How Do Pensions Affect Household Wealth Accumulation?

David M. Blau
The Ohio State University

Work in progress; preliminary and incomplete

First draft: March 2009
This draft: October 2009

Financial support under grant R01-AG02199 from the National Institute on Aging is gratefully acknowledged. I appreciate helpful comments on this and related work by Rob Alessie, Cristian Bartolucci, Luc Behaghel, Andrew Clark, Robert Clark, Gary Engelhardt, Robert Willis, and participants at the conference on Aging of the Population and Its Economic Consequences, Ramon Areces Foundation, Madrid, the conference on Saving for Old Age in Financial Turnover," CeRP, Collegio Carlo Alberto, Torino, the 2009 Population Association of America Annual Meetings, and seminar participants at Australian National University and Paris School of Economics. Helpful advice from Gary Engelhardt, Alan Gustman, and Tom Steinmeier on pensions, and from Wilbert van der Klaauw on model solution is also appreciated. None of the above are in any way responsible for the contents of this paper. Comments are welcome; contact the author at blau.12@osu.edu.

How Do Pensions Affect Household Wealth Accumulation?

Abstract

Empirical analysis of the effects of pensions on saving behavior is usually based on a stylized version of the life cycle model, with a fixed retirement age, a perfect capital market, no uncertainty, and no institutional restrictions on pension claiming. This model predicts one-for-one crowdout of household wealth by pension wealth over the life cycle. Empirical estimates of crowdout are often closer to zero than to -1. It is not well understood whether such estimates are accurate indications of the true magnitude of crowdout or whether the strong assumptions imposed in the analysis result in bias due to misspecification. In this paper, I specify a richer version of the life cycle model in which several of the key restrictions of the simplest model are relaxed. The effects of public and private pensions on household wealth are analyzed by solving and simulating the model. The simulated data are then used to estimate regressions like those typically found in the literature. Preliminary simulation results suggest that crowdout of household saving by public and private pensions in the more realistic life cycle model is relatively small, less than -$0.25 per dollar of pension wealth. Regression estimates of crowdout using the simulated data are close to this magnitude for private pensions, but much larger for Social Security. The magnitude of crowdout is quite sensitive to model specification.

1. Introduction

Public and private pensions are the most important sources of retirement income for most households in high income countries. An important question in the design of pension programs and evaluation of their welfare effects is how pensions affect household saving behavior. Pensions can be considered a form of "backloaded compensation," or forced saving. Economic reasoning based on the life cycle framework implies that workers will respond to forced saving in a pension plan by saving less in other forms. Thus pensions may "crowd out" household saving. The extent of crowdout may be affected by imperfections in the capital market (inability to borrow against future benefits or earnings), the tax-favored treatment of pension savings, legal restrictions and penalties on early access to pension benefits, labor supply responses to forced saving, and the importance of precautionary and other savings motives.

Analysis of pension effects on savings are usually based on a very stylized version of the life cycle model, with a fixed retirement age, a perfect capital market, no uncertainty, and no institutional constraints on pensions (tax treatment, restrictions on receiving a benefit while employed at the firm, etc.). Analytic solutions for optimal consumption and wealth profiles can be derived, and the model predicts complete crowdout: a one dollar increase in the present discounted value of pension benefits causes an increase in consumption in each period, with a present discounted value of one dollar. The increased consumption is financed by reducing household wealth (Gale, 1998). Under the assumptions of the model, one can compute a measure of the capitalized value of future pension benefits, referred to as "pension wealth," and use it as an explanatory variable in a regression of household wealth or saving.

It has long been understood that intuition about crowdout may not hold when the very strong assumptions of the simple life cycle model are relaxed. Feldstein (1974) showed that when the timing of retirement is a choice, the effect of Social Security wealth on private saving is ambiguous in sign (see also Crawford and Lilien, 1981, and Munnell, 1974). If an increase in pension wealth causes earlier retirement, then the additional saving needed to finance additional retirement years could result in a net increase in saving. In the presence of a borrowing constraint, the effect of future pension benefits on current household saving depends on the liquidity of the pension and the severity of the borrowing constraint (Gale, 1998). Uncertainty about future earnings, asset returns, and medical expenditures, and length of life will induce precautionary savings, which could affect the magnitude of pension crowdout. Other savings motives such as for bequest purposes could affect the extent of crowdout.

Despite awareness of these issues, empirical studies of the effects of pensions on household wealth not only base their intuition on the simplest version of the life cycle model, they also *impose* the assumptions of the model in the empirical analysis. The reason for this is that relaxing these assumptions leads to an intractable empirical model. It is very difficult to compute a measure of pension wealth if retirement age is a choice, if there is a liquidity constraint, or if there is uncertainty about future earnings, medical expenditure, etc. Empirical tractability is obviously an important concern, but it is not clear how to interpret empirical results estimated under such strong assumptions if the assumptions do not hold. The simple life cycle model predicts complete crowdout: a coefficient of -1 in a regression of household wealth on pension wealth (adjusting for stage of the life cycle). Coefficient estimates closer to zero than to -1 are common in practice. Researchers typically interpret such findings as indicating less than complete crowdout. But this finding is inconsistent with the simple life cycle model, and implies that the regression is misspecified, so it is not clear that such an inference is warranted. This is a highly policy-relevant issue: empirical estimates of pension crowdout effects are a key element in analysis of optimal pension design and the welfare effects of pensions.

In this paper, I specify a richer version of the life cycle model, in which several of the key restrictions of the simple model are relaxed. The model incorporates an employment decision, a liquidity constraint, realistic tax treatment of pensions, institutional constraints on pension claiming, and various forms of uncertainty. The parameters of the model are calibrated using data from the Health and Retirement Study (HRS). Pension plan formulas from actual pension plans in the HRS pension provider data base are used in solution of the model, along with formulas that determine Social Security benefits. The effects of pensions and Social Security on household saving, retirement, and pension claiming decisions are analyzed by solving the model and simulating behavior under alternative pension scenarios. The simulation results are used to measure the "true" magnitude of crowdout, and are also used to estimate wealth regressions like those found in the literature. Two measures of pension wealth are computed from the simulations: one that is consistent with the model and does not impose the usual strong assumptions of the simplest life cycle model, and a second measure that, incorrectly, does impose these assumptions. Using the "incorrect" measure is analogous to the approach of the empirical crowdout literature: strong assumptions are imposed, while the data are generated from a process in which these assumptions probably do not hold. The "correct" pension wealth measures can be used in the regression analysis to determine whether a linear model can capture behavior generated by a highly nonlinear model. I analyze the sensitivity of crowdout to alternative specifications of the model. Finally, the simulations are used to compute the value of pensions to households, as measured by the compensating variation: the change in private household wealth that equates the expected present discounted value of optimized lifetime utility with and without a pension.

Preliminary simulation results indicate that crowdout of household saving by Social Security and employer-provided pensions in the more realistic life cycle model is less than -$0.25 per dollar of pension wealth, much smaller than the value of -1 predicted by the simple model. Regression estimates of crowdout using the simulated data are close to this magnitude for private pensions, but are much larger for Social Security, indicating that empirical estimates of Social Security crowdout could be misleading. The magnitude of crowdout is quite sensitive to model specification. Welfare calculations indicate that, conditional on Social Security, employer-provided pensions are of considerable value to households, but are worth less than their wealth equivalents. In the absence of employer-provided pensions, Social Security is valued by households at more than three times its wealth equivalent value.

The following section of the paper provides a bit of background on pensions in the U.S., briefly characterizes the stylized life cycle model used as the basis for most empirical analysis, and summarizes the empirical literature. Section 3 describes the richer life cycle model and the solution and simulation approach. Section 4 describes the data, parameter calibrations, and initial conditions. The results are discussed in section 5, and conclusions in section 6.

2. Background, Stylized Life Cycle Model, and Previous Findings

A. Background

Defined Benefit (DB) pension plans provide employees with the promise of a monthly pension benefit determined by a formula that includes age, length of service at the firm, and pre-retirement earnings.[1] A typical DB plan has a normal retirement age, at which the maximum possible monthly benefit is available, and an early retirement provision that allows a worker to retire before the normal retirement age with a reduced monthly benefit. DB plans typically provide a strong incentive to remain with the firm until the early retirement age, and relatively little incentive to remain after this benchmark. Benefits are usually defined in terms of a lifetime (nominal) annuity, but some DB plans allow the worker to take a lump sum payment or a series of installment payments. Pension benefits are taxable income to the worker, but the firm's payments into the pension plan are tax deductible to the firm. The firm bears the risk of market fluctuations in financing pension payments, although default as a result of bankruptcy is not uncommon.

Defined Contribution plans specify the amount or percent of the worker's earnings contributed to the worker's pension account by the employer and the employee. The employer's contribution is usually a match of a fraction of the worker's contribution. The worker often has some choice over the percentage of earnings he contributes. The employer and employee contributions are made on a pre-tax basis. The assets in the account are allocated by the worker among the various investment options allowed by the plan. The interest, dividends, and capital gains or losses accrue to the account tax free. The funds in the account become available to the worker upon retirement from the firm.[2] All DC plans allow the retiree to take receipt as a lump

---

[1] See BLS (2003 and 2005) and Purcell (2007) for additional information.

[2] There is a tax penalty for withdrawing funds from the account before age 59½, and withdrawal of funds must begin by age 70½.

sum payment or to roll over the funds into a tax-sheltered Individual Retirement Account (IRA). Some DC plans also provide annuity options and installment payments.

Participation in pensions plans in the U.S. has remained relatively flat at about half the labor force in the past quarter century, but there has been a major shift away from DB plans toward DC plans in the private sector. Measured in terms of annual contributions by employers, DB plans accounted for 60% of total private sector employer contributions in 1980, and only 13% in 2000 (Poterba, Venti, and Wise, 2007).

