

Gary Burtless, Editor

**Work, Health, and Income
among the Elderly**

THE BROOKINGS INSTITUTION
Washington, D.C.

Robert A. Moffitt

Life-Cycle Labor Supply and Social Security: A Time-Series Analysis

Between the end of World War II and the present, there have been marked reductions in the employment rates, hours of work, and other measures of labor supply for men in the United States. These reductions have occurred at all ages but have been particularly strong at older ages. The reductions for older men are widely believed to be a result of the increasing generosity of the social security system, the benefits and coverage of which have grown over this period. The system was created in 1935, and benefits began to be paid in 1940, but benefit levels and reciprocity rates were low for the first decade or so of the system. In the 1950s, however, coverage rates, reciprocity rates, and benefit levels began to grow, and in the late 1960s and early 1970s large increases in real benefits were legislated by Congress. The labor supply of older men dropped sharply in the 1970s as well, lending further specific support to the hypothesis that the social security system has caused the decline in labor supply.

The resolution of the empirical question of the magnitude of the contribution of social security to these declines in labor supply is important for social security policy. A specific example of its importance lies in the cost implications of advancing or lowering the retirement ages of the elderly. If the growth rate of real social security benefits declines in the future, as it must, it is possible that retirement ages of the aged will eventually creep upward. This advancement of the retirement age in turn will have a significant effect on the cost consequences of planned increases in "normal" retirement ages under social security in the next century. The magnitude of

I thank Joseph Quinn, the participants of the Brookings Conference on Retirement and Aging, and the participants in a workshop at McMaster University for comments. The able assistance of Ralph Buckley, Suk-Jae Noh, and Anuradha Rangarajan is also gratefully acknowledged.

such consequences will clearly depend critically on the responsiveness of retirement decisions to the level of social security benefits.

Nevertheless, the hypothesis that the social security system has induced the time-series decline in the labor supply of the aged is based primarily on studies of individuals' behavior in the late 1960s and 1970s, both cross-sectional studies of individuals and panel-data studies following individuals over time. Evidence from the 1950s and early 1960s is sparse, other than aggregate time-series evidence, and the time-series data have rarely been used in the economic literature. A good example of the type of study that has been done is the early study of Boskin, who estimated what was essentially a cross-sectional model with data on individuals from 1968-72 and then used the results to "backcast" time-series trends in labor supply.¹ Boskin's results indicated that social security was the most important explanatory factor for the reductions in time-series labor supply. Other studies have used panel data sets on individuals from the Retirement History Survey or the National Longitudinal Survey, but these again cover only the late 1960s or the 1970s.² Such data are appropriately used for a direct study of the effects of the benefit increases in the late 1960s and early 1970s, but they are not sufficient for an examination of the effect of social security in the earlier years of the postwar period.

This study provides a direct examination of the postwar time-series aggregate data. Annual data on labor supply of men in the United States are used in conjunction with measures of the social security system to obtain estimates of the effect of the system on work effort. The study can thus provide, at minimum, an estimate of the extent to which the individual-data studies just mentioned provide accurate indicators of the entire postwar experience, and of the extent to which time-series trends of the 1950s and early 1960s are consistent with those studies.

1. Michael J. Boskin, "Social Security and Retirement Decisions," *Economic Inquiry*, vol. 15 (January 1977), pp. 1-25.

2. A few examples are two studies by Gary Burtless and Robert A. Moffitt, "The Effect of Social Security Benefits on the Labor Supply of the Aged," in Henry J. Aaron and Gary Burtless, eds., *Retirement and Economic Behavior* (Brookings, 1984), pp. 135-74, and "The Joint Choice of Retirement Age and Postretirement Hours of Work," *Journal of Labor Economics*, vol. 3 (April 1985), pp. 209-36; Peter A. Diamond and Jerry A. Hausman, "The Retirement and Unemployment Behavior of Older Men," in Aaron and Burtless, *Retirement and Economic Behavior*, pp. 97-134; and Michael D. Hurd and Michael J. Boskin, "The Effect of Social Security on Retirement in the Early 1970s," *Quarterly Journal of Economics*, vol. 99 (November 1984), pp. 767-90. A literature review is provided in Henry J. Aaron, *Economic Effects of Social Security* (Brookings, 1982).

Because the primary motivation for the paper rests in the argument that the time-series data are worthy of examination—if they may not be preferable in some ways to cross-sectional data and short panels—a brief review of the relative advantages and disadvantages of cross-sectional and panel data is necessary.

Cross-sectional individual data currently tend to be favored over time-series, aggregate data because of the large number of degrees of freedom in cross-sectional data. But use of cross-sectional data requires making statistical inferences solely from comparisons of individuals, a difficulty particularly important for the analysis of social security. For social security, the law is the same for all people at any given time; consequently, all cross-sectional variation in social security benefits or any other measure of the system must arise from cross-sectional variation in earnings received over the lifetime, in family size and the number of dependents, in marital status, and in other such variables. That is, there is no variation in the law itself. The potential difficulty of course is that the variables for which variation is available may have independent effects on labor supply; hence there is a fundamental identification problem in cross-sectional data, a problem that can only be overcome by making restrictions in functional form of one kind or another. This identification problem is only partially overcome by existing micropanel, for most of these are too short to capture many of the important changes in the law.³

Time-series data do not have this problem but do have problems associated with low degrees of freedom. Because of the small number of observations and because many variables trend together and hence are highly correlated, inferences made from time-series data are inherently weaker than those made from cross-sectional data. As a consequence, only simple tests have much statistical power; it is virtually impossible to discriminate

3. The microdata panels can capture the major changes in the law in 1972 and the late 1960s. Indeed, at least two studies (Hurd and Boskin, "Effect of Social Security on Retirement," and Gary Burtless, "Social Security, Unanticipated Benefit Increases, and the Timing of Retirement," *Review of Economic Studies*, forthcoming) have treated these law changes as unexpected shocks, just as in this study, and have directly examined the labor supply response to the shock. Relative to those studies, the present study considers a longer time span and a greater number of law changes. There is some question, however, whether the 1972 law did constitute a "shock" (see below). The time-series evidence has been used previously to examine the effects of the social security system on saving. See Martin Feldstein, "Social Security, Induced Retirement, and Aggregate Capital Accumulation," *Journal of Political Economy*, vol. 82 (September–October 1974), pp. 905–26; and Dean R. Leimer and Selig D. Lesnoy, "Social Security and Private Saving: New Time-Series Evidence," *Journal of Political Economy*, vol. 90 (June 1982), pp. 606–29.

between competing hypotheses of great sophistication and subtlety. Therefore, only a simple model of the effect of social security on life-cycle labor supply will be examined here—social security will be assumed simply to set off wealth effects that may decrease labor supply over the lifetime. Even in this simple model, however, the high correlation between trends in real income and in social security wealth makes it quite difficult, if not impossible, to estimate the separate effects of these two variables on labor supply. Indeed, an alternative hypothesis for the observed reductions in time-series measures of labor supply is that they have been a result of rising private income levels, not social security.

This identification problem in the time-series data provides a point of focus for the paper. Solution of the problem is found here by estimating the wealth elasticity of labor supply only from variations in *unexpected changes* in net social security wealth (that is, net of taxes) over the life cycle. Neither the value of private wealth nor the value of net social security wealth at the beginning of a cohort's life cycle is used to identify the elasticity; these variables are not even included in the estimating equations.

In the next section a brief overview of the time-series trends in labor supply and social security is provided. That first provides the background for the rest of the paper and documents the trend discussed here in the introduction. The second section contains a discussion of the life-cycle model of labor supply, on which the estimating equations in the study are based, as well as further discussion of the identification problem just described. The third section presents the existing time-series data and shows how limiting those data are. That section provides an illustration of the general point made about the inability of time-series data to discriminate between subtle hypotheses. Although the data are sufficiently weak to sustain legitimate doubt that an analysis can be undertaken, it is argued that the estimating equations can be reformulated in such a way as to provide a sound and potentially fruitful test of the effect of social security on labor supply. The equations represent little more than simple correlational analysis but capture correctly the intuition of how social security affects labor supply. The fourth section discusses how unexpected changes, or "shocks," in social security wealth are calculated and presents several alternative methods of calculation. There are major surprises in this section, for the results do not correspond to expectations. The fifth section, which presents the estimating results, confirms these unexpected findings. The final section offers a summary and conclusions.

Trends in Labor Supply and Social Security since World War II

Figures 1 and 2 show the trends in the labor-force participation and in annual hours of work of men over the postwar period separately by age group. Both participation rates and hours of work show declines after the late 1960s. For younger men there were increases in the measures in the 1950s, but for older men there is a continuous decline over the entire period. Note that the trends for the two older male categories are not the same. For the oldest age category, males 65 years old and older, participation rates seem to have declined the most in the 1950s and 1960s, less so in the 1970s. The hours worked by these men show a fairly steady decline, although somewhat larger in the early 1950s and 1970s than in the 1960s. For men aged 45-64, however, there is an accelerating rate of reduction in participation rates starting in the early 1960s.

An unpublished appendix to this paper, available from the author on request, shows trends for employment rates (that is, employment-to-population ratios) and for weeks worked over the year, as well as the trends in all four measures for women. The trends in employment rates and weeks worked for men follow those in figures 1 and 2, as should be expected. For women, increases in labor supply have been the rule over the postwar period; of interest, however, is that the rates of increase have declined for women over time. The labor supply levels of older women have actually declined, in marked contrast to those of other women.

There is a *prima facie* case that these declines in labor supply are a result of the social security system. Real benefit levels have increased, and coverage of the system has been broadened over time. Regressions of participation rates and hours of work for older age groups, performed on the real benefit, yield coefficients that are strongly negative and significant, reflecting their obvious negative correlation over time. At the same time, real wage rates have risen as the result of normal economic growth. Regressions of the same labor supply measures on real wage rates again yield significantly negative coefficients, a result in general interpreted as evidence of backward-bending labor supply. Separating the effects of benefits and wages is difficult in the presence of the high correlation between them.


Fortunately for the purposes of analysis, the rising benefit levels have not occurred steadily. Although coverage rates and reciprocity rates have trended upward gradually, and although social security tax rates were steady until the 1970s, benefit levels underwent dramatic changes. After

Figure 1. Labor-Force Participation Rate of Men, by Age Group, 1947-82

Participation rate
(percent)



Source: Data from the Bureau of Labor Statistics (BLS).

