

**DISCOUNT RATE HETEROGENEITY
AND SOCIAL SECURITY REFORM**

Andrew A. Samwick

Working Paper 6219

NBER WORKING PAPER SERIES

DISCOUNT RATE HETEROGENEITY
AND SOCIAL SECURITY REFORM

Andrew A. Samwick

Working Paper 6219
<http://www.nber.org/papers/w6219>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1997

This paper was originally presented at the 9th Annual Inter-American Seminar on Economics, Buenos Aires, Argentina, November 18-19, 1996. I would like to thank Sebastian Edwards, Martin Feldstein, Bill Gale, Alan Gustman, George McCandless, Jim Poterba, Jon Skinner, Steve Venti and seminar participants at IASE, the 1997 AEA meetings, and the 1997 NBER Summer Institute for helpful comments. This work has been supported by a grant from the National Institute on Aging to the National Bureau of Economic Research. Any errors are my own. This paper is part of NBER's research programs in Aging and Public Economics. Any opinions expressed are those of the author and not those of the National Bureau of Economic Research.

© 1997 by Andrew A. Samwick. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Discount Rate Heterogeneity and Social
Security Reform
Andrew A. Samwick
NBER Working Paper No. 6219
October 1997
JEL Nos. H55, E21
Aging and Public Economics

ABSTRACT

As many countries consider the privatization of existing pay-as-you-go Social Security systems, the option to make participation in the new system voluntary may appeal to policy makers who need to obtain the political support of their workers. A critical issue in evaluating such a reform and its economic consequences is the unobserved heterogeneity in households' preferences for consumption. This paper estimates the distribution of rates of time preference from the wealth data in the Survey of Consumer Finances 1992 and a flexible life-cycle model of consumption under income uncertainty. The estimated distribution is then applied to a variety of reform proposals that incorporate a voluntary choice of how much to contribute to a dedicated retirement account and a rebate of the existing payroll tax that increases with the magnitude of the contribution. The main finding is that an appropriate menu of reform plans can induce the voluntary buy out of 84 percent of existing payroll taxes at an immediate cost to national saving of less than 0.25 percentage point.

Andrew A. Samwick
Department of Economics
Dartmouth College
6106 Rockefeller Hall
Hanover, NH 03755-3514
and NBER
samwick@dartmouth.edu

I. Introduction

Social Security reform proposals have emerged in a variety of countries around the world in recent years. In developed countries, the combination of fiscal crises, aging populations, and maturing unfunded social insurance systems has sparked an interest in whether transitions to funded systems can restore stability to retirement income programs. In developing countries, the legacies of corrupt fiscal institutions and the perceived need to establish more advanced financial systems as a precursor to economic development have led many countries to follow the successful lead of Chile in privatizing existing state schemes.¹

A key element of any privatization is the explicit recognition of claims on the government that were previously implicit. As those claims are divested to private accounts, they will begin to earn higher rates of return than the old claims in the unfunded system.² The real return on the new accounts may also be higher due to favorable tax treatment and a lower risk premium associated with privately held assets. The introduction of opportunities to save at higher rates of return will generate competing income and substitution effects for saving. Aggregate saving will rise only to the extent that the latter outweighs the former.

This paper starts from the premise that maintaining the saving rate is a desired policy objective and considers ways in which this goal can be achieved while privatizing the public retirement income program. One way to ameliorate the reduction in saving is

¹ The Chilean privatization is discussed in Diamond and Valdés-Prieto (1994) and Edwards (1996). The current state of the U.S. system is discussed in Board of Trustees (1995), and a study of the problems of aging populations around the world is World Bank (1994).

² Feldstein (1995) shows that a higher return on private savings than on the unfunded pay-as-you-go contributions is one of the necessary conditions for a privatization to raise economic welfare.

to make the benefits of privatization an increasing function of the amount of saving that the worker does. A simple method of implementing this relationship is to have workers contribute to illiquid retirement accounts and then have the government rebate a portion of the existing payroll tax as a share of the amount contributed. Workers therefore reap the benefits of the privatization in proportion to how much they save, and, in effect, “buy themselves out” of the public program through higher levels of saving. Allowing the worker to have some choice over the extent to which his retirement benefits are privatized may also help policy makers gain support for the reform.

Designing and implementing fiscal policies such as a privatization that maintains the level of saving through a “buy out” plan with a voluntary component requires a detailed understanding of the saving behavior of the workers covered by the existing state program. Recent theoretical and empirical work in savings has shown that the standard life-cycle model under uncertainty is flexible enough to accommodate a variety of patterns of wealth accumulation, depending on the parameterization of the utility function and the intertemporal budget constraint.³ Carroll and Samwick (1997) have shown that much of the variation in the way wealth is accumulated over the life cycle can be summarized by the value of the household’s rate of time preference. The purpose of this paper is to demonstrate the feasibility of a buy out program by solving a life-cycle model of consumption twice; first to estimate the distribution of household rates of time preference and second to use those estimates to simulate the savings response to a buy out.

³ The recent theoretical literature is due to Kimball (1990). Engen (1994), Hubbard, Skinner, and Zeldes (1994, 1995), Carroll (1992, 1997), Carroll and Samwick (1995, 1997), Samwick (1995), and Gourinchas

The paper is organized as follows. Section II describes a simple voluntary privatization scheme in which workers buy themselves out of the existing state program. Section III develops the life-cycle model that will be used to estimate the distribution of rates of time preference, and Section IV applies that model to the data on wealth holdings in the Survey of Consumer Finances, 1992. Section V simulates the take-up and saving effects of a variety of buy out plans. The possibility of offering more than one type of buy out plan simultaneously is discussed. Section VI discusses the differences in the way such a reform could be implemented in developing and transitional economies compared to developed economies. Section VII suggests directions for future research and concludes.

II. Voluntary Privatization Scheme

This section illustrates a simple voluntary scheme to privatize some or all of the current Social Security program. The scheme is structured to allow workers to “buy themselves out” of the current program by saving in a dedicated retirement account.⁴ The concept underlying this scheme is similar to the feature of the Social Security system in the United Kingdom in which employees may use occupational or personal pensions to “contract out” of the earnings-related portion of the state benefits.⁵

Currently, the payroll tax is levied at a constant rate of t_p on all income Y below the maximum taxable earnings (MTE). Under the voluntary scheme, the tax rate will be

and Parker (1996) provide empirical evidence for various parameterizations of the life-cycle model under uncertainty.

⁴ The idea for this type of buy out plan grew out of work with Martin Feldstein on Social Security privatization. See Feldstein and Samwick (1996, 1997).

lowered for any worker who chooses to save in a dedicated retirement account in excess of an income-related threshold. The worker's benefits from Social Security will be prorated in accordance with elections to opt out of the system. An example of a new tax schedule is:

$$(1) \quad T(Y, s) = \begin{cases} (\alpha + \beta)Y, & s \leq s^*(Y) \\ (\alpha + \beta)Y - \gamma(s - s^*(Y)), & s^*(Y) < s < s^*(Y) + (\beta/\gamma)Y \\ \alpha Y, & s^*(Y) + (\beta/\gamma)Y < s \end{cases}$$

where the parameters can be summarized as follows:

The minimum possible payroll tax rate allowable for a worker who chooses to take full advantage of the opting out provisions is α . A convenient benchmark is to choose a value of $\alpha = 0$, which allows for the entire payroll tax to be eliminated. The maximum possible rate reduction in the payroll tax for a worker who fully buys out of the system is β . A natural choice for β is to set $\alpha + \beta = t_p$, which guarantees that even those who do not buy out at all do not pay higher payroll taxes. Alternatively, if staying with the old system is something to be discouraged or if additional revenues are needed, the payroll tax rate on those who do not buy out can be higher.

The rate at which a dollar of saving reduces the amount of payroll taxes paid is γ . Alternatively, this parameter is inversely related to the amount of saving that must be done to achieve fully opted-out status. The range for γ is between 0 and 1. If $\gamma = 0$, then there is no buy out. If $\gamma = 1$, then the buy out is completely free (i.e., it requires no additional contributions beyond the existing payroll tax). The higher is γ , the more

⁵ See Disney and Whitehouse (1993), Disney (1995), and Budd (1996) for recent analyses of the

tempting it is for individuals to buy themselves out. The benchmark value of this parameter is taken to be 0.5.

The savings threshold below which buying out is not possible is $s^*(Y)$. Along with the payroll tax and the contribution to the dedicated retirement account, the threshold is specified in units of annual income. This function should rise with income, but the relationship need not necessarily be linear. In general, it would be desirable to structure it so that workers at every income level have a marginal incentive to buy out some of their benefits. Providing such an incentive may also require increasing the MTE for the payroll tax. A schedule that rises faster than linearly with income may also reduce the size of the transfer to high income workers who currently save more than βY (in addition to their payroll taxes) on their own.

Table 1 provides some examples of $T(Y,s)$ for various combinations of the parameters $\{\alpha, \beta, \gamma, s^*(Y)\}$ and the choice variable (denoted as s) of how much to contribute to the account. In each cell in the right panel of the table, the top number is the amount of $T(Y,s)$, the payroll tax rate paid for the specified amount contributed to the dedicated retirement account. The number in parentheses below is the amount of $s + T(Y,s)$, or the total amount of retirement 'saving' that is done. Based on the current OASDI system, the value for t_p is set to 0.12.

The first row of the table shows the results when workers can buy themselves out of the full payroll tax ($\alpha = 0$; $\beta = 0.12$), contributions to the dedicated account offset the payroll tax dollar-for-dollar ($\gamma = 1$), and the buyout begins with the first dollar contributed ($s^*(Y) = 0$). The total payroll tax plus dedicated contribution is simply equal to the

contracting out feature of Social Security in the U.K.

dedicated contribution. In the second row, a threshold of 6 percent of income is set for $s^*(Y)$, and this increases the total tax plus contribution one-for-one if workers still choose to buy themselves out of the program. The third and fourth rows show that reducing γ from a dollar-for-dollar buy out to a “50 cents on the dollar” buy out will also increase the amount of dedicated contributions. The second group of parameterizations shows that increasing α to a nonzero value such as 0.03 has the effect of increasing taxes plus contributions for those scenarios in which the buy out rate was very generous ($\gamma = 1$) and workers were buying themselves entirely out of the scheme.

