# How Do Pensions Affect Household Wealth Accumulation?

David M. Blau
The Ohio State University

Work in progress; preliminary and incomplete

April 10, 2009

Financial support under grant R01-AG02199 from the National Institute on Aging is gratefully acknowledged. I appreciate detailed and helpful comments on related work from Rob Alessie, Gary Engelhardt, participants at the conference on Aging of the Population and Its Economic Consequences, Ramon Areces Foundation, Madrid, and seminar participants at Australian National University. Helpful advice on pensions from Gary Engelhardt, Alan Gustman, and Tom Steinmeier is 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 highly stylized version of the life cycle model, with a fixed retirement age, a perfect capital market, no uncertainty, and no institutional constraints on pensions. The model predicts one-for-one crowd out of household wealth by pension wealth over the life cycle. Empirical estimates of crowd out are usually much closer to zero, and it is not well understood whether such estimates are accurate indications of the true magnitude of crowd out or whether the strong assumptions imposed in the analysis result in severe misspecification. In this paper, I specify a richer life cycle model in which several of the key restrictions of the simple model are relaxed. The effects of pensions on household wealth are analyzed by solving and simulating the model. The compensating variation associated with pensions is treated as a measure of crowd out. The simulated data are then used to estimate regressions like those typically found in the literature. Preliminary results indicate that regression estimates of the effects of pensions on wealth accumulation estimated under the assumptions of the typical stylized life cycle model are quite misleading when those assumptions do not hold. Specifically, the extent of crowd out is substantially underestimated in the regression approach. This finding may help explain the common although not universal empirical finding of small crowd out.

## 1. Introduction

Employer-provided pensions are a form of "backloaded compensation," or forced saving. Economic reasoning based on the life cycle framework suggests 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 crowd out 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, and other channels of response such as labor supply. Alternatively, it has been suggested in the behavioral economics literature that pensions may serve as a commitment device, providing a means to avoid the problems of self control associated with time-inconsistent preferences (Benartzi and Thaler, 2007). In this scenario, pension saving could supplement rather than crowd out private saving.

Most analyses of pension effects on saving behavior are based on a simple highly 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 benefit while employed at the firm, etc.). In this model, an analytical solution for wealth can be derived, and the model predicts complete crowd out over the life cycle. Under the assumptions of the model, one can compute a measure of the capitalized value of the future benefits from a pension to use as an explanatory variable in a wealth regression.

It is well understood that intuition about crowd out may not hold when these very strong assumptions are relaxed. Feldstein's (1974) analysis of the effect of Social Security on saving demonstrates that when the timing of retirement is a choice, the sign of the effect of Social

Security wealth on private saving is ambiguous (see also Crawford and Lilien, 1981). 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 on the severity of the borrowing constraint (Rust, 1998). The concept of "pension wealth" may have little meaning when pensions are illiquid and borrowing is constrained. Uncertainty about future earnings, asset returns, and medical expenditures will induce precautionary savings, which could affect the magnitude of pension crowd out. Other savings motives such as for bequest purposes could affect the extent of crowd out. Finally, pension effects on saving in married-couple households are usually analyzed under the assumption of "unitary" preferences, implying that the effect of pension wealth is independent of which spouse receives it, but the "collective" model of household behavior with limited commitment implies that this approach can be misleading (Mazzocco, 2007).

Despite awareness of these issues, almost all 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 empirical analysis. The reason for this is that relaxing these assumptions leads to an intractable empirical model. Computing a measure of pension wealth is extremely complicated if retirement age is a choice or if there is a liquidity constraint or 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 clearest illustration of this problem is that the simplest life cycle model predicts complete crowd out: a coefficient of -1 in a regression of household wealth on pension wealth (adjusting for stage of the life cycle). Coefficient estimates less than -1 in absolute value are very common in practice. Researchers typically interpret such findings as indicating less than complete crowd out, implying that the simple life cycle model is incorrect. But in this case the regression is misspecified and it is not clear that such an inference is warranted.

In this paper, I specify a richer version of a 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, taxes, institutional constraints on pension claiming, and various forms of uncertainty. The effects of pensions on household wealth are analyzed by solving and simulating the model. I use the compensating variation associated with pensions to compute a measure of crowd out from the simulations without relying on measures of pension wealth computed under the strong assumptions of the simple model. I also compute such measures from the simulated data and use them to estimate regressions like those typically found in the literature. The reliability of such regression estimates can be evaluated by comparing the results to the "true" crowd out measures derived from the simulations. The results indicate that regression estimates of the effects of pensions on wealth accumulation estimated under the assumptions of the typical stylized life cycle model are quite misleading when those assumptions do not hold. Specifically, the extent of crowd out is substantially underestimated in the regression approach. This finding may help explain the common although not universal empirical finding of small crowd out. In ongoing work, I impose restrictions on the model such as a fixed retirement age, no uncertainty, and no liquidity constraint. I re-solve and simulate the restricted versions of the model in order to determine which assumptions matter the most to the finding that regression estimates of crowd out are so different from "true" crowd out.

The following section of the paper provides some background on pensions in the U.S., briefly characterizes the stylized life cycle model used as the basis for most empirical analysis, and summarizes previous empirical findings. 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 work in progress is described in section 6.

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

### *A. Background*

*Defined Benefit* pension plans provide employees with a promise of a pension benefit based on a formula that includes age and length of service at the date of retirement from the firm, and pre-retirement earnings.[1] A DB plan typically has a normal retirement age, at which the maximum possible benefit is available. Often, further service beyond the normal retirement age does not result in a higher benefit. Many DB plans also have an early retirement provision, defined by age or a combination of age and years of service, that allows a worker to retire before the normal retirement age with a reduced benefit. Many DB plans have strong incentives to remain with the firm until the early retirement age, and little incentive to remain after this benchmark. Benefits are 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. Workers are typically not "vested" in the DB plan until they have worked a certain number of years, with five years as the upper limit allowed by law. A worker who leaves the firm after becoming vested but before the early retirement age is entitled to a "vested deferred benefit" payable upon reaching

[1] This material is drawn largely from BLS (2003 and 2005) and Purcell (2007).

the early or normal retirement age.

*Defined Contribution* plans specify the amount or percent of the worker's earnings contributed to an 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 usually can choose the percentage of earnings he contributes. The contributions are made on a pre-tax basis, and are deposited in the worker's individual account. The assets in the account are allocated among the various investment options allowed by the plan. These options typically include relatively low-risk choices such as government securities, higher-risk options such as bond and stock funds, and in some notorious cases the high risk option of equity shares in the company that sponsors the plan. Workers are usually allowed to make the allocation decision, and can revise it periodically. 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 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. By one measure, participation in DB plans declined from about 25% of private sector employment in 1990 to 19% in 2004, while participation in DC plans rose from about 27% to 42% in the same period (Poterba, Venti, and Wise, 2007).