Figure 2. Annual Hours of Work for Men, by Age Group, 1950-81^a

Source: BLS data.

a. Annual hours are defined as workers' hours per week times weeks worked per year.

being initially set in 1937, benefits were held constant in nominal terms until 1949, with the result that real benefits were reduced over that period. Then in 1950 benefits were increased by up to 60 percent in a single year. Subsequently, benefits rose steadily but without dramatic change until the late 1960s and early 1970s, when benefits were increased frequently. That series of increases culminated in 1972 legislation that not only increased benefits by about 20 percent in a single year but also indexed benefits for the first time. Since then benefits have again risen only steadily.

The trends in figures 1 and 2 correspond partially, but not fully, to the expectations generated by the timing of these shocks. Although as a general rule the figures show that labor supply declined more in the late 1960s and 1970s than it did previously, this is not the case for every group (for example, not for males 65 years and older). Furthermore, although the 1950 benefit increase was equal to or larger in percentage terms than the 1972 increase, there is no indication of any significant post-1950 decline in participation rates (except possibly for the group over age 65) or in any of the participation rates or hours worked for groups of younger men.

Clearly there may be other time-varying explanatory factors for the differences in labor supply trends over the period, thus leading to multivariate analysis. But the statistical leverage provided by the sudden changes in the structure of the social security program will still be used as the primary identifying causal influence.

Life-Cycle Labor Supply, Social Security, and the Identification Problem

In this section I shall discuss a simple life-cycle model theory of labor supply and shall present some basic results showing how unexpected changes in the value of lifetime wealth affect the life-cycle profile of labor supply. The model will then be made concrete by specifying a particular form (Stone-Geary) for the utility function and by introducing social security as a simple wealth transfer that undergoes unexpected alterations. Equations are derived for labor supply at each age that are a function of, among other things, the entire past sequence of wealth shocks. The problem of identification arises because the level of private assets is not observed and because the trend in this level is correlated with that of social security wealth. The identification problem can be solved by simple first-differencing of the equation, by the introduction of cohort dummies, or by both. In all cases, the coefficient measuring the wealth elasticity of labor supply—which also measures the effect of social security—is identified in these equations solely by the values of the social security shocks.

The Theory of Wealth Shocks

Here it will be shown that unexpected changes in wealth have a greater negative effect on labor supply than do expected changes in wealth, and that such shocks have larger effects as the age at which the shock occurs increases. The reader who wishes to do so may skip this section because the subsequent subsection, on a specific model, does not depend for understanding on the material presented here.

In the life-cycle theory of labor supply, a person chooses labor supply in each period of life in response to values of wages and prices over a lifetime and to the value of exogenous wealth. As a first approximation in this framework, the social security system can be modeled as providing an increment to the value of exogenous wealth. This treatment is only an

approximation because the system also alters the relative price of leisure over the lifetime, but it can be safely assumed that the wealth effects of social security dominate its substitution effects in time series.⁴ Increases in the value of exogenous wealth decrease labor supply in all periods of the lifetime, if it is assumed that leisure is a normal good in all.

Unanticipated wealth increments have the same direction of effect but differ in their magnitude and pattern over the life cycle. Such increments in wealth decrease labor supply more than do expected increments because past labor supply was chosen on the basis of an expectation of lower lifetime wealth. When a wealth shock occurs, a person must replan his labor supply profile and make up for lost time.

In formal terms, suppose that the lifetime utility function is of the general nonseparable form:

$$(1) \quad U(H_1, H_2, \dots, H_T, C_1, C_2, \dots, C_T),$$

where H_a is a labor supply at age a , C_a is consumption at age a , and T is the length of adult life. For convenience, hours of work are used in the utility function instead of hours of leisure. The lifetime budget constraint at $a = 1$ is

$$(2) \quad A + \sum_{a=1}^T r_a W_a H_a = \sum_{a=1}^T r_a C_a,$$

where A is exogenous wealth at the start of the lifetime; r_a is the market discount rate, defined as one, divided by one plus the interest rate raised to

4. Although the social security system sets off both substitution and income effects (see Gary Burtless and Robert A. Moffitt, "Social Security, Earnings Tests, and Age at Retirement," *Public Finance Quarterly*, vol. 14 [January 1986], pp. 3-27, for a simple graphical exposition), econometric estimates of the elasticities of substitution are not large (Burtless and Moffitt, "Effect of Social Security Benefits" and "Choice of Retirement Age"). Moreover, the elasticities of substitution would have to be extraordinarily high to outweigh the large wealth effects set off by social security transfers, especially in time series, since the shocks in the 1950s, late 1960s, and early 1970s were predominantly due to changes in benefit levels rather than in tax rates. The one time-series study of relative price effects is that of Richard V. Burkhauser and John A. Turner, "A Time-Series Analysis on Social Security and Its Effect on the Market Work of Men at Younger Ages," *Journal of Political Economy*, vol. 86 (August 1978), pp. 701-15, which offered the interesting hypothesis that the retirement test induces younger men to work longer hours than older men. Burkhauser and Turner's empirical evidence, however, was based on trends in hours of work for all men, whereas the evidence given above indicates that men aged 45-64 (most of whom do not receive benefits) reduced their labor supply.

the a^{th} power; and W_a is the hourly wage rate at age a . The price of consumption goods is normalized to unity. The utility-maximizing labor supply functions can then be written in general form as⁵

$$(3) \quad H_a = f_a(W_1, W_2, \dots, W_T, A), \quad a = 1, 2, \dots, T.$$

It is expected that an increase in A will decrease all H_a .

Now suppose that there is an unanticipated increment to wealth in the amount ΔA at age q . If $q = 1$, and thus if the individual is aware of the increment from the start, then labor supply in all future periods is

$$(4) \quad H_a^* = f_a(W_1, W_2, \dots, W_T, A + \Delta A), \quad a = 1, 2, \dots, T.$$

But if $q > 1$, then the individual must remaximize the utility function:

$$(5) \quad U(\bar{H}_1, \dots, \bar{H}_{q-1}, H_q, \dots, H_T, \bar{C}_1, \dots, \bar{C}_{q-1}, C_q, \dots, C_T),$$

where an overbar denotes a quantity that is given. The values of H_a and C_a for all $q \geq a$ must be chosen subject to the constraint

$$(6) \quad \left[(A + \Delta A) + \sum_{a=1}^{q-1} r_a (W_a \bar{H}_a - \bar{C}_a) \right] + \sum_{a=q}^T r_a W_a H_a = \sum_{a=q}^T r_a C_a.$$

Here the term in brackets is the value of "exogenous" wealth (discounted back to $a = 1$) at period q , given that $\bar{H}_a$ and $\bar{C}_a$ have been chosen up to q . The solution to the maximization will not give functions of the form f_a used above, so it is a bit awkward to compare labor supply before and after the shock. But one may consult the literature on rationing to use the function f_a after all, by defining the "virtual wages," $W'_1, W'_2, \dots, W'_{q-1}$, as the wage rates that would have generated $\bar{H}_1, \bar{H}_2, \dots, \bar{H}_{q-1}$ as utility-maximizing choices had wealth been equal to $A + \Delta A$ at the start of the life cycle.⁶ These wages are obtainable by solving the $q - 1$ equations for

5. The discount rates r_a are suppressed in all subsequent functions in the interest of notational simplicity.

6. See Angus Deaton and John Muellbauer, *Economics and Consumer Behavior* (New York: Cambridge University Press, 1980), for a textbook treatment of the theory of rationing.

$\bar{H}_a$ as a function of the wages and of $A + \Delta A$. In this formulation one may write labor supply after $a = q$ with the functions f_a :

$$(7) \quad H_a^{**} = f_a(W'_1, W'_2, \dots, W'_{q-1}, W_q, \dots, W_T, A + \Delta A),$$

$$a = q, \dots, T.$$

Thus expressions have been formulated for labor supply in the absence of the wealth shock— H_a in equation 3; labor supply in the presence of the wealth shock if it had been known at the beginning of the lifetime— H_a^* in equation 4; and actual labor supply given that the shock occurred in year q — H_a^{**} in equation 7. The following propositions can then be demonstrated (proofs have been omitted), with ΔA assumed to be greater than zero:

—Proposition 1:

$$(8) \quad H_a^* \leq H_a;$$

—Proposition 2:

$$(9) \quad H_a^{**} \leq H_a^*;$$

—Proposition 3:

$$(10) \quad \partial D / \partial q < 0, \quad \text{where } D = H_a^{**} - H_a^*, \quad a \geq q.$$

Proposition 1 simply states that wealth effects on labor supply are negative. Proposition 2 states that labor supply will be reduced more by an unanticipated wealth shock than by one that is anticipated at the beginning of the life cycle. Because it was not known at the beginning that the wealth increment would later occur, labor supply and consumption prior to q were "misplanned." The reductions in labor supply generated by the wealth increment would have been spread over the entire lifetime if all had been known at the beginning. Proposition 3 is suggested by proposition 2, for proposition 3 states that this "excess reduction" caused by the early misplanning is greater, the later the shock occurs in the lifetime. All these propositions, including the third, are in principle testable. If a social security shock at a given time increases the wealth of all cohorts by the same increment, and if the cohorts are alike in all other respects, then the reduction in labor supply induced by the shock will be greater for those who are older. Note that this effect is implied even if a cohort has not yet received any benefits at the time of the shock.

A Specific Model

The specific functional form used will be the Stone-Geary utility function, which generates the linear expenditure system.⁷ The utility function is assumed to be

$$(11) \quad U = \sum_{a=1}^T p_a [\beta \ln(\alpha - H_a) + \gamma \ln(C_a - \delta)],$$

where $p_a = s_a/(1 + \rho)^a$, the probability of survival to age a (s_a) divided by one plus the subjective rate of discount (ρ) raised to the a th power. The lifetime budget constraint at $a = 1$ is

$$(12) \quad A = \sum_{a=1}^T r_a (C_a - W_a H_a),$$

where $r_a = s_a/(1 + i)^a$, the survival rate divided by one plus the interest rate raised to the a th power. The labor supply functions expressed in the form of earnings are⁸

$$(13) \quad W_a H_a = \alpha W_a - \beta \eta^a Y,$$

where

$$(14) \quad \eta = [(1 + i)/(1 + \rho)]$$

and

$$(15) \quad Y = A + \sum_{a=1}^T r_a (W_a \alpha - \delta).$$

Substituting equation 15 into equation 13 gives

$$(16) \quad W_a H_a = \alpha W_a - \beta \eta^a A - \beta \eta^a \alpha \sum_{\tau=1}^T r_{\tau} W_{\tau} + \beta \eta^a \delta \sum_{\tau=1}^T r_{\tau}.$$

7. See Orley Ashenfelter, "Macroeconomic and Microeconomic Analyses of Labor Supply" (Princeton University, 1984), for a discussion of the Stone-Geary model in a life-cycle context.