This simple buy out scheme allows for considerable flexibility in designing a privatization policy. For example, a country might choose to keep α positive in order to ensure that existing benefits in a pay-as-you-go system can be at least partially maintained.⁶ If policy makers know that workers with high incomes are more likely to be saving heavily for retirement beyond what is provided by the state system, then a schedule of $s^*(Y)$ that is increasing and convex in Y would be appropriate. This schedule ensures that the income effect of a cheaper means of funding retirement consumption does not generate a reduction in total savings.⁷ Similarly, if it is believed that workers will respond substantially to a change in the after-tax rate of return to saving, then γ should be kept low to provide a marginal incentive to save over a wider range.

⁶ In the U.K., the counterpart to α is greater than zero in order to finance a flat-rate benefit that is unrelated to earnings.

⁷ The use of nonlinear schedules for $s^*(Y)$ raises the possibility that if the buy out decision is made repeatedly, workers will opt out only in certain years when they can obtain the full tax advantage and not in others. For $s^*(Y) = 0.03$, $\beta = 0.09$, and $\gamma = 0.5$, a worker who was originally saving only 0.03 each year could choose to fully buy out once every seven years without increasing total savings at all.

The examples in Table 1 illustrate a wide range of potential responses to the reform proposal and suggest two important issues that will be discussed below. The first is the extent to which households will choose each option and what their choices imply about the Social Security system that will remain after the buy out. The second is the extent to which households will finance their chosen contributions by reducing other saving rather than current consumption. The heterogeneity in household budget constraints and preferences will be critical in both cases. The remainder of the paper specifies a model of consumption under uncertainty in order to assess this heterogeneity and simulate its impact on the reform.

III. Estimation Procedure

The workers' response to such a scheme depends critically on their preferences for future versus current consumption. This section characterizes these preferences by estimating the distribution of time preference rates from household wealth data and a model of optimal consumption under uncertainty. The methodology treats the household's rate of time preference as a residual parameter that reconciles the observed value of the household's wealth-to-income ratio given the values of the other parameters of the household's intertemporal budget constraint.

It should be noted at the outset that the voluntary nature of the scheme proposed in Section II is not essential for the distribution of time preference rates to be of interest. Allowing for a choice of how much to participate in the new system leads naturally to a question of what types of households make the choice. The answer below will be given based on the rate of time preference. However, the analysis of the savings response to

any change to the household's budget constraint must take into account the household's preferences, regardless of whether the household has the choice of what type of saving to do in addition to the choice of how much to save. For example, Cifuentes and Valdés-Prieto (1997) have used the distribution estimated here to simulate the aggregate saving response to a mandatory savings program.

Theoretical Framework

As a starting point, consider the following model of intertemporal consumption based on Samwick (1995):

$$(2) \quad V_s(A_s) \equiv \max_{\{C_t\}} E_s \sum_{t=s}^T (1+\delta)^{s-t} u(C_t)$$

such that:

- a) $A_t = (1+r)(A_{t-1} + Y_{t-1} - C_{t-1})$
- b) $Y_t = P_t V_t$
- c) $P_t = G_t P_{t-1} N_t$
- d) $A_s \geq 0$
- e) $A_T \geq 0$
- f) $C_t \leq A_t + Y_t \quad \forall t$

The value function $V_s(A_s)$ is the expected discounted utility of consumption in each period from the current period s to the final period T , discounted by a factor of $(1+\delta)^{-1}$ each period, where δ is the time preference rate. The date of retirement is exogenously fixed at period n . If the within-period utility function is then taken to be additively separable in consumption and leisure, the latter can be omitted from the maximization problem. The utility function has the Constant Relative Risk Aversion form:

$$(3) \quad u(C) = \frac{C^{1-\rho}}{1-\rho}$$

so that there is a precautionary motive for saving.⁸

Conditions (a)-(c) describe the accumulation of household assets and income over time. Next period's assets are equal to the current period's assets (A) plus income (Y) less consumption (C) multiplied by an interest rate $(1+r)$ which is nonstochastic and constant over time. Income is equal to permanent income (P) multiplied by a transitory shock (V). Permanent income in the current period is equal to permanent income in the previous period multiplied by a nonstochastic growth factor (G) and a stochastic shock (N).⁹ The stochastic components to income $\{N_t, V_t\}$ are assumed to be independently and identically distributed jointly lognormally with zero means and variances of the underlying normal distributions equal to zero and $\{\sigma_n^2, \sigma_v^2\}$, respectively.

The inequalities in (d)-(f) place restrictions on the value of the choice variable (consumption) and the state variable (assets). The initial value of assets is assumed to be nonnegative, as is the terminal value of assets at period T. Restriction (f) requires that the value of consumption be no greater than the current value of assets plus income received; that is, the household is liquidity constrained in that borrowing against even the nonstochastic component of future income (or pensions) is not allowed.

In this framework, retirement entails two changes to the income process. First, G_{n+1} is multiplied by a replacement rate factor such as 0.50, signifying that permanent income falls by 50 percent when the household stops working. This scenario describes a

⁸The requirement for a precautionary motive for saving is that marginal utility is a convex function, i.e. $u''(c) > 0$. See Kimball (1990) for an exposition of the theory of precautionary saving.

⁹An alternative parameterization of the stochastic component of income would be an AR(1) process as in Engen, Gale, and Scholz (1994) and Hubbard, Skinner, and Zeldes (1994, 1995) rather than the two shocks depicted here. Empirical estimates of ARMA(1,.) income processes have typically been above 0.90 and are often not significantly different from 1, yielding the random walk assumption imposed here. See Topel and Ward (1992).

worker with a combined pension and Social Security income that replaces half of pre-retirement wages during retirement. Setting G to a number less than one after retirement could incorporate the lack of complete indexation of retirement income to inflation or productivity. Second, income uncertainty in both the permanent and transitory components falls to zero after retirement, implying that once a retiree begins collecting pension and Social Security income, its real value is no longer uncertain.

The dynamic programming problem described above can be solved recursively back from the last period of life. The Euler Equation relating consumption in adjacent periods is given by:

$$\begin{aligned}
 (4) \quad u'(C_t^*[X_t]) &= (1+r)(1+\delta)^{-1} E_t u'(C_{t+1}^*[X_{t+1}]) \\
 X_t &\equiv A_t + Y_t \\
 X_{t+1} &= (1+r)(X_t - C_t) + Y_{t+1}
 \end{aligned}$$

where X_t corresponds to “cash on hand” defined by Deaton (1991), who shows that the optimal consumption amount, C_t^* is solely a function of X_t^* . The expectation is taken with respect to the values of the two future income shocks $\{N_{t+1}, V_{t+1}\}$ which comprise Y_{t+1} . Because the utility function is CRRA, numerical solutions for C_t^* must be found. The Euler equation allows the optimal consumption function in the current period to be deduced from the optimal consumption function in the following period. Because of liquidity constraints, C_t^* will be the lesser of the value given by the Euler equation and that given by condition (f).

Condition (e) and the absence of bequests imply that condition (f) will hold with equality in period T , generating the optimal consumption function $C_T^*[X_T] = X_T$. This

function can then be substituted into the Euler equation and into condition (f) to solve for $C_{T-1}^* [X_{T-1}]$ for an array of possible values of X_{T-1} . This process is repeated back to the first period of life to get $C_1^* [X_1]$, where $X_1 = A_1 + Y_1$ and A_1 is the initial value of the problem specified according to condition (d).

Representative consumption profiles are presented in Figure 1 for a household in which the head is 35 years old, the standard deviations of both shocks to income are 10 percent, and the replacement rate at retirement is 50 percent. The three curves represent different possible values of the discount rate. The axes are denominated in units of permanent income. For low values of cash on hand, households face a binding liquidity constraint and consume all available resources. As more cash on hand becomes available, households begin to consume less than their income. Households with higher discount rates have larger marginal and average propensities to consume than those with lower discount rates.

Predicted Age-Wealth Profiles

The optimal consumption rules can then be used to generate distributions of wealth holdings by age, as described in Samwick (1995), who used the implied wealth distributions to predict the extent to which higher retirement wealth entitlements offset other personal saving. For the present analysis, these predicted distributions can be used to infer the household's preferences. In particular, Equation 2 suggests that the household's expected level of wealth at a particular age will be a function of the following six parameters:¹⁰

¹⁰ Expected wealth also depends on the presence of liquidity constraints and the extent to which they are binding. With a priori information on which households face such constraints, the age-wealth profile could

- i) The expected growth rate of income (G)
- ii) The variances of permanent and/or transitory shocks to income (σ^2)
- iii) The replacement rate on retirement income (G_{n+1})
- iv) The interest rate earned on savings (r)
- v) The coefficient of relative risk aversion (ρ)
- vi) The rate of time preference (δ)

The first four parameters reflect the household's budget constraint and can therefore be specified based on empirical evidence. The growth rate of income and the variances of shocks to income can be estimated in a regression framework based on a panel of earnings data. The replacement rate on retirement income can be estimated from household data on pension and Social Security coverage, combined with a forecast of future earnings. As a first pass, the interest rate can be assumed to be constant across households and estimated from historical data on various portfolios of assets.