---

[2]There is a tax penalty for withdrawing funds from the account before age 59½, and withdrawal of funds must begin by age 70. Some plans allow the participant to borrow from the account before age 59½ with interest for a specified period of time.

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.[3]

*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[4]. There is no uncertainty, no income tax, and no institutional constraints on pension claiming. There is a perfect capital market with a rate of return $r$ in each period. 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 his employer each contribute a proportion $c$ of his salary to his DC account in each period, and the individual pays a proportional payroll tax $p$ to finance Social Security. The individual chooses consumption $C_j$, $j = t,...,T$, 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

---

[3]Pension coverage can be measured in several national data sources, including the Current Population Survey, the Survey of Income and Program Participation, the National Compensation Survey, the National Income and Product Accounts, and IRS Form 5500 data. Levels and trends often differ in the different sources, but all show the same trend away from DB and toward DC plans.

[4]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.

beginning of the period, and $\kappa_t \in (0, 1)$ is an adjustment factor for time until death that depends on $r$, $\delta$, $T$, and parameters of the utility function.[5] 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), 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 $t$, we obtain 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 pension benefits, and $S_t$ is the PDV of 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 ($\kappa_t$ approaches 1 as $t$ approaches $T$). $\kappa_t$ can be computed under assumptions about $r$, $\delta$, $T$, and the form and parameters of the utility function. The model thus implies a dynamic regression specification of the form[6]

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

where

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

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

---

[5]The perfect capital market assumption implies that the DC balance is perfectly fungible.

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

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 $\beta_4$, $\beta_5$, and $\beta_6$ are equal to -1, and coefficient estimates less than 1 in absolute value are taken as evidence of incomplete crowd out. 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 can be estimated. However, as noted above, a finding that the β's of interest are not equal to -1 may not be interpretable as evidence of incomplete crowd out, 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. The first panel shows results for studies that use household wealth as the dependent variable. 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 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 Health and Retirement Study to estimate a model like Gale's, using several estimation approaches. Their median regression estimate is roughly .15, significantly different from zero, and their Instrumental Variable median regression estimate is roughly .10, not significantly different from zero.[7] 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 crowd out is .012. Kapteyn et al. (2005) report a median regression estimate of -.115 for crowd out of private wealth by Social Security in Holland.

The studies by Attanasio and Rohwedder (2003) and Attanasio and Brugiavini (3003) listed in the lower panel of Table 6 use household expenditure data and estimate models of household saving.[8] They use public pension reforms in the U.K. and Italy as a source of exogenous 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 (SERPS) at ages 55-64 is -.75 (.24), indicating substantial crowd out. For Italy, the effects are -.49 (.11) for ages 46-55, -.21 (.14) for ages 56-60, and -.11 (.04) for ages 61-65. Their estimates are more likely than most to be free of bias from unobserved heterogeneity, so the large magnitude of their estimates for some age groups suggests significant crowd out by public pensions.

---

[7]Their IV approach uses pension wealth computed from the Health and Retirement Study pension provider data base for a set of "synthetic individuals" as an instrument for the pension wealth regressor, which is based on respondent-reported information.

[8]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.

To summarize, estimates of crowd out vary widely, from substantial crowd out, to no crowd out, 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 model specification.

**3. A Richer Life Cycle Model**

The model proposed 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.[9] 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) employment with a new employer, and (2) employment with the period $t$-1 employer. The latter 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 individuals can leave the pension-employer, collect the pension benefit, and work for another employer.[10] The length of a period is one year. The last

[9]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).

[10]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.

age at which a consumption decision is made 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.

The individual may face other choices as well. If he 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.[11] There is no separate claiming decision for DB pensions: age and job tenure at the time of exit fully determine whether the individual will receive a benefit and the benefit start date. It is assumed that the DB benefit must be taken as a 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, $\mu_k$ is an employer-specific wage component that is fixed for the duration of employment with an employer, and $\varepsilon_{kt}$ 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

[11]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. 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.

$\varepsilon_{3t}$ 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.[12] 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 $t$ is $\lambda(x_t\beta_4)$, where $\lambda$ is a binomial logit function.

An individual who is in bad health, chooses non-employment in a given period, and has accumulated a sufficient number of quarters of coverage (20) is eligible to apply for a benefit from the Social Security Disability Insurance (SSDI) program. Application is assumed to be costless for an individual in bad health who chooses non-employment. The outcome of the application is not determined until the following period, reflecting the fact that there is a waiting period for benefits (which in reality is five months rather than one year, as assumed here). At the beginning of $t$+1, an individual who applied for SSDI in period $t$ is accepted to the program with probability $\psi$. Once enrolled in SSDI, individuals remain in the program even if their health subsequently improves or until they re-enter employment or reach their Social Security Full Retirement Age. The SSDI benefit is calculated using Social Security rules, as described below. This is a crude approximation to the complicated application and acceptance process.

An individual is eligible to claim his Social Security retired-worker (OASI) benefit if he has accumulated a sufficient number of quarters of coverage (40) and has reached the Social

[12]Health status affects medical expenditure, wage offers, and mortality risk. It could also affect preferences, but for simplicity this is not allowed here.

Security early retirement age (62). The OASI benefit, $s_t$, upon initial entitlement is determined by the function $s_t = s(AIME_t, fe, E_t, a_t)$, 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, $fe$ is the age at which the individual first became entitled to the benefit, and $a_t$ is age in period $t$. Age and current earnings matter because there is an earnings test at some ages. $H_t$, is assumed to be 1,000 for part-time employment and 2,000 for full-time employment.

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 began (which could be different from the age at which the individual left 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 ($t$ - $a_b$). A sample of DB pension formulas from the Health and Retirement Study is used in solution of the model, as described in the next section.

DC pension plans are characterized by the account balance, and the individual 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.[13]

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 $\theta_t$ is an idiosyncratic household-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 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 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.

---

[13]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 a liquidity constraint.

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

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

where $\tau_t$ includes federal income and payroll taxes, calculated using the rules in effect for 1992. 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 $A_t^* = A_t + I_t - c_t$, (subject to the conditions described in the previous paragraph), 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:

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

where $W = 1$ if employed, zero otherwise; $NJ = 1$ if a new job is chosen, and zero otherwise; $\varepsilon_{ct}$ is an *iid* shock to the marginal utility of consumption; and $\varepsilon_{\ell t}$ is an *iid* shock to the marginal utility of leisure. $\gamma_1$ is the utility from working, $\gamma_2$ is the additional utility from working if the individual was not employed in the previous period, and $\gamma_3$ is the additional utility 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. There is no bequest motive.