8. The normalization imposed is $\sum_{a=1}^T p_a (\beta + \gamma) = 1$.

Thus earnings at age a are a linear function of the wage at age a and of the variable Y —or, in equation 16, a linear function of the current wage, the initial level of assets, and the (discounted) sum of wage rates over the lifetime. The intertemporal wage elasticity for this function is $(\alpha/H_a) - 1$. This function has never been estimated in this form because the initial level of assets A (say, at age 20 or 25) is never available in time-series or individual data sets. But if A is left out of the equation, the wealth effect β cannot be identified.

This difficulty leads to the identification problem if social security is added to the model. If someone has a fixed stock of social security wealth known at the beginning of the life cycle, equal to the discounted sum of the difference between benefits and taxes and denoted S , then the variable $A + S$ would appear in place of A in the equation. If S were unobserved, one would again be unable to identify β . If a value of S were calculable (as it will be below), then one would be able to identify β . But if A is unobserved, then the effects of A and S cannot be disentangled; omission of A from the equation would bias the estimate of β . Note that it is the unobservability of assets, not of wages, that creates these difficulties.

To identify β instead from a social security shock, consider an unexpected net wealth increment from the social security system in the amount Q at age q . Then the solution to the replanning problem yields the earnings function:

$$(17) \quad W_a H_a = \alpha W_a - \beta \eta^a (Y + k_q Q),$$

where

$$(18) \quad k_q = \sum_{\tau=1}^T p_\tau / \sum_{\tau=q}^T p_\tau,$$

and where Y now includes $A + S$. The wealth increment Q appears with a coefficient that is greater than unity and that increases with q , as discussed earlier. If there are continual unexpected net wealth increments of Q_q at each q , the earnings function becomes

$$(19) \quad W_a H_a = \alpha W_a - \beta \eta^a \left(Y + \sum_{\tau=2}^a k_\tau Q_\tau \right).$$

Here there is a "weighted" supernumerary income, in parentheses, that is equal not to the straight sum of wealth over the lifetime but rather to the sum of such increments weighted by the age at which they occur.⁹ For empirical purposes, the two parameters η^a and k_q will be approximated by their first-order Taylor-series expansions:

$$(20) \quad \eta^a \approx 1 + a \ln \eta = 1 + a\eta'$$

$$(21) \quad k_q \approx T/(T - q + 1) = k'_q,$$

thus giving the earnings function:

$$(22) \quad W_a H_a = \alpha W_a - \beta \left(Y + \sum_{\tau=2}^a k'_\tau Q_\tau \right) - \beta \eta' a \left(Y + \sum_{\tau=2}^a k'_\tau Q_\tau \right).$$

If Y is unobserved, then equation 22 cannot be estimated in this form. But if $\eta' = 0$ (that is, $i = \rho$), then a first-differenced equation would give

$$(23) \quad W_{a+1} H_{a+1} - W_a H_a = \alpha (W_{a+1} - W_a) - \beta (k'_{a+1} Q_{a+1}).$$

Regressing the change in earnings on the change in the wage rate and the $(a + 1)^{\text{st}}$ weighted net wealth increment thus allows both α and β to be identified.¹⁰ In particular, the wealth effect can be identified solely from the response to social security wealth shocks, not from the values of either

9. Certainty equivalence is assumed here; having to specify the effect of the variance of uncertain social security wealth is thus avoided. Note also that uncertainty in variables other than social security wealth is ignored. The first-differencing techniques used below are, in some circumstances, approximately equivalent to those that would be used in the case of uncertainty.

10. Ashenfelter and Ham also treated supernumerary income in a Stone-Geary model as a fixed effect. Orley Ashenfelter and John Ham, "Education, Unemployment, and Earnings," *Journal of Political Economy*, vol. 87 (October 1979), pp. S99-S116. Note that this treatment is quite different from first-differencing the Frisch demands. See Thomas E. MaCurdy, "An Empirical Model of Labor Supply in a Life-Cycle Setting," *Journal of Political Economy*, vol. 89 (December 1981), pp. 1059-85; and Martin Browning, Angus Deaton, and Margaret Irish, "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle," *Econometrica*, vol. 53 (May 1985), pp. 503-43. In that case, no wealth measures appear in the first-differenced equation, since the marginal utility of wealth is held constant. That formulation is not convenient for the application here because it is the uncertainty in wealth that is of direct interest; in Frisch models the uncertainty in wealth is instead subsumed in the (time-varying) function for the marginal utility of wealth.

private wealth or the level of social security wealth at any previous period. The problem of collinearity between unobserved private wealth and observed social security wealth is therefore eliminated.

If $i \neq \rho$, the first-differenced equation is

$$(24) \quad W_{a+1}H_{a+1} - W_aH_a = \alpha(W_{a+1} - W_a) - \beta(k'_{a+1}Q_{a+1}) \\ - \beta\eta'(ak'_{a+1}Q_{a+1}) \\ - \beta\eta'Y - \beta\eta'\left(\sum_{\tau=2}^{a+1} k'_\tau Q_\tau\right).$$

Here the parameter η' is identified from the coefficient on the interaction between age and the weighted wealth shock, but the unobservable Y also appears. To estimate Y , one can assume formally that Y is a fixed effect—that is, a fixed value for each cohort—and replace Y by a set of dummy variables for the cohorts in the sample.¹¹ This procedure allows the consistency of the other coefficients to be retained but does not alter the sources of parameter identification. But if Y is replaced by dummy variables here, it may also be replaced in the level equation 22, regardless of the value of η' , with the same consequences. Hence equations 22–24 can all be estimated, and in each the wealth effect is identified solely by the presence of unexpected increments in social security net wealth.

Data and Empirical Implementation

The estimating equations for the paper are equations 22–24 above, which contain a dependent variable for earnings (hours of work times the wage rate) and independent variables for the hourly wage rate, age, and the weighted social security wealth (the weight is merely a function of age; see equation 21). The variable Y is to be proxied by dummy variables for cohort (that is, birth year). By the nature of the model, the variables for hours of work, wage, and social security must be broken down by age.¹²

11. Browning and others, "A Profitable Approach," followed this same strategy.

12. Note here that the dependent variable in the model is hours of work, not the retirement date. No data on retirement dates are available in time series; in any case, the retirement date is not unambiguously defined even in microdata. If individuals work a fixed number of hours both before and after retirement, and if the only life-cycle decision is when to retire and move from one level of work hours to another, the model used here would yield lower elasticities than it would if hours were flexible. Because the hours variable used is a

Labor Supply and Wage Data

Information on hours of work by age is available for the postwar period only from the *Current Population Survey* (CPS). The March survey has always collected data not only on hours of work but also on weeks worked in the preceding year. To construct the variable for annual hours needed here, the average number of hours worked over each calendar year was multiplied by the number of average weeks worked in the same year.¹³

Calculating hourly wage rates in the time series is more difficult. In what must rank among the most frustrating discoveries this researcher has experienced, I found that the published and unpublished CPS data do not include any summary information on earnings by age. Information on individual income is widely available, but the difference between income and earnings is quite significant for the aged—and probably for younger segments of the population as well. As a substitute, earnings data by age were taken from the published reports of the Social Security Administration. The disadvantages of such data are that they are based only on workers covered and that the published reports were forced to rely on an imputation procedure to obtain earnings above the taxable maximum.¹⁴ In any case, these are the only data with which to work. Divided by annual hours, the requisite hourly wage is obtained.

The data impose another important limitation on implementation of the model because they are not available by single years of age, but rather by age categories. Unfortunately, hours of work are not available by age at all for years before 1955 (except from the decennial censuses) and exist only in four broad age categories for years after 1955: 25–34, 35–44, 45–64, and 65+. Thus one can study only the young, the youthful (for lack of a better term), the middle-aged, and the elderly. The more serious implica-

mean and would include both retirees and nonretirees, the trend in the dependent variable in the equations would reflect only the changing percentage of retirees in each age group. There is nothing particularly disadvantageous, however, in the use of time-series data in this case; microdata equations of the same type could be estimated and would also capture trends in retirement, even though with microdata one is more likely to model directly the probability of retirement for each individual.

13. Although some data on hours worked and weeks worked are available from published *Current Population Survey* (CPS) reports, better age detail is available from unpublished data from the Bureau of Labor Statistics. The data used here were obtained from these unpublished files.

14. The social security earnings data were obtained from various issues of the *Social Security Bulletin, Annual Statistical Supplement*, which contains data on earnings by age and by sex for all years from 1937 to the present.

tion of this age grouping is that exact cohorts cannot be identified and, hence, the first-differencing procedure described in the preceding section cannot be implemented. Instead, the best that can be done is take first differences of a particular age grouping for adjacent years, a procedure that will leave some cohort effects in the equation.¹⁵ Unfortunately, an additional problem created by the presence of cohort effects in all the equations is that the approach of using cohort dummy variables becomes impractical. For the ages 25–84 (using 84 as the upper age limit) over the years in the sample (1955–81), there are 91 birth cohorts. But there are only 108 observations in the entire sample because of the age grouping, not enough to estimate these cohort dummy variables.

To address these problems, the estimating equations 22–24 must be modified in two ways. First, those equations need to be modified to reflect the grouping over different ages, with the effect that average values of the variables must be used. Second, the variable Y must be proxied not by a set of cohort dummy variables but by some simple form of trend in Y that permits estimation. The reader interested only in the results of these modifications can directly inspect equations 28 and 29 below.

To solve the second problem, the unobservable cohort effect Y will be assumed to grow at a constant proportional rate over time. With this assumption, a variable denoting the year of the cohort will appear in the equations instead of the cohort dummy variables. Let Y_c be the value of Y for cohort c . Let c denote the year at which the cohort becomes 25, the earliest age in the data, and let $c = 0$ correspond to 1935, the year in which social security was put into effect. Then, the assumption is that Y_c follows the path

$$(25) \quad Y_c = (1 + g_1 + g_2 D)^c Y_0,$$

where D is a dummy variable equal to unity if c is greater than or equal to zero. Thus Y_c is assumed to grow at different rates depending on whether the cohort began its adult life before or after the introduction of social security. A different rate must be assumed because Y_c includes initial social security wealth as well as initial private assets.