Having specified these aspects of the budget constraint, only the two parameters of the household's utility function remain. For simplicity, the coefficient of relative risk aversion will be fixed at either of two values, 1 or 2, leaving the rate of time preference for each household to be estimated from the data.¹¹ The general procedure is to solve the model in Equation 2 based on the household-specific budget constraint under a range of values of δ and generate the corresponding age-wealth profiles. Each profile will give a predicted value of the average wealth-to-income ratio at the household's age. The

be estimated separately for each group. In the absence of such information, this analysis will assume that the constraints are in place.

¹¹ An alternative strategy would be to fix δ and estimate ρ from the data. In the CRRA model, the estimation is more sensibly based on δ because an increase in ρ has countervailing effects on the optimal consumption rules. Because $1/\rho$ is the intertemporal elasticity of substitution, an increase in ρ will reduce the expected growth rate of consumption. However, because ρ also indexes the strength of the precautionary saving motive, an increase in ρ will increase the expected growth rate of consumption. The only effect of an increase in δ is to lower the expected growth rate of consumption.

household's estimated rate of time preference is the one that matches the predicted with the observed wealth holdings.

Figure 2 graphs the average age profiles of wealth for the three consumption functions shown in Figure 1. The household with the high discount rate of 10 percent maintains only a small buffer stock of wealth during most of its working lifetime. It begins to save for retirement only around age 55 and accumulates less than two years worth of income by age 65. The household with the low discount rate of 2 percent does not engage in buffer-stock behavior, steadily accumulating roughly seven years of income before reaching retirement. The household with a discount rate of 5 percent is intermediate between the other two; it engages in buffer-stock behavior only until age 40 and then saves actively for retirement. Variations in the discount rates permit the life-cycle model under uncertainty to encompass a broad range of age-wealth profiles. Linearly interpolating between the two values of δ whose predicted values for the average wealth-to-income ratio bracket the observed ratio generates the value of δ^* , the household's estimated rate of time preference.¹²

Empirical Issues

The wealth data to be used to estimate the distribution of δ^* are drawn from the Survey of Consumer Finances (SCF). Starting with the first wave in 1983 and conducted every three years since, the SCFs represent the most comprehensive household survey of wealth for the United States. The SCF is comprised of approximately 4000 households drawn from an area-probability sample and a high-income sample based on tax returns.

¹² Any randomly observed value of the wealth-to-income ratio is a consistent estimate of the mean of the predicted distribution of the wealth-to-income ratio. The linear interpolation method of estimating the true

Oversampling households with high income enables the SCF to provide a more accurate characterization of the distribution of assets whose ownership is highly concentrated. This analysis uses the SCF 1992, but could be applied to any of the waves of the survey.¹³ The sample is comprised of all households with positive labor income and a head of household between the ages of 25 and 65.¹⁴ Two measures of household wealth will be used: total financial assets and total net worth.¹⁵ In order to match the predictions of the theoretical model, the household's wealth-to-income ratio will be calculated, adjusting for a flat 20 percent income tax on labor income for all households.

The SCF provides labor market information for the head and spouse of each household that can be used to characterize the household's intertemporal budget set. Carroll and Samwick (1995) document significant differences in income uncertainty across different occupation, education, and industry groups using estimates of the "Equivalent Precautionary Premium" from Kimball (1990) and the Panel Study of Income Dynamics. For each household in the SCF, the head and spouse were assigned three values of income uncertainty based on the group-specific means in the appendix of Carroll and Samwick (1995). The three values were averaged for each worker to get a measure of uncertainty. The household's income uncertainty value is just the average of the values for the two spouses, weighted by each spouse's labor income. For the

value of the time preference rate will therefore also be consistent. A formal justification for this procedure is provided in the Appendix.

¹³ See Kennickell (1996) for documentation of the SCF 1992.

¹⁴ There are 2177 such observations, representing 48.9 million out of 95.9 million households in 1992. The head of household is simply the spouse with the higher labor income, or, in the case of the same nonzero labor income, the older spouse.

¹⁵ Total financial assets includes the amounts in employer-sponsored defined contribution pension plans (adjusted for an expected tax rate of 20 percent for all households). The household's entitlements from defined benefit pension plans and the existing Social Security system are captured in the retirement replacement rate calculation.

purposes of the simulation, the distribution of household uncertainty values was divided into three equally sized groups (high, middle, and low) and assigned values of $\{\sigma_\eta, \sigma_v\}$ equal to $\{0.15, 0.15\}$, $\{0.10, 0.10\}$, and $\{0.05, 0.05\}$, respectively.

The SCF also reports information on household's expected replacement rates from employer-provided pensions at retirement. Based on these responses, income-weighted average replacement rates were constructed analogously to the uncertainty values. To compute an expected replacement rate for Social Security, the benefit formula for 1992 was applied to each spouse's current labor income.¹⁶ The household's Social Security replacement rate was calculated as the income-weighted average of the two, or the value given by the dependent spouse benefit formulas if higher. The pension and Social Security replacement rates were added together and then assigned values of 0.25, 0.50, and 0.75 if the actual values were $[0.000 - 0.375]$, $[0.375 - 0.625]$, and $[0.625 - 1.000]$, respectively.

In order to keep the number of different simulations manageable, other aspects of the budget constraint were held constant across households. Income growth was assumed to be 1.5 percent per year before retirement and 0 percent per year after retirement. The interest rate was assumed to be a value of 4 percent for all households in all years, approximating the real, after-tax return on a portfolio of equity and debt over the last twenty years. Figure 2 shows that under some reasonable parameterizations of the model in Equation 2, households save only to maintain a buffer stock until retirement is only a few years away. Consequently, the household's "age" for the purposes of the simulations is increased by 65 minus the age at which the head of household reports that he or she

will retire. Every simulation is constructed so that the household begins economic life at 25, retires at 65, and dies at 85.

IV. Estimates of δ^*

The age-wealth profiles in Figure 2 form the basis of the estimation of the discount rates. For each combination of assumptions about income uncertainty, retirement replacement rate, and risk aversion, the average age profiles were generated for the following ten values of δ : $\{-0.15, -0.05, 0.00, 0.01, 0.02, 0.03, 0.04, 0.05, 0.10, 0.20\}$. For each household and measure of wealth, a value of δ^* was assigned by linearly interpolating the mean values of predicted wealth at the household's age. Summary statistics for the estimated distributions of δ^* are presented in Table 2. The six rows of the table in both the top (financial assets) and bottom (net worth) panels correspond to the six different combinations of risk aversion coefficients ($\rho = 1, 2$) and initial assets ($A_1 = 0, 1, 4$).¹⁷

The middle panel of the table gives the median value of δ^* and the 25th and 75th percentiles. For the baseline case in which $\rho = 2$ and $A_1 = 0$ and total financial assets is the measure of wealth, the median rate of time preference is 7.63 percent per year. One quarter of the time preference rates lie below 2.93 percent and another quarter lie above 14.66 percent. The last two columns show that 5.32 percent of the households had wealth-to-income ratios so high that a value of δ^* less than -15 percent would be needed

¹⁶This calculation makes the simplifying assumption that the individual's income growth is identical to the growth rate in average covered wages for the whole Social Security system in each year.

¹⁷The value of initial assets is denominated in units of permanent income in the first year of economic life.

to reconcile them and 14.17 percent had ratios so low that a value of δ^* greater than 20 percent would be required.¹⁸

The probability density function of the first parameterization is presented in Figure 3. The figure shows that the (truncated) distribution of estimated time preferences is skewed to the left, with very few households intermediate between 0 and -15 percent. Most of the time preference rates are within a few percentage points of the interest rate of 4 percent per year, and they are distributed fairly symmetrically around the median value of 7.63 percent.

Returning to Table 2, the next five rows show the effect of varying the level of initial assets and the coefficient of relative risk aversion on the estimated distribution. For a value of $\rho = 2$, increasing the level of initial assets to one or four times the initial value of permanent income shifts the entire distribution of time preference rates higher. If the household is assumed to start out with a positive amount of assets, then it has to depress consumption by less in order to achieve a particular wealth-to-income ratio at a given age. Since households with higher time preference rates choose higher consumption in the current period, this results in a higher estimate of δ^* for positive initial assets. For a given level of initial assets, lowering the degree of risk aversion to $\rho = 1$ shifts the distribution of δ^* lower. If a household is assumed to be less risk averse, it must have a lower rate of time preference in order to generate a particular wealth-to-income ratio at a given age.

¹⁸ Negative rates of time preference are unrealistic. A more reasonable explanation of the behavior of households for whom $\delta^* < 0$ is that they have received an inheritance or plan to leave a sizable bequest and do not plan to consume all of this wealth before they die. Although it would be preferable to be able to

The bottom panel of the table reports the estimated distributions when total net worth is used as the measure of wealth. Because financial assets are a subset of the assets that make up total net worth, only households with very large debts have net worth lower than financial assets. As a result, the estimated values of δ^* are lower when net worth is used. For most parameterizations, the median is between 3 and 4 percent, and there is more censoring from below and less from above than when the estimates are based on financial assets. For tractability, the theoretical model in Section III has only one asset and one form of consumption. To the extent that households utilize their non-financial wealth to support retirement consumption or insulate consumption from shocks to income, then net worth is the appropriate concept and financial assets will yield estimated δ^* that are too high. However, if households primarily use their housing and other non-financial assets for other purposes, then including those assets will result in estimated δ^* that are too low. The simulations below will use the first parameterization in Table 2, with financial assets as the measure of wealth, $\rho = 2$, and $A_1 = 0$.

The estimates in Table 2 do not lend themselves easily to a “goodness of fit measure,” since the method for computing the δ^* will assign a value for any possible wealth-to-income ratio. One way to provide an external check on the accuracy of the estimates is to determine whether they actually reflect household preferences for consumption in the current period relative to consumption in future periods. The SCF asks respondents to report their most and least important financial planning horizons.¹⁹

model bequests directly, the behavioral response to Social Security reforms of households with strong bequest motives can still be approximated by the estimated time preference rate.

¹⁹ The questions are “In planning your saving and spending decisions, which of the time periods listed on this page is most (least) important to you?” The set of possible responses is listed in Table 3.