Social Security is the public pension program in the U.S. The benefit is provided as a real lifetime annuity, determined by a formula that depends on age at claiming, lifetime average earnings, and birth cohort. It is financed by employer and employee payroll taxes.

B. A Stylized Life Cycle Model

Consider an employed individual of age t who will work until age R-1 > t, retire at age R, and remain retired until death at age T+1[3]. There is no uncertainty, no income tax, and no institutional constraints on pension claiming. The capital market is perfect, with a rate of return r. The individual will earn salary W_j, $j = t,..., R$-1, and at age R will begin receiving fixed real annuities of amount b from a DB pension and amount s from Social Security (SS). In addition, at age R he will receive a lump sum distribution D_R from a DC plan. The individual and employer each contribute a proportion c of his salary to his DC account in each period, and pays a proportional payroll tax p to the Social Security system. The individual chooses consumption C_j,

[3]The simulation model developed below is in discrete time, so for ease of comparison I present a discrete time version of the simple model. See Alessie, Kapteyn, and Klijn (1997) and Attanasio and Rohwedder (2003) for expositions of the discrete time framework, and Gale (1998) for an exposition of the continuous time framework. I follow Alessie et al. here.

$j = t,...,T$, to maximize the Present Discounted Value (PDV) of remaining lifetime utility given a rate of time preference δ. If preferences over consumption are intertemporally additive and homothetic, the consumption decision rule is

$$C_t = \kappa_t[(1+r)(A_{t-1} + D_{t-1} + 2cW_{t-1}) \;+\; \sum_{j=t}^{R-1} W_j(1 - c - p)(1+r)^{t-j} \;+\; \sum_{j=R}^{T}(b + s)(1+r)^{t-j}]$$

where A_{t-1} is private wealth held at the end of period t-1, asset returns are realized at the beginning of the period, and $\kappa_t \in (0, 1)$ is an adjustment factor for time until death that depends on r, δ, T, and parameters of the utility function (κ_T=1, and κ_t declines as the remaining lifetime increases). The term in square brackets is total remaining lifetime wealth, which consists of private wealth plus the DC balance at the beginning of period t (the first term, which includes the DC balance, since it is perfectly fungible), the present value of future net earnings (the second term), and the present value of future pension and Social Security benefits (the third term). Using the budget constraint

$$A_t + D_t = (1+r)(A_{t-1} + D_{t-1} + 2cW_{t-1}) + W_t(1 - c - p) - C_t$$

to solve for wealth in period t yields an equation for private wealth plus the DC balance held at the end of period t:

$$A_t + D_t = (1-\kappa_t)[(1+r)(A_{t-1} + D_{t-1} + 2cW_{t-1}) \;+\; W_t(1 - c - p)] \;-\; \kappa_t[E_t + B_t + S_t],$$

where E_t is the PDV of remaining lifetime earnings, B_t is the PDV of DB future pension benefits, and S_t is the PDV of future Social Security benefits, discounted to t.

The implication of this simple framework is that over the lifetime, private wealth accumulation is crowded out dollar for dollar by public and private pensions. κ_t can be computed under assumptions about r, δ, T, and the form of the utility function. The model thus implies a dynamic regression specification of the form[4]

$$A_t = X_{1t} + X_{2t} - X_{3t} - X_{4t} - X_{5t} - X_{6t}$$

where

$$X_{1t} = (1-\kappa_t)[(1+r)(A_{t-1} + D_{t-1} + 2cW_{t-1})], \qquad X_{2t} = (1-\kappa_t)W_t(1 - c - p),$$

$$X_{3t} = \kappa_t E_t, \qquad X_{4t} = \kappa_t B_t, \qquad X_{5t} = \kappa_t S_t, \qquad X_{6t} = D_t.$$

As specified, there is nothing to be estimated: the model implies that the regression coefficients are equal to 1 or -1. This is obviously not a very interesting specification, so the typical specification in the literature is

$$A_t = \beta_1 X_{1t} + \beta_2 X_{2t} + \beta_3 X_{3t} + \beta_4 X_{4t} + \beta_5 X_{5t} + \beta_6 X_{6t}.$$

A null hypothesis of interest is that β_4, β_5, and β_6 are equal to -1, and coefficient estimates less than 1 in absolute value are taken as evidence of incomplete crowdout. The virtue of this framework is its empirical tractability: under the assumptions of the model, the right hand side variables can be computed and the regression estimated. However, as noted above, a finding that the β's of interest are not equal to -1 may not be evidence of incomplete crowdout, but rather a misspecified model.

C. Previous Empirical Findings

Table 1 summarizes findings from several previous studies of the effects of pensions and Social Security on household saving. Alessie et al. (1997) estimate that the effect of occupational pensions in Holland is *positive* and the effect of Social Security is close to zero. They interpret these findings as evidence of unobserved heterogeneity and re-estimate the model using a

[4]The more common cross section regression model used in the literature is derived in a continuous time framework in Gale (1998). The implications are the same.

household fixed effect estimator. The positive pension effect persists, while the SS effect changes drastically to -2.10, significantly different from zero but not from -1.

Gale's (1998) median regression estimate using cross section data from the Survey of Consumer Finances and a pension measure that combines DB, DC, and SS is -.77, significantly different from zero but not from -1. Engelhardt and Kumar (2007) use data from the 1992 wave of the HRS to estimate a model like Gale's, using several estimation approaches. Their median regression estimate is roughly .15, significantly different from zero, implying *crowd in*, and their Instrumental Variable median regression estimate is roughly .10, not significantly different from zero.[5] They estimate regressions for other quantiles as well, and find coefficient estimates that are positive at all quantiles using ordinary quantile regression. Using IV-quantile regression, their estimates become negative at quantiles above .65. Gustman and Steinmeier (1999) used the HRS 1992 cross section to estimate a specification like the one used by Gale (1998) and Engelhardt and Kumar (2007). Their median regression estimate of pension crowdout is .012. Kapteyn et al. (2005) report a median regression estimate of -.115 for crowdout of private wealth by Social Security in Holland.

The studies by Attanasio and Rohwedder (2003) and Attanasio and Brugiavini (2003) listed in the lower panel of Table 1 use time series of cross section household expenditure survey data and estimate models of household saving. The coefficient on pension wealth in their model of the saving rate has the same interpretation as in models in which wealth is the dependent variable. They use public pension reforms in the U.K. and Italy as a source of exogenous

[5]Their IV approach uses pension wealth computed from the HRS pension provider data base for a set of "synthetic individuals" as an instrument for pension wealth calculated from respondent-reported information.

variation to identify the effects of public pensions on household saving. The estimates for the U.K. for the effect of the earnings-related public pension at ages 55-64 is -0.75, indicating substantial crowdout. For Italy, the estimates are -0.49 for ages 46-55, -0.21 for ages 56-60, and -0.11 for ages 61-65. Their estimates are more likely than most to be free of bias from unobserved heterogeneity, so the large magnitude of the estimates for some age groups suggests substantial crowdout by public pensions. However, despite a plausibly exogenous source of identification, their estimates are based on the assumptions of the simple life cycle model.

To summarize, estimates of crowdout vary widely, from substantial crowdout, to no crowdout, to substantial crowd in. Different data sources, time periods, and institutional features may explain some of this variation. Estimates may differ also due to varying treatment of unobserved heterogeneity, measurement error, and model specification. In the remainder of the paper I focus on specification of the underlying life cycle model.

3. A Richer Life Cycle Model

The model developed here characterizes the employment and saving decisions of an unmarried individual from the middle to the end of the life cycle, taking as given the individual's characteristics when first observed in middle age.[6] The individual makes a discrete employment choice j_t and a continuous consumption choice c_t in period t, where consumption is defined as net of out-of-pocket medical expenditure. The employment choice set is: (0) non-employment, (1) a job with a new employer, and (2) remaining on the job with the period t-1 employer. The latter

[6]A version of the model for married-couple households has been developed, and simulation results will be included in a future draft of the paper. Many of the non-pension features of the model are based on the model of Van der Klaauw and Wolpin (2008).

alternative is available only if the individual was employed in period t-1 and not laid off at the end of the period. A job offer from a new employer is assumed to be available in every period, but new jobs do not offer pension coverage. Allowing job switching is important because pensions are usually employer-specific, and in some cases it is feasible to leave the pension-employer, collect the pension benefit, and work for another employer.[7] The length of a period is one year. The last age at to which the individual can survive is denoted T, and the last age in which employment is an option is $T^* < T$. In the analysis, $T = 100$ and $T^* = 75$. The employment choice is eliminated after age 75 in order to speed up solution of the model.

If the individual is eligible for a retirement benefit from Social Security (Old Age and Survivors Insurance, or OASI) and has not yet claimed the benefit, then he makes an entitlement choice in period t. Thus, employment and Social Security claiming are distinct decisions. An individual who chooses to leave a firm in which he is covered by a DC pension, or has previously left the firm and has not yet claimed the balance in the pension account, makes a claiming decision. The options are to allow the DC account balance to continue to accumulate tax free, or to claim the balance in the current period, as a lump sum.[8] There is no separate claiming decision for DB pensions: age and job tenure at the time of exit fully determine both

---

[7]There is also an hours-of-work choice: full time or part time. However, this distinction is unimportant in the results presented here, so I do not discuss it.

[8]There is a tax penalty of 10% of the balance if the pension is claimed before age 59½ (60 in the model). The balance must be claimed by age 70, consistent with legal requirements. There is no installment payment option or annuity option in the model. The model does not incorporate Individual Retirement Accounts (IRAs), but the option to let the account balance continue to accumulate tax free after leaving the pension-employer is equivalent to rolling over the balance into a tax sheltered IRA. Also, the model does not allow purchase of annuities in the private market.

whether the individual will receive a benefit and the benefit start date. It is assumed that the DB benefit must be taken as a nominal life annuity.