Grouping equation 22 (the basic earnings function) over the ages in each age interval gives the approximate result

15. An alternative is to match up each ten-year age category across ten-year intervals, and the twenty-year age category across twenty-year intervals. Unfortunately, this would reduce the sample to about fifteen observations.

$$(26) \quad W_{\bar{a}}H_{\bar{a}} = \alpha W_{\bar{a}} - \beta Y_{\bar{c}} - \beta \sum_{\tau=2}^{\bar{a}} k'_{\tau} Q_{\tau\bar{c}} - \beta \eta' \bar{a} \sum_{\tau=2}^{\bar{a}} k'_{\tau} Q_{\tau\bar{c}},$$

where an overbar denotes the average value in the age interval. Note that a cohort subscript is now added to the social security wealth variable to represent explicitly differences in the wealth variable across cohorts. If the first-order Taylor-series approximation,

$$(27) \quad Y_{\bar{c}} = (1 + g_1 \bar{c} + g_2 \bar{c}\bar{D})Y_0,$$

is used, equation 26 becomes

$$(28) \quad \begin{aligned} W_{\bar{a}}H_{\bar{a}} = & \psi_0 + \psi_1 W_{\bar{a}} + \psi_2 \sum_{\tau=2}^{\bar{a}} k'_{\tau} Q_{\tau\bar{c}} + \psi_3 \bar{c} \\ & + \psi_4 \bar{c}\bar{D} + \psi_5 \bar{a} \sum_{\tau=2}^{\bar{a}} k'_{\tau} Q_{\tau\bar{c}} + \psi_6 \bar{a} \\ & + \psi_7 \bar{c}\bar{a} + \psi_8 \bar{c}\bar{D}\bar{a}, \end{aligned}$$

where

$$\begin{aligned} \psi_0 &= -\beta Y_0 & \psi_3 &= -\beta Y_0 g_1 & \psi_6 &= -\beta Y_0 \eta' \\ \psi_1 &= \alpha & \psi_4 &= -\beta Y_0 g_2 & \psi_7 &= -\beta Y_0 \eta' g_1 \\ \psi_2 &= -\beta & \psi_5 &= -\beta \eta' & \psi_8 &= -\beta Y_0 \eta' g_2. \end{aligned}$$

Thus the function contains the usual variables for wage rate and wealth shocks as well as the mean cohort value (essentially a time trend) and the fraction of those in the age grouping who reached age 25 after 1937. All variables also interact with age.

The first-difference version becomes

$$(29) \quad \begin{aligned} \Delta W_{\bar{a}}H_{\bar{a}} = & \phi_1 + \phi_2 \Delta W_{\bar{a}} + \phi_2 \Delta \left(\sum_{\tau=2}^{\bar{a}} k'_{\tau} Q_{\tau\bar{c}} \right) \\ & + \phi_3 \Delta(\bar{c}\bar{D}) + \phi_4 \Delta \left(\bar{a} \sum_{\tau=2}^{\bar{a}} k'_{\tau} Q_{\tau\bar{c}} \right) \\ & + \phi_5 \bar{a} + \phi_6 \bar{a} \Delta(\bar{c}\bar{D}), \end{aligned}$$

where

$$\begin{aligned}\phi_0 &= -\beta Y_0 g_1 & \phi_3 &= -\beta Y_0 g_2 & \phi_6 &= -\beta Y_0 \eta' g_2, \\ \phi_1 &= \alpha & \phi_4 &= -\beta \eta' \\ \phi_2 &= -\beta & \phi_5 &= -\beta Y_0 \eta' g_1\end{aligned}$$

where it is understood that the differencing takes place across adjacent years for the same age grouping.

New equations 28 and 29 represent the model in the form in which it will be estimated and illustrate both the limitations and the value of the analysis. In the equations earnings of an age group are assumed to be a function of the wage rate, age, a cohort trend (essentially a time trend), a variable for the alteration in the time trend of earnings after 1937, and a social security shock variable. By the standards of microdata analyses, the specification is thus extremely parsimonious. Because only 108 observations are in the sample (twenty-seven years for four age groups), any more elaborate specification would not be warranted. The model is indeed so simple that the estimate of the effect of social security on labor supply is not much more than a trend-adjusted simple correlation coefficient between the age-specific labor supply and social security variables, with an extra control variable for the wage rate trend.¹⁶

Nevertheless, the conventional wisdom that reductions in labor supply in time-series data have been largely a result of social security is not based so much on the cross-sectional elasticities and their backcasts, all of which have problems of their own. Rather, this conventional view rests on the strong *prima facie* evidence provided by the graphical evidence discussed in the second section and on knowledge of the trend in real social security benefits—in particular, that the two are in general negatively correlated and that in the 1970s the growth rate of real social security benefits accelerated at the same time that labor supply declines did. Equations 28–29 capture that two-variable correlation plus a little more and, hence, are adequate tests of what has been a very simple but also very persuasive hypothesis. The interesting aspect of equations 28–29 is that their estimation does not simply confirm the graphical intuition. As the next section (“Results”) will indicate, equations 28–29 are not so parsimonious as to preclude unexpected findings.

16. Some regressions were run that included the annual age-specific unemployment rate. The coefficients were weak in significance and did not alter the values of the other coefficients.

Calculation of Social Security Wealth

To implement the models that have been discussed, it is of course necessary that some assumption about expectations of future social security wealth be made so that deviations from those expectations may be calculated. This problem has been discussed at times in the literature on social security, although not in as formal a fashion as here.¹⁷ To estimate expected social security wealth here, I shall simply specify a separate equation for the wealth-generating process and assume that people use that equation to forecast their own wealth values. Deviations in those forecasts from year to year will then constitute the values of the Q_t . An alternative procedure would be to use the information in the legislation concerning future tax rates, benefit schedules, and so on, and to forecast what individuals would actually receive when they retire under the law in effect in each year. In addition to being a very large and complex undertaking, such a procedure assumes that individuals expect the law to remain fixed. Such an assumption is implausible, for the social security law is periodically changed by Congress. This potential for change was particularly evident in the period since the 1960s, during which the law was changed every year or two on average.

It is well known that the assumptions of economists about the forecasting equations for individuals are often inadequate descriptors of real-world forecasts, and that quite restrictive and simple forms of forecasting equations may be in serious error. This problem is addressed here in part by the use of several different forecasting equations, each constituting a slight deviation from a general form. The extent to which the estimates of the model are sensitive to different but equally defensible forecasting assumptions will be an indication of the stability of the results.

To specify a forecasting equation, the values of net social security wealth for all cohorts retiring to date must be calculated. Net social security wealth at each age is equal to the difference between the present value of benefits received by that age and the present value of taxes paid. The method of calculation is a modification of that used earlier by Moffitt.¹⁸

17. See Burtless, "Social Security, Unanticipated Benefit Increases"; Hurd and Boskin, "Effect of Social Security on Retirement"; and Kathryn H. Anderson, Richard V. Burkhauser, and Joseph F. Quinn, "Do Retirement Dreams Come True? The Effect of Unanticipated Events on Retirement Plans," *Industrial and Labor Relations Review*, vol 39 (July 1986), pp. 518-26.

18. Robert Moffitt, "Trends in Social Security Wealth by Cohort," in Marilyn Moon, ed., *Economic Transfers in the United States* (University of Chicago Press, 1984), pp. 327-53.

Data on benefits, taxes paid, coverage, and reciprocity rates by age were obtained for each cohort alive since 1937 and are used to calculate SB_{ca} , the present value of benefits received by cohort c by age a ; ST_{ca} , the present value of taxes paid by cohort c at age a ; and $S_{ca} = SB_{ca} - ST_{ca}$, net social security wealth of cohort c by age a . The details of the computation are reported in the appendix.

An important characteristic of this concept of social security wealth is that such wealth will rise over time not only if real benefits rise, but also if coverage rates and reciprocity rates rise (or even if life expectancy rises—this possibility is included in the calculation). The calculation simply involves dividing total benefits paid to a cohort minus total taxes paid by that cohort by the number of individuals in the cohort at age 20. It thus gives the average increment of net wealth per cohort member. Because this wealth value will be the measure of the generosity of the system used here, an implicit assumption in the analysis is that the effects of changes in benefits, tax rates, coverage rates, and reciprocity rates have the same effects on labor supply if they also have the same effects on average social security wealth per cohort member.¹⁹

ACTUAL VALUES OF WEALTH RECEIVED TO DATE. Figure 3 shows the calculated values of S at various ages for selected cohorts. The early cohorts reached the end of their lifetimes having obtained little wealth (for example, the cohort that was 20 years old in 1897 and was already 60 years old by 1937)—in part because benefits were low in the beginning and in part because these cohorts had accumulated little coverage; hence their reciprocity rate was extremely low (less than 10 percent). More recent cohorts had higher coverage and reciprocity rates as well as higher benefits. The figures show, as also found by Moffitt, that the absolute size of the intergenerational transfer (although not the rate of return) by the end of life has grown monotonically for each successive retiring cohort since the system began.²⁰ The increased tax burden is also clearly shown in the figures for the more recent cohorts, particularly at young ages.


The more important figures are figures 4 and 5, which illustrate how the present values of benefits, taxes, and benefits minus taxes at the end of life

19. Although there certainly may be some cases in which this equivalence of effects will be incorrect, such instances are likely to be of second-order importance. Moreover, the assumption has been maintained just as strongly in most of the microdata studies, in which measures of social security benefits or wealth, for example, are set equal to zero for those not covered. The equations estimated here are, aside from differences in functional form, linear aggregates of those microdata equations; these aggregates yield social security values that are weighted values in the microdata equations for covered and uncovered workers.

20. Moffitt, "Trends in Social Security Wealth."

Figure 3. Net Social Security Wealth, by Age and Cohort^a

Net social security wealth
(1972 dollars)




Source: Discussed in Robert Moffitt, "Trends in Social Security Wealth by Cohort," in Marilyn Moon, ed., *Economic Transfers in the United States* (University of Chicago Press, 1984).

a. Year denotes point at which cohort reached age 20; age in parentheses is that of cohort at entry into the system.

(taken to be 80 years) have changed over time. Figure 4 shows that the present value of benefits for cohorts reaching age 80 in each successive year has grown monotonically over time. Despite the sizable increases in benefits in 1950 and in the late 1960s and the 1970s, however, there is no dramatic change in the present value of benefits at those times. There is a slight increase in the rate of growth of the present value of benefits in the

Figure 4. Present Value of Social Security Benefits, Taxes, and Differences between Them, 1937-81^a

Present value
(1972 dollars)



Source: See figure 3.

a. Values in each year are those for the cohort that had reached 80 years of age in that year.


late 1960s and early 1970s, but this alteration in the trend is only barely perceptible visually.

The strong growth in the 1950s and early 1960s is simply a result of the maturation of the social security system. Each successive cohort had spent more years in the system and, hence, had higher covered earnings; the total value of benefits received would therefore have risen even if the benefit schedule had not been altered. In addition, as coverage was legislatively expanded in the 1950s, reciprocity rates increased even more. There consequently was tremendous growth in the present value of benefits in the 1950s, at rates only barely below those in the late 1960s and early 1970s.

