were available (consumption expenditures, labor income, social security benefits, and social security taxes). No aggregate was available for private transfers.
21. See, for example, Mariger (1986).
22. See Chapter 6 of Light (1985) for a discussion of public opinion surveys on Social Security issues.
23. A proportional specification is used in all of the cross-section analyses cited in note 1 above.
24. The same riskless discount rate (.03) and income growth rates (.02, or 0 for current beneficiaries' gross social security wealth) were used in this construction of prospective wealth as were used in the specification given in Column 1 of Table 1.
25. Feldstein (1983) and Bernheim (1987) estimate coefficients of social security wealth in asset accumulation equations that imply a more than a one-for-one substitution.
26. See the references to the cross-section studies in note 1 above. While the social security wealth coefficient estimates reported in these studies are in the same range as the estimate in our preferred specification, the standard error of our estimate is considerably smaller.
27. The effect of social security wealth on aggregate consumption was estimated by computing social security wealth and the associated induced consumption for prototypical persons of each age in 1983, using the estimates of k_a , ψ , and r_s from our preferred specification in Column 3 of Table 1. The consumption induced for prototypical persons of each age was then aggregated across ages using the estimated population by age in 1983 as weights. Prototypical persons were defined as having average social security taxes and benefits for persons of that age, as estimated from the 1982-83 CES.
28. Feldstein's finding of a significant and sizable social security wealth coefficient is not typical of published time series estimates using correct data for the social security wealth variable. Lesnoy and Leimer (1981) consider a wide variety of alternative perception assumptions, consumption function specifications, and time periods, but find little evidence in the time series data to support the hypothesis that social security has reduced personal saving.
29. See note 24 above.

30. This exercise also suggests that the negative social security wealth risk premium implied by our preferred specification corresponds to a 47% reduction in personal saving. It is interesting to note that this estimate falls within the 35-60% range simulated by Kotlikoff, Shoven, and Spivak (1987) as the potential partial equilibrium long run reduction in wealth associated with the annuity aspects of a social security system.
31. See Eisner (1983) for an elaboration of this point.
32. All of the cross-section analyses cited in footnote 1 except Mariger (1986) estimate an asset equation model in which assets at a given age are explained in part by social security wealth.
33. See note 10 above.
34. We derived this estimate from unpublished projections of the Office of the Actuary of the Social Security Administration using the intermediate (II-B) assumptions of the annual Trustees' Report to Congress for recent years. This estimate is also roughly consistent with projections made by Leimer and Petri (1981) as well as a more recent version of that model.

References

- Aaron, Henry, "The Social Insurance Paradox," *Canadian Journal of Economics and Political Science* 32 (August 1966), 371-374.
- Abel, Andrew B., "Precautionary Saving and Accidental Bequests," *American Economic Review* 75 (September 1985), 777-791.
- Barro, Robert J., "Are Government Bonds Net Wealth?," *Journal of Political Economy* 82 (November/December 1974), 1095-1117.
- Barro, Robert J., *The Impact of Social Security on Private Saving* (Washington, D.C.: American Enterprise Institute, 1978).
- Bernheim, B. Douglas, "The Economic Effects of Social Security: Toward a Reconciliation of Theory and Measurement," *Journal of Public Economics* 33 (August 1987), 273-304.
- Blanchard, Olivier Jean and N. Gregory Mankiw, "Consumption: Beyond Certainty Equivalence," *American Economic Review* 78 (May 1988), 173-177.
- Blinder, Alan S., Roger H. Gordon, and Donald E. Wise, "Social Security, Bequests and the Life Cycle Theory of Saving: Cross-sectional Tests," in Franco Modigliani and Richard Hemming (eds.), *The Determinants of National Saving and Wealth* (New York: St. Martin's Press, 1983).
- Darby, Michael R., *The Effects of Social Security on Income and the Capital Stock* (Washington, D.C.: American Enterprise Institute, 1979).
- Davidson, Russell and James G. MacKinnon, "Several Tests for Model Specification in the Presence of Alternative Hypotheses," *Econometrica* 49 (May 1981), 781-793.
- Eisner, Robert, "Social Security, Saving, and Macroeconomics," *Journal of Macroeconomics* 5 (Winter 1983), 1-19.
- Feldstein, Martin, "Social Security, Induced Retirement, and Aggregate Capital Accumulation," *Journal of Political Economy* 82 (September/October 1974), 905-926.
- Feldstein, Martin, "Social Security and Private Saving: Reply", *Journal of Political Economy* 90 (June 1982), 630-642.
- Feldstein, Martin, "Social Security Benefits and the Accumulation of Pre-retirement Wealth," in Franco Modigliani and Richard Hemming (eds.), *The Determinants of National Saving and Wealth* (New York: St. Martin's Press, 1983).
- Feldstein, Martin and Anthony Pellechio, "Social Security and Household Wealth Accumulation: New Microeconomic Evidence," *Review of Economics and Statistics* 61 (August 1979), 361-368.