The logarithm of the hourly wage offer from employer k ($k = 1$ for a new employer, $k = 2$ for the previous-period employer) in period t is $ln(w_{kt}) = x_t\beta_1 + \mu_k + \varepsilon_{kt}$, where x_t includes demographics and human capital characteristics, μ_k is an employer-specific wage component that is fixed for the duration of employment with an employer, and ε_{kt} is a first-order serially correlated shock. The logarithm of out-of-pocket medical expenditure is $ln(m_t) = x_t\beta_3 + \varepsilon_{3t}$, where ε_{3t} is a first-order serially correlated shock. Health status h_t is a categorical variable that can take on the values 0 (good), 1 (bad), and 2 (dead). Health status transitions are determined by a first order Markov process, with the probability that health status takes on the value h' in period t, conditional on health status h in period t-1 given by $\pi_{hh'}(x_t\beta_{hh'})$. The $\pi_{hh'}$ functions are assumed to be multinomial logits.[9] An employed individual is subject to the risk of being laid off from his job. The probability of a layoff at the end of period is specified as a logit function: $\lambda_t = exp(x_t\beta_4)/(1 + exp(x_t\beta_4))$.

An individual is eligible to claim his OASI benefit if he has accumulated a sufficient number of quarters of coverage (40) and has reached the Social Security early retirement age (62).[10] The OASI benefit, s_t, upon initial entitlement is determined by the function $s_t = s(AIME_t$,

[9]Health status affects medical expenditure, wage offers, and mortality risk (i.e., it is an element of x_t). Health could also affect preferences, but for simplicity this is not allowed here.

[10]An individual who is in bad health, chooses non-employment in a given period, and has accumulated a sufficient number of quarters of coverage in the Social Security system can apply for a benefit from the Social Security Disability Insurance (SSDI) program before becoming eligible for retirement benefits. A rough approximation of the complex SSDI approval process is included in the model, but it plays little role in the results discussed below, so description is omitted for brevity.

fe, E_t, a_t), where $AIME_t$ is Average Indexed Monthly Earnings, $E_t = H_t w_t$ is annual earnings in period t, H_t is annual hours of work (0, 1000, 2000), fe is the age at which the individual claims the benefit, and a_t is age in period t. Age and current earnings matter because there is an earnings test at some ages.

The benefit in a DB pension plan, b_t, is determined by a formula that depends on age, job tenure, and earnings history as of the date of exit from the pension-providing firm. The formula can be written in general as $b_t = b(E_p, a_e, s_e, a_b, t)$, where E_p is a summary statistic for the worker's earnings at the pension-providing firm, a_e and s_e are age and years of enrollment in the plan at the date of exit from the firm, and a_b is the age at which the benefit begins (which could be greater than the age at which the individual leaves the firm). Both the formula and the specific form of the earnings summary statistic E_p depend on the particular pension plan. The DB benefit is nominal, so its real value at age t, b_t, depends on the inflation rate and years since the benefit began. There is no risk of default by the pension plan.

DC pension plans are characterized by the account balance and the employee and employer contribution rates. While employed at the pension-providing firm, the individual and the firm contribute specified fractions of the individual's earnings to the pension account. These fractions are taken as given and fixed. If the individual remains with the pension-providing firm at the beginning of period t+1, the account balance is given by $B_{t+1} = (B_t + E_t(wc + ec))(1+r_{t+1})$, where wc and ec are the worker and employer contribution rates, respectively, and r_{t+1} is the rate of return on assets. If the individual has left the firm but has not yet claimed the account balance, then $B_{t+1} = B_t(1+r_{t+1})$. Borrowing from a DC pension account is not allowed.[11]

The rate of return earned on assets held at the end of period t, r_{t+1}, is realized at the beginning of period t+1. The rate of return is determined by a mean-reverting stochastic process specified as $1+ r_{t+1} = (1+\bar{r})exp\{\theta_{t+1}\}$, where $\bar{r}$ is the mean rate of return, and θ_t is an idiosyncratic individual-specific shock. Returns are defined to include capital gains, so r_{t+1}<1 corresponds to a capital loss. The rate of return is assumed to be the same for the DC pension account and the individual's other assets

The law of motion for assets held outside the DC account is $A_{t+1} = A_t^*(1+r_{t+1})$, where A_t^* is the stock of assets held at the end of t, and A_{t+1} is the stock of assets at the beginning of t+1. There is assumed to be a borrowing constraint, so that assets cannot be negative, and a consumption floor, $\bar{C} > 0$. The consumption floor is a very crude but simple way to account for income-and-asset-tested government programs such as Supplemental Security Income (SSI), Food Stamps, and Medicaid that allow individuals with no other sources of income to survive (Hubbard, Skinner, and Zeldes, 1995). If cash on hand is less than $\bar{C}$, the government provides the individual with consumption of $\bar{C}$ and confiscates all cash on hand.

Non-asset income net of out-of-pocket medical expenditure and taxes is

---

[11]Previous simulation studies of pension effects on savings have accounted for some of the features modeled here, but in a more limited way. Scholz, Seshadri, and Khitatrakun (2006) specify a DB plan in which the benefit is approximated as a function of job tenure and earnings in the last period before retirement, with the approximation function estimated with HRS data. Engen, Gale, and Uccello (1999) use the HRS to estimate the mean replacement rate for a DB pension, based on final earnings and education. Laibson, Repetto, and Tobacman (1998) specify a DC plan with fixed employee and employer contribution rates and a tax penalty for early withdrawal. These papers do not allow for choice of retirement age, but they do incorporate earnings and/or medical expenditure uncertainty and in some specifications a liquidity constraint.

$$I_t = E_t + b_t + s_t - m_t - \tau_t,$$

where τ_t includes federal income and payroll taxes, calculated using the rules in effect for 1992, and assuming the individual takes the standard deduction. The tax computation accounts for the tax-sheltered nature of the worker's contribution to the DC account and the rules governing taxation of Social Security benefits. Cash on hand at the beginning of period *t* net of out-of-pocket medical expenditure and taxes is $A_t + I_t$, and assets carried forward to the next period, before the return is realized, are $A_t^* = A_t + I_t - c_t$, (unless the liquidity constraint is binding), where c_t is non-medical consumption expenditure.

Utility is a function of consumption and employment. The functional form assumed here is isoelastic in consumption, separable in consumption and employment, and dynamic in employment:

$$u_t = [c^{1-\alpha}/(1-\alpha)]\exp\{\varepsilon_{ct}\} + (\gamma_1 + \gamma_2\, a_t)W_t + \gamma_3(1-W_{t-1})W_t + \gamma_4 W_{t-1} NJ_t + H_t \varepsilon_{\ell t}$$

where $W = 1$ if employed, zero otherwise; $NJ = 1$ if a new job is chosen, and zero otherwise; ε_{ct} is an *iid* shock to the marginal utility of consumption; and $\varepsilon_{\ell t}$ is an *iid* shock to the marginal utility of leisure. $\gamma_1 + \gamma_2 a_t$ is the disutility from working at age a_t, γ_3 is the additional disutility from working if the individual was not employed in the previous period, and γ_4 is the additional disutility of changing jobs. The dynamic features of the utility function are important in order to avoid excessive churning in employment choices in response to transitory shocks.[12] There is no bequest motive.

[12]In principle, the low rate of re-entry to employment and job switching at older ages could be captured in the model by depreciation of general human capital while not employed and loss of firm-specific human capital upon separation from a firm. However, reasonable estimates of the rate of depreciation of general human capital are much too small to explain the low re-entry rate. The model does allow tenure to affect wages, but empirically this effect is quite small.

The individual's goal is to choose employment and consumption (and OASI and DC claiming, if relevant), to maximize the expected present discounted value of remaining lifetime utility, with discount factor δ, subject to the constraints described above.

The model is solved by backward recursion on the value function, using the regression approximation and Monte Carlo integration approach developed by Keane and Wolpin (1994) and applied by van der Klaauw and Wolpin (2008) in a context similar to the present model. For each of a randomly selected set of values of the period T state variables (assets, benefits, health, and T-1 medical expenditure), the expected value at the end of period $T-1$ of the maximum, over the available choices (in period T consumption is the only decision variable) in T, of the period T value function, denoted *Emax*, is computed. The expectation is computed by Monte Carlo integration over the continuous period T disturbances (ε's, θ), which are assumed to be independently normally distributed with mean zero. A polynomial spline regression of the *Emax* on the state variables is estimated on the sample of state points for which the *Emax* was computed, and the regression coefficients are stored. A set of points in the period T-1 state space is randomly selected, and the *Emax* of the period T-1 value function is computed for the selected state points, using the period T regression coefficients to approximate the expected value of the period T value function for all feasible choices and realizations in T-1. Another regression function is estimated for period T-1, with the coefficients stored. The process is repeated recursively to period 1. A similar regression approximation approach is used to compute the EPDV of DB and SS benefits, accounting for uncertainty, the employment and claiming choices, and optimizing behavior. The regression approximation is necessary because if the benefit has not yet been claimed, its EPDV depends on future choice probabilities and realizations of shocks.

4. Data, Calibration, and Simulation

Data from the Health and Retirement Study (HRS) are used to calibrate the parameters of the model and specify the initial conditions. The HRS also provides a sample of DB pension formulas, which are used in solution and simulation. The HRS is a biennial longitudinal survey of a sample of U.S. households with individuals aged over 50. The survey began in 1992 with a sample of individuals born from 1931 to 1941, and their spouses. Additional cohorts have been added over time, but I focus on the original cohort, for which the most extensive data are available. The original cohort sample contains 12,652 individuals residing in 7,067 households. Various subsamples are used to estimate the parameters of the wage, medical expenditure, health, layoff, and SSDI acceptance functions.