Tax payments were of course also increasing, as shown in figure 4; indeed, they have increased exponentially. The net value of social security wealth is also shown. The general shape of the net wealth curve is the

**Figure 5. Logarithms of Net Present Value of Social Security Benefits
Minus Taxes, 1941-81**
1972 dollars

Log of present value
(1972 dollars)



Source: Author's calculations based on data in figure 4.

same as that of benefits, despite the increase in taxes: steady growth throughout the 1950s and early 1960s, followed by a slight increase in the growth rate in the late 1960s and early 1970s.

The importance of these findings of steady growth over the entire period lies in the realization that such growth rates would appear, at first glance, to be rather easy to forecast. It appears as though a simple linear trend forecast would be fairly accurate, for example, and that there would be little or no social security shock in the late 1960s and early 1970s as a result.

Figure 5, which shows the profile of the logarithm of social security wealth, is even more striking. The growth rate of social security wealth has fallen monotonically over time, and visually it is almost impossible to discern any interruption in the late 1960s and early 1970s. As documented

elsewhere,²¹ the absolute value of the net wealth increment has grown monotonically for all successive retiring cohorts, but the rate of growth has monotonically fallen, and the internal rate of return has monotonically fallen as well. Again, these trends are natural results of a maturing pay-as-you-go system; as the system ages, tax rates must be increased, and individuals have spent more time in the system and therefore have paid more taxes. Cohorts retiring in the late 1960s and early 1970s were actually getting less out of the system in this sense than had earlier retirees. The benefit increases in the late 1960s and early 1970s merely kept the rate of return from falling faster.

FORECASTS AND CONSTRUCTION OF SHOCKS. To generate forecasts of expected wealth at the end of the lifetime, a separate regression is run at each year t using the data on social security wealth at age 80 in periods $\tau \leq t$. Simple moving-average representations of the process are assumed. At each t the resultant coefficients are used to forecast the expected end-of-lifetime wealth for each age group a at time t , by forecasting the equation $(80 - a)$ years into the future. Values of the social security shock are then calculated as the deviations of these forecasts for the same cohort in successive t 's.

Let S_t be the value of social security wealth for the cohort of age 80 in year t (shown in figure 4). Let l be the length of the moving-average process, and let $t_0 = t - l$ be the beginning year of the process. Then, for each t the following two different functional forms, a linear and a logarithmic equation, are estimated:

$$(30) \quad S_\tau = b_t + c_t \tau, \quad \tau = t_0, \dots, t$$

$$(31) \quad \log(S_\tau) = d_t + e_t \log(\tau - t_0 + 1), \quad \tau = t_0, \dots, t.$$


In the linear form, equation 30, wealth is fit to a simple linear trend. In the logarithmic form, equation 31, the logarithm of wealth is fit to a logarithmic trend starting at the beginning of the moving-average process. Most work was performed with equation 30, which was estimated using five-, ten-, and fifteen-year moving-average frames. The logarithmic equation was then estimated with a ten-year frame as well.²² At each t the coeffi-

21. Ibid.

22. Five-year and fifteen-year moving-average logarithmic equations were later estimated as well, with similar results. Tables of the results are available from the author on request.

Figure 6. Size of Forecast Shock in Net Social Security Wealth, by Age, Year, and Type of Forecast, 1950-81

Net social security wealth
(1972 dollars)



Source: Author's calculations.

cients (b_t , c_t) or (d_t , e_t) are used to forecast end-of-lifetime wealth for each age group a , equal to $\hat{S}_{t+80-a}$. The previous year's forecast for the same cohort, $\hat{S}_{(t-1)+80-(a-1)}$, are based on the coefficients (b_{t-1} , c_{t-1}) or (d_{t-1} , e_{t-1}). The value of the shock is then computed as the difference.²³

Figure 6 shows the estimated values of shock at each t for two of the age groups, a young group (30 years old) and an aged group (70 years old). The figure shows the estimated shocks from the linear five-year moving-average forecast. The curves indicate positive shocks throughout the 1950s and early 1960s, a second positive shock in 1972, but negative shocks throughout the rest of the 1970s. The negative shocks are a result of the rate of growth of wealth in the later 1970s, which did not keep up with the relatively rapid growth in the late 1960s and early 1970s. The shocks are smaller in absolute size for the aged group simply because its members are much closer to age 80; hence a change in the time profile of wealth has less aggregate effect.

The shock values for the ten-year moving averages show considerably larger values in the 1950s but a smaller shock in 1972. The ten-year moving averages clearly indicate falling shocks after about 1965. The fifteen-year moving average goes further in this direction: 1972 shows up only as a brief slowdown in the rate of decline of the shocks. As longer moving-averages are used, the 1972 blip has less of an effect on the forecasted wealth values. The logarithmic forecast goes even further in this direction. There are positive shocks around the mid-1950s and mid-1960s, but the shocks quiet down in the 1970s (the profile settles down to a smooth logarithm), and the 1972 episode does not show up at all, as presaged by figure 5.

Thus, the implication of virtually all the forecasts is that the shocks were larger in the 1950s and early 1960s than later. This result should not be surprising given the discussion of figures 4 and 5, for those illustrations also showed graphically that the benefit increases of the late 1960s and early 1970s had slight discernible effect on the rate of growth of social

23. More sophisticated forecasting equations were also estimated, such as equations including lagged values of S_t as regressors. The fits from those equations invariably produced implausible results, either predicting stationary values of S_t four to five years in the future or predicting exploding values of S_t . Recall that the difficulty here, relative to some other economic models, arises because the S_t values must be projected as much as sixty years into the future. Because more complex equations than equations 30 and 31 were unable to forecast plausibly this far in time, only the simpler equations were used.

security wealth. Nor is this result at all implausible. In the 1940s there was little reason to expect that the social security system would increase in generosity at the rate that it did in the 1950s, when coverage rates grew dramatically, when benefits were increased dramatically in 1950 and at slower rates thereafter, and when reciprocity rates grew rapidly among the elderly. In essence, the growth of the system in the 1950s represents the change from virtually no system at all (that is, one trivial in magnitude) to one with sizable aggregate wealth transfers. Retirees in the 1950s no doubt were much better off than they had expected they would be.

The lack of growth in the late 1960s and early 1970s is also not implausible once the principles of a pay-as-you-go system are realized, as noted earlier.²⁴ The cohorts retiring in that period were among the first to have paid taxes into the system for their entire lives, and the aggregate tax bite was increasing exponentially for each successive retiring cohort. For this reason the growth rate of wealth in a maturing pay-as-you-go system *must* fall unless benefits are increased by extraordinary amounts. What the results here indicate is that the benefit increases in the late 1960s and early 1970s, although large by some standards, were not sufficiently great to offset the forces of system maturation. Indeed, the implication of the results here is that, had those benefit increases not been legislated, retirees would have felt negative shocks when they retired. They did "expect" the benefit increases to occur; hence those increases were not shocks. This outcome seems implausible only if one believes that people expect current law to remain constant—an expectation that does not seem likely and one that the analysis here obviously does not assume. If people instead are thought to make only extremely simple linear or logarithmic projections of the trends they see in the past five to fifteen years, the wealth values in the late 1960s and early 1970s were easily forecastable.

Results

Tables 1 and 2 show the results of estimating equations 28 and 29 in level and difference form, respectively, by ordinary least-squares meth-

24. In the 1950s the social security system was not pay-as-you-go, for a large trust fund was accumulated. The trust fund was gradually dissipated and reached a stationary minimum amount only in 1965, from which time the system has been pay-as-you-go.

ods.²⁵ To recapitulate, recall that equation 28 requires regressing earnings on the wage rate, a cohort trend, a deviation in the cohort trend after 1937, age, and a weighted social security wealth variable (see also footnote b in table 1 for the definition of this last term). The age variable may or may not be made to interact with all other variables on the right-hand side of equation 28 except the wage rate. Equation 29 is equation 28 in first-difference form.

A set of weighted wealth variables is constructed for each of forecast equations 30–31. For each type of forecast equation, the estimated coefficients from each year's equation plus those of the previous year's equation are used to compute a shock value for individuals at all single years of age at that year. As described above, the shock value is equal to the change in the predicted value of end-of-lifetime social security wealth from one year to the next. Each of these shocks (one for each single year of age in each year) is then weighted by k_q (see equation 21), which is merely a weight for age. Finally, a weighted wealth value is calculated for each of the 108 observations in the sample (twenty-seven years, four age groups) by summing the weighted shock values over all ages of the lifetime for an individual of the mean age $\bar{a}$ in each age group in each year. This weighted wealth value (see equation 28 or footnote b in table 1) is then used directly in the level equations in table 1. In the first-difference equations is entered the change in the variable, which is, aside from the adjustment for the change in cohort composition of adjacent age intervals, roughly equal to the weighted wealth shock for an individual of each mean age.

The coefficient on the wage rate is always the estimate of α , which is only of secondary interest here. The coefficient on the noninteracting

25. Two econometric issues are worth mentioning here. First, there is a question about the correctedness of the standard errors on the social security variable, since that variable is a predicted value from an auxiliary equation. As Pagan has shown, if the predicted values are constructed from first-stage residuals, as those here are, an ordinary least-squares package gives correct standard errors. Adrian Pagan, "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, vol. 25 (February 1984), pp. 221–47. The second issue relates to an error-in-variables problem pointed out by Angus Deaton, "Panel Data from Time Series of Cross-Sections," *Journal of Econometrics*, vol. 30 (October–November 1985), pp. 109–26. The model here treats the aggregate means as true means in the population, whereas they are actually drawn only from samples. Hence the implicit treatment of data for groups of individuals—in successive years at successive ages—as being measures of means from identical cohort populations is incorrect because (more-or-less) independent samples are drawn from the CPS each year. The only solution to this problem requires an estimate of the covariance matrix of the regressors in the sample, which was not undertaken here because microdata from the CPS are not available for years before 1968.