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 $\delta$, 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 points in the period *T* state space (which consists of 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 *T*, of the period *T* value function (the "*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 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.

## 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 data also provide a sample of DB pension formulas to use 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. The full HRS sample of DB pension plan formulas is used in solution and simulation. A subsample of the 12,652 cases is used to estimate the parameters of the wage, medical expenditure, health, layoff, and SSDI acceptance functions. This subsample excludes same-sex couples, married couples in which the age difference between the spouses exceeds 10 years, and cases with missing data on year of birth, leaving 10,610 individuals residing in 6,413 households as of the initial interview. A subsample of single men born from 1939-1941 who were employed in the first wave with earnings of at least $10K is used to generate initial conditions for the simulations.

*A. Pensions*. The HRS asks respondents a substantial battery of questions about DB pensions, including the ages of early and normal retirement, expected benefits at the early and normal retirement ages, and the expected age of retirement and benefit at the expected age. Some studies have used this information to construct a measure of DB pension wealth (Engelhardt and Kumar, 2007; Chan and Stevens, 2008), but the information is not sufficient to calculate benefits for all possible scenarios for retirement age and future salary. 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 the 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 developed an accurate plan-specific regression approximation of the benefit formulas, described in the Appendix, and used this approximation in solution and simulation.

*B. 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. 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 estimated mean squared error from the log wage regression serves as an estimate of the variance of the log wage shock.[14] 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 Tables 2 and 3, and a summary of the key parameters derived from the estimates is shown in the top panel of Table 4.

Second, several other key parameters were set arbitrarily. These include the coefficient of relative risk aversion ($\alpha$=2.5), the consumption floor ($\bar{C}$=3), the mean rate of return ($\bar{r}$=.02), and the rate of time preference ($\delta$=.03). The second panel of Table 4 shows these and the values of

[14]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 the log out-of-pocket medical expenditure distribution implied an implausibly high variance, so it was not used in the calibration. This variance was chosen arbitrarily.

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 utility of employment parameters ($\gamma$’s) and the variances of the value-of-leisure and value-of consumption shocks in the utility function 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 : one with no pension, a second with a DB pension, and a third with a DC pension. In all three cases the household is assumed to be eligible for Social Security. A fourth simulation eliminates Social Security. 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.[15] 300 simulations are computed for the no-pension scenario, and three simulations per pension plan for the DB and DC scenarios. The results are then averaged across simulations. The simulations do not incorporate deaths.

---

[15]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.

As noted above, an important advantage of the simulation approach used here is that the impact of pensions on household saving can be analyzed without computing measures of pension wealth that are valid only under very strong assumptions. However, it is not obvious how to characterize the extent of pension crowd out in this model. I compute the value of a pension to the household by determining how much of the household's initial asset endowment must be subtracted in order to provide the same EPDV of remaining lifetime utility with the pension and without the pension. This is a measure of the compensating variation associated with the pension. This is not a typical measure of crowd out, but it has the virtues of being interpretable in terms of household welfare and being well defined and computable in the framework of a rich life cycle model.

*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 mean observed 1991 assuming a one percent real growth rate.[16] Other initial conditions were generated from the subsample of single men in the HRS who were

---

[16]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.

employed at ages 51-53 with earnings of at least $10K in 1991. Mean values of job tenure, work experience, previous-year earnings, and AIME from this sample were used. 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 the HRS respondents.

The household is endowed with initial assets equal to the median value of assets in the single male sample at ages 51-53, $41K. 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 5 summarizes the initial conditions.

## 5. Results

To gauge the magnitude of crowd out implied by the model, I simulate the model for the no-pension case, and compute the mean EPDV of optimized remaining lifetime utility in the first period, for the given initial asset level. I then simulate for the DB pension case and find the value of initial assets that yields the same mean EPDV of optimized remaining lifetime utility as in the no-pension case. Let $A_{NP}$ be the initial asset level for no-pension case, and $A_{DB}$ the initial asset level for the DB pension case that yield the same lifetime utility. The proposed measure of crowd out is $(A_{DB}/A_{NP}) - 1$: the compensating variation as a proportion of initial no-pension assets. In the DC case, the crowd out measure also adjusts for the initial account balance endowment, $B$. The proposed measure of crowd out in this case is $(A_{DC} - B)/A_{NP} - 1$, where $A_{DC}$ is the initial asset level for the DC pension case that yields same lifetime utility as in the no pension case, conditional on the value of $B$. This can also be interpreted as a measure of compensating variation, treating the initial no-pension asset level as $A_{NP} + B$. This is not strictly correct because privately held assets and DC balances are not perfectly fungible. I simulate Social Security crowd out in a similar way, but in this case I compute the amount of additional assets required to compensate the no-pension household for the loss of Social Security. Letting $A_{NPNoSS}$ be the value of initial assets that provides the no-SS-no-pension household the same level of lifetime utility as the no-pension household, the crowd out measure is $(A_{NP}/A_{NPNoSS})$ -1. Recall that the initial asset level for the no-pension case is 41 (all monetary amounts are expressed in thousands of 1992 dollars), and the initial DC balance is 17.

The simulated crowd out measures are:

DB Crowd out = $(A_{DB}/A_{NP})$ - 1 = (23/41) - 1 = -.44

DC Crowd out = $(A_{DC} - B)/A_{NP}$ - 1 = (18-17)/41 - 1 = -.98

SS Crowd out = $(A_{NP}/A_{NPNoSS})$ -1 = (41/93) - 1 = -.56

Qualitatively, these magnitudes seem plausible. DB pensions and SS are illiquid, so crowd out of less than 100% is plausible. DC pensions are much more liquid, so crowd out of close to 100% is also plausible. Claiming a DC balance before age 60 results in a tax penalty of 10% of the balance, so one might expect crowd out closer to 90%. But the DC balance is claimed before age 60 in only 20% of the simulated cases, as shown below. 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.[17] It is intuitive that crowd out will be sensitive to initial conditions, and other simulation results not reported here verify

[17]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 further increase the liquidity of DC balances and could in principle result in crowd out greater than 100%.

this. Median assets among married couples are much higher than for single men (and single women), so the compensating variation measures computed here may not be representative of the population.

Figure 1 illustrates the consumption, employment, asset, and pension claiming profiles generated by the simulations. The calibration process focused only on producing a "reasonable" employment profile, and did not attempt to fit observed employment and asset patterns. The initial assets used in the DB and DC simulations are the values that equalize lifetime utility with the no pension case: 23K and 18K, respectively.