A. Pensions. The HRS asks respondents a substantial battery of questions about DB pensions, including the ages of early and normal retirement, expected benefits if the respondent left the firm at the early and normal retirement ages, and the respondent's expected age of retirement and expected benefit at that age. Some studies have used this information to construct a measure of DB pension wealth (Engelhardt and Kumar, in press; Chan and Stevens, 2008), but the information is not sufficient to calculate benefits for all possible retirement ages and future earnings realizations. Respondents are also asked for permission to allow the HRS to contact the employer directly in order to request a copy of the written Summary Plan Description and other relevant information about the pension. The benefit formulas and other plan features derived from these documents were coded by HRS staff and made available to researchers (under restricted-access conditions) in a data base, along with pension calculator software. This allows researchers to compute the benefit to which an individual would be entitled under any feasible scenario for age and tenure at exit and the salary profile. However, solution and simulation of the model requires many thousands of benefit computations, and it was not feasible to use the pension data base and software directly. I use an accurate plan-specific regression approximation of the benefit formulas, described in the Appendix, in solution and simulation.

B. Estimation and Calibration. The model was calibrated in three steps. First, the HRS data were used to estimate the parameters of the wage, medical expenditure, health, layoff, and SSDI acceptance functions.[13] The health transition models were specified as multinomial logits, conditional on the two-year lag of health status. There is no straightforward way to convert the coefficients of these two-year transition models to one-year transition coefficients, so the two-year transition parameter estimates are used in the analysis. The estimates are shown in the appendix, and the key parameters derived from the estimates are shown in the top panel of Table 2.

Second, several other important parameters were set arbitrarily. These include the coefficient of relative risk aversion (α=2.5), the consumption floor ($\overline{C}$=3), the mean rate of return ($\overline{r}$=.02), and the rate of time preference (δ=.03). The second panel of Table 2 shows these

---

[13]The log wage and log out-of-pocket medical expenditure functions are specified as AR(1) models, but the HRS data are biennial. These models were estimated by linear regression with a two-year lag of the dependent variable, and the square root of the coefficient on the lagged dependent variable was used to estimate the one-year serial correlation coefficient. The firm-specific component of the wage disturbance was not accounted for in the regression specification, so its variance was chosen arbitrarily. The estimated mean squared error from these regressions could be used as estimates of the variances of the wage and expenditure shocks. However, the estimated mean squared errors in both models implied excessively large variances. These variances were therefore chosen arbitrarily.

and the values of other arbitrarily chosen parameters.

Finally, the remaining parameters were calibrated so as to generate "reasonable" employment patterns. This was an informal trial-and-error calibration process without a specific criterion for "reasonable." The disutility of employment parameters (γ's) and the variances of the leisure and consumption preference shocks in the utility function and the wage and medical expenditure shocks were chosen in this manner.

C. Simulation. The model is simulated in order to examine the effects of pensions on wealth accumulation. The simulations compare three pension scenarios: one with no pension, a second with a DB pension, and a third with a DC pension. In all three cases the household is eligible for Social Security. A fourth simulation assumes no Social Security as well as no pensions. Rather than use an arbitrarily chosen DB pension, the pension simulations are computed for all 834 DB pension plans in the HRS data, and the results are averaged over the pensions. A similar approach is used for DC pensions. DC plans are characterized by the account balance as of the first period in the simulation, and the employer and employee contribution rates. The employer and employee contribution rates are plan-and-individual-specific and are reported by the HRS respondents. The simulations are computed for all 1,410 observed DC plans and averaged over the plans.[14] 600 simulated life cycles from age 51 to 100 are computed for the no-pension and no-Social-Security scenarios. Six simulated life cycles per pension plan were computed for the DB and DC scenarios. The results are averaged across simulations for each

[14]Another relevant feature of a DC plan is the set of investment options available in the plan. As described above, I assume that assets in the DC account balance earn the same mean return as assets held directly by the household, so this feature of DC plans is not considered here. Asset allocation decisions of individuals are also not considered here.

scenario.

D. Initial Conditions. The simulations are computed for an unmarried white man with a high school education, born in 1941, and age 51 in the first period of the simulation. The individual is assumed to be in good health and employed full time at the beginning of the first period. In the pension simulations, the individual is assumed to have enrolled in the plan at age 30, and in the DB case a salary history at the pension-providing firm was generated by backcasting from the observed mean 1991 salary, assuming a one percent real growth rate.[15] Other initial conditions were generated from the subsample of single white men in the HRS with a high school education who were employed at ages 51-53 with earnings of at least $10K in 1991. Mean values of work experience, previous-year earnings, and AIME from this sample were used.[16]

The household is endowed with initial assets equal to the median value of assets in the single male sample at ages 51-53: $41K (all monetary amounts are expressed in thousands of 1992 dollars). In the DC pension scenario, the household is endowed with a DC balance of $17K, which is the median value observed in the data. Table 3 summarizes the initial conditions.

---

[15]The entire salary history at the pension-providing form is needed because each pension plan uses a different function of the salary history to compute benefits. Some plans use the average of the most recent five years, others use the average over all years of employment at the firm, and there are many other variations. The HRS contains a limited salary history that includes the starting salary on the job held in the first wave, so this could be used to generate a salary history instead of assuming an arbitrary growth rate. Administrative Social Security Earnings Records and W2 records are available for a subsample of respondents, and these could also be used to generate a salary history at the pension firm. These options will be explored in future work.

[16]AIME is the lifetime average earnings measure that is the basis for determining the Social Security benefit. It is available from administrative Social Security Earnings Records of respondents who gave permission for the HRS to have access to these records.

5. Results

A. Basic Patterns

Figures 1-4 show average simulated life cycle patterns of employment, consumption, and assets for the four scenarios of interest: no pension (NP), DB pension, and DC pension, all with Social Security, and no-pension-no-Social Security (NPNSS). Employment declines rapidly beginning at age 55 in all three scenarios with Social Security, especially in the DB case. In the NPNSS case, employment remains very high until close to the assumed mandatory retirement age of 76. The employment rate in the DB scenario falls more rapidly than in the other scenarios at ages 56-60; these are common early retirement ages in DB plans. The employment rate reaches zero by age 64 in all three scenarios with Social Security, which is counterfactual, but the general pattern of decline is consistent with the data. There is a noticeable drop in the employment rate at age 62, the earliest age at which Social Security retirement benefits are available.

Consumption declines with age, because the effective rate of time preference (including both pure time preference and mortality risk) exceeds the mean interest rate. Consumption is similar in the four scenarios until the early 60s, at which point consumption in the NP and DC scenarios drops below consumption in the DB and NPNSS scenarios.

The wealth trajectory follows the standard life cycle pattern, rising until retirement and then falling, with two exceptions. First, the tax advantage of DC pensions provides an inducement to let the DC balance accumulate as long as possible, and in about 40% of the simulated DC life cycles the balance is claimed and transferred into household assets at the latest allowable age, 70. This results in a jump in DC assets at 70. The second exception is that the DB asset trajectory levels off and even rises a bit before eventually declining again in the mid 80s. This pattern can be seen more clearly in Figure 4 (note the different scales of the vertical axes in Figures 3 and 4). By age 80, DB assets are higher than DC and NP assets, after having been much lower in the 50s and mid 60s. This seemingly peculiar pattern is robust to many alternative parameter values and specifications, with few exceptions. It has little influence on the crowd out results discussed below, which are computed only for ages at which the pension or SS benefit has not yet been claimed.

Comparing the different scenarios, DB pensions increase the demand for leisure and consumption at most ages compared to the baseline NP scenario. These are the total effects of DB pensions, including wealth and substitution effects. The effects of DC plans on consumption and employment relative to the NP scenario are much smaller, due in part to the low initial balance (17K). The effect of Social Security on demand for leisure is very large (compare the NPNSS and NP scenarios), but optimal consumption is actually lower when Social Security is available. This suggests that Social Security may have significant substitution effects, toward leisure and away from consumption, as implied by many empirical studies.

B. Crowdout

Figures 5-8 illustrate the crowdout patterns implied by the simulations. Figures 5 and 6 display life cycle crowdout patterns, measured by the difference in mean household wealth between the DB and DC scenarios and the NP scenario. Figure 7 shows the difference in mean household wealth between the NP and NPNSS scenarios. Note that all households begin with the same initial wealth (41), so crowdout is zero by construction at the beginning of age 51. The figures also show the relevant pension wealth measure: the EPDV of future DB benefits, the current DC balance, and the EPDV of future SS benefits.[17] Figure 8 reports proportional crowdout for each case, measured by the dollar magnitude of crowdout divided by the relevant pension wealth measure. Figure 5 shows that assets in the DB case decline gradually relative to the NP case, reaching about -40 at age 60. DB pension wealth rises from 75 at age 51 to a peak of 145 at age 57. Figure 8 shows that the implied magnitude of crowdout as a fraction of pension wealth starts at zero (by construction) and reaches -0.35 in the early 60s. Crowdout by DC pensions is quite small, never exceeding -10, or -0.16 as a fraction of the DC balance. Crowdout by Social Security is about zero at ages 51-59, and then grows substantially to -85 at age 65, or -0.65 as a fraction of SS wealth.