Table 1. Earning Regressions by Type of Shock Forecast, Equation 28
Levels; moving averages^a

Variable	Without social security variables (1)	Linear, five-year		Linear, ten-year		Linear, fifteen-year		Logarithmic, ten-year	
		(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
W	2,319.9* (124.2)	2,446.8* (121.5)	2,556.3* (75.1)	2,612.2* (115.5)	2,402.1* (75.4)	2,834.8* (93.0)	2,112.8* (83.5)	2,121.9* (77.4)	2,650.2* (212.2)
$\bar{c}$	59.5* (14.1)	79.6* (14.3)	-104.5* (39.7)	76.7* (12.3)	-76.1* (38.1)	66.6* (9.3)	-58.4* (33.0)	74.3* (8.8)	-133.3* (35.9)
$\overline{cD}$	-43.3* (17.3)	-105.6* (27.0)	72.8* (40.7)	-111.6* (18.2)	59.1 (38.3)	-111.7* (12.8)	47.1 (33.2)	-115.1* (11.9)	67.0* (36.9)
$\hat{Q}^b$	...	-158.6* (41.6)	130.1* (34.1)	-215.9* (34.1)	204.8* (32.7)	-315.3* (27.0)	299.1* (37.9)	-458.7* (34.5)	-31.0 (58.6)
$\bar{a}^c$	...	...	-48.1* (16.4)	...	-20.2 (15.2)	...	-2.8 (3.5)	...	-101.6* (24.3)
$\hat{Q}\bar{a}$	...	...	-6.1* (1.8)	...	-9.1* (1.5)	...	-12.4* (1.4)	...	0.6 (3.0)
$\overline{ca}$	...	...	1.8* (0.7)	...	1.8* (0.7)	...	2.1* (0.6)	...	1.6* (0.9)
$\overline{cDa}$	...	...	-3.3* (0.8)	...	-3.5* (0.8)	...	-3.3* (0.7)	...	-1.9* (0.8)
Constant	-1,646.4	-766.6	-670.3	-909.5	-788.6	-1,186.6	-226.7	759.1	171.7
$R^2 \times 100$	93.9	96.5	99.3	95.6	99.4	97.4	99.6	97.7	99.2

Source: Author's calculations.

*Significant at 10 percent level.

a. Standard errors appear in parentheses; $n = 108$; R^2 is the coefficient of determination.

b. $\hat{Q} = \sum_{t=2}^T \kappa' Q_{\kappa}$ divided by 1,000.

c. $\bar{a}$ is average age minus 25.

social security wealth variable is the estimate of the negative of β , which is of direct interest here because it measures the wealth elasticity of labor supply and, hence, the effect of social security on labor supply. The parameter β is hypothesized to be positive, so the wealth coefficient is expected to be negative. When social security is made to interact with age, some of the social security effect enters through the interaction variable. In the level equation, equation 28, some of the other parameters (Y_0 , g_1 , and so on) can be identified from the other coefficients, but these parameters are of no direct interest.

Column 1 of table 1 shows the estimates of the level equation without any social security variables. This form is somewhat similar to some of the equations estimated by Browning (and others) and by MaCurdy.²⁶ The results indicate a positive and significant wage effect. At the mean hours of work of 1,900, the estimate of α implies an intertemporal substitution elasticity of 0.22. This wage elasticity is reasonably robust across all columns in tables 1 and 2, varying from 0.11 to 0.39 across the equations. Thus, positive but small intertemporal elasticities are present in the time-series data.

Columns 2 and 3 in table 1, whose shock variables employ the linear five-year moving-average specification, show estimates of equation 28 with and without age interactions (the model without age interactions assumes that $i = \rho$). The first equation shows a negative and significant coefficient on the wealth variable and implies that β is about equal to 0.16. This simple level equation thus gives some indication of the expected social security effects. These effects disappear, however, when age is made to interact with the wealth variable, as shown in column 3, for the coefficient on the noninteracting wealth variable has the "wrong" (that is, unexpected) sign and is also significant. The interaction coefficient between age and the social security shock is negative and significant, however, and implies that social security has a negative effect on labor supply for individuals over the age of 40. This result is not implausible, for it may be that only after age 40 do individuals begin to make decisions influenced by retirement expectations. The magnitudes of the coefficients imply, for example, that at age 65 a dollar increase in $\hat{Q}$ reduces earnings by about 0.11. The implication is, in turn, that the negative coefficient in equation 2 was due primarily to effects at the older age points.

26. See Browning and others, "A Profitable Approach," and MaCurdy, "An Empirical Model."

Table 2. Earning Regressions by Type of Shock Forecast, Equation 29
First differences; moving averages^a

Variable	Linear, five-year			Linear, ten-year		Linear, fifteen-year		Logarithmic, ten-year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
ΔW	2,125.3* (173.9)	2,116.3* (176.1)	2,097.2* (182.0)	2,049.4* (177.0)	2,046.8* (182.5)	1,993.5* (183.2)	1,958.2* (189.5)	2,128.7* (177.2)	2,100.2* (185.1)
$\Delta(\bar{cD})$	...	12.3 (27.8)	-9.9 (109.7)	29.6 (29.2)	20.9 (111.3)	30.9 (29.3)	4.0 (107.1)	4.1 (27.1)	-44.1 (110.6)
$\Delta \hat{Q}$	...	30.7 (24.0)	-9.0 (41.2)	-78.1* (36.6)	72.5 (73.5)	104.5* (47.5)	133.4 (87.9)	20.6 (34.5)	51.6 (60.8)
$\Delta[\hat{Q}\bar{a}]$	...	...	2.9 (2.5)	...	0.7 (3.8)	...	-0.4 (3.9)	...	-1.5 (3.3)
$\bar{a}$	...	...	-1.7 (2.3)	...	-1.3 (2.3)	...	-1.9 (2.3)	...	-1.7 (2.5)
$\Delta(\bar{cD}\bar{a})$	...	...	-0.0 ^b (2.9)	...	-1.1 (2.9)	...	-1.0 (2.9)	...	-0.1 (2.0)
Constant	-46.9	-55.0	-6.6	-66.6	-21.8	-69.0	3.8	-50.6	28.4
$R^2 \times 100$	59.4	60.1	61.2	61.2	62.1	61.3	62.6	59.6	60.5

Source: Author's calculations.

*Significant at 10 percent level.

a. Standard errors in parentheses; $n = 104$.

b. Less than 0.05 in absolute value.

Columns 4-9 in table 1 show the corresponding estimates obtained when the three other forecast equations are used to construct the wealth variables. The equations utilizing the other two linear forecasting equations yield results quite similar to those utilizing the five-year linear forecast. In both cases, the coefficients on the wealth variables are negative when no age interactions are included but positive when they are included. In both cases the interaction coefficients are again negative and significant and imply that negative effects begin among people in their 40s. At age 65 the ten-year and fifteen-year equations imply that an extra dollar of $\hat{Q}$ reduces earnings by 0.16 and 0.20, respectively.

The logarithmic equation in column 9 of table 1 gives quite different results. Whereas the wealth coefficient in column 8 is again negative and significant, in column 9 the noninteracting wealth coefficient is negative, but the interacting coefficient is positive; both are insignificant. These results therefore indicate a very weak negative response to social security that is spread over all ages more or less uniformly. Exactly why the results should differ in this way from the linear forecast equations is not clear. The logarithmic forecasts did imply, as noted previously, much smaller shocks than the linear forms, shocks that virtually disappear in the later years of the data set. These shocks are apparently much more weakly correlated with labor supply in general and with the labor supply values of older age groups in particular. Although there is no a priori basis for rejecting logarithmic expectations and accepting linear ones, the results based on the latter appear more plausible.

Table 2 shows the first-difference results, estimates of equation 29. As noted earlier, the wage coefficients are quite robust across all specifications. The estimates of the effects of social security, however, virtually disappear entirely in these sets of regressions. Among the four equations with no age interactions (columns 2, 4, 6, and 8), the shock coefficient is unexpectedly positive in three of four cases. In the one case in which it is negative (the ten-year moving-average equation), its value is about -0.08, smaller than, for example, the results for age 65 implied by the level equations. In one case, the fifteen-year moving average, the shock coefficient is not only positive but significant. When age is made to interact with the variables, the effects are further weakened. In only one case (the five-year moving-average equation) is the noninteracting coefficient of the expected negative sign, but there it has a standard error over four times its size. The interaction coefficients in the four specifications are negative only half the time and are insignificant in all cases. Even if

the standard errors in the equations are ignored, the point estimate of the effects of an extra dollar of $\hat{Q}$ is negative at age 65 only in the logarithmic specification, in which case it is only -0.008 .

Thus, the pattern of effects in the first-difference equations in table 2 is quite sensitive to the forecast equation used to generate the shock values and shows, in general, a virtually undetectable social security response. The null hypothesis that social security has no effect on labor supply could obviously not be rejected by the results in table 2. The reason for the difference in the results of tables 1 and 2 is relatively simple. In the level equations, the wealth coefficients reflect the adjusted correlation between the summed wealth shocks over an individual's lifetime and labor supply. Despite the finding that the pattern of wealth shocks differed from expectations between the 1950s and the 1970s, the summed wealth shocks show a rising profile over time. The negative coefficients on the wealth coefficients in the level equations thus arise naturally from the falling trend in labor supply and reflect the secular, negative correlation that has been discussed throughout this paper. In the first-difference equations, in contrast, only the adjusted correlation between the change in wealth values over time (that is, the shocks) and changes in labor supply influence the coefficient signs. As discussed in the last section, because of the pattern of shocks in the 1950s and 1970s the raw correlation between the shocks and labor supply is positive, not negative. Thus, although the hypothesis of a nonzero effect of social security on labor supply is consistent with the overall trends in total social security wealth over the postwar period, it is not consistent with the relative growth rates of wealth and labor supply in the 1950s and 1970s. This is the major finding of this study. It implies that, loosely speaking, the lower-bound estimate of the effect of social security implied by the time-series data is in essence zero.

Treating the estimates in table 1 loosely as upper-bound estimates allows the corresponding effects of social security on labor supply to be calculated. Consider a man 65 years old in 1955 and assume that a dollar of $\hat{Q}$ decreases earnings by 0.13. Then the effect of social security can be calculated from the *total* social security wealth for the individual at this age, which is equal to the sum of his initially forecast wealth level (not in the regressions) plus the sum of the weighted shocks $\hat{Q}$. For the man aged 65, the value of this total wealth in 1955 implies a reduction in annual hours of work of 3–5 percent across the four forecast specifications, each of which implies a slightly different value for the shocks. If he had known at age 25 that his social security wealth would be what it came to be in

1955, his hours would have been 2–3 percent lower at all ages of his life. As emphasized in the second section of the paper, the unexpected nature of the actual pattern of shocks causes a larger decrease. The lower labor supply effects are also those that would apply in a steady state in which social security wealth stabilized in value. For a man 65 years old in 1975, however, reductions in hours are 13–19 percent, and 9–14 percent had social security wealth been known in the beginning. The absolute value of social security wealth is, of course, larger for the cohort aged 65 in 1975 than for the cohort aged 65 in 1955.