Simulated wealth peaks quite early, at ages 54-56. DB pension holders expect to receive lifetime annuities from both SS and the pension, which reduces the need to save for retirement. Thus it is not surprising that the initial wealth gap of 17K between the no-pension and DB-pension cases grows to about 40K by the early 60s. Non-pension and DC households expect an annuity from SS only. Despite holding low wealth, the DB households enjoy higher consumption than the other households after age 75.[18] The wealth gap between the no-pension and DC-pension cases is roughly constant until age 70, the age at which DC balances that have not yet been claimed must be transferred into the household's private wealth. The pension claiming graph shows that 40% of DC cases wait until age 70 to claim the balance, consistent with the incentive to enjoy tax-free asset returns as long as possible. The spike in wealth for DC households at age 70 is used to finance a higher level of consumption for several years compared to no-pension households. The consumption profile is flat for NP and DB households until about

[18]Figure 2 shows the simulated income profiles. The Appendix shows employment and asset profiles of single men in the HRS.

age 60, after which it declines steadily, with the exception of the temporary surge for DC households in the early 70s. The DB consumption profile is generally similar but exhibits some swings in the mid 50s to mid 60s. The DB profile is also more erratic, with spikes of 3-4K in a few cases.

The simulated employment patterns show that employment declines with age, and DB pension holders retire earlier on average than other households. Many DB plans have an early retirement age of 55, and the large drop in DB employment at 55 is clearly related to the abrupt drop in the implicit returns to continued employment past the early retirement age. The employment rate of DB-holders falls more rapidly than for other households at ages 56-60; these are also common early retirement ages in DB plans. The DB employment rate reaches zero at age 65, which is counterfactual but not drastically inconsistent with the data. The employment rate of non-DB holders drops too smoothly in the simulations compared to patterns observed in the data. In particular, the simulations do not capture the sizeable drop in employment at age 62, the earliest age at which Social Security retirement benefits are available. In the trial-and-error process used to calibrate the employment parameters, there were many cases in which the simulations did yield a large drop in employment at age 62, suggesting that Social Security retirement incentives are not misspecified in the model. The simulations show a sizeable drop in employment at age 65 for DC and especially NP households. The Social Security disincentive to work past age 65 was fairly large for earlier birth cohorts, but has been eliminated for more recent birth cohorts. Thus, it is not clear why there is a large drop in employment at 65, since the simulation is for an individual born in 1941, the latest birth cohort in the HRS cohort.[19]

It is interesting to note that DB pension holders enjoy both greater leisure, and, after age 70, greater consumption than their NP and DC counterparts, despite equalized lifetime expected utility. DB consumption is lower in the mid 50s by a fairly small amount. It seems that in expectation, the relatively small up front lower consumption for DB households yields substantial later benefits. The effective rate of time preference, which includes a rising risk of mortality is apparently high enough to rationalize this pattern.

Now I use the simulated data to estimate wealth regressions like those described above. I compute the EPDV of lifetime SS and DB pension benefits, and of remaining lifetime earnings, using the simulated data. In order to compute these capitalized benefit and earnings values, I must assume fixed and known retirement and claiming ages (I use the observed retirement and claiming ages from the simulations), no uncertainty (except over the age of death), no liquidity constraint, and no institutional restrictions on pension claiming. These assumptions are necessary in order to make the calculations feasible.[20] But the data are generated from simulations of a model in which these assumptions were not imposed. This is analogous to the empirical crowd out literature: strong assumptions are imposed, while the data are generated from a process in which presumably these assumptions did not hold. The regressions on the simulated data yield estimates of crowd out that can be compared to those found in the literature and to the

[19]The incentive in Social Security to work past age 65 is a function of the Delayed Retirement Credit (DRC), the percent increase in benefit per year of delay in claiming. The DRC was small for earlier birth cohorts (3.5% for the 1931 birth cohort), but is roughly actuarially fair for more recent birth cohorts (8% for the 1941 cohort).

[20]The DC balance is generated directly from the simulations, so no additional assumptions are needed to calculate it.

compensating variation measures described above. The sample used in the estimation include periods in which a simulated individual remains employed on the job held at the beginning of the simulation.[21] I compute the life cycle adjustment factor ($\kappa_t$) using a continuous time approximation from Gale (1998).[22]

Estimates for several specifications are reported in Table 6. The baseline specification in column 1 includes lagged wealth, as suggested by the simple discrete time version of the life cycle model described in Section 2. The capitalized values of the DB pension benefit, the Social Security benefit, and future earnings have been adjusted by $\kappa_t$, and the current earnings and lagged wealth variables have been adjusted by $1-\kappa_t$. The estimates in column 1 imply crowd out of essentially zero by DB pensions, -.34 for DC pensions, -.14 for Social Security. The ordering of these estimates is similar to the ordering of crowd out measures based on the compensating variation: smallest for DB pensions, larger for SS, and largest for DC pensions. But the regression estimates are much smaller in absolute value.

The baseline specification includes three variables not implied by the simple life cycle model: age and dummies for DB and DC coverage. These variables may pick up features of behavior that are not well-captured by the capitalized pension measures, which can only be

---

[21]This is the approach taken in most previous studies. For the DB and no-pension cases, post-retirement periods could be used as well, although crowd out behavior following retirement could be quite different than during the wealth accumulation phase. For DC households, transfer of the DC account balance into private wealth would artificially generate large crowd out, so it would not be appropriate to use post-retirement (more precisely, post-claiming) observations for these households.

[22]The continuous time version of the adjustment factor is $\kappa_t = (e^{xS} - 1)/(e^{x\tau} - 1)$, where $x = (\bar{r} - \delta^*)/\alpha - \bar{r}$, $S$ = age - 30, $\tau = T - 30$ ($T$ = 100), and $\delta^* = \delta$ + the mortality hazard. The mortality hazard is approximated here by a constant value of .01. Age 30 is the assumed age of enrollment in the pension plan, so the discounting is to age 30, following Gale (1998).

computed under assumptions that are at odds with the data-generating process. The coefficients on these variables suggest that wealth grows at a faster rate with respect to age than can be explained by the other variables, and that wealth held by individuals covered by DC plans is about 10K higher than can be explained by these variables. These variables are dropped in the specification reported in column 2, and the resulting crowd out estimates are somewhat sensitive to specification: DB crowd out increases to -.10, DC falls to -.23, and SS increases to -.48. The $R^2$ falls from .93 in column 1 to .79 in column 2. The specification in column 1 is similar to those reported in the literature in the sense that many control variables are typically included.

The wealth distribution is highly skewed, and it is common in the literature to use quantile regression to estimate crowd out. Column 3 reports median regression estimates of the baseline specification. The estimated DB crowd out is very small, as in column 1, but the DC and SS crowd out estimates are quite a bit larger in absolute value compared to column 1: -.57 vs. -.34 for DC, and -.30 vs. -.14 for SS. The Appendix shows a kernel density estimate of the wealth distribution for the sample used in the regressions. It shows much less skewness than in real wealth data, but there is some skewness.