The responses, along with the fraction of the households giving that response and the average value of δ^* for those households, are presented in Table 3.

For the question pertaining to the most important horizon, the responses are evenly distributed across the categories, ranging from the “Next Few Months” to “Ten Years or More,” and the average values of δ^* decline steadily with the planning horizon. The differences across groups are statistically significant below the one percent level and provide some evidence that the application of the life-cycle model under uncertainty is capturing the essential elements of household financial decision-making. The tabulations for the least important horizon provide weaker evidence, largely because very few households give results that are intermediate between the two extremes. Households who place little weight on the extreme short term are estimated to have low values of δ^* (an average of 6.38 percent), and those that place little weight on the long term are estimated to have high values of δ^* (an average of 9.20 percent). These results are consistent with the life-cycle framework.

The plausibility of the estimated δ^* can also be checked by examining the variation in δ^* by other characteristics of the household's optimization problem in Section III. Because each δ^* is a residual parameter, the estimates should be uncorrelated with other determinants of consumption such as age, permanent income, retirement replacement rates, and income uncertainty if the model is correctly specified. Figures 4 and 5 graph the average estimate of δ^* by five-year age categories (denoted by the highest age in the category) and income deciles, respectively. The estimated rates of time

preference decline with both age and income level.²⁰ These patterns suggest that the consumer's optimal consumption profiles change in a way that resembles greater patience with age and income that is not modeled explicitly. Understanding the reason for this possible misspecification is a topic for future research. For example, the hyperbolic discounting model of Laibson (1997) predicts both of these patterns. Becker and Mulligan (1997) also derive several predictions for the endogenously determined variation in rates of time preference, including the positive correlation between permanent income and patience.

The estimates of the time preference rates based on the life-cycle model under uncertainty show considerable heterogeneity. Their distribution is related via the optimal consumption rules and the shocks to income to the distribution of financial assets in the SCF. The pattern of estimates is consistent with self-reported information on financial planning horizons. Given the assumptions on the household's intertemporal budget constraint and the other parameters of its utility function, the estimated δ^* is the piece of information that is required to simulate the household's response to a buy out plan for Social Security.

V. The Savings Response to Privatization

The buy out plan described in Section II introduced a dedicated retirement account, contributions to which serve to reduce the worker's required payroll tax. The contributions earn market rates of return, either the after-tax rate or the pre-tax rate if the

²⁰ Analogous graphs of average discount rates by uncertainty and replacement rate deciles revealed no systematic patterns. Graphs for other parameterizations shown in Table 2 were similar in shape to those presented in Figures 4 and 5.

government decides to encourage workers to buy out by giving an extra tax advantage. Even at the after-tax return, the contributions have an advantage relative to saving outside the account because the contributions reduce the payroll tax, which typically has an internal rate of return below market rates.²¹ If there were no liquidity constraints or income uncertainty in the consumption model of Section III, then such a plan would enjoy universal popularity. Every worker would buy out completely, borrowing if necessary to maintain current consumption.

The key feature of the model is that the presence of income uncertainty and liquidity constraints makes the dedicated retirement account an imperfect substitute for other savings, just as the current Social Security program is an imperfect substitute for discretionary saving. The disadvantage of the retirement account relative to other savings accounts is that it cannot be used prior to retirement if needed to buffer consumption against random falls in income. The household's rate of time preference estimated in Section IV determines precisely how willing it will be to commit the extra resources to the dedicated account in light of this disadvantage.

Simplified 3-period Model

The life-cycle model under uncertainty from Section III must be modified in order to simulate a household's willingness to buy itself out of an existing Social Security system in this manner. The main modification is that the household now must choose how much to contribute to the dedicated retirement account in addition how much to consume. Additionally, these two choices will be made as a function not only of the

²¹ See Boskin, Kotlikoff, Puffert, and Shoven (1987) for calculations of implicit rates of return on Social Security contributions by birth cohort. Recent cohorts are projected to receive a real return of less than 1.5 percent on lifetime contributions.

amount of cash on hand but of the amount in the retirement account and the expected replacement rate on the old system after the effect of the buy out.

The complications of the one additional choice variable and two additional state variables necessitate several simplifying assumptions about the budget constraint in order to keep the problem tractable. The first is the reduction of the number of periods from 60 to 3, with each new period corresponding to twenty years of the previous life cycle. All rates of return will be compounded over 20 years in the simulations. These periods correspond to a young worker, an older worker, and a retiree. The second is that income uncertainty will be parameterized solely in terms of a binomial shock to income in the second period rather than as a combination of permanent and transitory shocks from a lognormal distribution. Income has a 50 percent chance of going up or down by 10 percent per year. The third is that payroll tax payments are assumed to earn a constant 1 percent per year rate of return (approximately the rate of growth of the Social Security covered wage).²² The previous strategy of specifying a particular replacement rate is no longer applicable because the replacement rate depends on how much of the old system the household chooses to buy out.

Given these changes, the model is essentially a means of analyzing how workers with different preferences will respond to the chance to buy out their first period contributions to the existing state system. Because the workers are assumed to make their contributions after the income shock is realized, the second period decision is made under

²² Although the rate of return on contributions is far from uniform (see Feldstein and Samwick (1992)), the retirement replacement rate from other income sources (such as employer-provided pensions) is also assumed to be constant at a rate of 25 percent. In reality, this replacement rate is increasing in income, so that the two simplifications at least partially offset.

certainty. Although similar models have been solved in more detail, this framework suffices to simulate the essential economic tradeoff embodied in the buy out decision.²³

Simulated Buy Out Plans

Tables 4a and 4b present five examples of different buy out plans. In each case, the household is assumed to start the initial period of economic life with no assets and to receive a 25 percent replacement rate in the retirement period from sources other than the existing Social Security system. The coefficient of relative risk aversion is assumed to be 2 in all cases. Contributions to the dedicated retirement account are assumed to grow at a pre-tax real rate of 6 percent per year, unless otherwise specified. For simplicity, the level of income in the first period is normalized to a value of 1, so that all saving amounts will be denominated as percentages of first period income.

The first panel of the table simulates the savings response to a buy out plan in which the first dollar of contributions goes toward reducing the payroll tax ($s^*(Y) = 0$), the entire payroll tax is eligible for the buy out ($\alpha = 0$), and each dollar of contributions reduces the payroll tax by 50 cents ($\gamma = 0.50$). Each row represents the behavior of a household with the specified rate of time preference (δ).²⁴ The first column shows the level of household saving under the old system (not including the payroll tax). The level of saving declines from 45 percent of income for households with $\delta = -0.05$ to 7 percent for those with $\delta = 0.04$. Households with higher rates of time preference were previously

²³ See, in particular, Engen, Gale, and Scholz (1994) for a multiperiod model in which the household chooses the level of consumption and the amount of a tax-preferred contribution in a 401(k) pension or Individual Retirement Account.

²⁴ Households with values of δ below -0.05 always fully opted out of any system, and those with values above 0.10 never opted out of any system. These two groups are therefore omitted from the table.

not saving at all, and the liquidity constraint prevented them from borrowing against future income.

The next three columns calculate the level of private saving after the buy out plan is offered. Both the dedicated retirement contribution and other saving that the household might do outside of the retirement account decline with the rate of time preference. Note that for $\delta = 0.02$ or 0.04 , households do not buy themselves out fully, nor do they reduce their non-retirement saving to zero. This is because the dedicated retirement account cannot be used to directly support consumption in the second period, due to the liquidity constraints. To the extent that the initial level of saving was chosen for precautionary reasons, the household will continue to save outside of the retirement account, even if that means passing up a higher interest rate. Households with $\delta = 0.06$ or 0.08 are induced to begin saving in order to take advantage of at least some of the buy out, and all of their saving is therefore done in the dedicated retirement account.

Subtracting the initial level of saving from the total private saving under the buy out plan yields the change in private saving. This number is always nonnegative, but, due to the interactions of precautionary saving and liquidity constraints, it need not be monotonically decreasing in the rate of time preference.²⁵ The next column gives the “Buy Out Proportion,” or the share of the payroll tax reduction for which the household was eligible that the household actually received. The change in government saving is simply the Buy Out Proportion multiplied by the original 12 percent payroll tax.²⁶ The

²⁵ Note that this nonmonotonicity should be a feature of the saving response to *any* tax-advantaged but illiquid form of saving, including IRAs and 401(k) plans. In general, this point has been overlooked in the literature that measures the savings response to tax-deferred retirement accounts.

²⁶ The change in government saving does not include any future reductions in capital income tax receipts that may result from the shifting of some saving that was previously being done at the after-tax rate to the

change in national saving is the sum of the changes in private and government saving. For the households with the lowest δ , this change is large and negative. For those with slightly higher δ , the change is negative but smaller. For those with $\delta = 0.06$ or 0.08 , the change is positive and equal to half of the contribution to the retirement account (because of the 50 cents on the dollar buy out rate).

For the purposes of comparing buy out plans, the last column of the table reports the Equivalent Variation (EV) for the reform. Based on Hausman (1981), EV is calculated according to:

$$(5) \quad V_s(A_s + EV | \text{old system}) = V_s(A_s | \text{new system})$$

where as in Equation 2, $V_s(A_s)$ is the household's expected discounted utility of consumption as a function of initial assets. EV measures the amount of wealth that the household would require to forego the opportunity to participate in the new system. It therefore provides a way to compare increases in welfare across both reform proposals and households. Because all reform plans reward higher saving and households with lower discount rates do more saving, households with lower discount rates reap more gains from each reform. Households with $\delta = -0.05$ value the reform as equivalent to 35.7 percent of first period income, while those with $\delta = 0.08$ value the reform as equivalent to only 0.5 percent of first period income.