A separate question is whether these upper-bound estimates can explain the acceleration in the decline in hours of work during the 1970s. The answer is not quite as obvious as it may appear. On the one hand, none of the four forecasting equations predicts larger shocks in the 1970s than earlier—in general, the opposite is true. But the cohorts approaching retirement age in the 1970s had begun their adult lives after the start of the social security system and, hence, I had higher expected wealth values all along. If these values were sufficiently high, their magnitude could outweigh the smaller size of the shocks they faced over their lifetimes.

Table 3 indicates that this is not the case. Under the linear ten-year moving-average forecast, men 55 years old in 1975 had a quite small initial expected wealth (\$329), given that in 1945 the system was small and that a linear forecast kept it so. Thus, a \$5,336 (= \$7,446 – \$2,110) increase in the value of wealth expected by the cohorts aged 55 in 1955 and 1965 was still larger than the \$2,507 increase for the cohorts aged 55

Table 3. Values of Social Security Wealth of Men at Ages 55 and 75 in 1955, 1965, and 1975

Moving averages

Age and year	Linear, ten-year			Logarithmic, ten-year		
	Total wealth	Wealth at age 25	Weighted shocks	Total wealth	Wealth at age 25	Weighted shocks
55						
1955	2,110	0	2,110	2,453	0	2,453
1965	7,446	0	7,446	3,927	0	3,927
1975	9,953	329	9,624	5,433	2,348	3,085
75						
1955	1,263	0	1,263	1,129	0	1,129
1965	4,975	0	4,975	3,746	0	3,746
1975	7,637	0	7,637	6,351	0	6,351

Source: Authors' calculations.

in 1965 and 1975. (The latter figure can explain about 20 percent of the reduction in hours over 1965-75.) Under the logarithmic forecast, in contrast, a man 55 years old in 1975 would have begun with an expected wealth of \$2,348. As a consequence, expected wealth at age 55 increased by about the same amount in the two intervals under this forecast. Nevertheless, this finding still cannot explain the large declines in hours worked in the 1970s. The panel in table 3 for men aged 75 gives the same results.

Summary

This paper has been a time-series examination of the effect of social security on labor supply. The existing time-series data for the United States, although quite scanty and weak in nature, allow one to estimate adjusted correlations between labor supply and measures of the generosity of the social security system. The evidence adduced here is intended to complement the cross-sectional and panel-data studies of the labor supply effects of social security, which have covered only the experience in the late 1960s and 1970s, and to determine whether the results of those microdata studies are consistent with trends in the earlier postwar years.

The major finding of the study is that there is an important inconsistency between the microdata studies in the later periods and the time-series data for the entire postwar period. The microdata studies in general show nonzero effects of social security on labor supply, although the effects vary in magnitude from one study to another, and also in general support the hypothesis that the sharp acceleration in the decline of labor supply of men in the 1970s was in part a result of the benefit increases in those years. The time-series data are inconsistent with this pattern because they show that social security wealth grew faster in the 1950s than it did in the later periods. Furthermore, by any of the assumptions about expectation formation tested in this paper, the wealth increments in the 1950s were more unexpected than were those in the late 1960s and early 1970s. The 1950s was a period in which the social security system grew from one trivial in magnitude to one with sizable aggregate wealth transfers, the growth arising from both increased benefit levels and increases in coverage rates. Cohorts maturing in the late 1960s and early 1970s, in contrast, had paid taxes into the system for most of their working lives, and growth rates of coverage had stabilized by that time. As a result, the benefit increases of the late 1960s and early 1970s only served to keep social security wealth

more or less equal to what could have been easily forecast from earlier growth rates. This finding is in retrospect obvious. By the basic principles of a pay-as-you-go system, the internal rate of return of contributions must fall over time, as must the growth rate of the net wealth increment.

The results of the paper also suggest, however, that the overall negative correlation in the levels of social security wealth and labor supply in time-series data is preserved when only the unexpected wealth increments are used to identify the social security (that is, wealth) elasticity. In other words, even when a model is formulated that excludes the initial level of private assets and social security wealth that an individual faces at the beginning of his working life—thus removing some of the time-series correlation between the growth rates of private wealth and social security wealth—there is still an estimated disincentive effect of social security on labor supply. The results obtained using this method imply labor supply reductions of 3–5 percent for the aged in the 1950s and of 13–19 percent for the aged in the 1970s. As just noted, however, the results cannot explain much of the acceleration in the decline of labor supply in the 1970s: the maximum percentage of the 1965–75 labor supply decline that can be explained by social security in the estimates is 20 percent, smaller than the previous estimates of Diamond and Hausman or of Hurd and Boskin.²⁷

Appendix: Algorithm for Computation of Social Security Wealth

The algorithm for the present value of social security benefits received by an individual in cohort c by age a' is

$$(32) \quad SB_{ca'} = \frac{1}{N} \sum_{a=a_m}^{a'} \frac{R_{ca}(B_{ca}^m + 0.5B_{ca}^f)}{(1+r)^{a-24}},$$

where

N = size of cohort at age 20

R_{ca} = number of social security recipients of cohort c at age a

B_{ca}^m = mean real social security benefit for retired men of cohort c at age a

27. Diamond and Hausman, "Retirement and Unemployment Behavior," and Hurd and Boskin, "Effect of Social Security on Retirement."

B_{ca}^f = mean real social security benefit for wives of cohort c at age a
 a_m = 65 before 1961, 62 after 1961.

It is assumed that half of the men are married during receipt of benefits and have wives who are collecting under the men's earning records. The implicit reciprocity ratio in the calculation is R_{ca}/N . The value of this ratio is affected by the survival probabilities of the cohort, probabilities of employment and coverage during the working life, and probabilities of actually receiving benefits during retirement. SB is thus the wealth value for the average member of the cohort. The benefits are discounted to age 24 because the labor supply data used subsequently begin at age 25.

The algorithm for the present value of social security tax payments made by the average member of cohort c by age a' is

$$(33) \quad ST_{ca'} = \frac{1}{N} \sum_{a=20}^{a'} \frac{C_{ca}T_{ca}}{(1+r)^{a-24}},$$

where C_{ca} is the number of workers of cohort c at age a in covered employment, and T_{ca} is the mean real tax payment of a covered worker of cohort c at age a . The mean tax payment is obtained by applying the tax rates for employers and employees to the earning level of the group.²⁸ The implicit coverage ratio C_{ca}/N will change over time as coverage expands, as survival probabilities change, and as employment probabilities change.

The calculated value of net social security wealth at age a' is

$$(34) \quad S_{ca'} = SB_{ca'} - ST_{ca'}.$$

A real interest rate of 3 percent is used in the calculations.

Comment by Joseph F. Quinn

This paper by Robert Moffitt is a time-series analysis of the effect of social security on life-cycle labor supply. As are many of the papers in this volume, it really is two papers. The author first describes a common problem that exists in the time-series analysis of the effect of social security on labor supply. In the theoretical section of the paper, he sketches out a highly creative solution to this problem. He then describes the data construction needed to test the impact of social security on labor supply.

28. See Moffitt, "Trends in Social Security Wealth."

The second part of the paper presents the estimation and discusses the results.

The two parts of the paper are well integrated. The empirical work follows as closely as can be hoped from the theory. The first part is quite successful and interesting; the second, as I read it, is less so.

Below I will describe the problem at hand, and the nature of the solution that Moffitt has proposed, and discuss why I would have predicted trouble in the empirical section. In addition, I would like to use this opportunity to pose a more general question. What, if anything, can aggregate time-series analyses of this sort tell us about the nature of the labor supply decision, and about the effect of social security on retirement?

As we all know, the world is divided into two types of people—not “big-enders” and “little-enders,” as Swift would have us believe, but people who like microeconomic cross-sectional data and people who like aggregate time-series data. There are good reasons for being in either camp.

Moffitt begins his paper with a discussion of just these issues. Why cross-sectional data? Most current microdata sets have huge sample sizes that allow precise estimation of parameters and interesting disaggregation into subgroups. The microdata sets have extensive demographic and economic information on each person in the sample, in some cases including each person's entire social security record. Another advantage of this micro-cross-sectional approach is that there is extreme variation in the values of the demographic variables. There are people with no pension wealth and people with pension wealth in the hundreds of thousands of dollars. There are people who do not work at all and people who work 4,000 hours a year. There are people in good, poor, and mediocre health. They are all there in the microdata files.

Unfortunately, they are all there at the same time—the year of the cross section. This fact creates at least two problems with the cross-sectional approach. It is difficult to discern how a person will react to changes in his circumstances in the future on the basis of the responses of other people to different circumstances at a point in time. How I will react if you change my X_1 to X_2 may not be the difference between my current response to X_1 and someone else's current response to X_2 . Because it is the reaction to possible policy changes that is of primary interest, this problem can be a serious shortcoming of cross-sectional data.

The other problem is that, despite all the variation in the variables in the survey, there is no variation in certain common circumstances. At any

time, all Americans have the same president. We all have the same federal budget deficit. We all pay taxes or receive benefits under the same social security law. One cannot estimate the effect of factors that are the same for everybody in the cross section. With respect to social security, the subject of this paper, it is true that people under the same law have different benefits because they have different inputs (wage histories) to the common formula. The question then becomes how to differentiate between the direct effects of those inputs on labor supply and the effects of the output (benefits) that emerge from the common social security law.

An obvious solution to this problem is to examine time-series data, in which the institutional parameters do change. When one does so, another whole set of problems emerges, as Moffitt has discussed. First, one has samples of very small size, numbering now in the dozens rather than the thousands. For example, in Moffitt's research, hours of work for different age categories are available only since 1955. Second, one has data on many fewer explanatory variables. Many of these simply do not exist by age category, as Moffitt's research requires. Finally, many of the time-series variables that do exist are highly collinear with one other. This last problem with the time-series data is really the focus of Moffitt's paper.

Moffitt describes this as an identification problem. How can the analyst tease out, or identify, the individual effects of highly collinear explanatory variables? More specifically, Moffitt's hypothesis is that the decrease in the labor-force participation rates of men over recent decades has been at least in part a result of the increased generosity of the social security system. I am confident that he is correct. The problem, as he points out, is that at the same time that real social security benefits (or wealth) have risen, so have real incomes. An alternative hypothesis is that the rising real incomes are responsible for what we have observed, and that declines in labor supply might have occurred even if we had not had a social security system. Because social security wealth and incomes are highly correlated over time (as are many other macroeconomic variables), it is difficult to tell "who dunnit." If thunder and lightning always come together, which is scaring the dog? This is what Moffitt's paper is about.