Column 4 reports results from a commonly used specification in which the DC balance and the capitalized values of the DB pension and SS benefits are summed to form a single explanatory variable. The crowd out estimate in this case is -.19. This is presumably a weighted average of the effects of the individual components of the sum, but the frequency of the different pension types in the simulation is not necessarily representative of the population, so not too much weight should be put on this estimate. The specification in column 5 uses unadjusted variables. The crowd out estimates for DC and SS are smaller than in the baseline, as expected since typical values of $\kappa_t$ in the age range of the regression sample are .5 to .6.

The last two columns report results from static regressions, in which the lagged wealth variable is omitted. This is a common specification in the literature, but it may not be appropriate in this context since the initial assets of the DB and DC households have been set deliberately lower than for the no-pension household. For this reason, the sample in the next-to-last column is limited to ages 59 and above, allowing eight periods of adjustment in wealth. The results imply crowd *in* of .19 for DB wealth, and fairly sizeable crowd out of -.56 for the DC balance and -.42 for SS wealth. The age restriction is lifted in the last column, producing smaller estimates in absolute value for all three measures.

To summarize, the regression estimates of crowd out are substantially smaller than the amount of crowd out implied by the compensating variation calculations:

| | | |
|---|---|---|
| Estimated DB crowd out: | +.17 to -.10 | (CV = -.44) |
| Estimated DC crowd out: | -.19 to -.57 | (CV = -.98) |
| Estimated SS crowd out: | -.04 to -.48 | (CV = -.56) |

It is not obvious that this is an appropriate comparison, but compensating variation is a useful indication of the value of a pension in an environment in which "pension wealth" is not a meaningful concept.

## 6. Work in Progress

An important question is which features of the model lead to less than 100% crowd out for DB pensions and Social Security? The most obvious candidate is the liquidity constraint, but other features may be important as well. Another interesting question is which features of the model cause the wealth regressions to underestimate crowd out? Work on these questions is in progress. It is also important to examine the sensitivity of the results to alternative parameter values. The ultimate goal is to estimate the key parameters structurally, removing much of the arbitrariness in the calibrations. Finally, most middle aged households contain married couples, so the model should be extended to include couples if it is to be as relevant as possible empirically. Work on this extension, along the lines of Van der Klaauw and Wolpin (2008) is in progress.

## Appendix A: Calculating Defined Benefit Pension Benefits

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.

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.

Mazzocco, Maurizio. 2007. "Household intertemporal behavior: a collective characterization and a test of commitment," *Review of Economic Studies*, 74 (3): 857-895.

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: Parameter Estimates from the Health and Retirement Study

| | Log wage | Log medical expenditure | Laid off (logit) |
|---|---|---|---|
| Intercept | .461 (.522) | -7.837 (.519) | -3.305 (.542) |
| age | -.0009 (.016) | .065 (.005) | .005 (.008) |
| age sq. | -.0002 (.013) | | |
| experience | .0012 (.0010) | | |
| experience sq | -.00026 (.00015) | | |
| tenure | .010 (.001) | | -.038 (.004) |
| tenure sq | -.0015 (.0003) | | |
| bad health | -.028 (.013) | .487 (.042) | |
| education | .022 (.002) | .099 (.005) | -.038 (.013) |
| lagged dependent variable | .730 (.008) | .307 (.007) | |
| black | -.032 (.015) | -.328 (.053) | |
| 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 | 26,793 |

Notes: standard errors in parentheses.

Table 3: 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*.

Table 4: Key Parameters in Calibration

| Parameter | Value |
|---|---|
| **Fixed arbitrarily** | |
| Coefficient of relative risk aversion ($\alpha$) | 2.5 |
| Consumption floor ($\bar{C}$) | 3 |
| Meal real rate of return ($\bar{r}$) | .02 |
| Rate of time preference ($\delta$) | .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 |
| **Derived from HRS** | |
| Serial correlation in log wage shock | .85 |
| Serial correlation in out-of-pocket medical expenditure shock | .55 |
| SSDI acceptance probability | .01 |
| Variance of log wage shock | .151 |
| **Calibrated to fit employment patterns** | |
| Utility of employment ($\gamma_1$) | -.001 |
| Additional utility of moving from non-employment to employment ($\gamma_2$) | -.23 |
| Additional utility of changing employers ($\gamma_3$) | -.036 |
| Variance of value-of-leisure shock | .00012 |
| Variance of value-of-consumption shock | .01000 |

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 Table 2. These approximate the one-year serial correlation coefficients used in the model.

Table 5: 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 | 14 | 158 |
| Work experience | 22 | 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.

Table 6: Regression Models of Wealth Using Simulated Data

| | baseline | no extra controls | Median regression | DB + DC + SS | No adjustment ($\kappa_t$) | No lag, age≥59 | No lag, all ages |
|---|---|---|---|---|---|---|---|
| 1. EPDV DB benefit | .01 | -.10 | -.002 | | .02 | .17 | .10 |
| 2. DC Balance | -.34 | -.23 | -.57 | | -.19 | -.56 | -.39 |
| 3. EPDV SS benefit | -.14 | -.48 | -.30 | | -.04 | -.42 | -.21 |
| 1 + 2 + 3 | | | | -.19 | | | |
| Current annual earnings | 1.17 | .54 | 1.42 | 1.22 | .42 | 1.21 | 1.79 |
| EPDV of remaining lifetime earnings | -.02 | -.10 | -.02 | -.40 | -.01 | -.04 | -.04 |
| Lagged wealth | 2.49 | 2.69 | 2.48 | 2.55 | .87 | | |
| Age | 3.74 | | 3.77 | 3.46 | .56 | 3.0 | 5.0 |
| DB dummy | .50 | | 1.20 | 14.1 | -3.5 | -27 | -30 |
| DC dummy | 10.8 | | 10.5 | 3.9 | 3.8 | -7 | -13 |

Notes: Monetary variables are in units of thousands of 1992 dollars. EPDV = Expected Present Discounted Value, DB = Defined Benefit, DC = Defined Contribution, SS = Social Security.