The large decrease in national saving for the low δ households is potentially detrimental for the success of a reform. One way to modify the buy out plan to mitigate this saving reduction is to raise $s^*(Y)$ above 0, so that contributions reduce the payroll tax

dedicated retirement accounts that earn the pre-tax rate. Nor does it include the corresponding increases in

only after some threshold amount. Reform #2 raises $s^*(Y)$ to 6 percent of income but leaves the remainder of the plan unchanged. Households with $\delta = -0.05$ do raise their saving slightly, resulting in a reduction in national savings of only 8.16 percent. These households' expected utility has increased because by raising $s^*(Y)$, the reform has allowed the household a greater access to the pre-tax rate of return. The change in national saving is also higher and even slightly positive for intermediate values of δ , but those households with the highest values of δ have now chosen not to participate in the buy out at all. Expected utility is lower for all households with values of δ above -0.05, since the decision to contribute more to the retirement account (at the pre-tax rate of return) is not inframarginal for them and the threshold has forced them to save more for the same benefit.

Another attempt to mitigate the lower savings of the low δ households is presented in Reform #3, which returns the saving threshold ($s^*(Y)$) to zero but makes the buy out rate less generous with $\gamma = 0.25$ rather than 0.50. Lowering the buy out rate causes the household with $\delta = -0.05$ to save even more, causing a fall of only 3.37 percent in national saving. This plan is preferred to #1 because it allows the pre-tax rate of return to be earned on a larger range of contributions, but its expected utility is lower than #2 because it forces the household to save more to get that benefit. Under this plan, these households buy themselves out of only 82 percent of their existing benefits. Households with higher values of δ have changes in national saving similar to those in #1, but with lower expected utilities due to the lower buy out rate.

consumption tax receipts.

Reform #4 raises α to 3 percent of income, corresponding to a buy out plan in which the government needs to maintain at least some amount of revenue from the payroll tax. It is very similar to Reform #1, with the only differences being that households with values of $\delta = -0.05$ or 0, who were fully bought out of the existing system under #1, are made worse off. On average, the change in national savings for these groups together is also the same as in #1.

The first four reforms have focused almost exclusively on the potential for large reductions in the national savings of households with low rates of time preference. Policy makers may also be concerned about getting a large fraction of the population to buy out of the existing system, just to make the reform more permanent. Reform #5 is designed to address the generally low buy out proportions for high δ households in the other reforms. It allows for nearly dollar for dollar reductions in the payroll tax with $\gamma = 0.95$. It also does not offer a tax advantage to contributions, so that contributions earn only the after-tax return of 4 percent. Under this plan, households have to give up only a tiny amount of current consumption to buy themselves out of their existing benefits.²⁷ Because they are finally participating, the households with $\delta > 0.06$ have higher expected utilities than in any of the other reforms. Since they were saving positive amounts initially and are now fully bought out of the existing system, households with lower rates of time preference have reductions in national saving. Because this buy out plan is limited

²⁷ Note that because everyone buys out of this plan, it is equivalent to a mandatory privatization in which payroll taxes are simply raised by 0.63 percentage points. Feldstein and Samwick (1996, 1997) calculate a transition cost of about 2 percentage points of covered payroll for a privatization of the entire U.S. Social Security system.

in scope and offers no access to the pre-tax rate of return, these households prefer the other four reforms to it.

One way for policy makers to achieve multiple objectives is to offer several plans simultaneously. The last panel of the table presents the results of a reform that offers the plans under #1, #3, and #5 simultaneously. That is, households can choose between buying out with pre-tax contributions at 50 cents on the dollar, buying out with pre-tax contributions at 25 cents on the dollar, and buying out with after-tax contributions at 95 cents on the dollar. Each household will choose the plan that gives it the highest expected utility. The first column shows that the households with low rates of time preference ($\delta = -0.05$) choose #3 because they value the opportunity to save at the pre-tax rate more than other features of a buy out. This is a good plan for them because it gives them a comparatively high expected utility with a low reduction in national savings. Households with intermediate rates of time preference ($\delta = 0, 0.02, 0.04$) choose #1, generating buy out rates of 44 to 100 percent and reductions in national savings less than 2 percent. Households with high rates of time preference ($\delta = 0.06, 0.08, 0.10$) choose #5, a full buy out with a small increase in national savings.

Aggregate Saving Responses

Table 4 is a useful guide to the tradeoff that policy makers confront between promoting national saving and achieving higher utility at the household level through voluntary buy out plans. In order to calculate the aggregate effect of a buy out plan, the responses of the different households (identified by hypothetical values of δ) must be weighted by the proportion of the population that was estimated in Section IV to have

each value of δ . Aggregation can also be done in terms of dollars by weighting those proportions by labor income.

These two sets of calculations are presented in Table 5 for each of the reforms in Table 4 for the estimated distribution of δ for the financial assets, $\rho = 2$, $A_1 = 0$ scenario in Table 2.²⁸ The first panel weights the distribution by the sample weights of the households at each point and reports the weighted average values of each column in Table 4. The second panel weights the distribution by the total income of all the households at each point, in order to analyze the flows of total dollars. Since Social Security taxes do not cover all labor income, the most appropriate weights lie intermediate between these two sets. The proportions at each δ value are:

	-0.05	0.00	0.02	0.04	0.06	0.08	0.10
Population Weighted	7.79	4.05	14.26	15.31	6.04	10.79	41.77
Income Weighted	8.68	4.52	16.45	17.12	7.34	12.13	33.75

Weighting by income does shift the distribution lower, especially at the highest values of δ . Given the similarity in the two distributions, only the income-weighted results in the second panel will be discussed.

The first four reforms are similar in their effects on aggregate saving. Starting from a base of 8.53 percent of income, total private saving increases by 2 to 3.5 percent as approximately 17 to 42 percent of the potential contributions to the retirement account are made. Given the declines in payroll tax revenues ranging from 2 to 4.2 percent of

²⁸ For the purpose of these tabulations, households with estimated rates of time preference less than -0.05 (greater than 0.10) are assigned a value of -0.05 (0.10). All other households are assigned to the nearest value of δ in the range specified in Table 4.

income, the declines in national saving under each plan are between 0.01 percent and 0.75 percent of income. With considerably less than half of the existing payroll tax privatized, it is not surprising that there is not much reduction to national saving.

The fifth reform was designed to combat the problem of low participation rates by offering a very generous buy out rate at the after-tax rate of return. Its participation rate is 100 percent, but it comes at the cost of a fairly high (1.28 percent) decline in national saving. However, Table 4b showed that much of that saving decline was due to the behavior of households with low values of δ . By offering reform proposals #1 and #3 simultaneously with #5, these households can be encouraged to switch to buy out plans that do not generate as large a loss in national saving. Under this combination plan, participation falls to 84 percent, but the decline in national saving is cut to 0.24 percent, or one fifth of the decline under #5 offered by itself. The combination plan is therefore able to achieve the high rates of participation of plan #5 along with the small declines in national saving of plans #3 and #1. It also delivers an average gain in expected utility equal to approximately 8 percent of first period income.

VI. Implementing Reform

The countries that are considering privatizations of their public retirement programs include developed economies, developing economies, and countries making the transition from socialism to more market-oriented economies. While the model presented here was calibrated to the U.S. economy and patterned after the "contracting out" feature of the U.K. social security system, it is also similar to the recommendations for the mandatory savings pillar of the system advocated by the World Bank (1994). With some

tailoring, it should be feasible in a wide range of countries in which the preconditions for a successful transition exist.

Any country considering social security reform must have a government that is sufficiently credible in its promises to support the new financial systems. The most important issue is that the reform will involve explicit recognition of the debt that was formerly implicit in the existing unfunded system. The government must be able to issue this debt, which has typically been done in the form of recognition bonds, as in Chile, Peru, and Columbia's reforms. The new financial system also requires the government to follow a dynamically consistent policy of keeping the tax burden on the dedicated retirement accounts low. It further requires an appropriate regulatory framework, in which the retirement accounts are not forced to hold excessive amounts of government debt and have access to international capital markets. In light of the distributional consequences of the reform shown in Table 4 and the correlation of low time preference rates with income shown in Figure 5, the government may also need to implement additional progressive taxation in order to maintain public support for the reform.

There are also aspects of the existing capital markets that will affect the way a privatization occurs. In developed countries, there are other institutions such as employer-provided pensions that already exist for tax-advantaged saving beyond what is provided by social insurance systems. Since pensions may be close substitutes for the dedicated retirement accounts, there is a risk that employers will terminate their pension plans in order to allow their employees greater access to the new account. To the extent that employer-provided pension balances are net additions to national saving, such terminations could make the aggregate (negative) impact of the buy out on national

savings larger.²⁹ In developing countries, old age support beyond public retirement programs is likely to be provided through informal channels such as the extended family. The risk in the capital market is not that existing saving gets switched to the new program but that there is insufficient opportunity for the dedicated retirement accounts to be invested in non-governmental securities. The best possible outcome in that case is the one that occurred in Chile, where the new social security system developed along side other financial market liberalization.

Another feature of capital markets that was not addressed in the simulations but is important in practice is the likely increase in administrative costs that would result from a privatization, especially one in which the new program might already include three different buy out options. Mitchell (1996) reports that annual administrative costs on the Social Security program (OASDI) in the U.S. are approximately 0.9 percent of current *benefits*, or about \$14 (in 1994 dollars) per participant. Tabulations of annual administrative expenses on private pension plans are 0.7 percent of fund *assets*, or approximately \$140 per participant. Administrative expenses in developing countries with less experience in managing retirement accounts are likely to be even higher. In general, administrative costs will be higher in systems with actively managed accounts, generous provisions for switching account managers, and the lack of mandatory annuitization of retirement balances.