The solution comes from a central feature of the life-cycle model that Moffitt describes. In a world of perfect certainty, a person makes a whole lifetime of decisions at the beginning, on the basis of initial conditions and what he knows is going to happen over the life cycle. An implication of this model is that an anticipated change in a variable way out in period t has a very small influence on contemporaneous decisions during period q since the effects of the change are spread over all the time periods.

If someone expects to win the John Bates Clark medal, for example—and many people do—and if that award were to come with a substantial cash prize, the recipient's consumption in all the time periods would be affected, not just consumption in the 40th year. An analogy is piles of ice cubes in a row, with each column representing an annual income. A year's consumption does not depend on that year's column of ice, but rather on the common level of water once all the ice cubes have melted. An additional cube in year q has only a small effect on the water level, and therefore on each year's consumption.

Now, here is the good news with respect to the problem at hand. *Unanticipated* changes in the state of the world can affect behavior only *after* the change. The later a given surprise or "shock" occurs, the larger the yearly effect of that change because there are fewer years left over in which to adjust.

Although social security wealth (from the growth in benefits) is highly correlated with many macroeconomic variables (such as real income), *unanticipated changes* in social security wealth are not highly correlated with these other variables. According to Moffitt, these unanticipated changes in social security wealth actually occurred in two big blips—one in 1950 and one in the late 1960s and early 1970s. The identification of the labor supply effect of social security comes from these unanticipated changes in social security wealth over the life cycle. This is a nice idea, and it is the primary contribution of the paper.

Let us look a little more closely at what the paper does. Moffitt begins with a description of labor force participation rates over time. There are declines since the 1960s for all age categories of men. The decline has been very modest for those under age 55 but dramatic for those aged 55 and older. For those 65 years old and over, the decline began earlier, before 1950. There is a similar story for annual hours of work, although it is not clear whether these hours of work are for all people or for only the people who are still working. The story is about the same in either case. There is not much new here; this decline is a well-documented and well-known phenomenon.

Moffitt then moves to a standard life-cycle consumption model. Utility depends on leisure and consumption in each year, subject to the lifetime budget constraint that initial wealth plus discounted earnings must equal discounted consumption.

Then comes the question. What happens if wealth A is augmented by ΔA ? All subsequent leisure and consumption demands rise. Moffitt demonstrates that the effect on leisure demand (or labor supply) is larger, the

larger is the increase in wealth and the later the increase comes in the life cycle. The intuition here is straightforward. The unanticipated ΔA means that one misplanned in the beginning. The later it occurs, the shorter is the time horizon over which one can adjust. In the extreme, if the ΔA comes in the last period, one has to consume it all at once.

Moffitt utilizes a Stone-Geary—Sir Stone-Geary now, I suppose—utility function, which is a good idea for at least two reasons. First, such a model yields a nice linear earnings function, and all Moffitt's empirical work is done with earnings rather than with hours of labor supply. Second, Sir Richard Stone himself is an appropriate test case for this theory because he has recently won the Nobel Prize in economics and has been awarded a huge ΔA .

In any case, the level of exogenous wealth in this model includes the traditionally defined assets, A , and a component called S . S is the net wealth that an individual (or a cohort, since aggregate data are used) will receive from the social security system; it is the difference between discounted benefits and discounted contributions. This social security component includes both an initial value and the unanticipated changes. Such changes are not simply added in but are weighted by a variable that depends on when the changes occurred. The later these surprises have occurred, the model predicts, the larger are the effects on behavior.

At this point Moffitt has produced a tractable linear earnings equation with initial wealth A and social security wealth S in it. Unfortunately, A is unobserved, so the equation as it stands cannot be estimated. Even if it could, only imprecise estimates of the two effects of A and S could be obtained because these variables are highly collinear in the aggregate.

Now here is the trick. The equation is first-differenced. In the first-differencing, traditional wealth and the initial level of social security wealth disappear from the equation. What remains is an equation in which the change in earnings from year to year is a function of the change in the wage rate and of the unanticipated social security shocks. The unobserved A and the potential collinearity disappear. The social security wealth effect is then identified by the unanticipated shocks that are left in the equation. This solution is very elegant and creative work.

Now the reader is prepared for the empirical work, and here is where one feels somewhat disappointed—like shopping in Poland, one cannot find any of the things one wants.

For example, hours of work are not available by age. Before 1955 there were no disaggregated age data available from the CPS. After 1955 such

data were available, but only for four broad age categories. The result is that, because of the data, one cannot follow a specific age cohort over time. If this is the case, then one cannot first-difference as one had hoped—one cannot first-difference the 51-year-olds in 1981 and the 50-year-olds in 1980. These would have been exactly the same people, which is what the theory was about.

What one can do is first-difference cohorts of those aged 45–64 in 1981 and 1980. But these are not exactly the same people. By 1981, the people 64 years old in 1980 have moved on to the next cohort and have been replaced by a new batch of 45-year-olds. So cohort effects began to sneak into the empirical work, after having washed out nicely in the theory. One cannot use dummy variables for each of the cohort-year combinations because there are too many of them, and these dummy variables would use up too many precious degrees of freedom. The solution is the introduction of an unobserved cohort factor that is assumed to grow steadily over time.

The second problem is that the CPS reports do not contain earnings by age—income, yes, but earnings, no. Moffitt therefore turns to published reports from the Social Security Administration. Unfortunately, these data include only the covered population, and the coverage changed dramatically over this time period. They also include earnings only up to the maximum taxable amount, necessitating imputation beyond that. These earnings estimates are then divided by estimates of hours of work to yield the wage rates that show up as an explanatory variable in the empirical work.

Moffitt next considers the social security wealth shocks, the key to the estimation procedure. One approach would have been to include dummy variables for when each of the big surprises occurred—around 1950 and again in the late 1960s and early 1970s. But this procedure would have implied that the size of the shock is the same for all cohorts, which it is not. The same change in the rule has different effects on different cohorts, both because it changes their wealth in different ways and because a given change in wealth has a larger effect on the behavior of older cohorts, who have fewer years left to adjust to the change. It is unlikely that a dummy variable would pick up these subtleties, so Moffitt has estimated the magnitude of the social security shocks for each of the cohorts.

The shock equals the change, from one year to the next, in the net aggregate social security wealth that a cohort expects to enjoy at age 80. Net social security wealth for a given cohort is defined as the difference between the present discounted value of past and future benefits, on the

one hand, and past and future tax contributions, on the other. In principle this difference will remain constant as long as a cohort's expectations about taxes and benefits remain unchanged.

The shock depends on what the cohort expected to receive at age 80. Rather than utilize the often changing legislation that regulates future benefits and taxes, Moffitt forecasts on the basis of what other cohorts of men 80 years old have enjoyed. He uses five-, ten-, and fifteen-year moving averages, in linear and logarithmic scale, and assumes that cohorts expect recent trends to continue.

A surprising finding emerges from these calculations. The net social security wealth of cohorts of 80-year-olds rose steadily between 1950 and 1980. (See figure 4 in the paper.) There are no jumps. *Natura* (nor social security) *non facit saltum*. The growth in the 1950s and early 1960s was primarily a result of expansions in coverage, and the maturation of a pay-as-you-go system. In the late 1960s and early 1970s the growth continued, but for a very different reason—the huge legislated increases in real benefits, and the subsequent indexing of benefits to inflation. From the perspective of an entire cohort (rather than simply those covered by the system), net social security wealth has been growing smoothly (although at a decreasing rate—see figure 5) for three decades. Moffitt's specifications implicitly assume that the labor supply effects of these two very different sources of increase in cohort wealth are the same—an assumption that may or may not make sense at the microeconomic level.

Figure 6 shows the shocks, calculated from this procedure, for a young and an old cohort. The most distinctive feature is the huge negative shock in the late 1970s, when the system failed to continue the real benefit increases of the early 1970s. The size of the disappointment, of course, is larger the shorter is the moving average, but the shock is still negative even with the fifteen-year window. Obviously, these negative shocks will make it difficult to explain the continuing decreases in the labor-force participation rate.

The results in the first-differenced earnings equations in table 2 of the paper bear this difficulty out. The shock coefficients are in general insignificant and often of the wrong sign. There is no evidence of a social security shock effect. (In the linear, *non*-first-differenced equations in table 1, a negative social security effect is found for men over age 40, but this, Moffitt explains, is probably due to the data collinearities discussed above.) The overall insight of the paper is this: the social security explana-

tion of recent retirement patterns, now much in vogue, is less obvious than it appears. The largest unexpected social security shocks occurred in the 1950s, when retirement changes were far less dramatic. Moreover, the shocks may have been negative during much of the 1970s, when the declines in labor-force participation rates accelerated.

What does a microdata fan (a little-ender) make of all this? I have two somewhat contradictory reactions. Moffitt's paper is a very useful addition to the retirement literature because it is one of the few papers to throw sand in the path of the "social security" bandwagon. Perhaps social security changes are not the sole answer to the question of recent retirement patterns. Perhaps such changes are not even the primary answer. Maybe, as Moffitt's results imply, these changes are not an answer at all.

But I doubt it. I doubt it because of the inherent limitations of the macroeconomic time-series data. They are too aggregate; none of the variables available is what one really wants. One is constrained to the aggregations and categorizations of others. I do not question Moffitt's use of these data. His is a competent and extremely creative paper. But that creativity itself reflects the problem—one has to be extremely innovative to work with the published macrodata. Every variable derived has a huge element of measurement error, which tends to bias coefficient estimates toward zero. Can a single aggregate social security variable possibly capture all the subtle effects of social security legislation on the lifetime budget constraint—the changes in wealth, the kinks at various levels of labor supply, and the intertemporal changes in the wage rate? It is too much to ask. When one adds to this the limitation of relatively small numbers of time-series observations, precise estimation of these behavioral effects becomes nearly impossible.

None of this is news to Moffitt. He acknowledges these shortcomings at the beginning of the paper. Time-series analyses "have problems associated with low degrees of freedom." Because of the small number of observations, and because many variables are highly correlated, "inferences made from time-series data are inherently weaker than those made from cross-sectional data. As a consequence, only simple tests have much statistical power; it is virtually impossible to discriminate between competing hypotheses of great sophistication and subtlety."

I agree. Moreover, for these reasons I tend to trust the cross-sectional, microeconomic results. These findings suggest that the wealth increments and the implicit wage cuts that the social security system imposes on older

workers do affect individual behavior and have contributed to the dramatic changes in retirement patterns documented here. Admittedly, no method—micro or macro—has been able to explain all, or even most, of the changes that have occurred, and in this respect Moffitt's results are consistent with the literature. There is more here than meets the eye, and plenty of room for new and inventive research.