It is useful to summarize the magnitude of crowdout with a single number, as in most of the empirical literature. Initial assets are the same in all scenarios, so it is advisable to allow behavior to play out for at least a few years before measuring asset differences. However, it is useful to measure crowdout at an age at which pensions and Social Security have not yet been claimed, both because this is the typical approach in the literature, and because crowdout behavior could be quite different in the asset decumulation phase. I arbitrarily chose to use the latest age at which no more than 25% of simulated cases have claimed the pension or Social Security. For the no-pension case I use the criterion that no more than 25% of simulated case have left the initial job held at age 51.[18] The first column of Table 4 shows that mean DB

[17]The DB and SS pension wealth measures are computed under the assumptions of the model, allowing for uncertainty, the liquidity constraint, and optimal claiming behavior. Before the pension is claimed, pension wealth is computed by a backward-recursion regression approximation method similar to the one used to approximate the Emax.

[18]In the DB case, claiming means leaving the pension job, even if the benefit does not begin until later. When making comparisons across cases, for example comparing DB to NP, the

crowdout is -20, DC -8, and SS -27. As a proportion of the relevant pension wealth measure, crowdout is -0.16 for DB, -0.18 for DC, and -0.24 for SS.

These crowdout measures are relatively small, but they are within the range of estimates from the empirical literature reported in Table 1. They are closer to zero than to -1, the value implied by the simplest life cycle model. Qualitatively, the DB and SS crowdout magnitudes are plausible, since DB pensions and SS are illiquid. DC account balances are fairly liquid, so small crowdout is unexpected.[19] There are many arbitrary features of the model and parameters, so these results are merely illustrative. I examine the sensitivity of the results to initial conditions and model features, but much further sensitivity analysis remains to be done.

The initial assets of 41 and DC balance of 17 are based on HRS data for single men, so they are not entirely arbitrary. But the model is highly nonlinear, so it is useful to examine whether crowdout is sensitive to initial conditions. Intuition suggests that crowdout will be larger when initial wealth is higher. Using initial assets of 150 and an initial DC balance of 60 yields the results shown in the second column of Table 4: crowdout of -17 by DB, -22 by DC, and -64 by SS. The proportional crowdout measures of -0.15, -0.26, and -0.58 are substantially larger than the corresponding figures in column 1 only in the case of SS. This provides only weak evidence that crowdout may be larger further away from the zero asset bound.

---

criterion has to be satisfied in both cases.

[19]Claiming a DC balance before age 60 results in a tax penalty of 10% of the balance, and the DC balance is claimed before age 60 in only 25% of the simulated cases. This is sensible, since the tax advantage of a DC plan should encourage households to use privately held assets to smooth consumption, and resort to the DC balance only after private assets are exhausted. The model assumes an all-or-nothing claiming decision, but in reality households can borrow against the DC balance and/or claim it in installments. Accounting for these possibilities would be expected to increase the liquidity of DC balances.

One might expect larger crowdout if there was no liquidity constraint. Intuition from the simple life cycle model suggests that the ability to borrow against future pension benefits would reduce the need to save in order to finance consumption until the benefit can be claimed. To study this question, I removed the liquidity constraint and re-solved the model. Two issues complicate the interpretation, however. In this version of the model agents are allowed to be in debt at the time of death, but assets must be nonnegative in the last period in which an individual can be alive, 100. Allowing arbitrarily large debt in the first period results in a consumption splurge, which is an artifact of initiating the model at age 51. In order to avoid this sort of sharp and unrealistic discontinuity, I impose a lower bound on wealth of -50 at age 50, rising to zero at age 100.[20] The second issue is the consumption floor, which is equal to 3 in the simulations. This feature of the model is important in order to avoid difficulties in solution that can arise when it is possible for consumption to approach arbitrarily close to zero.[21] The existence of a consumption safety net dampens the impact of relaxing the liquidity constraint, since going into debt will never cause starvation. The third column of Table 4 shows results from the simulation without a liquidity constraint: crowdout of -0.25 for DB, -0.15 for DC, and -0.35 for SS. These figures are fairly close to the results with a liquidity constraint in column 1, suggesting that a liquidity constraint does not affect crowdout much in the presence of uncertainty and a consumption floor.

Column 4 of Table 4 shows results for a model with no liquidity constraint and no

[20]This is similar to the approach used by Van der Klaauw and Wolpin (2008). They estimated the lower bound in the initial period rather than imposing it arbitrarily as I do here.

[21]Risk averse consumers will avoid choosing very low consumption, but the fact that a future catastrophic outcome could result in very low consumption affects the solution of the model, causing it to become unstable.

uncertainty, except over length of life. Crowdout in this simulation is -0.37 for DB, 0.01 for DC, and -0.44 for SS. These figures are surprisingly small, especially in the DC case. In principle, the only remaining feature of the model that could result in less than 100% crowdout is the employment decision.[22] Eliminating the employment choice makes the model equivalent to the simple life cycle model. Columns 5-7 of Table 4 show results from specifications in which individuals must work through age 63 and retire at age 64. These specifications are comparable to rows 1, 3, and 4 except for eliminating the employment choice. The results indicate that crowdout is not systematically larger in the absence of an employment choice, and in some cases there is less crowdout. These results are difficult to rationalize, and suggest that the results presented here could be suspect, so please treat them as preliminary.

C. Compensating Variation

Crowdout of household wealth is a measure of behavior: how does household saving respond to pension wealth? Another question of interest is how much value households derive from pensions. The value of a pension to a household can be measured by the compensating variation (CV): the amount by which the initial assets of the household with a pension must be reduced in order to equate its EPDV of optimized lifetime utility (the value function) to that of a similar household without a pension. Dividing the CV by initial pension wealth provides a measure with the same scale as the proportional crowdout measure, but it measures the effect of pensions on well being rather than on saving behavior. Table 5 reports CV figures for the simulations on which the results reported in columns 1 and 2 of Table 4 are based. For the

[22]The model incorporates institutional restrictions on pension and Social Security claiming, but the ability to borrow against future benefits should allow households to undo the effects of these restrictions if they bind.

baseline case with initial assets of 41, providing a household with a DB pension requires initial wealth to be reduced from 41 to -12 in order to equate lifetime utility to the NP case. The CV is 41 - (-12) = 53. The EPDV of lifetime DB benefits at age 51 is 83. Thus the value of the pension to the household is 53/83 = 64% of pension wealth. In the DC case the CV is 105% of the DC balance at the end of age 51of 20, indicating that DC pensions are valued by households at roughly their dollar value. This may be due to their relatively high liquidity. For SS, the CV equates the EPDV of lifetime utility of the NPNSS case to that of the NP case. This requires increasing initial wealth of the NPNSS case to 360, yielding a CV of 360 - 41 = 319. This is 370% of the EPDV of SS wealth of 86 at age 51. Social Security is evidently of great value to a household that would otherwise have to fully finance its consumption expenditure in retirement through saving. This result may be due in part to the assumption that there is no market for annuities. It is probably also closely related to the findings that Social Security causes a very large reduction in labor supply, and little crowdout. One might expect that pensions would have less value to wealthier households, but the results in the lower panel of Table 5 using initial wealth of 150 do not support this conjecture.

D. Regression Estimates of Crowdout

The final issue of interest is whether regression estimates of crowdout based on the strong assumptions of the very simple life cycle model are accurate when the data are generated from a process that does not obey these strong assumptions. To investigate this issue, I estimate regressions on the simulated data using two different methods of calculating pension wealth. The first measure is the “correct” one used above in the crowdout and compensating variation calculations, computed under the assumptions of the model from which the simulated data were generated. However, the model is highly nonlinear while the regression approach imposes linearity. Thus, even with this “correct” measure of pension wealth there is no guarantee that the regression approach will accurately measure crowdout. The second measure of pension wealth assumes fixed retirement and claiming ages (I use the observed retirement and claiming ages from the simulations), no uncertainty (except over the length of life), no liquidity constraint, and no institutional restrictions on pension claiming. This measure is comparable to measures of DB and SS pension wealth used in many empirical studies. In the DC case, the pension wealth measure is just the account balance, so no assumptions are required to compute it.

The simulation approach used here is somewhat different from the approach used to generate the results in Tables 4 and 5. Those results are based on simulations of four cases (NP, DB, DC, and NPNSS), but a sample size of four is too small for the regression analysis. Instead I simulate several different scenarios for each of the four cases, varying initial assets, the initial DC balance, the salary history, AIME, and the length of service in the pension job. This resulted in 107 simulated scenarios. I choose the age at which to measure pension wealth and assets using the same criterion as in the crowdout analysis: the latest age at which no more than 25% of simulated cases have claimed the pension or Social Security.

Table 6 reports coefficient estimates on pension and Social Security wealth in household wealth regressions. I computed the life cycle adjustment factor (κ_t; see page 7) using a continuous time approximation from Gale (1998). The other regressors are those implied by the theory: the EPDV of remaining lifetime earnings[23], current period earnings, and household

---

[23]The EPDV of remaining lifetime earnings was calculated using the “incorrect” approach: it is based on realized earnings, under the assumption of perfect foresight and a known date of retirement. It is “expected” only in the sense of adjusting for mortality risk.

wealth in the previous period. This dynamic specification has the advantage of removing to some extent the influence of the initial conditions; the estimates represent the impact of an additional dollar of pension wealth on household wealth, holding constant wealth in the previous period. The estimates in column 1using the correct pension wealth measures indicate crowdout of -.18 and -.19 by DB and DC pensions, almost exactly equal to the "true" crowdout measures reported in column 1 of Table 4. The estimated SS crowdout is -1.01, far larger than the "true" crowdout of -0.24. The estimates are not very sensitive to the calculation of pension wealth: using the incorrect pension wealth measures, the results in column 2 are similar to those in column 1.

For comparison with the more common cross section regression models used in the literature, columns 3 and 4 present estimates from a specification that does not control for lagged wealth. This specification yields systematically larger crowdout estimates. Interpreting these estimates in terms of the simple life cycle model described in Section 2 would require solving out for lagged wealth.