- Gallant, A. Ronald, "Testing a Subset of Parameters of a Nonlinear Regression Model," *Journal of the American Statistical Association* 70 (December 1975), 927-932.
- Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy* 86 (December 1978), 971-987.
- Hubbard, R. Glenn, "Pension Wealth and Individual Saving," *Journal of Money, Credit, and Banking* 18 (May 1986), 167-178.
- Hubbard, R. Glenn and Kenneth L. Judd, "Social Security and Individual Welfare: Precautionary Saving, Borrowing Constraints, and the Payroll Tax", *American Economic Review* 77 (September 1987), 630-646.
- Ibbotson Associates, *Stocks, Bonds, Bills and Inflation: 1989 Yearbook* (Chicago, 1989).
- Kotlikoff, Laurence J., "Testing the Theory of Social Security and Life Cycle Accumulation," *American Economic Review* 69 (June 1979), 396-410.
- Kotlikoff, Laurence J., John B. Shoven, and Avia Spivak, "Annuity Markets, Savings, and the Capital Stock," in Zvi Bodie, John B. Shoven, and David A. Wise (eds.), *Issues in Pension Economics* (Chicago: The University of Chicago Press, 1987).
- Kurz, Mordecai, "The Life-cycle Hypothesis as a Tool of Theory and Policy," in George R. Feiwel (ed.), *Arrow and the Foundations of the Theory of Economic Policy* (New York: New York University Press, 1987).
- Leimer, Dean R. and Selig D. Lesnoy, "Social Security and Private Saving: New Time-Series Evidence," *Journal of Political Economy* 90 (June 1982), 606-629.
- Leimer, Dean R. and Peter A. Petri, "Cohort-Specific Effects of Social Security Policy," *National Tax Journal* 34 (March 1981), 9-28.
- Lesnoy, Selig D. and Dean R. Leimer, "Social Security and Private Saving: New Time Series Evidence with Alternative Specifications," Working Paper No. 22, (Social Security Administration, Office of Research and Statistics, Division of Economic Research, September 1981).
- Light, Paul, *Artful Work: The Politics of Social Security Reform* (New York: Random House, 1985).
- Mariger, Randall P., *Consumption Behavior and the Effects of Government Fiscal Policies* (Cambridge: Harvard University Press, 1986).
- Marquardt, Donald W., "An Algorithm for Least-Squares Estimation of Nonlinear Parameters," *Journal for the Society of Industrial and Applied Mathematics* 11 (June 1963), 431-441.

- Merton, Robert C., "On the Role of Social Security as a Means for Efficient Risk Sharing in an Economy Where Human Capital is Not Tradable," in Zvi Bodie and John B. Shoven (eds.), *Financial Aspects of the United States Pension System* (Chicago: The University of Chicago Press, 1983).
- Munnell, Alicia H., *The Effect of Social Security on Personal Saving* (Cambridge: Ballinger, 1974).
- Sibley, David, "Permanent and Transitory Income Effects in a Model of Optimal Consumption and Wage Income Uncertainty," *Journal of Economic Theory* 11 (August 1975), 68-82.
- Yaari, Menahem E., "A Law of Large Numbers in the Theory of Consumer's Choice Under Uncertainty," *Journal of Economic Theory* 12 (April 1976), 202-217.
- Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations From Certainty Equivalence," *Quarterly Journal of Economics* 104 (May 1989), 275-298.