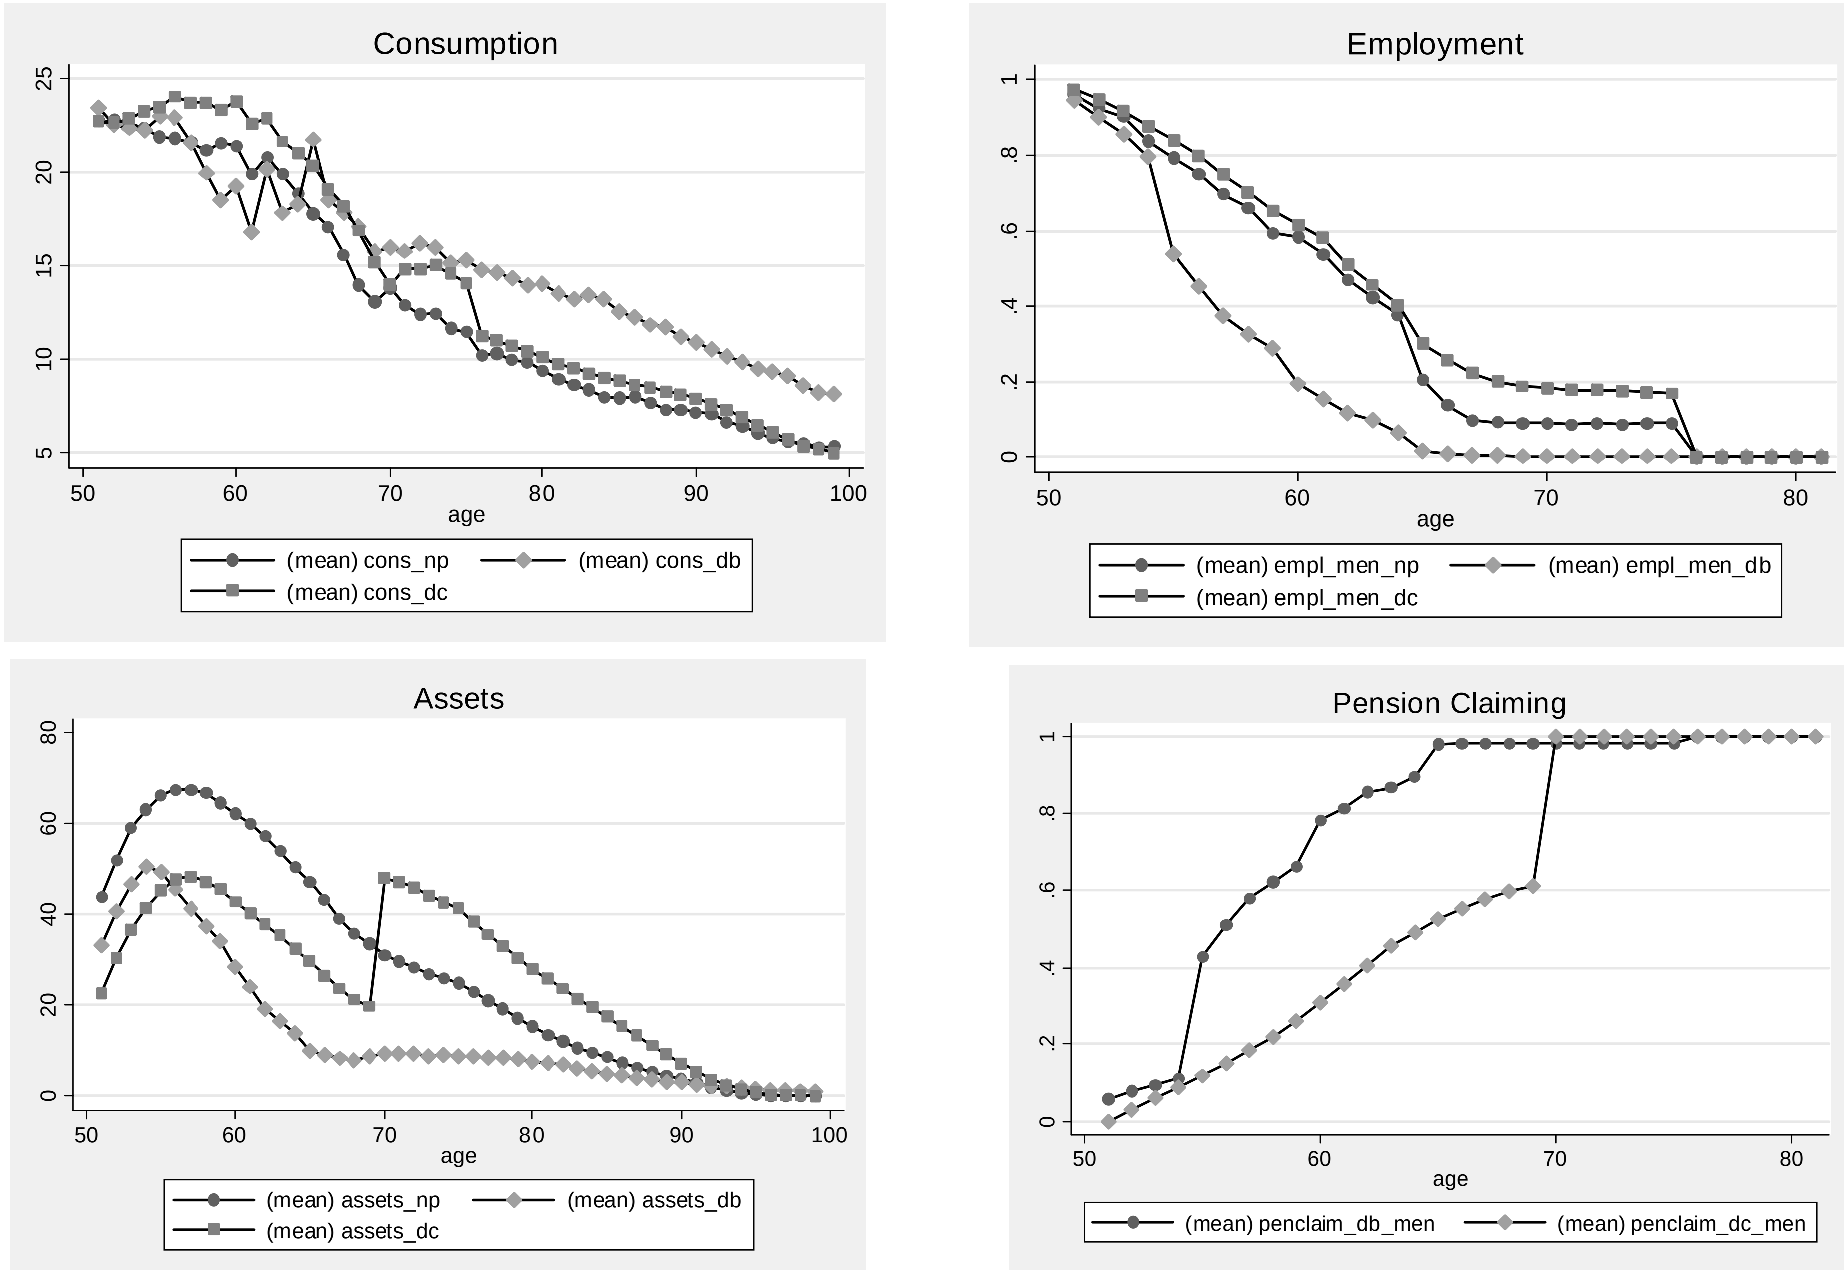


Figure 1: Simulated Profiles of Consumption, Employment, Assets, and Pension Claiming Age, by Pension Coverage