6. Conclusions and Work in Progress

The impact of public and private pensions on household behavior and welfare is of considerable importance in a rapidly aging world. Empirical studies of crowdout in recent years have focused on issues of identification and specification, both of which are clearly important. Measurement of pension wealth and interpretation of crowdout estimates have both relied heavily on the simplifications rationalized by the stylized life cycle model. Here I take a different and complementary approach, solving and simulating a richer model in order to study crowdout in a more realistic environment. The results reported here are preliminary, and much further work remains to be done before any firm conclusions can be drawn about usefulness of this approach and its implications for the accuracy of empirical estimates of crowdout. In work in progress I focus on identifying the features of the model that lead to less than 100% crowdout, sensitivity of the results to alternative parameter values, and comparability of interpretation between the crowdout measures calculated directly from the model and regression estimates. The ultimate goal is to estimate the key parameters structurally, removing much of the arbitrariness in the calibrations.

Appendix A: Calculating Defined Benefit Pension Benefits

HRS respondents who reported any pension coverage at wave 1, either on the current job or a previous job, were asked for permission to contact the employer to obtain information on the pension plan. For respondents who gave permission and whose employers provided the requested information, the formulas that determine the pension benefit for each plan were coded by the HRS staff and provided to researchers on a restricted access basis, along with pension calculation software. These formulas determine the pension benefit for all possible scenarios involving birth date, age, years in the plan at the time of exit, and salary history. Rather than use the pension calculation software (which is coded in Visual Basic) to directly compute benefits for each individual, I used an approximation approach. This was done so that the benefit calculations could be easily computed in the Fortran program used to solve and simulate the model.

The approximation approach uses the pension calculation software to compute benefits for each DB and combination DB/DC plan in which any respondent is enrolled at wave 1 (834 plans, some covering more than one HRS respondent), for about 5,000 artificial individuals, with alternative combinations of birth date, hire date, and real salary level and growth rate. For each artificial individual and each plan, I computed the monthly pension benefit and the age at which the individual is first eligible for the benefit for every possible age at which the individual could quit from the year after the hire date through age 75.

I then ran three regressions, separately for each pension plan, using the 5,000 observations for each plan. The dependent variable in the first regression is a binary indicator for whether the individual will ever be eligible for a benefit, given the age at exit. The dependent variable in the second regression is the age at which the individual is first eligible for the benefit, conditional on ever being eligible. The dependent variable in the third regression is the monthly benefit, conditional on eligibility before age 75. Each regression is specified with a very flexible functional form, with dummies for age at exit, tenure at exit, and combinations of quit-age and tenure. For the benefit regression, the specification includes average salary in the most recent five years, the second most recent five years, and so forth, and interactions of the salary averages with age and tenure dummies. The coefficient estimates from these three regressions for each plan are stored, and used to predict benefits in the solution and simulation of the model. These regressions are generally very accurate in predicting outcomes. I compared the predictions from the regressions to the "true value" computed directly from the pension calculator. For the "ever eligible" regression, using the rule that the prediction is zero if the fitted value is less than 0.5 and the prediction is one otherwise, the regression predicts every one of the approximately 5,000 observations correctly for 78% of the plans, and never predicts more than 13% incorrectly for any plan. Two thirds of the first-age-of-eligibility regressions predict the correct age exactly for every observation, and the 5^{th} and 95^{th} percentiles of the rounded residual distribution are 1 and -1 respectively. Finally, for the annual benefit regressions, the mean prediction error is -2.7 (in thousands of dollars per year), the median error is -0.6, the 75^{th} percentile of the prediction error is 0.6, and the 25^{th} percentile is -8.5. Comparing the benefits predicted from this approach with the actual benefit reported by HRS respondents who retired during the panel, given actual quit dates, yields a mean prediction error of 3.0 and a median of 2.7.

References

Alessie, Rob J.M., Arie Kapteyn, and Frank Klijn. 1997. "Mandatory Pensions and Personal Saving in the Netherlands," *De Economist* 145 (3): 291-324.

Attanasio, Orazio P., and Agar Brugiavini. 2003. "Social Security and household saving," *Quarterly Journal of Economics* 188 (3): 1075-1120.

Attanasio, Orazio P., and Susann Rohwedder. 2003. "Pension wealth and household saving: evidence from pension reforms in the United Kingdom," *American Economic Review* 93 (5): 1499-1521.

Benartzi, Shlomo, and Richard H. Thaler. 2007. "Heuristics and biases in retirement savings behavior," *Journal of Economic Perspectives* 21 (3): 81-104.

BLS. 2003. "National Compensation Survey: Employee Benefits in Private Industry in the US, 2000," BLS Bulletin 2555, January.

BLS. 2005. "National Compensation Survey: Employee Benefits in Private Industry in the US, 2002-3," BLS Bulletin 2573, January.

Chan, Sewin, and Ann Huff Stevens. 2008. "What you don't know can't help you: pension knowledge and retirement decision-making," *Review of Economics and Statistics*, May.

Crawford, Vincent P., and David M. Lilien. 1981. "Social Security and the retirement decision," *Quarterly Journal of Economics* 96 (3): 505-29.

Engelhardt, Gary V., and Anil Kumar. 2007. "Pensions and household wealth accumulation," Working paper, Syracuse University, March. In press, *Journal of Public Economics*.

Engen, Eric M., William G. Gale, and Cri E. Uccello. 1999. "The Adequacy of Household Saving, *Brookings Papers on Economics Activity*, 2: 65-165.

Feldstein, Martin, 1974. "Social Security, Induced Retirement, and Aggregate Capital Accumulation," *Journal of Political Economy* 82 (5): 905-926.

Gale, William G. 1998. "The Effects of Pensions on Household Wealth: A Reevaluation of Theory and Evidence," *Journal of Political Economy* 106 (4): 706-723.

Gustman, Alan L., and Thomas L. Steinmeier. 1999. "Effects of pensions on savings: Analysis with data from the Health and Retirement Study," *Carnegie-Rochester Series on Public Policy* 50: 271-324.

Hubbard, R. Glenn. 1986. "Pension Wealth and Individual Saving," *Journal of Money, Credit,* *and Banking*, 18 (2): 167-178.

Hubbard, R. Glenn, Jonathan Skinner, and Steven P. Zeldes. 1995. "Precautionary Saving and Social Insurance, *Journal of Political Economy* 103 (2): 360-399.

Kapteyn, Arie, Rob Alessie, and Anamaria Lusardi. 2005. "Explaining the wealth holdings of different cohorts: productivity growth and Social Security," *European Economic Review* 49: 1361-91.

Keane, Michael, and Kenneth Wolpin K. 1994. The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence. *Review of Economics and Statistics* 77: 648-672.

Laibson, David I., Andrea Repetto, and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement," *Brookings Papers on Economics Activity*, 1, 91-172.

Munnell, Alicia H. 1974. The Effect of Social Security on Personal Savings Cambridge MA: Ballinger Press.

Poterba, James, Steven Venti, and David A.Wise. 2007. "The changing landscape of pensions in the United States," NBER Working Paper 13381, September.

Purcell, Patrick. 2007. "Pension Sponsorship and Participation: Summary of Recent Trends" Congressional Research Service Report RL30122, September.

Rust, John. 1998. ""Effects of pensions on savings: Analysis with data from the Health and Retirement Study: A Comment," *Carnegie-Rochester Series on Public Policy* 50: 325-335.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. "Are Americans Saving Optimally for Retirement?" *Journal of Political Economy* 114 (4): 607-643.

Van der Klaauw, Wilbert, and Kenneth I. Wolpin. 2008. "Social Security and the Retirement and saving behavior of low income households," *Journal of Econometrics*.

Table 1: Selected Estimates of the Effects of Pension Wealth on Household Net Worth

Study	Estimation approach	Data type and country	Coefficient on Pension wealth	Coefficient on Social Security wealth	Coefficient on Pension plus SS wealth
Dependent Variable: Household Net Worth					
Alessie et al. (1997)	OLS	Longitudinal, Holland	0.28 (.06)	.003 (0.3)	
	Fixed effects		0.16 (.11)	-2.10 (0.81)	
Gale (1998)	Median regr.	cross section, U.S.			-.77 (.24)
Engelhardt and Kumar (2007)	OLS[a]	cross section, U.S.			.27 (.17)
	IV[a]				-.50 (.33)
	Median regr.				~.15*
	IV-median				~.10*
Gustman and Steinmeier (1999)	Median regr.	cross section, U.S.			.012 (.003)
Hubbard (1986)	OLS	cross section, U.S.	-.16**	-.33**	
Kapteyn et al. (2005)	Median regr.	Panel, Holland		-.115 (.020)	
Dependent Variable: Household Saving Rate					
Attanasio and Brugiavini (2003), Table 6 column 4, ages 46-55	IV	time series of cross sections, Italy		-.49 (.11)	
ages 56-60				-.21 (.14)	
ages 61-65				-.11 (.04)	
Attanasio and Rohwedder (2003) ages 55-64, SERPS	OLS	time series of cross sections, UK		-.75 (.24)	

[a]Table 2, column 5. * Standard error estimates are unavailable. See the text for discussion of the statistical significance of these estimates. ** Computed from coefficient estimates that are significantly different from zero.