The precise objective of the policy makers will have an effect on which type of reform proposal is implemented. In countries with a legacy of inequity in the existing public scheme, a primary goal may be to get a very high rate of voluntary participation in

²⁹ See Samwick (1997) and the references cited therein for a further discussion of this point.

the new system. Countries with more severe fiscal positions may need to focus on maintaining a high level of saving during the transition to promote a stable macroeconomic environment for growth. Other countries may simply want to increase the welfare of their citizens through a more efficient fiscal policy.

VII. Conclusion

Many countries around the world have implemented or given serious policy attention to Social Security reform as a means of curing long-standing fiscal problems at the national level. While the long-term benefits of a privatization of Social Security include a higher savings rate, there is a very real danger of a reduction in national saving when the reform is implemented. Any reform that promises to make households better off may lead them to consume more immediately. If the country is already in fiscal troubles, the drain on national saving to finance the transition could spell the end of the reform. The analysis presented here shows that it is possible to construct a simple privatization scheme in which the amount by which households benefit is directly related to the amount they save in a dedicated retirement account.

The demonstration has two components, each requiring the solution of a life-cycle model under uncertainty. The first is to estimate the parameters of each household's intertemporal utility function, in order to know how household savings will respond to any proposed privatization. This distribution is clearly of interest in its own right as well. The second is to simulate the savings response for that population of households based on the estimated distribution. The methodology employed here has applications beyond the analysis of Social Security reforms.

The models presented here are meant to be illustrative of the types of reforms that could be undertaken. Several simplifying assumptions were made, and relaxing each one provides the basis for further work. The most obvious is the use of the stylized three period model to simulate the effects of the reform. There may be important issues related to the life-cycle pattern of opting out that deserve specific attention. A more careful modeling of a bequest motive to explain the very highest wealth-to-income ratios would also be preferable to the current estimation procedure, which assigns such households a negative rate of time preference. Another is that the redistributive element in the existing state systems has been largely ignored. There may be interesting dimensions of the joint buy out decision of two workers in a household that remain to be explored.

The simulation model also represents only the partial equilibrium analysis of the privatization. Because any change to savings flows will cause interest rates to change as well, the results presented in Tables 4 and 5 are likely the upper bounds of the immediate impact on national saving. More importantly, the analysis of the savings impact does not consider the effects of the reform on future income and consumption tax receipts. Although a full general equilibrium analysis is beyond the scope of the present paper, Cifuentes and Valdés-Prieto (1997) have incorporated discount rate heterogeneity into an OLG model under certainty and shown that a mandatory savings plan can have a sizable positive impact on saving in steady state.

Appendix: Estimating Discount Rates

Given a value of the rate of time preference, δ , and a sequence of random shocks to income, $\{\eta, v\}$ with cumulative distribution function $F(\{\eta, v\})$, the function $g(\delta; \{\eta,$

v) generates a value of the wealth-to-income ratio, w . This function incorporates the relevant information on the budget constraint and other parameters. The observed value of the wealth-to-income ratio, w^* , is a consistent estimate of the mean of the wealth distribution generated by the true value of the rate of time preference, δ_0 :

$$w^* \rightarrow \bar{w}_0 = G(\delta_0)$$

where:

$$G(\delta) \equiv \int_{(\eta, v)} g(\delta; \{\eta, v\}) dF(\{\eta, v\})$$

Define the inverse of this function, $h(w) = G^{-1}(w)$, so that:

$$h(\bar{w}_0) = \delta_0$$

$$h(w^*) = \delta^*$$

If $h(w)$ is a continuous function, then:

$$w^* \rightarrow \bar{w}_0 \Rightarrow \delta^* \rightarrow \delta_0$$

The actual value of δ^* is calculated by linear interpolation:

$$\delta^* = \delta_1 + \left(\frac{\bar{w}_1 - w^*}{\bar{w}_1 - \bar{w}_2} \right) \delta_2$$

where:

$$\delta_1 \equiv \max_{\delta_i} \text{ s.t. } \bar{w}_i \geq w^*$$

$$\delta_2 \equiv \min_{\delta_i} \text{ s.t. } \bar{w}_i \leq w^*$$

References

- Becker, Gary S. and Casey B. Mulligan (1997). "The Endogenous Determination of Time Preference," *Quarterly Journal of Economics* 112, 729-758.
- Board of Trustees, Federal Old-Age and Survivors Insurance and Disability Trust Funds (1995). *Annual Report*. Washington: U.S. Government Printing Office.
- Boskin, Michael, Laurence Kotlikoff, Douglas Puffert and John Shoven (1987). "Social Security: A Financial Appraisal Across and Within Generations," *National Tax Journal* 40, 19-34.
- Budd, Alan (1996). "The Pensions System in the United Kingdom," Manuscript, August.
- Carroll, Christopher D. (1992). "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence," *Brookings Papers on Economic Activity* 2:61-135.
- Carroll, Christopher D. (1997). "Buffer Stock Saving and the Permanent Income Hypothesis," *Quarterly Journal of Economics* 112 (February), 1-56.
- Carroll, Christopher D. and Andrew A. Samwick (1997). "The Nature of Precautionary Wealth," forthcoming in *Journal of Monetary Economics* 40.
- Carroll, Christopher D. and Andrew A. Samwick (1995). "How Important Is Precautionary Saving?" National Bureau of Economic Research, Working Paper No. 5194, July, forthcoming in *Review of Economics and Statistics*.
- Cifuentes, Rodrigo and Salvador Valdés-Prieto (1997). "Forced -Saving and Pension Reform in the United States," Manuscript, Harvard University, June.
- Deaton, Angus (1991). "Saving and Liquidity Constraints," *Econometrica* 59, 1121-1142.
- Diamond, Peter and Salvador Valdés-Prieto (1994). "Social Security Reforms," in B. Bosworth, R. Dornbusch and R. Labán (eds.) *The Chilean Economy: Policy Lessons and Challenges*. Washington, DC: The World Bank.
- Disney, Richard (1995). "Occupational Pension Schemes: Prospects and Reforms in the UK," *Fiscal Studies* 16, 19-39
- Disney, Richard and Edward Whitehouse (1993). "Contracting Out and Lifetime Redistribution in the UK State Pension System," *Oxford Bulletin of Economics and Statistics* 55, 25-42.

- Edwards, Sebastian (1996). "The Chilean Pension Reform: A Pioneering Program," National Bureau of Economic Research, Working Paper No. 5811, November.
- Engen, Eric M. (1994). "Precautionary Saving and the Structure of Taxation," Manuscript, Federal Reserve Board.
- Engen, Eric M., William G. Gale, and John Karl Scholz (1994). "Do Savings Incentives Work?" *Brookings Papers on Economic Activity* 1:85-151.
- Engen, Eric M., William G. Gale, and John Karl Scholz (1996). "The Effects of Tax-Based Saving Incentives on Saving and Wealth," National Bureau of Economic Research, Working Paper No. 5759, September.
- Feldstein, Martin S. (1995). "Would Privatizing Social Security Raise Economic Welfare?" National Bureau of Economic Research, Working Paper No. 5281, September.
- Feldstein, Martin S. and Andrew A. Samwick (1992). "Social Security Rules and Marginal Tax Rates," *National Tax Journal* 45, 1-22.
- Feldstein, Martin S. and Andrew A. Samwick (1996). "The Transition Path in Privatizing Social Security," National Bureau of Economic Research, Working Paper No. 5761, September.
- Feldstein, Martin S. and Andrew A. Samwick (1997). "The Economics of Prefunding Social Security and Medicare Benefits," National Bureau of Economic Research, Working Paper No. 6055, June.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker (1996). "Consumption Over the Lifecycle," Manuscript, Massachusetts Institute of Technology, March.
- Hausman, Jerry A. (1981). "Exact Consumer's Surplus and Deadweight Loss," *American Economic Review* 71, 662-676.
- Hubbard, R. Glenn, Jonathan S. Skinner, and Stephen P. Zeldes (1994). "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," *Carnegie-Rochester Conference Series on Public Policy* 40, 59-126.
- Hubbard, R. Glenn, Jonathan S. Skinner, and Stephen P. Zeldes (1995). "Precautionary Saving and Social Insurance," *Journal of Political Economy* 103, 360-399.
- Kennickell, Arthur (1996). *Codebook for 1992 SCF*. Washington, DC: Federal Reserve Board, April.
- Kimball, Miles S. (1990). "Precautionary Saving in the Small and in the Large," *Econometrica* 58, 53-73.

- Laibson, David (1997). "Hyperbolic Discount Functions and Time Preference Heterogeneity," Manuscript, Harvard University, July.
- Mitchell, Olivia S. (1996). "Administrative Costs in Public and Private Retirement Systems," National Bureau of Economic Research, Working Paper No. 5734, August.
- Samwick, Andrew A. (1995). "The Limited Offset Between Pension Wealth and Other Private Wealth: Implications of Buffer-Stock Saving," Manuscript, Dartmouth College, December.
- Samwick, Andrew A. (1997). "The Effects of Tax Reform on Pension and Nonpension Saving," Manuscript, Dartmouth College, May.
- Topel, Robert H. and Michael P. Ward (1992). "Job Mobility and the Careers of Young Men," *Quarterly Journal of Economics* 107, 439-479.
- World Bank (1994). *Averting the Old Age Crisis: Policies to Protect the Old and Promote Growth*. Policy Report Series, Oxford and New York: Oxford University Press for the World Bank.