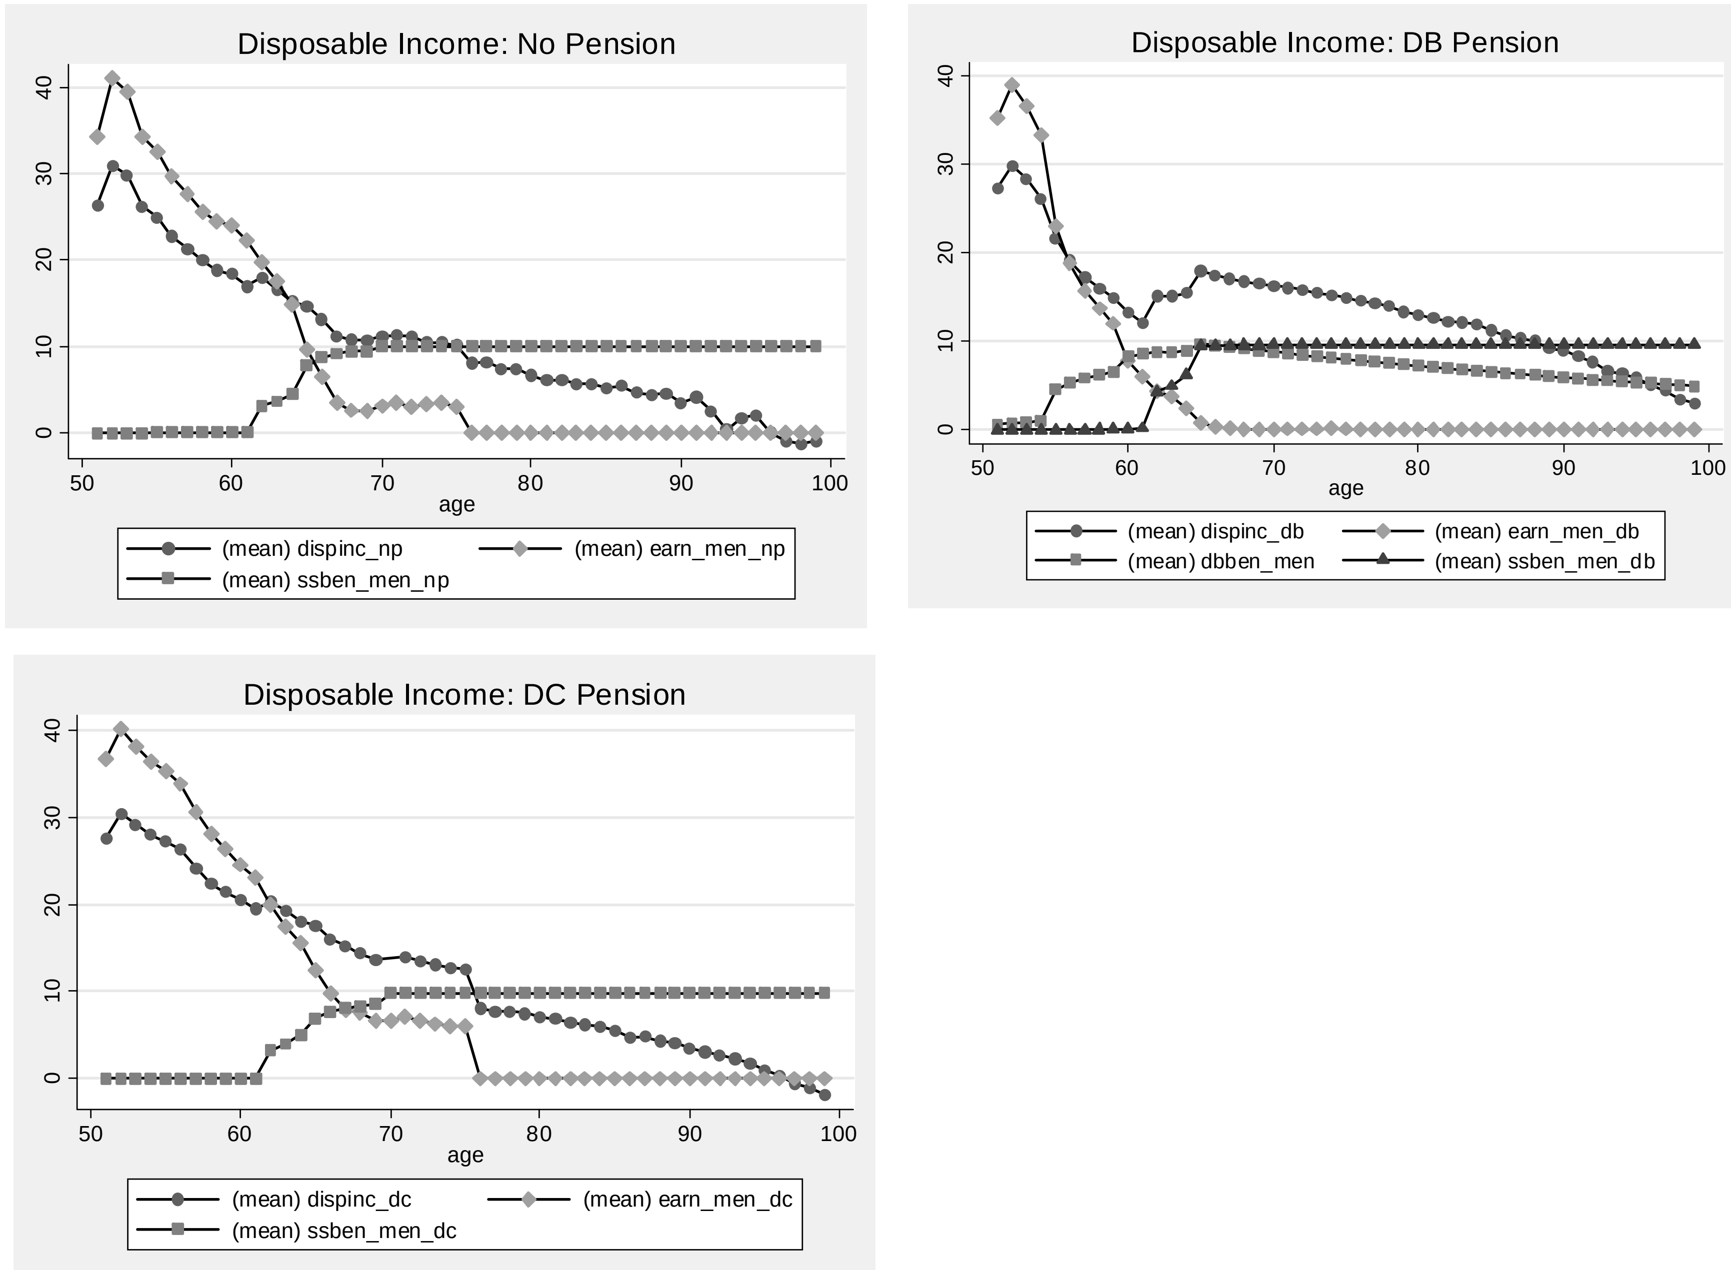


Figure 2: Sources of Disposable Income, by Pension Coverage

Note: disposable income equals earnings + benefits + asset returns - taxes - medical expenditure.

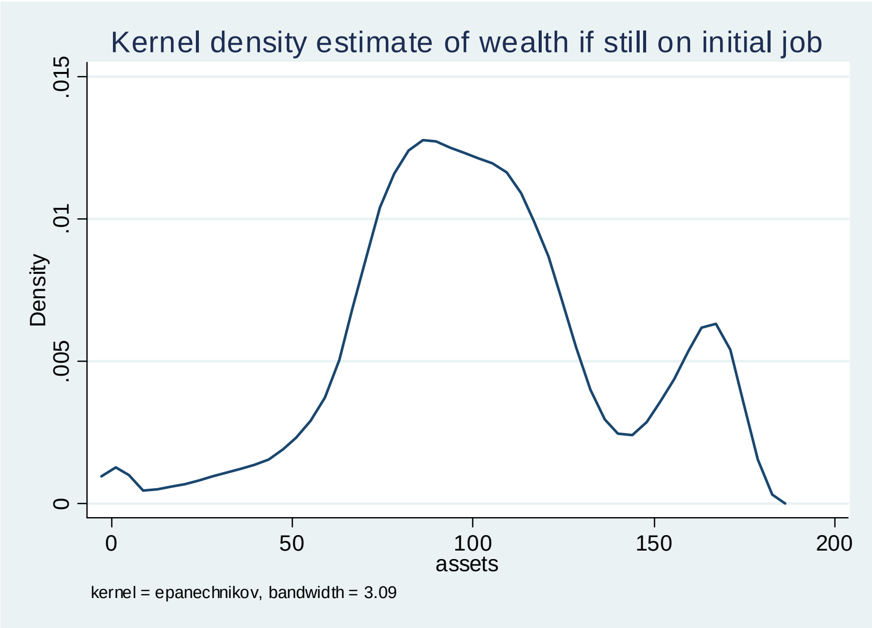


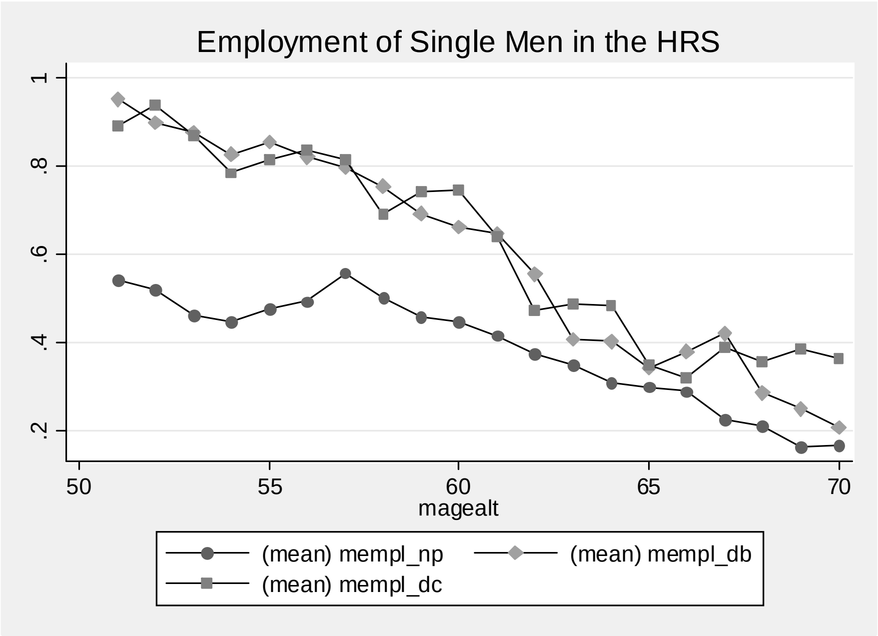


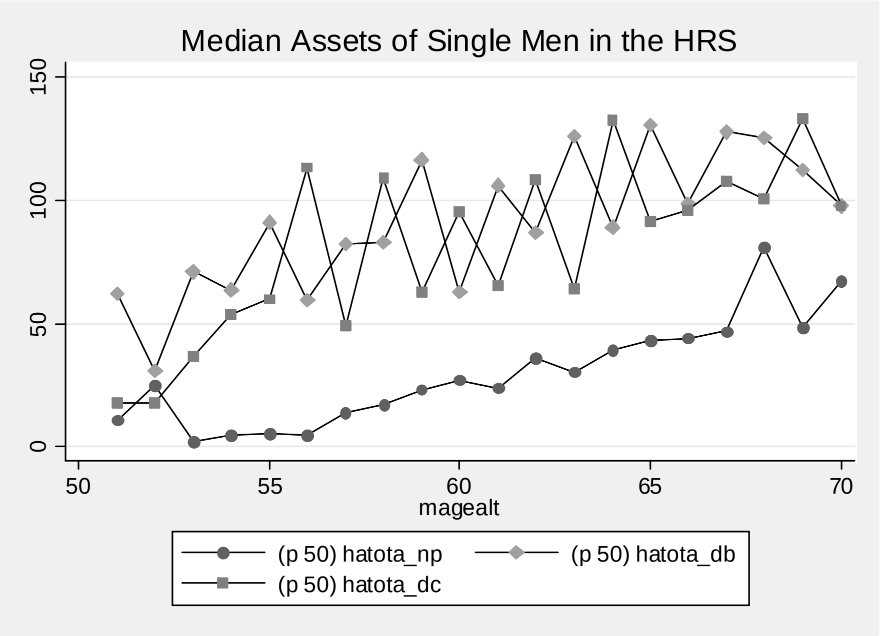


Appendix Figures