Table 2: Key Parameters in Calibration

Parameter	Value
Fixed arbitrarily	
Coefficient of relative risk aversion (α)	2.5
Consumption floor ($\bar{C}$)	3
Mean real rate of return ($\bar{r}$)	.02
Rate of time preference (δ)	.03
Variance of log out-of-pocket medical expenditure shock	1.0
Variance of rate of return shock	.00020
Variance of firm-specific log wage shock	.010
Variance of log wage shock	.030
Derived from HRS	
Serial correlation in log wage shock	.85
Serial correlation in out-of-pocket medical expenditure shock	.55
SSDI acceptance probability	.01
Calibrated to fit employment patterns	
Utility of employment intercept (γ_1)	-.003
Utility of employment slope (γ_2)	-.0004
Additional utility of moving from non-employment to employment (γ_3)	-.35
Additional utility of changing employers (γ_4)	-.08
Variance of value-of-leisure shock	.00091
Variance of value-of-consumption shock	.030

Notes: The SSDI acceptance probability is low because in the model it is assumed that application is costless if in bad health and not employed. The low acceptance rate is necessary in order to avoid an excessive level of SSDI enrollment.

Serial correlation coefficients are the square root of the coefficient on the lagged dependent variable in the wage and medical expenditure regressions in Appendix Table A1. These approximate the one-year serial correlation coefficients used in the model.

Table 3: Initial Conditions for Simulations

Variable	Value	Sample size
Age	51	
Education	12	
Age enrolled in pension plan	30	
Assets (median)	41	174
DC Balance (median)	17	40
Job Tenure	20[a]	
Work experience	24	158
Previous period annual earnings	35	174
Average Indexed Monthly Earnings (AIME)	30	135

Monetary amounts are in thousands of 1992 dollars. Computed from the HRS subsample of employed single white men aged 51-53 in 1992 with annual earnings of at least 10K.

[a] Mean job tenure for this sample is 14 years, but I use 20 years as the initial condition in order to make the value of job tenure at age 51 consistent with the assumed pension enrollment age of 30.

Table 4: Crowdout Computed Directly From Simulations

	1	2	3	4	5	6	7
Initial assets	41	150 (60)	41	41	41	41	41
Employment choice	yes	yes	yes	yes	no	no	no
Uncertainty?	Yes	yes	yes	no	yes	yes	no
Liquidity constraint?	yes	yes	no	no	yes	no	no
Defined Benefit							
Age	56	55	55	57	59	63	63
Dollar crowd out	-20	-17	-28	-47	2	-16	-53
Pension value	125	108	112	128	104	105	78
Proportional crowd out	-0.16	-0.15	-0.25	-0.37	0.01	-0.15	-0.68
Defined Contribution							
Age	58	56	56	61	63	63	63
Dollar crowd out	-8	-22	-5	1	-6	-7	-2
DC Balance	45	84	37	55	63	64	63
Proportional crowd out	-0.18	-0.26	-0.15	0.01	-0.11	-0.11	-0.02
Social Security							
Age	661	61	61	61	61	61	61
Dollar crowd out	-27	-64	-38	-49	-28	-42	-30
SS value	112	110	111	112	111	112	115
Proportional crowd out	-0.24	-0.58	-0.35	-0.44	-0.25	-0.37	-0.26

Notes: Age is the age at which the measurements are taken: the latest age at which no more than 25% of simulated cases have claimed the pension or Social Security. DB (DC) dollar crowdout is mean simulated household wealth in the DB (DC) case minus mean simulated household wealth in the no pension case. SS dollar crowdout is mean simulated household wealth in the no pension case minus mean simulated household wealth in the no pension/no SS case. DB pension value is the EPDV of future DB benefits, and SS value is the EPDV of future SS benefits. The proportional crowd out rows show dollar crowd out divided by pension value or DC balance. All figures are in thousands of 1992 dollars.

Table 5: Compensating Variation Associated with Pensions

	Initial Assets	Initial pension/SS wealth	Compensating Variation (CV)	CV/(initial pension/SS wealth)
Initial Assets 41				
NP	41	0		
DB	-12	83	53	0.64
DC	20	20	21	1.05
SS	360	0		
NP	41	86	319	3.70
Initial Assets 150				
NP	150	0		
DB	90	86	60	0.70
DC	97	64	53	0.84
SS	608	0		
NP	150	81	458	5.65

Notes: The initial assets column shows the amount of initial wealth that equates the EPDV of lifetime utility in the DB and NP cases, the DC and NP cases, and the NP and SS cases. Initial pension/SS wealth shows the EPDV of future benefits (or the DC balance) at age 51. CV is initial assets in the DB or DC case minus initial assets in the NP case, or initial NP assets minus initial NPNSS assets.

Table 6: Regression Models of Wealth Using Simulated Data

	Includes lagged wealth		Excludes lagged wealth	
	Correct pension/SS wealth measure	Incorrect pension/SS wealth measure	Correct pension/SS wealth measure	Incorrect pension/SS wealth measure
EPDV DB benefit	-0.18	-0.32	-.38	-.69
DC Balance	-0.19	-0.19	-.39	-.39
EPDV SS benefit	-1.01	-1.05	-1.44	-1.26
R^2	0.97	0.96	.81	.81

Notes: The entries are coefficient estimates on the indicated variables in regressions on household wealth. EPDV = Expected Present Discounted Value, DB = Defined Benefit, DC = Defined Contribution, SS = Social Security. The other explanatory variables, with coefficient estimates from the specification in the first column, are: current annual earnings (-3.18), EPDV of remaining lifetime earnings (.031), and (in the first two columns) lagged wealth (2.64). Initial assets are 41. Sample size is 107. Each observation is the mean over the random draws from a separate simulation, using the four pension/SS cases (NP, DB, DC, and NPNSS) in combination with alternative values of initial assets, the initial DC balance, the salary history, AIME, and the length of service in the pension job. One observation from each simulation is used: the one with the latest age at which no more than 25% of simulated cases have claimed the pension or Social Security.

Appendix Table A1: Parameter Estimates from the Health and Retirement Study

	Log wage	Log medical expenditure	Laid off (logit)	
			Full time	Part time
Intercept	3.03 (.522)	-8.237 (.505)	-8.371 (1.497)	-4.661 (2.947)
age	-.0009 (.016)	.065 (.005)	.056 (.016)	.032 (.030)
age sq.	-.0002 (.013)			
experience	.0012 (.0010)		-.011 (.006)	-.019 (.010)
experience sq	-.00026 (.00015)			
tenure	.010 (.001)		.040 (.006)	-.037 (.015)
tenure sq	-.0015 (.0003)			
bad health	-.028 (.013)	.487 (.042)		
education	.022 (.002)	.099 (.005)	-..037 (.016)	-.017 (.034)
lagged dependent variable	.730 (.008)	.307 (.007)		
black	-.032 (.015)	-.328 (.053)	.08 (.17)	.72 (.26)
birth year	-.0005 (.0018)	.047 (.006)		
part time	-.021 (.013)			
mean squared error	.151	6.065		
R squared	.646	.145		
Sample size	7,526	20,529	19,668	4,723

Notes: standard errors in parentheses.

Appendix Table A2: Multinomial Logit Models of Two-Year Health Transitions

	Good health in *t*		Bad health in *t*	
Health in *t*+2:	Bad	Dead	Good	Dead
Intercept	-2.402 (.722)	-7.709 (1.806)	-.780 (.942)	-4.065 (1.398)
Age	.030 (.007)	.078 (.018)	-.010 (.010)	.039 (.014)
Education	-.146 (.007)	-.077 (.019)	.067 (.009)	.048 (.014)
Birth year	.008 (.009)	-.007 (.022)	-.005 (.012)	-.023 (.017)
Black	.442 (.070)	.470 (.177)	-.222 (.084)	.155 (.117)
Mean of dependent variable	.115	.016	.265	.097
Sample size	16,553		5,161	

Notes: standard errors in parentheses. The first two columns report results from a multinomial logit model for cases in good health in period *t*. The outcomes are good health in *t*+2, bad health in *t*+2, and dead in *t*+2. The last two columns report results for the cases in bad health in period *t*.