Table 1: Payroll Taxes for Different Buy Out Plans and Contributions

Parameters of Buy Out Scheme:				Amount contributed (s):		
α	β	γ	$s^*(Y)$	0.06	0.12	0.18
0.00	0.12	1.0	0.00	0.06 (0.12)	0.00 (0.12)	0.00 (0.18)
0.00	0.12	1.0	0.06	0.12 (0.18)	0.06 (0.18)	0.00 (0.18)
0.00	0.12	0.5	0.00	0.09 (0.15)	0.06 (0.18)	0.03 (0.21)
0.00	0.12	0.5	0.06	0.12 (0.18)	0.09 (0.21)	0.06 (0.24)
0.03	0.09	1.0	0.00	0.06 (0.12)	0.03 (0.15)	0.03 (0.21)
0.03	0.09	1.0	0.06	0.12 (0.18)	0.06 (0.18)	0.03 (0.21)
0.03	0.09	0.5	0.00	0.09 (0.15)	0.06 (0.18)	0.03 (0.21)
0.03	0.09	0.5	0.06	0.12 (0.18)	0.09 (0.21)	0.06 (0.24)

Notes:

1) The top number in each cell reports the amount of payroll tax remaining after the buyout. The bottom number in each cell reports the sum of the remaining payroll tax and the contribution to the dedicated retirement account.

2) See the text for a discussion of the parameters.

Table 2: Distributions of Time Preference Rates

Parameterization:		Summary Statistics:			Percentages Censored:	
ρ value	A_1 value	Median	25th Percentile	75th Percentile	From Below	From Above
Financial Assets						
2	0	7.63	2.93	14.66	5.32	14.17
2	1	8.67	3.84	16.78	4.62	17.04
2	4	9.82	6.34	19.25	2.65	20.99
1	0	4.68	3.35	9.14	5.08	6.99
1	1	4.96	3.79	9.58	4.39	8.74
1	4	8.31	4.66	9.98	2.63	11.58
Net Worth						
2	0	3.16	-4.54	7.91	20.20	7.52
2	1	4.28	0.00	8.78	18.26	8.87
2	4	7.24	3.34	9.95	11.85	11.73
1	0	3.20	-0.72	4.76	19.31	5.65
1	1	3.75	1.71	5.00	18.04	6.44
1	4	4.75	3.52	8.46	11.83	8.82

Source: Author's estimates from the Survey of Consumer Finances, 1992.

Notes:

- 1) The top panel contains estimates when financial assets are used as the measure of wealth. The bottom panel estimates use total net worth.
- 2) The first two columns are the assumed values of the coefficient of relative risk aversion (ρ) and the household's initial assets (A_1) in the first period of economic life.
- 3) The middle three columns report the percentiles of the estimated distribution of time preference rates (in percentage points).
- 4) The last two columns are the percent of households whose estimates of δ were censored from below (at -15%) and from above (at 20%).

Table 3: Time Preference Rates by Financial Planning Horizons

Financial Planning Horizon	Horizon Cited As Most Important		Horizon Cited As Least Important	
	Average Value of δ^*	Percent of Population	Average Value of δ^*	Percent of Population
Next Few Months	10.43	18.22	6.38	49.55
Next Year	8.35	12.50	7.47	5.53
Next Few Years	7.77	25.51	9.46	3.85
Next Five to Ten Years	6.68	24.87	8.96	5.43
Ten Years or More	5.91	18.89	9.20	35.63

Source: Author's estimates from the Survey of Consumer Finances, 1992.

Notes:

1) The financial planning horizons are the households' responses to the questions, "In planning your saving and spending decisions, which of the time periods listed on this page is most (least) important to you?"

2) Values of δ^* were estimated using financial assets as the measure of wealth and values of $\rho = 2$ (risk aversion) and $A_1 = 0$ (initial assets).

3) Average values of δ^* include those households who are censored at -15 percent or 20 percent.

Table 4a: Simulated Effects of Reform Proposals on Private and National Saving

Time Preference Rate	Level of Initial Saving	Retirement Contrib.	Other Private Saving	Total Private Saving	Change in Private Saving	Buy Out Proportion	Change in Govt. Saving	Change in National Saving	Equivalent Variation
Reform #1: $s^*(Y) = 0, \alpha = 0, \gamma = 0.50, r^* = 0.06$									
-0.05	0.4576	0.2400	0.2432	0.4832	0.0256	1.00	-0.1200	-0.0944	0.357
0.00	0.2288	0.2400	0.0912	0.3312	0.1024	1.00	-0.1200	-0.0176	0.216
0.02	0.1408	0.1632	0.0479	0.2111	0.0703	0.68	-0.0816	-0.0133	0.135
0.04	0.0704	0.1056	0.0165	0.1221	0.0517	0.44	-0.0528	-0.0011	0.076
0.06	0.0000	0.0576	0.0000	0.0576	0.0576	0.24	-0.0288	0.0288	0.031
0.08	0.0000	0.0240	0.0000	0.0240	0.0240	0.10	-0.0120	0.0120	0.005
0.10	0.0000	0.0000	0.0000	0.0000	0.0000	0.00	0.0000	0.0000	0.000
Reform #2: $s^*(Y) = 0.06, \alpha = 0, \gamma = 0.50, r^* = 0.06$									
-0.05	0.4576	0.3000	0.1960	0.4960	0.0384	1.00	-0.1200	-0.0816	0.405
0.00	0.2288	0.2328	0.0734	0.3062	0.0774	0.72	-0.0864	-0.0090	0.181
0.02	0.1408	0.1512	0.0465	0.1977	0.0569	0.38	-0.0456	0.0113	0.102
0.04	0.0704	0.0936	0.0161	0.1097	0.0393	0.14	-0.0168	0.0225	0.044
0.06	0.0000	0.0000	0.0000	0.0000	0.0000	0.00	0.0000	0.0000	0.000
0.08	0.0000	0.0000	0.0000	0.0000	0.0000	0.00	0.0000	0.0000	0.000
0.10	0.0000	0.0000	0.0000	0.0000	0.0000	0.00	0.0000	0.0000	0.000
Reform #3: $s^*(Y) = 0, \alpha = 0, \gamma = 0.25, r^* = 0.06$									
-0.05	0.4576	0.3936	0.1287	0.5223	0.0647	0.82	-0.0984	-0.0337	0.391
0.00	0.2288	0.1824	0.0743	0.2567	0.0279	0.38	-0.0456	-0.0177	0.158
0.02	0.1408	0.1248	0.0472	0.1720	0.0312	0.26	-0.0312	0.0000	0.098
0.04	0.0704	0.0768	0.0164	0.0932	0.0228	0.16	-0.0192	0.0036	0.054
0.06	0.0000	0.0384	0.0000	0.0384	0.0384	0.08	-0.0096	0.0288	0.019
0.08	0.0000	0.0096	0.0000	0.0096	0.0096	0.02	-0.0024	0.0072	0.001
0.10	0.0000	0.0000	0.0000	0.0000	0.0000	0.00	0.0000	0.0000	0.000

Notes: See Text for Description of Columns.

Table 4b: Simulated Effects of Reform Proposals on Private and National Saving

Time Preference Rate	Level of Initial Saving	Retirement Contrib.	Other Private Saving	Total Private Saving	Change in Private Saving	Buy Out Proportion	Change in Govt. Saving	Change in National Saving	Equivalent Variation
Reform #4: $s^*(Y) = 0, \alpha = 0.03, \gamma = 0.50, r^* = 0.06$									
-0.05	0.4576	0.1800	0.3002	0.4802	0.0226	1.00	-0.0900	-0.0674	0.268
0.00	0.2288	0.1800	0.0948	0.2748	0.0460	1.00	-0.0900	-0.0440	0.205
0.02	0.1408	0.1656	0.0478	0.2134	0.0726	0.92	-0.0828	-0.0102	0.135
0.04	0.0704	0.1044	0.0166	0.1210	0.0506	0.58	-0.0522	-0.0016	0.076
0.06	0.0000	0.0576	0.0000	0.0576	0.0576	0.32	-0.0288	0.0288	0.031
0.08	0.0000	0.0252	0.0000	0.0252	0.0252	0.14	-0.0126	0.0126	0.005
0.10	0.0000	0.0000	0.0000	0.0000	0.0000	0.00	0.0000	0.0000	0.000
Reform #5: $s^*(Y) = 0, \alpha = 0, \gamma = 0.95, r^* = 0.04$									
-0.05	0.4576	0.1263	0.4369	0.5632	0.1056	1.00	-0.1200	-0.0144	0.083
0.00	0.2288	0.1263	0.1748	0.3011	0.0723	1.00	-0.1200	-0.0477	0.083
0.02	0.1408	0.1263	0.0874	0.2137	0.0729	1.00	-0.1200	-0.0471	0.083
0.04	0.0704	0.1263	0.0350	0.1613	0.0909	1.00	-0.1200	-0.0291	0.074
0.06	0.0000	0.1263	0.0000	0.1263	0.1263	1.00	-0.1200	0.0063	0.048
0.08	0.0000	0.1263	0.0000	0.1263	0.1263	1.00	-0.1200	0.0063	0.020
0.10	0.0000	0.1263	0.0000	0.1263	0.1263	1.00	-0.1200	0.0063	0.006
Reforms #1, #3, and #5 Simultaneously									
-0.05	#3	0.3936	0.1287	0.5223	0.0647	0.82	-0.0984	-0.0337	0.391
0.00	#1	0.2400	0.0912	0.3312	0.1024	1.00	-0.1200	-0.0176	0.216
0.02	#1	0.1632	0.0479	0.2111	0.0703	0.68	-0.0816	-0.0133	0.135
0.04	#1	0.1056	0.0165	0.1221	0.0517	0.44	-0.0528	-0.0011	0.076
0.06	#5	0.1263	0.0000	0.1263	0.1263	1.00	-0.1200	0.0063	0.048
0.08	#5	0.1263	0.0000	0.1263	0.1263	1.00	-0.1200	0.0063	0.020
0.10	#5	0.1263	0.0000	0.1263	0.1263	1.00	-0.1200	0.0063	0.006

Notes: See Text for Description of Columns.