Figure 1: Mean Simulated Employment Profiles

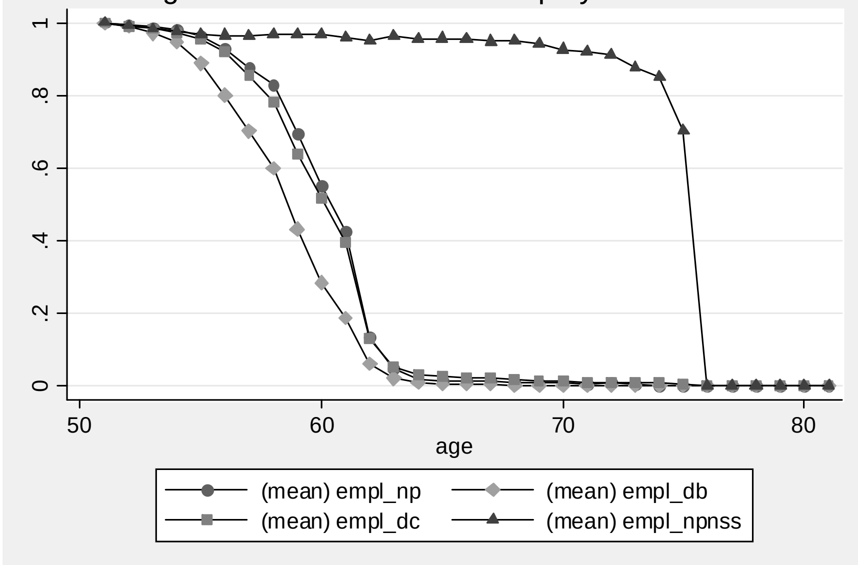

Figure 2: Mean Simulated Consumption Profiles

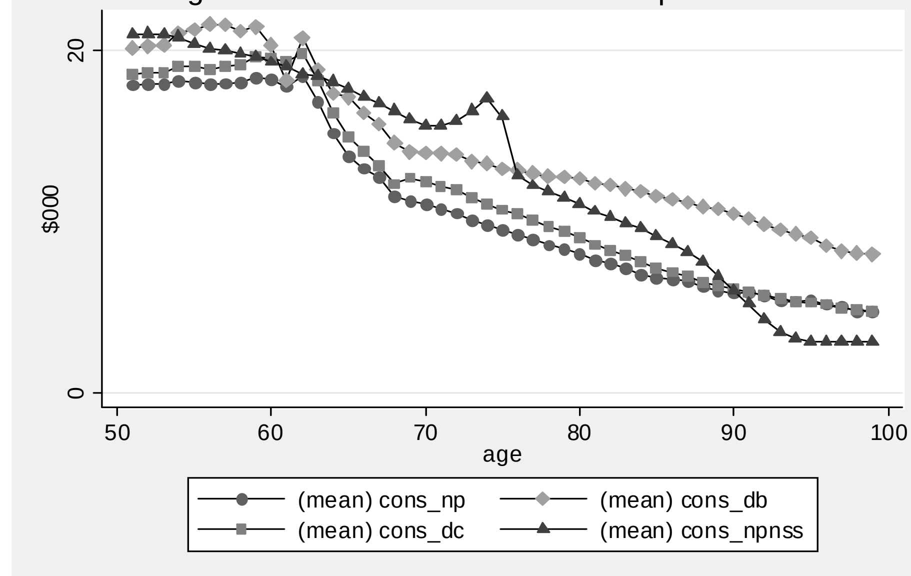

Figure 3: Mean Simulated Asset Profiles

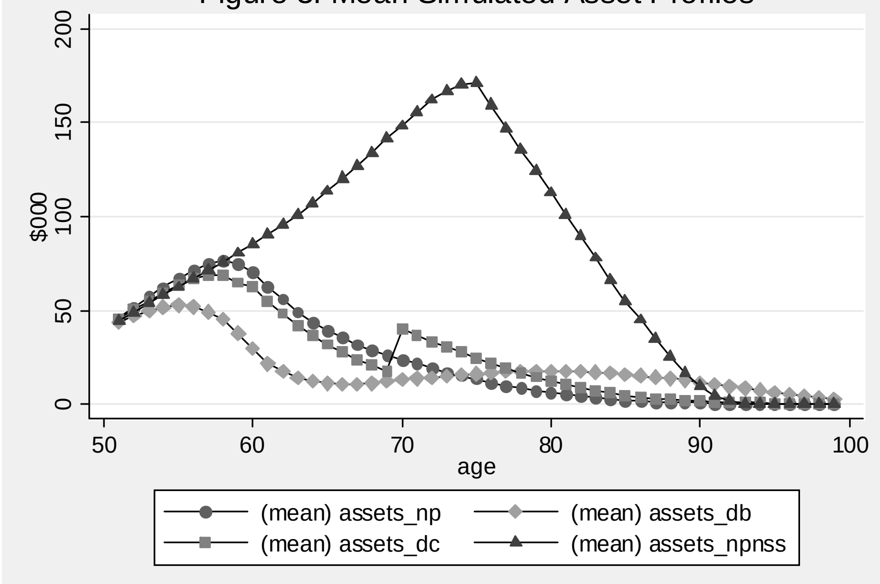

Figure 4: Mean Simulated Asset Profiles for NP DB and DC Cases

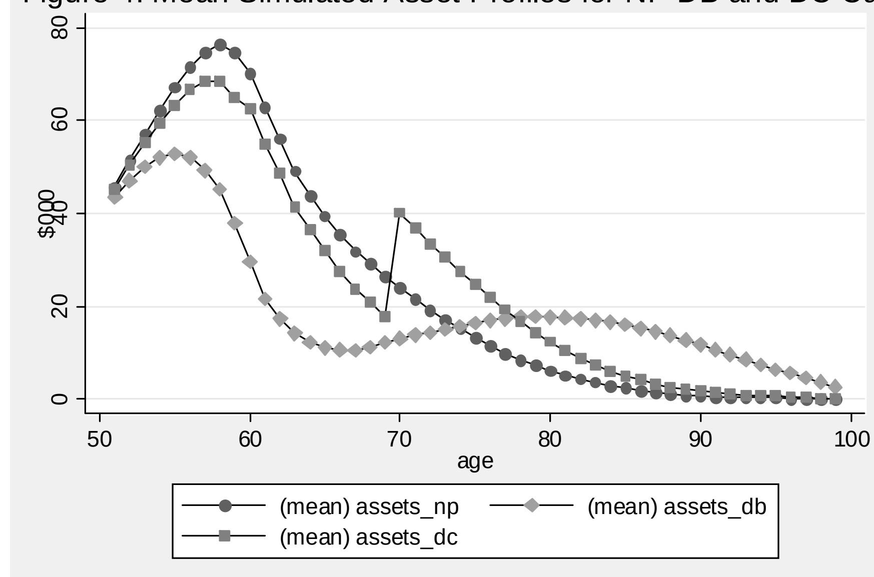

Figure 5: Defined Benefit Crowdout and EPDV of DB benefits

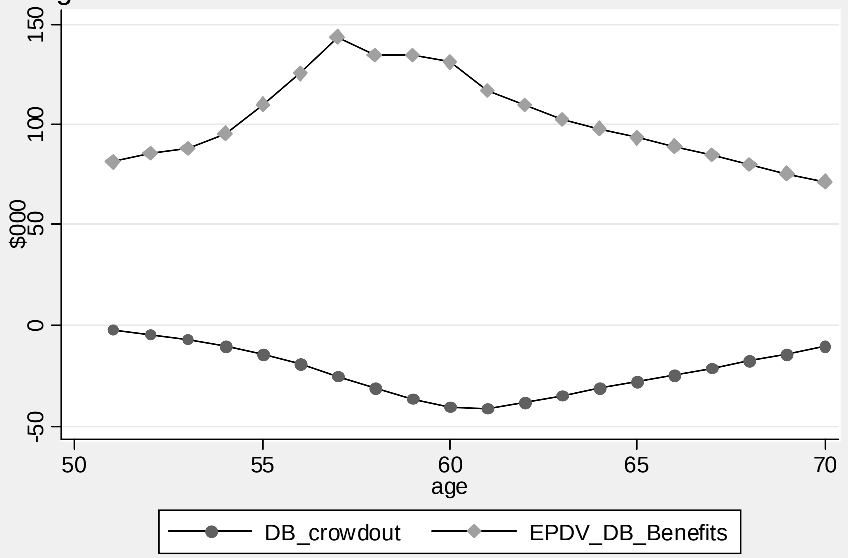

Figure 6: Defined Contribution Crowdout and DC Balance

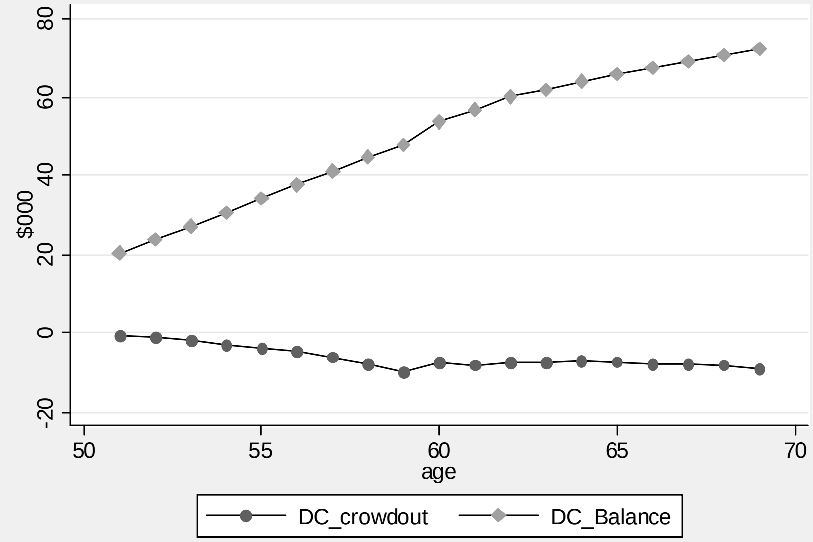

Figure 7: Social Security Crowdout and EPDV of SS benefits

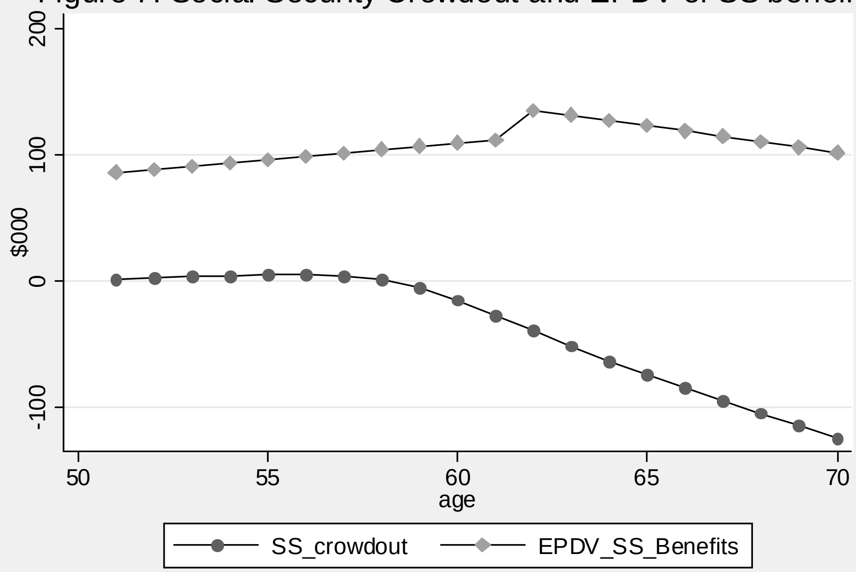

Figure 8: Proportional DB DC and SS Crowdout