Table 5: Aggregate Effects of Reform Proposals Based on Estimated Rates of Time Preference

Reform Number	Level of Initial Saving	Retirement Contrib.	Other Private Saving	Total Private Saving	Change in Private Saving	Buy Out Proportion	Change in Govt. Saving	Change in National Saving	Equivalent Variation
Weighted by Households									
#1	0.0758	0.0739	0.0320	0.1059	0.0301	0.31	-0.0370	-0.0068	0.0699
#2	0.0758	0.0687	0.0273	0.0960	0.0203	0.18	-0.0219	-0.0017	0.0602
#3	0.0758	0.0710	0.0223	0.0932	0.0175	0.15	-0.0177	-0.0003	0.0604
#4	0.0758	0.0671	0.0366	0.1037	0.0279	0.37	-0.0336	-0.0056	0.0625
#5	0.0758	0.1263	0.0589	0.1852	0.1095	1.00	-0.1200	-0.0105	0.0406
#1,#3,#5	0.0758	0.1538	0.0231	0.1769	0.1011	0.85	-0.1025	-0.0014	0.0777
Weighted by Household Income									
#1	0.0853	0.0838	0.0360	0.1197	0.0344	0.35	-0.0419	-0.0075	0.0789
#2	0.0853	0.0775	0.0308	0.1082	0.0229	0.21	-0.0247	-0.0018	0.0676
#3	0.0853	0.0801	0.0251	0.1052	0.0199	0.17	-0.0200	-0.0001	0.0680
#4	0.0853	0.0762	0.0411	0.1173	0.0319	0.42	-0.0381	-0.0062	0.0706
#5	0.0853	0.1263	0.0662	0.1925	0.1072	1.00	-0.1200	-0.0128	0.0453
#1,#3,#5	0.0853	0.1572	0.0260	0.1832	0.0979	0.84	-0.1003	-0.0024	0.0869

Notes: See Text for Description of Columns.

Figure 1

Consumption at Age 35 by Discount Rate

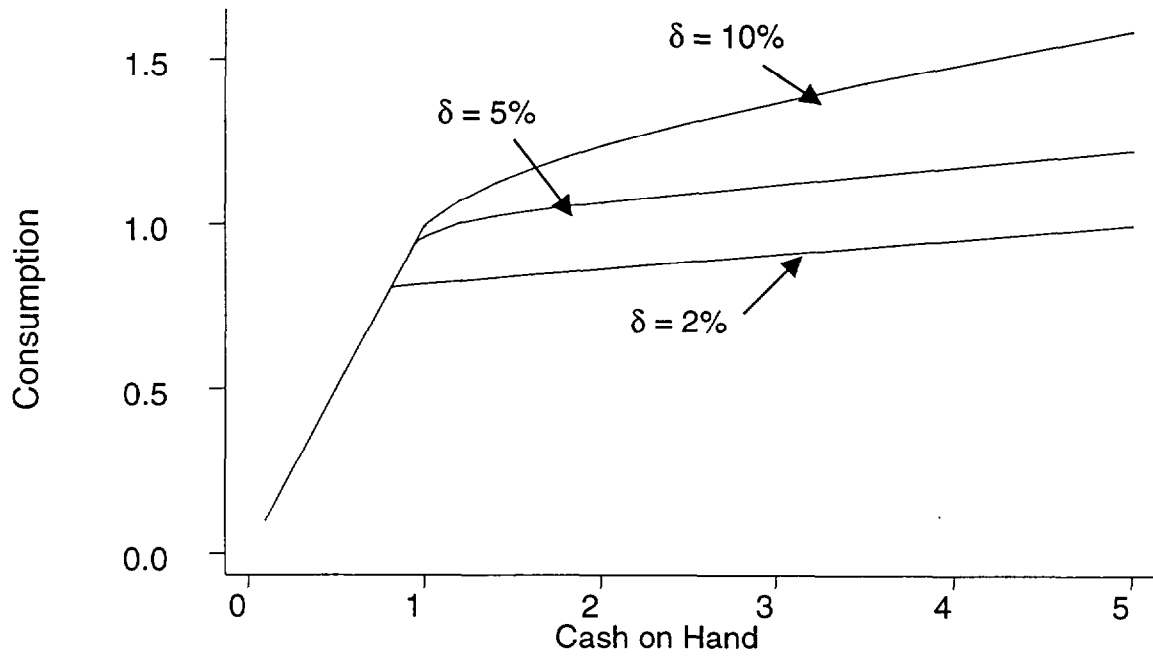


Figure 2

Age-Wealth Profiles by Discount Rate




Figure 3

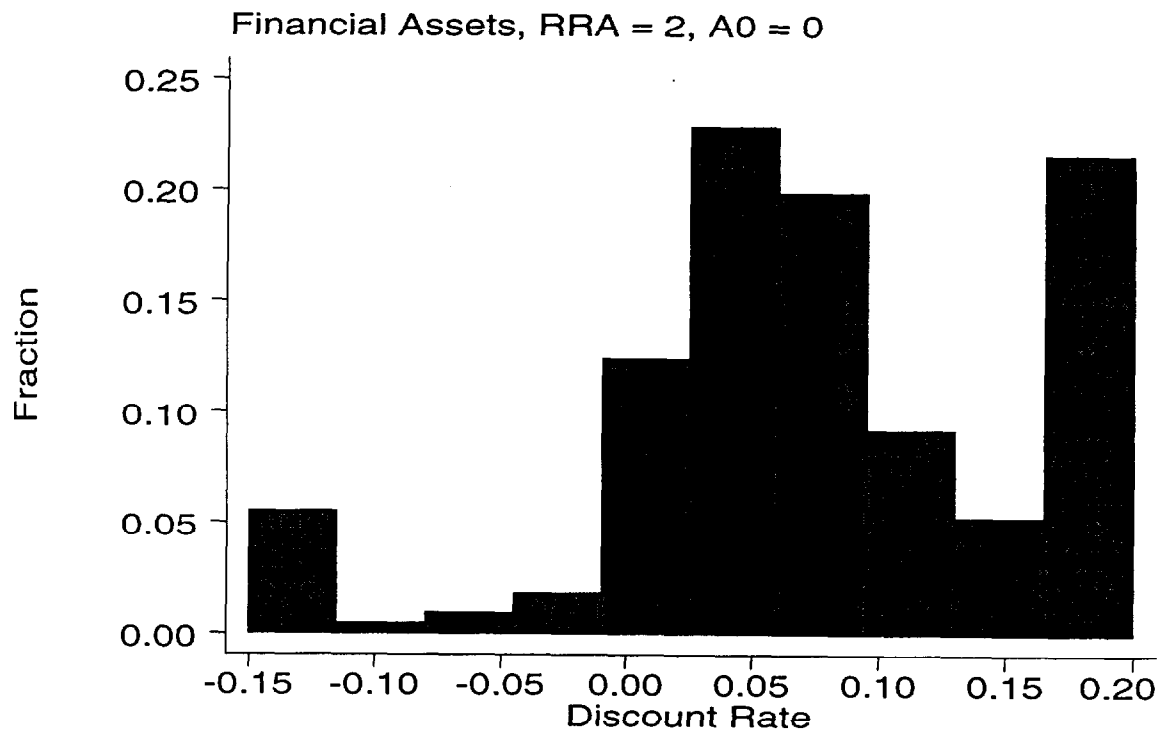


Figure 4

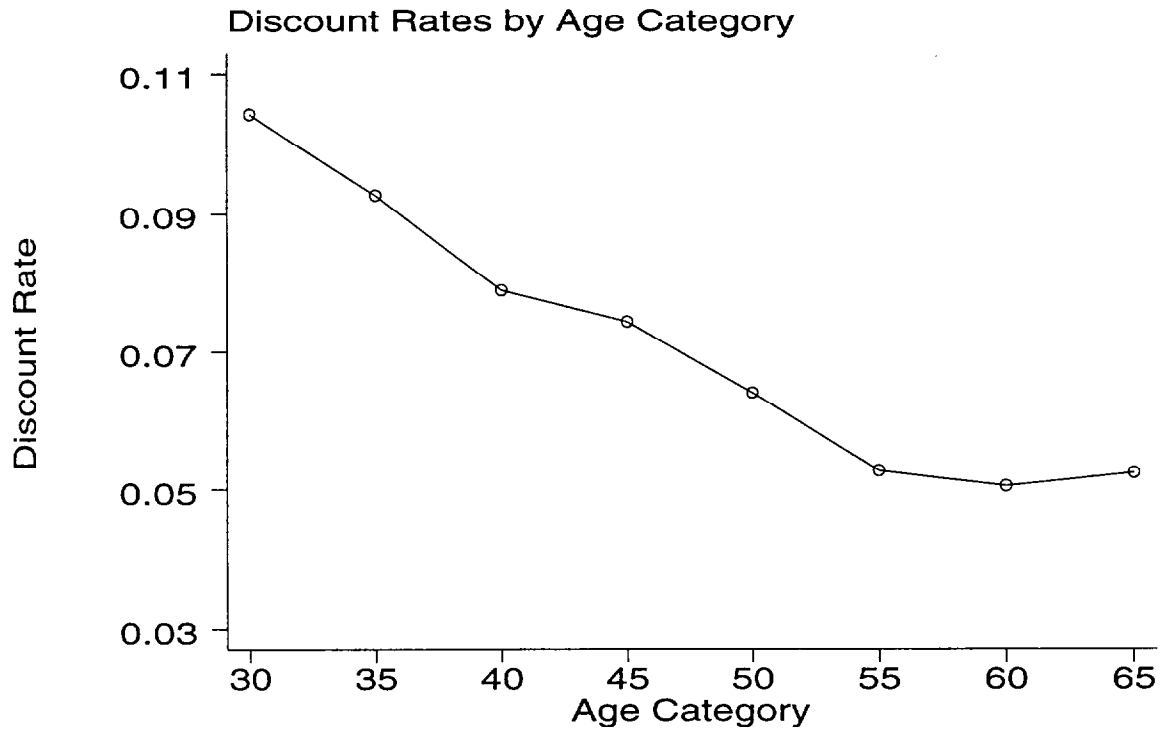


Figure 5

